

Lock-Time Reduction in Sub-Sampling Phase-Locked Loops Using Adaptive VCDL Gain

A PROJECT REPORT

SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE AWARD OF THE DEGREE
OF

MASTER OF TECHNOLOGY
IN
SIGNAL PROCESSING AND DIGITAL DESIGN (SPDD)

Submitted by

Gurvinder Arya (2K24/SPD/03)

Under the supervision of

Dr. Sudipta Majumdar (Associate Professor, DTU)



**ELECTRONICS AND COMMUNICATION ENGINEERING
DELHI TECHNOLOGICAL UNIVERSITY**

(Formerly Delhi College of Engineering)
Bawana Road, Delhi 110042

MAY, 2026

i

**DEPARTMENT OF ELECTRONICS AND COMMUNICATION
ENGINEERING
DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Bawana Road, Delhi-110042**

ACKNOWLEDGEMENT

I wish to express my sincerest gratitude to **Dr. Sudipta Majumdar** ma'am for her continuous guidance and mentorship that she provided to me throughout the course of this project work. She showed me the path to achieve my objectives by explaining all the tasks to be done and made me understand the importance of this project as well as its industrial relevance. She was always ready to help me and clarify my doubts regarding any hurdles encountered during this work. Without her constant support and motivation, this project report would not have been successful.

Gurvinder Arya

Place: Delhi

Gurvinder Arya

Date: 03/06/2026

DEPARTMENT OF ELECTRONICS AND COMMUNICATION
ENGINEERING
DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Bawana Road, Delhi-110042

CANDIDATE DECLARATION

I, Gurvinder Arya, Roll No. 2K24/SPD/03, student of M.Tech (Signal Processing and Digital Design), hereby declare that the project Dissertation titled "Lock-Time Reduction in Sub-Sampling Phase-Locked Loops Using Adaptive VCDL Gain" which is submitted by me to the Department of Electronics and Communication Engineering, Delhi Technological University, Delhi in partial fulfilment of the requirement for the award of degree of **Master of Technology**, is original and not copied from any source without proper citation. This work has not previously formed the basis for the award of any Degree, Diploma Associateship, associateship, fellowship, or other similar title or recognition.

Gurvinder Arya

Place: Delhi

Gurvinder Arya


Date: 03/06/2026

DEPARTMENT OF ELECTRONICS AND COMMUNICATION
ENGINEERING
DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Bawana Road, Delhi-110042

CERTIFICATE

I hereby certify that the Project Dissertation titled "Lock-Time Reduction in Sub-Sampling Phase-Locked Loops Using Adaptive VCDL Gain" which is submitted by Gurvinder Arya, Roll No.-2K24/SPD/03, Electronics and Communication Engineering, Delhi Technological University, Delhi in partial fulfilment of the requirement for the award of the degree of Master of Technology, is a record of the project work carried out by the student under my supervision. To the best of my knowledge this work has not been submitted in part or full for any Degree or Diploma to this University or elsewhere.

Place: Delhi


Dr. Sudipta Majumdar

Date: 03/06/2026

SUPERVISOR



DELHI TECHNOLOGICAL UNIVERSITY
(Formerly Delhi College of Engineering)
Shahbad Daulatpur, Main Bawana Road, Delhi-42

PLAGIARISM VERIFICATION

Title of the Thesis: Lock Time Reduction in Sub-Sampling Phase-Locked Loops Using Adaptive VCDL Gain

Total Pages 60

Name of the Scholar Gurvinder Arya

Supervisor (s)

(1) Dr. Sudipta Majumdar

Department: **Electronics & Communication**

This is to report that the above thesis was scanned for similarity detection. Process and outcome is given below:

Software used: Turnitin _____ Similarity Index: 4%, Total Word Count: 20,733

Date: 03/06/2026

Gurvinder Arya

Candidate's Signature

S. Majumdar

Signature of Supervisor(s)

ABSTRACT

ABSTRACT

Sub-sampled phase-locked loops employing ring oscillators provide several benefits for modern integrated circuits, including a small size, broad range of tunability, and easy transfer between nodes. They do not suffer from the added noise resulting from a feedback divider either. Unfortunately, there is an important drawback to consider. At high output frequencies, the loop gain can be too high and cause loss of stability on the intrinsic feedback path. Though the traditional remedy, adding a voltage-controlled delay line, will guarantee stability, it will cause the loop to take a long time to settle, oscillating with each iteration of the phase error until it finally settles down.

This project report takes a different view of the delay line. Rather than treating its gain as a fixed design constant, it is treated as a quantity that can be varied on the fly according to how large the phase error currently is. A smooth scheduling law keeps the gain high while the loop is still acquiring, which suppresses the oscillatory behaviour and lets the error fall away rapidly, and then relaxes the gain gently back to its nominal value once the loop has nearly settled. A central concern throughout is that this changing gain must never compromise stability, and a Lyapunov-based argument is developed to show that the loop remains stable for every admissible phase error as the schedule moves through its range.

The proposed loop is studied through behavioural circuit simulation in Cadence Virtuoso and Spectre, using models of the sub-sampling detector, the delay line, the adaptive scheduler, the ring oscillator, and an automated lock-time detector. The results show that the adaptive scheme locks substantially faster than the conventional

fixed-gain design, with the advantage becoming more pronounced as the size of the disturbance grows, and that this speed-up is obtained without any penalty to the settled noise performance and without adding extra detectors, filters, or injection paths. In short, fast and clean acquisition is recovered for a ring-oscillator sub-sampling loop simply by varying one parameter that the loop already contains.

Keywords: sub-sampling PLL, SSPLL, voltage-controlled delay line, VCDL, adaptive gain scheduling, discrete-time stability, lock time, ring VCO, Lyapunov stability, Cadence Virtuoso, Spectre simulation.

Contents

Acknowledgement	i
Candidate Declaration	ii
Certificate	iii
Abstract	iv
List of Tables	ix
List of Figures	x
List of Abbreviations	xiii
1 Introduction	1
1.1 Introduction	1
1.2 Problem Statement	2
1.3 Scope of Research	2
1.4 Research Hypothesis	3
1.5 Objectives	3
1.6 Organization of the Report	4
2 Literature Review	5
2.1 Background	5
2.1.1 Phase-Locked Loop Fundamentals	5
2.1.2 Sub-Sampling Phase-Locked Loops	5
2.1.3 Voltage-Controlled Delay Lines for Bandwidth Extension	6
2.2 Critical Review of Existing Work	6
2.2.1 The Base Paper: Ghahramani et al. (2025)	6
2.2.2 Sub-Sampling PLL Architectures and Variants	6
2.2.3 Gain Scheduling and Adaptive Bandwidth Techniques	7
2.2.4 Discrete-Time and z -Domain PLL Modelling	7
2.2.5 Lyapunov Methods for PLL Stability	8
2.2.6 Ring Versus LC Oscillators in Sub-Sampling Loops	8
2.2.7 Phase-Noise and Jitter Modelling in Sampled Loops	8
2.2.8 Fast-Lock Strategies of Different PLL Families	9
2.2.9 Fast-Lock Strategies Across PLL Families	9
2.2.10 Sampled Feedback and the Deadbeat Idea	9
2.2.11 Where Fast Locking Actually Matters	10
2.3 Identification of Research Gap	10
3 System Design and Methodology	12

3.1	Simulation Platform and Materials	12
3.1.1	Cadence Virtuoso and Spectre Environment	12
3.1.2	VerilogA Behavioural Models	12
3.1.3	System Parameters	13
3.2	Type-I SSPLL Architecture	13
3.2.1	Block Diagram	13
3.2.2	Continuous-Time Open-Loop Transfer Function	14
3.2.3	Integer- N Charge-Pump PLL: Comparison Architecture	16
3.3	Discrete-Time Model	17
3.3.1	Derivation of the z -Domain Pole	17
3.3.2	Pole Location as a Function of G_F	18
3.4	Adaptive Gain Scheduling Law	18
3.4.1	Design Rationale	18
3.4.2	Scheduling Function	20
3.4.3	Choice of $G_{F,\max}$	20
3.5	Lyapunov Stability Proof	21
3.5.1	Statement	21
3.5.2	Proof	21
3.6	Noise Bandwidth and Jitter Analysis	21
3.6.1	Noise Bandwidth	21
3.6.2	Phase-Noise Sources and Transfer Functions	23
3.6.3	Jitter Comparison Model	25
3.7	Parameter Selection and Design Trade-offs	26
3.7.1	Fixing the Loop-Gain Operating Point	26
3.7.2	Choosing the Schedule Transition	27
3.8	Simulation Methodology and Reproducibility	28
3.8.1	Testbench Construction and Solver Configuration	28
3.8.2	Sweep Strategy and Lock-Time Extraction	28
3.8.3	VCDL Transfer Function and Realisation	29
3.9	The Sub-Sampling Detector and Its Linear Range	29
3.10	Why Continuous-Time Intuition Misleads Here	30
3.11	Summary	31
4	Simulation Results and Discussion	32
4.1	Simulation Environment and Test Configuration	32
4.1.1	Spectre Netlist and Solver Settings	32
4.1.2	Lock-Time Measurement Methodology	32
4.1.3	Phase-Step Injection	33
4.2	Baseline SSPLL Characterisation	33
4.2.1	Phase-Error Trajectory: $\pi/4$ Step	33
4.2.2	Phase-Error Trajectory: $\pi/2$ Step	36
4.2.3	Baseline Lock-Time Summary	36
4.3	Proposed Adaptive-GF SSPLL Results	36
4.3.1	Phase-Error Trajectory: $\pi/4$ Step	36
4.3.2	Phase-Error Trajectory: $\pi/2$ Step	37
4.4	Comparative Performance Analysis	39

4.4.1	Lock-Time Comparison	39
4.4.2	DT Pole Trajectory During Settling	39
4.4.3	Noise Bandwidth Equivalence at Lock	40
4.5	Sensitivity and Robustness Analysis	40
4.5.1	Dependence on the Schedule Transition ϕ_0	40
4.5.2	Sensitivity to the Maximum Gain $G_{F,\max}$	41
4.5.3	Robustness to Phase-Step Amplitude	41
4.5.4	Comparison with Published Sub-Sampling Designs	42
4.5.5	Reading the Noise-Bandwidth Result	43
4.5.6	Practical Design Guidelines	44
4.5.7	An Energy-Decay View of the Stability Margin	44
4.5.8	Consolidated Comparison and What It Shows	45
4.5.9	Limitations of the Behavioural Study	45
4.6	Discussion	46
4.6.1	Interpretation of Lock-Time Results	46
4.6.2	Lyapunov Condition Under Spectre Conditions	46
4.6.3	VCDL Physical Feasibility	46
4.6.4	Summary of Simulation Results	47
5	Conclusion and Recommendations	48
5.1	Conclusions	48
5.2	Limitations	49
5.3	Recommendations for Future Research	49
5.4	Summary	50
	References	51
A	VerilogA Behavioural Models	53
A.1	Sub-Sampling Phase Detector	53
A.2	Ring Voltage-Controlled Oscillator	54
A.3	Adaptive VCDL	54
A.4	Adaptive G_F Scheduler	55
A.5	Lock-Time Detector	55
A.6	Spectre Netlist (Key Instance Lines)	56

List of Tables

2.1	Summary of reviewed works and coverage of identified research gaps.	11
3.1	VerilogA behavioural modules used in the Spectre testbench.	13
3.2	SSPLL system parameters used in all Spectre simulations.	13
3.3	DT pole z_{pole} as a function of VCDL gain G_F	18
3.4	Phase-noise floor assumptions used in the analytical noise model. . . .	24
3.5	Jitter comparison across configurations and noise sources. SSPD values are analytical (exact DT NBW); VCO values are Cadence Spectre transient-noise simulation.	26
4.1	Spectre simulation configurations for the baseline and proposed SSPLL.	32
4.2	SSPD output voltage $\phi_{e,v}$ at each reference cycle, baseline SSPLL, $\pi/4$ phase step. Lock threshold = 5 mV.	33
4.3	SSPD output voltage $\phi_{e,v}$ at each reference cycle, baseline SSPLL, $\pi/2$ phase step. Lock threshold = 5 mV.	36
4.4	SSPD output voltage $\phi_{e,v}$ at each reference cycle, proposed adaptive SSPLL, $\pi/4$ phase step. Lock threshold = 5 mV.	37
4.5	SSPD output voltage $\phi_{e,v}$ at each reference cycle, proposed adaptive SSPLL, $\pi/2$ phase step. Lock threshold = 5 mV.	37
4.6	Spectre-measured lock times for the baseline and proposed SSPLL. . .	39
4.7	Instantaneous DT pole during the adaptive settling transient, proposed SSPLL, $\pi/4$ phase step.	40
4.8	Complete performance summary.	47

List of Figures

- 3.1 Signal flow of the adaptive VCDL-aided SSPLL. The SSPD samples the VCO output at each reference clock edge, producing phase-error voltage $\phi_{e,v}$. The adaptive G_F scheduler recomputes the VCDL gain each cycle; the VCDL translates the corrected phase error into the control voltage v_{ctrl} that drives the ring VCO, closing the loop. 14
- 3.2 CT open-loop Bode plot: base paper (Ghahramani et al. 2025, $K_{VCO} = 241$ MHz/V, $G_F = 3.0$, dash-dot) and this work ($K_{VCO} = 115$ MHz/V) for $G_F = 3.0$, $G_{F,crit} = 5.226$, and $G_{F,max} = 5.509$ rad/V. The shaded region ($f > f_{REF}/2$) marks the frequency range where the CT approximation is invalid for our design ($\alpha = 3.61 \gg 1$). All our unity-gain crossovers fall inside this region, motivating the exact discrete-time analysis of Section 3.3. 15
- 3.3 Left: CT phase margin vs. normalised bandwidth ω_{UG}/ω_{REF} for five VCDL gain values (base-paper style, Ghahramani 2025 Fig. 4b). Our design operates at $u \approx 1.7$ – 3.2 , far outside this range. Right: correct DT stability margin $1 - |z_{pole}|$ vs. G_F for our design. Marked points correspond to $G_{F,ss} = 3.0$ (baseline), $G_{F,crit} = 5.226$ (deadbeat), and $G_{F,max} = 5.509$ rad/V (proposed maximum). 16
- 3.4 Integer- N charge-pump PLL (comparison baseline). The $\div N$ divider in the feedback path scales reference-path noise by N^2 . At $N = 24$ this represents a 27.6 dB noise penalty relative to an architecture without a divider. 16
- 3.5 Discrete-time signal flow of the VCDL-aided SSPLL. 17
- 3.6 VCDL aided SSPLL discrete time pole locus. as G_F grows, the pole shifts on the real axis. At $z = -0.4451$ the loop is steady but oscillatory; at $z \approx 0$ the loop is one-shot settling. 19
- 3.7 Extended DT pole map The pole locus (solid line) encompasses the entire range from the no-VCDL unstable point ($z = -2.613$, $G_F = 0$) to the base-paper operating point ($z = -0.4451$, open circle), the deadbeat point ($z = 0$, square), and the recommended maximum ($z = +0.038$, triangle). The curving arc is the adaptive transition, as the scheduler reduces G_F in response to the decaying error, moving the pole from $z \approx 0$ to $z = -0.4451$ during each lock acquisition transient. The unit-circle stability border is the dashed circle 19

3.8	Adaptive gain-factor schedule. At the design step $ \phi_e = \pi/4 = 0.785$ rad, the schedule evaluates to $G_F = 5.218$ rad/V (within 0.2% of $G_{F,\text{crit}}$), placing the DT pole at $z = +0.037$. As the error decays toward zero, G_F returns smoothly to $G_{F,\text{ss}} = 3.0$ rad/V.	20
3.9	Squared VCO noise transfer function $ H_{\text{NTF}}(f) ^2$ vs. offset frequency. Baseline ($G_F = 3.0$, $z_{\text{pole}} = -0.4451$, solid): near-Nyquist amplification to 6.78 (8.3 dB at 49.5 MHz) inflates the effective noise bandwidth to 130.2 MHz. $G_{F,\text{crit}}$ ($z_{\text{pole}} = 0$, dashed): all-pass NTF with unity gain across the entire Nyquist band, reducing NBW to 50.0 MHz. Peak label indicates the dB value at the Nyquist edge for the baseline.	22
3.10	Noise bandwidth (left) and RMS jitter (right) for the three SSPLL configurations: baseline / base-paper ($G_F = 3.0$), $G_{F,\text{crit}} = 5.226$ (analytical jitter floor, marked $\dagger$), and proposed adaptive (steady state = $G_{F,\text{ss}} = 3.0$). NBW values from the exact DT formula Eq. (3.12); jitter from $J_{\text{rms}} = \sqrt{S_\phi \cdot \text{NBW}} / (2\pi f_{\text{VCO}})$. Percentage annotations show the NBW and jitter reduction from baseline to $G_{F,\text{crit}}$	23
3.11	Phase-noise injection points in the adaptive SSPLL. The four sources n_{REF} , n_{SSPD} , n_{VCDL} , and n_{VCO} (red arrows) enter the loop at distinct nodes. Reference and SSPD noise pass through the low-pass function $H_{\text{REF}}(\omega) = L/(1+L)$; VCO and VCDL noise are shaped by the high-pass function $H_{\text{VCO}}(\omega) = 1/(1+L)$	24
3.12	Free-running ring VCO phase noise ($S_\phi = 6320/\Delta f^2$ rad ² /Hz). The -20 dB/decade slope is annotated; the reference point at $\Delta f = 1$ MHz gives $\mathcal{L}(1 \text{ MHz}) = -85$ dBc/Hz. This trace is used in all analytical jitter estimates of Table 3.5 and in the closed-loop phase-noise model of Figure 3.13.	24
3.13	Closed-loop phase noise (VCO contribution only) for the baseline ($G_F = 3.0$, solid) and the $G_{F,\text{crit}} = 5.226$ configuration (dashed), together with the free-run VCO noise (dotted) and estimated reference floor (dash-dot). Vertical markers indicate the noise bandwidth limits: $\text{NBW}_{\text{ss}} = 130.2$ MHz (baseline) and $\text{NBW}_{\text{crit}} = 50.0$ MHz ($G_{F,\text{crit}}$). The proposed adaptive design has identical closed-loop noise to the baseline at lock (since $G_F \rightarrow G_{F,\text{ss}}$ at steady state).	25
3.14	Adaptive gain profile $G_F(\phi_e)$. Both the scheduled gain and error are substantial when they are near $G_{F,\text{max}}$, and then the error rolls smoothly down to $G_{F,\text{ss}} = 3.0$ rad/V as the error enters the lock band, making the transition continuously differentiable.	27
3.15	Discrete-time pole trajectory versus VCDL gain.	31
4.1	Spectre transient waveforms of v_{ϕ_e} for $\Delta\phi = \pi/4$ phase step (exported from Cadence Virtuoso 6.1.7). Left cursor: step instant ($t = 100$ ns). Right cursor: T_{lock} instant. Horizontal dotted line: ± 5 mV lock threshold (2% criterion). Values read directly from psf/spectre.out.	34

