POLITECNICO DI TORINO

Corso di Laurea
Ingegneria Elettronica (Electronic Engineering)

Tesi di Laurea

Cryogenic CMOS Frequency Generation: LC-VCO Design and Integrated Control-Chip Implementation for Spin Qubits

Relatori
prof. Gianluca Piccinini
prof.ssa Mariagrazia Graziano
prof. Marco Vacca
prof. Fabrizio Riente

Candidato
Alberto Nesladek

Anno Accademico 2025-2026

Summary

Quantum computing is widely regarded as a promising paradigm for addressing classes of problems that are beyond the practical reach of classical computation. Among the different hardware platforms currently under investigation, semiconductor spin qubits stand out because of their compatibility with advanced microelectronic fabrication and their potential for large-scale integration. However, the realization of scalable spin-qubit processors does not depend only on qubit fabrication and coherence properties. It also requires the development of a classical electronic interface capable of generating, storing, and coordinating the bias, control, and readout signals needed for qubit operation within a cryogenic environment.

In present experimental setups, these functions are often performed by room-temperature instrumentation connected to the quantum device through a large number of cables. While this approach is acceptable for small-scale demonstrations, it becomes increasingly problematic as the number of qubits grows, due to wiring complexity, thermal load, footprint, and latency. For this reason, cryogenic CMOS electronics are increasingly considered a key enabling technology for scalable quantum computing, since they offer the possibility of relocating part of the control chain closer to the qubit plane while preserving a reasonable compromise between integration density, power dissipation, and system complexity.

Within this framework, the work presented in this thesis addresses the problem from a combined architectural, circuit-level, and modeling perspective. A first contribution consists in the conception of a reference cryogenic control-chip architecture for semiconductor spin qubits. The proposed architecture includes the main functional subsystems required for qubit operation, namely bias generation, local memory, digital sequencing, microwave generation, and readout support. In this sense, the thesis does not treat the oscillator as an isolated RF block, but places it within the broader context of a mixed-signal cryogenic controller intended for future scalable implementations.

Within the proposed architecture, particular attention is devoted to the RF front-end, since coherent spin control requires microwave bursts with high spectral purity. For this reason, the main circuit-level focus of the thesis is the design of a 10 GHz LC-VCO in TSMC 65 nm CMOS, intended as the first implemented RF building block of the proposed cryogenic control chip. The oscillator is designed through a simulation-driven methodology starting from system-level targets on operating frequency, tuning range, phase noise, output swing, and power consumption. The design flow progressively addresses resonant-tank dimensioning, active-core sizing, startup verification, bias-network implementation, continuous tuning with inversion-mode MOS capacitors, and discrete tuning with a switched-capacitor bank.

The oscillator is then physically implemented and verified at post-layout level. A complete layout is developed, validated through DRC and LVS, and simulated after parasitic extraction. The extracted results confirm stable oscillation, full tuning-range coverage,

and a realistic trade-off between power dissipation, output swing, and spectral performance. In order to connect circuit-level results more directly to the target quantum application, the simulated phase-noise spectra are also processed through the filter-function formalism, allowing the estimation of the average fidelity of a primitive X_π gate. This makes it possible to interpret the oscillator not only in terms of conventional RF metrics, but also in terms of its impact on qubit-control quality.

A second major contribution concerns the treatment of cryogenic transistor behavior. Since no qualified cryogenic compact model was available in the adopted industrial PDK, a dedicated BSIM4 re-centering procedure was developed in order to support the low-temperature interpretation of the oscillator results. The adopted approach preserves the original BSIM4 formulation and the standard Cadence/Spectre flow, while modifying a limited set of model parameters so as to reproduce the main cryogenic trends reported in the literature. The resulting design-oriented cryogenic model is validated against digitized transistor characteristics, including transfer curves, transconductance, output characteristics, and output conductance. It is then further assessed in an oscillator-like resonant testbench, showing physically meaningful startup and steady-state behavior under cryogenic assumptions.

Overall, the thesis provides a coherent contribution at multiple and complementary levels. At the system level, it proposes the architecture of a cryogenic control chip for semiconductor spin qubits. At the circuit level, it develops and validates a 10 GHz LC-VCO as the first concrete RF block within that architecture. At the modeling level, it introduces a practical BSIM4 re-centering methodology to interpret the active-device behavior under cryogenic conditions. Taken together, these results constitute a first step toward the realization of more integrated cryogenic control electronics for scalable silicon-quantum-computing platforms.

Acknowledgements

I would first like to express my sincere gratitude to Professor Gianluca Piccinini for his supervision throughout this work. I am especially grateful to him for having introduced me, through the course of Integrated System Technology, to the technological world behind electronics, and for later giving me the opportunity to approach and become fascinated by the field of quantum computing. His guidance played a fundamental role in shaping the direction of this thesis.

I would also like to sincerely thank Professor Mariagrazia Graziano for her constant support and encouragement. I am particularly grateful to her for introducing me to Claudio Capobianchi, with whom I had the opportunity to work closely and share many valuable discussions. I also wish to thank him for the fruitful collaboration, the exchange of ideas, and the joint work on the architectural aspects of this thesis.

I would also like to thank Professor Marco Vacca for his supervision and for the valuable advice he provided during the development of this work.

My sincere thanks also go to Professor Riente for his valuable contribution and the insightful discussions on quantum-computing topics.

Finally, I would like to thank Dr. Fabrizio Mo for having introduced me to this thesis work and for giving me the opportunity to engage with such an interesting and stimulating research topic.

Contents

List of Tables

List of Figures

List of Abbreviations

ADC	Analog-to-Digital Converter
ADE	Analog Design Environment
AM	Amplitude Modulation
BSIM4	Berkeley Short-channel IGFET Model 4
CMOS	Complementary Metal-Oxide-Semiconductor
DAC	Digital-to-Analog Converter
DC	Direct Current
DIBL	Drain-Induced Barrier Lowering
DRC	Design Rule Check
EDA	Electronic Design Automation
EDSR	Electric Dipole Spin Resonance
ESR	Electron Spin Resonance
FSM	Finite-State Machine
I/O	Input/Output
I/V	Current-to-Voltage
LDD	Lightly Doped Drain
LNA	Low-Noise Amplifier
LO	Local Oscillator
LUT	Look-Up Table
LVS	Layout Versus Schematic
MOS	Metal-Oxide-Semiconductor
MOSCAP	MOS Capacitor
NMOS	n-type Metal-Oxide-Semiconductor
PDK	Process Design Kit
PEX	Parasitic Extraction
PLL	Phase-Locked Loop
PM	Phase Modulation
Pnoise	Phase-Noise Analysis
PSS	Periodic Steady-State
PSD	Power Spectral Density
RF	Radio Frequency
RMS	Root Mean Square
RWA	Rotating-Wave Approximation
SET	Single-Electron Transistor
S/H	Sample-and-Hold
SP	Scattering-Parameter / Small-Signal Analysis
SSB	Single-Sideband
VCO	Voltage-Controlled Oscillator

Part I

Theoretical Background

Chapter 1

Introduction

1.1 Quantum computing

Quantum computers provide a means to solve specific classes of problems that are believed to be intractable for classical computers within any reasonable amount of time. In practice, they are expected to operate as specialized coprocessors: a classical microprocessor provides the instructions for the computation, while the quantum processor executes carefully designed sequences of operations by exploiting the quantum properties of its fundamental unit of information, the qubit, most notably superposition and entanglement.

The first and arguably most intuitive application of quantum computing is the simulation of quantum systems. Even relatively small many-body models can become computationally prohibitive for classical approaches because the Hilbert space grows exponentially with system size; for instance, already at the scale of ~ 100 lattice sites these simulations can become challenging [Svore and Troyer, 2016]. Efficient quantum simulation would have a major impact on materials science and chemistry, potentially accelerating the discovery of novel compounds, including room-temperature superconductors [Gibney, 2014]. Beyond simulation, quantum algorithms may also offer speedups for more traditional tasks such as cryptography, solving linear systems, and searching large datasets [Svore and Troyer, 2016].

However, scalability remains a central challenge. Any quantum system is inevitably coupled to its environment, resulting in decoherence and control/readout errors that progressively degrade the stored quantum information. As a consequence, practically useful machines require quantum error correction, in which logical qubits are encoded into many noisy physical qubits. Basic quantum-error-correction schemes may require from 10^2 to 10^4 physical qubits per logical qubit [Versluis, 2020]. Thus, even a quantum algorithm requiring only 100 logical qubits may demand 10^4 to 10^6 physical qubits. This strong overhead motivates the development of hardware platforms with extremely high qubit-integration density.

1.1.1 Qubits as two-level quantum systems

The elementary unit of quantum information is the qubit, i.e., a quantum two-level system. Mathematically, its state is represented by a vector $|\psi\rangle$ belonging to a two-dimensional Hilbert space. A convenient basis for this Hilbert space is provided by the two possible outcomes of a projective measurement performed on the qubit.

An electron in a static magnetic field oriented, for example, along the z axis, constitutes a two-level quantum system. In this case, the spin can be measured with respect to the field direction, and the two possible outcomes are the spin-up state $|\uparrow\rangle$ and the spin-down state $|\downarrow\rangle$. Before measurement, however, the qubit can exist in a coherent superposition of the two basis states:

$$|\psi\rangle = \alpha\,|\uparrow\rangle + \beta\,|\downarrow\rangle\,, \qquad \alpha,\beta \in \mathbb{C}, \quad |\alpha|^2 + |\beta|^2 = 1. \tag{1.1}$$

Born's rule states that the probabilities of obtaining the outcomes $|\uparrow\rangle$ and $|\downarrow\rangle$ are

$$P(\uparrow) = |\,\langle\uparrow|\psi\rangle\,|^2 = |\alpha|^2, \qquad P(\downarrow) = |\,\langle\downarrow|\psi\rangle\,|^2 = |\beta|^2.$$

Since these are the only possible outcomes in this basis, their sum must satisfy

$$P(\uparrow) + P(\downarrow) = 1 \quad \Rightarrow \quad |\alpha|^2 + |\beta|^2 = 1.$$

To describe measurements and control operations on a qubit, it is convenient to introduce linear operators acting on the same two-dimensional Hilbert space. Once the orthonormal basis $\{|\uparrow\rangle\,, |\downarrow\rangle\}$ is fixed, any linear operator can be represented by a 2×2 complex matrix. In particular, observables correspond to Hermitian operators. A standard and extremely useful operator basis is given by the identity and the Pauli matrices, which provide a compact description of single-qubit Hamiltonians, measurements, and rotations.

1.1.2 Pauli matrices, expectation values, and Bloch-sphere representation

The Pauli matrices $\sigma_x, \sigma_y, \sigma_z$ are defined, in the basis $\{|\uparrow\rangle\,, |\downarrow\rangle\} \equiv \{|\uparrow_z\rangle\,, |\downarrow_z\rangle\}$, as

$$\sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \qquad \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \qquad \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \tag{1.2}$$

They are Hermitian operators,

$$\sigma_i^\dagger = \sigma_i \qquad (i = x, y, z),$$

and therefore represent observables. Moreover, $\sigma_i^2 = I$, which implies that their eigenvalues are ± 1. Denoting by $|\uparrow_i\rangle$ and $|\downarrow_i\rangle$ the eigenstates of σ_i, one has

$$\sigma_i\,|\uparrow_i\rangle = +\,|\uparrow_i\rangle\,, \qquad \sigma_i\,|\downarrow_i\rangle = -\,|\downarrow_i\rangle \qquad (i = x, y, z).$$

Finally, the set $\{I, \sigma_x, \sigma_y, \sigma_z\}$ forms a basis for the real vector space of 2×2 Hermitian operators: any single-qubit observable $A = A^\dagger$ can be written uniquely as

$$A = a_0 I + a_x \sigma_x + a_y \sigma_y + a_z \sigma_z, \qquad a_0, a_x, a_y, a_z \in \mathbb{R}.$$

The Pauli matrices are particularly useful for describing expectation values and spin projections. For instance, preparing a qubit in the state $|\psi\rangle = \alpha\,|\uparrow_z\rangle + \beta\,|\downarrow_z\rangle$ and measuring the spin along the z axis, the possible outcomes are $+1$ and -1, obtained with probabilities $|\langle\uparrow_z|\psi\rangle|^2$ and $|\langle\downarrow_z|\psi\rangle|^2$, respectively. Repeating the measurement many times, the average outcome is

$$\langle\sigma_z\rangle_\psi = (+1)\,|\langle\uparrow_z|\psi\rangle|^2 + (-1)\,|\langle\downarrow_z|\psi\rangle|^2 = |\alpha|^2 - |\beta|^2.$$

More generally, for any Hermitian observable A, the expectation value in the state $|\psi\rangle$ is defined as

$$\langle A\rangle_\psi = \langle\psi|\,A\,|\psi\rangle\,.$$

The three quantities

$$\langle\sigma_x\rangle_\psi, \qquad \langle\sigma_y\rangle_\psi, \qquad \langle\sigma_z\rangle_\psi$$

represent the mean spin projections of the state $|\psi\rangle$ along the x, y, and z axes, respectively. Collecting them into a vector leads naturally to the Bloch-sphere representation of single-qubit states.

In quantum information, the computational basis is usually denoted by $\{|0\rangle, |1\rangle\}$. For a spin qubit, one may identify

$$|0\rangle \equiv |\uparrow\rangle\,, \qquad |1\rangle \equiv |\downarrow\rangle\,.$$

A convenient geometric representation of a pure single-qubit state is then provided by the *Bloch sphere*. Up to an unobservable global phase, the qubit state of Eq. (1.1) can be written as

$$|\psi\rangle = \cos\left(\frac{\theta}{2}\right)|0\rangle + e^{i\phi}\sin\left(\frac{\theta}{2}\right)|1\rangle\,, \tag{1.3}$$

where $\theta \in [0, \pi]$ and $\phi \in [0{,}2\pi)$ are spherical angles. This parametrization makes explicit that a pure qubit state is determined by two real parameters, which can be interpreted as the coordinates of a point on the surface of a unit sphere.

The mapping becomes clearer by introducing the *Bloch vector* $\vec{r}$, defined through the expectation values of the Pauli operators:

$$\vec{r} = \begin{pmatrix}\langle\sigma_x\rangle\\ \langle\sigma_y\rangle\\ \langle\sigma_z\rangle\end{pmatrix} = \begin{pmatrix}\langle\psi|\,\sigma_x\,|\psi\rangle\\ \langle\psi|\,\sigma_y\,|\psi\rangle\\ \langle\psi|\,\sigma_z\,|\psi\rangle\end{pmatrix}. \tag{1.4}$$

For the state represented as in Eq. (1.3), the Bloch vector is

$$\vec{r} = \begin{pmatrix}\sin\theta\cos\phi\\ \sin\theta\sin\phi\\ \cos\theta\end{pmatrix}, \tag{1.5}$$

which lies on the unit sphere, i.e., $\|\vec{r}\| = 1$. The north pole corresponds to $|0\rangle$ ($\theta = 0$), the south pole to $|1\rangle$ ($\theta = \pi$), while points on the equator ($\theta = \pi/2$) represent superposition states such as $(|0\rangle + e^{i\phi}\,|1\rangle)/\sqrt{2}$. This geometric representation is illustrated in Fig. 1.1.

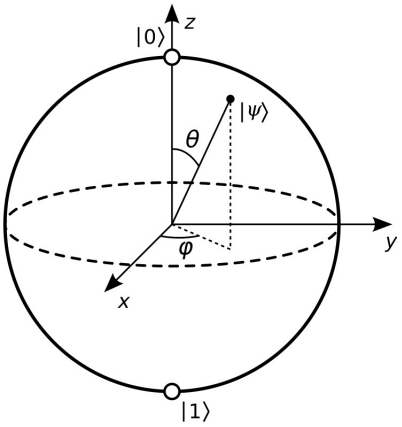

Figure 1.1. Bloch sphere representation of a single-qubit pure state.

1.1.3 Qubit evolution and single-qubit rotations

The time evolution of any closed quantum system prepared in a state $|\psi(t)\rangle$ is governed by the time-dependent Schrödinger equation

$$i\hbar\frac{\partial}{\partial t}\,|\psi(t)\rangle = \hat{H}(t)\,|\psi(t)\rangle\,, \tag{1.6}$$

where $\hat{H}(t)$ is the Hamiltonian operator, representing the total energy of the system.

Denoting by

$$|\psi(t_0)\rangle = \alpha(t_0)\,|0\rangle + \beta(t_0)\,|1\rangle$$

the state at the initial time t_0, its evolution is described by the unitary operator $U(t, t_0)$:

$$|\psi(t)\rangle = \hat{U}(t, t_0)\,|\psi(t_0)\rangle\,, \tag{1.7}$$

with

$$\hat{U}(t, t_0) = \mathcal{T}\exp\left[-\frac{i}{\hbar}\int_{t_0}^{t}\hat{H}(\tau)\,d\tau\right], \tag{1.8}$$

where $\mathcal{T}$ denotes time ordering.

A particularly important case is a time-independent Hamiltonian, $\hat{H}(t) = \hat{H}$, for which the propagator reduces to

$$\hat{U}(t, t_0) = \exp\left[-\frac{i}{\hbar}\hat{H}(t - t_0)\right]. \tag{1.9}$$

If $\{|n\rangle\}$ is an eigenbasis of the Hamiltonian, $\hat{H}\,|n\rangle = E_n\,|n\rangle$, and the initial state is expanded as $|\psi(t_0)\rangle = \sum_n c_n\,|n\rangle$, then the evolution is

$$|\psi(t)\rangle = \sum_n c_n\, e^{-\frac{i}{\hbar}E_n(t-t_0)}\,|n\rangle\,. \tag{1.10}$$

In particular, an energy eigenstate acquires only a phase factor in time, meaning that all measurable probabilities remain unchanged. Such a state is therefore called a *stationary state*.

The Bloch-sphere picture is particularly useful because single-qubit unitary operations correspond to rotations of the Bloch vector. In fact, any single-qubit unitary transformation can be expressed as

$$U(\hat{n}, \vartheta) = \exp\left(-i\frac{\vartheta}{2}\,\hat{n}\cdot\vec{\sigma}\right), \tag{1.11}$$

where $\hat{n}$ is a unit vector defining the rotation axis, ϑ is the rotation angle, and $\vec{\sigma} = (\sigma_x, \sigma_y, \sigma_z)$.

1.1.4 Electron-spin qubits in a magnetic field

A widely used physical realization of a qubit is the spin of a single electron confined in a semiconductor quantum dot, where quantum information is encoded in two Zeeman-split spin states in the presence of an external static magnetic field $\mathbf{B}_0 = B_0\hat{z}$ [Loss and DiVincenzo, 1998]. In the computational basis $\{|0\rangle, |1\rangle\} \equiv \{|\uparrow_z\rangle, |\downarrow_z\rangle\}$, chosen as eigenstates of σ_z,

$$\sigma_z|0\rangle = +|0\rangle, \qquad \sigma_z|1\rangle = -|1\rangle, \tag{1.12}$$

the Zeeman Hamiltonian of a single electron spin can be written as

$$\hat{H}_0 = \frac{\hbar\omega_0}{2}\,\sigma_z, \tag{1.13}$$

where the Larmor angular frequency is

$$\omega_0 = \frac{g\mu_B}{\hbar}\,B_0. \tag{1.14}$$

Here g is the effective electron g-factor in the host material and μ_B is the Bohr magneton. Equation (1.13) implies an energy splitting $\Delta E = \hbar\omega_0$ between the two qubit levels. In practice, the combination of Tesla-scale magnetic fields and $g \approx 2$ leads to qubit resonance frequencies in the multi-GHz to tens-of-GHz range in silicon devices [Veldhorst et al., 2014].

Larmor precession. The Zeeman Hamiltonian also provides a direct geometric interpretation of free evolution on the Bloch sphere. Consider a system prepared in one of the two eigenstates of σ_z, i.e., $|\psi(0)\rangle = |0\rangle$ or $|\psi(0)\rangle = |1\rangle$. The corresponding unitary operator is

$$\hat{U}_0(t) = \exp\left(-\frac{i}{\hbar}\hat{H}_0 t\right) = \exp\left(-i\frac{\omega_0 t}{2}\sigma_z\right) \equiv R_z(\omega_0 t). \tag{1.15}$$

The time evolution of the eigenstates of σ_z is then

$$|\psi(t)\rangle = \hat{U}_0(t)|0\rangle = e^{-i\omega_0 t/2}|0\rangle, \qquad |\psi(t)\rangle = \hat{U}_0(t)|1\rangle = e^{+i\omega_0 t/2}|1\rangle. \tag{1.16}$$

Thus, for a qubit prepared in an eigenstate of σ_z, the evolution consists only of a global phase accumulation, which is not observable.

Consider now a qubit prepared along the x axis,

$$|\psi(0)\rangle = |\uparrow_x\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}}.$$

Its time evolution under $\hat{H}_0$ is

$$\begin{aligned} |\psi(t)\rangle = \hat{U}_0(t)\,|+_x\rangle &= \frac{1}{\sqrt{2}}\left(e^{-i\omega_0 t/2}\,|0\rangle + e^{+i\omega_0 t/2}\,|1\rangle\right) \\ &= e^{-i\omega_0 t/2}\,\frac{1}{\sqrt{2}}\left(|0\rangle + e^{i\omega_0 t}\,|1\rangle\right), \end{aligned} \tag{1.17}$$

showing that the physically relevant effect is the accumulation of a relative phase $e^{i\omega_0 t}$. Consequently, the Bloch-vector components evolve as

$$\langle\sigma_x\rangle(t) = \cos(\omega_0 t), \qquad \langle\sigma_y\rangle(t) = \sin(\omega_0 t), \qquad \langle\sigma_z\rangle(t) = 0, \tag{1.18}$$

i.e., the Bloch vector precesses in the equatorial plane around the z axis at the Larmor frequency ω_0.

Under $\hat{H}_0$ alone, a generic qubit state acquires a relative phase between $|0\rangle$ and $|1\rangle$, corresponding to a rotation around the z axis on the Bloch sphere. An additional time-dependent drive term is needed to implement controlled rotations around axes in the equatorial plane, thereby enabling universal single-qubit control.

1.1.5 Driven control: ESR/EDSR and Rabi oscillations

Single-qubit gates are realized by applying an oscillating drive near the qubit resonance frequency, i.e., the Larmor frequency. This drive can be produced by an oscillatory magnetic field transverse to $\mathbf{B}_0$ (electron spin resonance, ESR) or, more commonly in scalable semiconductor devices, by an oscillating electric field that couples to the spin via spin-orbit interaction and/or a magnetic-field gradient (electric dipole spin resonance, EDSR) [Koppens et al., 2006]. In both cases, the driven dynamics can be modeled by an effective Hamiltonian of the form

$$\hat{H}_{\text{lab}}(t) = \frac{\hbar\omega_0}{2}\sigma_z + \hbar\Omega\cos(\omega_{\text{mw}}t + \phi)\,\sigma_x, \tag{1.19}$$

where ω_{mw} and ϕ denote the microwave angular frequency and phase, while Ω, determined by the drive amplitude and coupling mechanism, sets the rotation speed.

Rotating frame and rotating-wave approximation (RWA). To make Eq. (1.19) analytically tractable and connect directly to Bloch-sphere rotations, it is convenient to move to a reference frame rotating at ω_{mw} around z with the unitary transformation

$$\hat{U}_{\text{ref}}(t) = \exp\left(-i\frac{\omega_{\text{mw}}t}{2}\sigma_z\right). \tag{1.20}$$

The Hamiltonian in the rotating frame is

$$\hat{H}_{\mathrm{rot}}(t) = \hat{U}_{\mathrm{ref}}^{\dagger}(t)\,\hat{H}_{\mathrm{lab}}(t)\,\hat{U}_{\mathrm{ref}}(t) - i\hbar\,\hat{U}_{\mathrm{ref}}^{\dagger}(t)\,\frac{d}{dt}\hat{U}_{\mathrm{ref}}(t), \tag{1.21}$$

which contains a slowly varying term and fast oscillating components at approximately $2\omega_{\mathrm{mw}}$. When the drive is weak compared to the carrier, i.e., $\Omega \ll \omega_0$, the fast oscillating terms can be neglected (rotating-wave approximation), yielding a time-independent effective Hamiltonian:

$$\hat{H}_{\mathrm{RWA}} = \frac{\hbar\Delta}{2}\sigma_z + \frac{\hbar\Omega}{2}\left(\cos\phi\,\sigma_x - \sin\phi\,\sigma_y\right), \qquad \Delta = \omega_{\mathrm{mw}} - \omega_0. \tag{1.22}$$

Rabi frequency and pulse-area rule. On resonance ($\Delta = 0$), Eq. (1.22) reduces to a pure rotation in the equatorial plane. The propagator is

$$\hat{U}(T) = \exp\left[-i\frac{\Omega T}{2}\left(\cos\phi\,\sigma_x - \sin\phi\,\sigma_y\right)\right], \tag{1.23}$$

which matches the general rotation operator of Eq. (1.11) with rotation angle $\vartheta = \Omega T$ and axis $\hat{n} = (\cos\phi, -\sin\phi, 0)$. Therefore, the drive phase ϕ selects the rotation axis (X- versus Y-like rotations), while the product ΩT ("pulse area") sets the rotation angle. In particular,

$$T_{\pi} = \frac{\pi}{\Omega}, \qquad T_{\pi/2} = \frac{\pi}{2\Omega}. \tag{1.24}$$

Experimentally, sweeping the pulse duration T at fixed amplitude produces oscillations in the measured population, known as *Rabi oscillations* [Koppens et al., 2006].

Off-resonant drive. For finite detuning ($\Delta \neq 0$), the qubit rotates around a tilted axis with the generalized Rabi angular frequency

$$\Omega_{\mathrm{tot}} = \sqrt{\Omega^2 + \Delta^2}. \tag{1.25}$$

Assuming that the qubit is initialized in an energy eigenstate, the transition probability after a rectangular pulse of duration T takes the standard form

$$P_1(T) = \frac{\Omega^2}{\Omega_{\mathrm{tot}}^2}\,\sin^2\!\left(\frac{\Omega_{\mathrm{tot}}T}{2}\right), \tag{1.26}$$

showing that detuning both reduces the oscillation visibility and changes the oscillation frequency. This relation is routinely used to extract Ω from time-domain measurements and to calibrate the pulse duration for π and $\pi/2$ operations.

Equations (1.22)–(1.26) highlight that high-fidelity gates require accurate control of carrier frequency (small Δ), carrier phase ϕ, and pulse amplitude/duration (setting ΩT).

1.1.6 Phase noise, coherence, and gate fidelity

In practice, the resonant drive used to implement single-qubit rotations is generated by a local oscillator (LO) and therefore inherits its phase fluctuations. A convenient model is to include a stochastic phase term $\phi_N(t)$ in the carrier,

$$B_1(t) = B_1 \cos\big(\omega_{\mathrm{LO}} t + \phi + \phi_N(t)\big), \tag{1.27}$$

so that the driven qubit Hamiltonian, assuming the drive couples along x, becomes

$$\hat{H}(t) = \frac{\hbar\omega_0}{2}\sigma_z + \frac{\hbar\Omega}{2}\cos\big(\omega_{\mathrm{LO}} t + \phi + \phi_N(t)\big)\,\sigma_x, \tag{1.28}$$

where $\Omega = \gamma_e B_1$ is the on-resonance Rabi frequency. Transforming into the rotating frame at ω_{LO} and applying the rotating-wave approximation, LO phase noise appears as an effective dephasing term proportional to the instantaneous frequency fluctuation $\dot{\phi}_N(t)$, namely

$$\delta h_z^{(\mathrm{LO})}(t) = -\frac{1}{2}\,\dot{\phi}_N(t), \tag{1.29}$$

which is formally indistinguishable from a longitudinal noise bath that causes phase diffusion of the qubit state.

This relation connects oscillator specifications to qubit dephasing spectra. Denoting by $S^{(1)}_{\phi_N}(\omega)$ the unilateral power spectral density (PSD) of the LO phase fluctuations, the corresponding unilateral dephasing PSD entering filter-function calculations is

$$S_z^{(1)}(\omega) = \frac{1}{4}\,\omega^2 S^{(1)}_{\phi_N}(\omega). \tag{1.30}$$

Most LOs specify instead the single-sideband phase noise $L(\omega)$ in dBc/Hz, defined as $L(\omega) \equiv \frac{1}{2}S^{(1)}_{\phi_N}(\omega)$. Using $\tilde{L}(\omega) = 10\log_{10} L(\omega)$, one obtains the practical relation [Ball et al., 2016]

$$S_z^{(1)}(\omega) = \frac{1}{2}\,\omega^2\, 10^{\tilde{L}(\omega)/10}. \tag{1.31}$$

Finally, the impact on coherence and gate fidelity can be quantified through an overlap integral between $S_z^{(1)}(\omega)$ and a control-dependent filter transfer function [Green et al., 2013, Biercuk et al., 2011]. For generic single-qubit operations of duration τ, the average fidelity can be written approximately as

$$F_{\mathrm{av}}(\tau) \approx \frac{1}{2}\Big(1 + e^{-\chi(\tau)}\Big), \qquad \chi(\tau) = \frac{1}{\pi}\int_0^\infty d\omega\, \frac{S_z^{(1)}(\omega)}{\omega^2} \sum_{l\in\{x,y,z\}} G_{z,l}(\omega), \tag{1.32}$$

where the functions $G_{z,l}(\omega)$ depend on the specific control sequence, such as free evolution, Ramsey, a π-pulse, a driven X_π gate, or dynamically corrected gates. This framework makes explicit that LO phase noise degrades fidelity in a band-selective way: only the spectral components of $\tilde{L}(\omega)$ that overlap the filter function of the chosen operation contribute significantly to the error.

