

Multi-Phase Clock Generation for Programmable Broadband Delay Circuits

Armon H. James
`armon.james@alfheimdesigns.eu`

Department of Electrical and Information Technology
Lund University

Supervisor: Fabien Mesquita (Ericsson), Pallavi Paliwal (Ericsson),
Baktash Behmanesh (LTH) in cooperation with LTH and Ericsson

Examiner: Pietro Andreani

June 16, 2026

Acknowledgements

This master's thesis represents the culmination of two years of study and research in the field of integrated circuit (IC) design. This work was conducted at Ericsson AB, which provided essential resources, including computers, software, technical expertise, and guidance necessary for the research and testing phases.

I would like to thank my supervisors, Pallavi Paliwal, Fabien Mesquita, and Rehman Akbar, for their insight, guidance, and patience throughout this work. I am also grateful to Daniele Mastantuono for his confidence in selecting me for this project. My appreciation extends to my teaching assistant, Vilgot Snygg, for the extensive discussions on True Time Delay circuits and theory.

I extend special thanks to my partner, Dr Kaitlin Schaal, for her valuable insights and advice on writing and presenting, as well as to my friends for their enduring support over the years.

Finally, I would like to thank my late mother and my family, whose unwavering support throughout this journey has been indispensable to my achievements.

Statement of Originality

I hereby affirm that the Master's thesis I present here was researched and composed by myself, and that the work is my own except where explicitly stated and carefully referenced.

In preparing this thesis, I have used the company-provided GenAI Playground, which offered models from OpenAI's ChatGPT, Anthropic's Claude and others, to assist me in researching material and ensuring mathematical rigour. I have also used it to improve the language, flow and readability of the material presented. When using these tools, I have always carefully checked and reviewed the edited content, and I take full responsibility for the entire thesis.

Abstract

Many circuits in integrated-circuit design require a time-shifted or multi-phase clock. In time-interleaved ADCs, multiple phases enable a higher effective aggregate sampling rate without increasing the individual clock rate. In contrast, in SerDes links, bit throughput increases while keeping each lane at a lower base clock. In 5G and 6G wireless equipment, multiple phases can be used to drive a true time delay (TTD) block, enabling programmable broadband delays for applications ranging from beamforming to power-amplifier predistortion. This work targets the latter: a bank of phases at a fixed frequency clocks the sample-and-hold taps of a TTD network, where the selected phase sets the delay applied to the sampled signal. The achievable delay range and resolution are therefore set directly by the number of phases and their spacing, motivating a generator that produces many accurately spaced phases at RF. This thesis presents a duty-cycle-controlled multi-phase clock generator that achieves 90 fs cycle-to-cycle jitter at 5.64 mW total system power.

Popular Science Summary

Imagine you were at a game contest with 4 of your friends. To make it fun, the rules are weird.

- Each of your friends is on a swing and needs you to push them.
- Each time you push, a buzzer makes a specific sound.
- Each time a friend gets 5 pushes, a different buzzer sounds.
- No swing must stay still for more than 20 seconds. So you run around and push each friend 5 times each. Then start over.

Now think about it. Your buzzer sounds once every second, but each of your friends' buzzers sounds once every 20 seconds. Congratulations, you've just created 5 different clocks. You are the main clock, and each of your friends is a clock that is 4 times slower than yours. Even more, each of your friends' clocks are spaced in time. They have the same frequency, but they are operating spaced every 5 "ticks" of your clock.

If you could take this game further with more people, someone in another part of the game could do something every time they hear a specific buzzer. And since your friends' buzzers are all spaced in time, those people would be doing things at the same frequency just at different times.

This is called Multi-Phase Clocks. In electronics, these are an important building block used in various circuits. The basis of this thesis is the creation of such clocks and their accuracy when produced. Our fastest clock operates at 16GHz, that's you pushing a swing once every 62.5 ps or 16 billion times every second!

Table of Contents

1	Introduction	1
2	Multiphase Clock Applications	2
2.1	System Description of Hybrid Predistortion	2
2.2	Prior Art	5
2.3	System Specifications	8
3	Circuit Architecture and Design	10
3.1	N Phase Generation	10
3.2	Duty-Cycle Control	10
3.3	Duty Cycle Controller Circuit Simulation	17
3.4	Skew-Compensation Delay	20
3.5	M Phase Generation	22
3.6	N_d and M_d Phase Generation	23
4	Performance of the Clock Generation Circuit	28
4.1	Phase Error Analysis	28
4.2	Power Analysis	33
5	Layout Placement and Interconnect	35
5.1	Stage- N Interconnect Mismatch	35
5.2	Stage- M Interconnect Mismatch	37
6	Conclusions, Performance Comparison and Future Work	39
	References	42

List of Figures

2.1	A two-stage time-interleaved true time delay circuit[1].	4
2.2	An N sample and reconstruction branch on the left and the visualisation of the sample and hold mechanism on the right[1].	4
2.3	Block diagram of an ILRO with phase-correction feedback loop [2].	5
2.4	Block diagram of a DLL generating reference-aligned injection phases for an ILRO [3].	7
2.5	Block diagram of a flip-flop-based ring counter.	8
3.1	Block diagram of the stage- N phase generation chain: one-hot ring counter, per-phase duty-cycle control, and open-loop skew-compensation delay. Finally, the output of stage- N driving stage- M	11
3.2	Schematic of the four-stage one-hot ring counter, with reset and seed-pulse initialisation [4].	12
3.3	(a) Timing diagram of the one-hot ring counter outputs, showing four non-overlapping pulses.(b) Simulation showing the 4 GHz phases with a 25 % duty cycle.	12
3.4	Per-phase duty-cycle control circuit: rising edge set by N_1 , falling edge set by P_1 via an integrator/comparator loop referenced to V_{ref}	13
3.5	Simulated waveforms of the duty-cycle control circuit showing CLK_{IN} , the inverter output, and the PMOS control signal, confirming the absence of contention.	14
3.6	Schematics of the three integrator topologies evaluated: passive RC (A), active RC (B), and charge-pump (C).	14
3.7	Schematic of the charge-pump topology implemented.	16
3.8	StrongARM comparator shown with second gain stage removed.	17
3.9	Simulated duty cycle as a function of reference voltage V_{ref} , demonstrating the control range of the charge-pump integrator loop.	18
3.10	Simulated duty-cycle control loop showing the pulse, integrated ramp, V_{ref} , and PMOS control signal as they propagate around the loop.	19
3.11	Simulated pulse width of one stage- N phase as a function of V_{ref} , demonstrating monotonic duty-cycle control across the target range.	20
3.12	Schematic of the current-starved inverter delay cell. The bias PMOS sets a reference current that is mirrored into NMOS and PMOS sources feeding the starved inverter, ensuring symmetric rise and fall delays.	21
3.13	Waveform output of one stage- N phase delayed between 0 ps and 5 ps. The figure shows that the delay is not uniform across the control range.	22

3.14	Simulated output waveforms of the nine stage- M phases generated by the custom flip-flop ring counter, clocked by one stage- N output. Each phase is active for one period of the 4 GHz stage- N clock, giving a 444.4 MHz output with a 1/9 duty cycle and 250 ps coarse tap spacing.	24
3.15	Schematic of the stage- N_d and stage- M_d clock-generation blocks. Both stages are architecturally identical to stage- N and stage- M , respectively, and are driven from the same 4 GHz reference clock. Physical placement on the reconstruction side of the layout provides a routing distance to the N_d and M_d switches comparable to that of the sampling-side stages.	25
3.16	Simulated stage- N output waveforms for two different seed pulse positions, illustrating the change in phase sequence order.	27
4.1	Simulated phase noise of the stage- N output phases, obtained from PSS/PNoise analysis with interconnect parasitics back-annotated. The phase noise profile of the ADF4382 reference synthesiser at 16 GHz is shown for comparison [5].	30
4.2	Distribution of RMS cycle-to-cycle jitter for stage- N Clk0 across 100 Monte Carlo iterations incorporating process variation and device mismatch, with interconnect parasitics included. The dashed line indicates the 100 fs jitter budget. All samples meet the specification.	31
4.3	Distribution of RMS cycle-to-cycle jitter for stage- N Clk1 across 100 Monte Carlo iterations incorporating process variation and device mismatch, with interconnect parasitics included. The dashed line indicates the 100 fs jitter budget. All samples meet the specification.	31
4.4	Distribution of RMS cycle-to-cycle jitter for stage- N Clk2 across 100 Monte Carlo iterations incorporating process variation and device mismatch, with interconnect parasitics included. The dashed line indicates the 100 fs jitter budget. All samples meet the specification.	32
4.5	Distribution of RMS cycle-to-cycle jitter for stage- N Clk3 across 100 Monte Carlo iterations incorporating process variation and device mismatch, with interconnect parasitics included. The dashed line indicates the 100 fs jitter budget. All samples meet the specification.	32
4.6	Breakdown of simulated power consumption for the stage- N phase generation chain. The duty-cycle controller is the dominant contributor at 1.17 mW, followed by the output buffers at 0.712 mW and the delay cells at 0.492 mW, giving a stage- N total of 2.6 mW.	34
4.7	Breakdown of simulated power consumption for the duty-cycle controller. The dominant contributor here is the buffers at 456 μ W.	34
5.1	Keysight RFpro layout view of the length-matched stage- N interconnects. Routing was carried out within the Ericsson TTD floorplan and length-matched to within 10 μ m of the longest trace to minimise inter-phase skew. .	36
5.2	Place and Route as achieved by digital flow for stage- M	38

List of Tables

5.1	Inter-phase skew relative to stage- N_1 , before and after inclusion of the extracted interconnect parasitics.	35
6.1	Target and achieved specifications for Stage- N	39
6.2	Performance comparison with prior multi-phase generators. FoMs pair each design's single-output integrated RMS jitter with its complete-system power. $\text{FoM} = 10 \log_{10} [(\sigma_t/1 \text{ s})^2 \cdot P/1 \text{ mW}]$, $\text{FoM}_N = 10 \log_{10} [(\sigma_t/1 \text{ s})^2 \cdot P/(N \text{ mW})]$	41

Introduction

In digital systems, clock signals serve as the fundamental mechanism for synchronising operations and determining computational throughput. While increasing clock frequency generally enhances throughput, it also imposes significant demands on power consumption and circuit area, as all interfacing components must operate at an elevated clock speed. Multiphase clocking addresses these challenges by partitioning a single high-frequency clock into multiple lower-frequency phases offset in time, thereby optimising the trade-off between speed, power, and area.