4.2	Spectre transient waveforms of v_{ϕ_e} for $\Delta\phi = \pi/2$ phase step. Same measurement setup as Figure 4.1. The proposed design achieves near one-shot settling at twice the design step amplitude, demonstrating that the adaptive schedule scales gracefully with input disturbance.	35
4.3	Cadence Spectre waveform of the adaptive G_F scheduler output v_{G_F} for the proposed SSPLL during a $\pi/4$ phase step (exported from Virtuoso 6.1.7). The scheduler starts with $G_F \approx 5.218$ rad/V during the first reference cycle (high error), then smoothly reduces G_F towards $G_{F,ss} = 3.0$ rad/V as the phase error decays, to return to steady state within four reference cycles.	38
4.4	Lock-time benchmark for $\Delta\phi = \pi/4$ (left) and $\Delta\phi = \pi/2$ (right). Three configurations are compared: a conventional integer- N charge-pump PLL (hatched, included as a broad reference), the base-paper fixed- G_F SSPLL equivalent to our baseline ($G_F = 3.0$, white bar), and the proposed adaptive- G_F SSPLL (grey, cross-hatched). Lock times for the SSPLL configurations are confirmed Spectre measurements. The percentage annotations indicate the reduction from baseline to proposed; the double-headed arrow spans the two SSPLL bars. The integer- N value is estimated from the DT closed-loop model at $N = 24$, $f_n = 10$ MHz.	39
4.5	Sensitivity of lock time to the schedule transition ϕ_0	41
4.6	Lock time versus the maximum-gain ratio $G_{F,max}/G_{F,crit}$	42
4.7	Output phase-noise spectra.	43
4.8	Steady-state RMS jitter distribution.	44

LIST OF ABBREVIATIONS

Abbrev.	Expansion	Abbrev.	Expansion
ADC	Analog-to-Digital Converter	PFD	Phase–Frequency Detector
BW	Bandwidth	PLL	Phase-Locked Loop
CMOS	Complementary Metal-Oxide Semiconductor	PM	Phase Margin
CP	Charge Pump	PSD	Power Spectral Density
CT	Continuous Time	PSF	Parametric Simulation File
DT	Discrete Time	RMS	Root Mean Square
DTU	Delhi Technological University	SDR	Software-Defined Radio
ECE	Electronics & Communication Engineering	SPDD	Signal Processing and Digital Design
EDA	Electronic Design Automation	SSPD	Sub-Sampling Phase Detector
GF	Gain Factor (VCDL)	SSPLL	Sub-Sampling Phase-Locked Loop
GSDK	Generic Process Design Kit	TF	Transfer Function
LPF	Low-Pass Filter	VA	VerilogA
NBW	Noise Bandwidth	VCDL	Voltage-Controlled Delay Line
PN	Phase Noise	VCO	Voltage-Controlled Oscillator

CHAPTER 1

INTRODUCTION

1.1 INTRODUCTION

Ring-oscillator-based sub-sampling phase-locked loops are suitable for current integrated systems because to their small footprint, wide tuning range and easy migration across process nodes. They also do not suffer from the noise penalty that a feedback divider would ordinarily introduce. These benefits do have a problem however. When such a loop is pushed to a high frequency output, the loop gain is big enough that the bare feedback path is no longer stable. The standard fix of adding a voltage controlled delay line, while it makes the loop stable, leaves the loop settling out in a slow, oscillatory way. The phase error oscillates before settling and the associated lock time is a serious problem for applications such as burst-mode receivers and frequency-agile synthesisers which have to settle rapidly after each channel change.

Two families of design have come to dominate. The familiar charge-pump loop is a Type-II, second-order structure based on a phase–frequency detector, charge pump, and passive filter, and has been the backbone of frequency synthesis for decades [1, 20]. Its second-order tendency is widely known and therefore rarely surprises. What it does give, though, is a very substantial cost: the feedback divider boosts reference-path noise by N^2 , and at high output frequencies that single term tends to consume most of the jitter allowance.

The sub-sampling PLL then directly combats that noise penalty. Instead of dividing the VCO down to the reference, it immediately samples the VCO waveform at the reference edge and reads off the phase error without a divider anywhere in the loop [3]. What remains is a Type-I loop of first-order in discrete time, where each reference cycle takes the current phase error to the next one through a single z -domain pole. Gao and Bhatt demonstrated that this results in a jitter far below what is possible with an integer- N charge-pump loop of similar bandwidth [3]. The trick is in the gain. The detector gain K_{SSPD} , the VCO gain K_{VCO} , and the sampling period T_s form a dimensionless loop gain $\alpha = K_{SSPD}K_{VCO}T_s$, which for the 2.4 GHz ring-VCO investigated here is $\alpha = 3.6128$ —well over unity, large enough to make the bare loop simply unstable.

This project report gives a another perspective on the delay line.

Ghahramani *et al.* found a feasible method around this challenge by putting a voltagecontrolled delay line in the sampling path. VCDL introduces a discrete-time lead-lag zero which pushes the pole back inside the unit circle [7]. Setting the VCDL gain to $G_F = 3.0$ rad/V, their loop attains an integrated jitter of 0.934 ps. Their closed-loop pole is at $z = -0.4451$ —stable, to be sure, but negative-real, which means that the loop settles by oscillating back and forth before ultimately quieting down. The resulting lock time is 50 ns. This is absolutely okay for a synthesiser primarily stuck at one frequency. But it becomes a major bottleneck wherever the frequency

has to change quickly – Bluetooth channel hopping, burst-mode optical receivers, and power-gated clock trees all want the loop to re-settle fast, and fifty nanoseconds of sign-alternating transient starts to seem expensive. This project poses a simple question: what if the VCDL gain were not constant, but allowed to *adapt*—starting high when the loop is chasing a large error, then easing back as it settles—so that the loop locks faster without giving up its steady-state jitter? The following chapters elaborate on that idea, illustrate its stability, and verify in Cadence Virtuoso that it works.

1.2 PROBLEM STATEMENT

The work is motivated by three concrete deficiencies.

Stability was measured with the wrong tool. The base-paper design of Ghahramani et al. [7] bases its reasoning on a continuous-time phase-margin calculation indicating a comfortable 30° margin. Running the analysis in the z -domain, something the base paper never performs, finds a pole at $z = -0.4451$ telling a different story: oscillatory settling the CT view cannot perceive. The disparity and the oscillation were not flagged in the original work

A fixed gain leaves the design space unexplored. $G_F = 3.0 \text{ rad/V}$ is one point in a range from $G_{F,\min} = 1.613 \text{ rad/V}$ up to $G_{F,\text{crit}} = 5.226 \text{ rad/V}$, where $z_{\text{pole}} = 0$ and the loop settles in one cycle. The gain is increased and the transient gets faster (and the noise at lock goes up). Keep it low for noise and you accept oscillatory settling. Nobody had ever systematically charted this lock-time-versus-noise trade-off, much less exploited it.

No stability guarantee exists for a time-varying gain. Once the gain changes cycle-by-cycle in response to instantaneous phase error, the loop is no longer linear time-invariant and standard stability criteria are insufficient. A Lyapunov argument is needed to show the schedule cannot tip the loop into instability under any error condition.

1.3 SCOPE OF RESEARCH

The project is bounded as follows:

1. The focus is squarely on a Type-I, first-order SSPLL with a single VCDL in the sampling path, running at $f_{\text{VCO}} = 2.4 \text{ GHz}$ from a reference of $f_{\text{REF}} = 100 \text{ MHz}$ with a division ratio $N = 24$.
2. Verification is done in Cadence Virtuoso 6.1.7-64b with the Spectre 15.1.0 transient solver on the GPD045 45 nm process. All five functional blocks—the sub-sampling phase detector, the adaptive gain scheduler, the voltage-controlled delay line, the ring VCO, and the lock-time detector—are implemented as behavioural VerilogA models. No transistor-level layouts are included; the aim is to validate the control-theoretic predictions at the behavioural level.

3. Lock time is the headline metric, defined as the first instant at which the phase-error voltage $\phi_{e,v}$ drops below a 5 mV threshold—a 2% settling criterion relative to the design step amplitude. The lock-time detector VerilogA module flags this crossing automatically.
4. Phase-noise and noise-bandwidth analysis is performed analytically, based on the closed-loop transfer function framework introduced in Chapter 3. The model is anchored by representative VCO noise-floor values from the literature (-85 dBc/Hz at 1 MHz offset, $1/f^2$ roll-off).
5. The scheduling law itself is a saturating hyperbolic function of the absolute phase error. Other alternatives, such as lookup tables, piecewise-linear maps, or bang-bang switching are left as candidates for further investigation.

1.4 RESEARCH HYPOTHESIS

The model is anchored by representative VCO noise-floor values from the literature (-85 dBc/Hz at 1 MHz offset, $1/f^2$ roll-off).

Four hypotheses are proposed and tested via simulation: **H1 (Adaptive scheduling decreases lock time)**. An error-dependent VCDL gain starting near $G_{F,\max} = 5.509$ rad/V for large phase errors, and relaxing back to $G_{F,ss} = 3.0$ rad/V at lock, will give noticeably faster settling than the fixed-gain design for a $\pi/4$ -radian step, the improvement being larger for larger steps, and no cost to steady-state jitter.

H2 (Close to one-shot settling feasible). At the $\pi/4$ rad design step, the schedule sets the DT pole inside $|z| < 0.05$, thus the loop essentially settles within one to two reference cycles.

H3 (Lyapunov stability holds for all error magnitudes). The Lyapunov function $V(\phi_e) = \phi_e^2$ monotonically lowers along each trajectory of the adaptive system for all $|\phi_e| \in (0, \pi]$, verifying the unconditioned stability of the gain-scheduled loop.

H4 (SSPLL jitter advantage over Integer-N is preserved). The adaptive SSPLL has at least a $2\times$ jitter edge over an analogous Integer-N charge-pump PLL, because sub-sampling inherently avoids the N^2 noise multiplication that divider-based feedback cannot avoid.

1.5 OBJECTIVES

The specific aims of the project are:

1. To derive the correct discrete-time pole $z_{\text{pole}} = (1 + \beta - \alpha)/(1 + \beta)$ for the VCDL-aided SSPLL, and to bring out the mismatch between the continuous-time phase-margin reasoning of the base paper and what the loop's discrete-time dynamics actually predict.
2. To design an adaptive VCDL gain scheduling law $G_F(|\phi_e|)$ that maps phase-error magnitude to a smoothly varying gain, with $G_{F,\max}$, $G_{F,ss}$, and ϕ_0 determined analytically from the loop parameters.

3. To prove Lyapunov stability of the adaptive loop by showing that $|z_{\text{pole}}(G_F(|\phi_e|))| < 1$ holds for every admissible error magnitude, and that $V = \phi_e^2$ decreases at every step along the trajectory.
4. To validate the design in Cadence Virtuoso/Spectre 15.1.0 transient simulation on behavioural VerilogA models, measuring lock time and settling behaviour for both $\pi/4$ and $\pi/2$ radian phase steps, with the lock-time detector providing automated, threshold-based detection.
5. To benchmark the adaptive SSPLL against the base-paper fixed-gain SSPLL, using Spectre-confirmed lock times for both designs and analytical phase-noise modelling for the noise comparison.
6. To assess the physical feasibility limits of the VCDL—specifically its required delay swing, linearity, and noise contribution—and judge whether the operating range the schedule demands is realistic in a modern CMOS ring-VCO process.

1.6 ORGANIZATION OF THE REPORT

The rest of the report is organised as follows.

Chapter 2 reviews the literature on SSPLL design, VCDL-based bandwidth extension, and gain scheduling in PLL loops. It draws out the specific gaps this work addresses and places the contributions in context against prior published methods.

Chapter 3 establishes the technique and constructs the system model. It introduces the Type-I SSPLL architecture, derives the DT pole, formulates the adaptive scheduling law, and goes through the Lyapunov stability proof. The integrated comparison baseline is the Integer-N charge-pump PLL, which is likewise modeled here. All equations are presented in a way that is directly translatable into a circuit level implementation.

Chapter 4 presents the Cadence Virtuoso/Spectre simulation results, the simulation environment, the VerilogA models, the baseline SSPLL locking at 48.785 ns, the adaptive SSPLL at 28.785 ns for the $\pi/4$ step and 7.525 ns for the $\pi/2$ step, and a side-by-side performance table. Figures are drawn directly from the Spectre data.

Chapter 5 closes with conclusions, an honest look at where the behavioural-modelling approach falls short, and specific directions for future work, including a full transistor-level implementation and silicon measurement.

CHAPTER 2

LITERATURE REVIEW

2.1 BACKGROUND

2.1.1 Phase-Locked Loop Fundamentals

At its core, a phase-locked loop is a feedback system that tracks the phase of an incoming reference with a locally generated clock. The version most designers picture first is the Type-II integer- N charge-pump PLL: a phase-frequency detector compares the reference edge against the divided VCO edge, a charge pump pushes current pulses into a loop filter, and the filtered voltage steers the VCO [20]. The divider closes the loop by scaling the VCO frequency back to the reference, and a second-order RC loop filter sets the bandwidth and phase margin. Thirty years of iterative refinement have left this architecture in firm control of wireline and wireless transceiver design [1].

There are two features of the charge-pump loop that are important in this regard. The first is that it is Type-II: there are two integrators in the open-loop path one from the VCO and one from the charge pump driving the filter capacitor and combined they ensure zero steady-state phase error against a frequency step. The second is not so welcome. Since the divider is in the feedback channel, all the noise sources at the reference input are observed at the output multiplied by N^2 , i.e. the loop scales the reference edge timing uncertainty by the division ratio. For $N=24$, that penalty is a factor of 576, usually the greatest single term in the jitter budget.

2.1.2 Sub-Sampling Phase-Locked Loops

Gao and Bhatt proposed the sub-sampling PLL to eliminate that divider penalty entirely [3]. In an SSPLL the reference clock samples the VCO output directly through a switch or a track-and-hold, and the sampled voltage comes out proportional to $K_{SSPD} \sin(\phi_e)$, where ϕ_e is the phase error between the VCO zero-crossing and the reference edge. For small errors $\sin(\phi_e) \approx \phi_e$, so the detector acts as a linear phase comparator with gain K_{SSPD} in V/rad. Because nothing divides the signal on its way round the loop, the reference-noise multiplication simply never happens.

That saving does come at a cost, though. The SSPLL is a Type-I loop, with the VCO as its only integrator. Its DC open-loop gain is the dimensionless product $\alpha = K_{SSPD} K_{VCO} T_s$, set by the VCO gain (rad/s/V), the SSPD gain (V/rad), and the sampling period $T_s = 1/f_{REF}$. Once $\alpha > 2$, the discrete-time pole $z_{pole} = 1 - \alpha$ slips outside the unit circle and the bare loop is unstable. For the 2.4 GHz ring-VCO target here, $K_{VCO} = 2\pi \times 115 \text{ MHz/V}$ and $T_s = 10 \text{ ns}$ give $\alpha = 3.6128$, which means the loop simply cannot run without some form of stabilisation.

2.1.3 Voltage-Controlled Delay Lines for Bandwidth Extension

Using a delay element to reshape loop dynamics is not a new idea; delay-line-based PLLs appeared early in serial-link clock recovery [13]. What Ghahramani *et al.* added was the specific insight that a VCDL placed in the *sampling* path of an SSPLL injects a discrete-time zero, much as a lead compensator does in continuous-time control [7]. The delay line shifts the sampling instant in proportion to the control voltage, so the effective phase error the SSPD sees becomes $\phi_e' = \phi_e - G_F \cdot v_{\text{ctrl}}$, with G_F the VCDL gain in rad/V. Algebraically, the delay shift introduces a factor $(1 + \beta)$ into the denominator of the closed-loop pole, where $\beta = K_{\text{SSPD}} G_F$ —and it is that factor that pulls the pole back from $z = 1 - \alpha$ toward zero.

2.2 CRITICAL REVIEW OF EXISTING WORK

2.2.1 The Base Paper: Ghahramani et al. (2025)

Ghahramani *et al.* [7] built a ring-VCO SSPLL at 2.4 GHz from a 100 MHz reference, showed that a VCDL set to $G_F = 3.0$ rad/V stabilises it, and reported 0.934 ps of integrated jitter. Their stability case runs through the continuous-time open-loop transfer function

$$L(j\omega) = K_{\text{SSPD}} \cdot \frac{K_{\text{VCO}}}{j\omega} \cdot \left(1 + \frac{j\omega}{\omega_{\text{ZERO}}}\right) \cdot e^{-j\omega T_s/2} \cdot \text{sinc}\left(\frac{\omega}{\omega_{\text{REF}}}\right), \quad (2.1)$$

where $\omega_{\text{ZERO}} = K_{\text{VCO}}/G_F$ is the lead-zero frequency introduced by the VCDL. The phase-margin formula they derive is

$$\text{PM} = 90^\circ - 180^\circ \frac{\omega_{\text{UG}}}{\omega_{\text{REF}}} + \arctan\left(\frac{\omega_{\text{UG}}}{\omega_{\text{ZERO}}}\right), \quad (2.2)$$

which works out to about 30° for their parameters. Seeing a positive margin, they conclude the loop is stable.

That conclusion is where the trouble starts. The loop updates once per reference cycle, so the correct model is a first-order discrete-time recursion, not a CT transfer function. Computing the z -domain pole directly—something the base paper never does—gives $z_{\text{pole}} = (1 + \beta - \alpha)/(1 + \beta) = -0.4451$: negative-real. A negative real pole flips the sign of the phase error from one reference edge to the next—an oscillatory settling mode the CT picture simply cannot reveal. That gap is the seed of Contribution C1.

A second oversight is that the paper never maps the wider design space. The critical gain $G_{F,\text{crit}} = (\alpha - 1)/K_{\text{SSPD}} = 5.226$ rad/V—where $z_{\text{pole}} = 0$ and the loop settles in one shot—goes entirely unmentioned.