In realistic devices, the qubit is never perfectly isolated: coupling to the environment leads to relaxation and loss of phase coherence, which ultimately limit the time available for coherent control and set stringent requirements on gate fidelity. These effects are commonly quantified by three characteristic times. The energy-relaxation time T_1 describes the decay of population from the excited state to the ground state, while the coherence time T_2 characterizes the decay of off-diagonal elements of the density matrix. In addition, experiments often observe a shorter inhomogeneous dephasing time T_2^*, which captures the effect of quasi-static fluctuations, such as slow magnetic-field or electrostatic noise, that vary between experimental repetitions. Typically,

$$T_2^* \leq T_2 \leq 2T_1,$$

and T_2 can be extended beyond T_2^* using echo and dynamical-decoupling sequences [Green et al., 2013]. In the frequency domain, dephasing manifests as a broadening of the qubit resonance line: as an order-of-magnitude relation, the inhomogeneous linewidth satisfies

$$\Delta f \sim \frac{1}{\pi T_2^*},$$

directly linking coherence to the spectral selectivity of driven control and to the susceptibility of gates to detuning and phase/frequency noise [Veldhorst et al., 2014].

1.2 Silicon qubits

Spin states of single electrons confined in quantum dots can be exploited to realize a qubit [Loss and DiVincenzo, 1998]. Among the various qubit technologies, silicon-based qubits are particularly promising. First, the availability of isotopically purified silicon substrates significantly enhances coherence times by suppressing hyperfine-induced dephasing [Veldhorst et al., 2014]. Second, silicon benefits from a mature CMOS manufacturing ecosystem, offering a clear route toward scalable, high-density integration [Zwanenburg et al., 2013]. Finally, silicon qubits are naturally compatible with CMOS control electronics, and the possibility of operating them in the 1–4 K range is being actively investigated [Petit et al., 2018]. This may enable tighter co-integration of the quantum device and its classical control stack, for example by exploiting heterogeneous integration [Vahidpour et al., 2017].

1.2.1 Coherence and isotopic purification

The main source of decoherence for electron spins in silicon is the coupling to the bath of nuclear spins in the host material [Zwanenburg et al., 2013]. Natural silicon contains only a 4.7% concentration of the spin-carrying ^{29}Si isotope. This concentration can be further reduced, and isotopically purified substrates consisting almost entirely of spinless ^{28}Si can be obtained; for example, [Andreas et al., 2011] reports a ^{29}Si concentration of 5×10^{-5}. In such purified substrates, coherence times as long as $T_2 = 120$ μs have been demonstrated [Veldhorst et al., 2014].

1.2.2 CMOS compatibility and manufacturing scalability

Silicon benefits from the world's most mature semiconductor technology base, namely CMOS fabrication. The possibility of defining quantum dots and electrostatic gates using multi-layer metal stacks, aligned lithography, and well-controlled process modules suggests a credible path toward high-density qubit arrays and more reproducible devices.

Spin qubits in semiconductors may also be integrated with classical integrated-circuit technology, including processing, memory, and signal routing. Integration on chip is particularly natural because quantum-dot qubits rely on gate electrodes in a way that is conceptually similar to standard field-effect devices. Integration may also occur at the system level, with multiple chips communicating with one another [Vandersypen and Eriksson, 2019].

1.2.3 Typical sequence of silicon-qubit operations

From the control perspective, the operation of a semiconductor spin qubit follows a well-defined sequence. First, the device is biased through suitable DC gate voltages in order to electrostatically define the quantum dot and establish the desired working point. The qubit is then initialized into a known state. Once the operating regime has been set, fast voltage pulses can be applied to modify the electrostatic configuration, for instance by changing detuning or barrier conditions, while resonant microwave bursts are used to perform coherent single-qubit rotations. Finally, the qubit state is measured through an appropriate readout mechanism, often based on spin-to-charge conversion and RF reflectometry.

Since relaxation, dephasing, and control non-idealities progressively degrade the stored quantum information, the entire sequence must be executed within a timescale compatible with the coherence properties of the device. This sequence of operations makes clear that the qubit cannot be operated without a dedicated electronic interface capable of locally generating and coordinating bias voltages, pulsed control signals, microwave excitation, and readout support [Koppens et al., 2006, Veldhorst et al., 2014, Park et al., 2021].

Part II

Control-chip Architecture

Chapter 2

Cryogenic Control Chip for Semiconductor Spin Qubits

After introducing the physical principles of semiconductor spin qubits in Chapter 1, this chapter focuses on the classical integrated-circuit architecture required to operate them. The goal is to describe the electronic interface that generates, stores, and coordinates the signals needed for biasing, manipulation, and, more generally, control of the quantum device.

In the considered scenario, the qubit is fabricated separately, while the controller is implemented in TSMC 65 nm CMOS and operated at cryogenic temperature, around 4 K. This architectural choice is motivated by one of the main obstacles to the scalability of quantum processors: although qubits operate at cryogenic temperature, their control is still often performed through bulky room-temperature instrumentation. Such an approach is acceptable for few-qubit experiments, but it becomes increasingly inefficient as the system scales up because of the large number of cables, the associated thermal load, the footprint of the instruments, and the latency between the classical electronics and the quantum processor [Sebastiano et al., 2017].

For this reason, a major research direction consists in moving at least part of the classical interface closer to the qubit plane, by exploiting cryogenic CMOS electronics as an intermediate layer between room-temperature digital control and the milliKelvin quantum device [Patra et al., 2018, Geck et al., 2020, van Dijk et al., 2019]. In semiconductor spin-qubit platforms, such an interface must provide several classes of signals: accurate DC gate voltages to define the electrostatic working point, fast low-amplitude voltage pulses for tuning and control, microwave bursts for coherent qubit rotations, and support for the readout path [Vandersypen et al., 2017].

In the reference architecture considered in this thesis, the target specifications include tuned gate DC voltages up to approximately ± 1 V, gate pulses of a few millivolts, and microwave bursts with carrier frequency around 10 GHz. These signals must be generated with sufficient precision and stability to fit within the coherence window of silicon spin qubits, which can reach 120 μs in isotopically enriched devices [Veldhorst et al., 2014]. Within this broad architecture, the main block developed in this work is the RF source,

namely a cryogenic 10 GHz LC-VCO with 9.5–10.5 GHz tuning range, but it is essential to first place it in the context of the complete control system.

2.1 Architectural motivation from cryogenic-controller literature

The idea of a cryogenic control chip does not arise only from circuit-integration convenience, but from a broader system-level scalability problem. As discussed by Sebastiano *et al.*, the conventional approach based on room-temperature instrumentation becomes increasingly impractical as the number of qubits grows, because of the wiring burden between cryogenic quantum devices and external electronics. For this reason, cryogenic CMOS interfaces have been proposed as an intermediate layer, with most of the electronic interface operating around 4 K, where the available cooling power is significantly larger than at the qubit stage [Sebastiano et al., 2017].

This architectural direction has progressively evolved toward increasingly integrated cryogenic controllers. A particularly relevant example is the fully integrated cryo-CMOS SoC reported by Park *et al.*, which combines state manipulation, readout, high-speed gate pulsing, frequency multiplexing, and local digital supervision in a single platform for spin-qubit control. This demonstrates that a partitioned architecture including RF generation, gate-voltage control, memory resources, and digital coordination is fully consistent with state-of-the-art integrated solutions [Park et al., 2021].

At the same time, the literature also shows that the RF path is itself a critical building block. Bardin *et al.* presented a cryogenic 4-to-8 GHz pulse modulator for scalable quantum computing, highlighting that cryogenic control circuits must simultaneously satisfy high spectral performance and stringent power constraints. For this reason, in the present thesis the first circuit-level block developed within the broader control-chip architecture is precisely the RF source [Bardin et al., 2019].

These examples motivate the architectural approach adopted in this work and justify the use of a partitioned cryogenic controller in which bias generation, memory, sequencing, RF generation, and readout support are treated as coordinated subsystems.

2.2 Top-down definition of the control architecture

The proposed control chip is defined following a top-down approach, starting from system-level requirements and considering the following design constraints:

- operating temperature of 4 K;
- implementation in TSMC 65 nm CMOS technology;
- nominal supply voltage of 1.2 V;
- additional 2.5 V supply domain limited to the qubit-gate interface;
- target die area of approximately 1 mm^2;

- fewer than 20 available external pins, including power, I/O, and debug connections;
- low-power operation compatible with the limited cooling budget available at 4 K (few Watts);
- chip wire bonding as the packaging/interconnection solution;
- room-temperature digital control;
- reduced I/O speed, in the MHz range, to mitigate signal-integrity issues at the cryogenic interface;
- readout based on a single-electron transistor (SET).

These constraints naturally favor a partitioned mixed-signal architecture in which configuration data are stored locally, analog bias values are generated on chip, and the microwave control path is synthesized close to the qubit array. In particular, the limited pin budget motivates the adoption of a serial communication interface and local on-chip memory, so that the relevant control data can be loaded without requiring continuous high-speed communication between the room-temperature controller and the cryogenic signal-generation blocks [Vandersypen et al., 2017, Patra et al., 2018, Geck et al., 2020, van Dijk et al., 2019].

At a high level, the chip is divided into five main functional domains: a control unit, a memory subsystem, a bias-generation subsystem, a qubit-rotation subsystem, and a readout subsystem. A digital interface at room temperature communicates with the cryogenic controller and provides the data required to configure the system. This organization is shown in Fig. 2.1.

This partition is motivated by the distinct nature of the signals required by the qubit. The memory stores the configuration data and pulse information associated with qubit operation. Because qubit control requires low-latency and deterministic access to bias and waveform data, local on-chip memory is adopted to avoid continuous fast communication between the room-temperature digital controller and the cryogenic signal-generation blocks. The control unit supervises timing, addressing, and enable signals, so that the digital word is transferred to the bias-generation subsystem according to the operation being performed. The bias-generation subsystem converts stored digital words into the analog voltages applied to the gate electrodes. The qubit-rotation path generates the microwave signals required for ESR/EDSR-based operations. Finally, the readout section supports measurement of the qubit state.

The architectural considerations discussed above translate into a controller with the following target features:

- support for qubit rotations through a microwave-control path, with X-rotations implemented by ESR and compatibility with electrically controlled static operating points;
- separate analog generation paths for quantum-dot and SET biasing (two DACs);

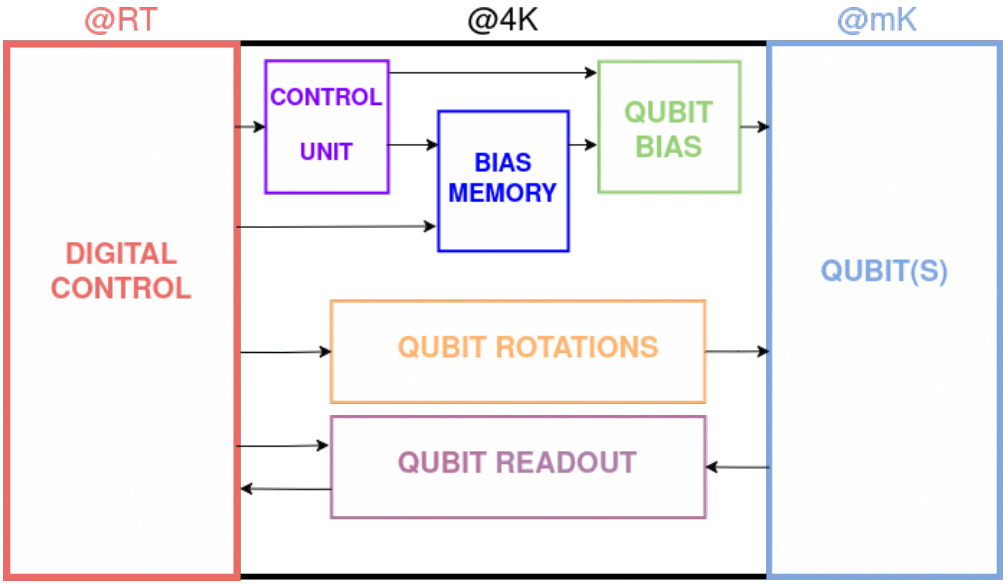

Figure 2.1. High-level architecture of the cryogenic control chip. The controller includes a control unit, local memory, bias generation, qubit-rotation circuitry, and readout support.

- local storage of DAC input words, loaded once at system startup and then reused during operation, in order to avoid continuous high-speed communication with room-temperature electronics;
- sample-and-hold stages for all QD and SET gates, periodically refreshed to preserve the applied analog voltages;
- a readout chain in which the SET current is converted into a voltage, amplified, sampled, and finally digitized before being transmitted back to the digital controller;
- serial input communication converted on chip into parallel words for DAC driving, and parallel ADC outputs serialized before transmission to room temperature;
- room-temperature digital supervision;
- a compact control unit based on finite-state-machine logic and a Johnson counter, in order to reduce the number of control signals required from outside the cryogenic stage;
- a microwave path based on a carrier around 10 GHz, generated through a PLL-based chain and then windowed or modulated according to the required qubit operation;
- support for a limited but programmable set of gate voltages associated with the different qubit operations.

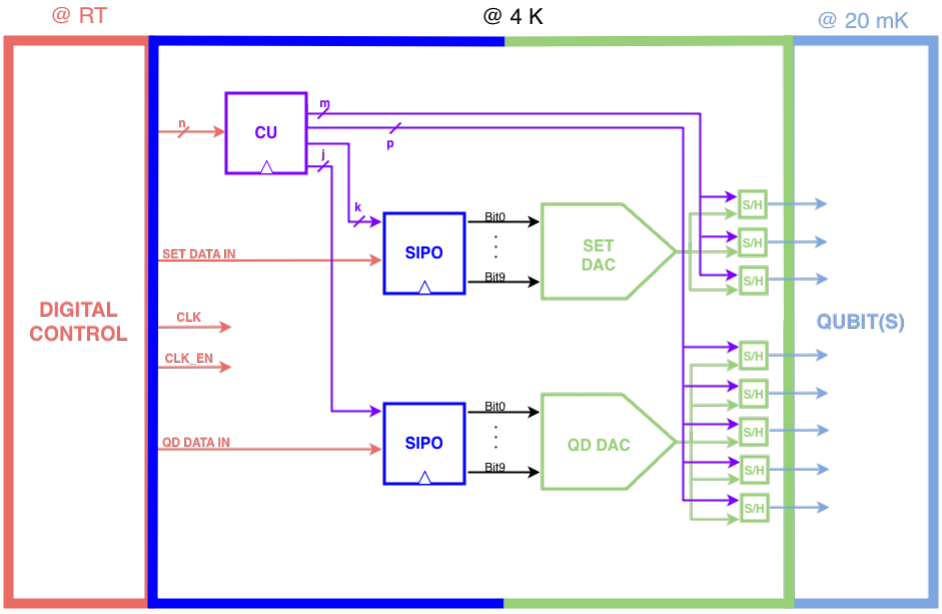

Figure 2.2. Reference architecture for memory, control, and bias generation. The DAC outputs are followed by sample-and-hold stages that drive the qubit electrodes.

2.3 Bias-generation subsystem

Besides the RF chain, the controller must generate the static and quasi-static voltages required to electrostatically define the qubit. This task is performed by the bias-generation subsystem, which translates digitally stored values into analog voltages applied to the gate electrodes.

The reference architecture for memory, control, and bias generation is shown in Fig. 2.2. In this scheme, a higher-level digital source provides the input data, a cryogenic control unit handles addressing and enable signals, serial-in/parallel-out structures receive the configuration words, and the two DAC outputs are connected to sample-and-hold stages that locally maintain the voltages applied to the qubit gates.

This organization has two major advantages. First, a single DAC can be reused to periodically refresh multiple gate electrodes, thus reducing area and power consumption. Second, the sample-and-hold stages preserve the local analog values seen by the qubit. This design philosophy is also supported by the literature, where DAC-plus-sample-and-hold solutions are often adopted in cryogenic qubit controllers because of the high input impedance of the gate electrodes [Geck et al., 2020].

From a design viewpoint, the bias-generation subsystem must satisfy several requirements simultaneously: sufficient voltage range, adequate resolution for fine tuning, low noise, and compatibility with cryogenic operation.

2.4 Memory and data transfer

A central aspect of the architecture is the use of local on-chip memory. In the considered system, the values required by the different qubit gates are not continuously streamed from room temperature during operation. Instead, the relevant data are first loaded into local memory and then reused by the cryogenic circuitry. This approach greatly reduces the need for continuous high-speed communication between the digital controller and the cryogenic analog front-end, which would otherwise complicate the interface and increase the number of required connections.

The use of local memory is especially beneficial because the information required by the qubit evolves on heterogeneous time scales. Bias data are relatively static and mainly determine the operating point, whereas pulse and sequence data are more dynamic and are associated with specific qubit operations. Storing both classes of information locally allows the cryogenic controller to respond with low latency and deterministic timing, while avoiding repeated communication with room-temperature electronics.

The conceptual organization of the memory is shown in Fig. 2.3. The architecture is based on serial-in/parallel-out structures and multiplexed readout logic, which provide a reasonable compromise between routing simplicity and rapid access to the stored words. The words needed to generate the bias voltages are organized in rows: each row contains one specific word, while words related to the same gate are placed sequentially and then multiplexed. The first layer of multiplexers selects the operation to be performed, whereas the second multiplexing layer selects the target gate (e.g., plunger or barrier) on which the operation is applied.

This approach is well aligned with the cryogenic-controller literature. For instance, Geck *et al.* proposed separate memory resources for bias values and RF sequences, relying on local storage to avoid continuous streaming of control data from outside the cryogenic stage [Geck et al., 2020].

2.5 Control unit and local sequencing

The control unit is the supervisory digital block of the chip. Its role is to coordinate memory access, select the proper data associated with the current operation, generate the enable signals for the analog subsystems, and ensure deterministic timing for the complete control sequence.

A simplified representation of the control unit is shown in Fig. 2.4. The architecture uses a LUT-based structure together with a Johnson counter to generate gate-selection and operation-selection signals.

In this framework, the LUT stores the correspondence between digital commands and the internal control actions required by the chip, while the Johnson counter provides the sequencing mechanism. Although the control unit is not the main design contribution of this thesis, it is an essential architectural block because it defines the timing framework within which the analog and RF subsystems operate.

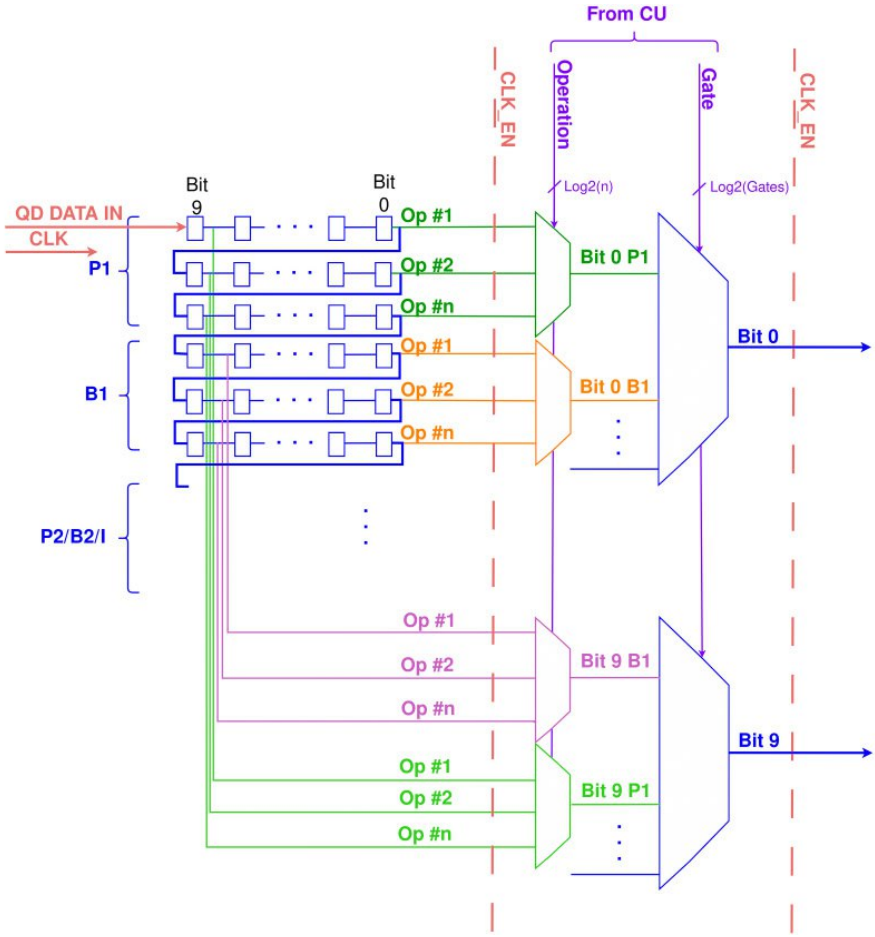

Figure 2.3. Conceptual memory organization based on serial-in/parallel-out structures and multiplexed readout.

2.6 RF front-end for qubit rotations

Among all subsystems of the controller, the RF front-end is the most relevant for this thesis, since it contains the block that has been designed and analyzed in detail: the cryogenic 10 GHz LC-VCO. The purpose of the RF path is to generate a spectrally clean microwave carrier, condition it through the frequency-synthesis chain, and modulate it so as to obtain the bursts required for qubit manipulation.

2.6.1 General RF chain

A simplified representation of the qubit-rotation path is shown in Fig. 2.5. The digital system controls the frequency multiplication and provides the reference timing signal. A PLL-based synthesis chain generates the carrier, the signal is then conditioned and

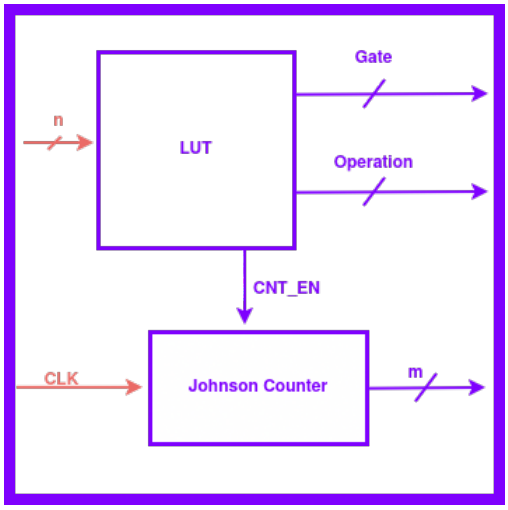

Figure 2.4. Simplified architecture of the control unit.

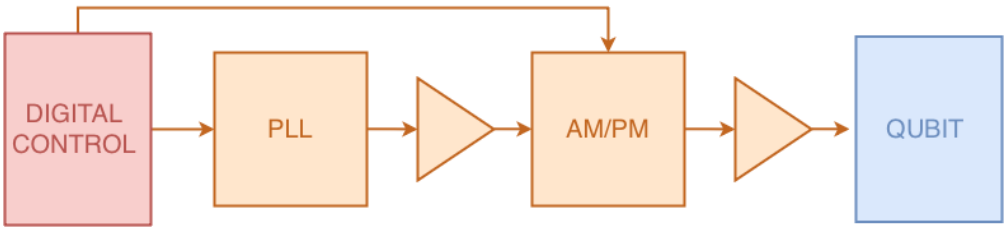

Figure 2.5. Simplified RF chain for qubit manipulation.

amplitude-modulated, and the resulting microwave burst is delivered to the qubit.

This structure separates two conceptually different functions: carrier synthesis and pulse formation. The frequency-generation stage must provide the correct oscillation frequency with sufficiently low phase noise, while the amplitude-modulation stage must shape the burst according to the desired gate operation. This modularity is advantageous because it allows each block to be optimized according to its dominant specification, while also leaving room for more advanced pulse-generation schemes, such as Y-rotations, Gaussian envelope shaping, or frequency multiplexing.

2.6.2 PLL-based frequency synthesis

A practical way to generate the microwave carrier is to adopt a phase-locked loop. The PLL locks a tunable oscillator to a stable reference and therefore enables frequency programmability while preserving long-term stability. This is particularly important in spin-qubit systems, where calibration and compensation of process variations, temperature shifts, or frequency multiplexing may require tuning of the local-oscillator frequency.

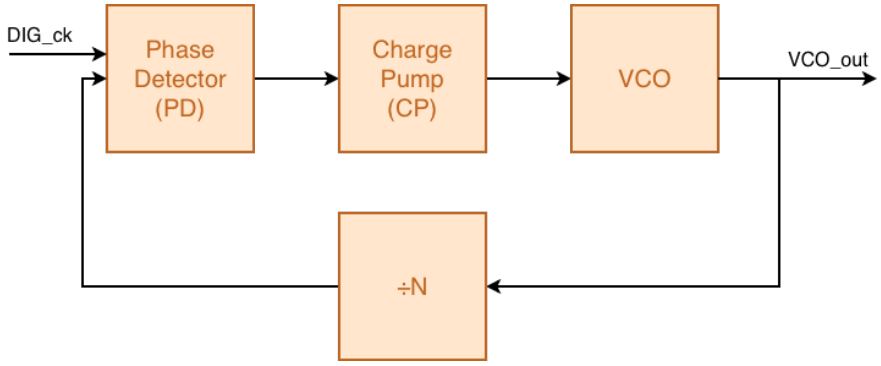

Figure 2.6. Conceptual PLL scheme adopted for microwave carrier synthesis.

A conceptual representation of the PLL is shown in Fig. 2.6. The architecture includes the standard phase detector, charge pump, divider, and VCO blocks, in line with conventional integrated RF synthesis schemes [Razavi, 2020].

Within the full controller, the PLL acts as the backbone of the microwave generation path. Its role is not only to produce the required carrier, but also to provide a programmable interface between the digital control domain and the analog RF oscillator. This is especially useful when dealing with silicon qubits, whose device characteristics may deviate from room-temperature expectations.

2.6.3 LC-VCO as the key RF building block

The core of the PLL is the voltage-controlled oscillator. In the considered architecture, an LC-VCO implemented in TSMC 65 nm CMOS is adopted because LC-based oscillators generally provide better phase-noise performance than ring-based alternatives, especially in applications where spectral purity is a first-order requirement [Hajimiri and Lee, 1998, Rael and Abidi, 2000]. This is precisely the case for spin-qubit manipulation, where phase fluctuations of the local oscillator translate into dephasing of the driven qubit and hence into gate infidelity [Ball et al., 2016, Green et al., 2013, van Dijk et al., 2018].

The adoption of an LC-VCO is therefore motivated not only by frequency-generation needs, but also by the requirement to preserve qubit fidelity as much as possible. In this thesis, the LC-VCO represents the first concrete circuit-level realization within the broader control architecture. Its target operating frequency is approximately 10 GHz, consistent with microwave-driven control of silicon spin qubits. More importantly, its role is not limited to frequency generation: because of the direct link between oscillator phase noise and qubit fidelity, the VCO is one of the blocks whose non-idealities most directly affect the feasibility of the overall controller.

2.6.4 Amplitude modulation and burst generation

The microwave carrier produced by the PLL and VCO must then be converted into bursts suitable for qubit manipulation. This is achieved through the amplitude- and phase-modulation stage, denoted as AM/PM in the architectural scheme. Phase modulation is used to select the rotation axis, thereby enabling different operations such as X- and Y-rotations. Amplitude modulation, on the other hand, is used to gate and shape the microwave burst; for a given drive strength, the pulse duration sets the rotation angle according to the Rabi frequency.

At the conceptual level, two main approaches are possible. A first option consists in directly gating the RF signal through a switch, which is attractive when a small number of qubits is controlled and only relatively simple pulse shapes are needed. A second option consists in using a mixer-based modulation stage, which enables more flexible envelope shaping and potentially more advanced schemes such as Gaussian pulses or single-sideband modulation. The latter becomes particularly relevant when moving toward larger qubit arrays, where channel sharing and spectral selectivity become more important.

More generally, coherent qubit manipulation requires not only a microwave source, but an RF subsystem able to coordinate waveform generation, carrier synthesis, and burst delivery. Even though the modulation block is not the main focus of this thesis, it plays an important architectural role because it transforms a continuous-wave local oscillator into a physically meaningful control signal for the qubit.

2.7 Readout subsystem

Besides bias generation and coherent control, the cryogenic controller must also support qubit-state measurement. In semiconductor spin-qubit platforms, the readout is generally not performed by directly sensing the spin degree of freedom. Instead, the spin information is first converted into a charge-dependent event through a spin-to-charge conversion mechanism, and the resulting charge variation is then detected by a nearby charge sensor.

In the reference architecture considered here, the readout path is based on a single-electron transistor (SET) used as a charge sensor. SET-based sensing is well established in silicon quantum-dot devices, where it enables sensitive detection of charge transitions down to the single-electron level and can therefore be exploited as the front-end of spin readout after spin-to-charge conversion [Yang et al., 2011, Veldhorst et al., 2014]. More specifically, once the qubit state has been mapped onto a charge configuration, the corresponding variation in the electrostatic environment modulates the SET current. The role of the electronic readout chain is then to convert this weak analog signal into a form suitable for further processing and transmission to the room-temperature controller.

