

UNIVERSITÀ
DI PAVIA

UNIVERSITY OF PAVIA
FACULTY OF ENGINEERING
DEPARTMENT OF ELECTRICAL, COMPUTER AND BIOMEDICAL ENGINEERING

MASTER'S DEGREE IN ELECTRONIC ENGINEERING

MASTER'S THESIS

Study and comparison of high-speed clock distribution techniques

Studio e confronto di tecniche di distribuzione di segnali di clock ad alta velocità

Candidate:
Giuseppe Bonarrigo

Supervisor:
Prof. Danilo Manstretta

Co-supervisors:
Simone Giroletti
Gabriele Minoia
Enrico Temporiti

A.Y. 2025/2026

Abstract

This thesis was carried out in collaboration with Marvell Technology, an American semiconductor company focused on high-performance integrated circuits for data infrastructure. The present work compares two clock distribution techniques for high-speed clock signals spanning a frequency band from 25 GHz to 30 GHz. A conventional non-resonant CMOS implementation is first analyzed as a reference solution in terms of power consumption, signal attenuation, bandwidth, jitter, and jitter amplification. After a theoretical introduction of the main metrics of interest for evaluating high-frequency clock signals, an inductive element is then added to make the clock distribution network resonate, and the resulting structure is re-evaluated under the same operating conditions to assess the impact of resonance on its overall performance. Circuit simulations, carried out in a 3 nm FinFET technology, show that, within the target operating band, the resonant approach reduces power consumption, jitter, and jitter amplification, at the cost of changing the clock distribution network from a broadband solution to a narrowband one, operating over a limited frequency range.

Contents

Abstract **i**

1 Introduction **1**
1.1 Motivation . . . 1
1.2 Evolution trends in high-speed communication systems . . . 2
1.3 Optical communication systems . . . 3
1.4 High-speed SerDes architecture . . . 5
1.5 Clock distribution in high-speed systems . . . 7

2 Theoretical Background **9**
2.1 Objective of the work . . . 9
2.2 Clock metrics . . . 10
2.3 Jitter definition . . . 13
2.3.1 Jitter characterization in clock signals . . . 13
2.3.2 Voltage-to-time conversion . . . 15
2.3.3 Jitter decomposition . . . 15
2.4 Jitter amplification . . . 19
2.5 Transmission-Line Modeling for Clock Distribution . . . 22
2.5.1 Transmission-line theory . . . 24
2.5.2 Equivalent circuit representation of the interconnect . . . 26
2.5.3 Impedance transformation through a transmission line . . . 28

3 Non-resonant clock distribution **31**
3.1 Analysis of the reference architecture . . . 31
3.1.1 Driver sizing analysis . . . 32
3.1.2 Effect of transmission-line length . . . 38
3.2 Design methodology for non-resonant clock distribution . . . 42
3.2.1 Tapered-chain design . . . 43
3.2.2 Driver sizing with transmission-line loading . . . 46
3.2.3 Clock distribution network design procedure . . . 50
3.3 Case Study: 1 mm Transmission Line . . . 51

3.3.1 Single-segment implementation . . . 52
3.3.2 Two-segment implementation . . . 55
3.3.3 Four-segment implementation . . . 57
3.3.4 Comparison of line implementations . . . 63
3.4 Effect of Back-to-Back Inverter Placement . . . 64
3.4.1 Placement at the transmission-line input . . . 65
3.4.2 Placement at the transmission-line midpoint . . . 68
3.4.3 Placement at the transmission-line output . . . 69
3.4.4 Summary of the Back-to-Back configurations . . . 71

4 Resonant clock distribution **73**
4.1 Motivation for resonant clock distribution . . . 73
4.2 Resonant architectures and operating principle . . . 74
4.2.1 Near-end resonant configuration . . . 75
4.2.2 Midpoint resonant configuration . . . 81
4.2.3 Far-end resonant configuration . . . 85
4.2.4 Comparison of the resonant placements . . . 88
4.3 Case study: 1 mm transmission line . . . 92
4.3.1 Near-end resonant implementation . . . 92
4.3.2 Midpoint resonant implementation . . . 99
4.3.3 Far-end resonant implementation . . . 104
4.3.4 Comparison of the three resonant approaches . . . 108
4.4 Summary . . . 109

5 Comparison and conclusions **111**
5.1 Comparison between non-resonant and resonant approaches . . . 111
5.2 Future developments . . . 117
5.3 Conclusions . . . 118

Bibliography **121**

Chapter 1

Introduction

1.1 Motivation

In recent years, the amount of data being created, processed, and transferred has drastically increased, as illustrated in Figure 1.1. This trend is particularly evident in data centers, where demanding applications such as artificial intelligence and cloud computing require large numbers of CPUs and GPUs to work together through high-speed interconnections [1, 2].

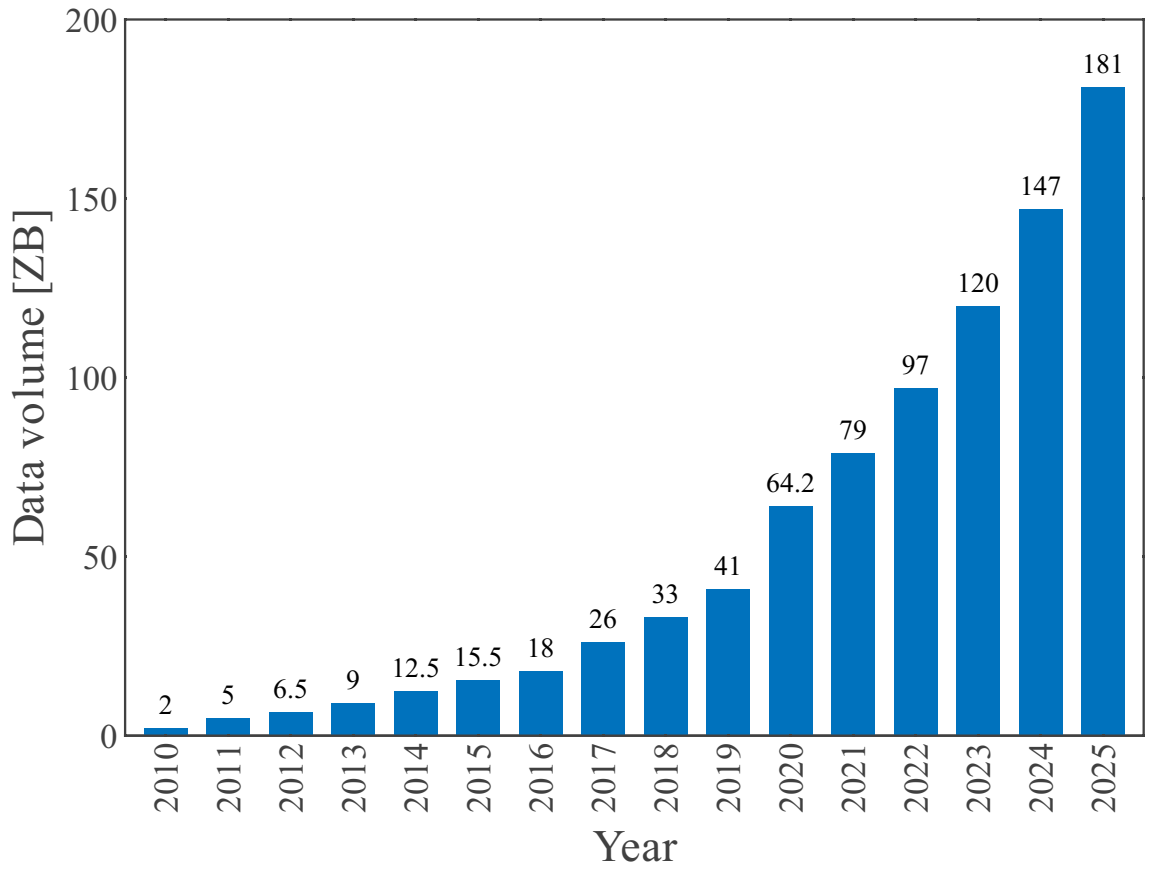


Figure 1.1: Worldwide data-volume growth from 2010 to 2025 [3].

As data volumes scale, the movement of data, both across geographically distributed data centers and among processing units within a single facility, becomes a critical bottleneck, potentially constraining overall system throughput and performance. In this context, to send more data over a small number of physical channels, modern wireline networks are increasingly pushed toward high per-lane data rates. Recent high-speed serial links operate at data rates of 100 Gb/s per lane and above. Even greater aggregate throughputs, reaching several hundreds of Gb/s, can be achieved by parallelizing multiple lanes, as demonstrated in modern 800 Gb/s optical transceivers [1, 2]. Consequently, the continuous scaling of per-lane and aggregate data rates has made the design of high-speed communication links a key limiting factor in large-scale computing infrastructures.

This trend brings significant challenges at both architectural and circuit level, especially in terms of bandwidth limitations and timing accuracy [4, 5].

1.2 Evolution trends in high-speed communication systems

The architectural organization of data centers provides a useful framework for contextualizing the evolution of modern high-speed communication systems. In data centers, servers are physically grouped into structures called racks. The servers inside a rack are locally connected to a specific switch device, commonly referred to as a top-of-rack switch, which collects the traffic within the rack and provides connectivity to other racks through higher-level switching stages [6].

A schematic representation of a data center interconnection network is shown in Figure 1.2. In this respect, while communication links are confined within the same rack, dedicated links provide connectivity across different racks, distinct areas within the same data center, or even between geographically distributed data centers. Consequently, these links must support heterogeneous transmission distances, while simultaneously satisfying the demand for high data rates driven by the increasing data volume.

Among the various approaches proposed in the literature, the data rate of interconnection systems can be enhanced either by increasing the number of physical channels or by raising the data rate per channel [7]. These two strategies are not mutually exclusive but can be concurrently leveraged within the same system architecture. The first approach enables high aggregate throughput by increasing the degree of parallelism, while maintaining relatively moderate per-channel data rates. This approach is limited by the available number of I/O pins, which in turn depends on package complexity, routing density, and cost [7]. Therefore, modern communication systems are increasingly pushed toward high per-lane data rates.

In this respect, the effectiveness of increasing the data rate is intrinsically limited by

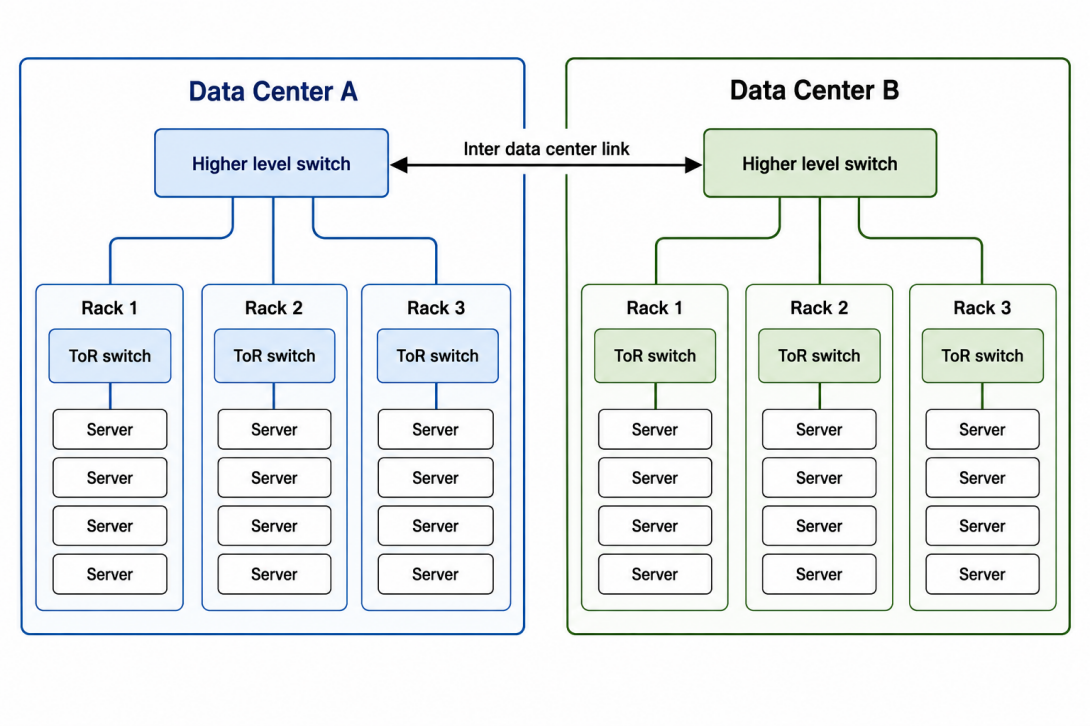


Figure 1.2: Schematic representation of data-center interconnections.

the physical constraints of the transmission channel, which in turn affect the integrity of the transmitted data. Indeed, attenuation, dielectric losses, reflections, impedance discontinuities, and crosstalk degrade the signal and make the transmission link more sensitive to noise. In addition, the attenuation accumulated along the channel increases with transmission distance, resulting in pronounced degradation of the transmitted signal.

Depending on the target link distance, different solutions can be employed. For short-reach interconnects, electrical links remain attractive due to their ease of integration with electronic circuitry and cost efficiency, particularly when channel losses are limited and can be compensated with moderate circuit complexity [1]. As the required distance and data rate increase, loss mechanisms in copper interconnects significantly degrade signal integrity [8]. In this scenario, optical interconnects represent a more suitable solution as they exhibit significantly lower loss per unit length compared to copper [7].

As a result, optical communication systems are currently the dominant solution for long-distance data transmission [9]. An overview of these systems is provided in Section 1.3.

1.3 Optical communication systems

From a general perspective, even in optical communication systems, data are generated and processed in the electrical domain. Indeed, the electrical signal is first converted into an optical signal and then transmitted through the fiber. At the receiver side, the opposite

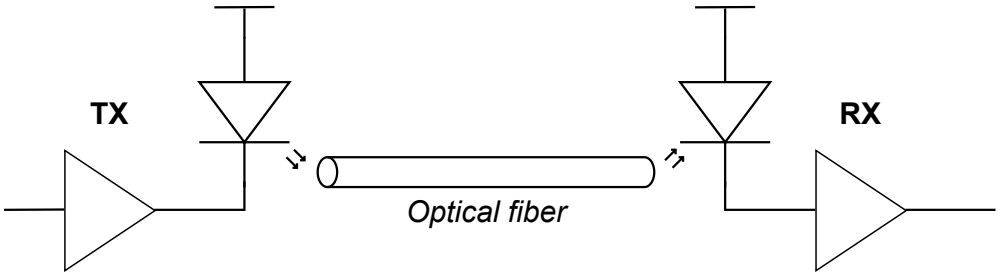


Figure 1.3: Simplified optical communication system.

conversion is performed, so that the optical signal is converted back into the electrical domain to be processed by the dedicated electronic systems. A simplified representation of this electrical-optical-electrical conversion is shown in Figure 1.3.

In a generic optical communication system, three main functional blocks can be identified:

- **Transmitter (TX):** it receives the electrical data stream and generates the high-speed driving signal applied to the optical modulator. The electrical-to-optical conversion can be obtained either by directly modulating a laser source, such as a Vertical Cavity Surface Emitting Laser (VCSEL), or by driving an external optical modulator [10]. In the first case, the laser acts as both the optical source and the modulating device, since the optical data stream is generated through variations in the laser driving current [10]. This leads to a simpler transmitter architecture than the externally modulated alternative, since no additional optical modulator is required. However, the achievable speed is affected by laser dynamics and parasitics [11]. In the second case, when an external modulator is adopted, the laser provides the optical carrier, while the modulator imposes the data on the optical signal, e.g., through absorption or phase modulation [10]. This solution can be more suitable for high-speed operation than direct modulation approaches, at the cost of an additional optical device and a dedicated driver [10]. In both cases, the transmitter must generate a high-speed electrical signal compatible with the selected optical interface.

- **Optical fiber channel:** it carries the optical signal from the transmitter to the receiver. Compared to copper-based interconnections, the optical fiber has lower distance-dependent losses and a much larger available bandwidth, thus making it suitable for long-reach high-speed links [7]. However, the channel cannot be considered ideal, since coupling losses, dispersion, and bandwidth limitations in both optical and electrical interfaces can still affect the integrity of the transmitted signal.

- **Receiver (RX):** it converts the optical signal back into the electrical domain. This operation is typically performed by a photodiode, which generates a current pro-

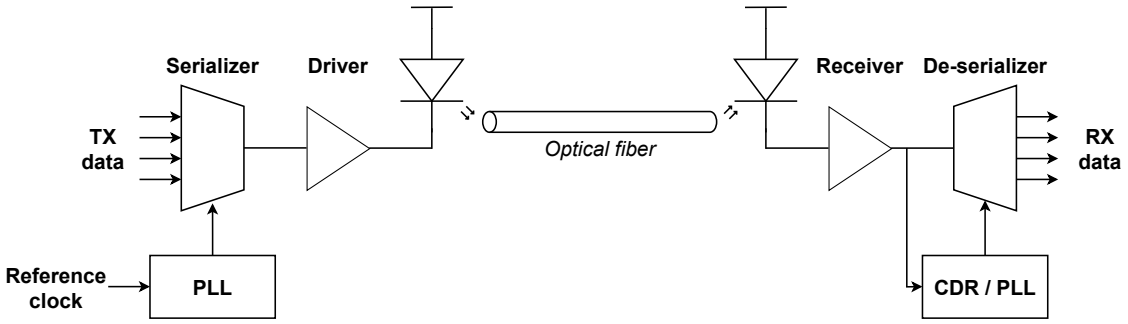


Figure 1.4: Simplified optical SerDes architecture with clock generation (PLL) and recovery systems (CDR).

portional to the received optical power [7]. The generated current is then converted into a voltage and amplified by a transimpedance amplifier (TIA) [7]. Usually, the TIA is a critical block of optical receivers, because its noise, gain, and bandwidth directly affect the sensitivity and quality of the recovered signal [1].

Ultimately, the performance of an optical communication system also depends on the electronic and optoelectronic interfaces placed at the two ends of the channel. A more detailed overview of a generic transceiver (i.e., transmitter and receiver) architecture adopted for data transmission on optical links is provided in Section 1.4.

1.4 High-speed SerDes architecture

In practical implementations, while dealing with bi-directional data flow, the transmitter and the receiver are usually implemented within the same chip. In this framework, the transmitter-receiver structure is usually referred to as transceiver. As discussed in Section 1.2, modern communication systems require several processing units to exchange large amounts of data. These data are typically generated and processed in parallel in the digital domain by a Digital Signal Processing (DSP) engine [1, 7]. At the transmitter side, the parallel data are converted by a serializer into a high-speed serial stream and sent through the channel. A deserializer recovers the serial data stream and converts it back to a parallel data flow for signal post-processing.

A simplified representation of this architecture is shown in Figure 1.4. The figure shows the basic organization of a generic optical transceiver[1] link and highlights the main functional blocks along the signal path. Details that depend on the specific circuit implementation, such as equalization, driver topology, biasing, and receiver decision circuitry, are omitted to keep the diagram focused on the overall architecture.

In this respect, in the transmitter side, the serializer converts the parallel input data into a high-speed serial stream. In this way, the information generated and processed in

[1]Due to its serializer/deserializer function, the transceiver architecture is also commonly referred to as SerDes.

parallel inside the electronic system, at a relatively low data rate (typically up to a few Gb/s), can be transmitted over a single physical channel at an increased data rate. This conversion is usually performed through a sequence of multiplexing stages, by combining the parallel input data into a serial stream, which progressively increase the data rate up to the final line rate. The serialized signal is then fed to the output driver, which is responsible for properly driving the optical source or modulator. In practical high-speed transmitters, equalization is included in the transmitter data path to compensate for the channel non-idealities and preserve the integrity of the transmitted signal [2, 4, 7, 12].

The optical-to-electrical conversion is performed at the receiver side by a photodiode, which generates a current related to the received optical power. This current is processed by the receiver front-end circuit, typically including a transimpedance amplifier, which converts the photodiode current into a voltage and provides the first signal conditioning. Then, the recovered high-speed serial stream is applied to the receiver back-end. In the simplified architecture of Figure 1.4, this block is represented as a deserializer, which converts the serial data back into a parallel format for subsequent digital processing. Depending on the receiver architecture, an Analog-to-Digital Converter (ADC) can also be used at this stage to convert the incoming analog signal into the digital domain.

To accomplish the aforementioned operation, at the transmitter side, the serializer requires a high-speed clock to control the multiplexing stages. This clock is generated by a phase-locked loop (PLL) starting from a lower-frequency reference clock. Since the serialization is usually performed in several stages, the clock must be delivered to different parts of the transmitter path with the proper frequency and phase [4, 5, 7].

Consequently, at the receiver side, the incoming serial stream can be correctly interpreted only if the sampling clock is properly aligned with the received data. Indeed, if the clock edge is not aligned with the data eye diagram, errors in the signal sampling process can arise, leading to incorrect bit recovery and information corruption. The alignment between the received data and sampling clock signals is performed by the clock and data recovery circuit (CDR), which extracts the timing information from the received signal and generates the sampling clock used to sample the incoming serial data and drive the subsequent deserialization stages. To this end, the CDR continuously aligns the sampling phase to sample close to the center of the eye diagram[2], where the timing margin is maximized while the error probability is minimized [1, 7, 12].

As the data rate increases, the timing margin of the link reduces. Indeed, for a given timing error, a short unit interval leads to a large reduction of the horizontal eye opening. Therefore, if the clock phases used by the serializer are not correctly positioned, or if the receiver samples too close to a data transition, the available margin is reduced and the

[2]An eye diagram is obtained by superimposing many consecutive bit intervals of the received signal over a single unit interval. In this way, all possible signal transitions are shown together. The empty region that remains in the middle of the plot is called the eye opening, because its shape resembles an eye.

probability of recovering an incorrect bit increases [1, 5, 7].

Therefore, once generated, the clock must also be distributed across the relevant TX and RX blocks while preserving timing accuracy and signal integrity. As a result, this requirement makes high-frequency clock distribution an increasingly critical and challenging task in modern transceiver circuits.

1.5 Clock distribution in high-speed systems

As discussed in Section 1.4, a high-speed clock is used both at the transmitter side to control the serializer and at the receiver side to define the sampling instants for data recovery. After being generated or recovered, the clock signal must reach the corresponding circuit blocks through interconnects and buffers. The distribution network introduces parasitic loading effects, contributing to increased propagation delay, phase skew, and waveform distortion, and resulting in overall clock degradation [1, 5].

For short clock routes, the interconnect can often be approximated as a lumped capacitive load. As the physical length of the route increases, especially at multi-GHz frequencies, this approximation becomes less representative of the actual behavior of the line. Consequently, propagation delay, attenuation, phase shifts, and distributed effects modify the clock signal along the distribution path. A more detailed discussion of these effects is provided in Section 2.5.