The report begins in Chapter 1 with an overview of the system architecture and its relevant application to programmable broadband circuits. Then, in Chapter 2, a comprehensive review of the state of the art and various methods for generating multiphase clocks, along with their advantages and limitations, is presented, as well as the system specifications. In Chapter 3, the proposed architecture and the rationale for the choice are presented, followed by a detailed circuit-level analysis explaining the methodology and functionality of each sub-circuit. The results of the simulations are presented in Chapter 4, which analyses performance metrics such as phase accuracy and power consumption. In Chapter 5, the layout is presented and details the effect of the interconnects on the system. Finally, a summary of the findings and a conclusion are presented in Chapter 6.

For this thesis, Ericsson AB provided the necessary CAD tools: Cadence Virtuoso and ADE for circuit simulations and finally Layout XL suite for layout. The work was carried out in 22-nm fully depleted silicon-on-insulator CMOS technology. Ericsson AB also provided office space and computers.

Multiphase Clock Applications

2.1 System Description of Hybrid Predistortion

Power amplifiers (PAs) operating in linear classes such as Class A and Class AB exhibit efficiency that increases with output power, and are therefore preferably driven close to saturation to maximise average efficiency [6]. However, maximising efficiency comes at the cost of linearity: as the PA approaches saturation, the gain compresses (AM–AM distortion) and the output phase shifts (AM–PM distortion). This causes unwanted emissions in adjacent frequency bands and degrades modulation accuracy, both of which are critical for modern communication standards. To address this, predistortion is used to linearise the output by applying an intentionally predistorted signal at the input. The predistortion is shaped to be the inverse of the PA’s nonlinearity. If the PA compresses large signals, the predistorter expands them. If the PA introduces phase lag at higher amplitudes, the predistorter applies an equivalent phase lead.

Digital Predistortion (DPD) applies this intentional distortion to the input signal at baseband. DPDs also require a feedback path that samples the PA output through an Analogue-to-Digital Converter (ADC), allowing the correction to adapt as the PA’s behaviour changes over time. This is particularly important for wideband signals, where the PA distortion varies across the signal bandwidth [7]. Distortion also depends not just on the current input but also on the recent history of the input signal, an effect known as *memory effects* [8]. Memory effects are caused by, but not limited to, thermal effects, charge trapping in the semiconductor, and frequency-dependent bias and matching networks, as shown in [8][7]. DPDs compensate for memory effects in the correction algorithm, so that the predistortion depends on both the present input and the recent history of the PA output.

DPDs, however, require advanced signal-processing resources and techniques, such as lookup tables, memory polynomials, or Volterra series models, to account for AM–AM, AM–PM, and memory-related distortion. The more accurate the model, the greater the silicon area, power consumption, and implementation complexity of the digital memory, high-speed DACs, and ADCs. Latency must also be carefully managed, as processing time is required for the digital filtering, polynomial computations, and DAC and ADC interfaces, in addition to significant firmware and algorithm development effort.

Analogue Predistortion (APD) is an alternative technique that typically applies the correction in the RF or intermediate frequency (IF) domain. Applying correction at RF means the predistorter sits close to the PA, reducing the correction path length and latency. APDs typically employ variable gain cells to adjust for AM–AM, phase shifters or vector modulators for AM–PM, and passive networks for matching. Another component relevant to this thesis, described in Section 2.1.1, is a True Time Delay (TTD) circuit,

which provides a constant group delay across the signal bandwidth and is therefore well-suited to wideband memory-effect compensation.

APDs, however, have their own disadvantages. Passive APDs, such as diode predistorters, which provide amplitude-dependent gain expansion, are low-cost and low-power compared with a DPD. Still, they have a limited tuning range and are strongly temperature-dependent. Transistor APDs offer both gain and phase-shift capability and are electronically tunable, but they introduce additional noise, consume more power, and can produce their own distortion without careful design.

A hybrid approach that combines the strengths of digital and analogue predistortion techniques aims to overcome their respective limitations. By using DPD to handle slowly varying, wide-range corrections and an APD with a TTD to handle fast memory-effect compensation, the hybrid system can achieve high linearity while reducing the required order of the digital memory model and, therefore, the digital complexity. This thesis focuses on the implementation of the clocking circuitry that drives the TTD used for memory-effect compensation in such a hybrid PA linearisation system.

2.1.1 True Time Delay

Using a True Time Delay (TTD) circuit is a practical approach to achieving adjustable phase shifting in the analogue domain. It operates by sampling an incoming signal and re-emitting it after a predetermined delay. Each TTD block consists of a sample-and-hold circuit that samples the signal, stores it for a programmed interval, and then releases it to the following stage. The delay interval is chosen to span the PA's dominant memory time constants. The delayed signal is combined with the undelayed input within the predistorter, producing a correction signal whose response depends on both the present input and its recent history, which is then applied to the PA. TTDs produce a flat group delay across the signal bandwidth, a feature necessary for PA predistortion, since any variation in group delay translates directly into uncompensated distortion.

To increase the achievable delay range for larger memory-effect compensation, and to reduce group-delay ripple, a two-stage, time-interleaved TTD is employed, as shown in Figure 2.1 with the structure of the N-Path shown in Figure 2.2. The two stages are driven by separate clock-generation circuits operating at a fixed frequency ratio of $f_M = f_N/9$, where f_N and f_M are the sampling rates of the first stage (Stage N) and second stage (Stage M), respectively. Within each stage, the sampling frequency is fixed, and the realised delay is determined by the order in which the sample-and-hold taps are released to the following stage; reordering the release sequence shifts the effective output instant and, therefore, the delay through the cascade. The duty cycle of the clocks driving stage- N is controlled independently of the delay setting. It is used to manage the overlap between adjacent phases at the higher sampling rate, so that successive switches hand over without a change in impedance, thereby preserving linearity.

In the system we present here, our baseband frequency is 800 MHz, and our RF frequency is 5 times that at 4 GHz.

2.1.2 TTD Clock Generation Requirements

In the initial stage, denoted stage N , the switches operate at 4 GHz with a pulse width ranging from 62.5 ps to 125 ps. These pulses undergo duty-cycle adjustment to produce controlled overlap between adjacent phases. Controlled phase overlap is critical for suppressing spurs in the reconstructed signal at the TTD output. To ensure that phase overlap occurs at the correct instants within each cycle, the jitter on the leading edge of each clock phase is constrained to a maximum of 100 fs.




Figure 2.1: A two-stage time-interleaved true time delay circuit[1].




Figure 2.2: An N sample and reconstruction branch on the left and the visualisation of the sample and hold mechanism on the right[1].

In the second stage, denoted stage M , the switches operate at 444.4 MHz with a duty cycle of 11.1 percent with no overlap. Stage M is synchronised such that each M -phase switch activates before the relevant N -phase switch, ensuring that each stage M tap overlaps four stage N taps.

2.2 Prior Art

2.2.1 DLLs

Delay Locked Loops (DLLs) are a commonly cited approach for generating multiphase clocks. A voltage-controlled delay line (VCDL) cascades N cells and is fed by a reference clock. This produces N evenly spaced phases at the reference frequency f_s , with adjacent taps separated by $1/(Nf_s)$. The DLL feedback loop adjusts the cell delay such that the total delay across all N cells equals exactly one reference period of f_s . All output phases are at the reference frequency f_s , since the DLL performs neither multiplication nor division. However, on their own, DLLs tend to exhibit higher inter-tap phase error. This arises from delay-cell mismatch and duty-cycle distortion along the chain, which accumulate from cell to cell. The most significant contributor, however, is that the feedback constrains only the total delay, sampled at the endpoint of the line, and does not correct the mismatch between the intermediate taps.

2.2.2 Injection Locked Ring Oscillator

Injection Locked Ring Oscillators (ILROs), shown in Figure 2.3, are a popular architecture for generating multiple phases in the analogue domain. They consist of a ring oscillator whose frequency and phase are controlled by an externally injected signal.




Figure 2.3: Block diagram of an ILRO with phase-correction feedback loop [2].