A conceptual block diagram of the adopted readout path is shown in Fig. 2.7. The current generated by the SET sensor is first converted into a voltage by a current-to-voltage (I/V) stage. The SET operating point is established by a dedicated DAC, following the same general biasing philosophy adopted for the qubit-control electrodes. The resulting voltage is then amplified by a low-noise amplifier (LNA), sampled by a sample-and-hold block, and finally digitized by an ADC. The digital output can then be serialized and transmitted back to the external control electronics. In parallel, control signals such

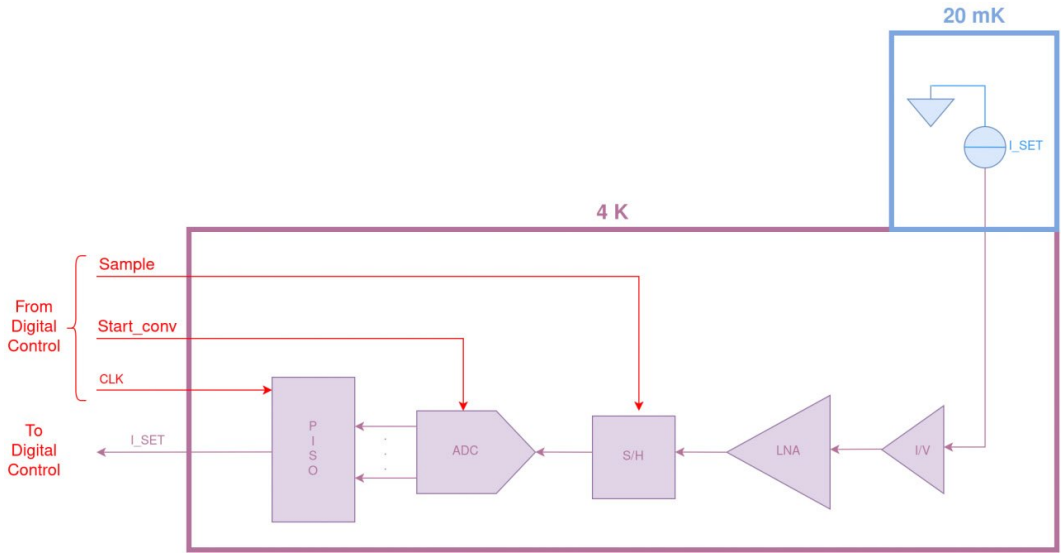

Figure 2.7. Conceptual readout chain of the cryogenic controller. After spin-to-charge conversion in the quantum device, the SET current is converted into a voltage, amplified, sampled, and digitized before transmission to the digital controller.

as the sampling clock and the start-conversion command are provided from the digital domain to coordinate the readout sequence.

From an architectural viewpoint, this organization separates the quantum transduction mechanism from the classical signal-conditioning chain. The quantum device and the SET operate at the milliKelvin stage, where the charge-sensitive measurement takes place, while the first stages of analog conditioning and data acquisition are placed at the 4 K controller level. This partition is attractive because it reduces the need to route very weak analog signals all the way to room temperature, thereby mitigating bandwidth limitations, interference, and wiring complexity.

Although the detailed design of the readout electronics is outside the scope of this thesis, its inclusion in the architectural discussion is important for completeness. In fact, a practical spin-qubit controller must not only generate DC biases, pulses, and microwave bursts, but must also provide a compatible measurement back-end able to acquire the charge-sensor response with sufficient sensitivity, timing accuracy, and scalability. For this reason, the readout path is considered here as an essential subsystem of the overall cryogenic control chip.

2.8 Technology choice and cryogenic operation at 4 K

The complete control architecture is conceived in TSMC 65 nm CMOS and intended to operate around 4 K. This temperature should be regarded as a practical compromise between close integration and thermal budget. Placing the electronics at room temperature leads to the scalability issues discussed above, while placing all control circuitry directly

at the milliKelvin stage would impose extremely severe dissipation constraints. The 4 K stage therefore emerges as a realistic intermediate level for near-term cryogenic control electronics [Vandersypen et al., 2017, Patra et al., 2018, van Dijk et al., 2019, Park et al., 2021].

The choice of a 65 nm CMOS node is also consistent with the mixed-signal nature of the controller. It offers a suitable compromise between integration density, RF capability, supply-voltage headroom, and design maturity, making it a practical technology platform for the co-integration of digital logic, bias circuitry, and microwave building blocks.

However, cryogenic CMOS operation is not transparent from a design viewpoint. As will be discussed in Chapter 4, device parameters may differ significantly from their room-temperature values due to threshold-voltage shifts, mobility changes, incomplete dopant ionization, and other low-temperature effects. These deviations are particularly critical in analog and RF blocks. The implemented LC-VCO is intended to operate under cryogenic conditions and must therefore be analyzed using suitably adapted models. The oscillator is not treated as a generic room-temperature RF circuit, but as a building block whose behavior must remain meaningful within the final 4 K controller.

2.9 Role of the proposed chip within this thesis

The architecture presented in this chapter should be interpreted as a reference control-chip architecture for semiconductor spin qubits, rather than as a chip already fully implemented in all its subsystems. Its purpose is to define the system-level framework within which the first blocks are currently being developed.

Within this broader architecture, the main contribution of the thesis is the RF source, namely the cryogenic 10 GHz LC-VCO with 9.5–10.5 GHz tuning range. This choice is motivated by the fact that microwave-driven qubit control places particularly stringent requirements on spectral purity, and therefore on oscillator design. Among all the subsystems of the control chip, the VCO is one of the blocks whose circuit-level imperfections most directly translate into reduced qubit fidelity.

At the same time, presenting the full controller is essential because it shows that the oscillator is not intended as a stand-alone demonstrator. On the contrary, it is conceived as the first validated subsystem of a wider mixed-signal cryogenic platform that will eventually include local memory, bias generation, modulation, and digital control. This system-level perspective justifies the architectural discussion developed in the present chapter and motivates the detailed focus on the RF block in the following chapters.

2.10 Conclusions

This chapter presented the reference cryogenic control-chip architecture considered for the operation of semiconductor spin qubits. Starting from system-level constraints such as cryogenic operation, reduced wiring, limited area, and power dissipation, the discussion introduced a partitioned mixed-signal solution including local memory, a supervisory control unit, a bias-generation front-end, a qubit-rotation subsystem, and readout support.

The electronic interface is a central element in the scalability of semiconductor quantum processors. Moving part of the classical control closer to the qubit plane can reduce wiring complexity, thermal load, and latency, but it also requires the design of cryogenic integrated circuits capable of locally storing control data, generating accurate bias voltages, and synthesizing spectrally pure microwave signals.

Within this architecture, particular attention was devoted to the RF front-end, since it contains the main block developed in this thesis: a 10 GHz LC-VCO in TSMC 65 nm CMOS with 9.5–10.5 GHz tuning range for cryogenic operation. The next chapters will therefore focus on the design methodology, cryogenic modeling assumptions, circuit implementation, and post-layout validation of this oscillator, which represents the first concrete building block of the proposed control chip.

Part III

Design, Modeling and Validation

Chapter 3

Cryogenic LC-VCO Design for Spin-Qubit Control

3.1 System Requirements and Design Targets

3.1.1 Motivation and context

High-fidelity manipulation of semiconductor spin qubits relies on phase-coherent microwave bursts in the GHz-to-tens-of-GHz range. In view of large-scale qubit integration, it is desirable to place part of the control electronics as close as possible to the qubit plane, ideally operating at cryogenic temperature, where cooling power is limited. In this context, the voltage-controlled oscillator (VCO) becomes a key building block, since its spectral purity ultimately bounds the phase stability of the local oscillator used for qubit drive and readout. The oscillator design flow was initially developed at room temperature, while cryogenic operation was addressed in a subsequent validation step. In particular, a dedicated BSIM4 re-centering procedure was introduced afterwards to verify whether the active-device behavior remained physically consistent under low-temperature conditions, as discussed in Chapter 4.

3.1.2 Technology node and chosen topology

This work targets a cryogenic implementation in a commercial CMOS node (TSMC 65 nm). Among integrated oscillator families, differential LC oscillators are preferred. Compared with ring oscillators, this topology is better suited to achieving high frequencies together with lower phase noise [Razavi, 2020]. Indeed, LC-based oscillators are also adopted in cryogenic quantum-control architectures [Patra et al., 2018].

3.1.3 Target frequency and tuning range

The VCO is designed to operate around a target carrier frequency of 10 GHz, with a tuning range spanning 9.5 GHz to 10.5 GHz. This choice is motivated by the requirements of

spin-qubit control, where microwave excitation in the X-band is commonly employed to drive electron-spin resonance (ESR) transitions. In practice, the exact resonance frequency depends on the applied static magnetic field and device-specific parameters; therefore, a 1 GHz tuning span is intentionally included to provide margin for qubit-to-qubit variability, calibration procedures, and frequency shifts due to process dispersion and cryogenic operation. From a circuit-design standpoint, the target band sets the LC tank resonance through $f_0 \approx 1/(2\pi\sqrt{LC_{\text{tot}}})$ and directly determines the required total-capacitance tuning ratio $C_{\max}/C_{\min} = (f_{\max}/f_{\min})^2$, which guides the partitioning between coarse tuning (capacitor bank) and fine tuning (MOS varactors).

3.1.4 Output swing constraint and saturation of the cross-coupled pair

Besides frequency range and phase-noise performance, the output swing of the oscillator must also be limited in order to keep the cross-coupled NMOS pair in saturation during the oscillation cycle. Let V_o denote the single-ended peak amplitude of the oscillation around the common-mode level. At the oscillation peak, one drain voltage can be approximated as

$$V_D = V_{DD} - V_o,$$

while the corresponding gate voltage, being cross-coupled to the opposite tank node, is

$$V_G = V_{DD} + V_o.$$

If the common-source node is at V_S, the saturation condition for the conducting NMOS is

$$V_{DS} > V_{GS} - V_{TH}. \tag{3.1}$$

Substituting the instantaneous node voltages gives

$$V_{DD} - V_o - V_S > V_{DD} + V_o - V_S - V_{TH}, \tag{3.2}$$

from which

$$2V_o < V_{TH}. \tag{3.3}$$

Therefore, the single-ended peak amplitude should remain below $V_{TH}/2$, which is equivalent to requiring the differential peak amplitude to stay below V_{TH}. Neglecting short-channel effects and assuming that the tank common-mode remains close to V_{DD}, this condition provides a simple and useful upper bound for the oscillator swing.

This derivation should be regarded as a first-order large-signal estimate. In particular, it neglects short-channel effects, finite output resistance, and possible common-mode variations of the tank nodes during the oscillation cycle. It is therefore used here as a practical design guideline rather than as a strict large-signal boundary.

In the present design, with an extracted threshold voltage of approximately $V_{TH} \approx$ 427 mV, Eq. (3.3) suggests a maximum single-ended peak amplitude of about 213.5 mV. Equivalently, the differential peak amplitude should remain below approximately 427 mV, corresponding to a differential peak-to-peak swing below about 854 mV. For this reason, the oscillator is intentionally designed to operate with a moderate output swing, avoiding excessive large-signal compression of the cross-coupled core.

3.1.5 Cryogenic operation and power budget

The main drawback of operating the control electronics at 4 K (or below) is the imposition of power constraints, since the available cooling power at the cold stage is limited and system-level scaling targets a large number of control channels. For this reason, low power consumption is regarded as an important design objective throughout the development of the oscillator.

At the architectural level, a sub-mW VCO budget was considered desirable. However, this value was not treated as a strict circuit-level constraint during the actual implementation. In practice, the bias current had to be chosen as a trade-off between power dissipation, reliable start-up, oscillation amplitude, tuning-network loading, and phase-noise performance.

During the schematic-design phase, a bias current of approximately 1.2 mA was adopted as a reasonable compromise to guarantee oscillation robustness and good phase noise while keeping the power consumption limited. After layout implementation and parasitic extraction, additional loading and tank degradation required a further increase in the bias current. In particular, post-layout operation required at least 1.5 mA, while the results presented in this chapter are reported for a bias current of 1.9 mA, unless otherwise stated.

The sub-mW figure should therefore be interpreted as an architectural low-power objective rather than as a hard design constraint. Cryogenic operation further motivates careful device modeling, as MOSFET parameters (e.g., threshold voltage, mobility, and noise behavior) deviate substantially from room-temperature conditions; a dedicated modeling and characterization flow is therefore required and will be detailed in Chapter 4.

3.1.6 Phase-noise requirements and the role of filter functions

Phase noise impact on qubit control

Oscillator phase noise can be interpreted as random fluctuations of the instantaneous phase $\phi(t)$ of the carrier, which translate into timing and phase errors on driven qubit rotations. A general and physically insightful framework models an autonomous oscillator as a periodically time-varying system, where noise sources project onto the phase of the signal to a time-varying sensitivity function. This viewpoint explains how stationary and cyclostationary device noise contribute to phase diffusion and clarifies the design levers available to suppress close-in phase noise [Hajimiri and Lee, 1998]. For differential LC oscillators, additional physical insight can be obtained by tracking how thermal noise from tank losses, the active pair, and tail circuitry is converted into phase-modulation sidebands, and how flicker noise may upconvert through bias-dependent frequency-modulation mechanisms [Rael and Abidi, 2000].

From qubit-level metrics to oscillator-level specifications

To translate qubit-level performance targets into circuit-level specifications, many works adopt the filter-function formalism [van Dijk et al., 2018]. Within this framework, the error induced by classical dephasing noise can be expressed through an overlap integral between the noise power spectral density and a control-dependent filter function. In the

weak-noise regime, the average fidelity decay can be written as [Green et al., 2013]

$$F_{\mathrm{av}}(\tau) \approx \frac{1}{2}\left[1 + e^{-\chi(\tau)}\right], \tag{3.4}$$

where the decay exponent is

$$\chi(\tau) = \frac{1}{\pi}\int_0^\infty \frac{d\omega}{\omega^2}\, S_z(\omega)\, F_z(\omega). \tag{3.5}$$

Here, $S_z(\omega)$ denotes the dephasing-noise power spectral density, while $F_z(\omega)$ is the dimensionless filter function associated with the applied control sequence. Equation (3.5) shows explicitly that only the spectral components of the noise lying within the pass-band of $F_z(\omega)$ contribute significantly to the accumulated error [Green et al., 2013].

Two complementary references are useful in this context. At the system level, Patra *et al.* impose an RMS frequency-noise target of $\sigma_f < 1.9\,\mathrm{kHz_{rms}}$ for spin-qubit control. Under their assumptions, including a typical PLL bandwidth of 300 kHz, this translates to an in-band PLL phase-noise requirement of approximately $-115\,\mathrm{dBc/Hz}$; however, this value refers to the PLL-level in-band noise and should not be interpreted as the phase-noise requirement of a stand-alone free-running VCO [Patra et al., 2018].

At the oscillator level, a more direct fidelity-oriented reference is given by van Dijk *et al.*, who report that preserving a 99.9 % fidelity threshold for a resonantly driven single spin qubit at 1 MHz Rabi frequency requires the oscillator phase noise to be below approximately $\mathcal{L}(1\,\mathrm{MHz}) < -106\,\mathrm{dBc/Hz}$, assuming a $-20\,\mathrm{dB/dec}$ phase-noise slope [van Dijk et al., 2018]. For this reason, the present work adopts the latter value as the most relevant oscillator-level reference when discussing the phase-noise performance of the proposed LC-VCO.

Phase-noise and frequency-noise relations

Single-sideband (SSB) phase noise $\mathcal{L}(f)$ (in dBc/Hz) is related to the phase-fluctuation PSD $S_\phi(f)$ (in $\mathrm{rad^2/Hz}$) by [Ball et al., 2016]

$$S_\phi(f) \approx 2 \cdot 10^{\mathcal{L}(f)/10}, \tag{3.6}$$

under the standard small-noise assumption. The instantaneous frequency fluctuation is

$$\delta f(t) = \frac{1}{2\pi}\frac{d\phi(t)}{dt},$$

hence in the frequency domain

$$S_f(f) = f^2\, S_\phi(f), \tag{3.7}$$

and the RMS frequency noise integrated over $[f_L, f_H]$ is

$$\sigma_f^2 = \int_{f_L}^{f_H} S_f(f)\, df = \int_{f_L}^{f_H} f^2 S_\phi(f)\, df. \tag{3.8}$$

Equations (3.6)–(3.8), together with Eqs. (3.4)–(3.5), provide the link between circuit-level phase-noise simulations and qubit-level specifications.

Table 3.1. Top-level VCO targets adopted in this work.

Parameter	Target or reference value
Carrier frequency f_0	10 GHz
Tuning range	9.5–10.5 GHz
Supply V_{DD}	1.2 V
Temperature T	≈ 4 K
Power consumption	few-mW range
Offsets of interest	1–10 MHz
Phase-noise reference	$\mathcal{L}(1\text{ MHz}) \lesssim -106$ dBc/Hz
Output swing	differential peak amplitude $\approx$ 420 mV

3.1.7 Summary of target specifications

Table 3.1 summarizes the preliminary design targets. The tuning range Δf is left as a parameter to be refined once the frequency plan (PLL bandwidth, division ratios, and calibration margin) is finalized.

3.2 Preliminary design (top-down sizing)

This section translates the system-level targets of Section 3.1 into first-order circuit constraints. The goal is not to predict the final performance with closed-form equations, but to obtain a consistent starting point for the simulation-driven flow of Section 3.3 and, ultimately, for the detailed design verification in Section 3.4.1–3.4.8. In particular, the target carrier frequency and tuning span determine the required tank capacitance range, while the expected tank losses set the minimum transconductance and bias current needed for reliable start-up.

3.2.1 From frequency targets to tank capacitance range

The oscillation frequency of a differential LC-VCO is primarily set by the resonant tank, which can be approximated by

$$f_0 \approx \frac{1}{2\pi\sqrt{LC_{\text{tot}}}}. \tag{3.9}$$

Given the target band 9.5 GHz to 10.5 GHz and a candidate inductance L, the total capacitance required at a generic frequency f is

$$C_{\text{tot}}(f) = \frac{1}{(2\pi f)^2 L}. \tag{3.10}$$

In this work, a center-tapped inductor with $L = 778.9$ pH is selected from the PDK library, showing $Q \approx 16.2$ around 10 GHz. Using (3.10), the nominal capacitance at 10 GHz is approximately

$$C_{\text{tot}}(10\text{ GHz}) \approx 325\text{ fF}, \tag{3.11}$$

and the capacitance span required to cover the target band becomes

$$C_{\mathrm{tot}}(9.5\ \mathrm{GHz}) \approx 360\ \mathrm{fF}, \qquad C_{\mathrm{tot}}(10.5\ \mathrm{GHz}) \approx 295\ \mathrm{fF}. \tag{3.12}$$

Hence, the tuning mechanism must provide a total variation of about $\Delta C_{\mathrm{tot}} \approx 65$ fF, corresponding to roughly a 22% change around the nominal capacitance. This requirement drives the subsequent partitioning between discrete (bank) and continuous (varactor) tuning.

3.2.2 Tuning partitioning: coarse bank for range, varactor for resolution

To achieve a wide and robust tuning range while keeping control sensitivity moderate, the total capacitance is decomposed as

$$C_{\mathrm{tot}} = C_{\mathrm{fix}} + C_{\mathrm{var}}(V_{\mathrm{ctrl}}) + C_{\mathrm{bank}}(\mathbf{b}), \tag{3.13}$$

where C_{fix} includes fixed capacitances and parasitics, C_{var} is provided by a voltage-controlled MOSCAP in inversion, and C_{bank} is a digitally (bit $\mathbf{b}$) controlled capacitor bank.

The coarse tuning is assigned to the bank to absorb process/model shifts and guarantee band coverage, while the varactor is reserved for fine control and calibration. A first-order estimate of tuning sensitivity follows from differentiating (3.9):

$$K_{\mathrm{VCO}} = \frac{df}{dV_{\mathrm{ctrl}}} \approx -\frac{f}{2C_{\mathrm{tot}}} \frac{dC_{\mathrm{var}}}{dV_{\mathrm{ctrl}}}. \tag{3.14}$$

Equation (3.14) clarifies the design intent: a smaller $\frac{dC_{\mathrm{var}}}{dV_{\mathrm{ctrl}}}$ (and therefore a moderate K_{VCO}) reduces the conversion of tuning-line noise into phase noise and is thus preferred, at the expense of requiring the bank to cover most of the total ΔC_{tot}.

3.2.3 Losses, start-up constraint, and current budget

The resonator losses are commonly represented by an equivalent parallel resistance R_{p}. A first-order estimate relates R_{p} to L, Q, and ω_0:

$$R_{\mathrm{p}} \approx Q\,\omega_0\,L. \tag{3.15}$$

This estimate provides an initial sense of the required negative conductance that must be generated by the cross-coupled pair. The start-up condition can be expressed as

$$|G_{\mathrm{neg}}| \geq \frac{1}{R_{\mathrm{p}}}, \tag{3.16}$$

and a design margin is typically enforced through $|G_{\mathrm{neg}}| = \alpha/R_{\mathrm{p}}$ with $\alpha \approx 2$–3. Since cryogenic operation changes g_m, V_{TH}, and output conductances, this constraint must ultimately be verified with compact models validated at low temperature; for this reason,

in the following sections R_{p} and the required g_m are extracted and confirmed through dedicated simulations rather than relying solely on (3.15).

Finally, power dissipation remains an important design consideration because cryogenic operation limits the allowable dissipation per control channel. For this reason, the bias current cannot be increased arbitrarily and must be chosen as a trade-off between start-up margin, oscillation amplitude, phase-noise performance, and overall power consumption. In the present design, this trade-off evolves throughout the implementation flow: an initial current of about 1.2 mA is used during schematic-level design, whereas higher current levels are required after including realistic parasitics.

3.2.4 Simulation-driven verification

The equations above define the initial design envelope: the target band sets the required C_{tot} range, tuning is partitioned between C_{bank} and C_{var} to balance range and sensitivity, and the estimated losses set a minimum negative conductance under a constrained current budget. The remainder of the chapter follows a simulation-driven methodology: the tank losses are extracted from the actual PDK devices (Section 3.4.1), the active core is sized from DC operating-point sweeps (Section 3.4.2), start-up and steady-state operation are validated in transient and PSS (Sections 3.4.3 and 3.4.4), and realistic biasing and tuning networks are integrated progressively until the full VCO implementation is verified (Sections 3.4.5–3.4.8).

3.3 Simulation Setup and Verification Methodology

The VCO design and verification are carried out in **Cadence Virtuoso** using **Spectre** as the simulation engine. The workflow progressively replaces ideal blocks with realistic implementations (bias mirror and tuning network) and validates the design through a consistent set of analyses. The objective of this section is to document the simulation environment and, most importantly, the analyses and extracted metrics used in the remainder of the chapter.

3.3.1 Simulation environment

The VCO design and verification activities are carried out in *Cadence Virtuoso*, using *Spectre* as the circuit simulator and *ADE* (*Analog Design Environment*) as the main interface to configure analyses, manage parametric sweeps, and collect outputs. The schematic-level workflow is complemented by *Calibre* for physical verification (*DRC*/*LVS*) and post-layout extraction (*PEX*) once the layout stage is reached.

To initialize the EDA environment on the server, the tool setup scripts are sourced prior to launching the applications. In particular, *Cadence Virtuoso* is initialized by running:

```
source /path/to/eda/scripts/init_cadence_20_21
```

and *Calibre* is initialized by running:

```
source /path/to/eda/scripts/init_calibre
```

These scripts configure the required environment variables and paths so that the process design kit (PDK) and simulation models are correctly found by the tools.

Within *ADE*, the correct PDK model libraries are selected for the considered process corners (e.g., `tt_lib`, `ss_lib`), and the simulator options are set consistently across all runs to ensure fair comparisons between design iterations. Since the design targets cryogenic operation, the room-temperature BSIM4 models are augmented with the cryogenic parameter set described in Chapter 4. In practice, this requires verifying that the low-temperature model files are properly referenced by the simulator and that the selected corners are available in the cryogenic model set.

Finally, before running extensive sweeps (tuning curves, PSS/Pnoise, and corner checks), a short "sanity check" simulation is performed on a minimal testbench to confirm that the environment is correctly configured (model inclusion, library access, and simulator execution). This avoids systematic errors due to missing paths or inconsistent model selections.

3.3.2 Overview of analyses and extracted metrics

Several complementary analyses are employed to validate the LC-VCO at different abstraction levels, from passive-tank characterization to full oscillator phase-noise assessment. Each analysis is associated with a specific design question and a set of extracted metrics:

- **SP/AC (small-signal) analysis for passive elements:** used to characterize the resonator impedance and extract the equivalent parallel resistance R_{p}, as well as to evaluate tuning elements (varactor and switch parasitics) around the operating frequency. In the tank testbench, the differential impedance $Z_{\mathrm{diff}}(j\omega)$ is reconstructed from a differential excitation, and $R_{\mathrm{p,diff}}$ is extracted from the peak of $\Re\{Z_{\mathrm{diff}}\}$ at resonance. In the varactor testbench, the admittance $Y(j\omega)$ is used to derive $C_{\mathrm{var}}(V_{\mathrm{ctrl}}) = \Im\{Y\}/\omega$ and estimate the tuning gain.

- **DC operating-point (dcOp) analysis:** used to inspect bias conditions and extract device small-signal parameters (e.g., g_m, g_{ds}, region of operation) through `dcOpInfo`. This is the main tool for simulation-driven core sizing, where the number of fingers (or multiplicity) is swept to meet the required start-up margin while limiting parasitic capacitances.

- **Transient (tran) analysis with noise enabled:** used to verify start-up and time-domain convergence to the steady-state limit cycle. Noise is enabled and the noise bandwidth is extended well beyond the target frequency so that oscillations can build up from noise without artificial initial conditions. The main outputs are the differential waveform $v_{\mathrm{diff}}(t)$, the achieved steady-state amplitude, and qualitative start-up robustness under worst-case conditions (e.g., `ss_lib`).

- **Periodic Steady-State (PSS):** used to compute the autonomous periodic solution and extract the oscillation frequency f_0 and harmonic content under realistic steady-state conditions. PSS provides a robust basis for noise analyses and is also used to

generate frequency-versus-control curves for tuning verification (sweeping V_{ctrl} and capacitor-bank code).

- **Phase noise (Pnoise) around PSS:** used to compute the single-sideband phase-noise spectrum $\mathcal{L}(\Delta f)$ around the PSS solution. This analysis is applied consistently across design iterations (ideal bias, bias mirror, full tuning network) to quantify phase-noise degradation due to added non-idealities. When available, the *Noise Summary* tool in ADE is used to identify dominant contributors (e.g., tail current source).

3.3.3 Measurement conventions

All results are reported using differential quantities. The differential output voltage is defined as $v_{\text{diff}} = v_x - v_y$, where v_x and v_y are the two tank nodes. The oscillation frequency f_0 is extracted from PSS as the fundamental of the periodic solution. For tuning characterization, f_0 is evaluated as a function of V_{ctrl} (continuous tuning) and capacitor-bank code (discrete tuning).

The tuning gain is defined as

$$K_{\text{VCO}}(V_{\text{ctrl}}) = \frac{df_0}{dV_{\text{ctrl}}}, \tag{3.17}$$

and is either estimated from the varactor small-signal characterization through (3.22) or extracted directly from the slope of the PSS tuning curves. Since the MOSCAP characteristic is nonlinear, K_{VCO} is inherently voltage-dependent; therefore, a representative operating region is selected when reporting a single value.

3.3.4 Process-corner strategy

Most tuning and phase-noise characterizations are first performed in the nominal corner (`tt_lib`) to establish a reference behavior and speed up iterations. A transient start-up verification is additionally performed in the slow–slow corner (`ss_lib`) as a conservative reduced-g_m condition, ensuring that oscillation onset is preserved under unfavorable device drive strength.

3.4 Circuit design

3.4.1 Tank AC simulation and extraction of R_{p}

The passive tank is first dimensioned around the target oscillation frequency by selecting an inductor from the PDK library with the aid of the *Inductor Finder*, using the preliminary values derived in Section 3.2.1. The search is constrained to (i) the target inductance value, and (ii) area minimization (without sacrificing the quality factor). This step provides a realistic starting point for L and its loss profile (series resistance and substrate loss), which cannot be captured accurately with hand calculations alone. Once L is fixed, a metal-insulator-metal capacitor (`mimcap`) is used as the main fixed capacitance

element in the tank, with a reference unit sized to 12.5 μm $\times$ 12.5 μm to obtain a squared layout. The obtained capacitance value is $C_{fixed} = 323.695$ fF. `Mimcap` devices are preferred to metal-oxide-metal capacitors (`momcap`) because they offer higher capacitance per unit area. To validate the loss model used in (3.15) and obtain a technology-accurate estimate of the equivalent parallel resistance, the differential tank is simulated in AC with a dedicated impedance testbench (Fig. 3.1). The center tap is biased at V_{DD}, while the two tank nodes are excited in differential mode through a small-signal source, so that the differential impedance can be directly reconstructed as

$$Z_{\mathrm{diff}}(j\omega) \;=\; \frac{V_{\mathrm{diff}}(j\omega)}{I_{\mathrm{diff}}(j\omega)}. \tag{3.18}$$

In practice, V_{diff} is measured between the two tank nodes and I_{diff} is the current delivered by the differential excitation (using an AC current source simplifies post-processing because $Z_{\mathrm{diff}} = V_{\mathrm{diff}}/I_{\mathrm{AC}}$).

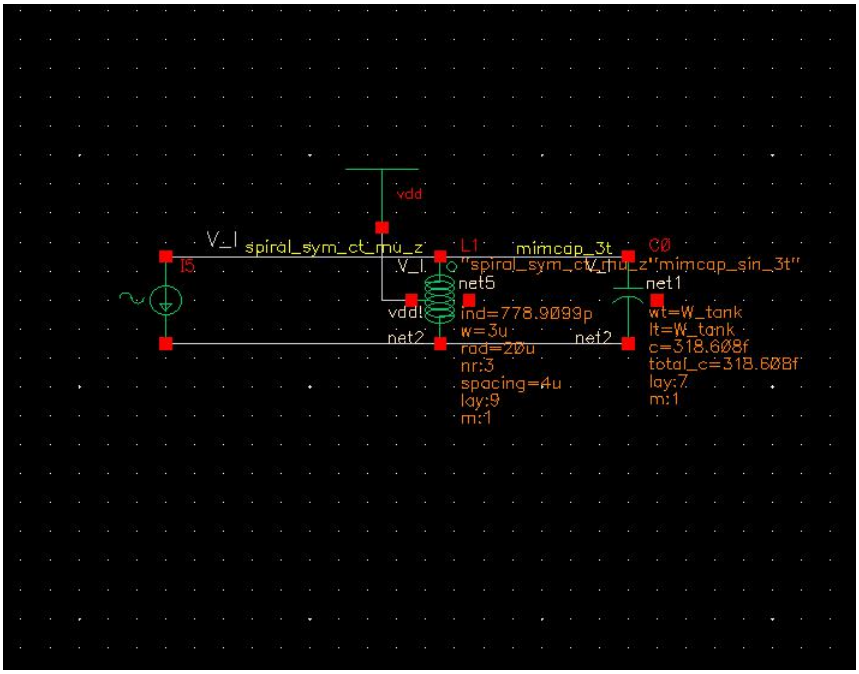

Figure 3.1. AC testbench used to excite the center-tapped differential LC tank and measure the differential impedance.