Therefore, the clock delivered to the target circuit blocks may differ significantly from the clock generated by the source. Depending on the length of the distribution path and on the requirements on voltage swing, rise/fall time, and jitter, additional circuit stages may be needed to drive the interconnect and restore the clock signal. However, these stages might introduce additional loading and device non-idealities, and can increase the overall power consumption and complexity of the clock distribution network. Therefore, high-speed clock distribution networks require specific circuit solutions to deliver the clock while limiting its degradation introduced along the path [1, 5].

In the literature, several implementations have been proposed to distribute a high-speed clock [1, 5]. A common approach is based on CMOS inverter chains, which provide full-swing clock signals and are compatible with CMOS logic blocks using the clock itself. Alternatively, the clock can be distributed using current-mode logic (CML) drivers. This approach leads to different trade-offs in terms of voltage swing, power consumption, supply-noise rejection, and interfacing with the following CMOS logic or sampling blocks [5]. Other approaches have also been proposed in the literature, including transmission-line-based and resonant clock distribution networks [1, 5]. This work will focus on CMOS clock distribution and on the comparison between non-resonant and resonant behavior of the clock line.

From the viewpoint of the distribution network, a clock line can be implemented

either as a non-resonant transmission line or as a resonant structure. In the resonant implementation, an inductive element is introduced so that the clock distribution network resonates around the required operating frequency. In this condition, part of the energy required to charge and discharge the clock-line capacitance is exchanged with the inductive element, reducing the current that must be provided by the driver to achieve the required clock swing. Therefore, resonance can reduce the current consumption of the driver and can also provide advantages in other clock-distribution metrics [1, 5, 7].

Accordingly, Chapter 2 introduces the main clock metrics and the transmission-line theory required for the analysis of the considered clock-distribution techniques.

Chapter 2

Theoretical Background

2.1 Objective of the work

The scope of the present thesis is to investigate and compare two high-speed clock distribution techniques suitable to operate in the 25–30 GHz frequency range, with the aim of identifying their main design trade-offs and deriving practical guidelines for their implementation. The comparison is carried out through circuit-level simulations in a 3 nm FinFET technology node.

The first technique investigated consists of a conventional non-resonant CMOS clock distribution network, where CMOS inverters are used to drive the transmission line and regenerate the clock signal along the interconnect. This implementation is adopted as a reference solution and is analyzed in terms of power consumption, signal integrity, bandwidth, jitter, and jitter amplification.

The second technique that has been investigated builds upon the aforementioned clock distribution technique by introducing an inductive element in the clock path to resonate with the interconnect capacitance at the target operating frequency. This solution is examined under the same conditions to assess the impact of resonance on network behavior and to determine the conditions under which it can provide advantages over the conventional non-resonant approach.

The comparison between the two clock-distribution strategies is based on a common clock-distribution case study. The reference scenario involves distributing a high-frequency clock signal over a distance of 1 mm, reflecting a realistic implementation case. More in details, a supply voltage around 0.75 V is assumed at typical operating conditions, together with a die temperature of 85 °C and an input clock frequency ranging from 25 GHz to 30 GHz. The complete set of design specifications of the defined case study is summarized in Table 2.1.

The remainder of the thesis is organized as follows:

- Chapter 2 provides the theoretical background required for the subsequent chap-

Table 2.1: Design specifications of the clock-distribution case study.

Parameter	**Acronym**	**Sim. conditions**	**Min**	**Typ**	**Max**	**Unit**
Clock-distribution length	LENGTH	–	–	1	–	mm
Die temperature	TDIE	Typical	–	85	–	°C
Supply voltage	AVDD	Typical	–	750	–	mV
Current consumption	IDD	Typical	–	–	25	mA
Input clock frequency	FREQ	Typical	25	–	30	GHz
Duty cycle	DC	Typical	45	–	55	%
Input voltage swing	dVI	Typical	–	750	–	mVpp
Output voltage swing	dVO	Typical	80% AVDD	–	–	mVpp
Output rise/fall time	TRF_O	Typical	–	–	10	ps
Random jitter	RJ	Typical	–	–	30	fs-rms
Supply-induced jitter	PSIJ	Typical	–	–	40	fs/mV
Jitter amplification	JA	Typical	–	–	1	fs/fs

Note: The input voltage swing is specified for each single-ended clock branch, whereas the output voltage swing, rise/fall time, and jitter-related specifications are evaluated differentially.

ters, including the main conventional clock-signal metrics, jitter-related effects, and transmission-line models used to describe clock propagation along on-die interconnects.

- Chapter 3 analyzes the conventional non-resonant CMOS implementation, first by discussing its operating principle and then by applying it to the 1 mm interconnect case study.

- Chapter 4 extends the analysis to the resonant implementation, first investigating how the introduction of the inductive element modifies the clock-distribution network behavior and the resulting design trade-offs. Specific results for the case study under consideration are then presented.

- Chapter 5 provides a comparison of the clock-distribution techniques discussed in Chapters 3 and 4, summarizes the main conclusions, and discusses potential future developments.

2.2 Clock metrics

A clock signal is a periodic waveform that repeats itself after a time interval T_{clk}, referred to as the clock period. Therefore, within one second the waveform repeats $f_{\mathrm{clk}} = 1/T_{\mathrm{clk}}$ times, where f_{clk} is the clock frequency.

Ideally, a clock signal alternates between two voltage levels: a low level V_{LOW} and a high level V_{HIGH}. The difference between these levels defines the voltage swing of the clock waveform, which is usually denoted as the peak-to-peak voltage V_{pp}:

$$V_{\text{pp}} = V_{\text{HIGH}} - V_{\text{LOW}}. \tag{2.1}$$

In high-speed circuits, clocks are often distributed in differential form to reject common-mode disturbances and improve the effective signal-to-noise ratio. In fact, for the same single-ended swing on each branch, the differential voltage swing is doubled, while uncorrelated noise contributions from the two branches add in rms value, resulting in a 3 dB SNR improvement under ideal balanced conditions [13, 14]. Therefore, the differential clock is carried by two complementary single-ended branches, whose voltages are denoted as $V_p(t)$ and $V_n(t)$. The differential clock signal $V_{\text{diff}}(t)$ can then be defined as:

$$V_{\text{diff}}(t) = V_p(t) - V_n(t), \tag{2.2}$$

where the common-mode voltage is defined as

$$V_{\text{CM}}(t) = \frac{V_p(t) + V_n(t)}{2}. \tag{2.3}$$

As a result, the voltage swing can be specified both for single-ended clock signals and for the differential ones. In this respect, if the two single-ended clock branches are driven with the same peak-to-peak swing, the differential peak-to-peak voltage is twice the single-ended value, provided that the two branches switch in opposite directions.

During each clock period, the signal remains at the low level for a time t_{LOW} and at the high level for a time t_{HIGH}. As a result,

$$T_{\text{clk}} = t_{\text{LOW}} + t_{\text{HIGH}}. \tag{2.4}$$

The ratio between the high-level duration and the clock period defines the duty cycle,

$$D = \frac{t_{\text{HIGH}}}{T_{\text{clk}}}. \tag{2.5}$$

Ideally, for a symmetric clock waveform, the duty cycle is 50%, meaning that the signal remains at the high level and at the low level for the same amount of time (i.e., 50% of the clock period T_{clk}) within each period.

However, in a real circuit, a clock transition is not instantaneous since it requires the capacitances along the clock path to be charged or discharged. These capacitances are associated with the drivers, the interconnects, and the load. As a result, the rise and fall times are limited by the effective time constants and by the bandwidth of the clock-distribution path [15, 16]. The low-to-high transition of a signal is commonly referred to as the rising edge, while the high-to-low transition is called falling edge. To quantify the duration of the low-to-high transition, the rise time is defined as the time required for the clock signal to transition from 20% to 80% of its voltage swing. In general, rise and

fall times are often defined using the 10%–90% voltage levels, which characterize most of the signal transition [17]. In this thesis, however, the 20%–80% definition is adopted consistently because the analysis is mainly focused on the central part of the clock edge, where the timing crossing used for jitter evaluation is located.

Figure 2.1 provides a visual definition of the aforementioned parameters.

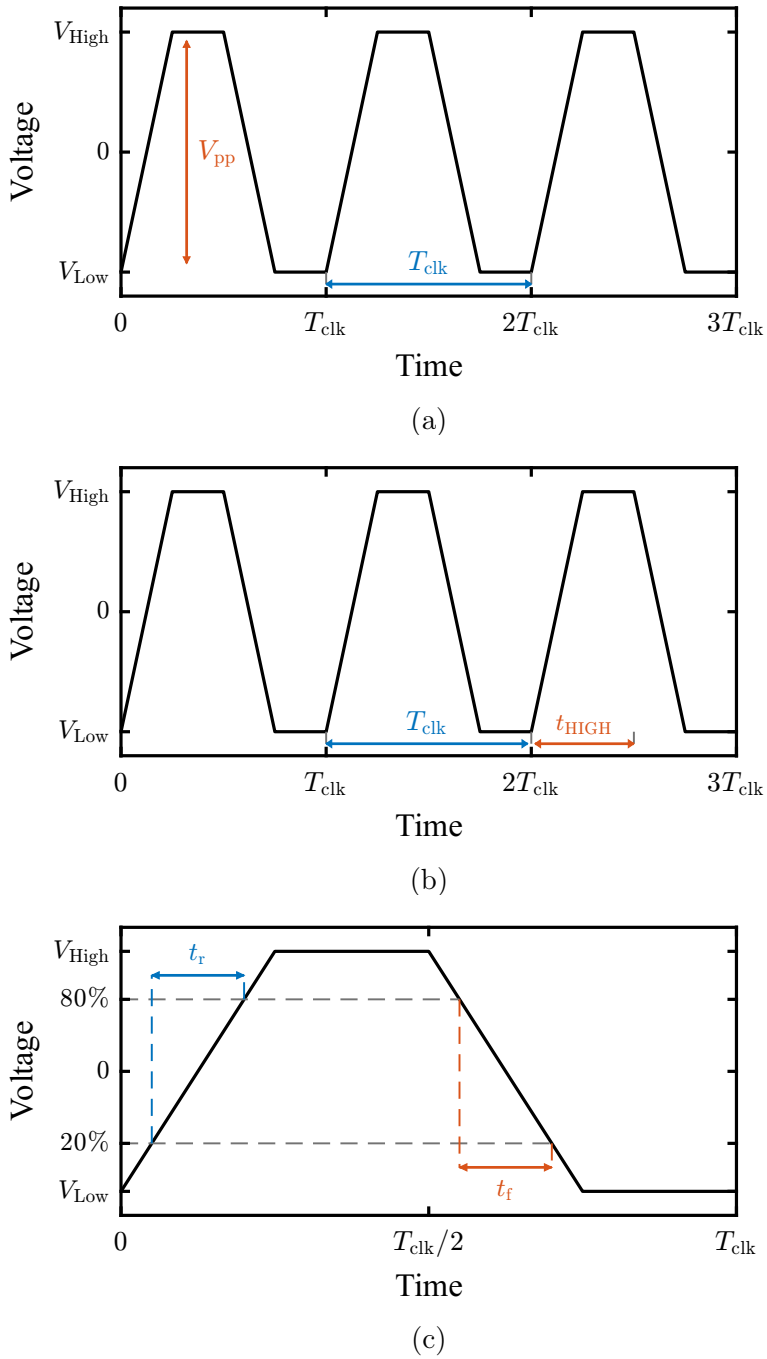


Figure 2.1: Ideal differential clock waveform illustrating the relevant clock metrics: (a) clock voltage swing V_{pp} and period T_{clk}; (b) clock duty cycle; (c) rise time t_{r} and fall time t_{f}, evaluated between the 20% and 80% voltage levels.

2.3 Jitter definition

Jitter is defined as the deviation of the timing position of a signal edge from its ideal position in time [17]. Depending on the timing quantity used as reference, different jitter metrics can be defined. In this work, the main quantity of interest is the edge jitter, which describes the displacement of a single clock edge with respect to its ideal timing position.

Since the edge transition is not instantaneous, the timing position of the signal edge must be associated with a specific point on the transition. According to [17], the transition instant is identified at the 50% reference-level crossing between the low and high voltage levels. This threshold is relevant because, for an ideal symmetric receiver, it represents the nominal decision point between the two logic states. Therefore, a displacement of this crossing directly corresponds to a timing error. For a symmetric differential clock, the same condition corresponds to the zero crossing of the differential voltage.

2.3.1 Jitter characterization in clock signals

Different jitter metrics can be defined depending on the timing quantity considered. In the literature, edge jitter, period jitter, and cycle-to-cycle jitter are commonly used to characterize clock signals [18].

The timing quantities involved in the definition of the three jitter metrics are shown in Figure 2.2. For the n-th edge, t_n denotes the measured crossing time, while $t_{\text{ideal},n}$ denotes the corresponding ideal crossing time expected by a jitter-free signal. The nominal clock period is denoted by T_{clk}. Starting from these quantities, the three jitter metrics are defined as follows.

- *Edge jitter* is the timing displacement of a single clock edge with respect to its ideal position and is defined as

$$J_{\text{edge}}(n) = t_n - t_{\text{ideal},n}. \tag{2.6}$$

  This metric identifies the timing error of one clock edge with respect to its ideal crossing time.

- *Period jitter* represents the timing error of one clock period. Considering consecutive edges with the same polarity, the period of the n-th cycle is

$$T_n = t_n - t_{n-1}. \tag{2.7}$$

  Therefore, the period jitter is defined as

$$J_{\text{per}}(n) = T_n - T_{\text{clk}} = (t_n - t_{n-1}) - T_{\text{clk}}. \tag{2.8}$$

  This metric identifies the error in the duration of a single clock period.

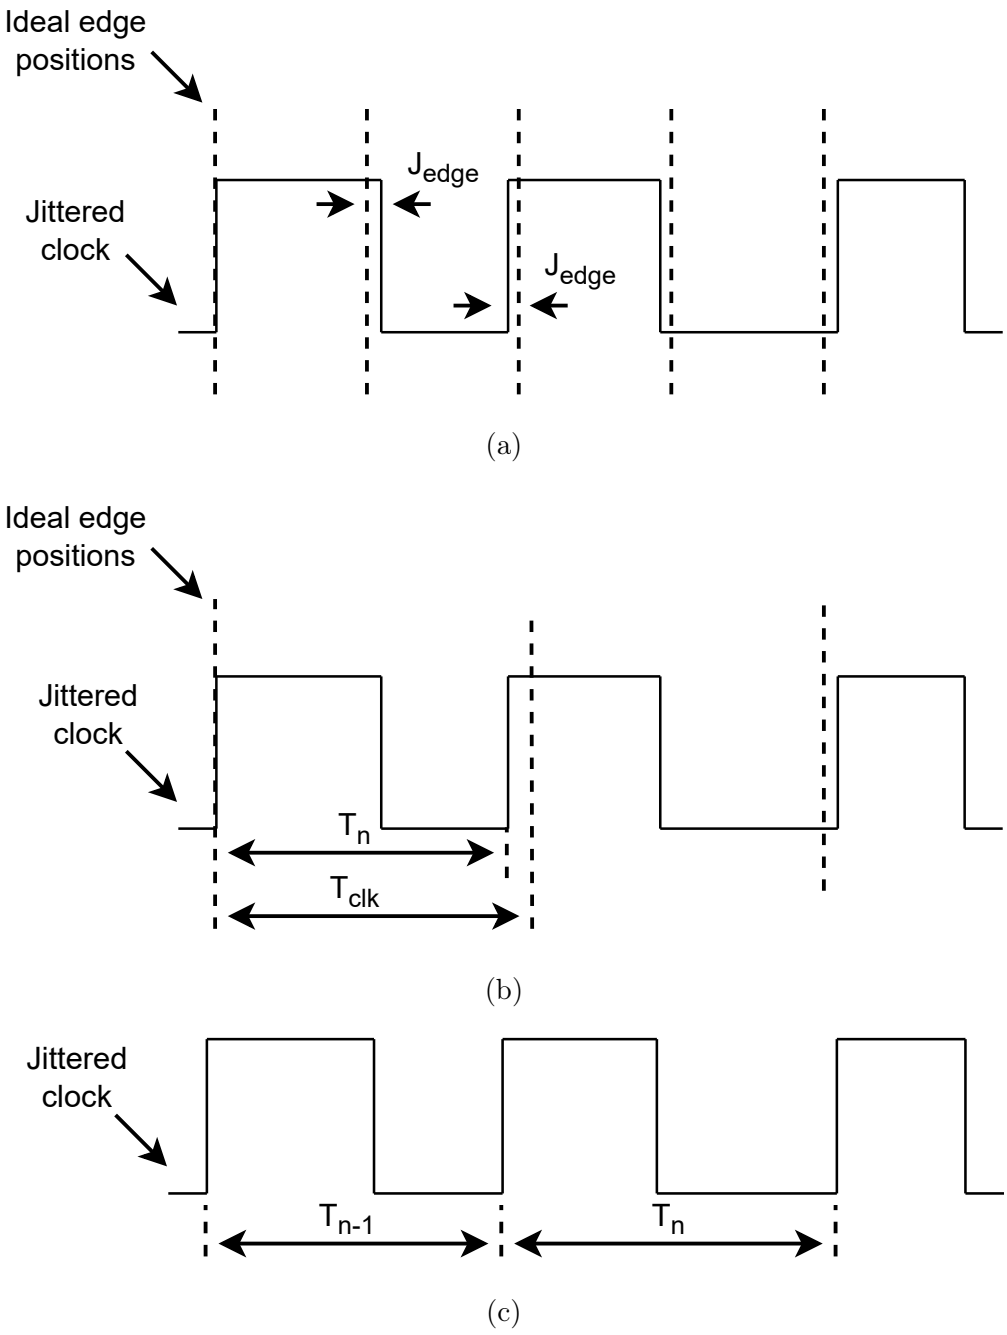


Figure 2.2: Theoretical representation of edge (a), period (b), and cycle-to-cycle (c) jitter evaluation methods.

- *Cycle-to-cycle jitter* represents the difference between two consecutive clock periods. It is defined as

$$J_{\mathrm{cc}}(n) = T_n - T_{n-1} = (t_n - t_{n-1}) - (t_{n-1} - t_{n-2})\,. \tag{2.9}$$

This metric tracks the change in clock period between two consecutive cycles.

In this work, the main focus is on edge jitter, because the analyzed clock-distribution networks are evaluated in terms of the timing displacement of their propagated clock

edges. From this point on, unless otherwise specified, the term jitter will refer to edge jitter.

2.3.2 Voltage-to-time conversion

Since the edge position is defined by the threshold crossing, any perturbation around this point can shift the measured timing of the edge. For instance, voltage noise adds a vertical displacement to the signal, making it cross the reference level earlier or later than in the ideal case.

Edge jitter magnitude depends on both the voltage noise at the reference crossing point and the signal slope. This dependence comes from the voltage-to-time conversion induced by the finite edge slope: for the same voltage perturbation, a steep transition reduces the resulting timing displacement, whereas a slow transition produces large jitter.

Since the voltage perturbation is a random quantity, its effect on the edge position is also described statistically. Therefore, the jitter contribution produced by voltage noise is expressed in rms value, which represents the standard deviation of the timing displacement around the ideal crossing time. According to [5], the rms timing jitter σ_t can be defined as

$$\sigma_t \simeq \frac{\sigma_v}{|S|} \tag{2.10}$$

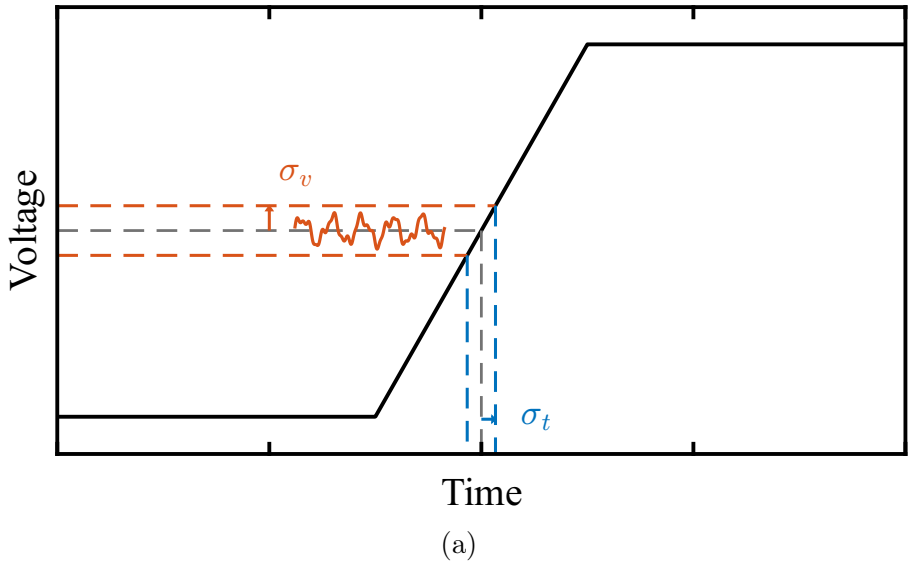
where σ_v is the rms voltage noise at the reference crossing point and $S = dv/dt$ is the local slope of the waveform at the same point.

Thus, for a given clock voltage swing, increasing the rise time reduces the slope around the reference crossing point and therefore increases the timing displacement produced by the same voltage perturbation. This dependence is qualitatively shown in Figure 2.3, where the same rms voltage noise is applied to two clock waveforms featuring different rise times. In Figure 2.3a, assuming a given noise level affecting the clock signal, a steep signal edge reduces the crossing-time spread, thus resulting in low jitter levels. Conversely, in Figure 2.3b, assuming the same noise level, a slow edge produces a large crossing-time spread, resulting in high jitter levels.

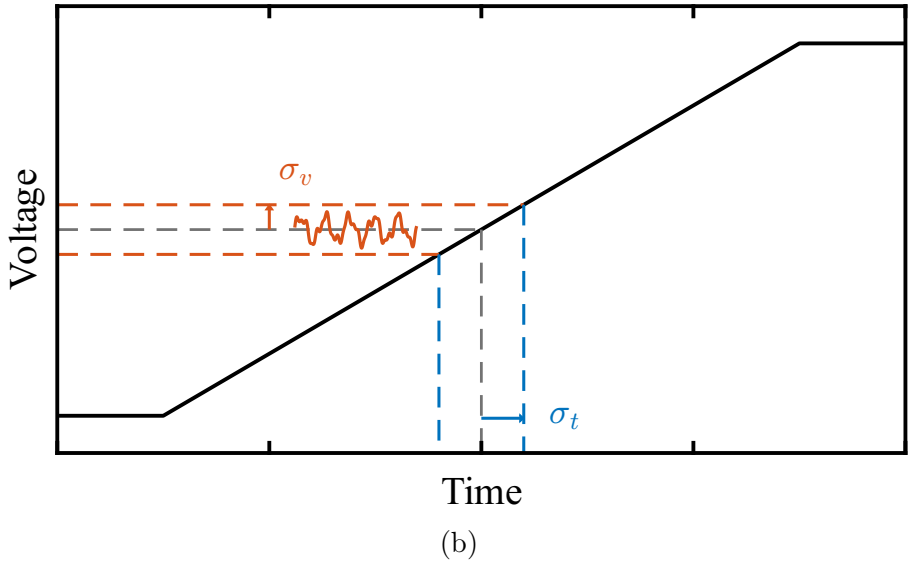
2.3.3 Jitter decomposition

In real systems, the observed jitter is composed of different contributions. For this reason, total jitter can usually be decomposed into separate components to distinguish random timing variations from repeatable or bounded effects.