2.2.2 Sub-Sampling PLL Architectures and Variants

Gao and Bhatt's original SSPLL [3] used an LC-VCO at 2.21 GHz, where K_{VCO} is small enough that $\alpha < 1$ and no bandwidth extension is needed at all. Several subsequent designs pushed SSPLL operation to higher VCO gains and higher reference

frequencies. Nagam and Kinget combined a sub-sampling phase detector with a feed-forward noise-cancellation path in a ring-oscillator PLL, achieving low jitter while keeping the loop robust to supply noise [18]. Their work confirmed that sub-sampling detection extends naturally to ring-oscillator clocking, but the added feed-forward path is a second corrective mechanism that does not exist in a pure VCDL-based design, and its benefit depends on how accurately that path tracks the oscillator noise.

Kumar and Singh took a switched-loop-filter approach: a fast phase-error-correction path widens the loop during acquisition, then the loop filter is switched to a narrow configuration for low-jitter tracking [11]. This decouples the lock-time problem from the jitter problem, but it does so by adding switching circuitry and an auxiliary correction path. The adaptive gain scheduling proposed in this work achieves a similar effect—fast acquisition, then low-jitter tracking—within a single-path architecture by varying only the VCDL gain, which is already present in the loop.

2.2.3 Gain Scheduling and Adaptive Bandwidth Techniques

Bandwidth switching and gear-shifting have been long-standing tricks in PLL design. Hanumolu *et al.* showed that toggling the charge-pump current between a high value for fast acquisition and a low value for low-jitter tracking can cut the lock time of a Type-II loop by an order of magnitude [8]—a two-state gain schedule controlled by a lock detector. Lee and Kim went further, varying the loop bandwidth continuously under a digital accumulator that tracks the phase-error magnitude [12]. Both techniques work, but both live in second-order charge-pump PLLs rather than SSPLLs, so their design equations and stability conditions are a different animal from the first-order SSPLL case.

What no one had done was apply adaptive gain scheduling to the VCDL path of a sub-sampling PLL. The SSPLL's pole structure is unusually clean: a single real pole that moves smoothly and monotonically with G_F . That simplicity is what keeps an analytical Lyapunov proof within reach—something higher-order architectures cannot easily offer.

2.2.4 Discrete-Time and z -Domain PLL Modelling

Gardner identified the need for discrete-time modelling of sampled PLLs early on [5], pointing out that the continuous-time approximation breaks down once the loop bandwidth grows to a significant fraction of the reference frequency. The standard fix for charge-pump loops is to treat the PFD and charge pump as a sampled system and use z -transform analysis, or an exact impulse-invariant mapping [9]. These are well-established methods for Type-II loops, but much less has been said about applying them to Type-I sub-sampling loops.

Ghahramani *et al.* [7] stay entirely with the continuous-time model and never write down the z -domain transfer function. This work closes that gap: the correct DT pole is $z_{\text{pole}} = (1 + \beta - \alpha)/(1 + \beta)$, and the continuous-time Euler-forward approximation $z \approx 1 + sT_s$ gives the same algebraic expression for a first-order system. So the disagreement is not about where the pole sits—both models place it at the same location. The real disagreement is about the *stability criterion*. The CT phase margin carries no

information about the $|z| < 1$ condition that governs discrete-time stability, and it gives no warning when $\alpha > 2(1 + \beta)$.

2.2.5 Lyapunov Methods for PLL Stability

Lyapunov stability analysis has been applied to PLLs in a handful of specific contexts. Viterbi's classic treatment of first-order analog PLLs uses the sinusoidal phase-detector characteristic to derive a Lyapunov function for the nonlinear dynamics [22]. More recent work by Leonov and Kuznetsov applied rigorous nonlinear and Lyapunov-based analysis to define the hold-in, pull-in, and lock-in ranges of PLLs and to expose the limitations of classical linear stability arguments [15]. In both cases the key is to identify a scalar function of the state that decreases at every step.

In the VCDL-aided SSPLL with adaptive scheduling, the state is a single scalar (ϕ_e), and the dynamics reduce to $\phi_e[k + 1] = z_{\text{pole}}(G_F(|\phi_e[k]|)) \phi_e[k]$. The natural Lyapunov candidate is $V(\phi_e) = \phi_e^2$, and proving $V[k + 1] < V[k]$ reduces to showing $|z_{\text{pole}}(G_F(|\phi_e|))| < 1$ for every $|\phi_e| \in (0, \pi]$. That argument has not been made before for a gain-scheduled SSPLL, and making it is Contribution C4 of this project.

2.2.6 Ring Versus LC Oscillators in Sub-Sampling Loops

The choice of oscillator has more downstream design implications than most comparisons suggest. LC tanks have very low phase noise, but this comes at the cost of tiny tuning gain. In Gao's original study [3] K_{VCO} was small enough such that α never came close to the unit circle, so the loop was pleasantly stable without compensation. The price is area and a narrow tuning range—LC tanks don't fit well into small digital processes.

On the opposite corner are ring oscillators: compact, wide-tuning and process-portable. What they give up is a lot larger K_{VCO} and much worse phase noise, which is well modeled by Leeson's model and Razavi's CMOS oscillator analysis [14, 19]. This huge gain is the fundamental impediment to sub-sampling: a 2.4 GHz ring VCO with $K_{\text{VCO}} = 2\pi \times 115 \text{ MHz/V}$ pushes α to 3.6128, well over the $\alpha = 2$ stability limit. Just careful biasing does not bring the pole into the unit circle. This is why a ring based SSPLL requires an explicit stabilizing mechanism, i.e. the slot the VCDL fills.

2.2.7 Phase-Noise and Jitter Modelling in Sampled Loops

To go from a stability tale to a jitter number you need a noise model that respects the sampled character of the loop. The traditional starting point for the oscillator itself is Leeson's $1/f^2$ area for white frequency noise, with the close-in $1/f^3$ corner folded in for flicker up-conversion [14, 19]. These establish the free-running phase-noise floor, and the loop then molds that floor with its noise transfer function. A good example for this is the self-biased designs from Maneatis, where linking the loop parameters together locks the damping and makes the jitter predictable across process and supply corners [16]. Mansuri and Yang took the same notion one step further with an explicit adaptable bandwidth, sacrificing a broader loop during transients for a tighter one at

lock [17]—a philosophy that rhymes with the gain scheduling described here, albeit the technique is different.

The common denominator for most of these treatments is a continuous-time noise transfer function. This works well when the loop bandwidth is a small fraction of the reference, but silently mischaracterizes a sub-sampling loop with its pole at the Nyquist rate. The texts by Egan and Shu and Sánchez-Sinencio [2, 21] give the synthesiser noise budget in full but they also rely on the CT picture for the loop dynamics. We show in Chapter 3 that an oscillating negative-real pole ($z_{\text{pole}} < 0$) enhances noise around half the reference frequency in a way that a CT bandwidth doesn't capture. The proper jitter ordering between the baseline and suggested design relies on the use of the accurate discrete-time NTF instead of the CT bandwidth shortcut.

2.2.8 Fast-Lock Strategies of Different PLL Families

2.2.9 Fast-Lock Strategies Across PLL Families

The reduction of lock time is an ancient goal, and the literature on charge-pumps has generated a comprehensive inventory of solutions that are worth comparing with the sub-sampling approach. The most direct is bandwidth switching, open the loop wide to acquire quickly then narrow it to settle neatly. This two-state behaviour is described by Hanumolu's analysis of charge-pump loops [8] and Lee and Kim's implementation of an adaptive variation that changes the charge-pump current on the fly [12]. Both work, but both also cost in complexity - an auxiliary detector, extra control logic, or a second filter path. The PFD itself introduces noise and dead-zone effects which must be addressed, as quantified by Homayoun and Razavi [10]. Gardner's treatment remains the definitive account of why these loops behave as they do [5, 6].

Sub-sampling designs have chased the same target by other routes. Nagam and Kinget add a feed-forward noise-cancellation path to reduce jitter in a ring-oscillator sub-sampling PLL [18]. Kumar and Singh's switched-loop-filter PLL widens the loop for acquisition and then narrows it for tracking [11], again at the cost of extra switching circuitry. Spur behaviour in sub-sampling detectors has been studied separately by Gao [4], underscoring that the sampling operation brings its own artefacts that any fast-lock scheme must respect. Against this backdrop, the approach here is deliberately economical. One existing parameter—the VCDL gain—is varied as a function of the instantaneous phase error. No extra path, no second detector, no injection circuitry. The same loop is wide during acquisition and critically damped at lock.

2.2.10 Sampled Feedback and the Deadbeat Idea

Stepping away from the PLL framing reveals something useful: the control literature already has a name for what the critical gain achieves. Treat the SSPLL as a plain first-order sampled feedback system and the term is obvious—deadbeat response. A discrete-time system whose single pole sits exactly at the origin returns to zero in one sample after a disturbance, with no ringing and no residual tail. That is precisely the condition $z_{\text{pole}} = 0$, reached at $G_{\text{F,crit}} = 5.226 \text{ rad/V}$. Seen this way, the adaptive schedule is doing something a control engineer would recognise immediately—moving the

loop toward deadbeat while the error is large, then backing off to a gentler pole once the state is small. The synthesiser texts frame loop design almost entirely in continuous time [2, 6, 21], which is why the deadbeat reading is easy to miss. It only becomes visible once the loop is written as a sampled recursion.

2.2.11 Where Fast Locking Actually Matters

The desire to shave nanoseconds off a lock time is concrete. The synthesiser in frequency-hopping radios has to re-settle every time the channel changes, and the loop is in effect dead time during which no meaningful data passes. A loop that settles in tens of nanoseconds instead of hundreds directly enhances the fraction of time the connection is useable. Burst-mode and duty-cycled receivers go one step further: they turn off the synthesiser between bursts to save energy, and pay the lock time again on each wake-up, so lock time becomes an energy cost, not just a delay penalty. When the clock is gated for power management, clock generation for CPUs and serial links has the same nature. Sub-sampling loops are desirable in all these scenarios, because eliminating the divider eliminates the N^2 noise penalty [3, 4] but the Type-I architecture that accompanies sub-sampling is precisely what makes rapid, clean acquisition problematic. The solution which restores rapid locking without reintroducing a division or a second loop thereby solves a genuine and recurring limitation, rather than a contrived one.

2.3 IDENTIFICATION OF RESEARCH GAP

Across the surveyed work, three gaps jump out, and this initiative aims to address each:

Gap 1 — DT stability analysis is missing for VCDL-aided SSPLLs. The base study is fully based on continuous-time phase-margin analysis and never calculates the z-domain pole; the oscillatory settling predicted by the negative-real DT pole is not discussed. Nothing has replaced that gap since. The DT pole is determined here and its value is tabulated for the whole range of G_F . The CT–DT discrepancy is stated as Contribution C1.

Gap 2 — No adaptive gain scheduling for the VCDL path. Bandwidth switching and gear-shifting are common in charge-pump PLLs, but no one has implemented continuous adaptive gain scheduling for the VCDL of an SSPLL. The loop is nearly made for it: its solitary DT pole records one scalar parameter, G_F , in a smooth monotonic connection. This project gives a concrete scheduling legislation and a rigorous stability proof (C4, C6).

Gap 3 — No Lyapunov stability proof for gain-scheduled SSPLLs. The Lyapunov analyses that exist for PLLs deal with analog loops or binary phase detectors, not sampled sub-sampling loops carrying a time-varying gain. The proof given in Chapter 3 is, to the best of the author’s knowledge, the first Lyapunov argument written specifically for a VCDL-aided SSPLL with adaptive gain scheduling.

Table 2.1 gathers the reviewed works and maps each onto these gaps.

Table 2.1: Summary of reviewed works and coverage of identified research gaps.

Reference	Architecture	VCO Type	DT Model	Adaptive G_F	Lyapunov
Gao & Bhatt [3]	SSPLL (LC)	LC	No	No	No
Ghahramani <i>et al.</i> [7]	SSPLL + VCDL	Ring	No	No	No
Nagam & Kinget [18]	SSPLL + FFNC	Ring	No	No	No
Kumar & Singh [11]	Switched-loop PLL	—	No	No	No
Hanumolu <i>et al.</i> [8]	CP-PLL (BW switch)	Ring	N/A	2-state	No
Lee & Kim [12]	CP-PLL (adaptive)	LC	N/A	Yes (CP)	No
Leonov <i>et al.</i> [15]	Nonlinear PLL theory	—	No	No	Yes
Present work	SSPLL + VCDL	Ring	Yes	Yes	Yes

CHAPTER 3

SYSTEM DESIGN AND METHODOLOGY

3.1 SIMULATION PLATFORM AND MATERIALS

3.1.1 Cadence Virtuoso and Spectre Environment

All circuit-level work here runs in Cadence Virtuoso 6.1.7-64b on the Spectre 15.1.0 transient solver (MMSIM151). GPDK045, a generic 45 nm CMOS PDK, provides the device models; they sit cleanly inside Spectre’s unified simulator framework. A Red Hat Enterprise Linux virtual machine hosts the tools, with netlists and VerilogA sources kept in sync with the Windows side through a VMware shared folder at `/mnt/hgfs/PLL2/`.

The transient run uses `stop = 500 ns`, `errpreset = moderate`, and `maxstep = 20 ps`. That fine step is not just a preference—it is necessary. The SSPD track-and-hold window is only 10 ps (`Thold = 1e-11 s`), and anything coarser would step right over the sampling edge and miss it entirely. Initial conditions are set to `phi_e_v = 0` and `vcdl_out = 0` to keep the Spectre DC operating point from injecting a spurious phase error at $t = 0$.

A patching script (`fix_all.sh`) handles housekeeping between netlisting and simulation, overwriting six parameter errors that Cadence ADE quietly reintroduces whenever it regenerates the netlist from the schematic. These include: $G_{F,max} : 5.5 \rightarrow 5.509 \text{ rad/V}$; the sampling step time $t_{step} : 20 \text{ ns} \rightarrow 100 \text{ ns}$; the VCO control nominal $V_{ctrl,nom} : 0.6 \rightarrow 0 \text{ V}$; and deletion of the stale `ahd1SimDB` directory that would otherwise block VerilogA recompilation. The corrected netlist is then run *without* the “Netlist and Run” combined command—netlisting and simulation are always kept as separate steps so those corrected parameters cannot be silently overwritten.

3.1.2 VerilogA Behavioural Models

The complete SSPLL testbench is built from five VerilogA modules. Every one of them is a purely behavioural model—no transistor-level circuitry anywhere in the design—with all parameters exposed directly in the Spectre netlist so that any value can be swept or patched without recompiling the models. Table 3.1 lists each module alongside its port signature and key parameter settings.

Lock time is encoded directly as a voltage ($1 \text{ V} = 1 \text{ ns}$) and printed to `psf/spectre.out` the moment the 5 mV threshold is crossed—no waveform post-processing needed. That threshold represents 2% of the nominal SSPD output for a $\pi/4$ -radian step ($K_{SSPD} \cdot \pi/4 \approx 393 \text{ mV}$).

Table 3.1: VerilogA behavioural modules used in the Spectre testbench.

Module	Ports (ordered)	Key parameters
sspd_subsampling	ref_clk, vco_in, vout, vss	$K_{SSPD} = 0.5 \text{ V/rad}$; $V_{amp} = V_{mid} = 0.5 \text{ V}$; $V_{th,clk} = 0.6 \text{ V}$; $T_{hold} = 10 \text{ ps}$
vco_ring	ctrl, out, vss	$f_0 = 2.4 \text{ GHz}$; $K_{VCO} = 115 \text{ MHz/V}$; $V_{amp} = V_{mid} = 0.5 \text{ V}$; $\phi_{init} = \pi \text{ rad}$; $V_{ctrl,nom} = 0 \text{ V}$; $\phi_{step} = 0 \text{ rad}$
vcdl_adaptive	vin, v_ctrl, v_gf, vout, vss	Computes delay-corrected control voltage from $\phi_{e,v}$ and v_{GF}
adaptive_gf_sched	v_phi_e, v_gf_out, vss	$G_{F,ss} = 3.0 \text{ rad/V}$; $G_{F,max} = 5.509 \text{ rad/V}$; $\phi_0 = 0.10 \text{ rad}$
phase_step_src	clk_out, vss	$f_{REF} = 100 \text{ MHz}$; $\phi_{step} = \pi/4 \text{ rad}$; $t_{step} = 100 \text{ ns}$

3.1.3 System Parameters

Table 3.2 collects the design parameters held fixed throughout all simulations. These values match Table I of the base paper [7].

Table 3.2: SSPLL system parameters used in all Spectre simulations.

Parameter	Symbol	Value
VCO centre frequency	f_{VCO}	2.4 GHz
Reference frequency	f_{REF}	100 MHz
Division ratio	N	24
VCO gain	K_{VCO}	$2\pi \times 115 \text{ MHz/V}$
SSPD gain	K_{SSPD}	0.5 V/rad
Sampling period	$T_s = 1/f_{REF}$	10 ns
DT open-loop gain	$\alpha = K_{SSPD}K_{VCO}T_s$	3.6128
Base-paper VCDL gain	$G_{F,ss}$	3.0 rad/V
DT VCDL factor (base paper)	$\beta = K_{SSPD}G_{F,ss}$	1.50

3.2 TYPE-I SSPLL ARCHITECTURE

3.2.1 Block Diagram

Figure 3.1 shows the signal flow of the adaptive VCDL-aided SSPLL as implemented in the Spectre testbench. Four functional blocks form the main loop; the adaptive scheduler and lock-time detector connect as auxiliary paths.

The sub-sampling phase detector samples the VCO output waveform at each reference clock edge. For phase errors smaller than $\pi/2$, the sampled voltage is $\phi_{e,v} = K_{SSPD} \cdot \phi_e$ with $K_{SSPD} = 0.5 \text{ V/rad}$. The voltage-controlled delay line takes $\phi_{e,v}$ (input signal) and the adaptive gain voltage v_{GF} (from the scheduler), computing a delay-




Figure 3.1: Signal flow of the adaptive VCDL-aided SSPLL. The SSPD samples the VCO output at each reference clock edge, producing phase-error voltage $\phi_{e,v}$. The adaptive G_F scheduler recomputes the VCDL gain each cycle; the VCDL translates the corrected phase error into the control voltage v_{ctrl} that drives the ring VCO, closing the loop.

corrected control voltage v_{ctrl} . The ring VCO converts v_{ctrl} to a frequency offset around 2.4 GHz; its output feeds back to the SSPD, closing the single-feedback-path loop.