This injected signal forces the free-running oscillator to synchronise with it, provided that the signal is sufficiently strong and its frequency is close enough to the ring oscillator's natural oscillation frequency. In a differential ring oscillator with N inverting

stages, the propagation delay of each stage produces a phase-shifted output tap, with adjacent phases separated by approximately $2\pi/N$ for an ideal oscillator, providing N evenly spaced clock phases. The free-running frequency of the ring oscillator is set by the per-stage delay and is given by,

$$f_{\text{osc}} = \frac{1}{2Nt_d} \quad (2.1)$$

where N is the number of inverting stages and t_d is the propagation delay of a single stage.

The strength of the injected signal is crucial to determining the ILRO's locking range, as described by Adler's equation.

$$\Delta\omega_{\text{lock}} = \frac{\omega_0}{2Q} \cdot \frac{I_{\text{inj}}}{I_{\text{osc}}} \quad (2.2)$$

where ω_0 is the free-running oscillation frequency, Q is the oscillator quality factor, and I_{inj} and I_{osc} are the effective injection and oscillator signal amplitudes at the injection node.

$$Q_{\text{eff}} = \frac{\omega_0}{2} \left| \frac{d\phi}{d\omega} \right| \quad (2.3)$$

The injection of a signal into the ring oscillator forces the ring to adjust its phase to minimise the error between itself and the injected signal. At steady state, $d\theta/dt = 0$, where θ is the phase difference between the injection signal and the oscillator. This mechanism aligns the ILRO as a whole to the injection signal but does not correct the mismatch between stages. To reduce the overall phase error, an ILRO can be paired with a feedback network consisting of a phase detector, a charge pump, and a loop filter. This loop compares the reference clock to the ILRO output and adjusts accordingly. The oscillator's free-running frequency is such that the bulk phase error between the reference and the ILRO is continuously corrected. Operating on a single control variable, the loop suppresses low-frequency phase drift and in-band jitter relative to the reference. Still, an ILRO does not by itself correct the static phase mismatch between individual stages of the ring because the injection occurs only at one tap [2][9].

ILROs are widely used owing to their relative simplicity and low power consumption, with figures of merit often below 0.2 mW/GHz per phase [2].

2.2.3 DLL + Injection Locked Ring Oscillator

A widely adopted improvement pairs a DLL with an ILRO, shown in Figure 2.4. In this configuration, a DLL generates reference-aligned phase taps, each of which is then injected into the ILRO, which produces the final output clock phases. The motivation is that each block, taken alone, is strong in one respect and weak in the other, and the two weaknesses are complementary.

A standalone DLL produces low-jitter phases, since its delay line is locked directly to the reference clock and does not accumulate phase noise across stages in the way a free-running oscillator does. Its phase accuracy, however, is limited as each delay cell experiences different input and output amplitudes, producing a static per-tap phase error analogous to AM-to-PM conversion [3]. A standalone ring oscillator has the opposite character. Its strong injection-locked feedback gives it high phase accuracy. But, on its own, the absolute phase positions and jitter are governed by the oscillator's dynamics — free-running frequency, injection strength, delay-cell mismatch, and layout asymmetry — rather than by a clean reference.

Combining the two allows each to correct the other. The DLL, locked to the reference, sets the system's jitter. The ILRO, driven by the DLL's taps, applies its strong feedback to correct the per-tap phase errors of the DLL without significantly degrading the jitter it has already established [3]. The result is a hybrid in which the jitter is inherited from the DLL and the phase accuracy from the ring, achieving a lower total output phase error than either block could reach alone.




Figure 2.4: Block diagram of a DLL generating reference-aligned injection phases for an ILRO [3].

2.2.4 Flip-Flop Dividers

In the digital domain, frequency division is commonly achieved using Flip-Flops (FFs). In particular, a one-hot ring counter, shown in Figure 2.5, generates N non-overlapping phases, each with a $1/N$ duty cycle, by dividing a higher-speed reference clock.

The output frequency and phase spacing are given by:

$$f_{\text{out}} = \frac{f_{\text{in}}}{N} \quad (2.4)$$

$$\Delta t = \frac{T_{\text{out}}}{N} = T_{\text{in}}, \quad \Delta \phi = \frac{2\pi}{N} \quad (2.5)$$

where N is the number of FFs. In [4], the output is re-timed with the reference clock so that the output edge is defined by a clean reference rather than the accumulated divider delay, reducing jitter at the final output. Compared to ILRO-based generators, these circuits are simple to design and can often be implemented using standard cells from the technology PDK, simplifying physical design and reducing development time.

The limiting factors are the input clock rate, which must be N times the output frequency, and the combined clk-to-Q and setup time of the FFs, which must fit within one input clock period.

$$f_{\text{in}} < \frac{1}{t_{\text{cq}} + t_{\text{su}}} \quad (2.6)$$

where t_{cq} is the clk-to-Q delay of one FF and t_{su} is the setup time of the next FF in the chain, with additional margin required in practice for clock skew and jitter. This sets the technology-dependent upper bound on f_{in} and therefore on N .




Figure 2.5: Block diagram of a flip-flop-based ring counter.

In addition, the one-hot ring counter requires explicit initialisation: a reset signal and an initial seed pulse are needed. At the device level, mismatches between FFs cause clk-to-Q variation, and the high-speed clock tree distributing the reference to all FFs introduces additional skew, which is exacerbated by asymmetric capacitive loading at the FFs' clock inputs. Unlike a free-running ILRO, the FF divider does not accumulate jitter around a feedback loop, since each output edge is triggered directly by an input clock edge; the dominant jitter contributor is therefore the input clock itself [4]. Finally, frequency division by N also reduces the input-referred phase noise by $20 \log_{10}(N)$ dB at the divider output [10].

The rationale for selecting a one-hot ring counter over the more common Johnson counter is discussed in Chapter 3.1.

2.3 System Specifications

This section summarises the design specifications, which were derived from internal system simulations at Ericsson and refined during the early phase of the thesis work.

The design must satisfy three primary specifications:

1. A maximum cascaded delay of approximately 2.25 ns.
2. Non-overlapping output pulses by default, with controlled overlap as an optional mode.
3. Configurable tap-release ordering to support programmable delay sequences.

The first specification, the maximum cascaded delay, is realised by the sum of a coarse and a fine component,

$$t_{\text{delay}} = t_{\text{coarse}} + t_{\text{fine}} = \frac{m_{\text{delay}}}{f_N} + \frac{n_{\text{delay}}}{N \cdot f_N}, \quad (2.7)$$

where f_N is the stage- N sampling frequency, N is the number of phases generated within stage- N , $M = 9$ is the stage- M divider ratio (so that $f_M = f_N/M$ as defined in Section 2.1), and $m_{\text{delay}} \in \{0, 1, \dots, M-1\}$ and $n_{\text{delay}} \in \{0, 1, \dots, N-1\}$ are the integer tap indices that select the coarse and fine delay steps respectively.

The 16 GHz input reference clock is divided by four at stage- N , giving $f_N = 4$ GHz and a fine delay step of $1/(N \cdot f_N) = 62.5$ ps. Stage- M provides $M = 9$ coarse taps spaced by the stage- N period, $1/f_N = 250$ ps, so that the full coarse cycle repeats at $f_M = f_N/M \approx 444.4$ MHz. The maximum realisable delay is therefore

$$t_{\text{delay, max}} = \frac{M-1}{f_N} + \frac{N-1}{N \cdot f_N} \approx 2.19 \text{ ns.} \quad (2.8)$$

The maximum index is $M-1$ rather than M because incrementing the coarse counter past $M-1$ wraps the ring counter back to its initial state, giving zero net delay relative to the reference. The achievable maximum is therefore a structural consequence of the modular ring-counter architecture rather than a design margin, and the resulting 2.19 ns satisfies the system-level requirement of approximately 2.25 ns.

The second specification requires non-overlapping pulses as the default operating mode. This is consistent with the wider TTD literature, in which non-overlapping clocks are typically used to drive the sampling switches [11][12]. As a requirement from Ericsson, the design is initialised with non-overlapping pulses. Then it applies controlled overlap to suppress output spurs arising from the discontinuous sampling inherent in non-overlapping operation.

The third specification requires that the tap-release order be directly programmable. Since the realised delay through the cascade is determined by the order in which taps are released onto the following stage (Section 2.1.1), this programmability translates directly into delay control without additional multiplexing hardware.

Together, the second and third specifications motivate the use of a one-hot ring counter rather than an DLL or ILRO-based generator, as discussed in the architectural comparison in Section 2.2. DLLs/ILROs and their derivatives produce a 50% duty cycle by default, which corresponds to a 50% overlap between adjacent phases; reaching the non-overlapping baseline would therefore require the duty-cycle controller to be continuously active. The one-hot FF approach instead starts from non-overlap and increases the pulse width to the desired value, which is the preferred direction. Furthermore, the FF release sequence is set by the initial state loaded into the ring, making reprogramming straightforward. In contrast, equivalent control in a DLL/ILRO-based generator would require additional multiplexing logic to reorder the output phases.

Two other timing requirements are imposed on the stage- N phase generator, both of which originate from Ericsson system-level studies. As the target system is at present a concept rather than a realised design, these figures represent design targets rather than measured constraints.

The first is a **cycle-to-cycle jitter** limit on the 4 GHz clock of 100 fs maximum, bounding the random timing uncertainty of successive clock edges.

The second is a **maximum inter-phase skew** of 3 ps between adjacent phases of the N -phase output after routing interconnects. Unlike the cycle-to-cycle jitter, this is a deterministic, systematic error arising from a mismatch between the phase-generation paths, and is the quantity targeted by the programmable delay cell (Section 3.4).