According to the jitter model described in [19], total jitter can be first separated into two main categories: random jitter (RJ), which represents the unpredictable components mainly associated with random noise sources, and deterministic jitter (DJ), which rep-



(a)



(b)

Figure 2.3: Voltage-to-time conversion produced by the finite clock-edge slope. The same voltage-noise variation results in a small timing displacement for the steep transition in (a) and a large timing displacement for the slow transition in (b).

resents repeatable or bounded components. More in details, the deterministic jitter can be further discussed in terms of duty-cycle distortion (DCD), bounded uncorrelated jitter (BUJ), and data-dependent jitter (DDJ). DCD describes a deviation of the duty cycle from its nominal value, BUJ represents bounded timing variations that are not correlated with the transmitted data pattern, such as disturbances produced by coupling or crosstalk and DDJ is a pattern-dependent timing variation mainly associated with inter-symbol interference.

In addition to this statistical decomposition, it is useful to identify the physical mechanisms that generate timing variations in CMOS clock distribution networks. Among them, power-supply-induced jitter (PSIJ) is particularly relevant, since supply-voltage fluctuations can modulate the delay of clock buffers and shift the timing of the propagated clock

edges [5]. In the following, the principal jitter components will be briefly presented.

2.3.3.1 Random jitter

Random jitter is the unpredictable part of the timing variation. Indeed, it does not show a repeatable pattern and is mainly associated with random noise sources in the circuit, such as thermal noise, shot noise, flicker noise, and the combined effect of several uncorrelated noise sources. For modeling purposes, random jitter is often described by a Gaussian distribution [19]. This assumption is justified by the central limit theorem, since the combined effect of many uncorrelated noise sources tends toward a Gaussian distribution. Specifically, in CMOS clock distribution circuits, transistor thermal noise in the clock buffers is an important physical source of random jitter [5].

Since the Gaussian distribution is unbounded, random jitter does not have a fixed peak-to-peak value. Therefore, the jitter peak-to-peak value depends on the number of observed samples and on the observation time. Thus, random jitter is usually specified by its rms value, which corresponds to the standard deviation of the timing distribution.

2.3.3.2 Deterministic jitter

Deterministic jitter represents the component of timing variation that originates from repeatable mechanisms. Therefore, unlike random jitter, the deterministic jitter is bounded, and its peak-to-peak value can be observed or estimated with high confidence from a finite data set [19]. In the decomposition considered in this section, deterministic jitter is assumed to include duty-cycle distortion, bounded uncorrelated jitter, and data-dependent jitter. Instead, power-supply-induced jitter is discussed separately below as a physical source of timing error in CMOS clock distribution networks.

- **Duty-cycle distortion**

  Duty-cycle distortion is the deviation of the duty cycle from its nominal value. The nominal value corresponds to the intended or specified duty cycle of the clock signal, which may be 50% for a symmetric waveform but is not limited to this case. When duty-cycle distortion is present, one portion of the period becomes longer than the other.

  This effect can originate from unmatched rising and falling slew rates, or from a decision threshold that is not centered with respect to the signal levels [19]. In both cases, the threshold-crossing instants of rising and falling edges are shifted by different amounts, producing a repeatable timing error. This is why duty-cycle distortion is classified as a deterministic jitter component.

- **Bounded Uncorrelated Jitter**

Bounded uncorrelated jitter is a timing variation not correlated with the transmitted data pattern. It is often associated with coupling or crosstalk from nearby channels or circuit blocks [19]. In clock distribution networks, similar bounded disturbances can arise from electromagnetic coupling or other deterministic aggressors that are not directly related to the clock pattern.

- **Data-Dependent Jitter**

  Although data-dependent jitter is mainly relevant in serial data links rather than in periodic clock distribution, it is briefly discussed here for the sake of completeness. Data-dependent jitter is a deterministic timing variation correlated with the transmitted bit sequence. It occurs when the timing of a transition depends on the symbols that preceded it, thus nominally identical transitions may cross the decision threshold at different instants depending on the previous data pattern.

  A common cause of data-dependent jitter is inter-symbol interference, which arises when the channel or device bandwidth is not sufficient to let each symbol settle independently before the next one is transmitted. In this case, residual contributions from previous symbols affect the waveform of the current symbol. As a result, the threshold-crossing instant becomes pattern-dependent, producing a deterministic timing error associated with the data sequence [18, 19].

  For this reason, data-dependent jitter is mainly relevant in serial data links, where arbitrary bit sequences propagate through bandwidth-limited channels. In a clock signal, instead, the sequence is fixed and alternates regularly between the two logic states, corresponding to a repeated pattern. Therefore, the transition timing is not data-dependent in the same sense as in an arbitrary serial data sequence. Bandwidth-limited propagation may distort the clock waveform and shift the clock edges, but this effect is interpreted as signal distortion or jitter amplification rather than data-dependent jitter.

2.3.3.3 Power-supply-induced jitter

Power-supply-induced jitter is a timing variation caused by fluctuations of the supply voltage. In CMOS clock distribution networks, these fluctuations modulate the propagation delay of clock buffers and shift the timing of the propagated clock edges. This mechanism plays a crucial role in inverter-based clock distribution, since CMOS buffers exhibit stronger sensitivity to supply noise than supply-isolated clocking structures such as CML buffers [5].

Although power-supply-induced jitter is not included among the jitter categories used in the decomposition model adopted from [19], it is a significant physical source of timing error in CMOS clock distribution networks. Indeed, depending on the spectral content of

the supply disturbance, the resulting jitter can appear as a deterministic contribution at the clock distribution output node, thus impairing the expected clock signal performance.

2.4 Jitter amplification

Following the identification of the main jitter contributions, the propagation of jitter in a clock distribution network becomes a relevant aspect. Each block along the clock path might affect the timing of the propagated clock edges, both by introducing additional jitter and by transferring the input jitter to the output with a certain jitter transfer behavior.

In this respect, jitter amplification across a clock distribution network occurs when the jitter at the output of a clock path is larger than the jitter applied at its input. This effect should be distinguished from jitter generation. Indeed, while jitter generation is caused by noise sources inside the circuit, such as transistor thermal noise or power-supply-induced delay variations, jitter amplification, instead, describes the increase of an already existing timing perturbation as it propagates through the clock path. Therefore, it can occur even in a noiseless distribution network [5].

In the following, the clock signal applied to the clock-distribution path will be referred to as the input clock, while the corresponding signal after propagation through the path will be referred to as the output clock.

Since buffers and interconnects have finite bandwidth, the overall clock-distribution path is bandwidth-limited, and the output node requires a finite time to settle after each edge. This effect is shown in Figure 2.4, where the input clock is compared with the corresponding clock signal after propagation through a low-pass model of the clock path.

The impact of this finite settling time on jitter propagation depends on the time avail-

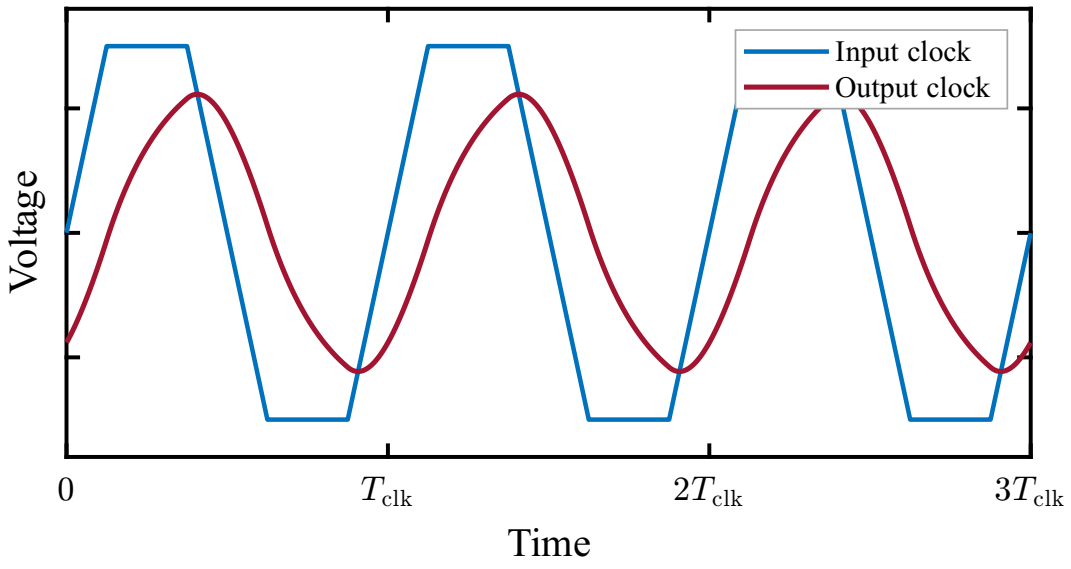


Figure 2.4: Effect of a bandwidth-limited channel on the propagated clock signal. The input clock is compared with the corresponding output clock after in-channel propagation.

able before the following clock transition. If the settling time is significantly smaller than the clock half-period, the clock signal reaches its final value before the following transition occurs. In this case, the effect of one edge is mostly independent of the previous one, and the input timing error is transferred without significant amplification. Conversely, if the clock signal does not have enough time to settle, the transition starts from a voltage level that still depends on the previous edge.

This implies that a timing displacement applied to one edge can modify the initial condition of the following transition, since the output signal may not have fully settled before the next edge occurs. This mechanism is illustrated in Figure 2.5, where the output clock, obtained after applying a timing displacement to one input edge, is compared with the corresponding nominal clock signal.

As shown in Figure 2.5, a delayed input edge reduces the settling time available before the following opposite transition. The resulting change in the initial condition can shift the subsequent threshold crossings, so that a single input jitter event produces a sequence of output timing errors. This coupling between consecutive edges is the time-domain origin of jitter amplification in bandwidth-limited clock paths [5].

The sequence of output edge displacements produced by a single input jitter event can be represented by the jitter impulse response (JIR). The JIR describes the output timing displacement associated with each edge following the applied input jitter event. For the bandwidth-limited clock path considered here, incomplete settling causes the perturbation to affect both the corresponding output edge and the subsequent transitions. The resulting

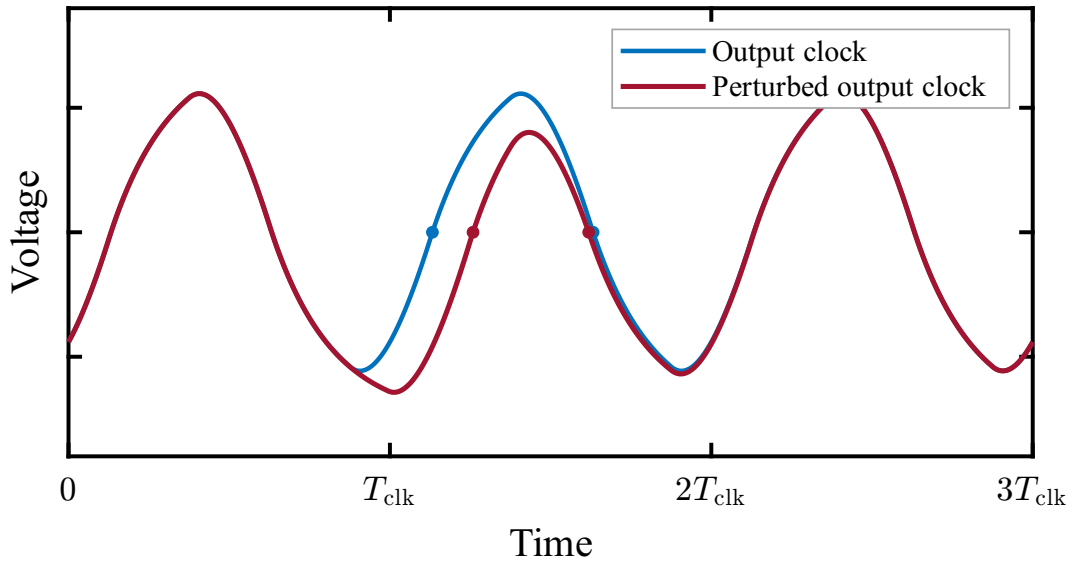
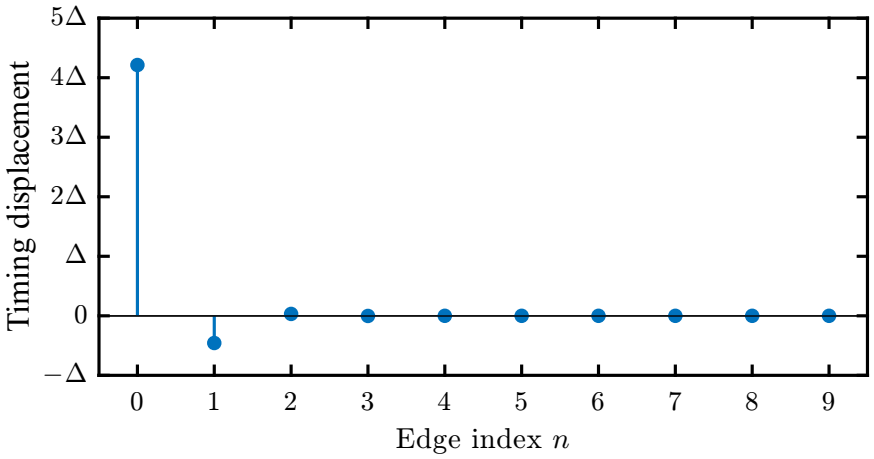
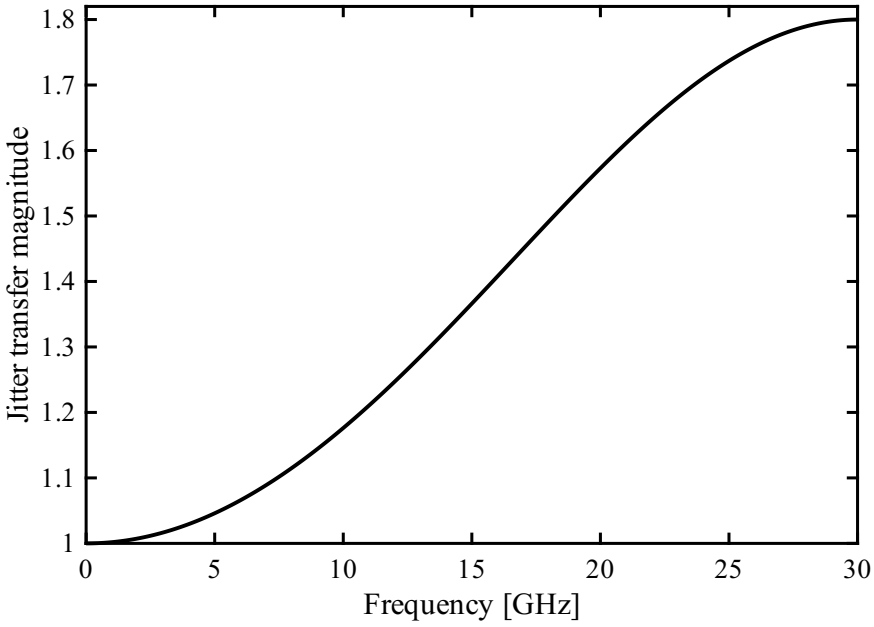


Figure 2.5: Propagation of a single input jitter event through a bandwidth-limited clock path. The nominal output clock is compared with the perturbed output clock obtained after applying a timing displacement to one input edge. The perturbation modifies the corresponding output transition and affects the settling conditions of the following transitions.



(a)



(b)

Figure 2.6: Time-domain (a) and frequency-domain (b) representations of the jitter-transfer behavior of the considered bandwidth-limited clock path. In (a), Δ denotes an arbitrary reference unit for timing displacement. The frequency-domain representation is obtained from the discrete-time Fourier transform of the JIR.

JIR is shown in Figure 2.6(a).

Since the JIR is a discrete sequence defined over the edge index, its discrete-time Fourier transform gives the jitter transfer function (JTF). The JTF represents the same jitter-transfer behavior in the frequency domain and describes how jitter components at different frequencies are transferred from the input to the output of the clock path. The JTF corresponding to the considered JIR is shown in Figure 2.6(b) [5, 18, 20].

A fundamental aspect is that the jitter transfer characteristics may differ from the voltage transfer behavior of the clock path. In this respect, a bandwidth-limited clock element may have a low-pass voltage response, while its jitter transfer response can exhibit

high-pass behavior [18]. Indeed, on one side, low-frequency timing variations are mainly transferred as slow phase shifts, whereas high-frequency edge-to-edge timing variations interact with the incomplete settling between consecutive transitions and can therefore be amplified [18].

Therefore, jitter amplification is closely related to bandwidth limitation. Indeed, the finite bandwidth of buffers and interconnects makes the clock signal depend on previous transitions. When the available settling time is not sufficient, the clock path adds an additional edge displacement at the output. As a result, high-frequency jitter components can be amplified even though the voltage transfer response is low-pass.

By assuming the distribution network as noiseless, a scalar jitter amplification factor A_J can be defined as the ratio between the output rms jitter $J_{\text{out,rms}}$ and the input rms jitter $J_{\text{in,rms}}$:

$$A_J = \frac{J_{\text{out,rms}}}{J_{\text{in,rms}}}. \tag{2.11}$$

When $A_J > 1$, the clock path amplifies the input jitter. When $A_J < 1$, the path attenuates the input jitter. This ratio is useful for comparing different clock distribution solutions under the same input jitter conditions, while the JTF provides a more detailed description of how jitter components at different frequencies are affected by the clock distribution network.

Jitter amplification can also be studied from the viewpoint of phase modulation. In this respect, a jittered clock can be represented as a phase-modulated clock signal. When this clock signal propagates through a lossy channel, the clock fundamental and the jitter-related sidebands are not attenuated by the same amount. The resulting imbalance modifies the phase modulation at the output and can increase the observed timing jitter [21]. This interpretation is consistent with the time-domain view based on incomplete settling: both describe how a non-ideal, bandwidth-limited propagation path can modify the timing perturbation carried by the clock.

In the present thesis, the time-domain interpretation is particularly useful because the analyzed networks are composed of CMOS buffers, interconnect segments, and tapping loads. These elements have an impact on the bandwidth and settling behavior of the propagated clock signal. Therefore, the clock distribution networks considered in the following chapters are evaluated not only in terms of voltage swing, attenuation, delay, and power consumption, but also in terms of jitter amplification features.

2.5 Transmission-Line Modeling for Clock Distribution

From a general perspective, when the physical length of an interconnect is much smaller than the wavelength at the frequency of interest, the interconnect can be modeled using a lumped-circuit approximation rather than be considered as a distributed element. In

this simplified condition, the phase variation along the wire is small, and voltage and current can be considered approximately uniform along its length, thus making the discrete approach a good approximation of the actual system behavior.

This approximation becomes less accurate as the operating frequency increases, since the wavelength associated with the signal becomes shorter compared to the physical length of the interconnect. From a general perspective, for a signal propagating with phase velocity v_p, the wavelength is defined as

$$\lambda = \frac{v_p}{f}. \tag{2.12}$$

For a fixed physical length l, the electrical length θ of the interconnect increases with frequency, according to

$$\theta = \beta l = \frac{2\pi l}{\lambda} = \frac{2\pi f l}{v_p}. \tag{2.13}$$

where β is the phase constant.

As the electrical length of the interconnect increases, the lumped-circuit approximation used in ordinary circuit theory becomes less accurate. In this condition, voltage and current can no longer be treated as uniform along the whole interconnect, but must be described as functions of position and time, $v(z,t)$ and $i(z,t)$ respectively. Following transmission-line theory, the interconnect is therefore treated as a distributed-parameter network, whose local electrical behavior is described through quantities defined per unit length [22]. This description is the basis for deriving the propagation constant and the characteristic impedance of the line, defined as

$$\gamma = \alpha + j\beta = \sqrt{(R + j\omega L)\,(G + j\omega C)} \tag{2.14}$$

and

$$Z_0 = \sqrt{\frac{R + j\omega L}{G + j\omega C}}, \tag{2.15}$$

which are then used to describe attenuation, phase propagation, and the relation between voltage and current waves along the interconnect.

The following subsections introduce the transmission-line model used to describe the clock link at circuit level, assuming a distributed model approach. Indeed, starting from the distributed per-unit-length formulation, the propagation constant and the characteristic impedance of the line are defined. The same distributed description is then related to an equivalent lumped-section representation, which is useful for circuit-level interpretation. Finally, the input impedance of a terminated line is introduced in order to evaluate how the interconnect and the receiving load determine the impedance seen by the clock-distribution driver, for both the non-resonant and resonant networks considered in the

case study under analysis.

2.5.1 Transmission-line theory

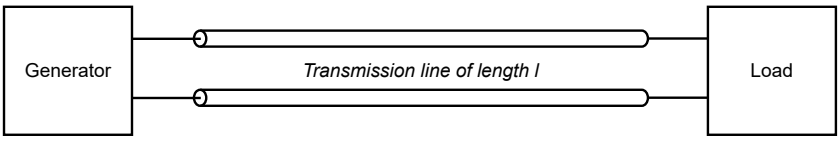
In the transmission-line formulation, the interconnect is represented as a distributed circuit rather than as a single lumped element. The voltage and current are therefore functions of the propagation coordinate z, denoted in the time domain as $v(z,t)$ and $i(z,t)$. A simplified view of a transmission line is shown in Figure 2.7, together with an equivalent circuit of an infinitesimal line section [22].

In this respect, in the elementary transmission-line model, a uniform line can be represented locally by per-unit-length series and shunt parameters. The series resistance R, inductance L, shunt conductance G, and capacitance C are defined per unit length. The series resistance R and inductance L account for the distributed series impedance of the line, while the shunt conductance G and capacitance C account for the distributed shunt admittance.

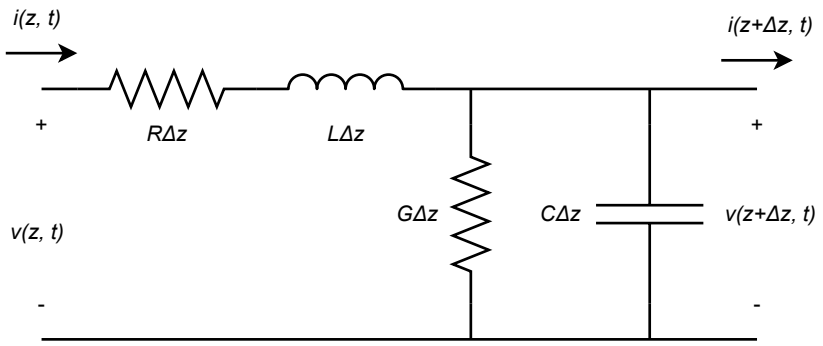
By applying Kirchhoff's voltage and current laws to the infinitesimal line section shown in Figure 2.7(b), the telegrapher equations, in the time domain, can be expressed as

$$\frac{\partial v(z,t)}{\partial z} = -Ri(z,t) - L\frac{\partial i(z,t)}{\partial t}, \tag{2.16}$$

$$\frac{\partial i(z,t)}{\partial z} = -Gv(z,t) - C\frac{\partial v(z,t)}{\partial t}. \tag{2.17}$$



(a)



(b)

Figure 2.7: Transmission-line representation of an interconnect. (a) Complete transmission line of length l, terminated by a generic load. (b) Equivalent circuit of an infinitesimal line section of length Δz, represented by distributed series and shunt elements.