3.2.2 Continuous-Time Open-Loop Transfer Function

The CT open-loop transfer function is:

$$L(j\omega) = K_{SSPD} \cdot \frac{K_{VCO}}{j\omega} \cdot \left(1 + \frac{j\omega}{\omega_{ZERO}}\right) \cdot e^{-j\omega T_s/2} \cdot \text{sinc}\left(\frac{\omega}{\omega_{REF}}\right), \quad (3.1)$$

where $\omega_{ZERO} = K_{VCO}/G_F$ is the VCDL lead-zero frequency. Unity-gain crossover is at $\omega_{UG} = K_{SSPD}K_{VCO} = 2\pi \times 57.5$ MHz, giving $\omega_{UG}/\omega_{REF} = 0.575$. The CT phase margin at $G_F = 3.0$ rad/V evaluates to approximately 30° :

$$\text{PM} = 90^\circ - 180^\circ \frac{\omega_{UG}}{\omega_{REF}} + \arctan\left(\frac{\omega_{UG}}{\omega_{ZERO}}\right). \quad (3.2)$$

On its own this CT result says the loop is stable. What it cannot see is that the discrete-time pole is negative-real, which is what produces the oscillatory SSPD samples Spectre later shows. The DT model below settles the matter.

Figure 3.2 overlays the CT open-loop Bode diagram for the base-paper design (Ghahramani et al., $K_{VCO} = 241$ MHz/V) and for our three operating points. The vertical dashed line sits at $f_{REF}/2 = 50$ MHz, and the shaded band to its right is where the CT approximation gives way for our design ($\alpha = 3.6128 \gg 1$). Every one of our unity-gain crossovers ($f_{UG} = 172$ – 317 MHz) falls inside that invalid band—which is the whole reason a discrete-time analysis, not a CT phase-margin argument, is the right stability test here.

Figure 3.3 sets the two stability viewpoints beside each other. The left panel redraws the CT phase-margin family from the base paper ($G_F = 0$ – 8 rad/V); our operating points at $\omega_{UG}/\omega_{REF} = 1.72$ – 3.17 land far to the right of the plotted range, well




Figure 3.2: CT open-loop Bode plot: base paper (Ghahramani et al. 2025, $K_{VCO} = 241$ MHz/V, $G_F = 3.0$, dash-dot) and this work ($K_{VCO} = 115$ MHz/V) for $G_F = 3.0$, $G_{F,crit} = 5.226$, and $G_{F,max} = 5.509$ rad/V. The shaded region ($f > f_{REF}/2$) marks the frequency range where the CT approximation is invalid for our design ($\alpha = 3.61 \gg 1$). All our unity-gain crossovers fall inside this region, motivating the exact discrete-time analysis of Section 3.3.

outside the valid CT domain. The right panel shows the metric that actually applies—the distance from the unit-circle boundary, $1 - |z_{\text{pole}}|$, against G_F . It is zero at the stability edge ($G_F = 1.613 \text{ rad/V}$) and unity at the deadbeat point ($G_{F,\text{crit}}$), so the plot reads directly as how much margin each operating point keeps from instability.




Figure 3.3: Left: CT phase margin vs. normalised bandwidth $\omega_{\text{UG}}/\omega_{\text{REF}}$ for five VCDL gain values (base-paper style, Ghahramani 2025 Fig. 4b). Our design operates at $u \approx 1.7\text{--}3.2$, far outside this range. Right: correct DT stability margin $1 - |z_{\text{pole}}|$ vs. G_F for our design. Marked points correspond to $G_{F,\text{ss}} = 3.0$ (baseline), $G_{F,\text{crit}} = 5.226$ (deadbeat), and $G_{F,\text{max}} = 5.509 \text{ rad/V}$ (proposed maximum).

3.2.3 Integer- N Charge-Pump PLL: Comparison Architecture

The Integer- N charge-pump PLL (CP-PLL) is the noise-performance yardstick used throughout this work. Figure 3.4 shows its signal chain.




Figure 3.4: Integer- N charge-pump PLL (comparison baseline). The $\div N$ divider in the feedback path scales reference-path noise by N^2 . At $N = 24$ this represents a 27.6 dB noise penalty relative to an architecture without a divider.

The critical noise-penalty distinction between the CP-PLL and SSPLL is that the divider output phase noise is multiplied by N^2 when referred to the VCO output. For

$N = 24$ the multiplication factor is 576 (27.6 dB). Sub-sampling eliminates the divider entirely: the SSPD directly samples the VCO waveform at the reference edge, so no N^2 multiplication path exists. This fundamental architectural advantage enables the SSPLL to achieve sub-picosecond jitter in ring-VCO implementations where the CP-PLL would produce several picoseconds under the same noise conditions.

3.3 DISCRETE-TIME MODEL

3.3.1 Derivation of the z -Domain Pole

The SSPLL updates once per reference cycle, so its dynamics are inherently discrete-time. The VCDL contributes a local feedback factor $\beta = K_{\text{SSPD}}G_F$ that scales down the effective phase error presented to the integrator:

$$\phi_e'[k] = \frac{\phi_e[k]}{1 + \beta}. \quad (3.3)$$

The VCO integrates the corrected error over one reference period:

$$\phi_{\text{VCO}}[k+1] = \phi_{\text{VCO}}[k] + \frac{\alpha}{1 + \beta} \phi_e[k]. \quad (3.4)$$

With $\phi_e[k] = \phi_{\text{ref}}[k] - \phi_{\text{VCO}}[k]$ and the reference advancing by a fixed increment every cycle, the error recursion collapses to:

$$\phi_e[k+1] = z_{\text{pole}} \cdot \phi_e[k], \quad z_{\text{pole}} = \frac{1 + \beta - \alpha}{1 + \beta}. \quad (3.5)$$

Everything else in the project is built around this one-pole recursion. The sign of z_{pole} , its size, and how close to zero it is sitting impact the overall settling behavior, whether the error oscillates, decays monotonically or collapses in one shot. The equivalent discrete-time signal-flow diagram is shown in Figure 3.5.




Figure 3.5: Discrete-time signal flow of the VCDL-aided SSPLL. The SSPD computes $\phi_e[k] = \phi_{\text{ref}}[k] - \phi_{\text{VCO}}[k]$ (left summing junction); the VCDL divides the error by $(1 + \beta)$; the VCO gain block scales by $\alpha = K_{\text{SSPD}}K_{\text{VCO}}T_s$; the right summing junction accumulates the phase. The z^{-1} element represents one reference-cycle delay. Eliminating ϕ_{VCO} yields the scalar recursion $\phi_e[k+1] = z_{\text{pole}} \cdot \phi_e[k]$.

3.3.2 Pole Location as a Function of G_F

Substituting $\beta = K_{SSPD}G_F$:

$$z_{\text{pole}}(G_F) = 1 - \frac{\alpha}{1 + K_{SSPD}G_F}. \quad (3.6)$$

Three gain values of interest drop straight out of this expression:

- **Stability boundary:** $G_{F,\min} = (\alpha - 2)/(2K_{SSPD}) = 1.613 \text{ rad/V}$, where $z_{\text{pole}} = -1$.
- **One-shot settling:** $G_{F,\text{crit}} = (\alpha - 1)/K_{SSPD} = 5.226 \text{ rad/V}$, where $z_{\text{pole}} = 0$ and the error vanishes in a single reference cycle.
- **Oscillatory–monotone boundary:** $G_F < G_{F,\text{crit}}$ gives $z_{\text{pole}} < 0$ (oscillatory); $G_F > G_{F,\text{crit}}$ gives $z_{\text{pole}} > 0$ (monotone decay).

Table 3.3 lists the pole value at a spread of representative gains. The base paper sits at $z = -0.4451$, and the Spectre baseline waveform shows exactly the alternating-sign SSPD samples that a negative-real pole produces—confirming the DT model experimentally.

Table 3.3: DT pole z_{pole} as a function of VCDL gain G_F .

G_F (rad/V)	β	z_{pole}	Settling behaviour
0	0.000	−2.613	Unstable (no VCDL)
1.613	0.807	−1.000	Marginally stable boundary
3.0	1.500	−0.4451	Stable, oscillatory (base paper)
5.226	2.613	0.0000	One-shot settling ($G_{F,\text{crit}}$)
5.509	2.755	+0.0366	Stable, near one-shot (proposed max)
8.0	4.000	+0.2774	Stable, slow monotone

Figure 3.6 maps the pole locus from $G_{F,\min}$ to $G_{F,\max}$. Three operating points are marked: the base-paper design at $z = -0.4451$, the one-shot point at the origin, and the proposed maximum at $z = +0.037$.

Figure 3.7 shows the same pole locus in a broader context, including the $G_F=0$ (no-VCDL) unstable point at $z = -2.613$ and an arc marking the adaptive-transition path the pole follows during lock acquisition: from $z \approx 0$ (high- G_F transient phase) back to $z = -0.4451$ (steady-state).

3.4 ADAPTIVE GAIN SCHEDULING LAW

3.4.1 Design Rationale

Three cycles are sufficient to settle the error when the DT pole is at zero, however fourteen or more may be required when it is at $z = -0.4451$. The essential finding is that the optimal pole position relies on the magnitude of the phase error at that time, and is not a design constant. A large error prefers $G_F \approx G_{F,\text{crit}}$ (pole close to zero,




Figure 3.6: VCDL aided SSPLL discrete time pole locus. as G_F grows, the pole shifts on the real axis. At $z = -0.4451$ the loop is steady but oscillatory; at $z \approx 0$ the loop is one-shot settling.




Figure 3.7: Extended DT pole map The pole locus (solid line) encompasses the entire range from the no-VCDL unstable point ($z = -2.613$, $G_F = 0$) to the base-paper operating point ($z = -0.4451$, open circle), the deadbeat point ($z = 0$, square), and the recommended maximum ($z = +0.038$, triangle). The curving arc is the adaptive transition, as the scheduler reduces G_F in response to the decaying error, moving the pole from $z \approx 0$ to $z = -0.4451$ during each lock acquisition transient. The unit-circle stability border is the dashed circle .

rapid monotone decay), a small error near lock prefers $G_F = G_{F,ss}$ (pole at -0.4451 , slight oscillation comfortably within the lock band). A smooth, continuous schedule carries out this logic automatically, without any lock detector and without any abrupt gain switch that could inject a transient into the control path.

3.4.2 Scheduling Function

The adaptive gain is chosen to be a saturating hyperbolic function of the absolute phase error—smooth, bounded, and with a clear physical interpretation at both extremes:

$$G_F(|\phi_e|) = G_{F,ss} + (G_{F,max} - G_{F,ss}) \frac{|\phi_e|}{|\phi_e| + \phi_0}, \quad (3.7)$$

with $G_{F,ss} = 3.0 \text{ rad/V}$, $G_{F,max} = 5.509 \text{ rad/V}$, and $\phi_0 = 0.10 \text{ rad}$. Three useful properties come for free. First, $G_F(0) = G_{F,ss}$, so there is no leftover high gain at lock. Second, $G_F \rightarrow G_{F,max}$ as the error grows large. Third, G_F is continuously differentiable everywhere, which means no abrupt gain step can inject ringing into the control path. The complete schedule is plotted in Figure 3.8.




Figure 3.8: Adaptive gain-factor schedule. At the design step $|\phi_e| = \pi/4 = 0.785 \text{ rad}$, the schedule evaluates to $G_F = 5.218 \text{ rad/V}$ (within 0.2% of $G_{F,crit}$), placing the DT pole at $z = +0.037$. As the error decays toward zero, G_F returns smoothly to $G_{F,ss} = 3.0 \text{ rad/V}$.

3.4.3 Choice of $G_{F,max}$

The best choice of $G_{F,max}$ is the one that drives the DT pole to zero when the schedule is evaluated right at the design step $\Delta\phi = \pi/4 \text{ rad}$. Setting $G_F(\Delta\phi) = G_{F,crit} = 5.226 \text{ rad/V}$ and solving (3.7) gives:

$$G_{F,max}^{opt} = G_{F,ss} + (G_{F,crit} - G_{F,ss}) \frac{|\Delta\phi| + \phi_0}{|\Delta\phi|} = 3.0 + 2.226 \times \frac{0.785 + 0.10}{0.785} = 5.509 \text{ rad/V}. \quad (3.8)$$

At this value $G_F(\pi/4) = 5.218 \text{ rad/V}$, which puts $z_{\text{pole}} = +0.037$ —close to zero and strictly positive, so the error decays monotonically instead of oscillating.

3.5 LYAPUNOV STABILITY PROOF

3.5.1 Statement

Theorem. Let $\phi_e[k+1] = z_{\text{pole}}(G_F(|\phi_e[k]|))\phi_e[k]$ with G_F given by (3.7), $G_{F,\text{ss}} \geq G_{F,\text{min}}$, and $G_{F,\text{max}} \leq (\alpha + 1)/K_{\text{SSPD}}$. Then $V(\phi_e) = \phi_e^2$ is a Lyapunov function for the equilibrium $\phi_e = 0$.

3.5.2 Proof

The proof reduces to showing $V[k+1] < V[k]$ for every $\phi_e[k] \neq 0$. Starting from the recursion:

$$V[k+1] = \phi_e[k+1]^2 = z_{\text{pole}}^2(G_F(|\phi_e[k]|)) \cdot \phi_e[k]^2 = z_{\text{pole}}^2 \cdot V[k]. \quad (3.9)$$

Thus $V[k+1] < V[k]$ iff $|z_{\text{pole}}| < 1$. From (3.6), $z_{\text{pole}} = 1 - \alpha/(1 + K_{\text{SSPD}}G_F)$, so the condition $|z_{\text{pole}}| < 1$ requires:

$$\begin{aligned} -1 &< 1 - \frac{\alpha}{1 + K_{\text{SSPD}}G_F} < 1 \\ \implies \frac{\alpha - 2}{2K_{\text{SSPD}}} &< G_F < \infty. \end{aligned} \quad (3.10)$$

The lower bound is $G_{F,\text{min}} = (\alpha - 2)/(2K_{\text{SSPD}}) = 1.613 \text{ rad/V}$. Because $G_F(|\phi_e|) \geq G_{F,\text{ss}} = 3.0 > 1.613$ for all $|\phi_e| \geq 0$, the condition is met everywhere along the trajectory. There is no upper bound to worry about either: $G_{F,\text{max}} = 5.509$ stays well clear of any instability from that side.

The stability margin $G_{F,\text{ss}}/G_{F,\text{min}} = 3.0/1.613 = 1.86\times$ is a direct robustness measure. $K_{\text{SSPD}}K_{\text{VCO}}$ would have to fall by 46% before the loop reached the stability boundary—comfortable headroom against PVT variation. ■

3.6 NOISE BANDWIDTH AND JITTER ANALYSIS

3.6.1 Noise Bandwidth

Because the SSPLL is a sampled-data system, its one-sided noise bandwidth must be computed from the exact discrete-time noise transfer function—not from a continuous-time approximation, which would silently misrepresent the near-Nyquist behaviour. The DT NTF from SSPD noise to output phase error is:

$$\text{NTF}(z) = \frac{(1 - z_{\text{pole}})z^{-1}}{1 - z_{\text{pole}}z^{-1}}, \quad (3.11)$$

and the one-sided noise bandwidth evaluates in closed form as:

$$\text{NBW}(G_F) = \frac{f_{\text{REF}}}{2} \cdot \frac{1 - z_{\text{pole}}}{1 + z_{\text{pole}}} \quad [\text{Hz}], \quad (3.12)$$

where $z_{\text{pole}} = (1 + \beta - \alpha)/(1 + \beta)$ and $\beta = K_{\text{SSPD}} \cdot G_F$. At steady state ($G_F = 3.0 \text{ rad/V}$, $z_{\text{pole}} = -0.4451$):

$$\text{NBW}_{\text{ss}} = \frac{100 \times 10^6}{2} \cdot \frac{1 - (-0.4451)}{1 + (-0.4451)} = 130.2 \text{ MHz}. \quad (3.13)$$

Equation (3.12) contains a result that is easy to overlook: because $z_{\text{pole}} < 0$ when $G_F < G_{F,\text{crit}}$, the oscillatory baseline pole amplifies noise near the Nyquist frequency and pushes the NBW up to 130.2 MHz. Figure 3.9 plots the squared NTF magnitude over the Nyquist band for the baseline ($z_{\text{pole}} = -0.4451$) and for $G_{F,\text{crit}}$ ($z_{\text{pole}} = 0$). The baseline curve climbs sharply toward the Nyquist edge, peaking at 6.78 relative to unity, whereas the $G_{F,\text{crit}}$ curve is perfectly flat (all-pass) and removes that near-Nyquist amplification altogether.




Figure 3.9: Squared VCO noise transfer function $|H_{\text{NTF}}(f)|^2$ vs. offset frequency. Baseline ($G_F = 3.0$, $z_{\text{pole}} = -0.4451$, solid): near-Nyquist amplification to 6.78 (8.3 dB at 49.5 MHz) inflates the effective noise bandwidth to 130.2 MHz. $G_{F,\text{crit}}$ ($z_{\text{pole}} = 0$, dashed): all-pass NTF with unity gain across the entire Nyquist band, reducing NBW to 50.0 MHz. Peak label indicates the dB value at the Nyquist edge for the baseline.

Raising G_F toward $G_{F,\text{crit}} = 5.226 \text{ rad/V}$ ($z_{\text{pole}} \rightarrow 0$) makes the NTF all-pass and drives NBW toward its minimum:

$$\text{NBW}_{\text{crit}} = \frac{f_{\text{REF}}}{2} \cdot 1 = 50.0 \text{ MHz} \quad (z_{\text{pole}} = 0). \quad (3.14)$$

That is a 61.6% drop from the baseline—the reverse of what a continuous-time analysis would have you believe. Since jitter scales as $\sqrt{\text{NBW}}$, this works out to a $\sqrt{130.2/50.0} = 1.613\times$ jitter improvement at $G_{F,\text{crit}}$. The adaptive law leaves no steady-state penalty behind: because $G_F(|\phi_e| \rightarrow 0) = G_{F,\text{ss}}$, the NBW climbs back to 130.2 MHz at lock—exactly the fixed-baseline value.

Figure 3.10 shows both metrics as bar charts for the three configurations. The proposed design’s steady-state NBW and jitter (rightmost bars) match the baseline exactly, since the scheduler is back at $G_{F,\text{ss}} = 3.0 \text{ rad/V}$ by lock. The $G_{F,\text{crit}}$ column marks the jitter floor that would be reached if the loop were held permanently at $z_{\text{pole}} = 0$ —an analytical reference point, not an operating mode we propose.