The final system-level requirement was an SFDR of 45 dB at an input signal frequency of $f_{\text{in}} = 800$ MHz and $f_s = 4$ GHz.

Circuit Architecture and Design

3.1 N Phase Generation

As concluded in Section 2.3, a one-hot ring counter is used to generate the four stage- N clock phases (the 4 GHz phases driving the 1st stage of TTD switches). The four phases are then passed through a per-phase duty-cycle control block (Section 3.2) and finally through an open-loop delay block (Section 3.4) that compensates for skew at the TTD switch loads. The complete signal chain is shown in Figure 3.1.

The digital blocks are implemented using standard cells from a 22 nm PDK. Each FF output is buffered, length-matched for load balancing, and retimed by the input reference clock. The retiming approach follows [4], where re-clocking the divider outputs by a clean reference suppresses the jitter accumulated through the divider chain.

The counter is implemented by cascading four FFs (Figure 3.2), with a reset and an initial seed pulse to begin the sequence in a prescribed order. The starting phase can be reordered by a multiplexer at the seed input, a feature used later for output reconstruction (Section 3.6). The counter circulates a single 1 around the loop once every four reference cycles, producing four non-overlapping phases as shown in Figure 3.3. Each phase runs at 4 GHz (period 250 ps) with a 25 % duty cycle, giving a pulse width of 62.5 ps.

Flip-flop sizing is chosen to minimise the output rise time within the available power budget. Since the TTD switches are triggered on the rising edge, jitter is evaluated at the rising-edge crossings only (Section 4.1).

3.2 Duty-Cycle Control

The duty-cycle control circuit follows the topology of [13] as shown in Figure 3.4. The first stage of the duty-cycle block is an inverter formed by NMOS N_1 and PMOS P_1 . N_1 produces the rising edge of the output phase, while P_1 sets the timing of the falling edge.

The pulse width of the inverter output is sensed by integrating that output, producing a ramp whose final value is proportional to the pulse width. This ramp is compared against a reference voltage V_{ref} that represents the target duty cycle. When the ramp crosses V_{ref} , the comparator triggers P_1 , pulls the output low and ends the high pulse. Thus, the falling edge —and hence the duty cycle —is set by V_{ref} .

The operational timing requirement is that the complete loop settles within one 4 GHz cycle (250 ps), i.e. before the next rising edge of the input phase. P_1 must finish driving the falling edge before that edge arrives, so that N_1 and P_1 are not simultaneously on, which would cause contention and a large crowbar current, as shown in Figure 3.5.




Figure 3.1: Block diagram of the stage- N phase generation chain: one-hot ring counter, per-phase duty-cycle control, and open-loop skew-compensation delay. Finally, the output of stage- N driving stage- M .




Figure 3.2: Schematic of the four-stage one-hot ring counter, with reset and seed-pulse initialisation [4].




Figure 3.3: (a) Timing diagram of the one-hot ring counter outputs, showing four non-overlapping pulses. (b) Simulation showing the 4 GHz phases with a 25% duty cycle.

3.2.1 Integrator

The integrator converts the inverter's pulse width into a voltage ramp whose slope is proportional to the input level. The comparator then detects the instant at which this ramp crosses V_{ref} , setting the falling edge of the output phase.

The integrator design is subject to three hard requirements. **Speed** is the first constraint: the complete duty-cycle loop must operate within one 250 ps cycle, of which the integrator consumes only a fraction; the remainder is taken by the comparator, P_1 drive, and per-cycle reset. In particular, the integrator must produce a usable ramp during the input pulse width, which is as short as 62.5 ps at 25% duty cycle. **Rail-to-rail input range** is also required, since the integrator input is driven directly by the




Figure 3.4: Per-phase duty-cycle control circuit: rising edge set by N_1 , falling edge set by P_1 via an integrator/comparator loop referenced to V_{ref} .

latch holding the inverter output, which swings rail-to-rail; the integrator must therefore accept the full supply range without clipping or distortion. Finally, **linearity** must be maintained to a sufficient degree that the comparator crossing time is a monotonic and predictable function of V_{ref} , though a strictly linear operation is not required.

Three topologies were considered in turn, as shown in Figure 3.6: a passive RC network (Section 3.2.1), an active op-amp integrator (Section 3.2.1), and a charge-pump integrator (Section 3.2.1). The charge pump was ultimately selected.

Passive RC Integrator

The simplest topology is a passive RC network, in which the input pulse charges a capacitor through a resistor. Simulations confirmed that with suitably sized R and C , the network can produce a ramp within the 62.5 ps pulse-width window. The output is also monotonic, satisfying the linearity requirement in principle.

However, the RC output follows an exponential charging curve rather than a linear ramp:

$$V(t) = V_{DD} \left(1 - e^{-t/RC} \right) \quad (3.1)$$

whose slope,

$$\frac{dV}{dt} = \frac{V_{DD}}{RC} e^{-t/RC}, \quad (3.2)$$

decays exponentially across the charging interval. In this design, a higher V_{ref} corresponds to a lower duty cycle, since the comparator triggers earlier in the charging interval. Near V_{SS} , the ramp slope is at its maximum, so the crossing-time sensitivity $\partial t_{\text{cross}} / \partial V_{\text{ref}}$ is low, giving poor resolution at high duty-cycle settings. Near V_{DD} , the slope flattens and the sensitivity becomes large, making low duty-cycle settings disproportionately susceptible to noise on V_{ref} and the supply.

In addition, since the feedback loop controls the falling edge of the output phase, ramp noise manifests as cycle-to-cycle duty-cycle jitter rather than phase jitter. The




Figure 3.5: Simulated waveforms of the duty-cycle control circuit showing CLK_{IN} , the inverter output, and the PMOS control signal, confirming the absence of contention.




Figure 3.6: Schematics of the three integrator topologies evaluated: passive RC (A), active RC (B), and charge-pump (C).

ramp node carries thermal noise of variance kT/C , independent of R . Reducing C to achieve a faster ramp slope, therefore, directly increases duty-cycle jitter at the output. A further constraint arises from the impedance of R : to achieve a suitably steep slope

within the 62.5 ps window, R must be small. However, a small R lowers the impedance of the integrating node, requiring the upstream latch and output buffer to source a proportionally larger charging current, thereby increasing their power budgets and that of the preceding digital cells. For these reasons, the passive RC topology was rejected in favour of an active approach.

Active RC Integrator

The second topology considered was an active RC integrator, in which an op-amp is configured in the standard inverting integrator arrangement, with a resistor at the input and a capacitor in the feedback path. This produces a true linear integral of the input:

$$V_{\text{out}}(t) = -\frac{1}{RC} \int_0^t V_{\text{in}}(\tau) d\tau \quad (3.3)$$

directly addressing the exponential nonlinearity of the passive RC topology. The high input impedance of the op-amp also removes the drive-strength constraint on R , as the op-amp drives the capacitor directly, decoupling the charging current from the preceding latch and buffer, and eliminating the power budget concern identified in Section 3.2.1.

However, the active integrator introduces a demanding speed requirement. To produce a usable output within the 62.5 ps pulse width window, the op-amp must have a unity-gain bandwidth of at least [14]:

$$f_{\text{UGB}} \geq \frac{1}{2 t_w} = \frac{1}{2 \cdot 62.5 \text{ ps}} = 8 \text{ GHz} \quad (3.4)$$

Furthermore, for the output to swing across the full supply range $V_{DD} = 0.8 \text{ V}$ within this window, the required slew rate is:

$$\text{SR} = \frac{V_{DD}}{t_w} = \frac{0.8 \text{ V}}{62.5 \text{ ps}} = 12.8 \text{ V/ns} \quad (3.5)$$

An op-amp simultaneously meeting an 8 GHz unity-gain bandwidth and a 12.8 V/ns slew rate would require a prohibitively large quiescent current, and is impractical to realise within the power constraints of this design. For this reason, the active RC topology was also rejected.

Charge-Pump Integrator

The third topology considered was a charge-pump integrator shown in 3.7, in which a capacitor is charged by a constant current source during the input pulse width. When the input is high, a PMOS current source of value I_P charges the integration capacitor C ; when the input is low, an NMOS switch resets the capacitor to V_{SS} in preparation for the next cycle. The resulting output voltage is a linear ramp:

$$V_{\text{out}}(t) = \frac{I_P}{C} t \quad (3.6)$$

where t is measured from the start of the pulse width window following each per-cycle reset. Since the charging current I_P is held constant by the current source, the ramp slope $dV_{\text{out}}/dt = I_P/C$ is independent of the output voltage, producing a truly linear ramp. This directly resolves the exponential nonlinearity of the passive RC topology identified in Section 3.2.1.

Furthermore, the charge pump requires no high-bandwidth amplifier. The ramp slope is set entirely by the ratio I_P/C , which can be sized to produce a suitable swing within the 62.5 ps pulse width window without the slew-rate and bandwidth constraints that




Figure 3.7: Schematic of the charge-pump topology implemented.

made the active RC topology impractical. For these reasons, The charge-pump topology was selected.

3.2.2 Comparator

To detect the instant at which the charge-pump ramp crosses V_{ref} , a comparator is required. A single-stage op-amp in a closed-loop configuration was first considered, with feedback providing sufficient gain to drive P_1 either directly or via a preceding buffer stage. However, the bandwidth required to settle within the 62.5 ps pulse width window, as derived in Section 3.2.1, demands an op-amp with a unity-gain bandwidth of 8 GHz. At this speed, the simulated quiescent current resulted in a power consumption of approximately 700 μW per phase, which exceeded the available power budget. A StrongARM latch comparator, as shown in Figure 3.8, was therefore selected [15].