For sinusoidal steady-state analysis, voltage and current can be represented by phasors $V(z)$ and $I(z)$. Therefore, the telegrapher equations can be written as

$$\frac{dV(z)}{dz} = -\left(R + j\omega L\right) I(z), \tag{2.18}$$

$$\frac{dI(z)}{dz} = -\left(G + j\omega C\right) V(z). \tag{2.19}$$

Combining these two first-order equations gives the wave equations for voltage and current:

$$\frac{d^2V(z)}{dz^2} - \gamma^2 V(z) = 0, \tag{2.20}$$

$$\frac{d^2I(z)}{dz^2} - \gamma^2 I(z) = 0. \tag{2.21}$$

The complex propagation constant γ, defined in Eq. (2.14), has real part α, which represents the attenuation constant, and imaginary part β, which represents the phase constant. The attenuation constant describes the exponential reduction of wave amplitude along the line, whereas the phase constant describes the phase variation per unit length.

The characteristic impedance is the voltage-to-current ratio of a single traveling wave. For a general lossy line, it is given by

$$Z_0 = \sqrt{\frac{R + j\omega L}{G + j\omega C}}. \tag{2.22}$$

Therefore, Z_0 is not determined only by the ohmic resistance of the interconnect, but by the distributed series impedance and shunt admittance of the line. In the lossy case, it can be complex and frequency dependent [22].

With these definitions, the voltage and current along the line can be written as the superposition of a forward-traveling wave and a reflected wave:

$$V(z) = V_0^+ e^{-\gamma z} + V_0^- e^{\gamma z}, \tag{2.23}$$

$$I(z) = \frac{V_0^+}{Z_0} e^{-\gamma z} - \frac{V_0^-}{Z_0} e^{\gamma z}. \tag{2.24}$$

where coefficient V_0^+ represents the amplitude of the wave traveling in the $+z$ direction, while V_0^- represents the amplitude of the wave traveling in the $-z$ direction. A reflected component appears when the line is terminated by an impedance different from Z_0.

In the ideal lossless case, $R = 0$ and $G = 0$. The propagation constant in Eq. (2.14) and the characteristic impedance in Eq. (2.15) respectively reduce to

$$\gamma = j\beta = j\omega\sqrt{LC}. \tag{2.25}$$

$$Z_0 = \sqrt{\frac{L}{C}}. \tag{2.26}$$

The corresponding phase velocity is

$$v_p = \frac{\omega}{\beta} = \frac{1}{\sqrt{LC}}. \tag{2.27}$$

Equations (2.25)–(2.27) show that, in the lossless case, propagation is determined by the distributed inductance and capacitance of the line.

In real on-chip interconnects, losses cannot generally be neglected. The series resistance R contributes to signal attenuation through conductor losses, while the shunt conductance G accounts for dielectric losses in the surrounding medium. As a result, the attenuation constant α becomes nonzero, causing the clock amplitude to decrease as the signal propagates along the interconnect. At the same time, the phase constant β determines the phase variation introduced by the line.

In the context of clock distribution, the attenuation constant α, the phase constant β, and the characteristic impedance Z_0 are the main parameters used to describe interconnect behavior. The attenuation constant is associated with swing reduction, the phase constant with propagation delay, and the characteristic impedance with loading, mismatch, and reflections. These parameters provide the basis for both the equivalent-circuit interpretation of the interconnect and the impedance-transformation analysis introduced in the following subsections.

2.5.2 Equivalent circuit representation of the interconnect

The purpose of this subsection is to introduce a compact lumped π-equivalent model that provides a circuit-level interpretation of the interconnect behavior. This model is not intended to describe the continuous variation of voltage and current along the line, but to represent the dominant parasitic effects introduced by the interconnect in a simple equivalent circuit.

Consistent with this purpose, the interconnect is modeled by the single-ended lumped π-equivalent circuit shown in Figure 2.8. This representation is used as a compact approximation to identify the main parasitic contributions of the interconnect, namely conductor loss, self-inductance, and capacitive loading. Its accuracy is expected to improve when the interconnect is electrically short over the frequency range of interest.

In a single-ended interconnect, the series branch includes the resistance and self-inductance associated with the conductor path. The resistance accounts for conductor loss, while the self-inductance represents the magnetic energy storage associated with the current flowing along the line. The shunt capacitors represent the capacitance from the signal conductor to the substrate and to nearby conductive structures.

In integrated clock distribution, clock lines switch continuously at high frequency and

can couple noise to nearby logic and circuit blocks. For this reason, they are often routed with ground or shielding metal structures targeting reduced coupling towards adjacent circuits. In the equivalent circuit, the shunt capacitance therefore includes the capacitive coupling toward the substrate, ground conductors, shields, and nearby metal layers. As a result, the π-equivalent model gives a compact circuit interpretation of the series loss, self-inductance, and capacitive loading introduced by the interconnect.

The single-ended model can be extended to a differential version, as shown in Figure 2.9. In this respect, each conductor can be associated with its own single-ended equivalent section, but the two conductors are also coupled to each other.

The capacitive coupling arises because the two signal conductors are separated by a dielectric. When a voltage difference exists between them, an electric field is established across this dielectric, and the two conductors behave as a coupled capacitive structure. In the proposed equivalent circuit, this effect is represented in the equivalent circuit by the coupling capacitance C_c. In addition, the current flowing in each conductor generates a

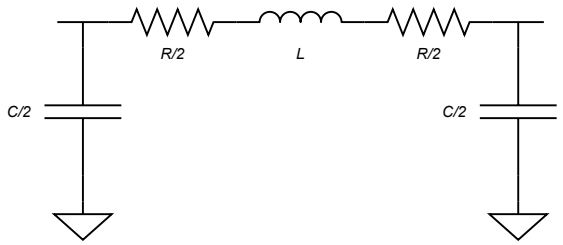


Figure 2.8: Proposed single-ended lumped π-equivalent model for the interconnect.

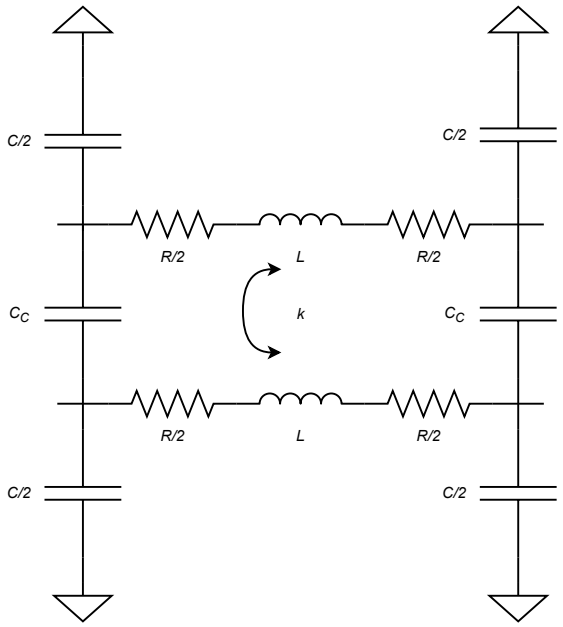


Figure 2.9: Lumped equivalent representation of a differential interconnect section, including capacitive coupling C_c and magnetic coupling represented by the coupling coefficient k, related to the mutual inductance M.

magnetic field. Part of the magnetic flux associated with one conductor can link the other conductor, producing magnetic coupling between the two lines. This effect is represented by the mutual inductance M, or equivalently by a coupling coefficient k. These coupling terms depend on the interconnect geometry, including line spacing, conductor width, metal layer, and the surrounding return or shielding structures. As a result, the layout directly affects the effective capacitance and inductance seen in differential operation, and therefore the differential impedance and propagation behavior of the line.

As long as the physical length of the interconnect remains small compared to the wavelength over the frequency range of interest, the proposed equivalent circuit can be used as a first-order approximation of the line behavior, including the dominant parasitic and coupling terms.

2.5.3 Impedance transformation through a transmission line

At the input of the transmission line, the driver sees an impedance Z_{in}, which depends on the load impedance Z_L, the characteristic impedance Z_0, and the electrical length of the interconnect.

For a lossless transmission line of characteristic impedance Z_0, phase constant β, and length l, terminated by a load Z_L, the input impedance is [22]

$$Z_{\text{in}}(l,\omega) = Z_0 \frac{Z_L(\omega) + jZ_0 \tan(\beta l)}{Z_0 + jZ_L(\omega)\tan(\beta l)}. \tag{2.28}$$

For a lossy transmission line, the propagation constant is a complex quantity, as defined in Eq. (2.14). In this case, the input impedance of a terminated line is given by [22]

$$Z_{\text{in}}(l,\omega) = Z_0 \frac{Z_L(\omega) + Z_0 \tanh(\gamma l)}{Z_0 + Z_L(\omega)\tanh(\gamma l)}. \tag{2.29}$$

Equations (2.28) and (2.29) describe the same impedance transformation in two different cases. The lossless form in Eq. (2.28) describes the case in which the propagation constant is purely imaginary, so the line introduces phase accumulation without attenuation. Conversely, the lossy form in Eq. (2.29) also includes attenuation during propagation, since the propagation constant has a real part α. Therefore, the lossy expression accounts for both phase shift and amplitude decay along the interconnect. In the following discussion, the lossless expression is used to show the basic transformation mechanism for a capacitive CMOS load. Moving from the lossless to the lossy case, the main difference is that the lossless approximation does not capture the real part of the transformed impedance. In fact, starting from an ideal capacitive load, the lossless expression gives a purely imaginary input impedance, whereas the lossy expression also introduces a real component associated with line losses.

In the clock-distribution case considered in this thesis work, the input of the receiving inverter is approximated, as a first-order model, by an effective capacitance C_L. If this capacitance were connected directly to the driver, its impedance would be

$$Z_L(\omega) = \frac{1}{j\omega C_L}. \tag{2.30}$$

Substituting this capacitive load into Eq. (2.28) gives

$$Z_{\text{in}}(l,\omega) = jZ_0 \frac{Z_0 \tan(\beta l) - \frac{1}{\omega C_L}}{Z_0 + \frac{\tan(\beta l)}{\omega C_L}}. \tag{2.31}$$

Equation (2.31) shows that the impedance associated with the physical capacitance C_L is transformed by the interconnect. In the lossless case, the transformed impedance remains purely imaginary, because neither the line nor the load introduces dissipation. However, its reactance differs from the reactance of the isolated capacitance C_L, since it would depend on frequency, on the characteristic impedance Z_0, and on the electrical length βl.

After the substitution, the input impedance obtained in the lossless case is still purely imaginary. Therefore, the transformed load can be interpreted from the sign of its input reactance: a negative reactance corresponds to capacitive behavior, whereas a positive reactance corresponds to inductive behavior. Since this sign depends on the term $\tan(\beta l)$, it varies periodically with the electrical length. As a result, the same physical capacitive load can appear at the input of the transmission line either as a different capacitive reactance or as an inductive reactance. This means that the finite-length interconnect modifies the impedance presented to the driver.

If the transformed impedance is capacitive, an equivalent capacitance seen at the input of the line can be defined from the input reactance. Writing

$$Z_{\text{in}}(\omega) = jX_{\text{in}}(\omega), \tag{2.32}$$

with $X_{\text{in}} < 0$, the equivalent input capacitance is

$$C_{\text{in,eq}}(\omega) = -\frac{1}{\omega X_{\text{in}}(\omega)}. \tag{2.33}$$

This quantity is not the physical capacitance of the receiving inverter. It is the capacitance that would produce the same input reactance at the considered frequency if it were connected directly to the driver.

If $X_{\text{in}} > 0$, the transformed load is instead inductive, and it can be described through an equivalent inductance

$$L_{\text{in,eq}}(\omega) = \frac{X_{\text{in}}(\omega)}{\omega}. \tag{2.34}$$

This case is relevant because a capacitive termination can therefore present an inductive input reactance after transmission-line transformation, depending on the operating frequency and line length.

For a lossy line, the general expression in (2.29) must be used. In that case, even if the load itself is purely capacitive, the input impedance can have both real and imaginary parts:

$$Z_{\text{in}}(\omega) = R_{\text{in}}(\omega) + jX_{\text{in}}(\omega). \tag{2.35}$$

The imaginary part describes the transformed reactive behavior of the load, while the real part is associated with losses introduced by the interconnect. Therefore, the lossless expression is useful to understand the reactive transformation mechanism, whereas the lossy expression gives the more complete impedance seen by the driver in a real clock link.

This impedance-transformation interpretation provides the link between the transmission-line model and the clock-distribution circuits analyzed in the following chapters. In the non-resonant case, it helps identify the actual load that the CMOS driver has to drive. In the resonant case, it helps determine the inductance value required to resonate at the target operating frequency.

Chapter 3

Non-resonant clock distribution

3.1 Analysis of the reference architecture

This chapter investigates the non-resonant implementation of a high-speed clock distribution network, aiming to identify its key characteristics, limitations and trade-offs and to establish an effective design methodology for different optimization objectives and implementation scenarios. To this end, the implementation shown in Figure 3.1 was developed and used throughout the study, with the clock signal provided by an ideal source and applied to a CMOS inverter that drives the transmission line.

However, in most practical implementations, a buffer stage alone cannot efficiently drive interconnects ranging from hundreds of μm up to a few mm, due to its limited driving capability and to the attenuation and loading effects introduced by the transmission line,

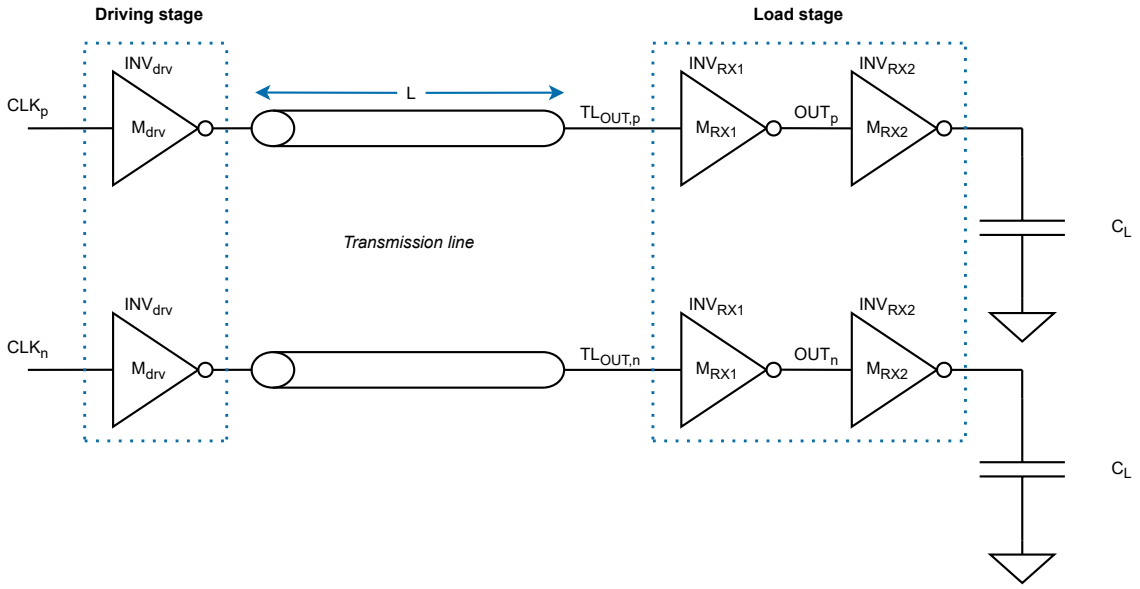


Figure 3.1: Circuit configuration used to evaluate the main design trends of a non-resonant CMOS clock distribution network. The driver multiplicity M_{drv} and the transmission-line length L are varied, while the load stage is kept fixed.

as discussed in Chapter 2.

For this reason, several approaches can be adopted, as will be discussed throughout this chapter. Among them, an additional CMOS inverter chain may be inserted between the driving stage and the transmission line. This chain is typically composed of cascaded inverter stages with progressively increasing size, commonly referred to as a tapered chain (see Section 3.2.1).

However, increasing the inverter size is not always beneficial. Large stages reduce the output resistance and improve the ability to charge and discharge the interconnect, but they also introduce large parasitic capacitances while increasing the current consumption of the distribution path. Therefore, simply maximizing the size of the transmission-line driver $\mathrm{INV_{DRV}}$ does not lead to an optimal non-resonant implementation. Indeed, the driving stages should be large enough to provide the required signal voltage swing and transition time, while avoiding unnecessary capacitance and power consumption.

In this respect, the simplified setup of Figure 3.1 includes a driving CMOS inverter stage, denoted as $\mathrm{INV_{DRV}}$, with multiplicity M_{drv} and a transmission line of length L. The effects of these two parameters are analyzed in the following subsections.

A pair of inverters, denoted as $\mathrm{INV_{RX1}}$ and $\mathrm{INV_{RX2}}$, with multiplicities M_{RX1} and M_{RX2}, respectively, is used as the load stage. The $\mathrm{INV_{RX1}}$ inverter provides clock regeneration at the transmission-line output and drives the $\mathrm{INV_{RX2}}$ stage, which represents the reference CMOS load assumed for this work. A lumped capacitance C_L is connected after $\mathrm{INV_{RX2}}$ as a simplified load.

In the following analysis, unless otherwise specified, the performance are evaluated differentially at the node pair OUT_p and OUT_n. Moreover, all simulation results presented in this chapter are obtained under typical simulation conditions reported in Table 2.1, with a supply voltage of 750 mV and a die temperature of 85 °C.

In the following subsections, the performance of the clock distribution network is first analyzed by investigating its dependence on the design parameters M_{drv} and L, representing the driving-stage size and transmission-line length, respectively.

3.1.1 Driver sizing analysis

The first evaluation is performed on the circuit configuration shown in Figure 3.1 by varying the driver multiplicity M_{drv} for three different transmission-line lengths (100 μm, 300 μm, and 500 μm) and a fixed load condition. This allows to evaluate the effect of drive strength on the output rise time, peak-to-peak voltage, jitter, jitter amplification, and current consumption for each selected line length.

As M_{drv} increases, the effective output resistance of the driver decreases and the current that can be delivered during a clock transition increases. As expected, at a fixed load, the clock signal rise time would decrease approximately with the driver resistance,

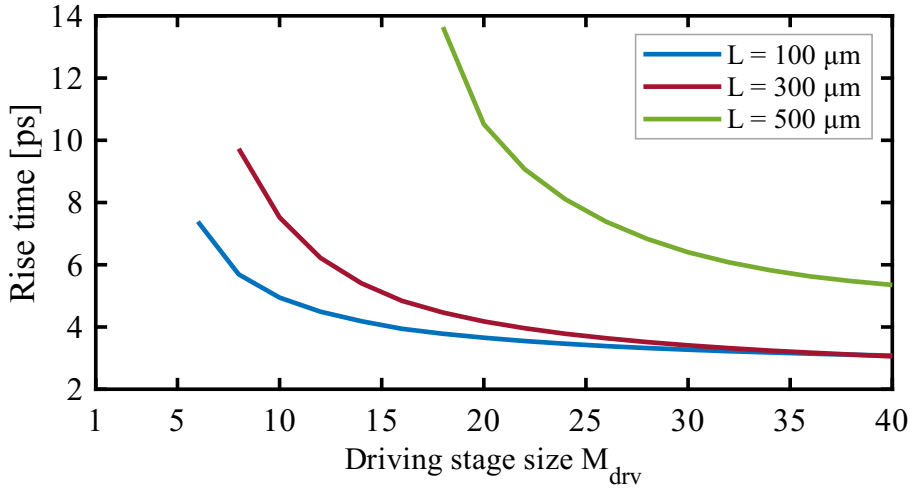


Figure 3.2: Output rise time as a function of the CMOS driving stage size M_{drv} for 100 μm, 300 μm, and 500 μm transmission line lengths.

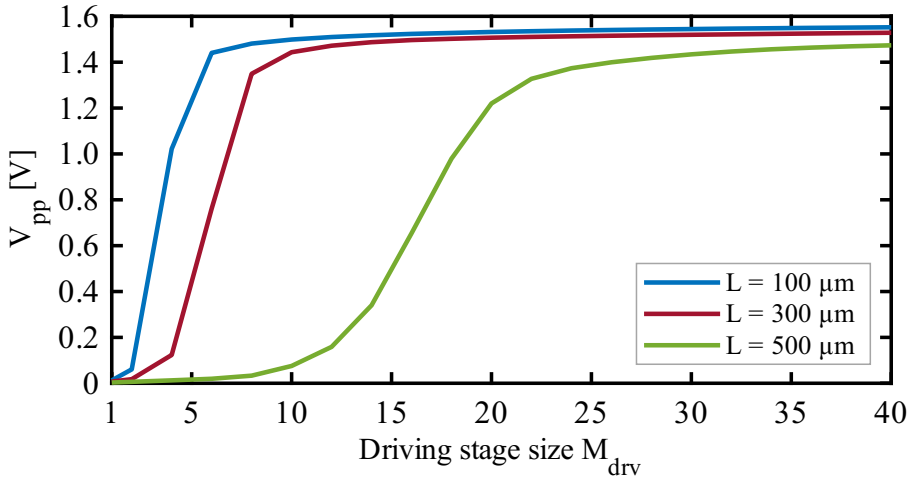


Figure 3.3: Differential output peak-to-peak voltage evaluated between OUT_p and OUT_n as a function of the CMOS driving stage size M_{drv} for 100 μm, 300 μm, and 500 μm transmission line lengths.

and therefore roughly with the inverse of the driver size.

This trend is shown in Figure 3.2, assuming a 30 GHz clock input signal at the CLK_p and CLK_n nodes with 50% duty cycle and 5 ps rise and fall times. The rise time reduction is strong for small driver sizes and progressively reduces as the driver size increases. This produces a hyperbola-like behavior, with the curves tending to flatten at large driver sizes, since the reduction in output resistance is progressively offset by the increase in the driver's parasitic capacitance.

This behavior is not an ideal inverse law because the capacitance involved in the transition is not constant while INV_{DRV} is resized. Indeed, while at a defined line length the interconnect contribution is given, the output parasitic capacitance of the driver increases with the driver size. This effect is more relevant for short lines, where the interconnect capacitance is small and the INV_{DRV} capacitance has a large relative weight. As a result,

part of the benefit obtained by reducing the driver resistance is offset by the increase in its capacitance, leading to progressively diminishing improvements in rise time. Thus, for long transmission lines, the interconnect capacitance represents a dominant fraction of the total load, making the approximation of a constant load capacitance more accurate.