Figure 3.2 reports $\Re\{Z_{\mathrm{diff}}\}$ and $\Im\{Z_{\mathrm{diff}}\}$ versus frequency. The resonance frequency is identified by the zero crossing of $\Im\{Z_{\mathrm{diff}}\}$, while R_{p} is extracted as the peak value of $\Re\{Z_{\mathrm{diff}}\}$ at resonance:

$$R_{\mathrm{p,diff}} \;\approx\; \max_{\omega} \Re\{Z_{\mathrm{diff}}(j\omega)\} \;\approx\; 400\Omega. \tag{3.19}$$

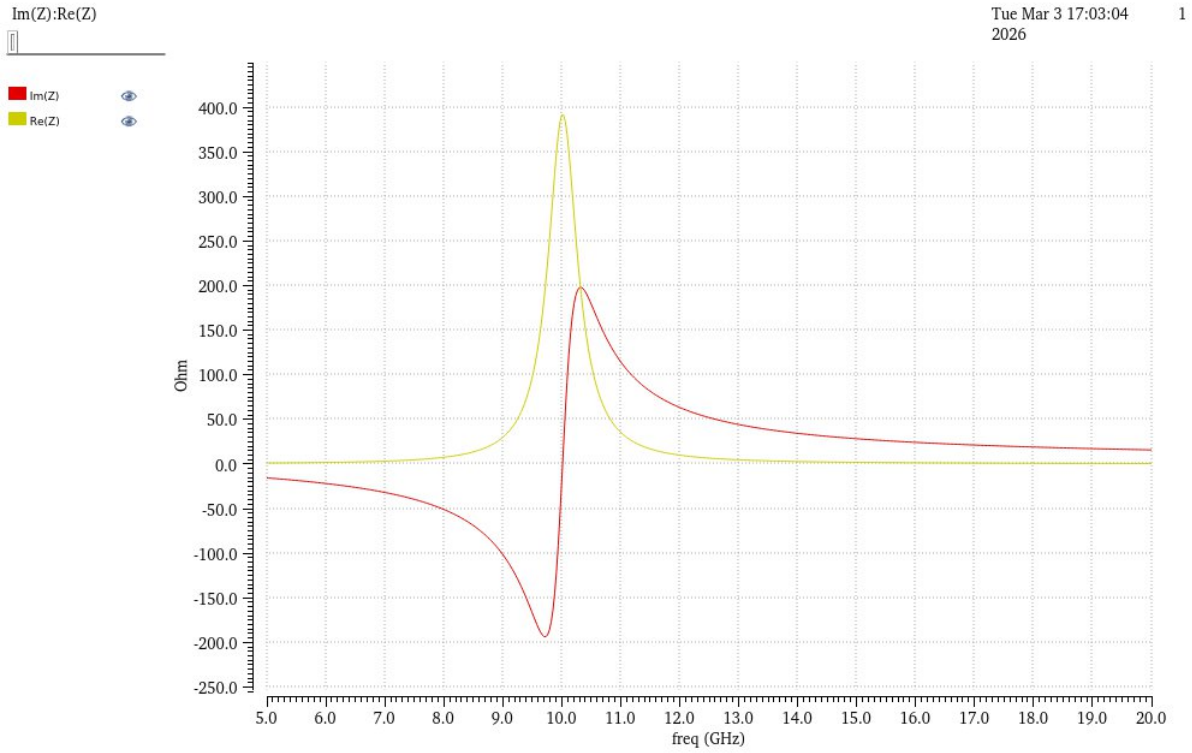

Figure 3.2. Differential tank impedance versus frequency. Resonance is at the $\Im\{Z\} = 0$ crossing; $R_{\mathrm{p,diff}}$ is extracted from the peak of $\Re\{Z\}$ at resonance.

3.4.2 Core sizing from DC operating-point sweeps

Once $R_{\mathrm{p,diff}}$ has been extracted from the actual tank, the active core is sized to guarantee start-up with the margin introduced in Section 3.2.3. The adopted approach is purely simulation-driven: instead of relying on closed-form g_m/I_D assumptions, the device transconductance is obtained from DC operating-point analysis under realistic bias conditions.

A unit NMOS of the cross-coupled pair is biased with controlled V_{GS} and V_{DS}, and its geometry is swept via the number of fingers (`nf`) at a fixed standard width of 800 nm. For each sweep point, the simulator provides g_m, g_{ds}, V_{OV} and region-of-operation flags in `dcOpInfo`. The resulting trend (Fig. 3.3) allows selecting the minimum multiplicity that meets the required negative conductance while keeping parasitics moderate.

The sizing criterion is formulated in terms of the effective negative conductance generated by the cross-coupled pair. Using the extracted tank resistance, the target transconductance is chosen with a safety factor $\alpha = 2$:

$$|G_{\mathrm{neg}}| \;\geq\; \alpha \frac{1}{R_{\mathrm{p,diff}}} \approx 5 \text{ mS}. \tag{3.20}$$

From this criterion, and according to Fig. 3.3, an initial value of 12 fingers is obtained.

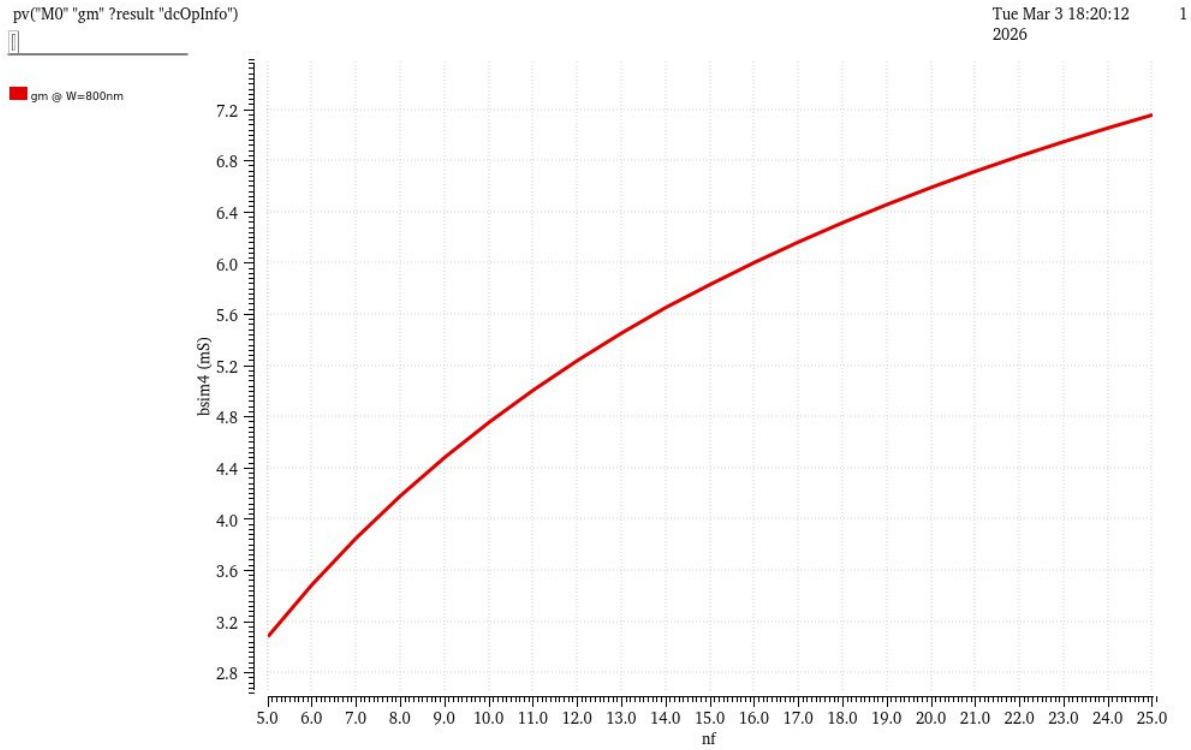

Figure 3.3. Extracted g_m versus number of fingers for a fixed unit width, obtained from `dcOpInfo`. This sweep directly supports selecting the minimum core size that satisfies the start-up margin.

3.4.3 Start-up verification and transient steady-state

The passive tank and the active core (including an ideal tail current source for the first iteration) are combined in the schematic shown in Fig. 3.4.

Start-up is verified through a transient analysis. In order to guarantee that the oscillation can build up from noise, device noise is enabled and the noise bandwidth is extended well beyond the target oscillation frequency. The adopted transient setup is reported in Fig. 3.5, where the simulation time is set to 100 ns and the maximum noise frequency to $F_{\max} = 50$ GHz.

The transient waveform of the differential output is shown in Fig. 3.6. The small initial noise component around the resonance frequency is progressively amplified until the oscillation reaches a steady-state differential amplitude of approximately 550 $\mathrm{mV_{pp}}$.

From the initial simulations, the effective parasitic capacitance in the core resulted higher than expected. Therefore, the number of fingers of the cross-coupled devices was reduced to 10 to limit the capacitive loading on the tank nodes.

To further assess robustness, the same transient analysis was repeated using the slow–slow corner library (`ss_lib`), which represents a conservative condition due to reduced device drive strength and consequently lower g_m. As shown in Fig. 3.7, oscillation still

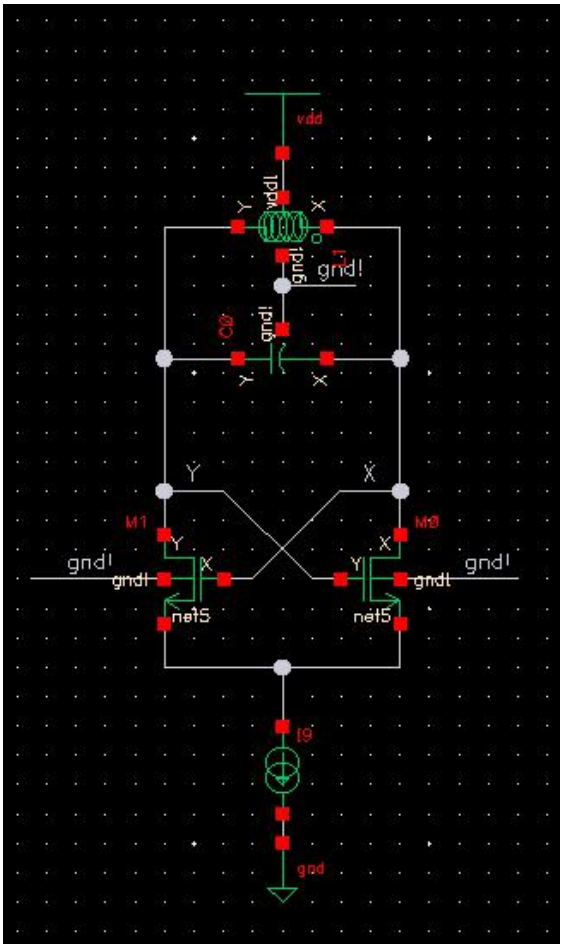

Figure 3.4. First implementation of the core oscillator using an ideal tail current source.

starts reliably, while both steady-state amplitude and oscillation frequency decrease with respect to the nominal case, as expected.

3.4.4 PSS and Pnoise simulations

After verifying start-up in the time domain, the steady-state behavior is characterized through Periodic Steady-State (PSS) and phase-noise (Pnoise) analyses. PSS is used to verify the oscillation frequency of the baseline oscillator. Pnoise analysis is exploited to establish a phase-noise reference that will be used as a benchmark to quantify the degradation introduced by subsequent circuit additions, such as the bias current mirror and the frequency-tuning network (varactors and switched-capacitor bank).

The PSS analysis is performed by using the oscillation frequency estimate obtained from transient simulations as an initial guess, allowing the solver to converge to the limit cycle without imposing an external excitation. Once convergence is reached, the

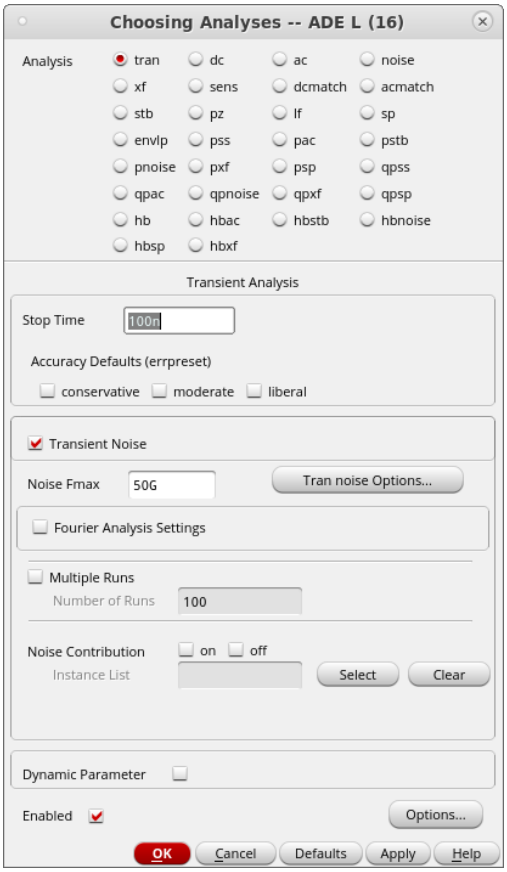

Figure 3.5. Transient simulation setup: total time 100 ns and noise bandwidth up to $F_{\max} = 50$ GHz.

fundamental frequency f_0 is extracted and compared against the target range, while the steady-state waveform is inspected to ensure that the differential swing remains compatible with the available headroom. The spectral components of the differential output voltage are reported in Fig. 3.8.

Phase noise is then evaluated with the Pnoise analysis around the PSS solution. The resulting $L(\Delta f)$ spectrum, shown in Fig. 3.9, provides the baseline noise profile of the core oscillator, which is subsequently reused as a reference when enabling additional blocks. In particular, the same analysis setup (bias current, load conditions, and observation node) is maintained across design iterations to ensure that any variation in phase noise can be attributed to the newly introduced elements rather than to changes in simulation assumptions.

The PSS/Pnoise workflow also supports identifying critical trade-offs between oscillation amplitude, effective tank losses, and noise upconversion mechanisms. These results guide the following design steps, where bias implementation and tuning circuitry are progressively integrated while preserving oscillation robustness and minimizing phase-noise

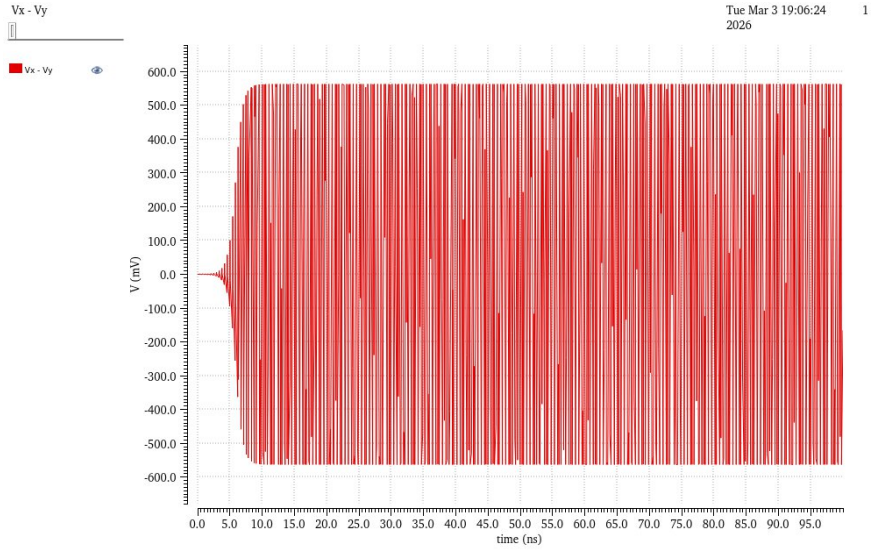

Figure 3.6. Transient start-up waveform of the differential output voltage (v_{diff}) in the nominal (`tt_lib`) case.

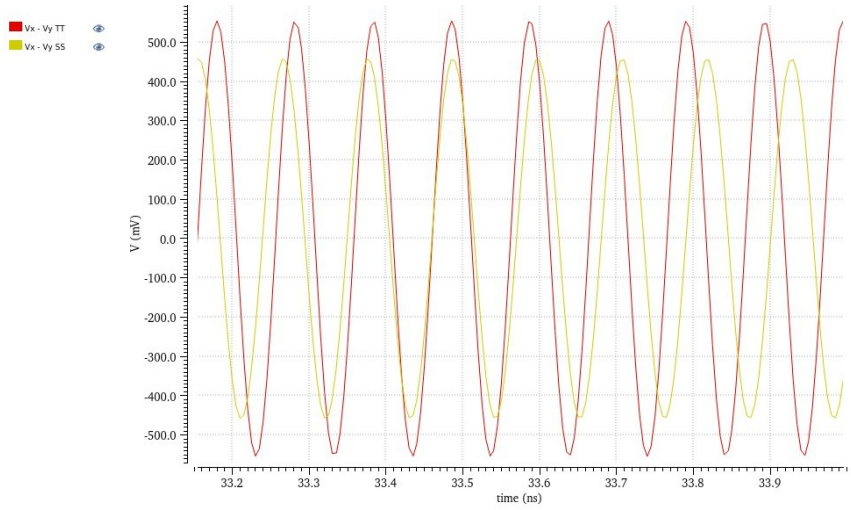

Figure 3.7. Differential output voltage $v_{\text{diff}}(t)$ in `tt_lib` and `ss_lib` corners (same bias and tuning configuration).

penalties.

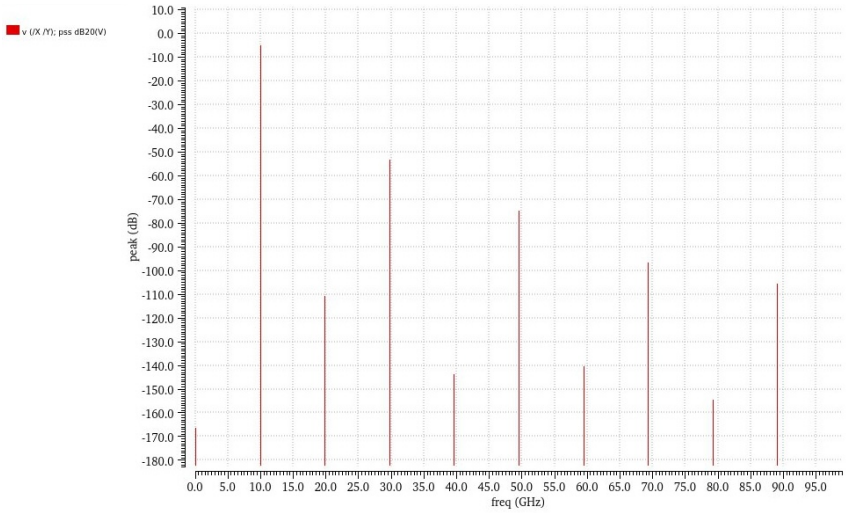

Figure 3.8. Spectrum (in dB) of the output differential voltage. The extracted fundamental frequency is $f_0 = 9.9$ GHz.

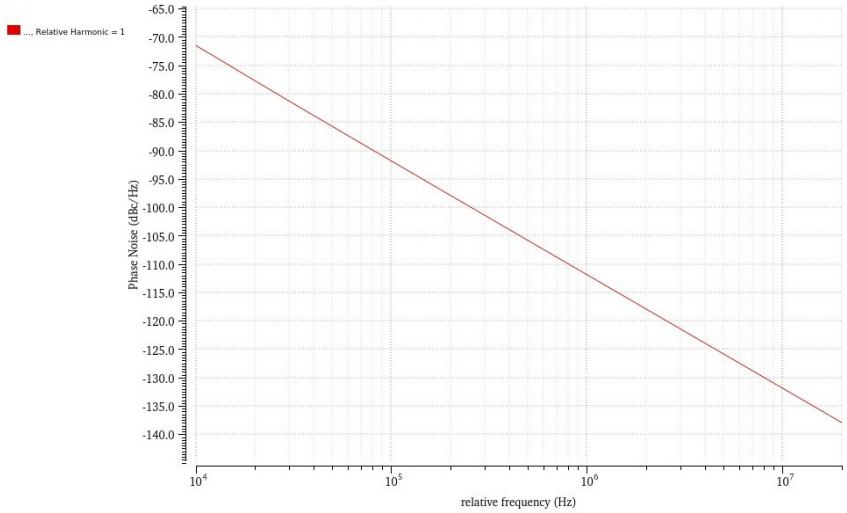

Figure 3.9. Phase noise of the first realization of the oscillator.

3.4.5 Tail current and bias mirror

The baseline oscillator is initially biased with an ideal tail current source to decouple the tank/core design from biasing non-idealities. In the next iteration, the tail current is generated on-chip by an NMOS current source driven by a current mirror. This step is necessary to obtain realistic DC operating conditions and, most importantly, to capture the additional phase-noise mechanisms introduced by a practical bias network.

As discussed by Razavi, the diode-connected device in the reference branch is a critical noise source because its noise directly modulates the gate voltage of the mirror and therefore the tail current of the VCO [Razavi, 2020]. Two distinct components must be addressed. The flicker-noise contribution is mitigated by increasing the effective device area of the diode-connected transistor (larger WL reduces $1/f$ noise). In practice, rather than using a single very wide device, a large effective area can be realized by stacking multiple identical transistors in series in the diode-connected branch, which improves the flicker-noise behavior while keeping layout and matching more controllable [Razavi, 2020]. The thermal-noise contribution, instead, is reduced by inserting a low-pass filter between the diode-connected node and the mirror gate node, so that the wideband noise from the reference branch is attenuated before it can modulate the tail current [Razavi, 2020].

Following these guidelines, the ideal tail source is replaced by a practical NMOS bias network based on a 4:1 current mirror, shown in Fig. 3.10. Both the diode-connected reference device and the tail device use a unit width of $W = 730$ nm to preserve matching, while the current ratio is implemented through device multiplicity: the reference branch employs 70 fingers, whereas the tail device uses 20 fingers. With this sizing, the mirror provides approximately 600 μA per branch (about 1.2 mA total), restoring the intended operating point adopted during the schematic-design phase after removing the ideal source.

The resulting steady-state behavior is then validated by repeating PSS and Pnoise analyses under the same observation conditions used for the baseline oscillator, so that any frequency shift, harmonic content variation, and phase-noise degradation can be directly attributed to the finite impedance and noise of the implemented tail current source and mirror network. No significant shift in the oscillation frequency or spectral content is observed, as shown in Fig. 3.11. The extracted fundamental frequency remains approximately $f_0 = 9.9$ GHz. As expected, however, the inclusion of a practical bias network leads to higher close-in phase noise. Figure 3.12 compares the phase-noise spectra of the same oscillator biased with an ideal current source and with the implemented mirror, highlighting a degradation in the near-carrier region. This result confirms that particular care must be taken in the design of the tail current source and bias mirror to minimize bias-induced noise modulation.

This observation is further supported by Fig. 3.13. Using the *Noise Summary* tool in *ADE L*, the tail transistor is identified as the dominant contributor to the overall phase noise. In the current schematic, this device corresponds to instance *M13*. This behavior is consistent with the considerations discussed by Razavi: the tail current source is not an ideal high-impedance element, and its device noise is partially translated into phase fluctuations through the nonlinear time-varying operation of the oscillator. In particular, flicker noise generated in the tail transistor modulates the bias current and the common-mode voltage, producing amplitude perturbations that are converted into phase noise (AM–PM conversion), while its finite output resistance reduces the isolation between the differential core and the supply/ground network. As a consequence, the close-in region becomes more sensitive to the tail device, which explains why the *Noise Summary* identifies the tail transistor (*M13*) as the dominant contributor in Fig. 3.13.

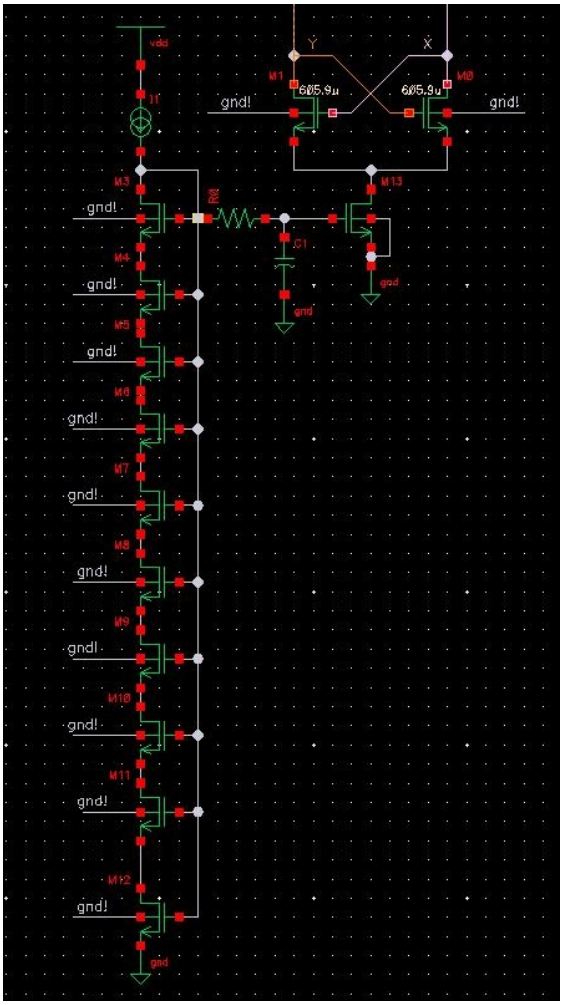

Figure 3.10. Schematic of the oscillator with bias mirror implementation.

3.4.6 Continuous tuning: inversion-mode MOSCAP sizing and verification

Continuous tuning is implemented through two inversion-mode `nmoscap` varactors connected across the differential tank nodes. The device dimensions are selected through simulation-driven sizing to obtain a monotonic tuning characteristic within a limited control-voltage window, while keeping the tuning sensitivity moderate and therefore reducing susceptibility to control-line noise. The varactor characterization testbench is shown in Fig. 3.14. Two coupling capacitors are included to ensure that the moscap terminal is tied to ground.

To evaluate the varactor behavior at the operating frequency, an SP analysis is performed by sweeping the control voltage V_{ctrl} and extracting the small-signal admittance

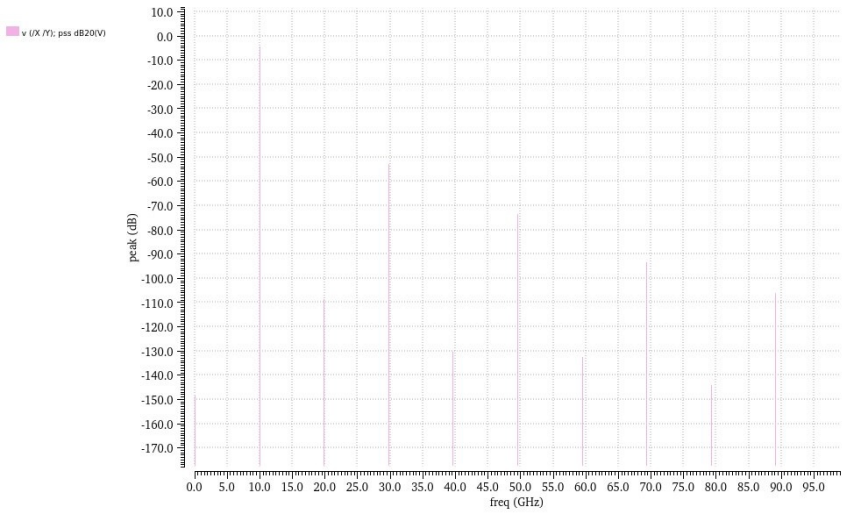

Figure 3.11. Spectrum (in dB) of the output differential voltage for the oscillator realization with the bias mirror. The extracted fundamental frequency is still $f_0 = 9.9$ GHz.

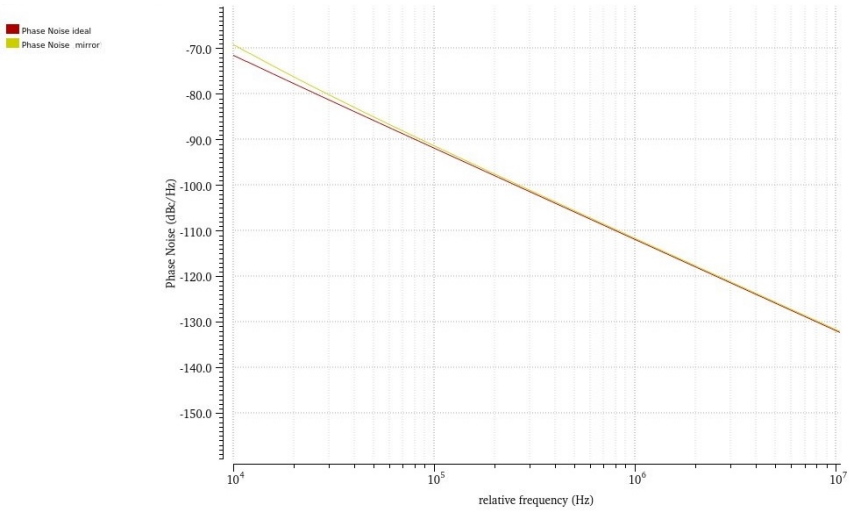

Figure 3.12. Phase noise of the oscillator realization with bias mirror (yellow) against the realization with the ideal current source (red). An increase of the $1/f$ part is observed.