The StrongARM topology is a dynamic regenerative comparator: it draws significant current only during the brief evaluation phase, after which the regenerative positive feedback rapidly resolves the output to a valid logic level. Its effective power consumption is therefore proportional to the clock frequency rather than to a continuous quiescent bias, making it well-suited to the per-cycle operation of the duty-cycle loop. Simulations showed that removal of the second gain stage reduced the total comparator propagation delay, thereby extending the achievable duty-cycle range towards lower values. Although the removal of the second gain stage introduces a modest increase in average power consumption, simulations confirm that this increase remains negligible at the operating clock frequency of 4 GHz.

The propagation delay of the comparator, however, is not constant; it varies with the overdrive voltage at the moment the ramp crosses V_{ref} . At higher values of V_{ref} , the




Figure 3.8: StrongARM comparator shown with second gain stage removed.

ramp crosses with a lower overdrive rate, resulting in a longer propagation delay and a corresponding shift in the effective output transition away from the ideal crossing instant. The ideal crossing time is given by:

$$t_{\text{cross}} = \frac{C \cdot V_{\text{ref}}}{I_P} \quad (3.7)$$

This voltage-dependent delay is the primary cause of the nonlinearity observed in Figure 3.9, which also confirms that the loop provides a monotonic control range across the target duty-cycle range. Provided the control characteristic remains monotonic, the nonlinearity itself is not a fundamental limitation; the DAC driving V_{ref} can be calibrated to account for the nonlinear mapping between code and duty cycle, so the achievable accuracy is determined primarily by the DAC resolution, provided the nonlinear characteristic is sufficiently stable across operating conditions.

3.3 Duty Cycle Controller Circuit Simulation

The final output of the loop is shown in Figure 3.10, which illustrates each signal as it traverses the loop: the input clock, the integrated ramp, the reference voltage V_{ref} , and the PMOS control signal at P_1 . Figure 3.11 shows the effect of varying V_{ref} on the pulse width of one stage- N phase.




Figure 3.9: Simulated duty cycle as a function of reference voltage V_{ref} , demonstrating the control range of the charge-pump integrator loop.




Figure 3.10: Simulated duty-cycle control loop showing the pulse, integrated ramp, V_{ref} , and PMOS control signal as they propagate around the loop.




Figure 3.11: Simulated pulse width of one stage- N phase as a function of V_{ref} , demonstrating monotonic duty-cycle control across the target range.

3.4 Skew-Compensation Delay

A programmable delay cell is inserted in each stage- N output path, providing a mechanism to deliberately retard each phase in time. This delay can be applied to compensate for either inter-phase mismatch between the four stage- N phases or routing-line mismatch in the interconnects feeding the stage- M inputs, so that the effective arrival times at the TTD switch loads are aligned to within the required 62.5 ps. Which of these the delay is allocated to, or how it is apportioned between them, depends on the dominant skew contributor observed in the fabricated device. As the measurement of inter-phase skew and interconnect mismatch is outside the scope of this thesis, this allocation is left as an area for future work; the delay cell nonetheless provides the means to correct for either source once the device has been fabricated, measured, and characterised. The required correction is estimated from simulation, as discussed in Section 4.1 and 5, to be no greater than 3 ps.

The architecture chosen for the delay cell is a current-starved inverter, illustrated in Figure 3.12, a topology widely used for voltage-controlled delay generation [16].

A bias PMOS transistor sets a reference current, which is mirrored into both an NMOS and a PMOS current source within the delay stage. These mirrored currents




Figure 3.12: Schematic of the current-starved inverter delay cell. The bias PMOS sets a reference current that is mirrored into NMOS and PMOS sources feeding the starved inverter, ensuring symmetric rise and fall delays.

symmetrically starve the pull-up and pull-down paths of the inverter, ensuring that the rise and fall delays are matched. The bias current I_{bias} is set by an analogue control voltage V_{ctrl} applied to the gate of the bias PMOS, and limits the rate at which the inverter can charge or discharge the load capacitance of the following stage. The propagation delay through the cell is approximately:

$$t_d \approx \frac{C_L \cdot V_{\text{DD}}}{2 I_{\text{bias}}} \quad (3.8)$$

where C_L is the load capacitance at the inverter output, increasing V_{ctrl} reduces I_{bias} and therefore increases the delay, giving a monotonic control characteristic.

The current-starved topology was selected for its simplicity and its suitability for analogue control. The control voltage V_{ctrl} is provided externally per phase via a DAC, allowing independent calibration of each path during testing.

The required skew range was determined by simulating the approximate skew introduced by the layout interconnects, as discussed in Section 4.1 and 5. The transistor sizes and their ratios within the cell set the achievable skew characteristic: the sizing of the bias PMOS and the mirror ratio into the starved inverter determine the range of I_{bias} , and hence the range of delay, attainable across the V_{ctrl} sweep, while the inverter loading capacitance C_L governs the fidelity, or resolution, of the delay step achievable for a given change in control voltage. These were adjusted together to provide sufficient range and resolution to cover the worst-case simulated skew, with the 0 ps to 5 ps range providing headroom above the 3 ps worst-case correction. Figure 3.13 shows the resulting delay of a single stage- N phase swept across the control range. The delay is monotonic but

non-uniform: the change in delay per unit V_{ctrl} is not constant, reflecting the inverse dependence of t_d on I_{bias} in Equation 3.8. This non-linearity is acceptable here because the cell is calibrated to a target delay per phase rather than driven as a linear actuator.




Figure 3.13: Waveform output of one stage- N phase delayed between 0 ps and 5 ps. The figure shows that the delay is not uniform across the control range.

3.5 M Phase Generation

As stated in Section 2.1, the target of phase generation is an interleaved TTD architecture, requiring both a stage- N and a stage- M clock generator. While the stage- M generator has relaxed timing requirements due to its lower operating frequency, the requirement for an architecture that defaults to non-overlapping pulses is the same as for stage- N . Both the stage- N and stage- M counters are clocked from the same 4 GHz reference clock, with buffering applied at the input to each counter to ensure sufficient drive strength.

Because the stage- M clock rate of 444.4 MHz is an order of magnitude lower than that of stage- N , the timing margins are sufficiently relaxed to make a fully synthesised digital implementation feasible. Two approaches were therefore evaluated. The first uses a standard digital flow, in which the one-hot counter is described in Verilog RTL and synthesised to standard cells from the available 22 nm PDK under user-defined timing constraints. The second mirrors the custom design used for stage- N , extended from

four to nine flip-flops to generate the required nine phases. The two approaches are described in the following subsections, with the trade-offs and final selection summarised in Section 3.5.3.

3.5.1 Digital Flow Implementation

The one-hot counter was described in Verilog RTL, synthesised with Cadence Genus, and placed and routed with Cadence Innovus. The timing constraints, written in SDC, were derived from system-level simulations provided by Ericsson. In its simplest form, a one-hot counter is interpreted by the synthesis tool as a shift register, which is the intended topology: a single logic-high token propagates through nine flip-flops, each output being active for one period of the 4 GHz reference clock. This produces a 444.4 MHz output with a 1/9 duty cycle per phase. A synchronous reset initialises the counter to a known single-active-output state, preventing the all-zero lockup condition inherent to shift-register counters.

3.5.2 Custom Implementation

The custom implementation follows the same architecture as the stage- N design described in Chapter 3.1, extended from four to nine flip-flops. The additional five stages increase the division ratio from $N = 4$ to $M = 9$, producing the required 444.4 MHz output phases with 250 ps coarse tap spacing. Since the operating frequency is significantly lower, the flip-flop setup and hold margins are considerably more relaxed than for stage- N . A simulation of the nine output phases is shown in Figure 3.14. As with the stage- N design, initialisation is handled by an external circuit that asserts a synchronous reset and injects a seed pulse, placing the ring in a known one-hot state before normal operation.

Unlike stage- N , duty-cycle control and inter-stage skew compensation are unnecessary for stage- M . The stage- M phases are generated in close physical proximity to the TTD switches they drive. In contrast, the stage- N phases must be routed over a longer distance and are therefore subject to greater interconnect skew. Given the lower operating frequency and short routing distances, any residual skew is small compared with the 250 ps coarse tap spacing and is therefore considered negligible. Furthermore, precise duty-cycle control is not required: the stage- M phases are used solely for coarse tap selection, and any duty-cycle error does not introduce reconstruction spurs in the interleaved TTD architecture, as confirmed by system-level simulations.

3.5.3 Comparison and Selection

The two implementations are evaluated based on layout area, power consumption, phase accuracy, and design effort. The synthesised implementation offers faster design iteration and direct timing closure through Cadence Genus, while the custom implementation provides a predictable topology consistent with the stage- N design methodology. A quantitative comparison based on post-layout results is presented once place-and-route is complete.

3.6 N_d and M_d Phase Generation

For reconstruction, it was decided to duplicate the stage- N and stage- M blocks and route the duplicated signals to the N_d and M_d switches. Two factors motivated this decision.




Figure 3.14: Simulated output waveforms of the nine stage- M phases generated by the custom flip-flop ring counter, clocked by one stage- N output. Each phase is active for one period of the 4 GHz stage- N clock, giving a 444.4 MHz output with a $1/9$ duty cycle and 250 ps coarse tap spacing.