While the rise time describes the speed of the transition, the differential peak-to-peak voltage V_{pp} is used to evaluate the settling of the output voltage $\mathrm{OUT_p} - \mathrm{OUT_n}$ within the available half-period of the clock, since incomplete settling prevents the signal from reaching its expected voltage levels, resulting in a reduced peak-to-peak voltage.

For very small driver sizes, the voltage at the input of $\mathrm{INV_{RX1}}$ does not move sufficiently beyond the inverter switching threshold within the available half-period. As a consequence, the voltage at the output of $\mathrm{INV_{RX1}}$ remains limited and the evaluated V_{pp} is small. For intermediate values of M_{drv}, the signal at the input of $\mathrm{INV_{RX1}}$ enters the switching region more effectively. Around this point, $\mathrm{INV_{RX1}}$ behaves as a high-gain amplifying stage, so a small variation at its input produces a large change in the output. Finally, for large driver sizes, $\mathrm{INV_{RX1}}$ is already driven close to complete switching between the available voltage levels. The output swing therefore approaches its maximum value, and further driver upsizing only produces a limited improvement.

This behavior is shared across all the considered test line lengths, as shown in Figure 3.3. As the interconnect length increases, a progressively large driver size is required to reach a defined V_{pp} value, consistently with the impedance-transformation effect discussed in Chapter 2. Indeed, depending on the line length, the capacitance of the load stage is transformed into a different equivalent load seen by the driver. In the considered range, long lines correspond to a large effective capacitive burden, so a strong driver is required to obtain the same output swing as for short transmission lines.

Regarding the jitter performance, they are correlated with the voltage-swing results discussed above. Indeed, at small sizes of the driving stage, the differential signal between $\mathrm{TL_{OUT,p}}$ and $\mathrm{TL_{OUT,n}}$ exhibits a strongly reduced voltage swing, especially for the longest line. This condition limits the switching of $\mathrm{INV_{RX1}}$, thus resulting in incomplete voltage transitions of the differential signal between $\mathrm{OUT_p}$ and $\mathrm{OUT_n}$. Consequently, the extracted crossing time is poorly defined. The large jitter values observed in this region can therefore serve as a guideline for driver sizing, helping to ensure that the output jitter meets the required specification.

To further investigate the jitter behavior, the jitter simulated on the differential waveform between OUT_p and OUT_n is compared with that simulated on the differential waveform between $TL_{\mathrm{OUT},p}$ and $TL_{\mathrm{OUT},n}$.

At the transmission-line output, the jitter decreases monotonically as the driver size increases, as provided in Figure 3.4. The jitter reduction is significant while the signal at the transmission-line output is not yet sufficiently settled within the available half-period. As M_{drv} increases, the jitter curves progressively flatten, indicating diminishing

benefits from further driver upsizing. This behavior is consistent with the rise-time trend: a strong driver leads to a fast signal transition into the line, so the differential signal between $\mathrm{TL_{OUT,p}}$ and $\mathrm{TL_{OUT,n}}$ exhibits a steep transition at its zero crossing.

As shown in Figure 3.4, following the initial steep decrease in jitter, the random jitter simulated on the differential signal between $TL_{\mathrm{OUT},p}$ and $TL_{\mathrm{OUT},n}$ approximately follows a $1/M_{\mathrm{drv}}$ dependence, so that doubling the driver multiplicity approximately halves the jitter.

The jitter simulated on the differential signal between OUT_p and OUT_n does not decrease in the same way as the one simulated between $TL_{\mathrm{OUT},p}$ and $TL_{\mathrm{OUT},n}$, as can be seen by comparing Figure 3.4 and Figure 3.5. This difference arises because the signal at the transmission-line output is subsequently processed by $\mathrm{INV_{RX1}}$ of the load stage. For small values of M_{drv}, the differential signal between $TL_{\mathrm{OUT},p}$ and $TL_{\mathrm{OUT},n}$ is still poorly established, so the switching of $\mathrm{INV_{RX1}}$ and consequently the crossing between OUT_p and OUT_n are strongly affected by the degraded input transition. For this reason, the jitter between OUT_p and OUT_n is much larger than that observed between $TL_{\mathrm{OUT},p}$ and $TL_{\mathrm{OUT},n}$.

This explains why the final output jitter eventually tends to a limiting value instead of following the reduction observed at the transmission-line output. Since the load chain is not resized throughout this analysis, increasing the driver size mainly improves the signal propagating through the line, but it does not remove the jitter contribution of the following fixed stages. The variation of the driver size therefore highlights a trade-off: sufficient drive strength is required to move from a degraded signal to a low-jitter clock, while excessive upsizing keeps increasing current consumption with progressively weaker benefits on the final output jitter.

Jitter amplification was introduced in Chapter 2 as the frequency-dependent ratio be-

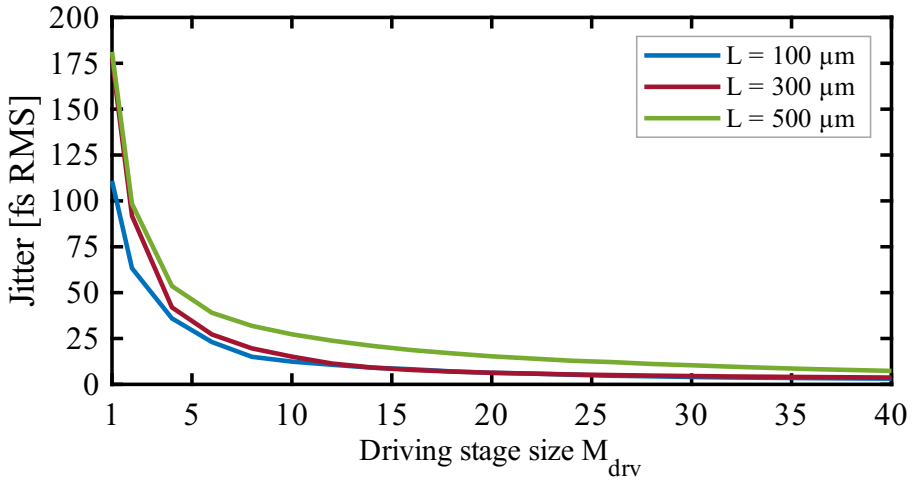


Figure 3.4: Random jitter simulated for the differential signal between $TL_{\mathrm{OUT},p}$ and $TL_{\mathrm{OUT},n}$ as a function of the CMOS driving stage size M_{drv} for 100 μm, 300 μm, and 500 μm transmission line lengths.

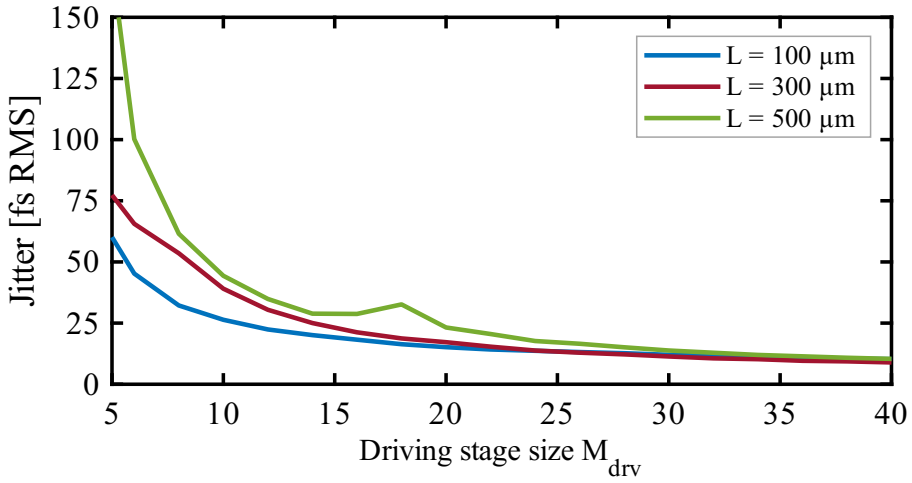


Figure 3.5: Random jitter simulated for the differential signal between OUT_p and OUT_n as a function of the CMOS driving stage size M_{drv} for 100 μm, 300 μm, and 500 μm transmission line lengths.

tween the output jitter and the applied input jitter. The jitter-transfer function therefore describes how the circuit transfers input jitter depending on its frequency. An example of the complete frequency-dependent jitter-transfer function is shown later in Figure 3.24, for the single-segment implementation. Since the design specification requires that the input jitter is not amplified at any frequency within the analyzed range, the maximum value of the jitter-transfer function is considered to evaluate the effect of driver sizing. In this respect, Figure 3.6 reports this value as a function of the driver size for the three considered transmission-line lengths. Values at or below unity indicate that the requirement is satisfied over the analyzed frequency range, whereas values above unity indicate jitter amplification at least at one frequency.

Jitter-amplification data are reported only when the output rise time can also be extracted. Otherwise, the output edge is not sufficiently well defined, and the corresponding crossing-based jitter-amplification value is not considered meaningful.

Within the valid range, the maximum jitter amplification is high for the small driver sizes and decreases as the driver size is increased. At the low end of this range, the output waveform does not fully settle before the following clock edge. Consequently, a timing perturbation affecting one edge changes the waveform state from which the subsequent transition starts and can therefore affect the timing of the following edges, as discussed in Chapter 2.

As the driver size increases, the waveform settles more rapidly within the available half-period, reducing the influence of previous edges. Consequently, the maximum jitter amplification approaches unity. Over the considered driver-size range, the jitter amplification does not show a clear reduction below unity. Therefore, increasing the driver size reduces the tendency of the path to amplify the input jitter, but does not provide significant jitter attenuation in this non-resonant configuration.

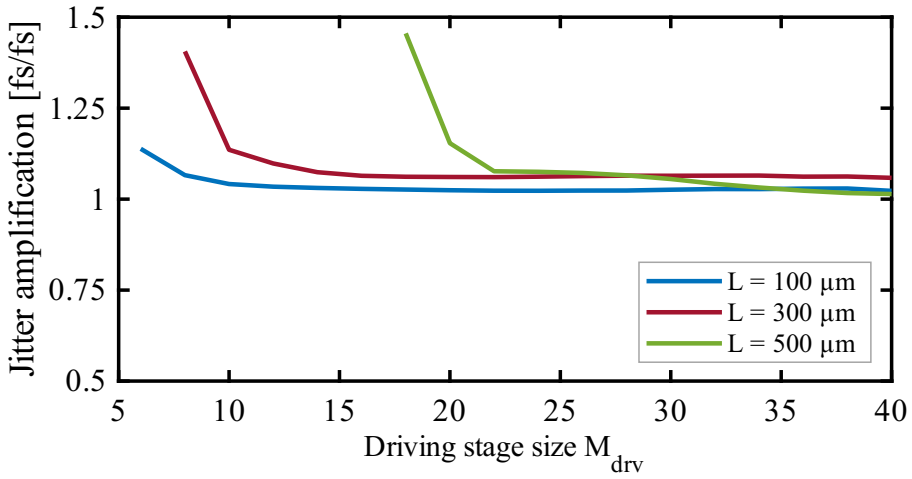


Figure 3.6: Jitter amplification as a function of the CMOS driving stage size M_{drv} for 100 μm, 300 μm, and 500 μm transmission lines. For each configuration, the plotted value corresponds to the highest value of the corresponding frequency-dependent jitter amplification. The plot includes only the driver sizes for which a valid output rise time is obtained.

After the voltage, timing, jitter, and jitter-amplification results, the average current consumption is evaluated as the main cost associated with driver size scaling, as shown in Figure 3.7. As expected, the current consumption increases with the driving stage size. The current-consumption trend also helps identifying whether an additional driver strength still provides a performance benefit or not.

The curves in Figure 3.7 exhibit two regions with different slopes. Before the change in slope, the current consumption is mainly associated with the equivalent capacitive load presented at the transmission-line input. Within the considered length range, increasing the transmission-line length shifts this change toward increasing driver multiplicities. This behavior is related to the impedance transformation introduced by the transmission line, which determines the effective load seen by the driver.

At large driver size, the parasitic capacitance introduced by the driver itself becomes comparable to, and progressively more significant than, the equivalent transmission-line capacitance. Therefore, further upsizing mainly increases the current required to charge and discharge the additional parasitic capacitance introduced by the driver upsizing itself, while the corresponding performance improvement progressively diminishes. Since these capacitances scale approximately with the driver multiplicity, the current exhibits a more regular, nearly linear dependence on driver size.

From a general viewpoint, increasing the CMOS driver size might be beneficial only up to a certain point. A large driver improves the output signal by reducing the rise time, increasing the voltage swing, lowering the output jitter, and reducing the maximum jitter amplification toward values close to unity. However, these improvements become progressively weak since further increases in M_{drv} produce limited reductions in rise time

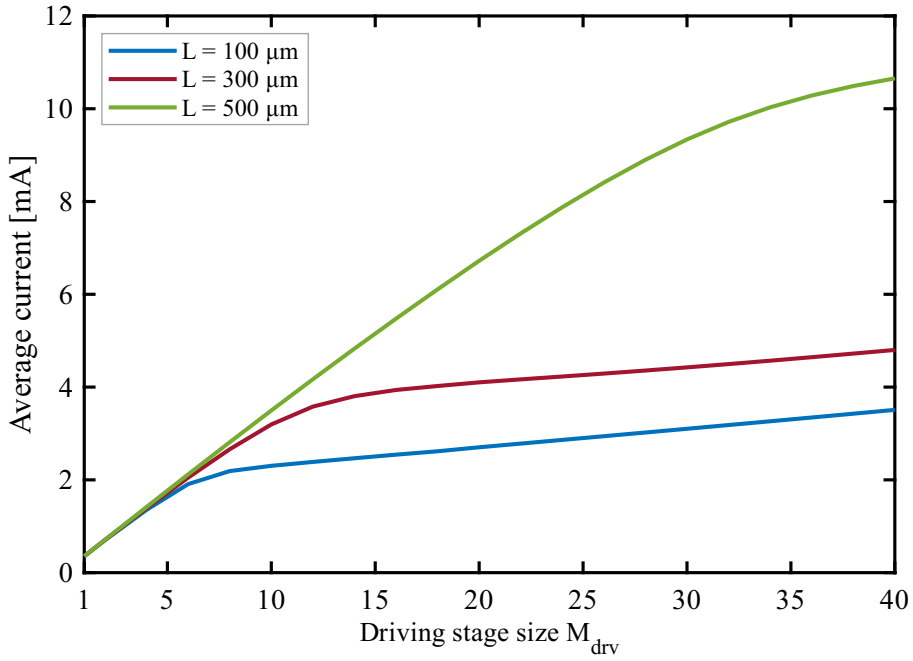


Figure 3.7: Average current consumption as a function of the CMOS driver size for 100 μm, 300 μm, and 500 μm transmission line lengths.

and output jitter, while the voltage swing approaches its limiting value and the average current consumption continues to increase. Therefore, the optimal driver size is a trade-off between clock-signal improvement and current consumption. In the next section, the impact of the transmission-line length on the clock performance is investigated, assuming a fixed driver configuration to isolate the effect of the interconnect itself.

3.1.2 Effect of transmission-line length

The transmission-line length is varied while keeping the CMOS driving chain sizing fixed. The analysis considers line lengths up to 500 μm. The transmission-line component used in these simulations relies on a lumped-element representation, which is adopted to investigate the main design trends over relatively short interconnect lengths. The complete 1 mm interconnect required by the design is analyzed later using a more accurate electromagnetic model extracted with EMX (electromagnetic extraction, an electromagnetic extraction tool from Cadence), including the relevant parasitic effects.

The effect of the line length on the output rise time is examined. Figure 3.8 provides the simulated rise time of the clock signal at different values of the transmission-line length. In this respect, for short transmission-line lengths, the output rise time remains almost unchanged. Indeed, in this region, the equivalent capacitance introduced by the line is still small compared with the capacitance associated with the driver output node. As a result, increasing the line length only weakly changes the total load driven by the CMOS chain, and the output transition time does not vary significantly.

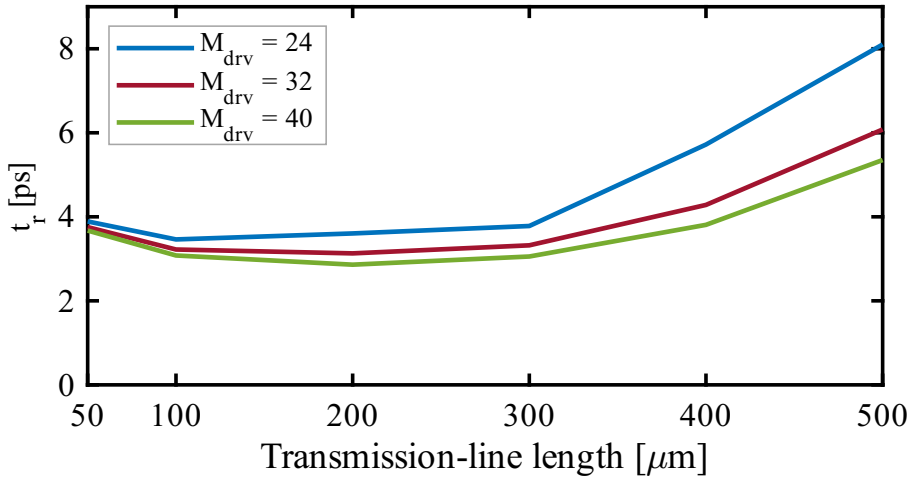


Figure 3.8: Output rise time evaluated between OUT_p and OUT_n as a function of the transmission-line length for different CMOS driver sizes.

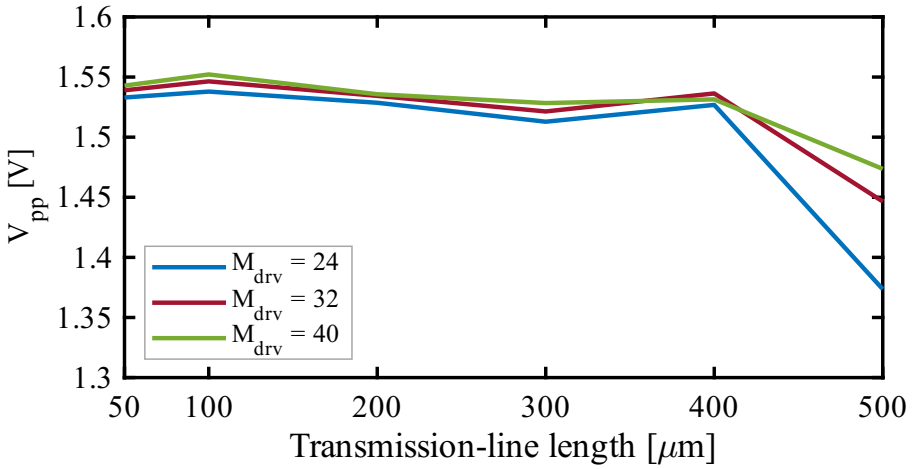


Figure 3.9: Peak-to-peak output voltage evaluated between OUT_p and OUT_n as a function of the transmission-line length for different CMOS driver sizes.

This effect mainly affects small driver sizes, where the available drive current is low. Conversely, large drivers are less affected because they can drive the increased equivalent capacitance more effectively, preserving a fast transition over the considered length range.

The peak-to-peak voltage indicates whether the waveform has enough time to settle close to its high and low levels within each clock period. Figure 3.9 shows that, as the transmission-line length is initially increased, the output swing remains almost constant, indicating that the differential output signal maintains a voltage swing close to the full differential swing.

As the transmission line becomes long, the effective load increases and the settling becomes slow. The waveform then does not reach the same high and low values within the available time, so the peak-to-peak voltage decreases. This reduction mainly affects small driver sizes, while large drivers preserve a high swing due to their higher current drive capability than small drivers.

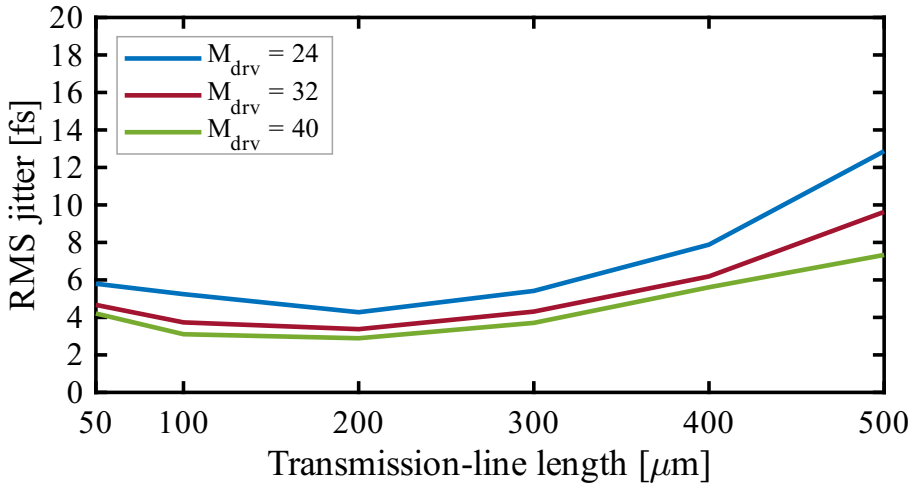


Figure 3.10: RMS jitter simulated on the differential signal $TL_{\mathrm{OUT},p}$ - $TL_{\mathrm{OUT},n}$ as a function of the transmission-line length for different CMOS driver sizes.

The jitter at the transmission-line output is reported in Figure 3.10. As the transmission-line length increases, so does the jitter, showing that the signal transmitted through the line progressively degrades. The effect is also strongly dependent on the driver size: large drivers produce low jitter at the line output, because they drive the transmission line with sharp transitions and reduce the sensitivity of the crossing time to noise and waveform distortion.

The jitter simulated at the $\mathrm{OUT_p} - \mathrm{OUT_n}$ differential node is reported in Figure 3.11. After the transmission line, $\mathrm{INV_{RX1}}$ partially restores the signal, but it also introduces its own noise and jitter. Therefore, the jitter of $\mathrm{OUT_p} - \mathrm{OUT_n}$ includes both the jitter already present on $\mathrm{TL_{OUT,p}} - \mathrm{TL_{OUT,n}}$ and the additional contribution introduced by $\mathrm{INV_{RX1}}$.

The same general tendencies observed on $TL_{\mathrm{OUT},p}$ - $TL_{\mathrm{OUT},n}$ are preserved on OUT_p - OUT_n: increasing the transmission-line length leads to an increase in jitter, while increasing the driver size reduces it. This confirms that the line-length increase affects the RMS jitter of the clock between OUT_p and OUT_n, and that increasing the driver size helps limit this degradation.

The maximum jitter amplification for the same driver sizes used in the previous plots is reported in Figure 3.12. For these configurations, the jitter amplification shows a limited increase with the transmission-line length over the considered range. Overall, increasing the line length is associated with a degradation of the transmitted signal, an increase in the absolute RMS jitter, and a limited increase in jitter amplification.