$Y(j\omega)$ at a frequency close to the nominal oscillation frequency. Different varactor widths are explored using the *Parametric Analysis* tool, and the final sizing is selected based on monotonicity, tuning leverage, and RF losses.

Assuming a predominantly capacitive behavior in the selected operating region, the equivalent capacitance is computed from the imaginary part of the admittance as

$$C_{\text{var}}(V_{\text{ctrl}}) = \frac{\Im\{Y(j\omega)\}}{\omega}, \qquad \omega = 2\pi f. \tag{3.21}$$

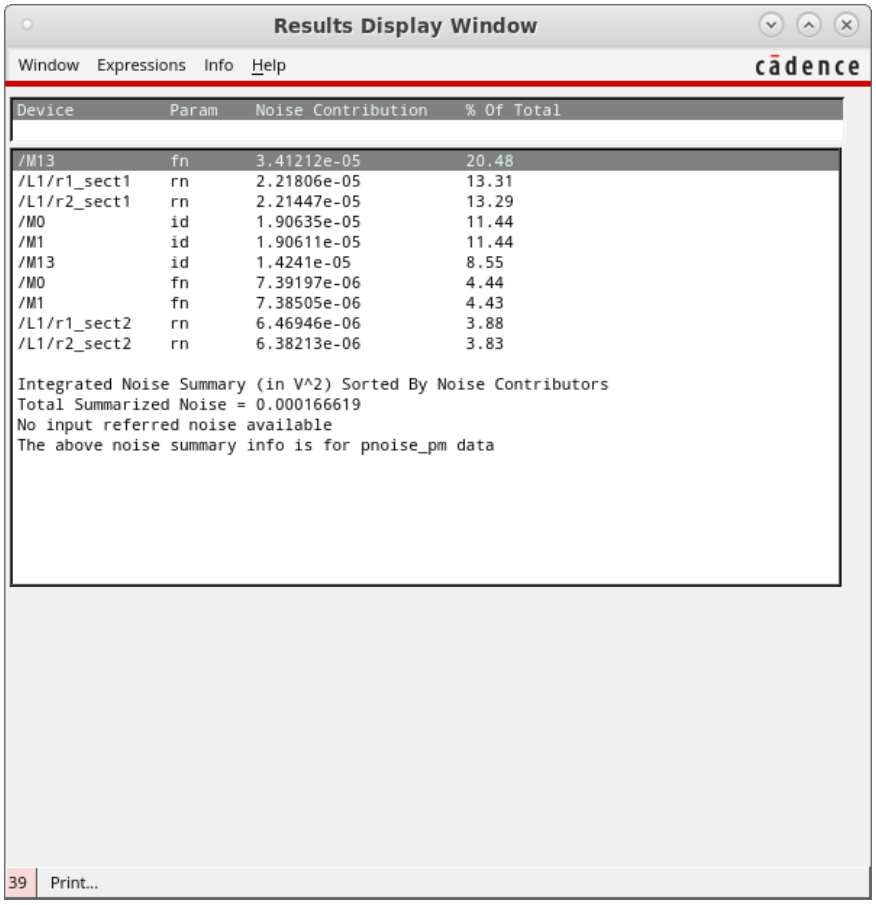

Figure 3.13. Phase noise contribution of the different instances of the circuit. *M13* is the tail transistor and largest noise source.

The final varactor sizing is $W = 3.5\ \mu$m and $L = 1\ \mu$m, with V_{ctrl} swept between 0 and 0.6 V. The capacitance characteristic of the selected varactor is shown in Fig. 3.15. The result is consistent with the preliminary design considerations of Section 3.2.2: over the control-voltage range 0–0.6 V, the varactor capacitance varies from approximately 26 fF to 33.5 fF, as highlighted in Fig. 3.16.

The impact of the varactor on the tuning gain is estimated by linearizing the resonance condition around the operating point. Since $f \simeq (2\pi\sqrt{LC_{\text{tot}}})^{-1}$, small capacitance variations produce

$$K_{\text{VCO}} = \frac{df}{dV_{\text{ctrl}}} \approx -\frac{f}{2C_{\text{tot}}}\frac{dC_{\text{var}}}{dV_{\text{ctrl}}}, \tag{3.22}$$

where C_{tot} is taken at a representative capacitor-bank setting. Because $C_{\text{var}}(V_{\text{ctrl}})$ is nonlinear, K_{VCO} is generally a function of V_{ctrl}; in this work it is obtained from the extracted $C_{\text{var}}(V_{\text{ctrl}})$ curve by evaluating a local numerical derivative in the intended operating region. The resulting $K_{\text{VCO}}(V_{\text{ctrl}})$ characteristic is reported in Fig. 3.17.

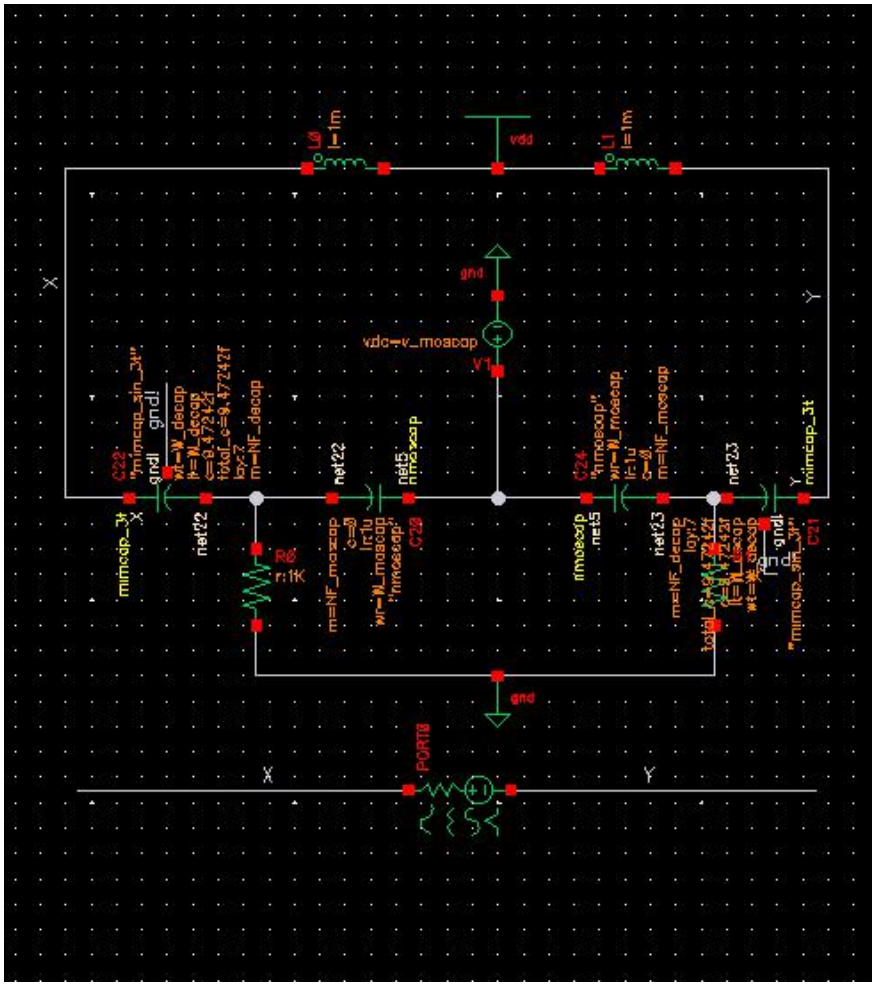

Figure 3.14. Varactor characterization testbench. Large-value inductors are used as RF chokes to provide DC biasing while isolating the RF nodes.

This tuning-gain estimate provides a reference for the expected fine control resolution and for the potential AM/PM conversion due to control-noise injection on the tuning node.

3.4.7 Discrete tuning: switch design and capacitor-bank unit verification

Coarse tuning is implemented through an 8-step switched-capacitor bank. Each step enables (or disables) a unit capacitive element by means of the NMOS switch cell shown in Fig. 3.18. The digital control bit `DIG_IN` is buffered by an inverter to generate a rail-to-rail gate drive for the switch. The inverter is kept at a standard sizing ($W_p = 1.5$ μm and $W_n = 1$ μm at minimum channel length) since it mainly behaves as a common-mode

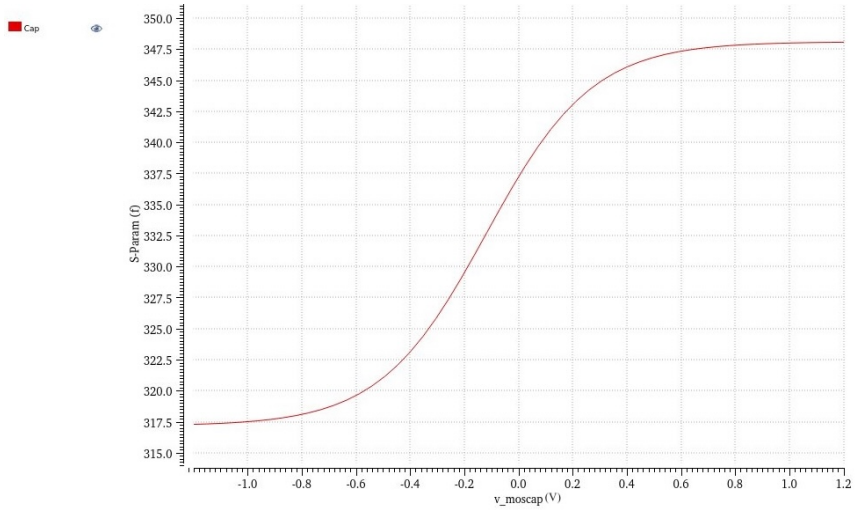

Figure 3.15. Extracted capacitance of the selected varactor as a function of the applied control voltage.

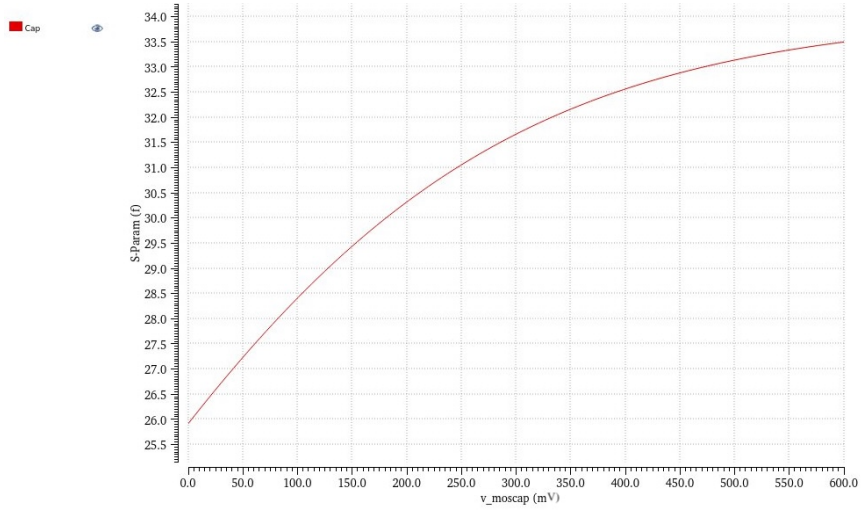

Figure 3.16. Varactor characteristic in the range 0–0.6 V.

digital driver and is not directly part of the RF signal path.

The unit capacitor associated with each step is implemented with two symmetric capacitors (one per differential side). Each capacitor has a nominal value of 20.2 fF and is connected to the tank through the switch; when integrated into the differential tank, the enabled unit contributes an effective series capacitance at the differential level.

Switch sizing is driven by a trade-off between on-state loss and parasitic loading. A larger NMOS device reduces the on-resistance R_{on} and therefore the additional damping

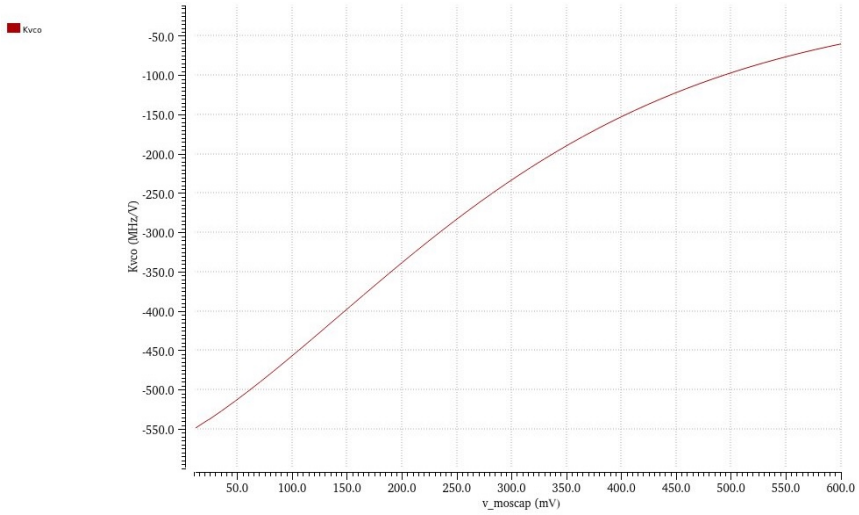

Figure 3.17. Estimated tuning gain $K_{\mathrm{VCO}}(V_{\mathrm{ctrl}})$ derived from the varactor SP characterization and (3.22).

introduced when the unit capacitor is enabled, but it also increases the off-state parasitic capacitance (junction and overlap terms) that loads the tank even when the unit is disabled. To identify a suitable operating point, R_{on} is extracted as a function of the number of fingers for a fixed width, as shown in Fig. 3.19. The curve exhibits a steep reduction at low multiplicities followed by diminishing returns beyond approximately 10-15 fingers, indicating that further enlargement would mainly increase parasitics rather than providing substantial R_{on} improvement. Based on this observation, the final NMOS switch is implemented at minimum channel length with $W_n = 1$ μm and a moderate number of fingers.

In addition, series gate resistors are introduced to control the switching edges and reduce the injection of digital switching transients into the sensitive tank nodes. The resistor connected to the switch-gate path is set to 10 kΩ, while the two auxiliary resistors in the cell (used to properly bias the switch) are set to 1 kΩ (Fig. 3.18).

This unit-level verification is a prerequisite to the full-bank integration, which is addressed in the combined tuning analysis where the effective tuning range and code-to-code overlap are verified on the complete oscillator.

3.4.8 Full VCO implementation

After validating the continuous varactor and the discrete bank at unit level, the complete VCO is assembled by integrating the LC tank, the cross-coupled core, the inversion-mode `nmoscap` tuning network, and the 8-step switched-capacitor bank. In the final topology, the two `nmoscap` devices are placed in series with each other (similar to the testbench of Fig. 3.14; the intermediate node is driven by V_{ctrl}, while the overall series combination

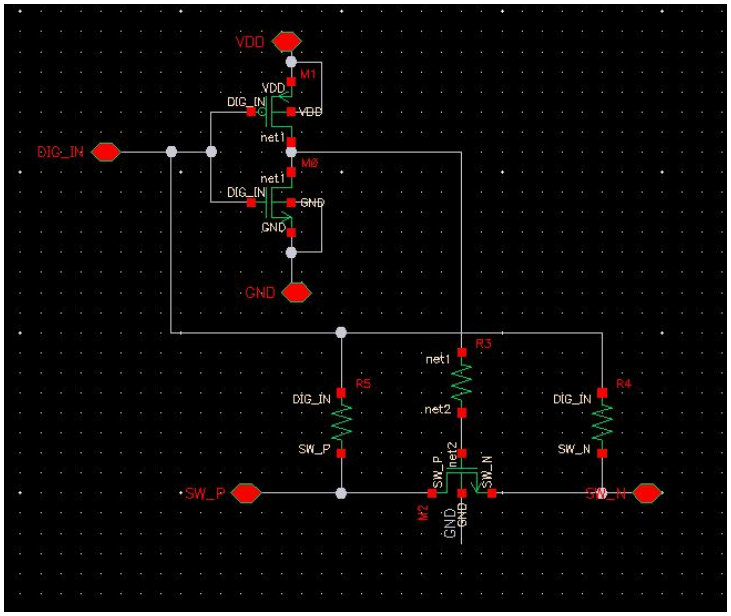

Figure 3.18. NMOS switch cell driven by a local inverter. The digital input `DIG_IN` enables/disables the capacitor unit for coarse tuning; series resistors are used to shape the switching edges.

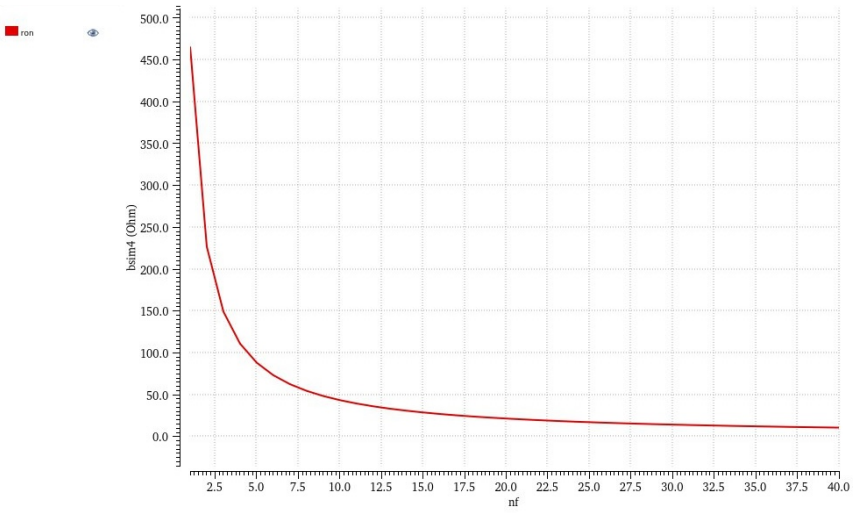

Figure 3.19. Extracted on-resistance R_{on} versus number of fingers for the NMOS switch (fixed width). The improvement saturates beyond $\sim$10–15 fingers, motivating a moderate sizing to limit parasitic capacitance.

is connected in parallel to the tank. This arrangement provides a monotonic and well-controlled fine-tuning action within the selected control window.

The tuning performance is verified in the `tt_lib` corner through PSS by sweeping V_{ctrl} from 0 to 0.6 V for each bank code. The resulting tuning map is reported in Fig. 3.20 and shows the expected downward frequency trend with increasing V_{ctrl} (increasing effective capacitance). Full coverage of the target band is achieved, spanning from $f_0 \approx 10.5$ GHz at code 7 and $V_{\mathrm{ctrl}} = 0$ V down to $f_0 \approx 9.5$ GHz at code 0 and $V_{\mathrm{ctrl}} = 0.6$ V. In addition, the curves corresponding to adjacent codes exhibit sufficient overlap, ensuring continuous coverage and avoiding tuning gaps when transitioning between discrete steps.

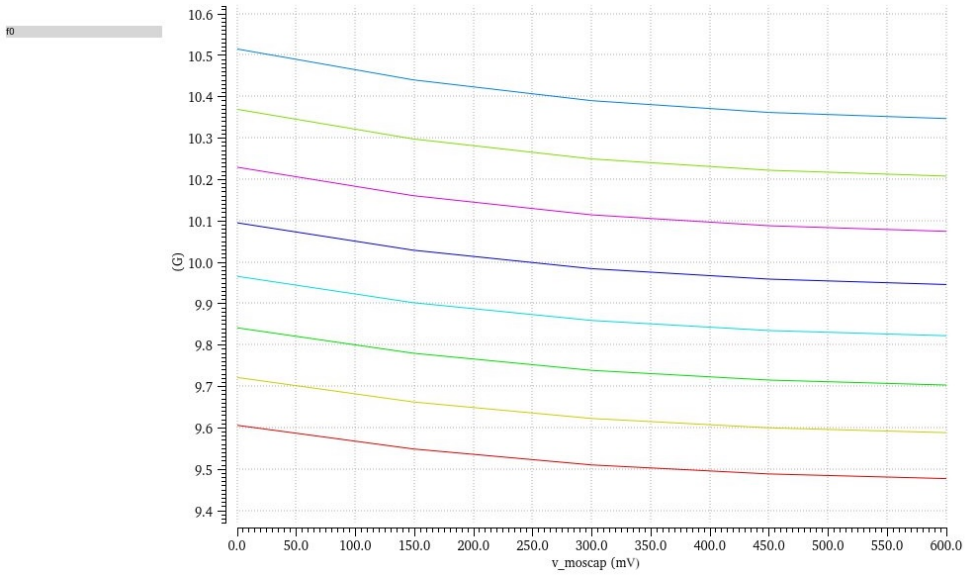

Figure 3.20. Full tuning verification (PSS, `tt_lib`): oscillation frequency versus V_{ctrl} for each of the 8 discrete bank codes. The plot demonstrates band coverage and overlap between adjacent curves.

The tuning gain is obtained consistently from the same sweep by evaluating the local slope $K_{\mathrm{VCO}}(V_{\mathrm{ctrl}}) = df/dV_{\mathrm{ctrl}}$ at a representative bank setting (Fig. 3.21). As expected from the nonlinearity of the MOSCAP characteristic, K_{VCO} varies across the control range, which motivates selecting an operating region where the tuning gain remains moderate to reduce sensitivity to control-line noise.

During full integration, minor re-optimization steps were required to compensate for the additional loading introduced by the tuning network and the realistic switch parasitics. The cross-coupled pair was resized to $W = 730$ nm with 18 fingers to restore the intended oscillation margin and steady-state swing once the complete tuning circuitry was connected. In the same phase, the tank fixed capacitance was slightly adjusted by resizing the `mimcap` to 10.4 μm $\times$ 10.4 μm, so that the nominal frequency is re-centered around the target band after including the bank and varactor contributions.

Finally, besides the `tt_lib` characterization, a transient start-up check is also performed in the `ss_lib` corner to confirm that oscillation onset is preserved under a conservative reduced-g_m condition.

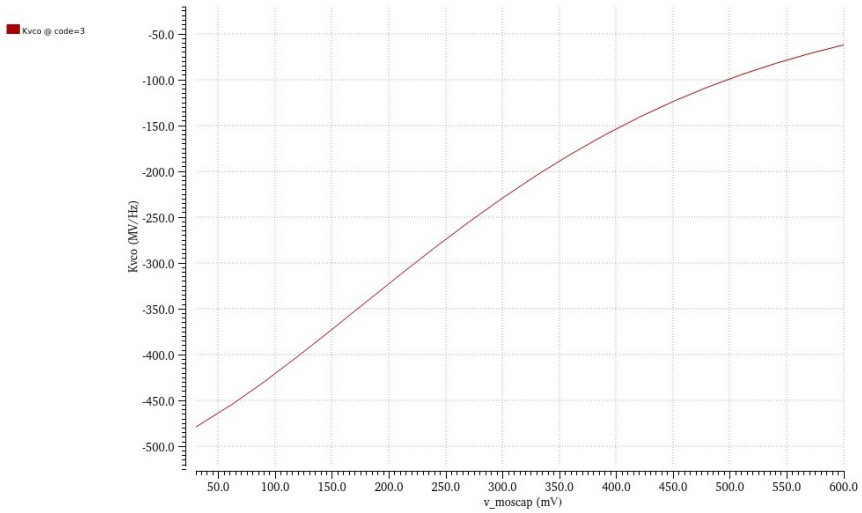

Figure 3.21. Tuning gain $K_{\mathrm{VCO}}(V_{\mathrm{ctrl}})$ extracted from the PSS frequency sweep at a representative bank code.

3.5 Layout implementation and post-layout verification

After completing the schematic-level design, the VCO has been fully implemented at layout level and verified through a standard physical sign-off flow. The implementation includes the LC tank, cross-coupled core, tail device, and the continuous tuning network, as well as the digital tuning support blocks (inverter and NMOS switch). To enable a consistent comparison between schematic-level and post-layout performance, the VCO is first realized as a dedicated layout cell, which is then instantiated in a top-level testbench for both *schematic* simulations and *PEX* simulations (extracted view).

3.5.1 Implemented blocks and VCO cell integration

Figure 3.22 shows the top-level VCO layout cell. The most RF-critical structures (tank and core) are laid out with a strict left–right symmetry around the differential nodes to reduce mismatch and preserve differential balance. The cross-coupled pair is implemented in a compact and symmetric arrangement (Fig. 3.23), and the tail device is placed to keep the connection to the source nodes short. The continuous tuning MOSCAP network is integrated in the same cell to minimize additional routing parasitics on the tank nodes, while keeping the V_{ctrl} routing separated from the digital tuning lines.

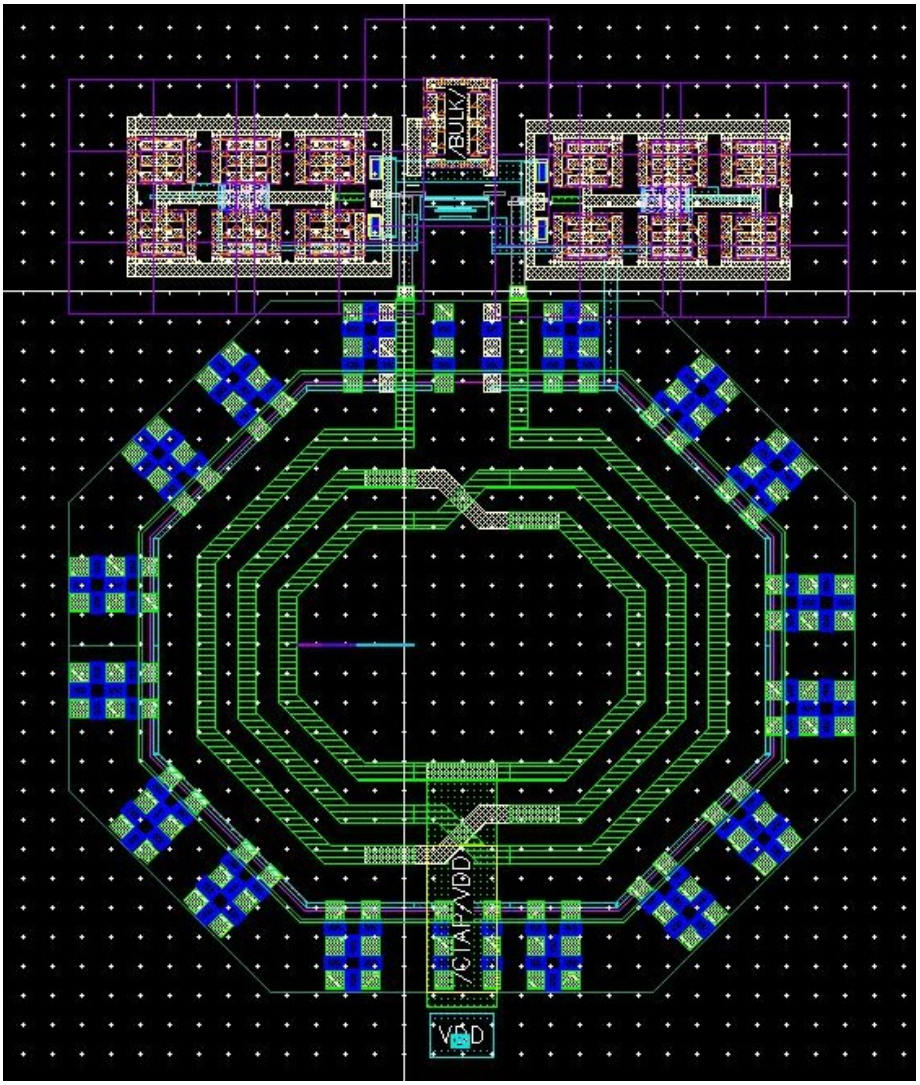

Figure 3.22. Top-level layout cell of the VCO including tank, core, tail device, and continuous tuning network. The cell is used as the reference implementation for schematic and post-PEX simulations.

3.5.2 DRC and LVS sign-off

Physical verification is performed using *Calibre*. First, the *Design Rule Check (DRC)* is run to ensure manufacturability with respect to the PDK constraints (minimum width/spacing, enclosures, via rules, and density/antenna constraints when applicable). Since the VCO layout represents a block-level cell rather than a full-chip layout, the default *ALL CHIP* option in the rule deck triggers density-related violations that are not meaningful for a standalone macro. For this reason, the *ALL CHIP* option is disabled in the rule file so that density checks do not report false positives during block-level verification. After

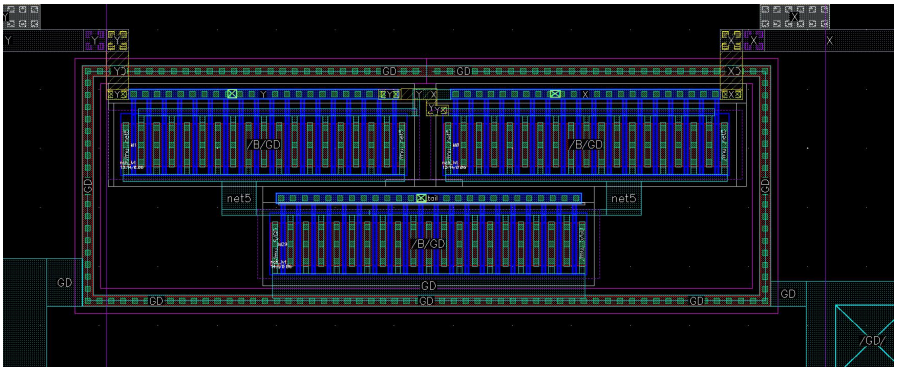

Figure 3.23. Detail of the cross-coupled core layout. A symmetric placement is adopted to minimize differential mismatch and unbalanced parasitic loading on the tank nodes.

this adjustment, the VCO cell passes DRC without violations (Fig. 3.24).

After DRC, *Layout Versus Schematic (LVS)* is executed to confirm electrical equivalence between layout and schematic. The final LVS report confirms a correct match at top level (Fig. 3.25), excluding connectivity errors as a root cause of any performance discrepancy between schematic and extracted simulations.