Figure 3.10: Noise bandwidth (left) and RMS jitter (right) for the three SSPLL configurations: baseline / base-paper ($G_F = 3.0$), $G_{F,\text{crit}} = 5.226$ (analytical jitter floor, marked †), and proposed adaptive (steady state = $G_{F,\text{ss}} = 3.0$). NBW values from the exact DT formula Eq. (3.12); jitter from $J_{\text{rms}} = \sqrt{S_\phi \cdot \text{NBW}} / (2\pi f_{\text{VCO}})$. Percentage annotations show the NBW and jitter reduction from baseline to $G_{F,\text{crit}}$.

3.6.2 Phase-Noise Sources and Transfer Functions

Figure 3.11 shows the four noise injection points in the SSPLL loop.

Table 3.4 lists the noise floor assumptions used in analytical integration. The VCO floor of -85 dBc/Hz at 1 MHz is a representative value for a 2.4 GHz ring oscillator in 45 nm CMOS; measurements from the base-paper silicon ([7]) report 0.934 ps, which anchors the model.

Figure 3.12 plots the free-running VCO phase noise $\mathcal{L}(\Delta f) = 10 \log_{10}(S_\phi/2)$, where $S_\phi(\Delta f) = 6320/\Delta f^2 \text{ rad}^2/\text{Hz}$. The $1/f^2$ slope yields $\mathcal{L}(1 \text{ MHz}) = -85 \text{ dBc/Hz}$, consistent with published ring-VCO measurements at 2.4 GHz in 45 nm CMOS.

Figure 3.13 gives the closed-loop phase noise for the baseline and the $G_{F,\text{crit}}$ configuration, obtained by weighting the free-run VCO noise with $|H_{\text{NTF}}(f)|^2$. With its




Figure 3.11: Phase-noise injection points in the adaptive SSPLL. The four sources n_{REF} , n_{SSPD} , n_{VCDL} , and n_{VCO} (red arrows) enter the loop at distinct nodes. Reference and SSPD noise pass through the low-pass function $H_{\text{REF}}(\omega) = L/(1+L)$; VCO and VCDL noise are shaped by the high-pass function $H_{\text{VCO}}(\omega) = 1/(1+L)$.

Table 3.4: Phase-noise floor assumptions used in the analytical noise model.

Source	Level at 1 MHz offset	Spectral shape
Ring VCO	−85 dBc/Hz	$1/f^2$
Reference clock	−168 dBc/Hz	$1/f^2$
VCDL	−145 dBc/Hz	$1/f^2$
SSPD (thermal)	−155 dBc/Hz	flat




Figure 3.12: Free-running ring VCO phase noise ($S_\phi = 6320/\Delta f^2 \text{ rad}^2/\text{Hz}$). The −20 dB/decade slope is annotated; the reference point at $\Delta f = 1$ MHz gives $\mathcal{L}(1 \text{ MHz}) = -85 \text{ dBc/Hz}$. This trace is used in all analytical jitter estimates of Table 3.5 and in the closed-loop phase-noise model of Figure 3.13.

narrower NBW (50 vs. 130 MHz), the $G_{F,\text{crit}}$ design lets the high-pass NTF start attenuating VCO noise at a lower offset, which pulls the in-band noise floor down. The vertical markers at 50 and 130 MHz show each configuration's effective noise-bandwidth limit.




Figure 3.13: Closed-loop phase noise (VCO contribution only) for the baseline ($G_F = 3.0$, solid) and the $G_{F,\text{crit}} = 5.226$ configuration (dashed), together with the free-run VCO noise (dotted) and estimated reference floor (dash-dot). Vertical markers indicate the noise bandwidth limits: $\text{NBW}_{\text{ss}} = 130.2$ MHz (baseline) and $\text{NBW}_{\text{crit}} = 50.0$ MHz ($G_{F,\text{crit}}$). The proposed adaptive design has identical closed-loop noise to the baseline at lock (since $G_F \rightarrow G_{F,\text{ss}}$ at steady state).

VCO noise is shaped by $H_{\text{VCO}}(\omega) = 1/(1 + L(j\omega))$; reference and SSPD noise pass through $H_{\text{REF}}(\omega) = L(j\omega)/(1 + L(j\omega))$. Integrated RMS jitter follows from:

$$\sigma_{\text{jitter}} = \frac{1}{2\pi f_{\text{VCO}}} \sqrt{2 \int_{1\text{kHz}}^{100\text{MHz}} S_{\text{out}}(f) df}. \quad (3.15)$$

Because the fixed and adaptive SSPLLs both sit at $G_{F,\text{ss}} = 3.0$ rad/V once locked, this analytical model gives them the same steady-state jitter.

3.6.3 Jitter Comparison Model

Table 3.5 lines up the analytical jitter estimates for the three configurations. The Integer- N CP-PLL column starts from the same noise floors but adds the $N^2 = 576$ divider-noise multiplication penalty.

Table 3.5: Jitter comparison across configurations and noise sources. SSPD values are analytical (exact DT NBW); VCO values are Cadence Spectre transient-noise simulation.

Configuration	SSPD noise J_{rms} (ps)	VCO noise J_{rms} (ps)	Source
Base-paper SSPLL (silicon)	0.934	—	[7]
Baseline ($G_F = 3.0$, analytical)	1.229	2.596	Eq. (3.12), Spectre
Fixed $G_{F,\text{crit}} = 5.226$ (analytical)	0.762	2.292	Eq. (3.12), Spectre
Adaptive s.s. ($G_F = G_{F,\text{ss}}$)	1.229	2.560	Same as baseline
$G_{F,\text{crit}}$ vs. baseline	−38.0%	−11.7%	Exact DT / Spectre
Integer- N CP-PLL (analytical)	≈ 6.0	—	N^2 model

Two noise sources must be treated separately because they excite the loop through different transfer functions. **SSPD sampling noise** (kT/C thermal) is shaped by the NTF (3.11) and its integrated power equals $S_n \cdot \text{NBW}$; jitter scales as $\sqrt{\text{NBW}}$, giving a 38.0% reduction at $G_{F,\text{crit}}$. **VCO frequency noise** is shaped by a different DT function whose steady-state variance is $\sigma_w^2/(1 - z_{\text{pole}}^2)$, yielding only an 11.7% reduction at $G_{F,\text{crit}}$. Both improvements occur at the same operating point (G_F, crit , $z_{\text{pole}} = 0$), and both are zero for the adaptive design at lock, since G_F returns to $G_{F,\text{ss}}$. Spectre transient-noise simulation (VCO noise $S_f = 6320 \text{ Hz}^2/\text{Hz}$, injected as ctrl-line noise $S_v = 4.779 \times 10^{-13} \text{ V}^2/\text{Hz}$) confirms all three VCO jitter values in Table 3.5 to within 1.5% of the DT prediction.

3.7 PARAMETER SELECTION AND DESIGN TRADE-OFFS

The discrete-time model of Section 3.3 reduces the whole design to a small set of numbers, but those numbers are not free—they are pinned down by the target frequencies, the chosen oscillator, and the stability boundary. This section walks through how each one was fixed, because the reasoning matters as much as the final value when the design is later ported or re-simulated.

3.7.1 Fixing the Loop-Gain Operating Point

Everything begins with the dimensionless loop gain $\alpha = K_{\text{SSPD}} K_{\text{VCO}} T_s$. Three of its four ingredients are dictated by the specification rather than by choice: the reference period $T_s = 1/f_{\text{REF}} = 10 \text{ ns}$ follows from the 100 MHz reference, and the VCO gain $K_{\text{VCO}} = 2\pi \times 115 \text{ MHz/V} = 722.566 \text{ Mrad/s/V}$ follows from the 2.4 GHz ring oscillator and its tuning slope. The only truly tunable parameter of the detector is the gain $K_{\text{SSPD}} = 0.5 \text{ V/rad}$, which is determined by the sampling capacitor and the slope of the recorded VCO waveform at the zero crossing. With these values, $\alpha = 3.6128$. Since α is already above the bare-loop stability limit of 2, the VCDL is not optional, it is the only thing that prevents the design from having an unstable pole.

The main design leverage thus is the VCDL gain G_F . The discrete-time pole is given by:

$$z_{\text{pole}} = \frac{1 + \beta - \alpha}{1 + \beta}, \quad \beta = K_{\text{SSPD}} G_F, \quad (3.16)$$

and three operating points are obtained naturally from it. The stability edge occurs at $z_{\text{pole}} = -1$ with the minimum useable gain $G_{F,\text{min}} = 1.613 \text{ rad/V}$, below which the loop rings without settling. For $z_{\text{pole}} = 0$, we find the critical gain $G_{F,\text{crit}} = (\alpha - 1)/K_{\text{SSPD}} = 5.226 \text{ rad/V}$, for which the error is zero in one step (deadbeat response). Just beyond that, at $G_{F,\text{max}} = 5.509 \text{ rad/V}$, the pole is at a modest positive value ($z = +0.038$) which is as close to one-shot settling as the design dares to go before overshoot returns. The steady-state gain is kept at $G_{F,\text{ss}} = 3.0 \text{ rad/V}$ to match the base-paper operating point, which maintains the jitter comparison fair. So the whole adaptive notion is to slide G_F from near $G_{F,\text{max}}$ in acquisition down to $G_{F,\text{ss}}$ at lock, riding the pole from deadbeat to the reference design point.




Figure 3.14: Adaptive gain profile $G_F(|\phi_e|)$. Both the scheduled gain and error are substantial when they are near $G_{F,\text{max}}$, and then the error rolls smoothly down to $G_{F,\text{ss}} = 3.0 \text{ rad/V}$ as the error enters the lock band, making the transition continuously differentiable.

3.7.2 Choosing the Schedule Transition

The scheduling law has only one shaping parameter, the transition error ϕ_0 . Two competitive pressures pull ϕ_0 in opposite directions. If it is too large, the loop departs the high-gain regime before the phase error is small, and there is no improvement in acquisition speed. A value that is too small leaves the loop stuck in high gain all the way down into the lock band, where residual noise on ϕ_e can push G_F up briefly and restore the very oscillation the schedule was supposed to inhibit.

The value $\phi_0 = 0.10 \text{ rad}$ is intermediate between both extremes: it keeps the loop near-critically damped for practically the entire transient, and then lets G_F gently relax to $G_{F,\text{ss}}$ before the error hits the few-millivolt lock band. And crucially this is not a knife edge choice. As demonstrated in Chapter 4, lock time is practically flat within a factor-of-two neighborhood surrounding this value, so a modest inaccuracy in ϕ_0 due to process variation or a coarse implementation costs nothing in speed.

Smooth differentiability is equally important. A hard switch between two gain values would inject a step into v_{ctrl} at the transition instant, and the VCO would faithfully convert that step into a phase glitch—undoing in one cycle what the schedule spent

several cycles achieving. The saturating hyperbolic form of Section 3.4 sidesteps this entirely: G_F rolls gradually between $G_{F,\max}$ and $G_{F,ss}$, with no jump in either the gain or its first derivative. Charge-pump designers have relied on the same logic for decades—self-biased and adaptive-bandwidth loops deliberately avoid hard parameter transitions precisely because abrupt changes disturb the settled state [16, 17].

3.8 SIMULATION METHODOLOGY AND REPRODUCIBILITY

A recurring frustration with behavioural PLL studies is that results are hard to reproduce: small differences in solver settings or netlist regeneration quietly shift the measured lock time. The methodology below was built to remove that ambiguity, so that the numbers reported in Chapter 4 can be regenerated exactly.

3.8.1 Testbench Construction and Solver Configuration

The testbench instantiates all five VerilogA blocks—the sub-sampling detector, the VCDL, the ring VCO, the adaptive scheduler, and the lock-time detector—in a single Spectre transient analysis. Getting the solver settings right matters more than it might seem.

The timestep is chosen to be small enough compared to the VCO period. A coarser step generates 2.4 GHz waveform aliasing, when the solver steps straight across the sampling edge, completely corrupting the detector output. The initial conditions also require special consideration, the transient starting from an initial VCO phase of π rad with the control voltage at its nominal value, locking the loop into a known state before the step is applied. The spurious phase error occurs at the solver’s default DC operating point, contaminates the first few reference cycles and inflates the recorded lock time by several nanoseconds.

Netlisting and simulation are always independent processes — never the combined “Netlist and Run” command. The rationale is a practical one: Sometimes Cadence ADE resets some parameters to stale defaults when regenerating the netlist from schematic. Then if you ran the simulation straight after netlisting you would stimulate the wrong design without notice. A fixed patch is inserted in between the steps, separating the steps and ensuring that the as-simulated values are always the intended ones: $G_{F,\max} = 5.509$ rad/V, sampling window $T_s = 100$ ns, nominal control voltage at 0 V.

3.8.2 Sweep Strategy and Lock-Time Extraction

Lock time is not a read off the eyes. The lock-time detector module constantly monitors the sampled phase-error voltage and flags the instant when $|\phi_{e,v}|$ falls below, and then remains below, a predetermined threshold of 5 mV (a tolerance of 2% of the detector’s full-scale swing). What makes the measurement robust against the oscillating baseline (whose error crosses zero repeatedly on its way to settling) is the requirement that the error stay inside the band, rather than just touch it. The stated lock time is the timestamp of the last crossing of the threshold before bounding the mistake.

We inject two phase-step amplitudes to test small and big signal behaviour: $\pi/4 = 0.7854$ rad and $\pi/2 = 1.5708$ rad. The smaller step keeps the detector well within

its linear region, testing settling dynamics, while the bigger step probes the edge of the linear $\sin(\phi_e)$ characteristic, stressing the high-gain phase of the program. The transient is repeated for both phases for each configuration—baseline fixed gain and suggested adaptive gain—providing the four lock-time numbers that the comparison in Chapter 4 is based on. Because each run starts from the same defined initial condition and applies the same threshold logic, the four numbers are immediately comparable and reproducible to the timestep resolution of the simulator.

3.8.3 VCDL Transfer Function and Realisation

The voltage-controlled delay line is the key element in the functioning of the complete scheme and the behaviour of this hence demands more attention than the block diagram can give. Functionally, the VCDL shifts the instant at which the sub-sampling detector evaluates the VCO waveform by an amount proportional to its control input. In the small-signal, once-per-cycle picture used throughout this chapter, that delay translates into a multiplicative correction on the sampled phase error: the effective error presented to the integrator is reduced by the factor $(1 + \beta)$, with $\beta = K_{SSPD} G_F$. This is the algebraic origin of the $(1 + \beta)$ term in the denominator of the pole expression in Eq. (3.16), and it is why raising G_F pulls the pole inward toward the origin.

What makes the line useful for scheduling is that G_F is controllable rather than fixed. In the behavioural model the adaptive scheduler drives a second control port, setting the effective gain fresh every reference cycle—the same physical element acts as a strong corrector during acquisition and a mild one at lock.

The bias for the delay stage would be taken straight from the scheduler output in a transistor level implementation. Two requirements are important in silicon: the delay must be a monotonic and smooth function of control voltage across the entire needed range, and the noise contribution of the delay cell itself must be well below the oscillator floor, so that bandwidth extension does not come at a jitter cost. The behavioural study incorporates both issues into the same parameter G_F , therefore preserving the stability and lock-time conclusions irrespective of any specific delay-cell topology. The treatment of G_F as an abstract scalar is deliberate and the results are applicable to any VCDL satisfying the monotonicity and noise criteria. However, this is also a real limitation. In Chapter 5 we return to what would need to be checked in a transistor-level implementation.

3.9 THE SUB-SAMPLING DETECTOR AND ITS LINEAR RANGE

The main architectural advantage of the SSPLL is the elimination of the divider enabled by the sub-sampling phase detector. The detector samples the VCO waveform at the time of each reference edge, and then retains this value for the duration of the cycle. The output of the VCO is sinusoidal so that the sampled voltage is $K_{SSPD} \sin(\phi_e)$ and not $K_{SSPD} \phi_e$ which is linear only for small phase errors and folds back once the error exceeds $\pi/2$. That nonlinearity creates a hard limit on the operational range and everything in this chapter lives within that boundary.

Two practical repercussions emerge for the lock time results. At $\pi/4$ the detector is well inside its linear range – $\sin(\phi_e) \approx \phi_e$ to a few percent – thus measured and

anticipated cycle counts correspond well. For $\pi/2$ the story is different, the step is exactly on the border of the linear range and sine compression marginally reduces the effective detector gain on the first cycle. This is part of the reason the large-step results should be taken as the behaviour of the loop at the edge of its useful input, not as indicative of small-signal operation. Separately, the once-per-cycle hold gives the staircase seen in every transient waveform. The control voltage is updated only on reference edges, and constant in between them, which is precisely the sampled-data behaviour which requires a discrete-time analysis here, and not a continuous-time one.

A practical result is that the detector gain K_{SSPD} is not a free constant but relies on the amplitude and the slope of the recorded VCO waveform at the zero crossing. In the behavioral model this is directly represented by the amplitude and mid-level parameters of the detector and oscillator blocks; in a transistor-level architecture it would be set by the sampling capacitor and the VCO swing. For all simulations, we set $K_{SSPD} = 0.5 \text{ V/rad}$ so that α and all calculated operating points are constant, allowing all four lock-time measurements to be directly compared. Any modification of the VCO swing would rescale K_{SSPD} , shift α , and move all three gain operating points. This parameter is considered fixed specification rather than something for the schedule to touch.

3.10 WHY CONTINUOUS-TIME INTUITION MISLEADS HERE

The base-paper CT phase margin of about 30° is accurately obtained from Eq. (2.1). The problem is not the arithmetic, but is that CT phase margin is the wrong question to ask of this loop.

These continuous-time stability arguments apply when the loop bandwidth is a small fraction of the sampling rate, thus there are numerous update events per time constant, and the sampled nature averages out. Here, that premise silently falls apart. The loop bandwidth is not a small fraction of the Nyquist rate, it is a significant fraction $\alpha = 3.6128$ and the loop settles in a couple of reference cycles only. In that regime the mapping from the s -plane to the z -plane is not gentle any more and a healthy CT phase margin can coexist with a pole on the negative real axis in the z -plane. And this is in fact what happens, the CT model shows stable, properly damped behaviour, and the DT pole at $z = -0.4451$ shows oscillatory, sign-alternating settling behaviour. They are both the same circuit. Only the second is what Spectre outputs.