Since the jitter amplification shows a limited increase with the transmission-line length for the driver sizes considered in Figure 3.12, the same analysis is repeated for small driver sizes, as shown in Figure 3.13. For these small driver sizes, the jitter amplification exhibits a clear increase with the transmission-line length, whereas the increase is limited for the

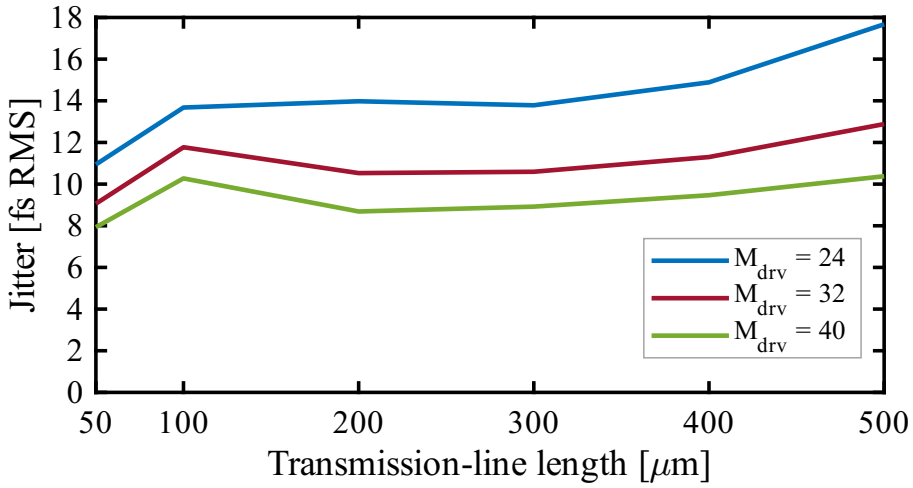


Figure 3.11: Final output RMS jitter simulated on the differential signal OUT_p - OUT_n as a function of the transmission-line length for different CMOS driver sizes.

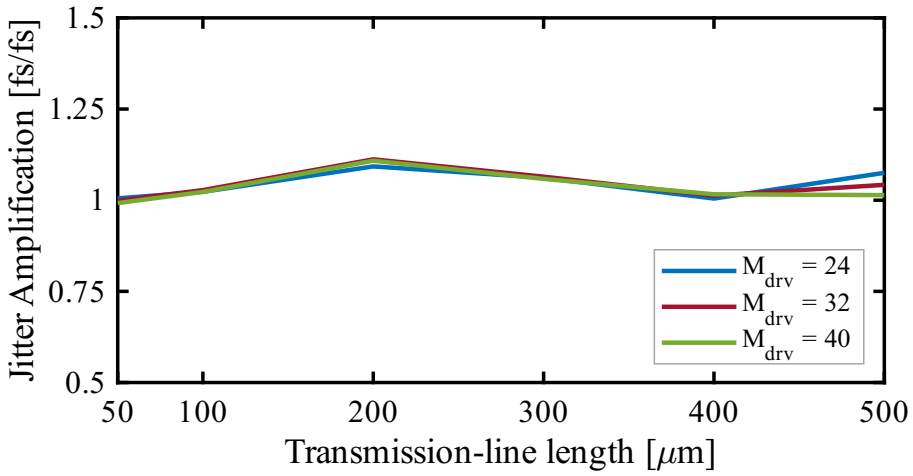


Figure 3.12: Jitter amplification as a function of the transmission-line length for different CMOS driver sizes. Each point represents the maximum magnitude of the jitter-transfer function over the analyzed frequency range.

driver sizes considered in Figure 3.12.

The driver current is reported in Figure 3.14. For the considered driver sizes, the current increases with the transmission-line length. This trend is consistent with the increase of the equivalent load seen by the CMOS driver at the input of the line: as this load increases, more charge has to be provided during each clock transition than for small loads, resulting in a high average current.

As the transmission-line length increases, the difference between the considered driver sizes becomes evident. Small drivers progressively approach their current-delivery limit and cannot supply the additional current required by the increasing load, whereas large drivers provide a high current-drive capability. Therefore, increasing the driver size helps preserve the signal integrity and reduce the timing degradation, but it also increases the current consumption, especially for long transmission lines and high load values.

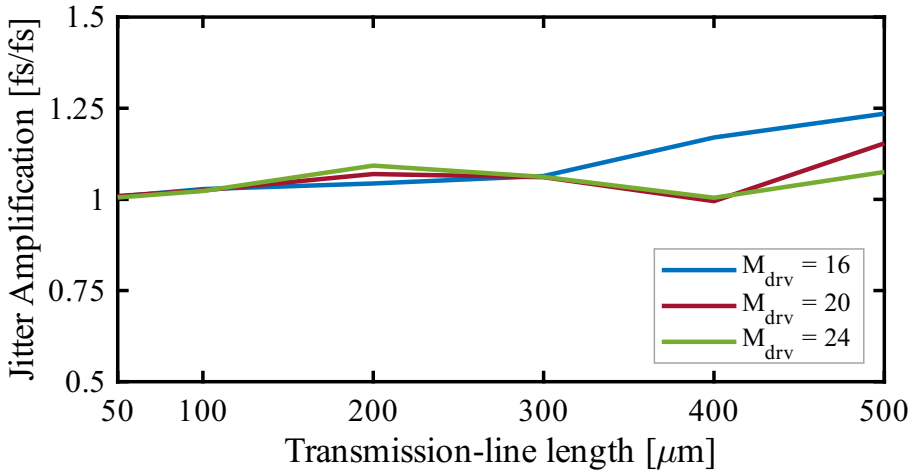


Figure 3.13: Jitter amplification as a function of the transmission-line length for small CMOS driver sizes. Each point represents the maximum magnitude of the jitter-transfer function over the analyzed frequency range.

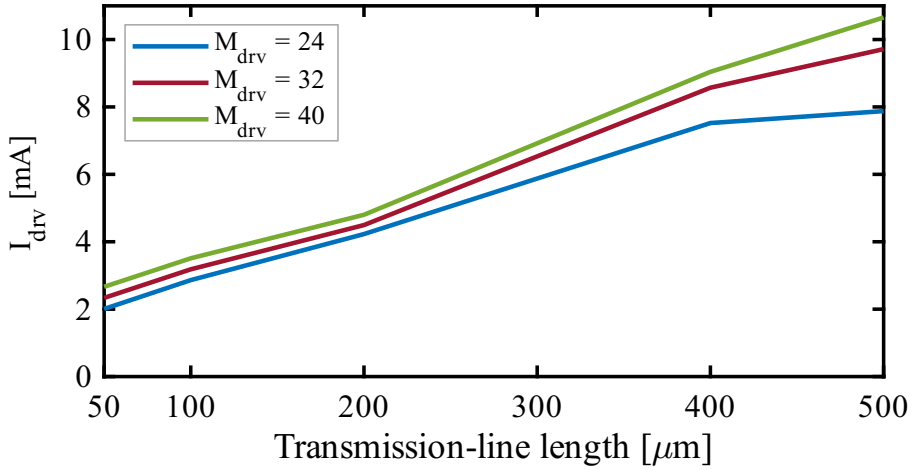


Figure 3.14: Driver current as a function of the transmission-line length for different CMOS driver sizes.

3.2 Design methodology for non-resonant clock distribution

Previous sections describe the operating principle of the non-resonant CMOS distribution network and show how changes in driver size and transmission-line length affect rise time, voltage swing, jitter, jitter amplification, and current consumption.

This section provides a methodology for sizing the entire non-resonant CMOS chain. The methodology is developed in two stages. First, a sizing criterion is derived for a generic tapered CMOS inverter chain. Then, the effect of the transmission line is included by replacing the physical receiver capacitance with the equivalent load seen at the line input. Combining these two steps provides a procedure for sizing the complete non-resonant CMOS chain.

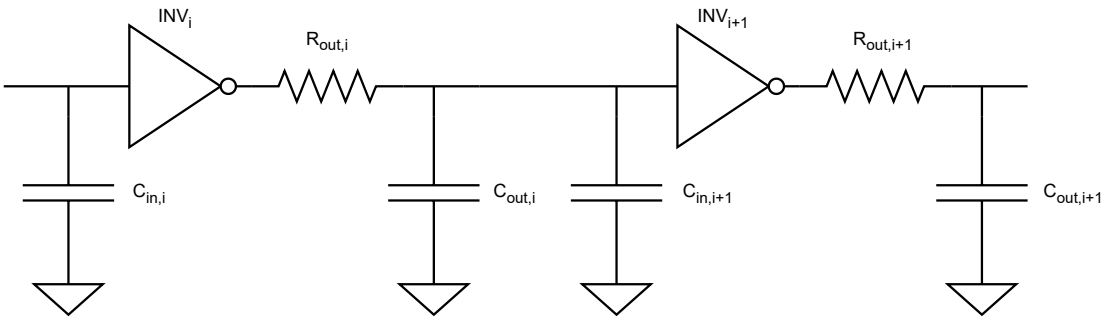


Figure 3.15: First-order RC model of two consecutive stages in a tapered CMOS inverter chain.

3.2.1 Tapered-chain design

The first step of the methodology concerns the sizing of the tapered CMOS inverter chain. The purpose of the tapered chain is to connect the source stage to the line driver through a sequence of intermediate inverter stages with progressively increasing drive strength.

A conventional tapered-chain design is often optimized for minimum propagation delay. In this work, however, the primary objective is to prevent jitter amplification rather than minimize propagation delay. Following the settling-time criterion proposed by [5], each stage must preserve a sufficiently fast transition so that the clock signal settles within one half of the clock period. Incomplete settling increases the dependence of each transition on the previous signal state and can therefore lead to jitter amplification. The sizing criterion is expressed in terms of the maximum time constant, or equivalently the minimum bandwidth, allowed for each stage. The objective is therefore to determine the maximum tapering ratio $r_{\max}$ that satisfies this bandwidth requirement at every interstage node.

At each interstage node shown in Figure 3.15, the output of the i-th inverter drives the input of the $(i+1)$-th inverter. The driving stage is modeled by an effective output resistance $R_{\text{out},i}$ and by its intrinsic output capacitance $C_{\text{out},i}$, which includes the parasitic capacitances connected to the transistor drains. The following inverter contributes its input capacitance $C_{\text{in},i+1}$, mainly associated with the gate capacitances of its NMOS and PMOS devices. The corresponding first-order RC model is shown in Figure 3.15.

Therefore, the total capacitance $C_{\text{node},i}$ connected to the node is approximated as

$$C_{\text{node},i} = C_{\text{out},i} + C_{\text{in},i+1}. \tag{3.1}$$

This capacitance is charged and discharged through the effective output resistance of the driving inverter $R_{\text{out},i}$, giving the stage time constant

$$\tau_i \simeq R_{\text{out},i} \left(C_{\text{out},i} + C_{\text{in},i+1} \right). \tag{3.2}$$

The same structure is repeated at every interstage node of the tapered chain, so the same first-order model can be applied to all stages.

An inverter stage with multiplicity M_i consists of M_i identical unit inverters connected in parallel. Let $R_{\text{out},0}$, $C_{\text{out},0}$, and $C_{\text{in},0}$ denote the output resistance, output capacitance, and input capacitance of the unit inverter, respectively.

Since the conducting paths of the unit inverters are connected in parallel, the effective output resistance of the i-th stage, $R_{\text{out},i}$, decreases with its multiplicity:

$$R_{\text{out},i} \simeq \frac{R_{\text{out},0}}{M_i}. \tag{3.3}$$

The drain parasitic capacitances of the unit inverters are connected to the same output node and add together. Therefore, at first order the output capacitance $C_{\text{out},i}$ should increase linearly with the multiplicity:

$$C_{\text{out},i} \simeq M_i C_{\text{out},0}. \tag{3.4}$$

The same rule applies to the input capacitance. Since the gate capacitances of the unit inverters are connected to the same input node, their contributions add together, thus:

$$C_{\text{in},i} \simeq M_i C_{\text{in},0}. \tag{3.5}$$

Therefore, the time constant of the i-th interstage node can be rewritten by substituting the scaling relations given in (3.3), (3.4), and (3.5):

$$\tau_i \simeq \frac{R_{\text{out},0}}{M_i}\left(M_i C_{\text{out},0} + M_{i+1} C_{\text{in},0}\right) = R_{\text{out},0}\left(C_{\text{out},0} + \frac{M_{i+1}}{M_i} C_{\text{in},0}\right). \tag{3.6}$$

The tapering ratio between the i-th inverter and the following stage is defined as

$$r_i = \frac{M_{i+1}}{M_i}. \tag{3.7}$$

Therefore, the stage time constant τ_i in (3.6) can also be expressed as

$$\tau_i \simeq R_{\text{out},0}\left(C_{\text{out},0} + r_i C_{\text{in},0}\right). \tag{3.8}$$

The absolute multiplicity M_i therefore cancels out, and the time constant depends on the tapering ratio between consecutive stages. To preserve the required bandwidth at every interstage node, the same time constant must be maintained throughout the chain. Since the time constant depends on r_i, this condition is obtained by using a constant tapering ratio,

$$r_i = r, \tag{3.9}$$

for all consecutive stages. Under the dominant-pole approximation[1], this result leads to approximately the same bandwidth at every node of the tapered chain.

In a realistic scenario, the maximum admissible time constant is selected according to the settling requirement at each interstage node. Therefore, the following derivation is kept general and expressed in terms of the maximum admissible time constant.

Under this first-order approximation, the stage bandwidth can be related to the rise time by considering the response of the equivalent RC circuit to an input step. With the output initially at $v(0) = 0$, the voltage follows $v(t) = V_f(1 - e^{-t/\tau_i})$, where V_f is the final voltage. Denoting by t_{20} and t_{80} the instants at which the output reaches 20% and 80% of its final value, respectively, the 20–80% rise time is given by

$$t_{r,20\text{–}80} = t_{80} - t_{20} = -\tau_i \ln(0.2) + \tau_i \ln(0.8) = \tau_i \ln\left(\frac{0.8}{0.2}\right) = \tau_i \ln(4). \tag{3.10}$$

Therefore, for a first-order response, the maximum admissible time constant can be related to the maximum allowed 20–80% rise time $t_{r,\max}$ as

$$\tau_{\max} = \frac{t_{r,\max}}{\ln(4)} \tag{3.11}$$

This approach follows the settling-time criterion discussed by Mo et al. [5]: the clock waveform should have sufficient time to settle before the following transition. Accordingly, the maximum admissible time constant must be selected in relation to the clock half-period, so that the residual effect of previous transitions is sufficiently small. This limits the interaction between successive edges and the resulting jitter amplification.

The maximum time constant corresponds to the minimum required bandwidth

$$B_{\text{req}} = \frac{1}{2\pi\tau_{\max}}. \tag{3.12}$$

The bandwidth requirement can therefore be expressed as

$$\tau_i \leq \tau_{\max}. \tag{3.13}$$

Since the stage time constant increases with the tapering ratio, the maximum admissible value of r is obtained by imposing $\tau_i = \tau_{\max}$. Using (3.8), the maximum tapering ratio is

$$r_{\max} = \frac{\tau_{\max}/R_{\text{out},0} - C_{\text{out},0}}{C_{\text{in},0}}. \tag{3.14}$$

The value $r_{\max}$ represents the largest admissible size ratio between two consecutive inverter stages. Each stage is implemented by connecting an integer number of unit inverters

[1]Each interstage node is modeled by the output resistance and output capacitance of the driving inverter, together with the gate capacitance of the following inverter. This first-order RC approximation describes the node through a single dominant pole.

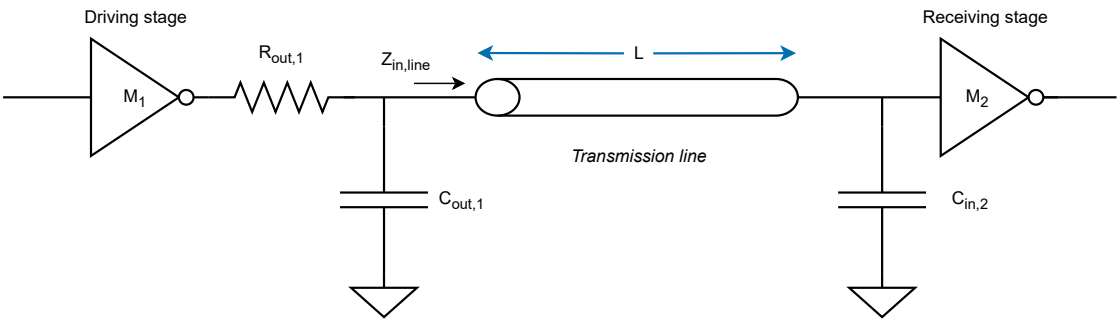


Figure 3.16: Circuit model of the driving and receiving inverter stages connected through a transmission line of length L.

in parallel. The intermediate stages must therefore be assigned integer multiplicities while ensuring that this limit is not exceeded.

The resulting chain provides an initial sizing estimate based on the first-order model. Transistor-level simulations may be necessary to refine the stage dimensions and assess whether the design requirements are met.

3.2.2 Driver sizing with transmission-line loading

The sizing criterion derived in the previous subsection can now be extended to the case in which a transmission line is inserted between the driving and receiving inverters. Let M_1 and M_2 denote their respective multiplicities, and let $C_{\text{in},1}$ and $C_{\text{in},2}$ denote their input capacitances. Since the input capacitance of an inverter scales approximately with its multiplicity,

$$\frac{C_{\text{in},1}}{C_{\text{in},2}} = \frac{M_1}{M_2}. \tag{3.15}$$

The corresponding circuit model is shown in Figure 3.16. The receiving inverter presents the physical input capacitance $C_{\text{in},2}$ at the far end of the transmission line. The line transforms this capacitance into an equivalent load seen at its input, expressed as

$$C_{\text{in,line}} = k(L)C_{\text{in},2}, \tag{3.16}$$

where $k(L)$ is the capacitance-transformation factor introduced by a transmission line of length L. In the limit case $L = 0$, no transformation occurs and $k(0) = 1$.

For a tapering ratio r, the equivalent load presented at the transmission-line input must satisfy

$$C_{\text{in,line}} = rC_{\text{in},1}. \tag{3.17}$$

Combining 3.16 and 3.17 gives

$$C_{\text{in},1} = \frac{C_{\text{in,line}}}{r} = \frac{k(L)}{r} C_{\text{in},2}. \tag{3.18}$$

Therefore, by using 3.15, the same relation can be written as

$$M_1 = \frac{k(L)}{r} M_2. \tag{3.19}$$

As a consequence, the dependence of the required driver size on the transmission-line length can be expressed through the sizing function

$$M_1 = g(L) M_2, \qquad g(L) = \frac{k(L)}{r}. \tag{3.20}$$

To obtain the minimum driver multiplicity compatible with the bandwidth requirement, the largest admissible tapering ratio is selected, $r = r_{\max}$. Therefore, from 3.19,

$$M_1 = \frac{k(L)}{r_{\max}} M_2. \tag{3.21}$$

For $L = 0$, the transmission line is absent and $k(0) = 1$. The output of the driving inverter is therefore directly connected to the input of the following inverter, giving

$$M_1 = \frac{M_2}{r_{\max}}, \qquad g(0) = \frac{1}{r_{\max}}. \tag{3.22}$$

Comparing (3.20) and (3.21) gives

$$g(L) = \frac{k(L)}{r_{\max}}. \tag{3.23}$$

The driver-sizing problem is therefore reduced to the determination of the capacitance-transformation factor $k(L)$ introduced by the transmission line.

The factor $k(L)$ is obtained by explicitly evaluating how the transmission line transforms the input capacitance of the receiving inverter. As a first-order approximation, the receiver is represented by its input capacitance $C_{\text{in},2}$, corresponding to the load impedance $Z_{\text{in},2}$ given by

$$Z_{\text{in},2}(\omega) = \frac{1}{j\omega C_{\text{in},2}}. \tag{3.24}$$

As discussed in Chapter 2, $Z_{\text{in},2}$ is transformed by the transmission line into an impedance seen at the line input. For the lossless transmission-line model, this impedance is

$$Z_{\text{in,line}}(L, \omega) = Z_0 \frac{Z_{\text{in},2}(\omega) + jZ_0 \tan(\beta L)}{Z_0 + jZ_{\text{in},2}(\omega) \tan(\beta L)}. \tag{3.25}$$

Conversely, for the lossy transmission-line model, the corresponding expression is

$$Z_{\text{in,line}}(L,\omega) = Z_0 \frac{Z_{\text{in,2}}(\omega) + Z_0 \tanh(\gamma L)}{Z_0 + Z_{\text{in,2}}(\omega) \tanh(\gamma L)}. \tag{3.26}$$

The transformed impedance $Z_{\text{in,line}}$ can equivalently be expressed in terms of the input admittance

$$Y_{\text{in,line}}(L,\omega) = \frac{1}{Z_{\text{in,line}}(L,\omega)}. \tag{3.27}$$

When the imaginary part of the input admittance is positive, the transformed load exhibits capacitive behavior. The corresponding parallel-equivalent capacitance is defined as

$$C_{\text{in,line}}(L,\omega) = \frac{\text{Im}\left\{Y_{\text{in,line}}(L,\omega)\right\}}{\omega}. \tag{3.28}$$

For a lossless transmission line, the transformed impedance is purely reactive. In this model, the transformed load is assumed to be capacitive, so the equivalent capacitance can be written as

$$C_{\text{in,line}}(L,\omega) = -\frac{1}{\omega \ \text{Im}\left\{Z_{\text{in,line}}(L,\omega)\right\}}. \tag{3.29}$$

The capacitance-transformation factor is then obtained as

$$k(L,\omega) = \frac{C_{\text{in,line}}(L,\omega)}{C_{\text{in,2}}}. \tag{3.30}$$

For $L = 0$, the transformed impedance coincides with the physical load impedance, so that $C_{\text{in,line}} = C_{\text{in,2}}$ and $k(0,\omega) = 1$, consistently with the directly connected case.

Although the transformation factor depends on both the transmission-line length and the operating frequency, the sizing procedure is performed at a fixed frequency. The frequency dependence is therefore omitted in the following, and the factor is simply denoted by $k(L)$.

To illustrate how the transmission line transforms the load impedance, Figure 3.17 shows the imaginary part of the transmission-line input impedance as a function of the line length at 30 GHz. At small line lengths, the input reactance is negative, indicating capacitive behavior. As the line length increases, the input reactance approaches zero, and then becomes positive, corresponding to a transition from capacitive to inductive behavior. At large lengths, the reactance turns back to negative values and the same trend repeats due to the periodic impedance transformation introduced by the transmission line.