1. **Routing Complexity:** Deriving stage- N_d from stage- N by routing the phases across the full layout and inserting multiplexers at each switch to implement the required sampling-order programmability were also considered; however, based on the estimated routing complexity and available layout area, this was judged impractical and not pursued. Instead, a duplicate stage- N block is placed on the opposite side of the layout, providing a routing distance to the N_d switches comparable to that of stage- N to the sampling switches. The same argument applies to stage- M_d , as described in Section 2.1. Duplicating both blocks incurs additional power and area overhead. Still, this cost was deemed acceptable given the reduc-

tion in routing complexity and the avoidance of the additional skew and corrective circuitry that the multiplexer-based approach would have required.

2. **Architecture:** Since stage- N drives stage- M and the first stage of sampling switches, the N clock phases are already routed to those switches. It is therefore natural to place stage- M in proximity to the switches it drives, with stage- N providing its clock input along the same route. The duplicate pair stages N_d and M_d follow the same arrangement on the reconstruction side of the layout.

The resulting clock-generation arrangement for the reconstruction path is shown in Figure 3.15.




Figure 3.15: Schematic of the stage- N_d and stage- M_d clock-generation blocks. Both stages are architecturally identical to stage- N and stage- M , respectively, and are driven from the same 4 GHz reference clock. Physical placement on the reconstruction side of the layout provides a routing distance to the N_d and M_d switches comparable to that of the sampling-side stages.

3.6.1 Phase Order Programmability

Stage- N_d and stage- M_d are architecturally identical to stage- N and stage- M , respectively. Stage- N and stage- N_d are both driven directly from the same 4 GHz reference clock, while stage- M and stage- M_d are driven by the outputs of stage- N and stage- N_d , respectively. The only circuit-level modifications required are the OR gates needed to give each flip-flop access to a seed pulse line; the only other distinction is physical placement and the independent assignment of the seed pulse at initialisation. The relative delay between the sampling and reconstruction paths is therefore determined entirely by the difference in seed pulse positions between the two counter pairs.

The propagation delay between stage- N and stage- M , and equivalently between stage- N_d and stage- M_d , is governed by the routing parasitics and buffers along the clock path between each pair of stages. Since the two pairs are architecturally identical and similarly placed, these delays are expected to be closely matched. However, a system-level timing requirement, discussed in Section 4 and 5, demands that the stage- M switches open before their corresponding stage- N switches. To satisfy this, a coarse delay block

is inserted into each stage- N output path to retard the N phases by the required margin relative to the M phases.

The ability to program the tap order of stage- N_d and stage- M_d is fundamental to the delay control of the TTD architecture. Without this programmability, the delay would be fixed by the initial state of each counter and could not be adjusted. As described in Chapter 3.1, the one-hot ring counter architecture enables this programmability through the seed pulse mechanism: rather than asserting the seed pulse on a fixed flip-flop at initialisation, it can be directed to any flip-flop in the ring, causing that phase to be the first active output and thereby setting the starting point of the tap sequence. Stage- N_d supports four distinct starting phases, corresponding to the four flip-flops in the ring, while stage- M_d supports nine.

Since both counters are clocked from the same reference clock, they remain mutually synchronised regardless of which flip-flop receives the seed pulse. The two counters may therefore be initialised independently, and the reconstruction delay can be programmed to any of the $N \times M = 36$ tap combinations by selecting the appropriate starting phase for each counter independently.

The seed pulse and reset are external signals, assumed to be driven by control logic in the wider system. A multiplexer or decoder is required to route the seed pulse to the desired flip-flop in each ring. The tap order is set once at startup and cannot be changed dynamically: to select a different delay, the system must be reset and a new seed pulse issued. This constraint is accounted for at the system level when defining the reconfiguration requirements of the TTD architecture. The effect of seed pulse placement on the output phase order is illustrated in Figure 3.16.




Figure 3.16: Simulated stage- N output waveforms for two different seed pulse positions, illustrating the change in phase sequence order.

Performance of the Clock Generation Circuit

4.1 Phase Error Analysis

A principal system requirement is the generation of clock phases with low phase noise. The relevant figure of merit is the phase noise of each stage- N and stage- M output, since its integral over the band of interest gives the RMS jitter on the rising edge, and it is this rising edge that defines the switching instant of the corresponding TTD switch. Timing uncertainty on the edge translates directly into uncertainty in the sampling instant, degrading the linearity of the reconstructed output. This is most critical for stage- N , given its narrow 62.5 ps pulse width.

The phase-error requirement is set by the system-level SFDR target of 45 dB at an input signal frequency of $f_{\text{in}} = 800$ MHz and $f_s = 4$ GHz (Section 2.3). The SFDR floor imposed by RMS jitter σ_t is

$$\text{SFDR} = -20 \log_{10}(2\pi \cdot f_{\text{in}} \cdot \sigma_t), \quad (4.1)$$

which, rearranged for the maximum tolerable RMS jitter, gives

$$\sigma_t \leq \frac{10^{-\text{SFDR}/20}}{2\pi \cdot f_{\text{in}}} = \frac{10^{-45/20}}{2\pi \times 800 \text{ MHz}} \approx 1.1 \text{ ps}. \quad (4.2)$$

4.1.1 Jitter

The design target for the rising-edge jitter was set substantially lower than this, at 100 fs. The large headroom is deliberate. As discussed in Section 5, the routing from the stage- N outputs to the TTD switches involves long interconnects whose parasitics add both jitter and inter-phase skew. In addition, custom transistor-level layout was outside the scope of this work, so the simulated results do not capture the device-level parasitics that such a layout would introduce. Targeting 100 fs ensures that, once these additional layout contributions and the effects of device mismatch and PVT variation are accounted for, the total timing error at the switch input remains well within the 1.1 ps budget of Equation 4.2.

The jitter on each phase is obtained by integrating its phase noise over a specific range. A key advantage of the one-hot ring counter is that phase noise does not accumulate within stage- N . Although each FF gets its data from the preceding stage, all FF's are clocked directly by the 4 GHz reference clock, which retimes every output each cycle. The AND gates further retimes this at each output. The phase noise on a given stage- N output is therefore set only by the reference-clock phase noise and the FF's contribution, unlike DLL or ILRO topologies in which phase noise accumulates with each stage (Section 2.2).

Stage- M does not share this advantage, as each stage- M FF is clocked by its stage- N phase and so inherits that phase's noise in addition to its own contribution.

For random phase noise, the dominant mechanism is thermal noise in the transistor, which disturbs the threshold-crossing time and so contributes to the high-frequency region of the phase-noise spectrum [17][14]. In this design, the noise summary analysis identified the buffers between system blocks as the dominant source of this contribution, rather than the flip-flops or current-starved delay cells. Supply noise contributes another phase-correlated component by modulating every transistor's drive strength simultaneously [18], which was mitigated through layout shielding (Section 5).

To improve simulation fidelity, the input reference clock was modelled using the measured phase noise profile of the ADF4382 synthesiser IC [5], a real-world clock source representative of the conceived system. This allowed the simulation to incorporate a realistic reference-clock phase-noise profile rather than an ideal clock source created by the tool, providing a more accurate estimate of the phase noise observed at the output phases.

4.1.2 Device Mismatch

The extracted RMS jitter was 90 fs, measured with the interconnect parasitics included. This meets the 100 fs design target and sits well within the 1.1 ps system budget of Equation 4.2, leaving margin for the device-level parasitics of a future custom layout as well as the mismatch and PVT contributions discussed above. To assess robustness to process variation and device mismatch, a Monte Carlo analysis of 100 iterations was performed. The resulting jitter distribution is shown in Figures 4.2, 4.3, 4.4, and 4.5. All samples remain within the 100 fs target, with a worst-case value of 98 fs, confirming that the design meets the phase-noise specification across the expected process spread. Power analysis is discussed in Section 4.2.

4.1.3 Interconnect Mismatch

Custom transistor-level layout was not performed within the scope of this work. Instead, the interconnect routing from the stage- N outputs to the TTD switches and stage- M was performed using Cadence Virtuoso Layout Editor, and parasitic extraction was performed on the resulting interconnect netlist, as discussed in Section 5. The extracted parasitics were back-annotated into the simulation, as interconnect parasitics were identified as the dominant contributor to inter-phase skew for stage- N . The results reported below, therefore, reflect post-layout interconnect conditions.

Phase noise was evaluated in the Cadence Virtuoso simulation environment using Periodic Steady-State (PSS) and Periodic Noise (PNoise) analyses, which characterise the steady-state noise response of periodically driven circuits and are the standard method for phase-noise estimation in mixed-signal simulation at Ericsson and in industry. As the stage- N outputs are digital oscillating signals, the PNoise analysis was run in **sampled (jitter)** mode with the measurement referenced to the threshold-crossing of each rising edge. This is the appropriate choice because the phase of a switching digital signal is most sensitive to noise during the transition through the threshold and far less sensitive during the settled high and low portions of the cycle [17]; sampling the edge therefore captures the noise at the instant that most affects the timing uncertainty of the TTD switch. The resulting phase-noise spectrum was integrated over 10 kHz to 1 GHz [3][2] to obtain the RMS jitter, which was then compared against the 100 fs target.

The primary mitigation applied was increasing the transistor width. A larger gate width reduces the thermal-noise current and increases drive strength, thereby reducing the output transition time and lowering the phase noise at the switching threshold.