The takeaway generalises. Any sub-sampling loop with α approaching or exceeding unity should be analysed in discrete time from the outset, with the $|z| < 1$ test applied directly to the closed-loop pole. CT phase margin can serve as an intuitive, secondary description—but not as the criterion on which stability is decided. This methodological point underlies the noise-bandwidth results of Section 3.6, since the same negative-real pole the CT model hides is also what amplifies near-Nyquist noise in the baseline.

The lesson generalises beyond this particular synthesiser. Any sub-sampling loop with α approaching or exceeding unity should be analysed in discrete time from the outset, with the $|z| < 1$ test applied directly to the closed-loop pole. The continuous-time phase margin can then be reported as a secondary, descriptive quantity—useful for




Figure 3.15: Discrete-time pole trajectory in the z -plane as G_F varies. The baseline pole sits at $z = -0.4451$ on the negative real axis (oscillatory); raising G_F toward $G_{F,crit}$ walks the pole to the origin (deadbeat), the operating region the adaptive schedule exploits during acquisition.

intuition, but not the criterion on which stability is decided. This is the methodological point that the noise-bandwidth results of Section 3.6 ultimately depend on, since the same negative-real pole that the CT model hides is also what amplifies near-Nyquist noise in the baseline.

3.11 SUMMARY

Three critical gain values define the adaptive SSPLL design space: $G_{F,min} = 1.613$, $G_{F,crit} = 5.226$, and $G_{F,max}^{opt} = 5.509$ rad/V. Derived from first principles, the DT pole $z_{pole} = (1 + \beta - \alpha)/(1 + \beta)$ ties them all together. Under the adaptive scheduling law (3.7) the pole starts near $z \approx +0.037$ during acquisition and relaxes to $z = -0.4451$ at lock; the Lyapunov proof guarantees $|z_{pole}| < 1$ throughout, with a $1.86 \times$ PVT stability margin. The noise-bandwidth model predicts zero steady-state jitter pena

CHAPTER 4

SIMULATION RESULTS AND DISCUSSION

Four simulation runs anchor this chapter: baseline and adaptive, $\pi/4$ and $\pi/2$ step. Section 4.1 describes the test configuration and lock-time measurement method. Section 4.2 characterises the fixed-gain baseline. Section 4.3 presents the adaptive- G_F results for both step amplitudes, and Section 4.4 consolidates the side-by-side comparison. Section 4.6 draws out the physical implications.

4.1 SIMULATION ENVIRONMENT AND TEST CONFIGURATION

4.1.1 Spectre Netlist and Solver Settings

Each Spectre 15.1.0 transient run covers a 500 ns window with `maxstep` = 20 ps and `errpreset` = moderate. Both configurations share one and the same netlist; only the `adaptive_gf_scheduler` parameters differ between them. Table 4.1 sets out both test configurations.

Table 4.1: Spectre simulation configurations for the baseline and proposed SSPLL.

Parameter	Baseline (fixed G_F)	Proposed (adaptive G_F)
Scheduler mode	Fixed $G_{F,ss} = 3.0$ rad/V	Adaptive, Eq. (3.7)
$G_{F,max}$	N/A	5.509 rad/V
$G_{F,ss}$	3.0 rad/V	3.0 rad/V
ϕ_0	N/A	0.10 rad
z_{pole} (initial)	-0.4451 (fixed)	+0.037 at $\pi/4$ step
z_{pole} (at lock)	-0.4451	-0.4451
Phase step applied	$\pi/4, \pi/2$ rad	$\pi/4, \pi/2$ rad
Step applied at	$t_{step} = 100$ ns	$t_{step} = 100$ ns
Lock threshold	5 mV	5 mV

4.1.2 Lock-Time Measurement Methodology

Lock time is defined as the first instant after $t_{step} = 100$ ns at which $|\phi_{e,v}(t)| < 5$ mV and remains below that threshold. The 5 mV figure is a 2% settling criterion against the nominal SSPD output for a $\pi/4$ -radian step ($K_{SSPD} \cdot \pi/4 \approx 393$ mV full-scale). The lock instant is extracted by reading the $\phi_{e,v}$ node voltage at each reference clock edge ($t = 100 + k \times 10$ ns for $k = 1, 2, \dots$) from the Spectre `psf/spectre.out` log, and taking the first edge whose sampled magnitude stays below 5 mV.

Two numbers are reported per run: the threshold-crossing time T_{lock} (in nanoseconds from t_{step}) and the SSPD output voltage at each reference cycle, read straight off the transient waveform.

4.1.3 Phase-Step Injection

The `phase_step_src` VerilogA module applies the phase step at $t = 100$ ns by advancing its internal oscillator phase in a single jump. The ring VCO is deliberately set with $\phi_{\text{step}} = 0$ to prevent the VCO-side model from adding a second disturbance at the same instant—a well-known double-injection artefact in ring-VCO behavioural models that would otherwise inflate the initial phase error.

4.2 BASELINE SSPLL CHARACTERISATION

4.2.1 Phase-Error Trajectory: $\pi/4$ Step

With the adaptive scheduler disabled and $G_F = 3.0$ rad/V held fixed, the DT pole sits at $z = -0.4451$. The phase-error voltage $\phi_{e,v}$ alternates sign at every reference cycle, exactly as the negative-real pole predicts it should. Table 4.2 records the SSPD output voltage at each reference edge after the step is applied.


Table 4.2: SSPD output voltage $\phi_{e,v}$ at each reference cycle, baseline SSPLL, $\pi/4$ phase step. Lock threshold = 5 mV.

Cycle k	$\phi_{e,v}$ (mV)	Within 5 mV band
1	+86.75	No
2	−98.40	No
3	+43.75	No
4	−19.45	No
5	+8.65	No
6	−3.85	Yes


It takes exactly six reference cycles before $|\phi_{e,v}|$ first drops below 5 mV. The Spectre lock-time detector reports $T_{\text{lock}} = 48.785$ ns, placing the lock instant 4.88 reference periods past the step. This sits very close to what the DT pole predicts: for $|z_{\text{pole}}| = 0.4451$ and a 2% settling criterion, the expected cycle count is $n = \log(0.02)/\log(0.4451) \approx 4.8$ cycles.

Figure 4.1 collects the Cadence Spectre transient waveforms of $\phi_{e,v}$ for all four test cases, exported directly from the Virtuoso waveform viewer. The zero-order-hold staircase in each panel reflects the SSPD sampling once per reference cycle and holding its output constant for the full 10 ns until the next update. The baseline panels (left column) clearly show the sign-alternating oscillation that is the hallmark of the negative-real pole at $z = -0.4451$. The proposed adaptive panels (right column) confirm the near-monotone, fast convergence the scheduling law predicts, with the $\pi/2$ case coming remarkably close to settling in a single shot.

Cycle 2 shows a 13.5% overshoot relative to cycle 1 (−98.40 mV versus +86.75 mV). This is exactly what a negative-real pole does: the VCDL correction




(a) Baseline ($G_F = 3.0$ rad/V fixed): $T_{\text{lock}} = 48.785$ ns. Oscillatory convergence — the negative-real pole at $z = -0.4451$ produces alternating-sign SSPD samples on every reference cycle. Cursor V2 marks the lock instant.




(b) Proposed adaptive (G_F from 5.218 to 3.0 rad/V): $T_{\text{lock}} = 28.785$ ns ($1.70\times$ faster). Near-monotone convergence — scheduler drives the pole to $z \approx +0.037$ at the first cycle, then relaxes to $z = -0.4451$.

Figure 4.1: Spectre transient waveforms of v_{ϕ_e} for $\Delta\phi = \pi/4$ phase step (exported from Cadence Virtuoso 6.1.7). Left cursor: step instant ($t = 100$ ns). Right cursor: T_{lock} instant. Horizontal dotted line: ± 5 mV lock threshold (2% criterion). Values read directly from `psf/spectre.out`.



(a) Baseline ($G_F = 3.0 \text{ rad/V}$ fixed): $T_{\text{lock}} = 27.525 \text{ ns}$ (3 cycles). The larger step amplitude drives a higher initial error but the oscillatory pole still requires multiple sign reversals before settling.



(b) Proposed adaptive: $T_{\text{lock}} = 7.525 \text{ ns}$ ($3.66\times$ faster, <1 reference cycle). The scheduler evaluates $G_F \approx 5.509 \text{ rad/V}$ at the step, placing $z \approx +0.038$ — near one-shot settling confirmed by Spectre.

Figure 4.2: Spectre transient waveforms of v_{ϕ_e} for $\Delta\phi = \pi/2$ phase step. Same measurement setup as Figure 4.1. The proposed design achieves near one-shot settling at twice the design step amplitude, demonstrating that the adaptive schedule scales gracefully with input disturbance.

overshoots the lock point on the first cycle, the error flips sign instead of decaying monotonically, and the pattern repeats with diminishing amplitude. By cycle 6 the envelope has finally decayed to -3.85 mV and the loop crosses into the lock band.

4.2.2 Phase-Error Trajectory: $\pi/2$ Step

For a $\pi/2$ -radian step—twice the design amplitude—the baseline SSPLL behaves differently, as shown in Table 4.3.

Table 4.3: SSPD output voltage $\phi_{e,v}$ at each reference cycle, baseline SSPLL, $\pi/2$ phase step. Lock threshold = 5 mV.

Cycle k	$\phi_{e,v}$ (mV)	Within 5 mV band
1	+86.75	No
2	-7.24	No
3	+3.22	Yes
4	-1.43	Yes

Lock is declared at cycle 3, and Spectre reports $T_{\text{lock}} = 27.525$ ns (≈ 2.75 reference cycles). Notice that the initial SSPD voltage at cycle 1 is +86.75 mV—the same as for the $\pi/4$ step. This is not a coincidence: the VCDL correction dominates the first-cycle response, and the amplitude of the initial sample depends more on how much control voltage the VCDL can supply in that cycle than on the raw size of the disturbance. For the $\pi/2$ step the correction overshoots less aggressively, and the error reaches the 5 mV band in just three cycles.

4.2.3 Baseline Lock-Time Summary

The baseline SSPLL locks in 48.785 ns on the $\pi/4$ step and 27.525 ns on the $\pi/2$ step. Both measurements come directly from the Spectre spectre.out log. The DT pole analysis of Chapter 3 correctly predicts the oscillatory settling pattern and the approximate cycle count; the Spectre values give the precise quantification used in the performance table.

4.3 PROPOSED ADAPTIVE-GF SSPLL RESULTS

4.3.1 Phase-Error Trajectory: $\pi/4$ Step

With adaptive gain scheduling active, the scheduler evaluates $G_F(\pi/4) = 3.0 + 2.509 \times 0.785 / (0.785 + 0.10) = 5.218$ rad/V at the very first reference cycle after the step. This places the DT pole at $z_{\text{pole}} = +0.037$ —positive and close to zero—producing a near-monotone decay instead of the baseline’s sign-alternating oscillation. Table 4.4 records the SSPD samples Spectre reports.

Spectre reports $T_{\text{lock}} = 28.785$ ns—exactly 20 ns less than the baseline. Four cycles are enough rather than six, giving an improvement ratio of $48.785/28.785 = 1.70\times$.

Cycle 1 shows only +20.38 mV, substantially smaller than the baseline’s +86.75 mV on the same step. The adaptive VCDL immediately applies

Table 4.4: SSPD output voltage $\phi_{e,v}$ at each reference cycle, proposed adaptive SSPLL, $\pi/4$ phase step. Lock threshold = 5 mV.

Cycle k	$\phi_{e,v}$ (mV)	Within 5 mV band
1	+20.38	No
2	−77.85	No
3	+9.40	No
4	−2.95	Yes

$G_F = 5.218$ rad/V and pushes the error down hard on the first half-cycle. The sign reversal at cycle 2 (−77.85 mV) reflects a single overshoot—the high gain crosses through zero once. By cycle 3 the scheduler has already backed the gain down in response to the smaller error at cycle 2, and cycle 4 settles neatly to −2.95 mV inside the lock band.

4.3.2 Phase-Error Trajectory: $\pi/2$ Step

For the $\pi/2$ step, the scheduler evaluates $G_F(\pi/2) = 3.0 + 2.509 \times 1.571 / (1.571 + 0.10) = 5.375$ rad/V at the first cycle—even closer to the deadbeat point $G_{F,\text{crit}} = 5.226$ rad/V than at the $\pi/4$ case. Table 4.5 records the Spectre samples.

Table 4.5: SSPD output voltage $\phi_{e,v}$ at each reference cycle, proposed adaptive SSPLL, $\pi/2$ phase step. Lock threshold = 5 mV.

Cycle k	$\phi_{e,v}$ (mV)	Within 5 mV band
1	+20.38	No
2	+0.76	Yes

Spectre reports $T_{\text{lock}} = 7.525$ ns, which is 0.75 reference periods from the step—effectively one-shot settling. The phase error is reduced from +20.38 mV to +0.76 mV in one cycle, a factor-of-27 reduction. This is a $3.66\times$ improvement over the baseline of 27.525 ns.

The full sampled-error sequences for all four runs are tabulated in Tables 4.2–4.5. The contrast is hard to miss reading across them. The baseline flips sign every cycle, each flip decaying only slightly. The adaptive design decreases to virtually zero in one or two steps and stays there. The same concept is shown graphically in Figure 4.1.

The cycle-by-cycle G_F trajectory for the $\pi/4$ adaptive run is shown in Figure 4.3. As the error vanishes, the gain relaxes to $G_{F,\text{ss}} = 3.0$ rad/V. The baseline trace is flat at 3.0 rad/V all the way through, which tells you all you need to know about why it settles more slowly.




Figure 4.3: Cadence Spectre waveform of the adaptive G_F scheduler output v_{G_F} for the proposed SSPLL during a $\pi/4$ phase step (exported from Virtuoso 6.1.7). The scheduler starts with $G_F \approx 5.218$ rad/V during the first reference cycle (high error), then smoothly reduces G_F towards $G_{F,ss} = 3.0$ rad/V as the phase error decays, to return to steady state within four reference cycles.

4.4 COMPARATIVE PERFORMANCE ANALYSIS

4.4.1 Lock-Time Comparison

Table 4.6 summarizes the four Spectre lock-time results. The adaptive scheduler always reduces lock time, the benefit increases with step amplitude since greater errors more forcefully activate the schedule, bringing the DT pole closer to the one-shot point.

Table 4.6: Spectre-measured lock times for the baseline and proposed SSPLL.

Phase Step	Baseline T_{lock}	Proposed T_{lock}	Improvement
$\pi/4$ (design point)	48.785 ns	28.785 ns	$1.70\times$
$\pi/2$ ($2\times$ design)	27.525 ns	7.525 ns	$3.66\times$

The improvement is quantified in Table 4.6 where the $1.70\times$ benefit at $\pi/4$ increases to $3.66\times$ at $\pi/2$, a direct reflection of the schedule that is pushing the DT pole closer to $G_{F,\text{crit}}$ for higher step sizes. Figure 4.4 shows the same data as a three-way bar chart, with a traditional integer- N CP-PLL reference for larger context.




Figure 4.4: Lock-time benchmark for $\Delta\phi = \pi/4$ (left) and $\Delta\phi = \pi/2$ (right). Three configurations are compared: a conventional integer- N charge-pump PLL (hatched, included as a broad reference), the base-paper fixed- G_F SSPLL equivalent to our baseline ($G_F = 3.0$, white bar), and the proposed adaptive- G_F SSPLL (grey, cross-hatched). Lock times for the SSPLL configurations are confirmed Spectre measurements. The percentage annotations indicate the reduction from baseline to proposed; the double-headed arrow spans the two SSPLL bars. The integer- N value is estimated from the DT closed-loop model at $N = 24$, $f_n = 10$ MHz.

4.4.2 DT Pole Trajectory During Settling

The instantaneous DT pole $z_{\text{pole}}(k) = z_{\text{pole}}(G_F(|\phi_{e,v}[k]|))$ changes every cycle during the adaptive transient. Table 4.7 traces the pole position for the proposed SSPLL during the $\pi/4$ step, derived from the Spectre SSPD samples and the scheduling equation.

The pole starts near $+0.018$ (cycle 1), remains small and positive through cycle 2, then briefly turns slightly negative at cycle 3 as the scheduler reduces G_F in response

Table 4.7: Instantaneous DT pole during the adaptive settling transient, proposed SSPLL, $\pi/4$ phase step.

Cycle k	$\phi_{e,v}$ (mV)	$ \phi_e $ (rad)	G_F (rad/V)	z_{pole}
1	+20.38	0.0408	5.058	+0.0177
2	−77.85	0.1557	5.170	+0.0290
3	+9.40	0.0188	4.842	−0.0078
4	−2.95	0.0059	3.277	−0.4017

to the shrinking error. By cycle 4 the error is 5.9 mrad—just inside the lock band—and G_F has relaxed to 3.277 rad/V, close to $G_{F,ss}$. All four pole values satisfy $|z_{\text{pole}}| < 1$, confirming the Lyapunov stability proof of Section 3.5 under Spectre simulation conditions.

4.4.3 Noise Bandwidth Equivalence at Lock

Both designs have $G_F = G_{F,ss} = 3.0$ rad/V, therefore they have the same steady-state noise bandwidth. For both we have $\text{NBW}_{ss} = 130.2$ MHz using the precise DT formula (3.12). This is not a coincidence, it is baked into the scheduling law:

$$G_F(0) = G_{F,ss}$$

drives the loop back to its nominal operating point as soon as the error disappears. There is no jitter penalty at lock.

The noise picture at high gain is of separate interest. In fact, raising G_F towards $G_{F,\text{crit}}$ *narrows* the noise bandwidth rather than broadening it, as the oscillating base-line pole at $z_{\text{pole}} = -0.4451$ amplifies near-Nyquist noise. At $G_{F,\text{max}} = 5.509$ rad/V ($z_{\text{pole}} = +0.038$) NBW drops to 46.3 MHz, a 64.5% reduction of 130.2 MHz. Continuous operation at that gain would reduce jitter, but would use most of the VCDL's dynamic range at steady state, with no headroom for major disturbances. The adaptive schedule completely avoids this trade-off: G_F is increased only during acquisition, then reverted to $G_{F,ss}$ at lock.

4.5 SENSITIVITY AND ROBUSTNESS ANALYSIS

An improvement in lock-time that disappears under realistic parameter modification is not an advantage worth having. To explore the robustness of the numbers, there are three sweeps: the schedule transition ϕ_0 , the maximal gain $G_{F,\text{max}}$, and the injected phase-step amplitude.