Within the capacitive region of the impedance transformation, the transformed load can be represented by the equivalent input capacitance defined in (3.28). The proposed sizing model is restricted to this region and is not applied when the transformed load becomes inductive. Figure 3.18 shows this capacitance as a function of the transmission-line length up to 1 mm. The vertical axis is limited to preserve the readability of the

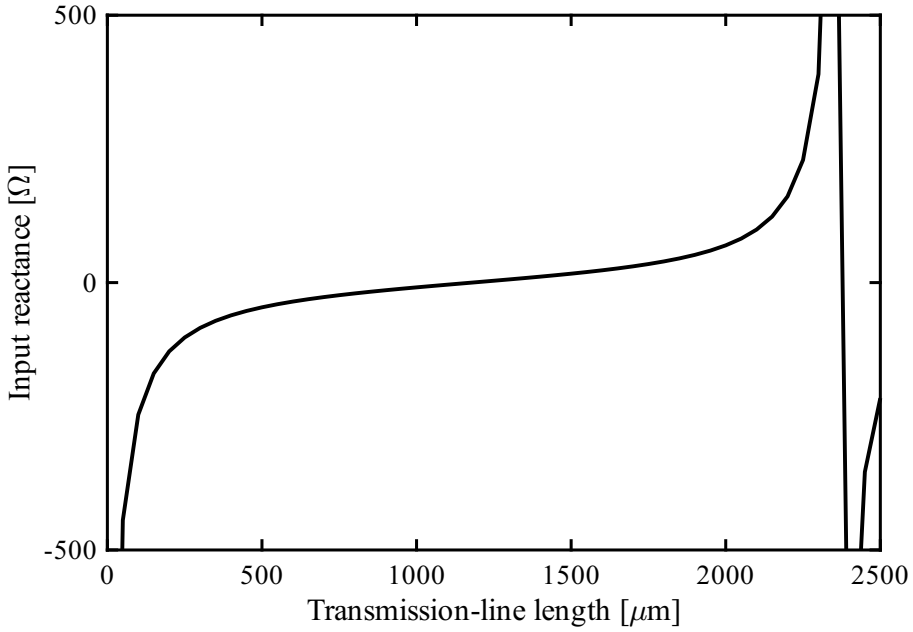


Figure 3.17: Input reactance as a function of the transmission-line length at 30 GHz.

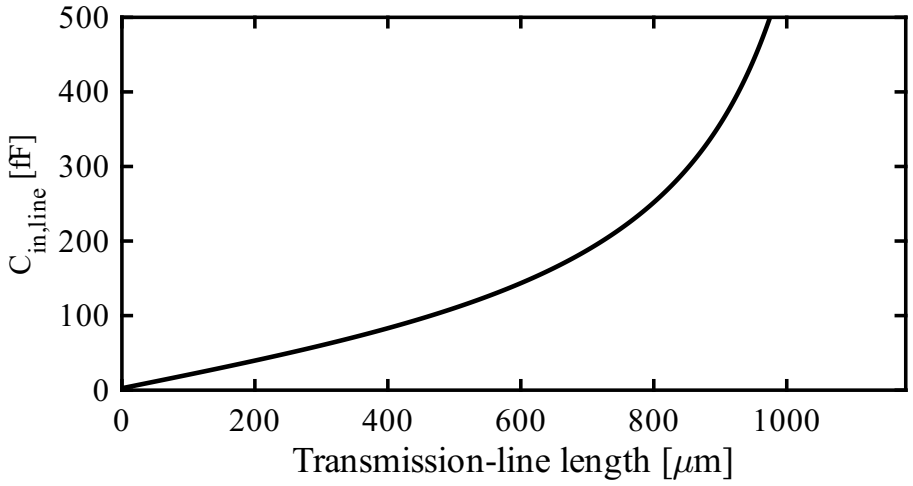


Figure 3.18: Equivalent input capacitance $C_{\text{in,line}}$ as a function of the transmission-line length at 30 GHz.

capacitive region, since some values have a large magnitude. At $L = 0$, the equivalent capacitance coincides with the physical input capacitance of the receiving inverter. As the line length increases, the equivalent capacitance increases, showing that the load seen at the transmission-line input appears to be progressively capacitive. The values beyond 500 μm are shown only to illustrate the impedance-transformation trend and are not used in the proposed sizing procedure.

Normalizing the equivalent capacitance with respect to the physical input capacitance of the receiving inverter gives the transformation factor $k(L)$, shown in Figure 3.19. By definition, $k(0) = 1$. Its increase with the transmission-line length quantifies how much larger the receiving load appears when observed from the line input.

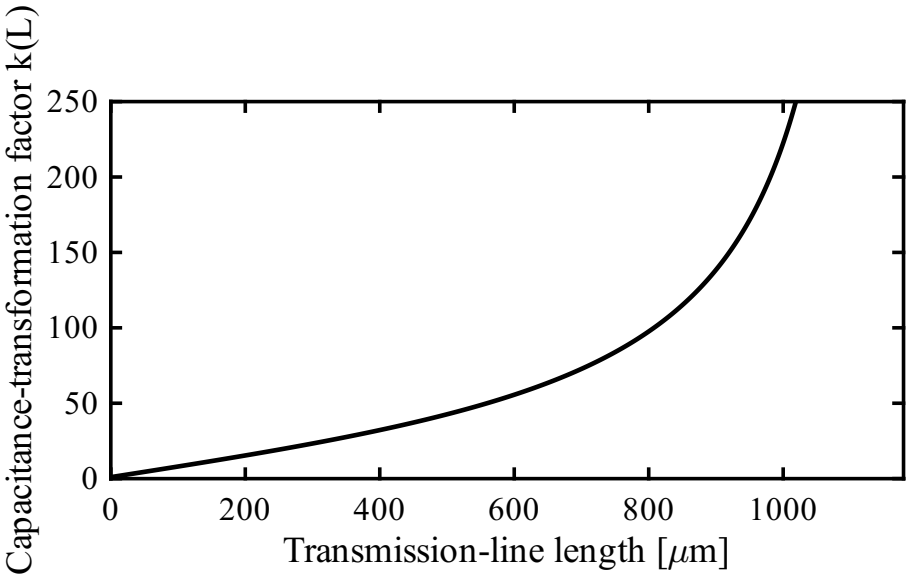


Figure 3.19: Capacitance-transformation factor $k(L)$ as a function of the transmission-line length at 30 GHz.

Once $k(L)$ has been determined, the minimum required driver multiplicity follows directly from (3.23). The transmission-line loading problem is therefore reduced to the tapered-chain sizing problem derived in the previous subsection.

3.2.3 Clock distribution network design procedure

The sizing criteria derived in the previous subsections can now be combined to design the complete non-resonant distribution network. The total transmission-line length can be implemented as a single section or divided into N_{seg} segments, where N_{seg} represents an architectural design parameter. For the segmented architecture considered here, equal segment lengths are adopted to obtain a modular and repeatable structure, with each transmission-line section driven by the same tapered CMOS inverter chain. The resulting modular structure is illustrated in Figure 3.20.

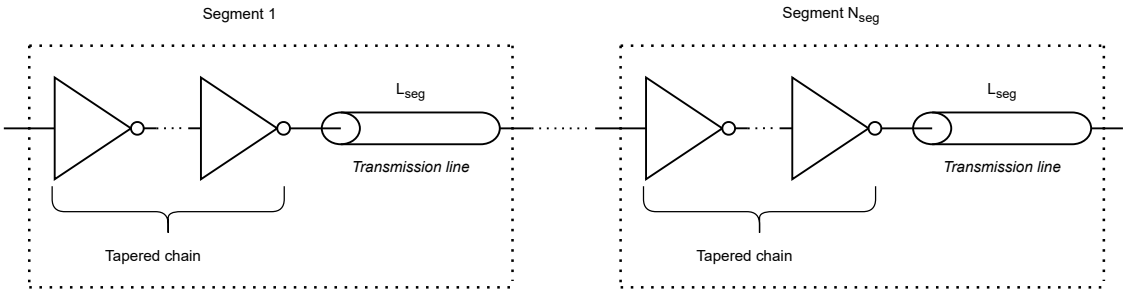


Figure 3.20: Modular clock-distribution structure with a tapered CMOS inverter chain and a transmission-line section of length L_{seg} in each segment. Intermediate segments are omitted for clarity.

The nominal segment length L_{seg} is

$$L_{\text{seg}} = \frac{L_{\text{tot}}}{N_{\text{seg}}}. \tag{3.31}$$

where L_{tot} represents the total transmission-line length and N_{seg} the number of segments. In the considered modular architecture, the source stage of the first segment has a multiplicity fixed by the system specifications. Each following segment starts from an inverter with the same multiplicity, connected to the output of the preceding transmission-line section. Under this assumption, the same initial sizing condition is restored at each segment, so that only one representative segment needs to be sized and the resulting driving chain can be repeated throughout the distribution.

For the selected value of L_{seg}, the load capacitance connected at the end of the segment is transformed toward the transmission-line input according to (3.16). The resulting equivalent capacitance is used to determine the required line-driver multiplicity. The intermediate inverter stages are then sized using the maximum tapering ratio $r_{\max}$ from (3.14), starting from the fixed source-stage multiplicity introduced above and progressively increasing up to the line driver.

The configuration in which the complete 1 mm transmission line is implemented as a single segment is also retained as a reference. Since this length lies outside the range adopted for the analytical sizing method, its line-driver multiplicity is selected through circuit-level simulations.

Changing the number of segments modifies both the load driven by each line driver and the total number of CMOS inverter stages along the clock path. Shorter segments reduce the local driving effort but require additional inverter chains, with a corresponding increase in current consumption, circuit area, device-noise contributions, and cumulative jitter transfer. The final architecture is therefore selected by comparing the complete simulated implementations.

3.3 Case Study: 1 mm Transmission Line

The design methodology developed in the previous section is now applied to the 1 mm clock distribution considered in this work. Three different implementations are evaluated: the complete transmission line driven as a single segment, a distribution divided into two segments, and a distribution divided into four segments. All of these configurations cover the same total length but differ in the size of the individual transmission-line sections and in the number of CMOS driving and regeneration blocks.

3.3.1 Single-segment implementation

The first configuration uses the complete 1 mm transmission line as a single segment. Its architecture and the multiplicities of the CMOS inverter stages are shown in Figure 3.21. The fixed source stage drives a tapered CMOS inverter chain, whose final and largest stage acts as the line driver. At the far end of the transmission line, a receiving inverter chain restores the clock waveform and drives the final load.

Since the transmission-line length exceeds 500 μm, the sizing method defined in the previous section cannot be reliably applied to determine the line-driver size. Extending the equivalent-capacitance approximation to the complete 1 mm line would provide an excessively large estimate of the required driver.

The line-driver multiplicity is therefore selected through transistor-level simulations. Once the driver size has been determined, the tapering stages connecting the fixed source inverter to the line driver are sized using the tapered-chain criterion derived in Subsection 3.2.1.

The differential peak-to-peak voltage and rise time evaluated between $\mathrm{OUT_p}$ and $\mathrm{OUT_n}$ are reported in Figure 3.22. Within the target operating range from 25 to 30 GHz, the output voltage decreases only slightly, while the rise time exhibits a limited increase. The differential peak-to-peak voltage slightly exceeds $2V_{DD}$ because of circuit parasitics producing a small transient overshoot at the inverter outputs. Since the peak-to-peak value is calculated from the maximum and minimum voltage levels, this overshoot affects the reported simulation results. Both quantities satisfy the specifications reported in Table 2.1, showing that the single-segment implementation meets the required differential output swing and rise time.

The output random jitter and average current consumption are reported in Figure 3.23.

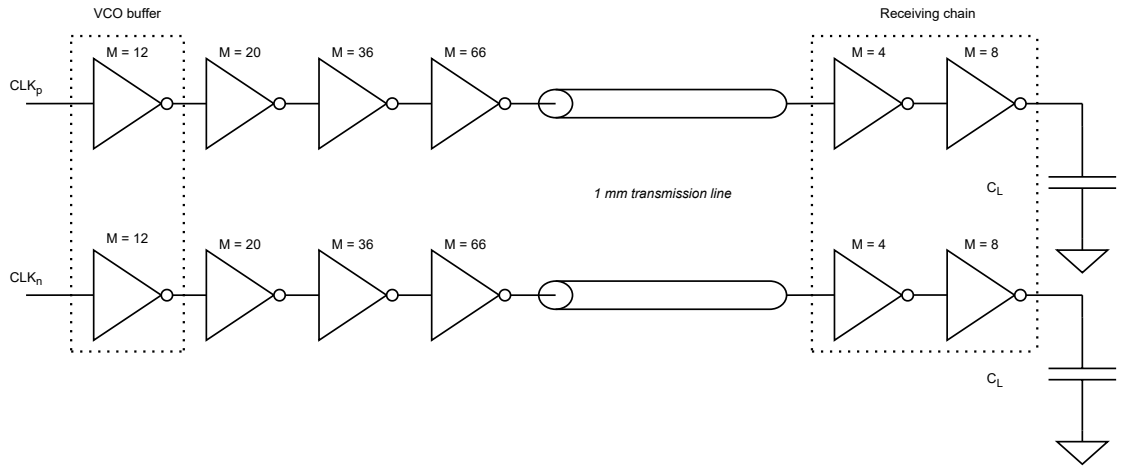
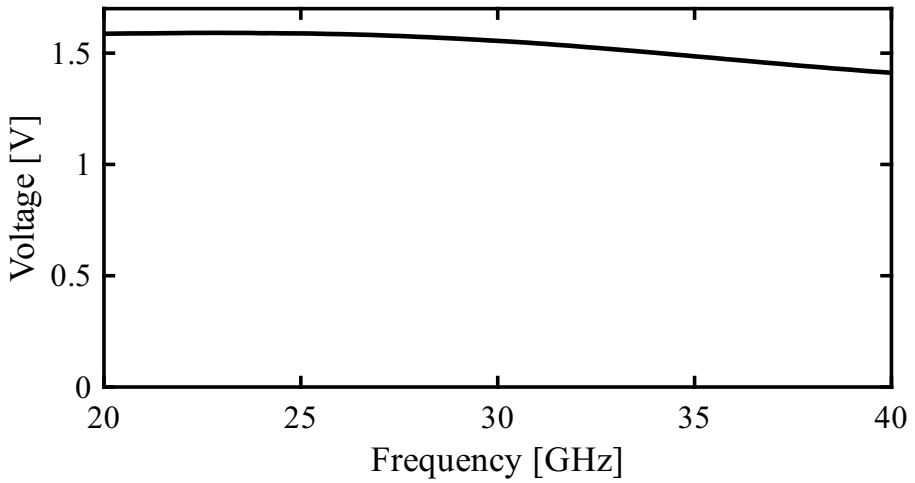


Figure 3.21: Schematic-level representation of the 1 mm transmission-line clock-distribution case-study implementation obtained with the design procedure developed in Section 3.2.3.

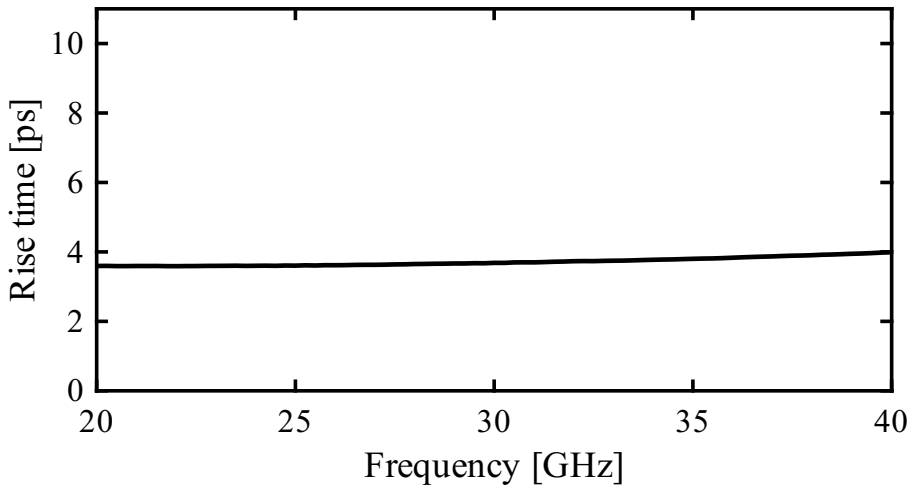
The random jitter remains relatively stable over the lower part of the simulated frequency range and gradually increases toward higher operating frequencies, while remaining below the corresponding specification. The average current presents a broad variation over the complete simulated range, but remains relatively stable within the target operating range. It therefore provides the reference for the subsequent comparison with the segmented implementations.

The jitter transfer functions are shown in Figure 3.24. The curves remain close to unity over the lower part of the analyzed frequency range and exhibit only a small amplification peak. Beyond this region, the transfer magnitude progressively decreases as the frequency increases. The peak amplification remains close to the corresponding specification limit.

Finally, the power-supply-induced jitter transfer function (PSIJ) is reported in Figure 3.25. It describes how a disturbance applied to the supply voltage is converted into jitter at the output of the distribution network. The sensitivity remains below the cor-



(a)

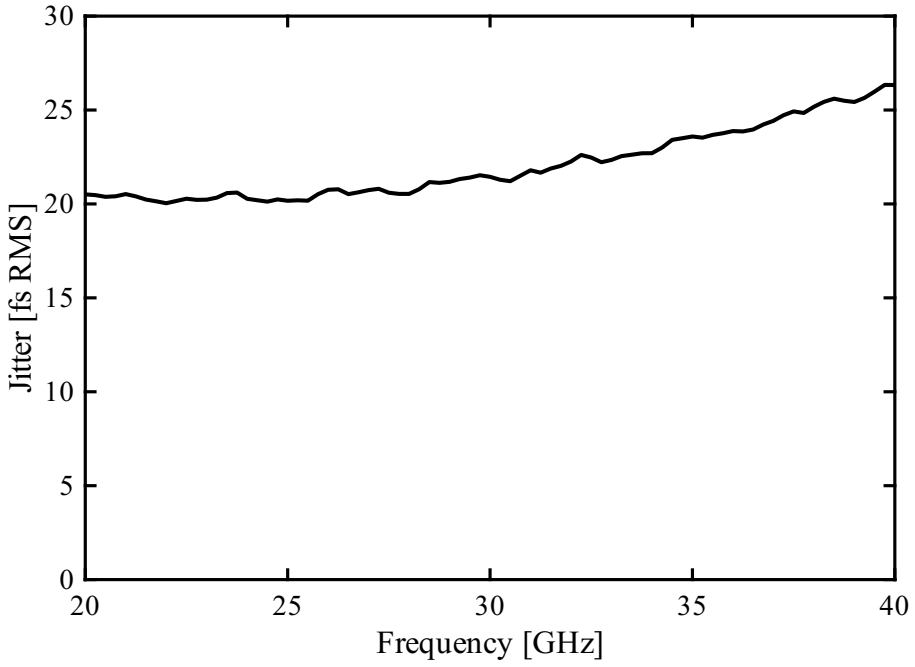


(b)

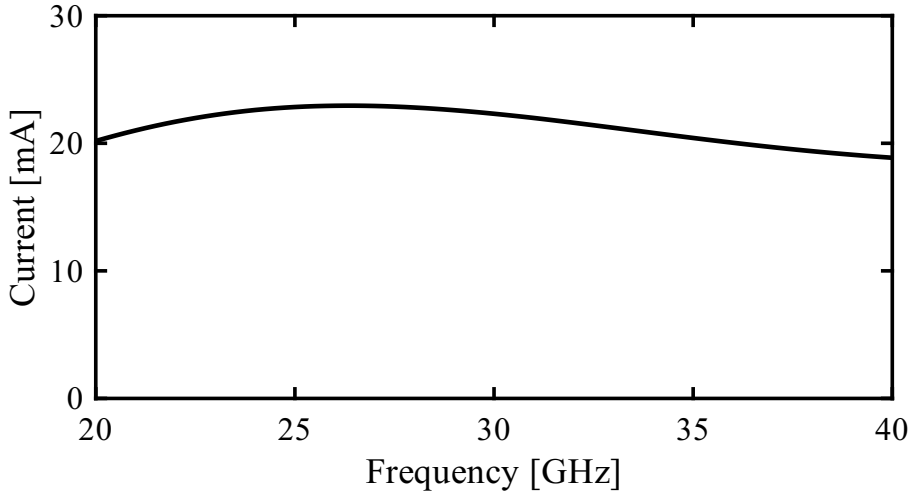
Figure 3.22: Output differential peak-to-peak voltage (a) and rise time (b) as functions of the operating frequency for the single-segment 1 mm transmission-line implementation.

responding specification over the analyzed frequency range. Nevertheless, even a supply disturbance of only a few millivolts can produce several femtoseconds of output jitter, making this contribution potentially critical.

Overall, the single-segment implementation satisfies the considered signal-quality requirements within the target operating range. The output voltage swing, rise time, random jitter, and power-supply-induced jitter retain sufficient margin, while the jitter amplification remains close to its specification limit and therefore represents the main constraint of this configuration. The following sections investigate whether segmenting the transmission line can improve this critical metric while preserving the other performance



(a)



(b)

Figure 3.23: Output random jitter (a) and average current consumption (b) for the single-segment 1 mm transmission-line implementation.

quantities.

3.3.2 Two-segment implementation

The first segmented implementation divides the complete 1 mm transmission line into two 500 μm sections. Its architecture and the multiplicities of the CMOS inverter stages are shown in Figure 3.26. The first section is driven by a tapered CMOS inverter chain starting from the fixed source stage. At the end of the first transmission-line section, an inverter with the same multiplicity as the fixed source stage is introduced to restore the same initial sizing condition for the second segment. An identical tapered chain can therefore be used to drive the second section, resulting in the same driving module for both segments.

Unlike the complete 1 mm transmission line, each 500 μm section falls within the range in which the equivalent-capacitance approach can provide a theoretical estimate of the required line-driver multiplicity. Since this procedure does not determine the final transistor dimensions exactly, small sizing adjustments are introduced through transistor-level simulations. Once the line-driver multiplicity has been selected, the intermediate inverter stages are sized using the tapered-chain criterion derived previously.

Compared with the single-segment implementation, each driver sees a shorter transmission line section and a smaller transformed load. This reduces the effort required from each individual driving stage, while the introduction of only one intermediate inverter chain limits the additional circuit complexity.

The output differential peak-to-peak voltage and rise time are reported in Figure 3.27. The output voltage remains comparable with that obtained for the single-segment implementation, while the rise time is slightly reduced. The intermediate inverter chain therefore improves the output transition without significantly affecting the voltage swing.

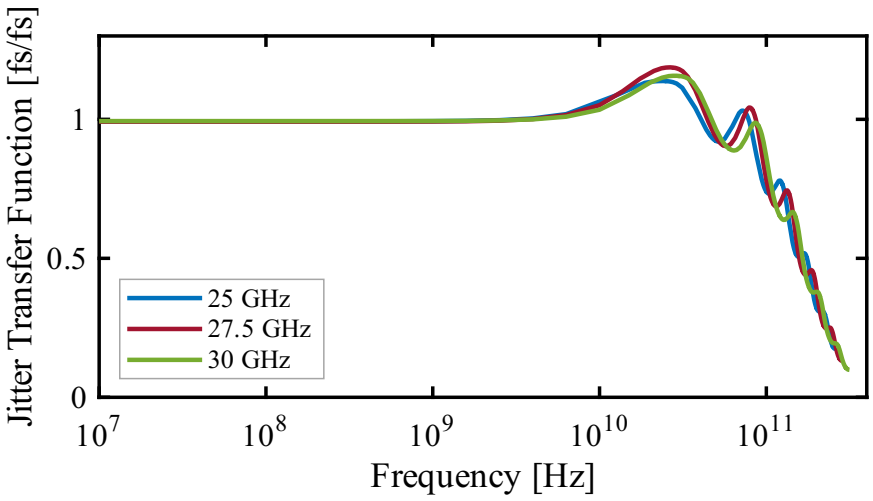


Figure 3.24: Jitter transfer functions at different clock frequencies for the single-segment 1 mm transmission-line implementation.