3.5.3 PEX and post-layout performance comparison

Once DRC and LVS are clean, *Parasitic Extraction (PEX)* is performed to generate the extracted view of the VCO cell including interconnect parasitics. The extracted netlist is then simulated using the same methodology adopted at schematic level (PSS for f_0 extraction and tuning sweeps), enabling a direct post-layout assessment of tuning coverage.

Figure 3.26 reports the post-PEX tuning map of the complete VCO, obtained by sweeping V_{ctrl} from 0 to 0.6 V for each of the 8 discrete tuning codes. The map confirms that the tuning behavior is preserved under realistic parasitic loading and provides the post-layout reference for band coverage and overlap between adjacent codes.

To quantify the frequency shift introduced by layout parasitics, Fig. 3.27 shows an overlay of the oscillation frequency as a function of V_{ctrl} for schematic and post-PEX simulations (representative tuning condition). The separation between curves directly reflects the additional parasitic capacitance and routing effects introduced by the physical implementation.

Overall, the post-layout flow confirms that the integrated layout (tank, core, tail, continuous tuning, and switch/inverter blocks) reproduces the expected tuning behavior observed at schematic level, while capturing the frequency shift introduced by realistic interconnect parasitics.

3.5.4 Post-layout gate-fidelity estimation

Figure 3.28 compares the single-sideband phase-noise spectra obtained from schematic-level and post-PEX simulations for the same oscillation condition (3 open switches and

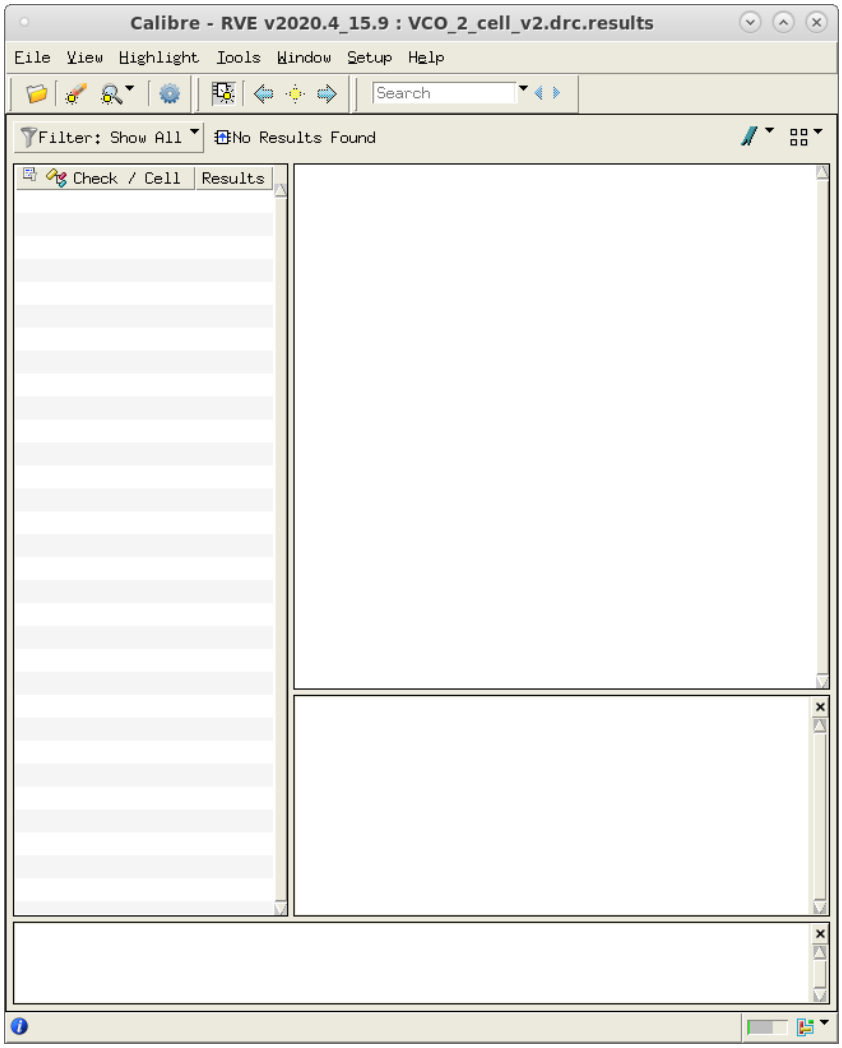

Figure 3.24. Calibre RVE report showing a DRC-clean result for the VCO layout cell (block-level verification with *ALL CHIP* disabled to avoid non-meaningful density errors).

no voltage applied to the moscaps). The two curves are close in magnitude, but parasitic extraction modifies the spectral distribution of phase noise across offset frequency rather than producing a uniform shift.

To complement the post-layout RF analysis, the phase-noise spectrum obtained from the extracted VCO is also translated into a qubit-level performance metric. As discussed in Section 3.1.6, the filter-function formalism provides the link between oscillator phase noise and the expected degradation of single-qubit gate fidelity. In particular, the average fidelity of a primitive X_π gate can be expressed through Eqs. (3.4) and (3.5), where the decay parameter $\chi(\tau)$ is determined by the overlap between the dephasing-noise spectrum and the control-dependent filter function.

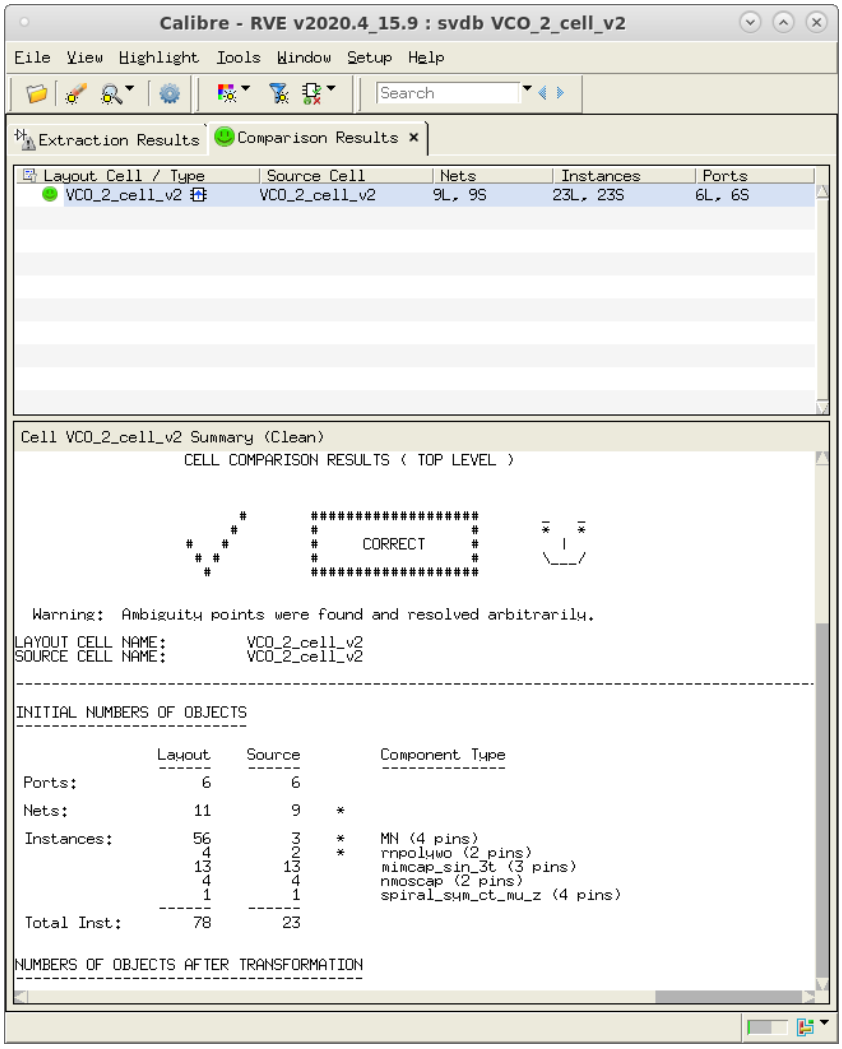

Figure 3.25. Calibre LVS comparison report for the VCO cell, showing a correct match between layout and schematic.

In the present work, the post-PEX single-sideband phase-noise spectrum $L(f)$ is exported from *Cadence* and post-processed in *MATLAB* using the primitive X_π filter function of Green *et al.*. The numerical implementation therefore provides a practical evaluation of the same theoretical framework introduced in Section 3.1.6, directly starting from the simulated phase-noise data of the extracted oscillator. The MATLAB routines used for the primitive X_π-gate fidelity estimation and for the comparison between schematic-level and post-PEX results are reported in Appendix A.

This step is especially important because the extracted view includes additional parasitic loading introduced by routing, interconnect capacitances, and layout asymmetries, which are not captured at schematic level. As a consequence, the post-layout phase-noise

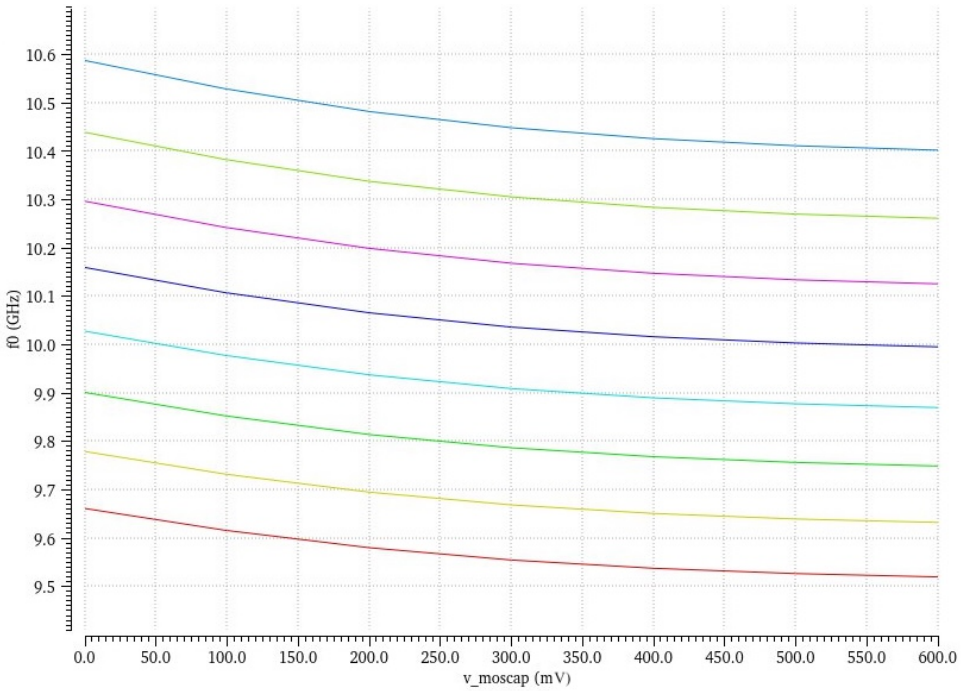

Figure 3.26. Post-PEX tuning verification (PSS): oscillation frequency versus V_{ctrl} for each of the 8 discrete bank codes, simulated on the extracted VCO cell.

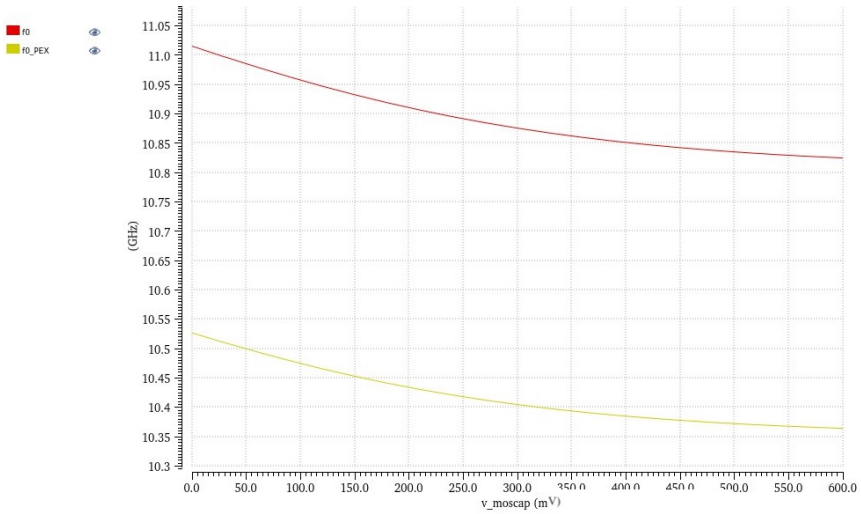

Figure 3.27. Overlay comparison of tuning curves (f_0 vs. V_{ctrl}) obtained from schematic simulations and post-PEX simulations of the extracted VCO cell. The difference between curves quantifies the frequency shift due to layout parasitics.

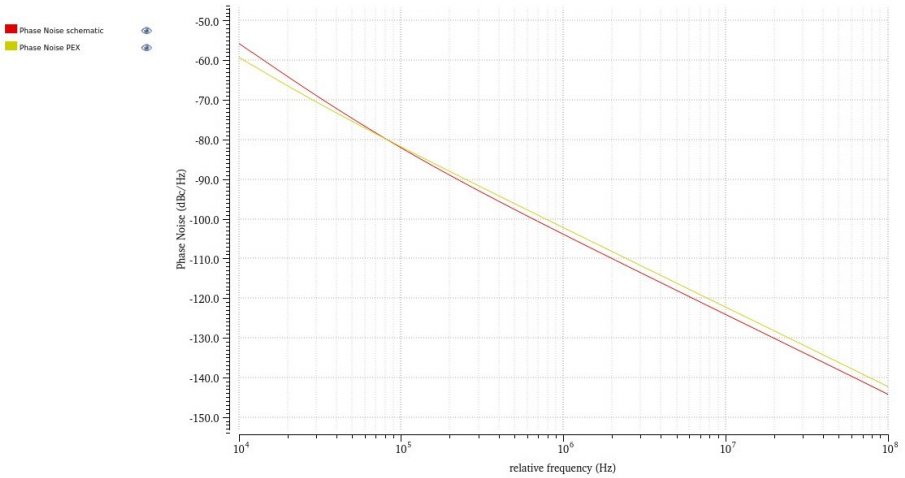

Figure 3.28. Comparison of the single-sideband phase-noise spectra obtained from schematic-level and post-PEX simulations under the same oscillation condition. Parasitic extraction does not introduce a uniform degradation over the whole offset-frequency range: in the present case, the extracted oscillator exhibits slightly different behavior at low and high offset frequencies, which later translates into a non-uniform effect on gate fidelity.

spectrum provides a more realistic estimate of the oscillator contribution to gate infidelity in the final integrated implementation.

Figure 3.29 reports the overlap integrand obtained from the post-PEX phase-noise spectrum for a representative X_π-pulse duration. This quantity directly reflects the spectral regions that dominate the accumulated error χ: only the frequency components for which the oscillator phase noise overlaps significantly with the control filter function contribute appreciably to gate infidelity. In the considered case, the numerical integration of the post-PEX spectrum yields $\chi = 0.0534$, corresponding to an estimated average gate fidelity of $F_{\text{av}} = 0.97398$.

The same procedure is then applied both to schematic-level and post-PEX phase-noise spectra over a range of π-pulse durations, in order to compare the resulting fidelity curves. The results are reported in Figure 3.30. This comparison shows that, over a wide range of π-pulse durations, the post-PEX VCO yields a lower gate fidelity than the schematic-level implementation. However, for pulse durations longer than approximately 4 μs, the two curves cross. This behavior is consistent with the filter-function interpretation: as the pulse duration increases, the fidelity becomes progressively more sensitive to phase-noise components closer to the carrier, so that the relative spectral redistribution introduced by parasitic extraction can change the ranking of the two cases.

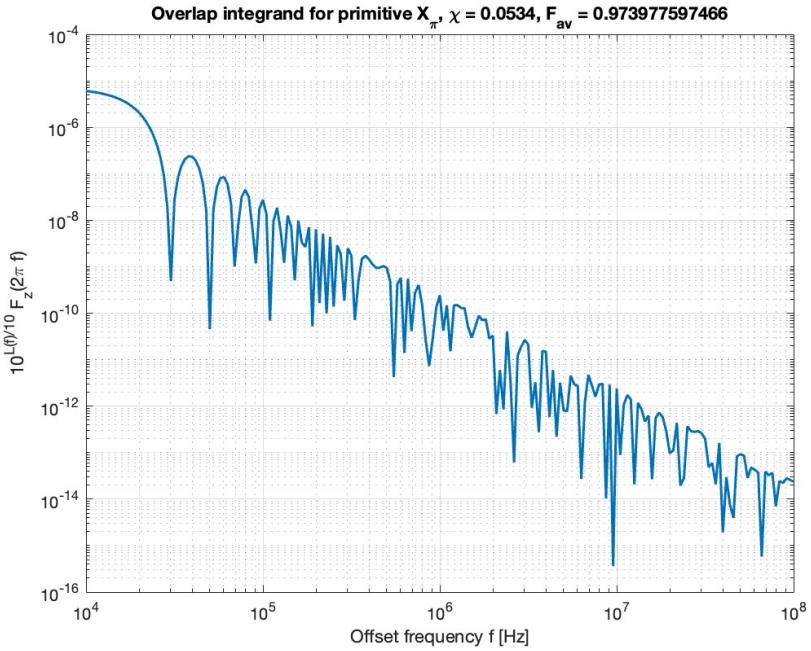

Figure 3.29. Overlap integrand $10^{L(f)/10}F_z(2\pi f)$ obtained from the post-PEX phase-noise spectrum for a representative primitive X_π pulse. The area under this curve determines the decay parameter χ, and hence the corresponding average gate fidelity.

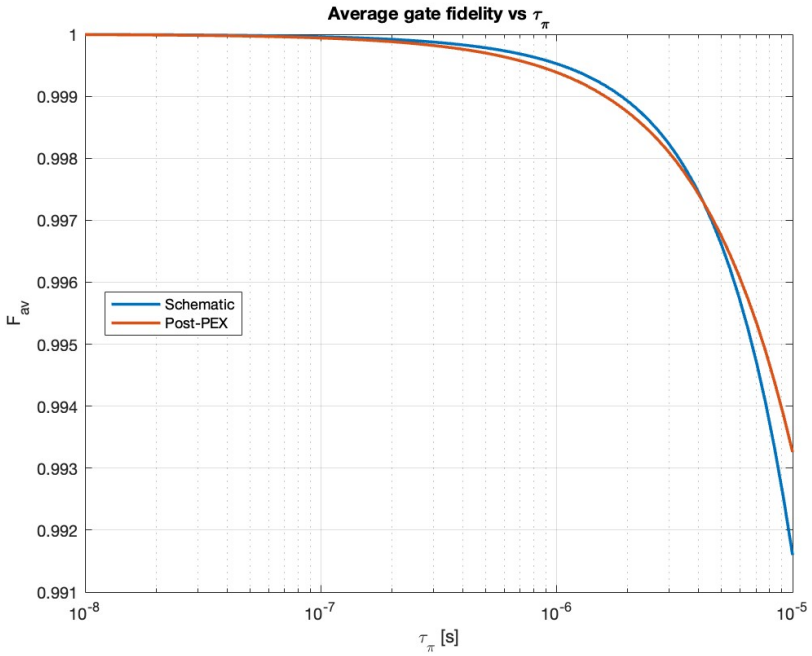

Figure 3.30. Average primitive X_π-gate fidelity as a function of the pulse duration τ_π, evaluated from schematic-level and post-PEX phase-noise spectra through the filter-function formalism. For short pulse durations, the post-PEX fidelity is slightly lower, indicating a stronger contribution from offset-frequency regions where the extracted oscillator is noisier. For longer pulse durations, the spectral weighting shifts toward frequencies closer to the carrier, and the relative ranking of the two curves changes accordingly.

3.5.5 Post-layout performance summary and bias-current trade-off

The post-layout results discussed in the previous sections can be summarized through a compact set of key performance metrics. In particular, the extracted VCO exhibits a clear trade-off between bias current, output swing, phase noise, and qubit-level fidelity. Increasing the bias current improves the oscillation amplitude and reduces the phase noise, at the cost of higher power dissipation. This behavior is illustrated in Fig. 3.31 and summarized in Table 3.2.

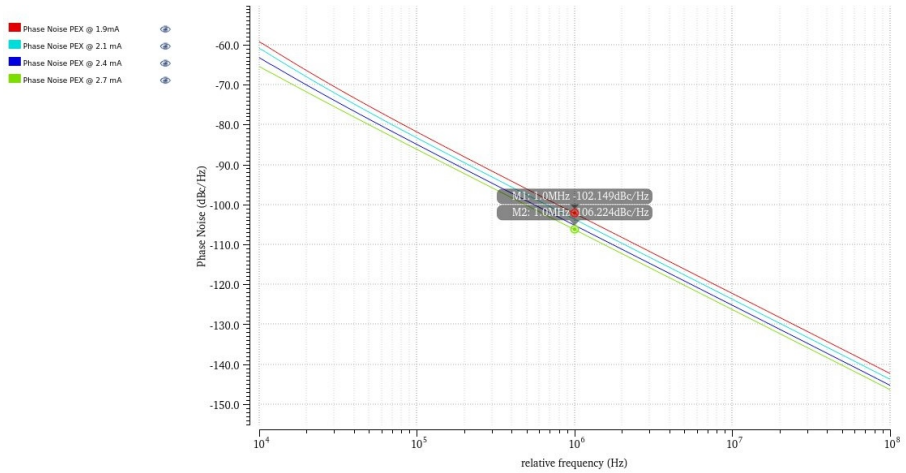

Figure 3.31. Post-PEX phase-noise spectra obtained at different bias-current values. Increasing the bias current improves the phase-noise performance over the whole offset-frequency range considered here, including the region around 1 MHz relevant for qubit-control specifications.

Post-layout simulations also showed that sustained oscillation requires at least approximately 1.5 mA, confirming that parasitic loading and tank degradation significantly tighten the bias-current requirement with respect to the initial schematic-level implementation.

The comparison highlights that the initial oscillator-level phase-noise target can be approximately recovered at post-layout level by increasing the bias current up to 2.7 mA. However, this improvement is obtained at the expense of higher power dissipation and a significantly larger output swing. For this reason, the 1.9 mA operating point is retained in this work as a representative trade-off between oscillation robustness, acceptable output swing, and realistic post-layout qubit-level performance.

Overall, the results of this chapter show that the proposed LC-VCO can achieve stable operation, full tuning-range coverage, and physically meaningful post-layout fidelity values within the considered design constraints. At the same time, they also highlight that the interpretation of low-temperature operation cannot rely only on the standard room-temperature PDK models. For this reason, the next chapter addresses the cryogenic re-centering of the active-device BSIM4 model, in order to support the low-temperature

Table 3.2. Summary of representative post-layout operating points of the proposed LC-VCO. The reported average gate fidelities are evaluated from the simulated phase-noise spectra using the filter-function formalism with $\tau_\pi = 50\ \mu$s.

Metric	**Post-PEX @ 1.9 mA**	**Post-PEX @ 2.7 mA**
Supply voltage	1.2 V	1.2 V
Power consumption	2.28 mW	3.24 mW
Differential peak amplitude	≈ 310 mV	≈ 520 mV
Phase noise @ 1 MHz	≈ -102.1 dBc/Hz	≈ -106.0 dBc/Hz
$\chi(\tau_\pi = 50\ \mu\text{s})$	0.05345	0.01470
$F_{\text{av}}(\tau_\pi = 50\ \mu\text{s})$	0.97398	0.99270

interpretation of the oscillator behavior discussed here.

Chapter 4

Cryogenic Modeling and BSIM4 Re-Centering

Although the VCO design methodology was introduced in Chapter 3, the cryogenic behavior of the active devices deserves a separate discussion because no qualified low-temperature compact model was available in the adopted PDK. The oscillator design flow was first developed at room temperature, while the analysis reported in this chapter was carried out afterwards to verify whether the active-device behavior remained physically meaningful under cryogenic conditions. This chapter therefore focuses on the BSIM4 re-centering procedure adopted to support the cryogenic interpretation of the oscillator results. The purpose of this activity was to obtain a design-oriented transistor model sufficiently accurate to support circuit-level simulations and verifications under cryogenic conditions while remaining fully embedded in the standard industrial CAD flow.

4.1 Cryogenic behavior of nanoscale MOSFETs

Cryogenic operation modifies MOSFET behavior through a combination of electrostatic, transport, and parasitic effects. Among the most relevant consequences are the increase in threshold voltage, the enhancement of carrier mobility, the reduction of leakage currents, the improvement of subthreshold swing, and the modification of short-channel and output-related phenomena [Incandela et al., 2018, Liu et al., 2021].

4.1.1 Threshold-voltage shift and electrostatics

A key effect of cryogenic operation is the increase in threshold voltage. This condition depends on the MOS electrostatics through the Fermi potential and the surface band bending. In first approximation, strong inversion is reached when the surface potential satisfies

$$\psi_s \approx 2\phi_F,$$

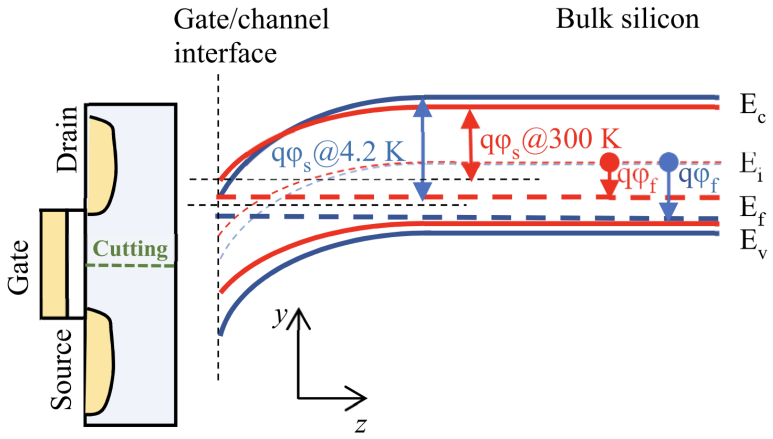

Figure 4.1. Energy-band schematic of an nMOS transistor at the onset of inversion, illustrating the stronger band bending required at cryogenic temperature and the resulting increase in threshold voltage. Reproduced from [Liu et al., 2021].

where

$$\phi_F = \frac{kT}{q} \ln\left(\frac{N_A}{n_i}\right).$$

As the temperature decreases, the intrinsic carrier concentration n_i drops dramatically. Although the thermal voltage kT/q also decreases, the reduction of n_i is so strong that the overall electrostatic condition for inversion shifts toward a larger required gate bias. In physical terms, fewer thermally generated carriers are available, and stronger band bending is needed to form a stable inversion layer at the silicon surface. As a result, the threshold voltage increases at cryogenic temperature [Beckers et al., 2018, Gutiérrez-D et al., 2001, Liu et al., 2021].

In addition, incomplete dopant ionization modifies the electrostatics of the substrate. At very low temperature, part of the dopants remains frozen out in the bulk, while near the surface they can become rapidly ionized due to field-assisted ionization when the gate voltage increases. Therefore, cryogenic operation cannot be interpreted simply as a condition in which fewer carriers must be removed. Rather, the relevant effect is that the channel-formation condition changes, and a higher gate voltage is typically required to reach strong inversion. Experimental characterization on nanometric bulk CMOS confirms this behavior and reports a clear threshold-voltage increase at 4.2 K [Incandela et al., 2018, Liu et al., 2021, Tada et al., 2024].

4.1.2 Mobility enhancement and transconductance

The increase in threshold voltage is partially compensated for by the enhancement of carrier mobility. At cryogenic temperature, phonon scattering is strongly reduced. This generally leads to higher transconductance and larger ON-state current, especially in moderate and strong inversion. However, the mobility improvement is not uniform across

all geometries. As discussed by Liu *et al.*, long-channel devices benefit more strongly from cooling, whereas in shorter and narrower devices the improvement is partially limited by impurity scattering, interface-related effects, and velocity-saturation-related mechanisms [Liu et al., 2021]. This is particularly relevant in VCO design, where the g_m parameter plays a central role in oscillator start-up.

4.1.3 Leakage, subthreshold, and output-related effects

Another beneficial consequence of cooling is the strong reduction of OFF-state leakage current and the improvement of subthreshold swing. The lower thermal energy suppresses subthreshold leakage, while the larger silicon bandgap at cryogenic temperature also helps reduce band-to-band tunneling-related leakage mechanisms, such as gate-induced drain leakage. In the same 55 nm technology, Liu *et al.* report several orders of magnitude reduction in I_{OFF} and a significantly steeper subthreshold swing at 4.2 K [Liu et al., 2021]. Similar trends are also widely reported in the cryogenic CMOS literature [Incandela et al., 2018, Beckers et al., 2018].

Cryogenic operation also affects short-channel and output-related phenomena, including DIBL, series-resistance contribution from LDD implantation, velocity saturation, and output conductance. These effects do not always evolve in the same direction across different technologies. In the general cryogenic MOS literature, parasitic series resistance and incomplete ionization in the bulk can become more relevant at low temperature, while mobility enhancement tends to increase current drive and transconductance [Gutiérrez-D et al., 2001, Beckers et al., 2018]. In the 55 nm bulk CMOS study by Liu *et al.*, the DIBL effect was observed to worsen at 4.2 K because freeze-out in the lower-doped substrate increases the depletion width, effectively worsening drain-side electrostatic control [Liu et al., 2021]. By contrast, other works on nanometric bulk CMOS report that some short-channel effects exhibit only a weak temperature dependence, highlighting that these trends remain technology-dependent [Gatti and Tavernier, 2024].