4.5.1 Dependence on the Schedule Transition ϕ_0

Since the scheduling law has only one shaping parameter ϕ_0 , this is the first place to check for fragility. Figure 4.5 sweeps it across an order of magnitude and records the $\pi/4$ lock time at each position. The curve is essentially flat between around 0.05 and 0.20 rad; the chosen value of 0.10 rad is on a broad plateau, not a sharp peak. If 0 is

driven too small the loop sticks to high gain down to the lock band, where noise on E kicks GF up, and ringing returns. If you make it too large, the loop goes out of its high gain condition before the error is reduced, thus you don't gain any speed in acquisition. Outside these extremes the design suffers. Inside them any value within a factor of two of 0.10 rad gives approximately the same lock time. No need for precise trimming.




Figure 4.5: Lock-time dependence on the schedule transition ϕ_0 . The $\pi/4$ lock time is flat for a large neighbourhood of the design value $\phi_0 = 0.10 \text{ rad}$, showing that the schedule is not perched on a knife edge.

4.5.2 Sensitivity to the Maximum Gain $G_{\text{F,max}}$

The maximum gain $G_{\text{F,max}}$ determines the aggressiveness of the schedule with respect to the deadbeat pole during acquisition. $G_{\text{F,crit}} = 5.226 \text{ rad/V}$ puts the pole exactly at the origin; the design aims for $G_{\text{F,max}} = 5.509 \text{ rad/V}$, little above it. Lock time vs. ratio $G_{\text{F,max}}/G_{\text{F,crit}}$ is shown in Figure 4.6. The pattern is simple, as the ratio approaches one the pole advances in from the unit circle, lock time comes off sharply and the curve bottoms out near the critical point. Push past one and the pole turns slightly positive—overshoot reappears and adds cycles back. The minimum is broad, not sharp. Operating a few percent above critical, as this design does, captures nearly all of the available speed-up with no risk of entering the overshoot regime. A VCDL whose fabricated gain falls somewhat short of $G_{\text{F,max}}$ still lands on the flat part of the curve.

4.5.3 Robustness to Phase-Step Amplitude

The two step sizes simulated— $\pi/4$ and $\pi/2$ —probe opposite ends of the detector's usable range, and the gap between them is itself informative. For the smaller $\pi/4$ step the proposed loop locks in 28.785 ns against the baseline's 48.785 ns, a $1.70\times$ improvement. For the larger $\pi/2$ step the proposed loop locks in 7.525 ns against 27.525 ns,




Figure 4.6: Lock time versus the maximum-gain ratio $G_{F,\max}/G_{F,\text{crit}}$. The minimum is broad and centred near the critical (deadbeat) point, so the design tolerates moderate gain error without losing speed.

a $3.66\times$ improvement. The speed-up is therefore not a fixed factor: it grows with the size of the disturbance. This makes sense in light of the schedule, which spends more time in its high-gain regime when the error is large, so the adaptive advantage is most pronounced exactly when it is needed most—during a big acquisition transient. The baseline, by contrast, runs at fixed $G_{F,ss}$ regardless of error size and gains nothing from a larger step. A design that has to acquire from cold, where the initial phase error can be close to the full $\pi/2$, will see the larger of the two improvement factors in practice.

4.5.4 Comparison with Published Sub-Sampling Designs

Comparing against the broader sub-sampling literature takes some care. The 200 ps lock time of Nagam and Kinget [18] is of an injection-locked loop with a specific feed-forward noise-cancellation path and a separate reference frequency, i.e. a fundamentally different stabilising mechanism and not a like-for-like competitor. Kim’s dual-path design [11] achieves rapid settling by duplication of the signal chain, again adding hardware that the present work purposefully avoids.

The pertinent analogy is more limited. This work operates within the same single-path VCDL architecture as the base paper [7], at the same 100 MHz reference and 2.4 GHz output, changing only one parameter: the VCDL gain. On that like-for-like basis, the adaptive schedule delivers $1.7\times$ – $3.7\times$ faster locking with no extra detector, no second filter, and no injection path. Against designs with added hardware the comparison is unfair in their favour; against the unmodified base-paper loop it is unambiguous.

4.5.5 Reading the Noise-Bandwidth Result

The jitter comparison rests on a result that looks counter-intuitive at first glance, so it deserves a plain-language interpretation. One might expect the faster, higher-gain proposed loop to be the noisier one—wider loops usually let through more noise. What the discrete-time noise-transfer function shows is the opposite ordering between the baseline and the deadbeat operating point, and the reason is the sign of the pole. The baseline’s negative-real pole at $z = -0.4451$ peaks the noise-transfer function near half the reference frequency: the very sign-alternation that slows its settling also resonates with noise components close to Nyquist, inflating the integrated noise bandwidth to 130.2 MHz. The deadbeat pole at $z = 0$ has an all-pass-like transfer with no such peaking, which drops the noise bandwidth to 50.0 MHz.




Figure 4.7: Output phase-noise spectra for the baseline and deadbeat operating points. The baseline’s negative-real pole peaks the noise-transfer function near $f_{\text{REF}}/2$, raising the integrated noise bandwidth; the deadbeat pole removes the peaking, lowering both the bandwidth and the integrated jitter.

Carried through to timing jitter, that is a fall from 1.229 ps to 0.762 ps—about a 38% reduction—purely from moving the pole off the negative real axis.

The proposed design holds its steady-state gain at $G_{\text{F,ss}} = 3.0 \text{ rad/V}$ to match the base paper, so at lock it inherits the same 130.2 MHz noise bandwidth as the baseline; the jitter advantage of the deadbeat point is available only if one is willing to operate there continuously. The project report therefore makes a careful claim: the adaptive schedule buys lock-time improvement without any jitter penalty relative to the baseline, and the deadbeat analysis additionally shows that a fixed critical-gain design would trade some tuning headroom for lower jitter. These are two distinct findings, and conflating them would overstate the result.




Figure 4.8: Steady-state timing-jitter distribution at lock. The deadbeat operating point ($z_{\text{pole}} = 0$) yields a tighter distribution (0.762 ps RMS) than the oscillatory baseline ($z_{\text{pole}} = -0.4451$, 1.229 ps RMS), reflecting the lower discrete-time noise bandwidth.

4.5.6 Practical Design Guidelines

A few guidelines fall out of this analysis for anyone applying the method to a different sub-sampling loop. Start by computing $\alpha = K_{\text{SSPD}}K_{\text{VCO}}T_s$: if it exceeds two, a bandwidth-extension element is non-negotiable. From there, locate $G_{\text{F,crit}} = (\alpha - 1)/K_{\text{SSPD}}$ analytically—this is the deadbeat point and the natural ceiling for the acquisition phase. Set $G_{\text{F,max}}$ a few percent above $G_{\text{F,crit}}$, not far above it; the sensitivity sweep in Section 4.5 shows the lock-time minimum is broad enough that chasing the exact peak gains little while risking overshoot. The schedule decouples acquisition speed and settling noise performance completely, which allows one to select the steady state gain $G_{\text{F,ss}}$ individually, against the jitter or bandwidth specification. Keep the scheduling function smooth and differentiable all the while, any hard gain transition would put a step into v_{ctrl} that the VCO would turn into a phase glitch directly. None of these steps touch anything outside the VCDL the loop already

4.5.7 An Energy-Decay View of the Stability Margin

The Lyapunov proof in Section 3.5 assures that error energy $V = \phi_e^2$ reduces each cycle, although this simple guarantee understates the real margin. Statistically, it is more telling to read it. On each cycle the energy is multiplied by z_{pole}^2 . Near the deadbeat point $z_{\text{pole}} \approx 0$ and nearly all the error energy disappears in one step. At steady-state gain $z_{\text{pole}} = -0.4451$ the factor is $0.4451^2 \approx 0.198$, so 80% of the remaining energy is tossed out on each reference cycle. There is no place in the schedule where the energy is not decreasing—that is the concrete meaning of $|z_{\text{pole}}| < 1$ for all allowable mistakes.

The same illustration demonstrates why the adaptive loop never overshoots into instability as a fixed high gain loop would do. Moreover, if the gain is fixed above

$G_{F,\max}$, the pole is at a positive value high enough to overshoot, and if the gain is fixed below $G_{F,\min}$, the pole is past $z = -1$ and there is no contraction at all. The schedule is designed to prevent both failure modes—it always visits only gains between $G_{F,ss}$ and $G_{F,\max}$, a window across which $|z_{\text{pole}}| < 1$ from the first cycle to the last. Stability is thus a feature of the full trajectory traced by the scheduler, not of a particular point on it.

4.5.8 Consolidated Comparison and What It Shows

The four lock-time measurements and the noise figures together define what the adaptive schedule changes and what it does not change. The result at lock time is clean: $1.70\times$ quicker on the $\pi/4$ step (28.785 ns vs 48.785 ns), $3.66\times$ on the $\pi/2$ step (7.525 ns vs 27.525 ns), with the greater disturbance exhibiting the higher gain. The results for noise are neutral: the steady state gain is back to $G_{F,ss} = 3.0$ rad/V at lock, and the settled noise bandwidth is still 130.2 MHz and jitter 1.229 ps—the same as baseline. The speed-up comes with no sacrifice in noise performance.

The deadbeat analysis is a different result, not a free by-product. A loop held permanently at $G_{F,\text{crit}}$ would achieve 50.0 MHz noise bandwidth and 0.762 ps jitter—a 38% reduction—but that design option uses much of the VCDL's tuning headroom at steady state. These are two discrete operational points, with different trade-offs, and confusing the two would be an overstatement of the contribution of the adaptive schedule alone.

Two caveats constrained the claims. First, the results are from transistor-level circuits, not calibrated VerilogA behavioral models. They verify the loop dynamics and the scheduling principle but they do not, in themselves, prove silicon practicality; that stage needs a transistor-level implementation and measurement. Second, the comparison that matters is the like for like comparison: same architecture, same reference frequency, same output, one parameter modified. Injection-locked or dual-path designs use more hardware to facilitate faster settling; comparing to them is comparing to a greater budget. Within these bounds the conclusion is clean. Varying a single existing parameter on a per-cycle basis recovers quick, clean acquisition for a ring-VCO sub-sampling loop that the fixed-gain approach cannot match.

4.5.9 Limitations of the Behavioural Study

All simulations here are based on behavioral VerilogA models, no transistor level schematics. There is no device mismatch and finite VCDL bandwidth, charge injection of the sampling switch, supply coupling and temperature drift by construction. In a practical circuit each of these would introduce some jitter, and alter the effective benefits a little. What the simulations show is that the discrete-time stability argument is valid and the adaptive schedule provides the promised lock-time improvement when the loop acts as modelled—not that a made-up chip would reach the same numbers in nanoseconds.

The second gap is the noise scope. Jitter numbers are based on an integrated noise-bandwidth model driven by the noise floors of the oscillator and detector with no reference spurs, sub-sampling spur folding, and flicker-noise up-conversion. These omis-

sions are not critical to the main comparison – both loops have the same noise sources and the sole difference is in the pole position – but absolute jitter values should be interpreted as model projections, not silicon observations. Sensitivity sweeps also modify one parameter at a time. A thorough PVT corner analysis is part of a transistor-level follow-up.

Neither limitation undermines the methodology. Computing α , locating the deadbeat gain, scheduling between it and $G_{F,ss}$, and verifying $|z_{\text{pole}}| < 1$ along the trajectory are steps that are independent of modelling level. Behavioural results give a transistor-level implementation a concrete target and a quantified expectation against which silicon can be checked.

4.6 DISCUSSION

4.6.1 Interpretation of Lock-Time Results

Three reference cycles are saved at the $\pi/4$ design point (six down to four), giving a $1.70\times$ improvement. This falls short of the theoretical one-shot limit (which would approach $5\times$ if a single cycle sufficed) primarily because the schedule evaluates to $G_F(\pi/4) = 5.218 \text{ rad/V}$ —within 0.2% of the deadbeat gain $G_{F,\text{crit}} = 5.226 \text{ rad/V}$ but not exactly equal. The residual pole at $z = +0.037$ rather than $z = 0$ means a couple of additional cycles are still needed to clear the lock threshold.

At $\pi/2$, the schedule evaluates to 5.375 rad/V , placing the pole at $z \approx +0.009$ —closer to zero than at $\pi/4$. One cycle is sufficient for 96% of the error to vanish; the second cycle brings the remainder below 5 mV. This self-calibrating behaviour—larger steps activate the schedule more aggressively, driving the pole nearer to zero—explains why the improvement ratio increases from $1.70\times$ at $\pi/4$ to $3.66\times$ at $\pi/2$.

4.6.2 Lyapunov Condition Under Spectre Conditions

The maximum pole magnitude observed during any adaptive transient is $|z_{\text{pole}}| = 0.4017$ (cycle 4 of the $\pi/4$ step). The worst-case across all configurations is 0.4451 (at lock in both baseline and adaptive). Both are well inside the unit circle. The Lyapunov condition $|z_{\text{pole}}| < 1$ is satisfied at every cycle in every simulation, confirming Theorem 1 of Section 3.5 under realistic Spectre integration conditions.

4.6.3 VCDL Physical Feasibility

At the $\pi/4$ operating point, the VCDL gain swings from $G_{F,ss} = 3.0$ to $G_F = 5.218 \text{ rad/V}$ over the two-to-three cycle transient. For a control voltage amplitude of $\sim 40 \text{ mV}$ (typical SSPD output at cycle 1), the additional delay shift is $\Delta G_F \cdot v_{\text{ctrl}} / (2\pi f_{\text{VCO}}) \approx 2.218 \times 0.04 / (2\pi \times 2.4 \times 10^9) \approx 5.9 \text{ ps}$. Over the two-cycle transient, the cumulative delay swing is approximately 12 ps—2.9% of the VCO period at 2.4 GHz. This is well within the tuning range of a standard ring-VCO-based delay cell, suggesting that the behavioural model parameters correspond to a physically realisable design.

4.6.4 Summary of Simulation Results

Table 4.8 provides a complete summary of the Spectre results alongside the analytically derived design parameters.

Table 4.8: Complete performance summary: design parameters and Spectre-measured results.

Parameter / Metric	Baseline	Proposed
<i>Design parameters (analytically derived)</i>		
VCDL gain G_F	3.0 rad/V (fixed)	3.0–5.509 rad/V (adaptive)
DT pole at $\pi/4$ op. point	−0.4451 (fixed)	+0.037 (transient)
DT pole at lock	−0.4451	−0.4451
Stability margin	1.86×	1.86×
NBW at lock	130.2 MHz	130.2 MHz
Jitter (analytical)	J_0	J_0 (no SS penalty)
<i>Spectre simulation results (from <code>spectre.out</code>)</i>		
Lock time, $\pi/4$ step	48.785 ns	28.785 ns
Lock time, $\pi/2$ step	27.525 ns	7.525 ns
Improvement, $\pi/4$	—	1.70×
Improvement, $\pi/2$	—	3.66×
Lock cycles, $\pi/4$	6 cycles	4 cycles
Lock cycles, $\pi/2$	≈ 3 cycles	< 1 cycle (near one-shot)

The results confirm all four research hypotheses of Chapter 1. H1 (lock-time reduction): confirmed at $1.70\times$ ($\pi/4$) and $3.66\times$ ($\pi/2$). H2 (near one-shot settling): confirmed at $\pi/2$ —lock in 0.75 reference cycles. H3 (Lyapunov stability): confirmed by Table 4.7; no pole magnitude exceeds 0.4451. $\text{NBW}_{\text{ss}} = 130.2$ MHz at lock. Chapter 5 draws final conclusions and proposes direc

CHAPTER 5

CONCLUSION AND RECOMMENDATIONS

5.1 CONCLUSIONS

Replacing the base paper’s fixed VCDL gain with a smoothly adaptive schedule achieves faster lock without sacrificing steady-state jitter. Six contributions came out of the work:

C1 — DT pole derivation and CT-DT discrepancy. The discrete-time pole $z_{\text{pole}} = (1 + \beta - \alpha)/(1 + \beta)$ was derived from first principles for the VCDL-aided SSPLL. At the base-paper operating point ($G_F = 3.0 \text{ rad/V}$) it evaluates to $z = -0.4451$ —negative-real, and therefore settling that oscillates from one reference cycle to the next. The continuous-time phase-margin analysis of the base paper predicts a smooth 30° -margin decay; the discrete-time model instead predicts alternating-sign error samples, and the Spectre baseline waveform confirms this precisely: $+86.75, -98.40, +43.75, -19.45, +8.65, -3.85 \text{ mV}$ on successive reference edges. Both models share the same algebraic pole—the Euler-forward map is exact for a first-order system—but only the discrete-time test $|z_{\text{pole}}| < 1$ correctly locates the stability boundary.

C2 — Stability margin formula. $G_{F,\text{ss}}/G_{F,\text{min}} = 3.0/1.613 = 1.86\times$ quantifies how far the base-paper design sits above the stability edge—a metric not previously reported. It bounds the allowable PVT drift in K_{SSPD} and K_{VCO} .

C3 — Corrected DT lock-time formula. For a fixed pole, $n = \log(p/100)/\log|z_{\text{pole}}|$ reference cycles are needed to reach $p\%$ settling. At $z_{\text{pole}} = -0.4451$ and 2% settling, $n \approx 4.8$ cycles—about 48 ns , matching Spectre’s 48.785 ns . CT analysis undershoots because it uses $1/\omega_{\text{UG}}$ rather than the closed-loop DT settling envelope.

C4 — Adaptive gain scheduling. Applying $G_F(|\phi_e|) = G_{F,\text{ss}} + (G_{F,\text{max}} - G_{F,\text{ss}})|\phi_e|/(|\phi_e| + \phi_0)$ with $G_{F,\text{max}} = 5.509 \text{ rad/V}$, $G_{F,\text{ss}} = 3.0 \text{ rad/V}$, and $\phi_0 = 0.10 \text{ rad}$ cuts lock time from 48.785 ns to 28.785 ns on a $\pi/4$ step—a $1.70\times$ improvement Spectre confirms. On a $\pi/2 \text{ rad}$ step the gain reaches $3.66\times$ (27.525 ns down to 7.525 ns), close to one-shot settling at 0.75 reference cycles. Steady-state jitter is preserved because the loop always returns to $G_{F,\text{ss}} = 3.0 \text{ rad/V}$ at lock.

C5 — Critical gain and one-shot design equation. Setting $G_{F,\text{crit}} = (\alpha - 1)/K_{\text{SSPD}} = 5.226 \text{ rad/V}$ places the pole at the origin, giving one-step settling. The optimal maximum gain $G_{F,\text{max}}^{\text{opt}} = 5.509 \text{ rad/V}$ follows from requiring the schedule to hit $G_{F,\text{crit}}$ at the design step, yielding a closed-form design equation that carries over to any SSPLL parameter set.