Maximum-width standard cells were used wherever it was possible to make a trade-off for performance over dynamic switching power dissipation. A noise summary report generated by Cadence Virtuoso was used to identify the dominant noise contributors, and device sizing was adjusted accordingly.

The simulated phase noise is shown in Figure 4.1. Both the input reference clock and the flip-flop chain output are flat at low offset frequencies and roll off at higher offsets. At the low offset integrated (10 kHz), the FF output sits approximately 3 dB below the reference, at roughly -123 dBc/Hz against the reference's -120 dBc/Hz, consistent with partial benefit from the frequency division though not as good as the theory of $20 \log_{10}(N)$ dB describes. At higher offsets, the two curves cross: the reference continues to fall to approximately -157 dBc/Hz, while the flip-flop output falls to approximately -148 dBc/Hz, sitting above the reference. The noise summary report identified the inter-block buffers — those buffering the flip-flop outputs to the duty-cycle control, the duty-cycle control to the delay cells, and the delay cells to the TTD switches and stage- M — as the dominant internal noise contributors across the entire offset range, rather than the flip-flops or delay cells themselves. The close-in output therefore still benefits partially from the frequency division, while at higher offsets the added buffer noise raises the output above the reference.




Figure 4.1: Simulated phase noise of the stage- N output phases, obtained from PSS/PNoise analysis with interconnect parasitics back-annotated. The phase noise profile of the ADF4382 reference synthesiser at 16 GHz is shown for comparison [5].




Figure 4.2: Distribution of RMS cycle-to-cycle jitter for stage- N Clk0 across 100 Monte Carlo iterations incorporating process variation and device mismatch, with interconnect parasitics included. The dashed line indicates the 100 fs jitter budget. All samples meet the specification.




Figure 4.3: Distribution of RMS cycle-to-cycle jitter for stage- N Clk1 across 100 Monte Carlo iterations incorporating process variation and device mismatch, with interconnect parasitics included. The dashed line indicates the 100 fs jitter budget. All samples meet the specification.




Figure 4.4: Distribution of RMS cycle-to-cycle jitter for stage- N Clk2 across 100 Monte Carlo iterations incorporating process variation and device mismatch, with interconnect parasitics included. The dashed line indicates the 100 fs jitter budget. All samples meet the specification.




Figure 4.5: Distribution of RMS cycle-to-cycle jitter for stage- N Clk3 across 100 Monte Carlo iterations incorporating process variation and device mismatch, with interconnect parasitics included. The dashed line indicates the 100 fs jitter budget. All samples meet the specification.

4.2 Power Analysis

A key system requirement and aim of this thesis was to realise a design within a total power budget of approximately 3 mW. Given the number of functional blocks required to support duty-cycle control and skew compensation across all four phases, meeting this target was challenging. The final simulated total power consumption of 2.6 mW for the stage- N chain and 221 μ W for the stage- M chain confirms that the design exceeds this requirement.

4.2.1 Stage- N

As shown in Figure 4.6, the duty-cycle controller is the dominant power consumer within the stage- N chain, followed by the output buffers and the delay cells. The one-hot ring counter itself contributes only 249 μ W to the total, confirming that the core digital divider is not a significant power burden.

Duty-Cycle Controller

The duty-cycle controller is the largest contributor to total power consumption at 1.17 mW, as shown in Figure 4.6 and further expanded on in Figure 4.7. The large transistor widths required in the comparator and charge-pump stages to meet the 4 GHz correction-loop requirement result in substantial gate capacitance. The dynamic power associated with charging and discharging this capacitance in each cycle is the primary driver of power consumption in this block. Further optimisation of transistor sizing within the duty-cycle controller is identified as an area for future work in Section 4.

Output Buffers

The output buffers are the second largest contributor, accounting for 0.712 mW of the total power. The buffers are required to drive the TTD switch loads of 10 fF and must present sufficient drive strength to maintain signal integrity at 4 GHz. As with the duty-cycle controller, the transistor widths required to achieve this drive strength lead to significant dynamic power dissipation at the operating frequency.

Delay Cell

The programmable delay cells contribute 0.492 mW to the total power consumption. Although the current-starved topology is inherently power-efficient due to the limited bias current, the four per-phase instances add up to a non-negligible total. The delay cell power consumption and its relationship to transistor sizing are discussed further in Section 3.4.

4.2.2 Stage- M

The stage- M chain consumes 221 μ W in total, broken down between the flip-flop chain at 106 μ W and the output buffers at 114 μ W. The lower power consumption relative to stage- N reflects the reduced operating frequency of approximately 444 MHz and the absence of duty-cycle control circuitry at this stage.




Figure 4.6: Breakdown of simulated power consumption for the stage- N phase generation chain. The duty-cycle controller is the dominant contributor at 1.17 mW, followed by the output buffers at 0.712 mW and the delay cells at 0.492 mW, giving a stage- N total of 2.6 mW.




Figure 4.7: Breakdown of simulated power consumption for the duty-cycle controller. The dominant contributor here is the buffers at 456 μ W.

Layout Placement and Interconnect

5.1 Stage- N Interconnect Mismatch

The design utilised standard cells instead of a custom transistor-level layout, as described in Section 4.1. Manual layout and extraction were performed for the interconnect routing, which was identified as the primary source of inter-phase skew. The critical nets include those connecting the stage- N outputs, the TTD switches, and the stage- M inputs. Any length mismatch in these traces results in skew at the switch inputs, independent of device matching. Routing was based on a TTD floorplan provided by Ericsson, which established the relative positions of the stage- N and stage- M blocks within the overall system.

The stage- N traces were length-matched to within $10\mu\text{m}$ of the longest trace, and ground shielding was implemented throughout to minimise coupling between adjacent lines. Electromagnetic (EM) analysis was conducted using Keysight RFpro, and the extracted subcircuit was back-annotated into the simulation netlist. The post-layout jitter results presented in Section 4.1 incorporate these extracted parasitics. Figure 5.1 illustrates the resulting layout. The layout in the figure shows an unbalanced tree due to the relative positions of the TTD switches in the provided floor plan.

The stage- N_1 trace was the longest, so the remaining traces were referenced to it when computing inter-phase skew. Table 5.1 summarises the skew before and after the interconnect parasitics are included. Before routing, the maximum skew is 2.02 ps, set solely by device mismatch. With the extracted parasitics back-annotated, it rises to 3.92 ps, exceeding the 1.1 ps phase-error budget set in Section 2.3.

Table 5.1: Inter-phase skew relative to stage- N_1 , before and after inclusion of the extracted interconnect parasitics.

	Pre-Interconnect	Post-Interconnect
Phase 2–1 skew	1.51 ps	2.47 ps
Phase 3–1 skew	2.02 ps	3.30 ps
Phase 4–1 skew	1.74 ps	3.04 ps
Maximum skew	2.02 ps	3.30 ps




Figure 5.1: Keysight RFpro layout view of the length-matched stage- N interconnects. Routing was carried out within the Ericsson TTD floorplan and length-matched to within $10\text{ }\mu\text{m}$ of the longest trace to minimise inter-phase skew.

5.2 Stage- M Interconnect Mismatch

5.2.1 Custom Layout

The stage- M routing is short due to its proximity to the switches and operates at a lower clock rate than the stage- N routing. Consequently, the interconnect mismatch between its phases is negligible and does not constitute a limiting factor.

5.2.2 Digital Flow Place and Route

Static timing analysis identifies the slow-slow, 0.72 V, 125 °C corner as the worst case, where the longest clock-to-output delay is 102.4 ps (3σ) at `clk[0]`. This output is an outlier because Innovus mapped its flip-flop to a slower standard-cell variant from the other eight phases. More significant for the multiphase generator is the inter-phase skew: at this corner, the nine rising-edge arrivals span roughly 12 ps. This figure could not be constrained directly, as the SDC was only able to bound the skew of each phase against the reference clock rather than between phases, leaving inter-phase alignment as an unconstrained outcome of the flow and possibly an area for future work. The final place and route are shown in Figure 5.2.




Figure 5.2: Place and Route as achieved by digital flow for stage- M .

Conclusions, Performance Comparison and Future Work

This work aimed to design a two-stage multiphase clock generator. Table 6.1 presents the original Stage- N targets alongside the achieved results. Table 6.2 compares the proposed design with two injection-locked references from the literature.

Table 6.1: Target and achieved specifications for Stage- N .

Parameter	Unit	Stage- N	
		Target	Achieved
Frequency	GHz	4.0	4.0
Phases (N)	–	4	4
Duty cycle	%	25–50	23–50
Peak phase error ^a	°	1.58	4.75
C2C jitter	fs	≤100	90
Power ^b	mW	1.5	2.6

^a Peak phase error expressed as the largest angular deviation of any output from its ideal 90°-spaced position at 4 GHz ($\theta = 360^\circ f \Delta t$). Target corresponds to $\Delta t = 1.1$ ps; achieved corresponds to $\Delta t = 3.30$ ps.

^b Single Stage- N path only. The fully duplicated system (Stage- N , Stage- N_d , Stage- M , Stage- M_d) consumes 5.64 mW in total.

Stage- N meets the cycle-to-cycle jitter target with margin, and the skew-correction circuit operates as intended. The duty-cycle controller surpasses its specification, achieving a 23–50% range rather than the required 25–50%. The only unmet target is phase error, which reaches 4.75° compared to the 1.58° target; this shortfall is discussed below. Stage- M achieved both its design goals: a phase count of 9 and an output frequency of 444.44 MHz.