Finally, non-ideal cryogenic effects such as the kink effect must also be considered. The kink effect at cryogenic temperature is a bulk-potential modulation phenomenon. At large V_{DS}, impact ionization generates a bulk current near the drain. Because the substrate resistance increases at low temperature due to carrier freeze-out, this current raises the local bulk potential. The resulting reduction in threshold voltage enhances the drain current abruptly, producing the kink observed in the output characteristics. In older or thicker-oxide technologies, kink-related anomalies may become significant due to larger V_{DS} and lower vertical fields and require explicit corrections in the compact model. However, in modern thin-oxide nanoscale bulk CMOS, such effects can be strongly mitigated or even absent depending on device structure and biasing conditions. This is consistent with the present work, where no explicit kink correction was required for the selected $0.06\,\mu$m/$10\,\mu$m LVT NMOS used for model re-centering [Incandela et al., 2018, Gatti and Tavernier, 2024].

4.2 BSIM4 re-centering methodology

Since no foundry-qualified cryogenic compact model was available for the adopted TSMC 65 nm CMOS technology, a dedicated transistor-level re-centering procedure was carried out before validating the final LC-VCO under cryogenic operating conditions. The objective of this step was to adapt the room-temperature BSIM4 model card provided by the foundry so that it could reproduce, with sufficient design accuracy, the behavior of a representative cryogenic device reported in the literature. This strategy is consistent with the modeling approach proposed for modern thin-oxide 65 nm bulk CMOS technologies, where cryogenic operation can be described effectively through parameter re-centering while preserving the original BSIM4 equations and simulator efficiency [Gatti and Tavernier, 2024, Incandela et al., 2018, Tada et al., 2024].

4.2.1 Reference device and fitting targets

The fitting activity focused on an NMOS LVT transistor with geometry $L = 0.06\,\mu$m and $W = 10\,\mu$m, corresponding to the wide, short-channel device reported in the reference paper with a geometry close to the one adopted in the oscillator design. The target curves were extracted from the published plots using WebPlotDigitizer and included:

- the transfer characteristic $I_D - V_{GS}$ at $V_{DS} = 50\,$mV;
- the corresponding transconductance $g_m - V_{GS}$;
- the output characteristics $I_D - V_{DS}$ for $V_{GS} = 0.75$, 0.975, and 1.2 V;
- the output conductance $g_{ds} - V_{DS}$ for $V_{GS} = 0.75$ and 0.975 V.

The bulk terminal was kept at $V_{BS} = 0$.

4.2.2 Re-centering strategy

Instead of replacing the foundry model with a custom cryogenic implementation, the adopted methodology consisted in re-centering the original BSIM4 card (bin 5 of the `nch_lvt` model) through parameter overrides. This choice preserves the native model equations, avoids the need for external Verilog-A implementations, and remains compatible with the standard Cadence/Spectre flow. Such an approach has been shown to be practical and robust for cryogenic compact modeling of 65 nm CMOS technologies, especially when the goal is design-oriented circuit simulation rather than full physical re-extraction [Gatti and Tavernier, 2024].

The procedure followed in this work was inspired by the cryogenic modeling flow reported in the literature. First, the simulator temperature was set to $-220\,^\circ$C in ADE as a practical workaround to emulate cryogenic devices, since the properties of Si nanoscale CMOS devices do not change significantly for $T < 50$ K [Gatti and Tavernier, 2024]. In parallel, the temperature-related dependencies of the original model were neutralized by modifying the model parameter `tnom` to $-220\,^\circ$C, so that the compact model would not extrapolate room-temperature temperature laws into a regime for which they were not

intended. A similar strategy is commonly adopted in cryogenic compact modeling, where the simulator temperature is forced to a low effective value and the nominal temperature dependencies of the original compact model are removed or neutralized before re-fitting the relevant parameters [Incandela et al., 2018].

After this preparation step, the cryogenic behavior was reproduced by introducing differential parameters with respect to the original LVT NMOS model. Most of these parameters were already available in the override flow and could therefore be changed externally. However, some additional delta-parameters, namely `da0n_lvt`, `dagsn_lvt`, `dpclmn_lvt`, and `dminvn_lvt`, were not exposed in the original file structure. For this reason, a local copy of the model files (named `crn65lp_2d5_lk_v1d7.scs` and `toplevel.scs`, along with the override file) was created and modified: the appropriate `tnom` entry was set to $-220\,^{\circ}\mathrm{C}$, and the required additional fitting knobs were exposed. The modified files were then included in the simulation library so that these extra re-centering parameters could also be injected into the final model card.

4.2.3 Manual fitting procedure

The fitting was performed manually and iteratively, starting from physically reasonable cryogenic trends rather than from a blind numerical optimization. This choice was motivated by the strong correlation among several BSIM4 parameters and by the need to preserve a meaningful physical interpretation of the final model. It was also motivated by the limited number of digitized target curves and by the need to preserve a physically interpretable parameter set for subsequent circuit-level use.

The procedure was organized in four main steps.

1) Threshold alignment. The first step targeted the horizontal shift of the transfer curve. The expected increase in threshold voltage at cryogenic temperature was introduced primarily through `dvthn_lvt` and `dcitn_lvt`. The initial value was chosen to reflect the threshold increase reported in the reference paper, which is on the order of a few tens of percent for LVT devices. The quality of this first alignment was evaluated on the I_D–V_{GS} characteristic at $V_{DS} = 50\,\mathrm{mV}$.

2) Current drive and transconductance adjustment. Once the threshold position was approximately matched, the low-field mobility and gate-field mobility degradation were tuned to reproduce both the slope of the transfer curve and the $g_m - V_{GS}$ peak and plateau. In practice, `du0n_lvt`, `duan_lvt`, and `dubn_lvt` were used to control the effective current drive in moderate and strong inversion. Additional refinement was obtained through `dminvn_lvt`, which helped shape the transition between weak/moderate inversion and the high-current region.

3) Output characteristics refinement. After the transfer behavior was aligned, the output curves $I_D - V_{DS}$ at different V_{GS} were used to refine the saturation region and the spacing among curves. In this phase, `drdswn_lvt` and `dvsatn_lvt` played a major role, since they affect the high-field current behavior and the effective compression of the

output curves at large V_{DS}. The pair `da0n_lvt`/`dagsn_lvt` was especially useful to control how the current scales across different gate biases, thereby improving the simultaneous agreement of the $V_{GS} = 0.75$, 0.975, and 1.2 V characteristics.

4) Output conductance validation. Finally, the fitting was checked against g_{ds}–V_{DS}, which provides a more sensitive indicator of saturation quality and residual mismatch in the output conductance. This last step was useful to verify that the model was not only reproducing the current level but also the correct drain-bias sensitivity in the saturation region.

The fitting quality was therefore assessed mainly through visual agreement across multiple complementary observables rather than through a single scalar error metric. Although no formal RMS optimization was performed, the combined use of I_D, g_m, and g_{ds} significantly reduced the risk of overfitting one curve while degrading the others.

4.2.4 Override implementation

The final cryogenic override file used in Spectre is reported below for completeness:

```
simulator lang=spectre
parameters dvthn_lvt     = +0.061
parameters du0n_lvt      = +0.9
parameters duan_lvt      = +35e-9
parameters dubn_lvt      = +265e-19
parameters drdswn_lvt    = -61
parameters dcitn_lvt     = +0.045
parameters dnfactorn_lvt = -0.2
parameters deta0n_lvt    = -0.07
parameters dvsatn_lvt    = +14000

//added in the main model file
parameters da0n_lvt      = -7
parameters dagsn_lvt     = -4
parameters dpclmn_lvt    = 0
parameters dminvn_lvt    = +0.4
```

Listing 4.1. Spectre override file used for cryogenic BSIM4 re-centering.

Table 4.1 summarizes the role of each fitted delta-parameter in the final re-centered model.

It should be noted that the entries `da0n_lvt`, `dagsn_lvt`, `dpclmn_lvt`, and `dminvn_lvt` were not originally exposed in the standard override structure. Therefore, they were explicitly added to a local copy of the model file and then included in the model library used during simulation. This modification allowed the fitting to remain fully embedded in the original BSIM4 flow while extending the set of accessible re-centering parameters.

Table 4.1. Cryogenic BSIM4 delta-parameters used for the LVT NMOS re-centering.

Parameter	Delta value	Main effect on fitting
Threshold and subthreshold		
dvthn_lvt	+0.061	Threshold increase; horizontal shift of I_D-V_{GS}
dcitn_lvt	+0.045	Interface-trap correction; subthreshold and moderate-inversion shaping
dnfactorn_lvt	−0.2	Ideality-factor correction; subthreshold slope refinement
dminvn_lvt	+0.4	Moderate-inversion transition shaping
Mobility and strong inversion		
du0n_lvt	+0.9	Low-field mobility increase; overall current drive and g_m
duan_lvt	$+35 \times 10^{-9}$	High-V_{GS} mobility-degradation shaping
dubn_lvt	$+265 \times 10^{-19}$	Strong-inversion mobility shaping
dvsatn_lvt	+14000	Velocity-saturation adjustment; high-V_{DS} current behavior
Output-region refinement		
drdswn_lvt	−61	Bias-dependent source/drain resistance; output-current compression
deta0n_lvt	−0.07	Drain-bias and short-channel sensitivity refinement
dpclmn_lvt	0	No extra channel-length modulation correction retained
Cross-bias scaling		
da0n_lvt	−7	Bulk-charge shaping; relative spacing of I_D-V_{DS} curves
dagsn_lvt	−4	Gate dependence of bulk-charge effect across different V_{GS}

4.3 Validation of the re-centered model

4.3.1 Fitting results

The final fitting results are shown in Figs. 4.2–4.5. Overall, the re-centered model reproduces the main cryogenic trends of the target device with satisfactory agreement for circuit-design purposes.

Figure 4.2 compares the transfer characteristic at $V_{DS} = 50\,\text{mV}$. The threshold shift is correctly captured and the simulated current follows the extracted curve over the full inversion range, with only a limited residual mismatch in the transition region. Figure 4.3 shows that the transconductance enhancement expected at cryogenic temperature is also reproduced well, including the rapid rise of g_m and the subsequent plateau in strong inversion. Particular attention was devoted to the fitting of the g_m characteristic, since the effective transconductance of the cross-coupled core directly determines the negative resistance available to compensate the tank losses. For this reason, an accurate reproduction of g_m was considered especially important for the subsequent oscillator-like validation.

The output characteristics reported in Fig. 4.4 confirm that the model is able to reproduce the drain current over a wide range of gate biases. The agreement is particularly good at $V_{GS} = 1.2\,\text{V}$, while a slightly larger mismatch remains for the intermediate and lower gate biases, especially in the transition from the linear to the saturation region. This behavior is consistent with the fact that a limited number of BSIM4 parameters had to capture several coupled cryogenic effects simultaneously, including threshold shift, mobility enhancement, and bias-dependent output shaping.

Finally, Fig. 4.5 provides an additional validation of the saturation region. The simulated g_{ds} tracks the target trend reasonably well and confirms that the re-centered model

captures not only the current level but also the evolution of output conductance with drain bias. This aspect is important for RF oscillator design, since an inaccurate g_{ds} would directly affect the negative-resistance condition, voltage swing, and phase-noise prediction.

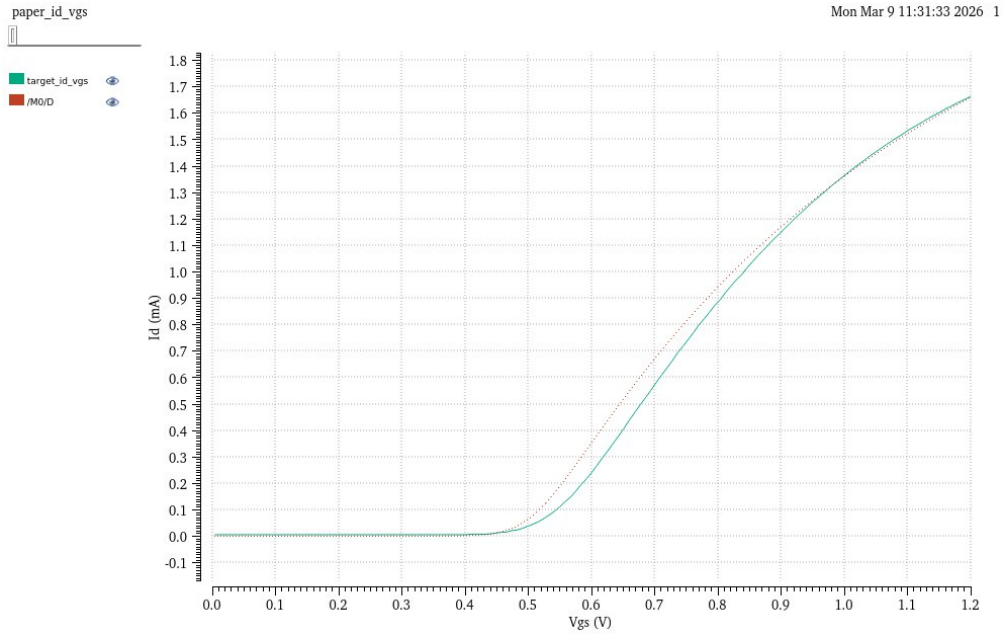

Figure 4.2. Comparison between the digitized cryogenic target curve (solid line) and the re-centered BSIM4 simulation (dotted line) for the transfer characteristic $I_D - V_{GS}$ at $V_{DS} = 50\,\mathrm{mV}$.

4.3.2 Circuit-level validation in an oscillator-like testbench

After transistor-level re-centering, the same cryogenic model was also verified in a simplified oscillator-like environment. The purpose of this step was not to validate the complete LC-VCO with all the original PDK passive devices under cryogenic conditions, but rather to confirm that the re-centered active-device model could sustain physically consistent oscillation when embedded in a resonant testbench representative of the intended application.

This distinction is important because the passive components available in the TSMC PDK were characterized for standard operating conditions, and their direct use at $-220\,^{\circ}\mathrm{C}$ would have produced non-meaningful extracted values. For this reason, the original foundry passives were not used for cryogenic oscillator validation. Instead, a surrogate resonant tank was reconstructed using *analogLib* components, so that the passive network could be ported to cryogenic temperature without relying on questionable temperature extrapolations of the original PDK passive models.

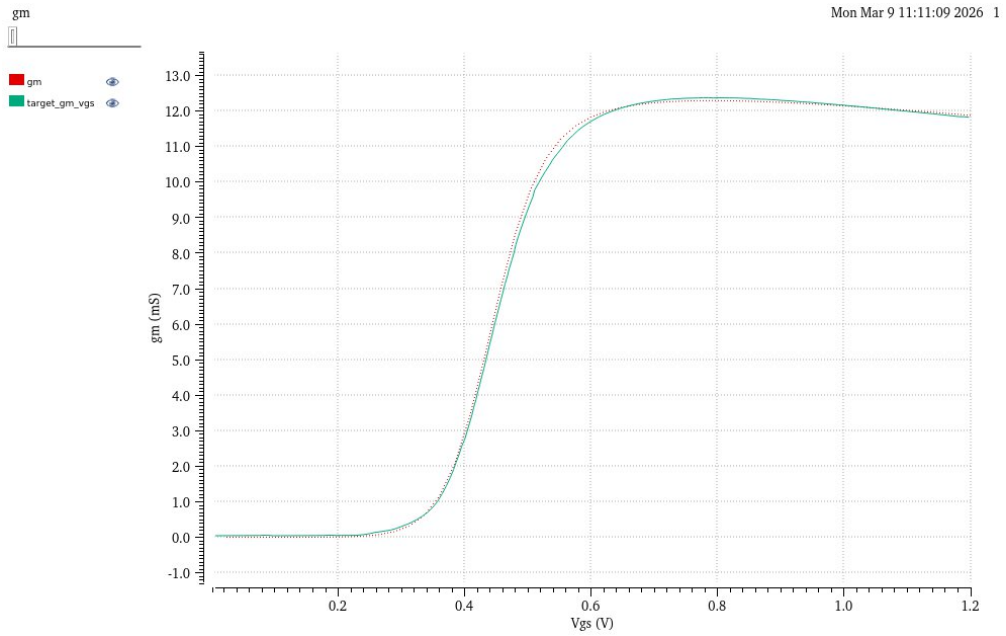

Figure 4.3. Comparison between target (solid line) and simulated (dotted line) transconductance $g_m - V_{GS}$.

The values of the surrogate tank were selected in order to reproduce a response close to that of the original TSMC tank around the target oscillation frequency. Figure 4.6 compares the impedance response of the original PDK tank and of the modeled surrogate tank, showing that the resonance frequency is preserved and that the overall impedance peak remains of the same order. In addition, the bias voltage was adjusted to 595 mV, yielding again a tail current of approximately 1.2 mA, consistent with the room-temperature operating point used in the schematic-level oscillator design.

Once this equivalent resonant environment was defined, the re-centered active-device model was inserted into the oscillator-like schematic shown in Fig. 4.7. In a first validation step, the equivalent parallel resistance of the tank was set to $R_p = 400\ \Omega$, i.e., close to the room-temperature value extracted from the original tank. In a second step, R_p was increased to 800 Ω in order to emulate an improved low-temperature loss condition. This assumption is physically motivated by cryogenic BEOL characterization reported in the literature for standard 65 nm CMOS, where line and via resistances were shown to decrease significantly at 4.2 K, while no relevant change in interconnect capacitance was observed [Tada et al., 2024]. In particular, the reported resistance reduction reaches approximately 75% for Cu interconnects and 20% for vias at 4.2 K .

Under both assumptions, the re-centered model led to stable and physically plausible oscillatory behavior. Figure 4.8 compares the differential output waveform for $R_p = 400\ \Omega$ and $R_p = 800\ \Omega$. As expected, the larger equivalent parallel resistance produces a larger steady-state oscillation amplitude because the active core needs to compensate lower tank

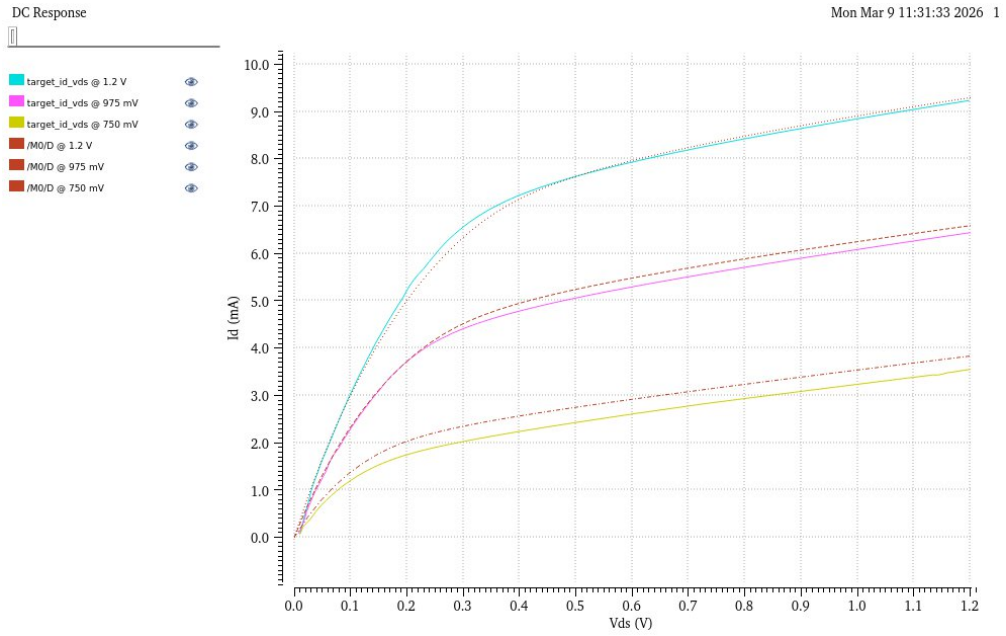

Figure 4.4. Comparison between target (solid line) and simulated (dotted line) output characteristics $I_D - V_{DS}$ for $V_{GS} = 0.75$, 0.975, and 1.2 V.

losses.

The corresponding output spectra are shown in Fig. 4.9. In both cases, the oscillator starts correctly and exhibits a well-defined spectral content, while the improved-loss case ($R_p = 800\ \Omega$) leads to a stronger fundamental component and, consistently, to a larger output swing.

This validation does not claim full cryogenic accuracy for the complete VCO, since the passive network was intentionally replaced by a surrogate tank. Rather, it demonstrates that the re-centered BSIM4 model is sufficiently consistent to support design-oriented circuit simulations in oscillator-like conditions, which is the level of accuracy required for the present work. The role of this validation is therefore not to replace the full circuit analysis of Chapter 3, but to confirm that the transistor-level cryogenic re-centering is compatible with oscillator operation in a representative resonant environment. Although this oscillator-like validation does not reproduce the full extracted VCO of Chapter 3, it remains consistent with the post-layout results discussed there, in the sense that the re-centered active-device model supports startup and physically meaningful oscillatory behavior in a representative resonant environment.

4.3.3 Methodological considerations and limitations

The proposed methodology has two main strengths. First, it is lightweight and immediately compatible with the original design environment, since it relies on parameter overrides rather than on custom model code. Second, it preserves the design flexibility of

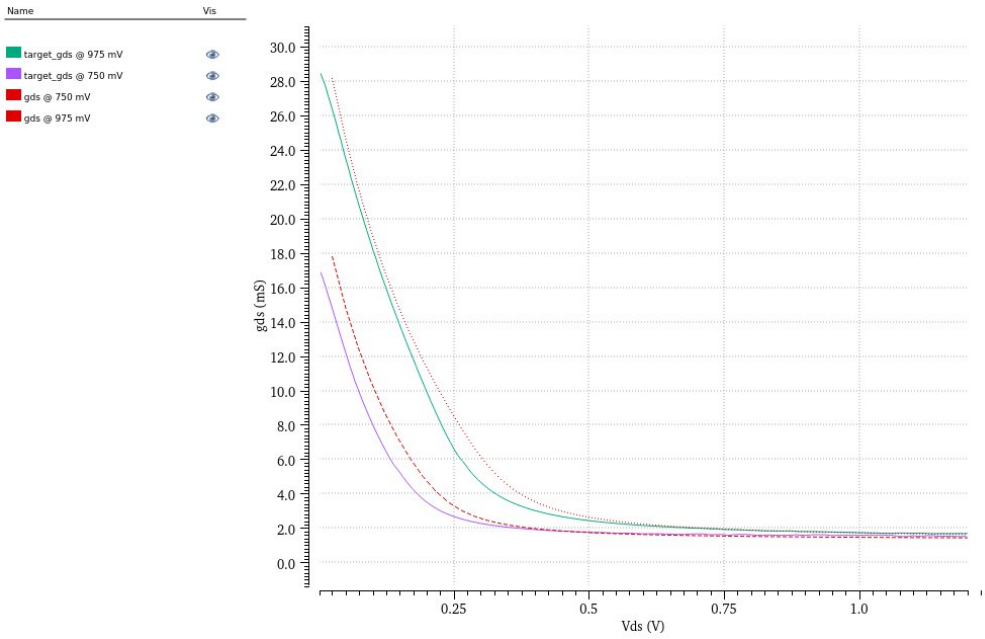

Figure 4.5. Comparison between target (solid line) and simulated (dotted line) output conductance $g_{ds} - V_{DS}$.

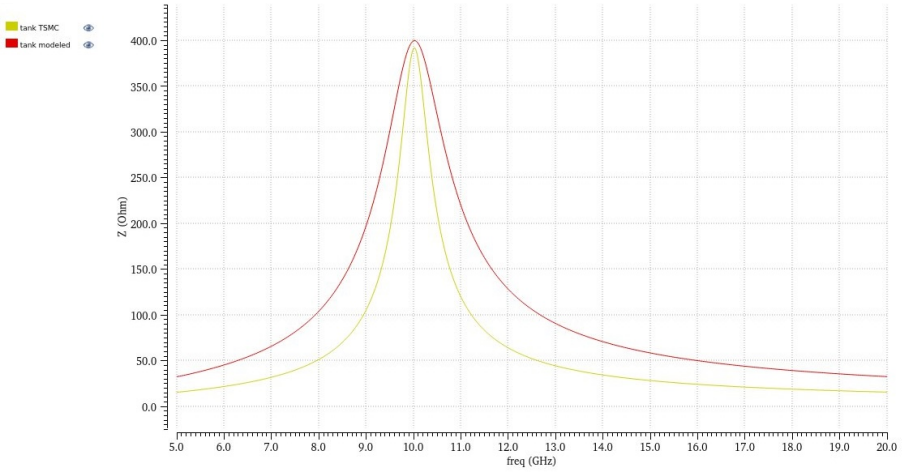

Figure 4.6. Comparison between the impedance response of the original TSMC tank and the surrogate tank reconstructed with *analogLib* components for cryogenic oscillator-like validation.

the PDK, allowing the cryogenic model to be reused directly in device- and circuit-level simulations.

At the same time, some limitations should be acknowledged. The fitting was based

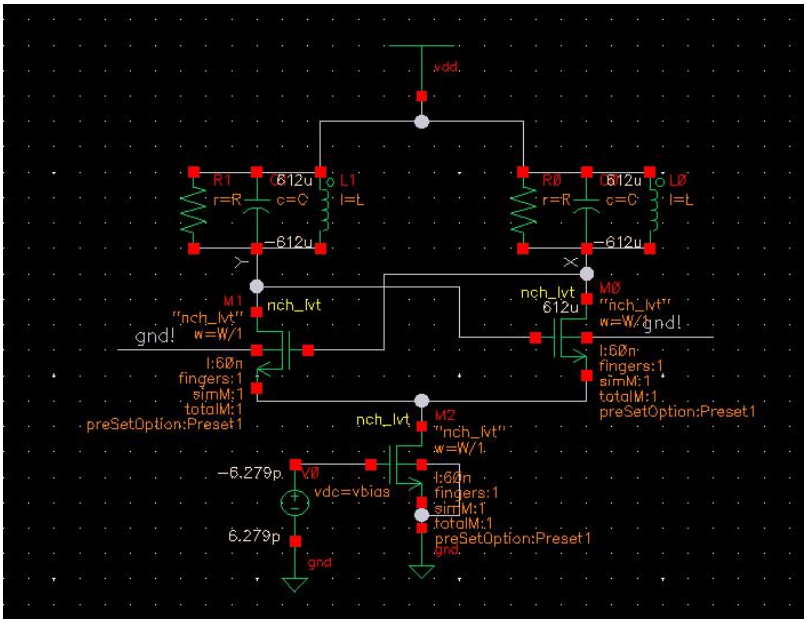

Figure 4.7. Oscillator-like cryogenic validation testbench using the re-centered NMOS model and a surrogate tank built with *analogLib* components. The bias voltage is adjusted to recover a tail current of approximately 1.2 mA.

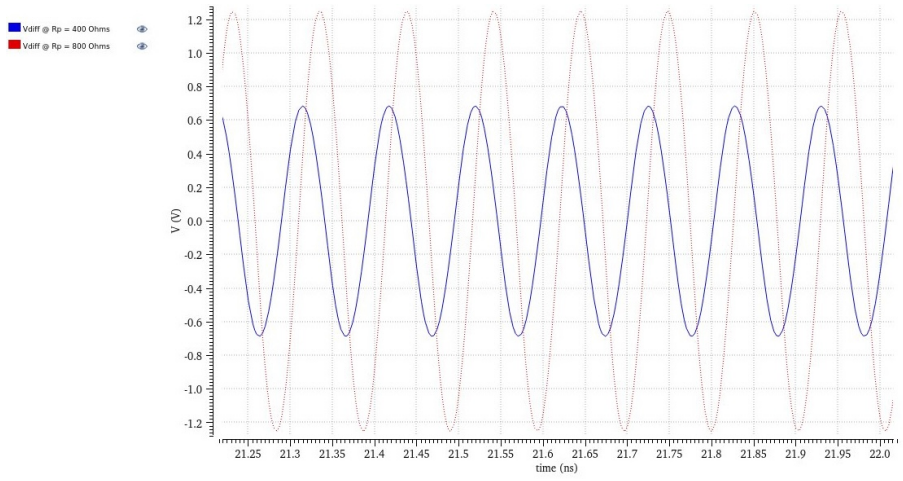

Figure 4.8. Differential output waveform obtained in the oscillator-like cryogenic validation testbench for two values of the equivalent tank resistance, $R_p = 400\ \Omega$ and $R_p = 800\ \Omega$.

on digitized curves from the literature rather than on direct measurements of the exact devices used in this work. Moreover, the validation was mainly visual and no formal multi-objective optimization metric was employed. Finally, the re-centering was performed

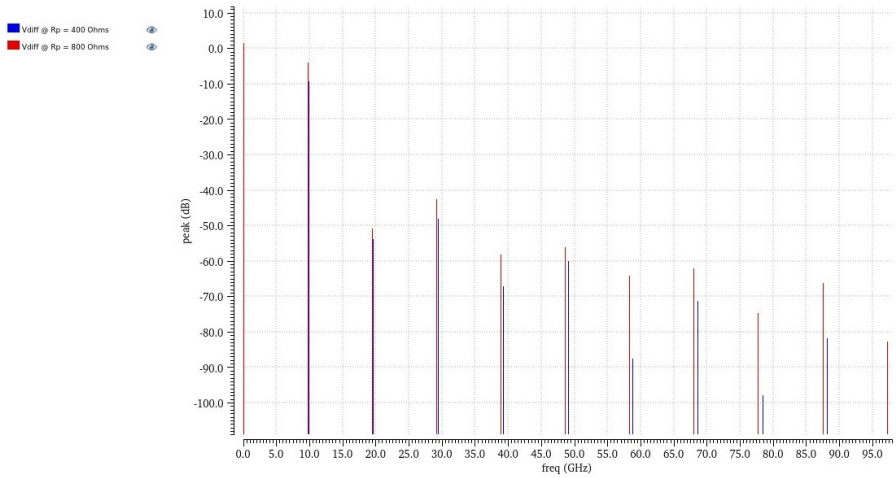

Figure 4.9. Output spectrum of the oscillator-like cryogenic validation testbench for $R_p = 400\ \Omega$ and $R_p = 800\ \Omega$.

on a representative transistor geometry, so its accuracy outside the fitted region should be considered with caution. Nevertheless, the achieved agreement was sufficient for the intended design task, namely the cryogenic validation of the LC-VCO active devices.