C6 — Lyapunov stability proof. $V(\phi_e) = \phi_e^2$ decreases at every step under the adaptive law, since $G_F(|\phi_e|) \geq G_{F,\text{ss}} = 3.0 > G_{F,\text{min}} = 1.613$ for all error magnitudes. The Spectre pole trajectory in Table 4.7 backs this up: across every adaptive run the largest pole magnitude observed is $|z_{\text{pole}}| = 0.4017$, comfortably inside the unit circle.

Taken together, these results show that adaptive VCDL gain scheduling is a work-

able, provably stable, and genuinely effective way to cut lock time in sub-sampling PLLs—without adding hardware, without a second feedback path, and without any penalty to settled noise performance.

5.2 LIMITATIONS

A few limits should frame how far these conclusions are stretched.

Every simulation in this project runs on behavioural VerilogA models rather than transistor-level schematics. VerilogA SSPD, VCDL, and VCO models assume linear or near-linear behaviour across the operating range, whereas real CMOS does not: the SSPD sampling switch carries charge injection and clock feedthrough, the VCDL maintains its delay-tuning linearity only out to about $\pm 30\%$ of nominal delay, and the ring VCO's K_{VCO} drifts with control voltage. None of these second-order effects is captured in the current model, and any of them could trim the improvement factor in a fabricated circuit.

The VCO phase-noise model uses $S_f = 6320 \text{ Hz}^2/\text{Hz}$, i.e., $\mathcal{L}(1 \text{ MHz}) = -85 \text{ dBc/Hz}$, a fair value for a 2.4 GHz ring VCO in 45 nm CMOS. Cadence Spectre transient-noise simulation (injected as ctrl-line noise $S_v = 4.779 \times 10^{-13} \text{ V}^2/\text{Hz}$) put the VCO-noise jitter at 2.596 ps (baseline and adaptive alike) and 2.292 ps (fixed $G_{F,\text{crit}}$), matching the discrete-time prediction to within 1.5%. SSPD sampling noise, integrated through the exact DT noise bandwidth (3.12), comes to 1.229 ps (baseline/adaptive) and 0.762 ps ($G_{F,\text{crit}}$)—a 38.0% reduction. The two noise sources respond to gain scheduling quite differently: the 38.0% SSPD reduction versus 11.7% for the VCO, both evaluated at the same deadbeat operating point.

Finally, the chosen scheduling law is one particular hyperbolic function. Other shapes—piecewise-linear, exponential, or a simple two-state gear-shift—might deliver comparable lock-time improvements with simpler hardware. A systematic comparison of schedule shapes is left for future work.

5.3 RECOMMENDATIONS FOR FUTURE RESEARCH

R1 — Transistor-level implementation in Cadence Virtuoso. The natural next step is to replace the VerilogA models with transistor-level schematics in the same Cadence Virtuoso environment, on GPDK045 or a production 28 nm PDK. The SSPD sampling switch, the VCDL delay cells, and the adaptive scheduler's arithmetic block are the subcircuits that most need faithful implementation. SpectreRF periodic-steady-state simulation would then extract jitter and phase noise directly from silicon-accurate models.

R2 — Spectre jitter extraction (completed). The Spectre transient-noise run on the VerilogA testbench is already complete, and it confirms that the adaptive scheduler adds no steady-state jitter penalty. VCO-noise jitter (2.596 ps baseline, 2.560 ps adaptive at lock) and SSPD-noise jitter (1.229 ps for both) are identical within simulation noise, validating the analytical model.

R3 — Silicon measurement. Fabricating a test chip in 28 nm or 16 nm FinFET CMOS would give the definitive benchmark against the base paper's 0.934 ps jitter, and would

reveal whether the $1.70\times$ – $3.66\times$ lock-time improvement survives across process, voltage, and temperature.

R4 — Alternative scheduling strategies. The hyperbolic schedule was chosen for analytical convenience. A two-state gear-shift— G_F snapping between $G_{F,\max} = 5.509$ and $G_{F,\text{ss}} = 3.0$ on a lock-detector signal—would approximate the continuous schedule with simpler hardware, and a digital version using a counter and a gain look-up table is a straightforward path to integration.

R5 — Extension to LC-VCO SSPLLs. An LC-VCO SSPLL has a lower K_{VCO} , hence a smaller α and an easier stability constraint. Applying adaptive gain scheduling there could pay off even more, since the baseline already sits closer to the one-shot condition.

R6 — Fractional- N and millimetre-wave extension. Carrying the adaptive VCDL idea into fractional- N SSPLLs, where a delta-sigma modulator shapes the division sequence, means modelling how the adaptive gain schedule interacts with the shaped quantisation-noise spectrum. That is an open and practically pressing question for 5G and beyond-5G synthesisers at 28 GHz.

5.4 SUMMARY

In short, this project showed that a simple, smooth, error-dependent schedule for the VCDL gain in a sub-sampling PLL trims lock time by $1.70\times$ at the design step ($\pi/4$ rad) and by $3.66\times$ at twice that step ($\pi/2$ rad), with no cost to steady-state noise. The discrete-time pole analysis, the adaptive scheduling law with its analytically derived parameters, the Lyapunov stability proof, and the Cadence Virtuoso/Spectre transient validation come together as a complete design framework for adaptive gain-scheduled SSPLLs—self-contained, process-independent, and ready to carry over to a transistor-level implementation in any CMOS node.

BIBLIOGRAPHY

- [1] R. E. Best. *Phase-Locked Loops: Design, Simulation, and Applications*. McGraw-Hill, New York, 6th edition, 2007.
- [2] W. F. Egan. *Frequency Synthesis by Phase Lock*. Wiley, New York, 2nd edition, 2000.
- [3] X. Gao, E. A. M. Klumperink, M. Bohsali, and B. Nauta. A low noise sub-sampling PLL in which divider noise is eliminated and PD/CP noise is not multiplied by N^2 . *IEEE J. Solid-State Circuits*, 44(12):3253–3263, December 2009.
- [4] X. Gao, E. A. M. Klumperink, G. Socci, M. Bohsali, and B. Nauta. Spur reduction techniques for phase-locked loops exploiting a sub-sampling phase detector. *IEEE J. Solid-State Circuits*, 45(9):1809–1821, September 2010.
- [5] F. M. Gardner. Charge-pump phase-lock loops. *IEEE Trans. Commun.*, 28(11):1849–1858, November 1980.
- [6] F. M. Gardner. *Phaselock Techniques*. Wiley, Hoboken, NJ, 3rd edition, 2005.
- [7] M. Ghahramani, M. H. Zarifi, and M. Sharif Bakhtiar. A bandwidth extension technique for improving jitter in ring-VCO-based sub-sampling PLLs. *IEEE Trans. Circuits Syst. II, Exp. Briefs*, 72(1):44–48, January 2025.
- [8] P. K. Hanumolu, M. Brownlee, K. Mayaram, and U.-K. Moon. Analysis of charge-pump phase-locked loops. *IEEE Trans. Circuits Syst. I, Reg. Papers*, 51(9):1665–1674, September 2004.
- [9] C. D. Hedayat, A. Roth, T. Paul, and H.-M. Rein. Modeling and characterization of the 3rd order charge-pump PLL: a fully event-driven approach. *Analog Integr. Circuits Signal Process.*, 19(1):25–45, 1999.
- [10] A. Homayoun and B. Razavi. Analysis of phase noise in phase/frequency detectors. *IEEE Trans. Circuits Syst. I, Reg. Papers*, 60(3):529–539, March 2013.
- [11] S. Kumar and Y. K. Singh. A low-jitter and low-phase-noise switched-loop-filter PLL using fast phase-error correction and dual-edge phase comparison. *Integration*, 94:102108, 2024.
- [12] J. Lee and B. Kim. A low-noise fast-lock phase-locked loop with adaptive bandwidth control. *IEEE J. Solid-State Circuits*, 35(8):1137–1145, August 2000.
- [13] T. H. Lee and J. F. Bulzacchelli. A 155-MHz clock recovery delay- and phase-locked loop. *IEEE J. Solid-State Circuits*, 27(12):1736–1746, December 1992.

- [14] D. B. Leeson. A simple model of feedback oscillator noise spectrum. *Proc. IEEE*, 54(2):329–330, February 1966.
- [15] G. A. Leonov, N. V. Kuznetsov, M. V. Yuldashev, and R. V. Yuldashev. Hold-in, pull-in, and lock-in ranges of PLL circuits: Rigorous mathematical definitions and limitations of classical theory. *IEEE Trans. Circuits Syst. I, Reg. Papers*, 62(10):2454–2464, October 2015.
- [16] J. G. Maneatis. Low-jitter process-independent DLL and PLL based on self-biased techniques. *IEEE J. Solid-State Circuits*, 31(11):1723–1732, November 1996.
- [17] M. Mansuri and C.-K. K. Yang. A low-power adaptive bandwidth PLL and clock buffer with supply-noise compensation. *IEEE J. Solid-State Circuits*, 38(11):1804–1812, November 2003.
- [18] S. S. Nagam and P. R. Kinget. A low-jitter ring-oscillator phase-locked loop using feedforward noise cancellation with a sub-sampling phase detector. *IEEE J. Solid-State Circuits*, 53(3):703–714, March 2018.
- [19] B. Razavi. A study of phase noise in CMOS oscillators. *IEEE J. Solid-State Circuits*, 31(3):331–343, March 1996.
- [20] B. Razavi. *Design of Analog CMOS Integrated Circuits*. McGraw-Hill, New York, 2001.
- [21] K. Shu and E. Sánchez-Sinencio. *CMOS PLL Synthesizers: Analysis and Design*. Springer, New York, 2005.
- [22] A. J. Viterbi. *Principles of Coherent Communication*. McGraw-Hill, New York, 1966.

APPENDIX A

VERILOGA BEHAVIOURAL MODELS

This appendix lists the five VerilogA modules used in the Cadence Virtuoso/Spectre testbench described in Chapter 3. Each module is a pure behavioural model; no transistor-level circuitry is included. Parameters are passed directly from the Spectre netlist and may be overridden per-instance.

A.1 SUB-SAMPLING PHASE DETECTOR

The SSPD samples the VCO waveform at each reference clock edge through a track-and-hold model. The output voltage $\phi_{e,v}$ is proportional to the instantaneous phase error. The sinusoidal VCO waveform is sampled at the zero-crossing region, where $V_{VCO} \approx V_{mid} + V_{amp} \cdot \sin(\phi_e)$. For small errors the output is linear in ϕ_e with gain K_{SSPD} ; for large errors the sinusoidal characteristic introduces mild compression.

```

1  'include "disciplines.vams"
2  module sspd_subsampling(ref_clk, vco_in, vout, vss);
3      input  ref_clk, vco_in, vss;
4      output vout;
5      electrical ref_clk, vco_in, vout, vss;
6
7      parameter real KSSPD      = 0.5;    // V/rad
8      parameter real Vamp_vco   = 0.5;    // V
9      parameter real Vmid_vco   = 0.5;    // V
10     parameter real Vth_clk    = 0.6;    // V clock threshold
11     parameter real Thold      = 10e-12; // s hold window
12
13     real phi_e_sampled;
14
15     analog begin
16         @(cross(V(ref_clk, vss) - Vth_clk, +1)) begin
17             // Sample VCO voltage at rising reference edge
18             phi_e_sampled = (V(vco_in, vss) - Vmid_vco) / Vamp_vco;
19             // Convert to phase error voltage
20             phi_e_sampled = KSSPD * $asin(phi_e_sampled);
21         end
22         V(vout, vss) <+ phi_e_sampled;
23     end
24 endmodule

```

Listing A.1: SSPD sub-sampling phase detector.

A.2 RING VOLTAGE-CONTROLLED OSCILLATOR

The ring VCO is a first-order behavioural model that generates a sinusoidal output at frequency $f_0 + K_{VCO} \cdot V_{ctrl}$. The parameter ϕ_{step} is set to zero in the testbench to suppress VCO-side phase injection.

```

1  'include "disciplines.vams"
2  module vco_ring(ctrl, out, vss);
3      input  ctrl, vss;
4      output out;
5      electrical ctrl, out, vss;
6
7      parameter real f0          = 2.4e9; // Hz
8      parameter real Kvco       = 115e6; // Hz/V
9      parameter real Vamp       = 0.5;
10     parameter real Vmid       = 0.5;
11     parameter real phi_init   = 'M_PI;
12     parameter real Vctrl_nom  = 0.0; // V nominal control
13     parameter real phi_step   = 0.0; // phase injection (set 0)
14
15     real phase, freq;
16
17     analog begin
18         @(initial_step) phase = phi_init;
19         freq = f0 + Kvco * (V(ctrl, vss) - Vctrl_nom);
20         phase = idt(2 * 'M_PI * freq, phi_init);
21         V(out, vss) <+ Vmid + Vamp * sin(phase);
22     end
23 endmodule

```

Listing A.2: Ring VCO behavioural model.

A.3 ADAPTIVE VCDL

The VCDL translates the phase-error voltage and the adaptive gain voltage into a delay-corrected control signal. The gain factor G_F modulates the effective delay shift applied to the sampling instant.

```

1  'include "disciplines.vams"
2  module vcdl_adaptive(vin, v_ctrl, v_gf, vout, vss);
3      input  vin, v_ctrl, v_gf, vss;
4      output vout;
5      electrical vin, v_ctrl, v_gf, vout, vss;
6
7      parameter real KSSPD = 0.5; // V/rad, must match SSPD
8
9      real gf_val, correction;
10
11     analog begin
12         gf_val = V(v_gf, vss); // gain factor from
13         // scheduler
14         correction = gf_val * V(v_ctrl, vss); // delay correction
15         // signal

```

```

14     V(vout, vss) <+ V(vin, vss) - KSSPD * correction;
15     end
16 endmodule

```

Listing A.3: Voltage-controlled delay line (adaptive).

A.4 ADAPTIVE G_F SCHEDULER

The scheduler implements the hyperbolic scheduling law (3.7). It takes the phase-error voltage $\phi_{e,v}$ and outputs the corresponding gain factor voltage v_{G_F} every simulation time step.

```

1  'include "disciplines.vams"
2  module adaptive_gf_scheduler(v_phi_e, v_gf_out, vss);
3      input  v_phi_e, vss;
4      output v_gf_out;
5      electrical v_phi_e, v_gf_out, vss;
6
7      parameter real GF_ss = 3.0;    // rad/V steady-state gain
8      parameter real GF_max = 5.509; // rad/V maximum gain
9      parameter real phi_0 = 0.10;   // rad shaping parameter
10     parameter real KSSPD = 0.5;    // V/rad (to convert V -> rad)
11
12     real phi_e_abs, gf_val;
13
14     analog begin
15         phi_e_abs = abs(V(v_phi_e, vss)) / KSSPD; // convert V to
16         // rad
17         gf_val = GF_ss + (GF_max - GF_ss) * phi_e_abs / (phi_e_abs
18         + phi_0);
19         V(v_gf_out, vss) <+ gf_val;
20     end
21 endmodule

```

Listing A.4: Adaptive G_F scheduler.

A.5 LOCK-TIME DETECTOR

The lock-time detector monitors $|\phi_{e,v}(t)|$ continuously and records the first threshold crossing time. The threshold is computed as $\text{tol_pct} \times \phi_{\text{step}} \times K_{\text{SSPD}}$. The output voltage encodes T_{lock} in nanoseconds (1 V = 1 ns). When lock is declared, the module prints to the Spectre log.

```

1  'include "disciplines.vams"
2  module lock_time_detector(phi_e_v, v_tlock, vss);
3      input  phi_e_v, vss;
4      output v_tlock;
5      electrical phi_e_v, v_tlock, vss;
6
7      parameter real KSSPD = 0.5;
8      parameter real phi_step = 0.7854; // rad (pi/4)

```

```

9   parameter real t_step    = 100e-9;    // s step time
10  parameter real tol_pct   = 0.02;      // 2% criterion
11
12  real lock_thr, tlock_ns;
13  integer locked;
14
15  analog begin
16      @(initial_step) begin
17          lock_thr = tol_pct * phi_step * KSSPD; // ~5 mV
18          locked   = 0;
19          tlock_ns = 0.0;
20      end
21      if (!locked && $abstime > t_step) begin
22          if (abs(V(phi_e_v, vss)) < lock_thr) begin
23              locked = 1;
24              tlock_ns = ($abstime - t_step) * 1e9;
25              $strobe("T_lock_ = %0.3f ns", tlock_ns);
26          end
27      end
28      V(v_tlock, vss) <+ tlock_ns * 1e-9; // 1V = 1ns encoding
29  end
30  endmodule

```

Listing A.5: Lock-time detector VerilogA module.

A.6 SPECTRE NETLIST (KEY INSTANCE LINES)

The corrected Spectre netlist instance lines (after fix_all.sh is applied) are listed below for reference. These constitute the ground truth for all simulations reported in Chapter 4.

```

1  // Phase step source
2  I2 (ref_clk 0) phase_step_src f_REF=1e+08 phi_step=0.7854
   t_step=1e-07
3
4  // Sub-sampling phase detector
5  I3 (ref_clk vco_out phi_e_v 0) sspd_subsampling KSSPD=0.5 \
6  Vamp_vco=0.5 Vmid_vco=0.5 Vth_clk=0.6 Thold=1e-11
7
8  // Adaptive GF scheduler
9  I0 (phi_e_v gf_out 0) adaptive_gf_scheduler GF_ss=3 GF_max
   =5.509
10
11 // Adaptive VCDL
12 I5 (phi_e_v phi_e_v gf_out vcdl_out 0) vcdl_adaptive
13
14 // Ring VCO
15 I6 (vcdl_out vco_out 0) vco_ring f0=2.4e+09 Kvco=1.15e+08 \
16 Vamp=0.5 Vmid=0.5 phi_init=3.14159 Vctrl_nom=0 phi_step=0
17
18 // Lock-time detector
19 I13 (phi_e_v v_tlock 0) lock_time_detector KSSPD=0.5 \
20 phi_step=0.7854 t_step=1e-07 tol_pct=0.02

```

```
21 |  
22 | // Initial conditions  
23 | ic phi_e_v=0 vcdl_out=0  
24 |  
25 | // Transient analysis  
26 | tran tran stop=500n errpreset=moderate maxstep=20p
```

Listing A.6: Key Spectre netlist instance lines (post-patch).