Stage- N consumes 2.6 mW, exceeding the 1.5 mW per-path budget implied by the 3 mW system target. This power overrun results from the architectural decision to duplicate Stage- N and Stage- M into the N_d and M_d paths, rather than routing a single block through multiplexers. The alternative approach imposed prohibitive routing complexity, as discussed in Section 3.6. Duplication was therefore a deliberate choice, approximately doubling the active logic but ensuring the required phase count without the infeasible

routing complexity of the multiplexed alternative. The fully duplicated system, comprising Stage- N , Stage- N_d , Stage- M , and Stage- M_d , consumes 5.64 mW in total.

In terms of raw jitter–power figure of merit (FoM), the generator outperforms the injection-locked ring oscillator by approximately 10 dB and is within 0.3 dB of the DLL-assisted design. The power and phase count for this work correspond to a single Stage- N path (2.6 mW, 4 phases), which is directly comparable to a single-block reference design. When the FoM is normalised to phase count, the per-phase efficiency is broadly comparable to the DLL-assisted reference and significantly better than the injection-locked ring oscillator.

The primary limitation of the design is phase error. At 4.75° (3.30 ps at 4 GHz), this value is the highest among the three compared designs and exceeds the 1.1 ps system budget derived from the 45 dB SFDR requirement (Equation 4.2). This outcome reflects the absence of phase-correction feedback, as the flip-flop ring counter maintains a fixed phase relationship and lacks a mechanism to correct skewed edges, unlike the injection-locking and DLL techniques referenced. Additional error arises from interconnect mismatch due to routing constraints and the absence of a custom transistor-level layout, which was beyond the scope of this work. The programmable delay cell (Section 3.4) offers a post-fabrication mechanism to reduce residual skew after characterisation, and a full-custom layout is expected to significantly reduce the interconnect-mismatch contribution. Its main strength is efficiency: at 0.163 mW/GHz per phase, it matches the DLL-assisted reference and is more than twice as efficient as the injection-locked ring oscillator.

Future work should focus on reducing phase error, which remains the primary limitation of the proposed architecture. The current design lacks phase-correction feedback, preventing active compensation for residual skew resulting from routing mismatch and fixed edge relationships. Incorporating a programmable delay cell or another calibration mechanism in future revisions could enable post-fabrication output correction. Implementing a full-custom layout is also expected to enhance device matching and minimise interconnect-induced phase error. Further optimisation of the duplicated Stage- N_d and Stage- M_d paths may reduce overall power consumption while maintaining the required phase count. Finally, silicon validation of the proposed architecture would provide a comprehensive assessment of its practical performance and confirm the simulated trends presented in this work.

Table 6.2: Performance comparison with prior multi-phase generators. FoMs pair each design's single-output integrated RMS jitter with its complete-system power.
 $\text{FoM} = 10 \log_{10} [(\sigma_t/1 \text{ s})^2 \cdot P/1 \text{ mW}]$,
 $\text{FoM}_N = 10 \log_{10} [(\sigma_t/1 \text{ s})^2 \cdot P/(N \text{ mW})]$.

Parameter	This work (Simulated)	[2]	[3]
Technology	22 nm FDSOI	22 nm CMOS	28 nm FDSOI
Configuration	Flip Flop Based	ILRO	DLL + ILRO
Supply	0.8 V	1.1 V	0.8 V
Frequency	4 GHz	8 GHz to 28 GHz	0.2 GHz to 11.7 GHz
Phases (N) ^a	4	8	8
Duty Cycle Control (N)	23%–50%	–	–
RMS jitter σ_t ^b	90 fs	98 fs	38 fs ^c
Integration band	10 kHz–1 GHz	10 kHz–1 GHz	1 kHz–1 GHz
Peak phase error	4.75°	1.5°	3°
System power	2.6 mW ^d @ 4 GHz	26 mW @ 8 GHz	15 mW @ 11.7 GHz
Power/freq	0.65 mW/GHz	3.25 mW/GHz	1.28 mW/GHz
Power/freq/ N	0.163 mW/GHz	0.41 mW/GHz	0.16 mW/GHz
Jitter–power FoM	–256.3 dB	–246.0 dB	–256.6 dB
FoM _{N} per phase	–262.3 dB	–255.1 dB	–265.7 dB

^a This work reports a single Stage- N path, corresponding to 4 phases. Reference designs report 8 phases each; the FoM _{N} normalisation accounts for this difference.

^b Jitter measured at a single Stage- N output, consistent with the per-output figures reported by the references; power covers the single Stage- N path in every case.

^c Integrated over 1 kHz–1 GHz rather than 10 kHz–1 GHz; the wider band captures additional low-offset noise, and this figure is not strictly comparable to the others.

^d Single Stage- N path only (2.6 mW). Fully duplicated system power (Stage- N , Stage- N_d , Stage- M , Stage- M_d) is 5.64 mW.

References

- [1] Snygg, Vilgot and Nilsson, Vejde, “True-Time Delay Cancellers for Full-Duplex,” 2024, Student Paper.
- [2] E. Monaco, G. Anzalone, G. Albasini, S. Erba, M. Bassi, and A. Mazzanti, “A 2–11 ghz 7-bit high-linearity phase rotator based on wideband injection-locking multi-phase generation for high-speed serial links in 28-nm cmos fdsoi,” *IEEE Journal of Solid-State Circuits*, vol. 52, no. 7, pp. 1739–1752, 2017.
- [3] Y. Tian, J. Gu, W. He, S. Liu, R. Yin, X. Shen, L. Kong, H. Xu, and N. Yan, “A 28-nm 8–28-ghz eight-phase clock generator using an injection-locked dual-feedback ring oscillator,” *IEEE Journal of Solid-State Circuits*, vol. 61, no. 1, pp. 47–62, 2026.
- [4] J. Heel, H. S. Bindra, S. M. Louwsma, A. Dezzani, and B. Nauta, “A 12.8-gs/s time-interleaved sub-sampling adc front end with 38-ghz input bandwidth and >39-db snr for 1–32 ghz in 22-nm fdsoi,” *IEEE Journal of Solid-State Circuits*, vol. 60, no. 12, pp. 4321–4335, 2025.
- [5] Analog Devices, “Adf4382 datasheet and product info,” <https://www.analog.com/media/en/technical-documentation/data-sheets/adf4382.pdf>.
- [6] C. S. C., *RF Power Amplifiers for Wireless Communications, Second Edition*. Artech House Publishers, 2006.
- [7] C. N. B. Jose Carlos Pedro, *Intermodulation Distortion in Microwave and Wireless Circuits*. Artech House Publishers, 2003.
- [8] J. Vuolevi, T. Rahkonen, and J. Manninen, “Measurement technique for characterizing memory effects in rf power amplifiers,” *IEEE Transactions on Microwave Theory and Techniques*, vol. 49, no. 8, pp. 1383–1389, 2001.
- [9] M. Raj, S. Saeedi, and A. Emami, “22.3 a 4-to-11ghz injection-locked quarter-rate clocking for an adaptive 153fj/b optical receiver in 28nm fdsoi cmos,” in *2015 IEEE International Solid-State Circuits Conference - (ISSCC) Digest of Technical Papers*, 2015, pp. 1–3.
- [10] B. Razavi, *RF Microelectronics*. Prentice Hall Press, 2011.
- [11] E. Zolkov, A. Madjar, R. Weiss, and E. Cohen, “Analysis and design of n-path true-time-delay circuit,” *IEEE Transactions on Microwave Theory and Techniques*, vol. 68, no. 12, pp. 5381–5394, 2020.
- [12] T. Forbes, B. Magstadt, J. Moody, J. Saugen, A. Suchanek, and S. Nelson, “A 0.2–2 ghz time-interleaved multistage switched-capacitor delay element achieving 2.55–448.6 ns programmable delay range and 330 ns/mm² area efficiency,” *IEEE Journal of Solid-State Circuits*, vol. 58, no. 8, pp. 2349–2359, 2023.

- [13] H.-C. Chow, “Duty cycle control circuit and applications to frequency dividers,” in *ICECS’99. Proceedings of ICECS ’99. 6th IEEE International Conference on Electronics, Circuits and Systems (Cat. No.99EX357)*, vol. 3, 1999, pp. 1619–1622 vol.3.
- [14] B. Razavi, *Design of Analog CMOS Integrated Circuits*. McGraw-Hill Education - Europe, 2016.
- [15] —, “The strongarm latch [a circuit for all seasons],” *IEEE Solid-State Circuits Magazine*, vol. 7, no. 2, pp. 12–17, 2015.
- [16] M. Svoboda and P. Horsky, “High-matching current-starved inverter for two-dimensional vernier time-to-digital converter,” in *2025 35th International Conference Radioelektronika (RADIOELEKTRONIKA)*, 2025, pp. 1–5.
- [17] A. Hajimiri and T. Lee, “A general theory of phase noise in electrical oscillators,” *IEEE Journal of Solid-State Circuits*, vol. 33, no. 2, pp. 179–194, 1998.
- [18] M. S. Illikkal, J. N. Tripathi, and H. Shrimali, “Analysing the impact of various deterministic noise sources on jitter in a cmos inverter,” in *2019 6th International Conference on Signal Processing and Integrated Networks (SPIN)*, 2019, pp. 208–211.
- [19] medwatt, “circuitikz_symbols: Circuitikz symbols for inkscape,” https://github.com/medwatt/circuitikz_symbols, 2023.