4.4 Conclusions

This chapter presented a design-oriented cryogenic re-centering procedure for the BSIM4 model of a TSMC 65 nm LVT NMOS transistor. Since no foundry-qualified cryogenic model was available, the room-temperature compact model was adapted through parameter overrides so as to reproduce the main low-temperature trends reported in the literature while preserving the original BSIM4 equations and compatibility with the standard Cadence/Spectre flow.

The fitting procedure was performed manually and iteratively, using digitized I_D–V_{GS}, $g_m - V_{GS}$, $I_D - V_{DS}$, and $g_{ds} - V_{DS}$ curves as complementary fitting targets. The final re-centered model reproduced the threshold shift, transconductance enhancement, and output-region behavior with sufficient accuracy for circuit-design purposes.

Finally, the model was also validated in a simplified oscillator-like testbench using a surrogate resonant tank, introduced because the original PDK passive models were not suitable for direct cryogenic use. Although this does not constitute a full cryogenic validation of the complete LC-VCO, it confirms that the re-centered transistor model is compatible with physically consistent oscillator behavior in a representative resonant environment. This level of accuracy was adequate to support a cryogenic, design-oriented interpretation of the LC-VCO results while remaining within a standard industrial CAD flow and without replacing the original BSIM4 model equations.

Chapter 5

General Conclusions

This thesis investigated the design of a cryogenic CMOS frequency-generation block for semiconductor spin-qubit control, with particular focus on a $9.5 - 10.5$ GHz LC-VCO intended as a first building block of a broader cryogenic control-chip architecture.

The work started from the physical and system-level requirements of semiconductor spin qubits. In these platforms, coherent manipulation relies on microwave control signals with high spectral purity, while large-scale integration requires moving at least part of the classical interface closer to the qubit plane. This motivates the adoption of cryogenic CMOS electronics as an intermediate layer between room-temperature digital control and the quantum processor. Within this framework, the thesis first introduced a reference mixed-signal cryogenic control architecture including bias generation, local memory, sequencing logic, RF generation, and readout support, showing how the proposed oscillator should be interpreted not as an isolated RF demonstrator, but as a functional subsystem of a wider scalable control platform.

The main design contribution of the thesis is the development of a 10 GHz LC-VCO in TSMC 65 nm CMOS with 9.5–10.5 GHz tuning range for cryogenic operation, consistent with microwave-driven control of silicon spin qubits. The design flow combined first-order analytical considerations with a simulation-driven methodology. Starting from the target frequency range, the required tank capacitance range was derived and partitioned between coarse switched-capacitor tuning and fine continuous tuning based on inversion-mode MOS varactors. The resonator losses were then characterized through AC simulations, the active core was sized from DC operating-point sweeps, and the oscillator was progressively verified through transient, PSS, and Pnoise analyses. Practical non-idealities were gradually introduced, including an implemented bias mirror, continuous tuning network, discrete capacitor bank, and finally the complete post-layout extracted view.

The obtained results show that the proposed VCO achieves the intended tuning range and preserves oscillation after full physical implementation. Post-layout simulations confirmed continuous tuning coverage over the target band and demonstrated stable operation under realistic parasitic loading. In addition, the phase-noise performance was translated into a qubit-level metric through the filter-function formalism, providing a direct link

between circuit-level oscillator behavior and single-qubit gate fidelity. This aspect is particularly important in the context of quantum-control electronics, since it shows that oscillator design quality cannot be evaluated only through conventional RF metrics, but should also be interpreted in terms of its impact on qubit coherence and gate error.

A further contribution of the thesis is the cryogenic interpretation of the active-device behavior through BSIM4 re-centering. Since no foundry-qualified compact model for low-temperature operation was available in the adopted PDK, a dedicated parameter-override methodology was developed to reproduce the main cryogenic trends reported in the literature while preserving compatibility with the standard Cadence/Spectre flow. The re-centered model was validated against digitized cryogenic transistor characteristics and then verified in an oscillator-like resonant testbench. Although this does not constitute a full cryogenic validation of the complete extracted VCO, it supports the physical plausibility of the active-device behavior under low-temperature conditions and therefore strengthens the interpretation of the oscillator results in a cryogenic-control context.

Overall, the thesis shows that a 10 GHz LC-VCO in a commercial CMOS technology can be designed, physically implemented, and evaluated in a way that is consistent with the constraints of semiconductor spin-qubit control. At the same time, the results highlight the main trade-offs that remain critical in this application domain, especially those involving phase noise, output swing, post-layout degradation, power dissipation, and low-temperature model reliability. In this sense, the proposed oscillator should be regarded as a first validated RF block within a broader cryogenic control strategy rather than as a complete final solution.

Several directions naturally emerge for future work. A first extension concerns the completion of the full frequency-generation chain, including the PLL and the modulation stage required to convert the continuous-wave carrier into qubit-control bursts. A second direction is the integration of the oscillator with the other subsystems of the proposed control chip, such as local memory, DAC-based bias generation, and digital sequencing logic. From the modeling viewpoint, future work should also address a more complete cryogenic validation of passive components and interconnect parasitics, ideally supported by dedicated low-temperature measurements. Finally, a broader system-level analysis of the trade-off between oscillator power consumption, spectral purity, and qubit-level fidelity would be especially valuable for assessing the scalability of cryogenic control electronics in large qubit arrays.

In conclusion, this work contributes to the development of cryogenic integrated electronics for semiconductor quantum computing by combining architectural analysis, RF circuit design, post-layout validation, and design-oriented cryogenic modeling. The proposed LC-VCO represents a concrete step toward the realization of compact and scalable cryogenic control hardware for spin-qubit platforms.

Appendix A

MATLAB Scripts for Phase-Noise-Based Fidelity Evaluation

This appendix reports the MATLAB routines used to estimate the average fidelity of a primitive X_π gate starting from simulated single-sideband phase-noise spectra exported from Cadence. The first function evaluates the fidelity for a single phase-noise dataset and also generates the corresponding overlap-integrand plot. The second function is a no-plot version used for automated sweeps. Finally, the third script compares the fidelity obtained from schematic-level and post-PEX phase-noise spectra as a function of the π-pulse duration.

A.1 Function for primitive X_π-gate fidelity estimation

Listing A.1 reports the MATLAB function used to estimate the average gate fidelity from a single phase-noise spectrum exported from Cadence. The routine imports the single-sideband phase noise, evaluates the primitive X_π filter function following the formulation discussed in Chapter 3, computes the overlap integral χ, and returns the corresponding average fidelity.

Listing A.1. MATLAB function used for primitive X_π-gate fidelity evaluation from simulated SSB phase-noise data.

```
function out = fidelity_from_phase_noise_green_Xpi(filename, tau)
%FIDELITY_FROM_PHASE_NOISE_GREEN_XPI
% Estimate the average fidelity of a primitive X_pi gate starting from
% single-sideband phase noise data exported from Cadence.
%
% The calculation follows the filter-function approach of:
% Green et al., "Arbitrary quantum control of qubits in the presence of
```

```
% universal noise", New J. Phys. 15, 095004 (2013).
%
% INPUTS
% filename : CSV file containing two columns:
% [offset_frequency_Hz, phase_noise_dBc_per_Hz]
% tau : pi-pulse duration [s]
%
% OUTPUT
% out : struct containing imported data, filter functions,
% overlap integral chi, and fidelity estimates.
%
% ASSUMPTIONS
% 1) The input CSV contains SSB phase noise L(f) in dBc/Hz.
% 2) The fidelity is evaluated using the primitive X_pi dephasing
% filter function Fz(omega).
% 3) The overlap integral is computed numerically as:
%
% chi = integral_0^inf [ 10^(L(f)/10) * Fz(2*pi*f) ] df
%
% and the average gate fidelity is approximated as:
%
% F_av = 0.5 * (1 + exp(-chi))
%
% NOTES
% - The point omega = Omega is treated numerically because the analytical
% response functions exhibit a removable singularity there.
% - This script operates directly on phase-noise data and does not
% reconstruct an intermediate voltage-noise quantity.

    if nargin < 2 || isempty(tau)
        tau = 10e-6;
    end

    %% Import data
    T = readtable(filename, 'Delimiter', ',', 'VariableNamingRule', 'preserve')
        ;

    f = T{:,1}; % offset frequency [Hz]
    L_dBc = T{:,2}; % SSB phase noise [dBc/Hz]

    % Ensure column vectors
    f = f(:);
    L_dBc = L_dBc(:);

    % Keep only valid rows
    valid = isfinite(f) & isfinite(L_dBc) & (f > 0);
    f = f(valid);
    L_dBc = L_dBc(valid);

```

```
    % Sort in ascending frequency
    [f, idx] = sort(f);
    L_dBc = L_dBc(idx);

    %% Angular-frequency variables
    omega = 2*pi*f; % offset angular frequency [rad/s]
    Omega = pi/tau; % Rabi angular frequency for a primitive X_pi pulse [rad/s]

    %% Primitive X_pi response functions (Green et al.)
    den = omega.^2 - Omega^2;
    common = exp(1j * omega * tau) + 1;

    Rzz = (omega.^2 ./ den) .* common;
    Rzy = (1j * omega * Omega ./ den) .* common;

    %% Numerical treatment of the removable singularity near omega = Omega
    tol = 1e-9 * Omega;
    near = abs(omega - Omega) < tol;

    if any(near)
        eps_rel = 1e-6;

        om_left = Omega * (1 - eps_rel);
        om_right = Omega * (1 + eps_rel);

        den_left = om_left^2 - Omega^2;
        den_right = om_right^2 - Omega^2;

        common_left = exp(1j * om_left * tau) + 1;
        common_right = exp(1j * om_right * tau) + 1;

        Rzz_left = (om_left^2 / den_left) * common_left;
        Rzz_right = (om_right^2 / den_right) * common_right;

        Rzy_left = (1j * om_left * Omega / den_left) * common_left;
        Rzy_right = (1j * om_right * Omega / den_right) * common_right;

        Rzz_lim = 0.5 * (Rzz_left + Rzz_right);
        Rzy_lim = 0.5 * (Rzy_left + Rzy_right);

        Rzz(near) = Rzz_lim;
        Rzy(near) = Rzy_lim;
    end

    %% Filter functions
    Gzz = abs(Rzz).^2;
    Gzy = abs(Rzy).^2;
    Fz = Gzz + Gzy;

```

```
105 %% Overlap integral
106 % Convert SSB phase noise from dBc/Hz to linear units
107 L_lin = 10.^(L_dBc / 10);
108
109 % Integrand of the overlap integral
110 integrand = L_lin .* Fz;
111
112 % Numerical integration over the simulated frequency grid
113 chi = trapz(f, integrand);
114
115 %% Fidelity estimates
116 F_av = 0.5 * (1 + exp(-chi));
117 F_lin = max(0, 1 - chi/2);
118
119 %% Store outputs
120 out = struct();
121 out.filename = filename;
122 out.tau = tau;
123 out.f = f;
124 out.omega = omega;
125 out.Omega = Omega;
126 out.L_dBc = L_dBc;
127 out.L_lin = L_lin;
128 out.Rzz = Rzz;
129 out.Rzy = Rzy;
130 out.Gzz = Gzz;
131 out.Gzy = Gzy;
132 out.Fz = Fz;
133 out.integrand = integrand;
134 out.chi = chi;
135 out.F_av = F_av;
136 out.F_lin = F_lin;
137
138 %% Plot: only overlap integrand
139 figure('Name', 'Primitive X_\pi overlap integrand', 'Color', 'w');
140 loglog(f, integrand, 'LineWidth', 1.5);
141 grid on;
142 xlabel('Offset frequency f [Hz]');
143 ylabel('10^{L(f)/10} F_z(2\pi f)');
144 title(sprintf('Overlap integrand for primitive X_\\pi, \\chi = %.3g, F_{av}
        = %.12f', chi, F_av));
145
146 %% Console output
147 fprintf('--- Primitive X_pi fidelity from SSB phase noise ---\n');
148 fprintf('File: %s\n', filename);
149 fprintf('tau = %.6g s\n', tau);
150 fprintf('Omega = %.6g rad/s\n', Omega);
151 fprintf('chi = %.12g\n', chi);
152 fprintf('F_av = %.12f\n', F_av);
```

```
153     fprintf('F_lin = %.12f\n', F_lin);
154 end
```

A.2 No-plot function for automated sweeps

Listing A.2 reports the no-plot version of the same fidelity-evaluation routine. This version is used for automated sweeps over the pulse duration without generating a figure at each iteration.

Listing A.2. MATLAB no-plot function used for automated primitive X_π-gate fidelity evaluation from simulated SSB phase-noise data.

```
1  function out = fidelity_from_phase_noise_green_noplot(filename, tau)
2  %FIDELITY_FROM_PHASE_NOISE_GREEN_NOPLOT
3  % Estimate the average fidelity of a primitive X_pi gate starting from
4  % single-sideband phase noise data exported from Cadence.
5  %
6  % The calculation follows the filter-function approach of:
7  % Green et al., "Arbitrary quantum control of qubits in the presence of
8  % universal noise", New J. Phys. 15, 095004 (2013).
9  %
10 % INPUTS
11 % filename : CSV file containing two columns:
12 % [offset_frequency_Hz, phase_noise_dBc_per_Hz]
13 % tau : pi-pulse duration [s]
14 %
15 % OUTPUT
16 % out : struct containing imported data, filter functions,
17 % overlap integral chi, and fidelity estimates.
18 %
19 % ASSUMPTIONS
20 % 1) The input CSV contains SSB phase noise L(f) in dBc/Hz.
21 % 2) The fidelity is evaluated using the primitive X_pi dephasing
22 % filter function Fz(omega).
23 % 3) The overlap integral is computed numerically as:
24 %
25 % chi = integral_0^inf [ 10^(L(f)/10) * Fz(2*pi*f) ] df
26 %
27 % and the average gate fidelity is approximated as:
28 %
29 % F_av = 0.5 * (1 + exp(-chi))
30 %
31 % NOTES
32 % - The point omega = Omega is treated numerically because the analytical
33 % response functions exhibit a removable singularity there.
34 % - This script operates directly on phase-noise data and does not
35 % reconstruct an intermediate voltage-noise quantity.
36
37     if nargin < 2 || isempty(tau)
```

```
38          tau = 10e-6;
39      end
40
41      %% Import data
42      T = readtable(filename, 'Delimiter', ',', 'VariableNamingRule', 'preserve')
            ;
43
44      f = T{:,1}; % offset frequency [Hz]
45      L_dBc = T{:,2}; % SSB phase noise [dBc/Hz]
46
47      % Ensure column vectors
48      f = f(:);
49      L_dBc = L_dBc(:);
50
51      % Keep only valid rows
52      valid = isfinite(f) & isfinite(L_dBc) & (f > 0);
53      f = f(valid);
54      L_dBc = L_dBc(valid);
55
56      % Sort in ascending frequency
57      [f, idx] = sort(f);
58      L_dBc = L_dBc(idx);
59
60      %% Angular-frequency variables
61      omega = 2*pi*f; % offset angular frequency [rad/s]
62      Omega = pi/tau; % Rabi angular frequency for a primitive X_pi pulse [rad/s]
63
64      %% Primitive X_pi response functions (Green et al.)
65      den = omega.^2 - Omega^2;
66      common = exp(1j * omega * tau) + 1;
67
68      Rzz = (omega.^2 ./ den) .* common;
69      Rzy = (1j * omega * Omega ./ den) .* common;
70
71      %% Numerical treatment of the removable singularity near omega = Omega
72      tol = 1e-9 * Omega;
73      near = abs(omega - Omega) < tol;
74
75      if any(near)
76          eps_rel = 1e-6;
77
78          om_left = Omega * (1 - eps_rel);
79          om_right = Omega * (1 + eps_rel);
80
81          den_left = om_left^2 - Omega^2;
82          den_right = om_right^2 - Omega^2;
83
84          common_left = exp(1j * om_left * tau) + 1;
85          common_right = exp(1j * om_right * tau) + 1;
```

```
86
87          Rzz_left = (om_left^2 / den_left) * common_left;
88          Rzz_right = (om_right^2 / den_right) * common_right;
89
90          Rzy_left = (1j * om_left * Omega / den_left) * common_left;
91          Rzy_right = (1j * om_right * Omega / den_right) * common_right;
92
93          Rzz_lim = 0.5 * (Rzz_left + Rzz_right);
94          Rzy_lim = 0.5 * (Rzy_left + Rzy_right);
95
96          Rzz(near) = Rzz_lim;
97          Rzy(near) = Rzy_lim;
98      end
99
100     %% Filter functions
101     Gzz = abs(Rzz).^2;
102     Gzy = abs(Rzy).^2;
103     Fz = Gzz + Gzy;
104
105     %% Overlap integral
106     L_lin = 10.^(L_dBc / 10);
107     integrand = L_lin .* Fz;
108     chi = trapz(f, integrand);
109
110     %% Fidelity estimates
111     F_av = 0.5 * (1 + exp(-chi));
112     F_lin = max(0, 1 - chi/2);
113
114     %% Store outputs
115     out = struct();
116     out.filename = filename;
117     out.tau = tau;
118     out.f = f;
119     out.omega = omega;
120     out.Omega = Omega;
121     out.L_dBc = L_dBc;
122     out.L_lin = L_lin;
123     out.Rzz = Rzz;
124     out.Rzy = Rzy;
125     out.Gzz = Gzz;
126     out.Gzy = Gzy;
127     out.Fz = Fz;
128     out.integrand = integrand;
129     out.chi = chi;
130     out.F_av = F_av;
131     out.F_lin = F_lin;
132 end
```

A.3 Script for schematic-versus-post-PEX fidelity comparison

Listing A.3 reports the MATLAB script used to compare the average primitive X_π-gate fidelity obtained from schematic-level and post-PEX phase-noise spectra over a sweep of π-pulse durations. The script calls the no-plot function of Listing A.2.

Listing A.3. MATLAB script used to compare schematic-level and post-PEX primitive X_π-gate fidelities as a function of the π-pulse duration.

```
% ============================================================
% Compare F_av(tau) for schematic and post-PEX phase noise
% ============================================================

clear; clc;

% --- Input files ---
file_sch = 'pn_schematic_300_bis.csv';
file_pex = 'pn_PEX_300_bis.csv';

% --- Tau sweep [s] ---
tau_vec = logspace(-8, -5, 80); % from 10 ns to 10 us

% --- Preallocate ---
F_sch = zeros(size(tau_vec));
F_pex = zeros(size(tau_vec));
chi_sch = zeros(size(tau_vec));
chi_pex = zeros(size(tau_vec));

% --- Sweep ---
for k = 1:length(tau_vec)
    out_sch = fidelity_from_phase_noise_green_noplot(file_sch, tau_vec(k));
    out_pex = fidelity_from_phase_noise_green_noplot(file_pex, tau_vec(k));

    F_sch(k) = out_sch.F_av;
    F_pex(k) = out_pex.F_av;

    chi_sch(k) = out_sch.chi;
    chi_pex(k) = out_pex.chi;
end

% --- Plot F_av(tau) ---
figure('Color','w');
semilogx(tau_vec, F_sch, 'LineWidth', 1.8); hold on;
semilogx(tau_vec, F_pex, 'LineWidth', 1.8);
grid on;
xlabel('\tau_\pi [s]');
ylabel('F_{av}');
title('Average gate fidelity vs \tau_\pi');
```

```
legend('Schematic', 'Post-PEX', 'Location', 'best');

% --- Optional: plot chi(tau) ---
figure('Color','w');
loglog(tau_vec, chi_sch, 'LineWidth', 1.8); hold on;
loglog(tau_vec, chi_pex, 'LineWidth', 1.8);
grid on;
xlabel('\tau_\pi [s]');
ylabel('\chi');
title('Decay parameter \chi vs \tau_\pi');
legend('Schematic', 'Post-PEX', 'Location', 'best');
```

Bibliography

B. Andreas, Y. Azuma, G. Bartl, P. Becker, H. Bettin, M. Borys, I. Busch, M. Gray, P. Fuchs, K. Fujii, H. Fujimoto, E. Kessler, M. Krumrey, U. Kuetgens, N. Kuramoto, G. Mana, P. Manson, E. Massa, S. Mizushima, A. Nicolaus, A. Picard, A. Pramann, O. Rienitz, D. Schiel, S. Valkiers, and A. Waseda. Determination of the avogadro constant by counting the atoms in a crystal. *Physical Review Letters*, 106(3):030801, 2011. doi: 10.1103/PhysRevLett.106.030801.

Harrison Ball, William D. Oliver, and Michael J. Biercuk. The role of master clock stability in quantum information processing. *npj Quantum Information*, 2:16033, 2016. doi: 10.1038/npjqi.2016.33.

Joseph C. Bardin, Evan Jeffrey, Erik Lucero, Trent Huang, Ofer Naaman, Rami Barends, Ted White, Marissa Giustina, Daniel Sank, Pedram Roushan, Kunal Arya, Benjamin Chiaro, Julian Kelly, Jimmy Chen, Brian Burkett, Yu Chen, Andrew Dunsworth, Austin Fowler, Brooks Foxen, Craig Gidney, Rob Graff, Paul Klimov, Josh Mutus, Matthew McEwen, Anthony Megrant, Matthew Neeley, Charles Neill, Chris Quintana, Amit Vainsencher, Hartmut Neven, and John Martinis. A 28nm bulk-cmos 4-to-8ghz <2mw cryogenic pulse modulator for scalable quantum computing. In *2019 IEEE International Solid-State Circuits Conference (ISSCC)*, pages 456–458. IEEE, 2019.

Arnout Beckers, Farzan Jazaeri, and Christian Enz. Cryogenic mos transistor model. *IEEE Transactions on Electron Devices*, 65(9):3617–3625, 2018.

Michael J. Biercuk, Andrew C. Doherty, and Hermann Uys. Dynamical decoupling sequence construction as a filter-design problem. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 44(15):154002, 2011. doi: 10.1088/0953-4075/44/15/154002.

Alberto Gatti and Filip Tavernier. Cryogenic small dimension effects and design-oriented scalable compact modeling of a 65-nm cmos technology. *IEEE Journal of the Electron Devices Society*, 12:369–377, 2024. doi: 10.1109/JEDS.2024.3394167.

Lukas Geck, Andreas Kruth, Natalia S. Molin, Filip Tavernier, Fabio Sebastiano, Robert Bogdan Staszewski, and Edoardo Charbon. A reconfigurable cryogenic platform for the classical control of scalable quantum computers. *Quantum Science and Technology*, 5(1):015004, 2020. doi: 10.1088/2058-9565/ab4c1d.

E. Gibney. Quantum computer quest. *Nature*, 2014.

Todd J. Green, Jarrah Sastrawan, Hermann Uys, and Michael J. Biercuk. Arbitrary quantum control of qubits in the presence of universal noise. *New Journal of Physics*, 15(9):095004, 2013. doi: 10.1088/1367-2630/15/9/095004.

E. A. Gutiérrez-D, M. J. Deen, and C. Claeys. Physics of silicon at cryogenic temperatures. In *Low Temperature Electronics*, pages 1–103. Academic Press, San Diego, CA, USA, 2001. doi: 10.1016/B978-012310675-9/50002-5.

A. Hajimiri and T. H. Lee. A general theory of phase noise in electrical oscillators. *IEEE Journal of Solid-State Circuits*, 33(2):179–194, February 1998. doi: 10.1109/4.658619.

Rosario M. Incandela, Lin Song, Harald Homulle, Edoardo Charbon, Andrei Vladimirescu, and Fabio Sebastiano. Characterization and compact modeling of nanometer cmos transistors at deep-cryogenic temperatures. *IEEE Journal of the Electron Devices Society*, 6:996–1006, 2018. doi: 10.1109/JEDS.2018.2821763.

F. H. L. Koppens, C. Buizert, K. J. Tielrooij, I. T. Vink, K. C. Nowack, T. Meunier, L. P. Kouwenhoven, and L. M. K. Vandersypen. Driven coherent oscillations of a single electron spin in a quantum dot. *Nature*, 442:766–771, 2006. doi: 10.1038/nature05065.

Yan Liu, Lili Lang, Yongwei Chang, Yi Shan, Xiaojie Chen, and Yemin Dong. Cryogenic characteristics of multinanoscales field-effect transistors. *IEEE Transactions on Electron Devices*, 68(2):456–463, 2021. doi: 10.1109/TED.2020.3041438.

D. Loss and D. P. DiVincenzo. Quantum computation with quantum dots. *Physical Review A*, 57:120–126, 1998.

Jongseok Park, Sushil Subramanian, Lester Lampert, Todor Mladenov, Ilya Klotchkov, Dileep J. Kurian, Esdras Juarez-Hernandez, Brando Perez Esparza, Sirisha Rani Kale, Asma Beevi K. T., Shavindra P. Premaratne, Thomas F. Watson, Satoshi Suzuki, Mustafijur Rahman, Jaykant B. Timbadiya, Saksham Soni, and Stefano Pellerano. A fully integrated cryo-cmos soc for state manipulation, readout, and high-speed gate pulsing of spin qubits. *IEEE Journal of Solid-State Circuits*, 56(11):3289–3303, 2021. doi: 10.1109/JSSC.2021.3115988.

Bishnu Patra, Rosario M. Incandela, Jeroen P. G. van Dijk, Harald A. R. Homulle, Liran Song, Mahdi Shahmohammadi, Robert Bogdan Staszewski, Andrei Vladimirescu, Masoud Babaie, Fabio Sebastiano, and Edoardo Charbon. Cryo-cmos circuits and systems for quantum computing applications. *IEEE Journal of Solid-State Circuits*, 53(1): 309–321, January 2018. doi: 10.1109/JSSC.2017.2737549.

L. Petit, J. M. Boter, H. G. J. Eenink, G. Droulers, M. L. V. Tagliaferri, R. Li, D. P. Franke, K. J. Singh, J. S. Clarke, R. N. Schouten, V. V. Dobrovitski, L. M. K. Vandersypen, and M. Veldhorst. Spin lifetime and charge noise in hot silicon quantum dot qubits. *Physical Review Letters*, 121(7):076801, 2018. doi: 10.1103/PhysRevLett.121.076801.

J. J. Rael and A. A. Abidi. Physical processes of phase noise in differential lc oscillators. In *Proceedings of the IEEE Custom Integrated Circuits Conference (CICC)*, 2000.

Behzad Razavi. *Design of CMOS Phase-Locked Loops: From Circuit Level to Architecture Level*. Cambridge University Press, 2020. ISBN 9781108494540. doi: 10.1017/9781108626200.

Fabio Sebastiano, Harald A. R. Homulle, Jeroen P. G. van Dijk, Rosario M. Incandela, Bishnu Patra, Mohammadreza Mehrpoo, Masoud Babaie, Andrei Vladimirescu, and Edoardo Charbon. Cryogenic cmos interfaces for quantum devices. In *2017 IEEE International Conference on Rebooting Computing (ICRC)*, pages 59–62. IEEE, 2017.

K. M. Svore and M. Troyer. The quantum future of computation. *Computer*, 49(9), 2016.

Munehiro Tada, Koichiro Okamoto, Takahisa Tanaka, Makoto Miyamura, Hiroki Ishikuro, Ken Uchida, and Toshitsugu Sakamoto. A 65nm cryogenic cmos design and performance at 4.2k for quantum state controller application. *IEEE Journal of the Electron Devices Society*, 12:28–33, 2024. doi: 10.1109/JEDS.2023.3340136.

Mehrnoosh Vahidpour, William O'Brien, Jon Tyler Whyland, Joel Angeles, Jayss Marshall, Diego Scarabelli, Genya Crossman, Kamal Yadav, Yuvraj Mohan, Catvu Bui, Vijay Rawat, Russ Renzas, Nagesh Vodrahalli, Andrew Bestwick, and Chad Rigetti. Superconducting through-silicon vias for quantum integrated circuits, 2017.

J. P. G. van Dijk, E. Kawakami, P. van Woerkom, et al. Impact of classical control electronics on qubit fidelity. *arXiv preprint*, 2018.

J. P. G. van Dijk, E. Charbon, and F. Sebastiano. The electronic interface for quantum processors. *Microprocessors and Microsystems*, 66:90–101, 2019. doi: 10.1016/j.micpro.2019.02.004.

L. M. K. Vandersypen, H. Bluhm, J. S. Clarke, A. S. Dzurak, R. Ishihara, A. Morello, D. J. Reilly, L. R. Schreiber, and M. Veldhorst. Interfacing spin qubits in quantum dots and donors—hot, dense, and coherent. *npj Quantum Information*, 3:34, 2017. doi: 10.1038/s41534-017-0038-y.

Lieven M. K. Vandersypen and Mark A. Eriksson. Quantum computing with semiconductor spins. *Physics Today*, 72(8):38–45, August 2019. doi: 10.1063/PT.3.4270.

M. Veldhorst, J. C. C. Hwang, C. H. Yang, A. W. Leenstra, B. de Ronde, J. P. Dehollain, J. T. Muhonen, F. E. Hudson, K. M. Itoh, A. Morello, and A. S. Dzurak. An addressable quantum dot qubit with fault-tolerant control-fidelity. *Nature Nanotechnology*, 9(12): 981–985, 2014. doi: 10.1038/nnano.2014.216.

R. Versluis. Quantum computers scale up. *IEEE Spectrum*, 57(4), 2020.

C. H. Yang, W. H. Lim, F. A. Zwanenburg, and A. S. Dzurak. Dynamically controlled charge sensing of a few-electron silicon quantum dot. *AIP Advances*, 1(4):042111, 2011. doi: 10.1063/1.3654496.

Floris A. Zwanenburg, Andrew S. Dzurak, Andrea Morello, Michelle Y. Simmons, Lloyd C. L. Hollenberg, Gerhard Klimeck, Sven Rogge, Susan N. Coppersmith, and Mark A. Eriksson. Silicon quantum electronics. *Reviews of Modern Physics*, 2013.