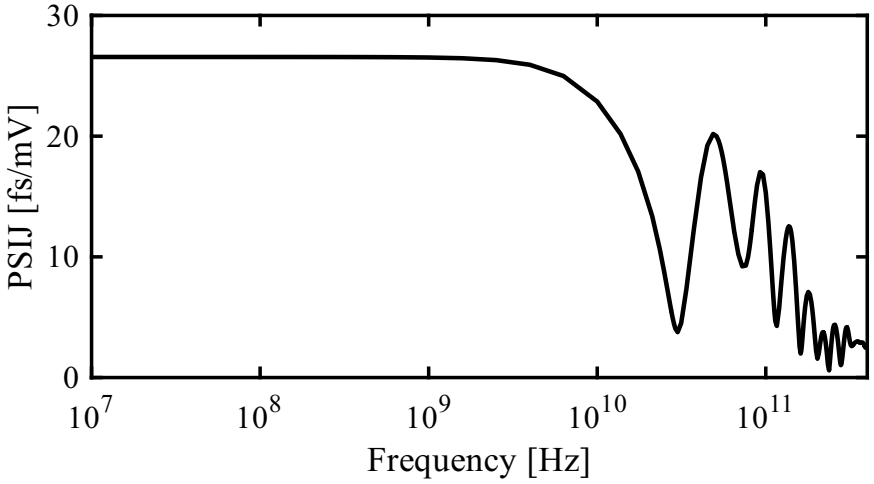


Figure 3.25: Power-supply-induced jitter transfer function for the single-segment 1 mm transmission-line implementation.

The output random jitter and average current consumption are reported in Figure 3.28. The random jitter remains comparable with the single-segment result. The additional inverter chain introduces further device-noise contributions; however, these are partially compensated by the improved waveform regeneration and by the reduced load driven by each stage. As a result, the output random jitter remains in the same range as in the single-segment implementation. The average current consumption instead increases because of the additional active stages introduced along the clock path.

The jitter transfer functions are reported in Figure 3.29. The transfer function magnitude remains close to unity over the analyzed range, showing a clear improvement compared with the single-segment implementation. The two-segment configuration therefore limits jitter amplification while preserving the signal quality along the complete distribution path.

Finally, the PSIJ transfer function is reported in Figure 3.30. Compared with the single-segment implementation, the sensitivity to supply disturbances is significantly increased and exceeds the corresponding specification over part of the analyzed frequency range. Together with high average current consumption, this increased sensitivity repre-

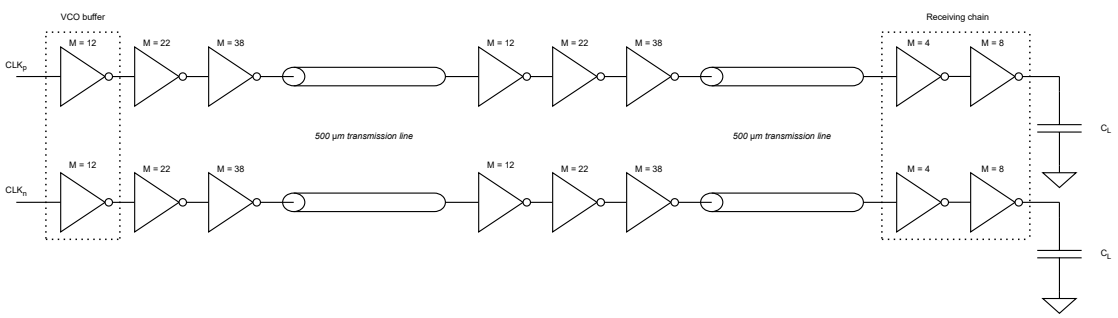


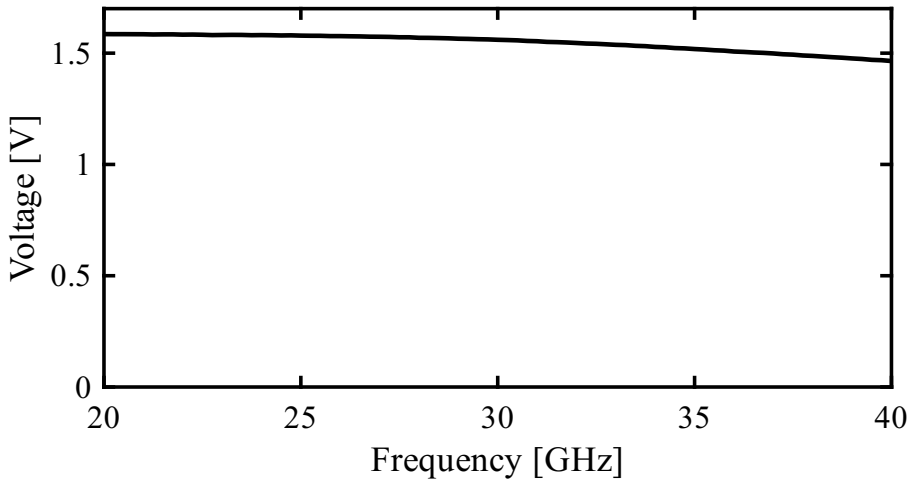
Figure 3.26: Schematic-level representation of the two-segment 1 mm transmission-line clock-distribution case-study implementation with the defined design procedure.

sents one of the main trade-offs introduced by the additional inverter chain.

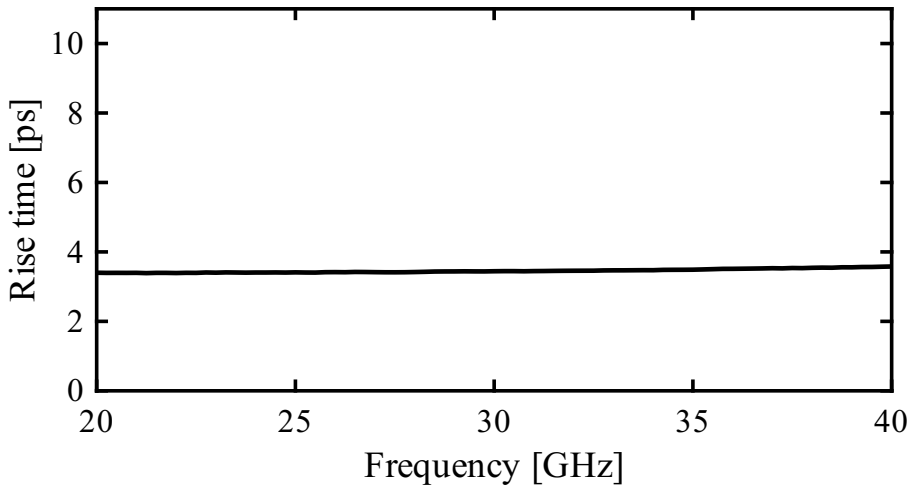
3.3.3 Four-segment implementation

The second segmented implementation divides the complete 1 mm transmission line into four 250 μm sections. Its architecture and the multiplicities of the CMOS inverter stages are shown in Figure 3.31. Each section is driven by the same tapered CMOS inverter chain. The first chain starts from the fixed source stage, while each following section starts from an inverter with the same multiplicity as the fixed source stage, restoring the same initial sizing condition at each segment. The same driving chain can therefore be repeated throughout the distribution. The shorter segment length further reduces the transformed load presented to each individual line driver.

The theoretical sizing procedure is applied to the 250 μm transmission-line sections to obtain an initial estimate of the required line-driver multiplicity. As in the two-segment



(a)



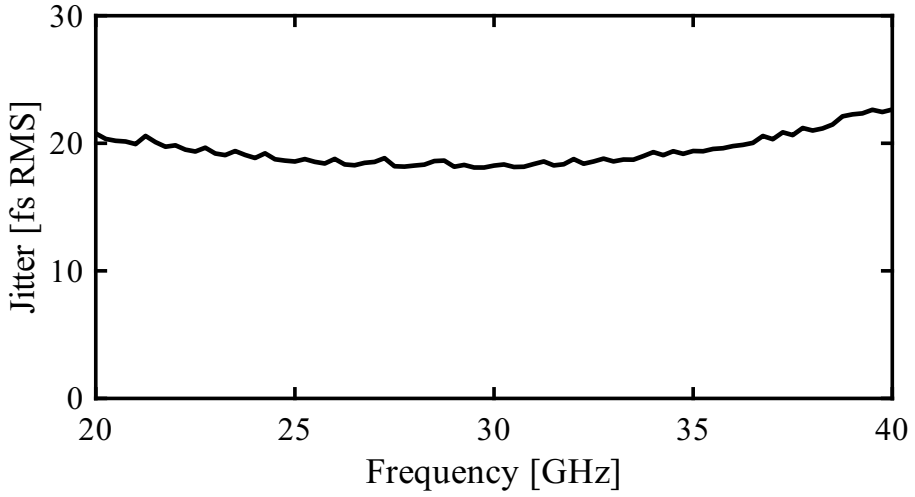
(b)

Figure 3.27: Output differential peak-to-peak voltage (a) and rise time (b) as functions of the operating frequency for the two-segment 1 mm implementation.

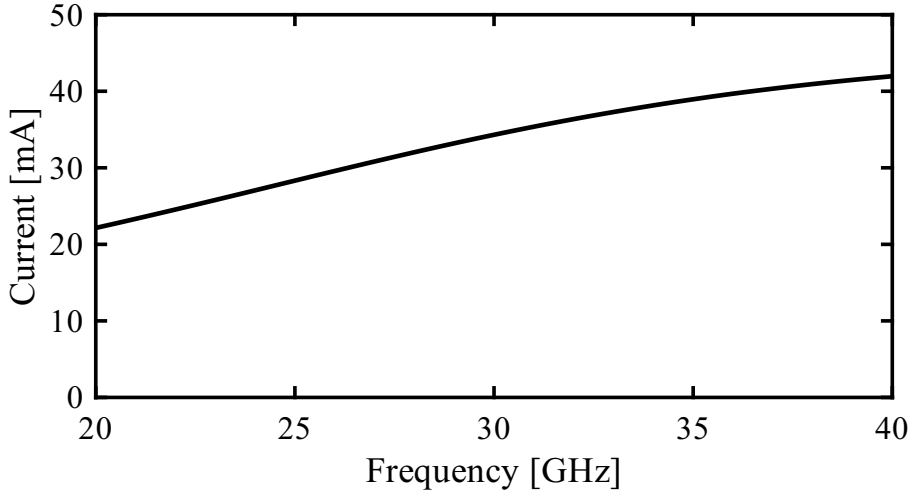
implementation, the resulting transistor dimensions are subsequently refined through transistor-level simulations, introducing small adjustments where necessary.

From the point of view of the individual driving chain, this configuration is less demanding than both the single- and two-segment implementations. Indeed, each line driver only has to charge and discharge the equivalent load associated with a 250 μm transmission-line section. However, the complete clock signal must pass through a larger number of inverter stages than in previous implementations before reaching the final output.

The output differential peak-to-peak voltage and rise time are reported in Figure 3.32. Both quantities remain comparable with those obtained for the previous implementations and vary only slightly within the target operating range. The output voltage exhibits a limited reduction as the frequency increases, while the rise time shows only a moderate increase. The four-segment architecture therefore preserves the output swing and transition



(a)



(b)

Figure 3.28: Output random jitter (a) and average current consumption (b) for the two-segment 1 mm transmission-line implementation.

speed despite the larger number of cascaded inverter stages.

The output random jitter is reported in Figure 3.33. It is significantly higher than in the single- and two-segment implementations, showing that the noise contributions introduced by the additional inverter stages accumulate along the complete clock path. Although the shorter transmission-line sections improve the local driving conditions, this advantage does not compensate for the larger number of active stages. The degradation becomes more evident toward the upper end of the frequency range, where the complete inverter chain operates under more demanding conditions.

The jitter transfer functions are reported in Figure 3.34. Although each transmission-line section is shorter and easier to drive than in aforementioned implementations, the complete distribution contains a large number of cascaded inverter chains. Small jitter-

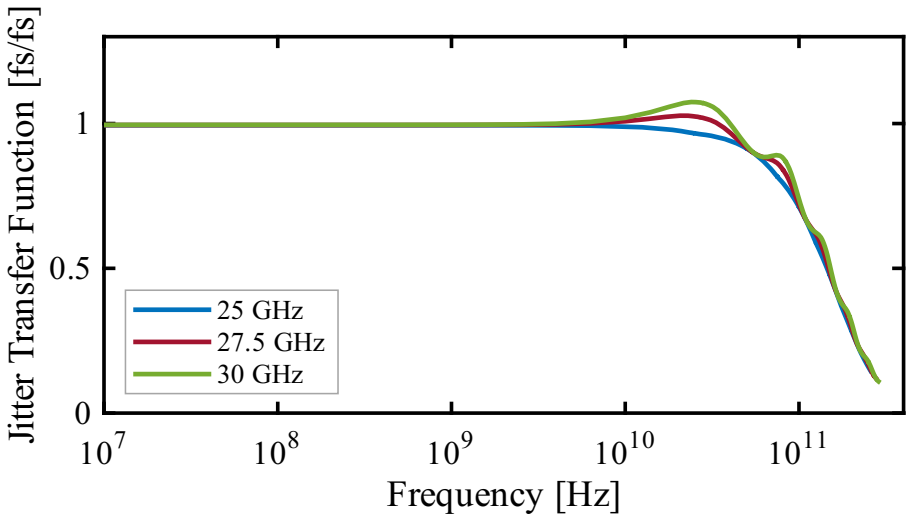


Figure 3.29: Jitter transfer functions for the two-segment 1 mm transmission-line implementation.

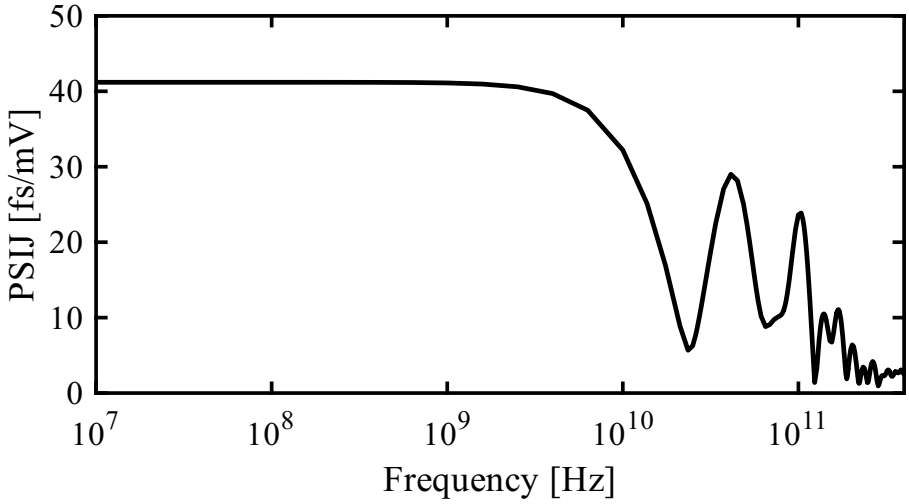


Figure 3.30: Power-supply-induced jitter transfer function for the two-segment 1 mm transmission-line implementation.

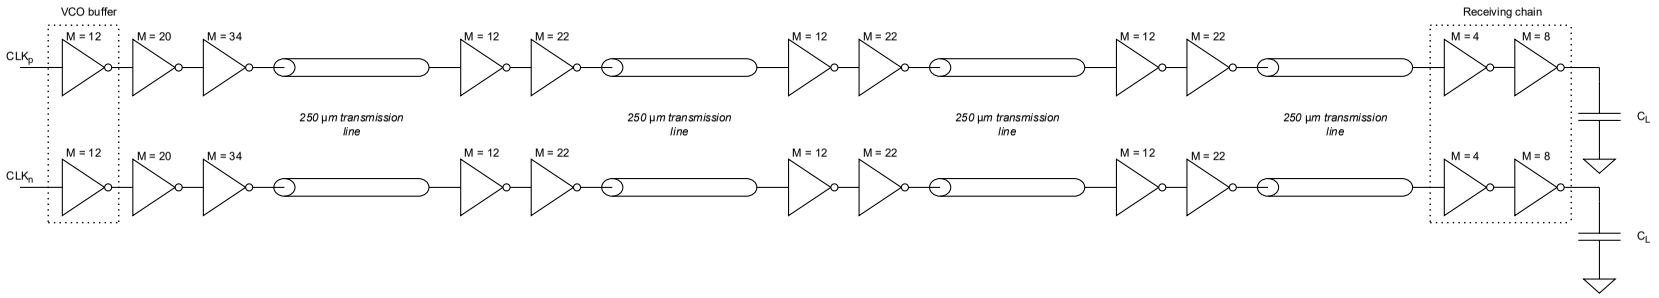
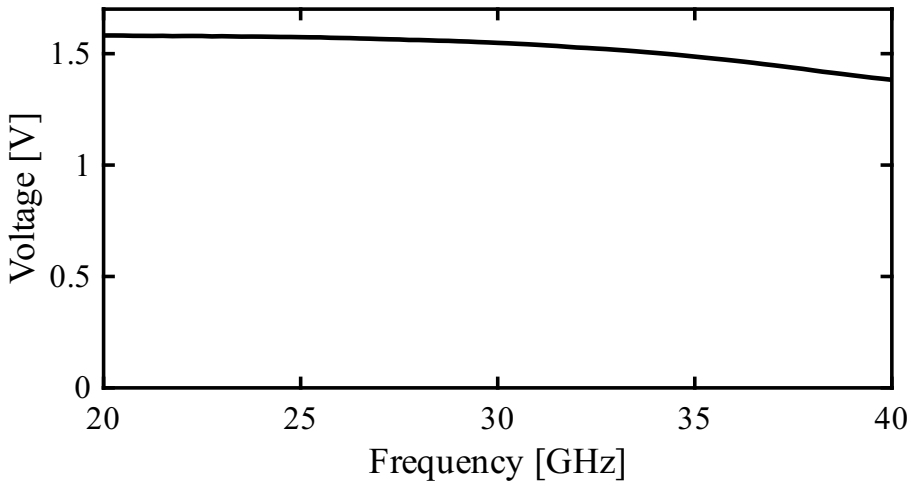
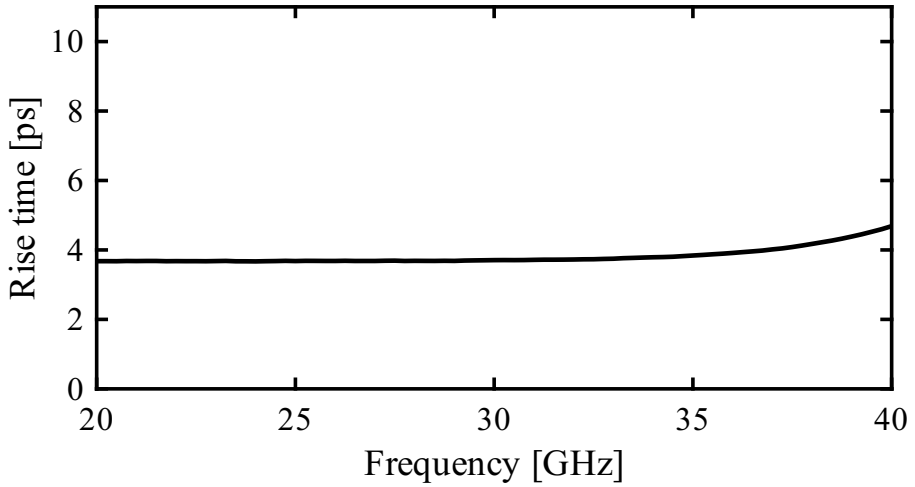


Figure 3.31: Schematic-level representation of the four-segment 1 mm transmission-line clock-distribution case-study implementation with the defined design procedure.



(a)



(b)

Figure 3.32: Output differential peak-to-peak voltage (a) and rise time (b) as functions of the operating frequency for the four-segment 1 mm implementation.

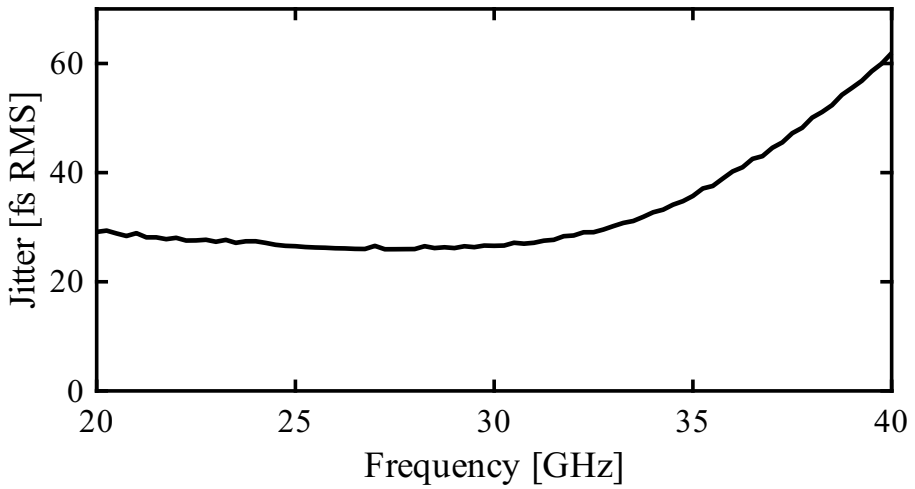


Figure 3.33: Output random jitter as a function of the operating frequency for the four-segment 1 mm transmission-line implementation.

transfer contributions introduced by the repeated stages accumulate along the clock path, resulting in the highest overall jitter amplification among the three implementations considered. This effect becomes increasingly evident at higher operating frequencies, where the transfer magnitude departs more significantly from unity.

The average current consumption is reported in Figure 3.35. Compared with both the single- and two-segment implementations, the four-segment configuration exhibits a substantial increase in current consumption, caused by the large number of inverter chains, the corresponding additional switching activity, and the increased total parasitic capacitance.

Finally, the PSIJ transfer function is reported in Figure 3.36. Compared with both the single- and two-segment implementations, the resulting sensitivity is significantly higher. This result indicates that the large number of cascaded inverter chains introduces an additional penalty in terms of supply-noise susceptibility.

Therefore, the four-segment implementation shows that reducing the length of the individual transmission-line sections does not produce a monotonic improvement in the complete distribution network. Although each local interconnect becomes easier to drive and the differential output voltage swing and rise time remain within their respective specification limits, the large number of cascaded inverter stages produces high random jitter, large jitter amplification, higher current consumption, and degraded sensitivity to supply disturbances. Dividing the line into too many sections is therefore not advantageous for the considered distribution network.

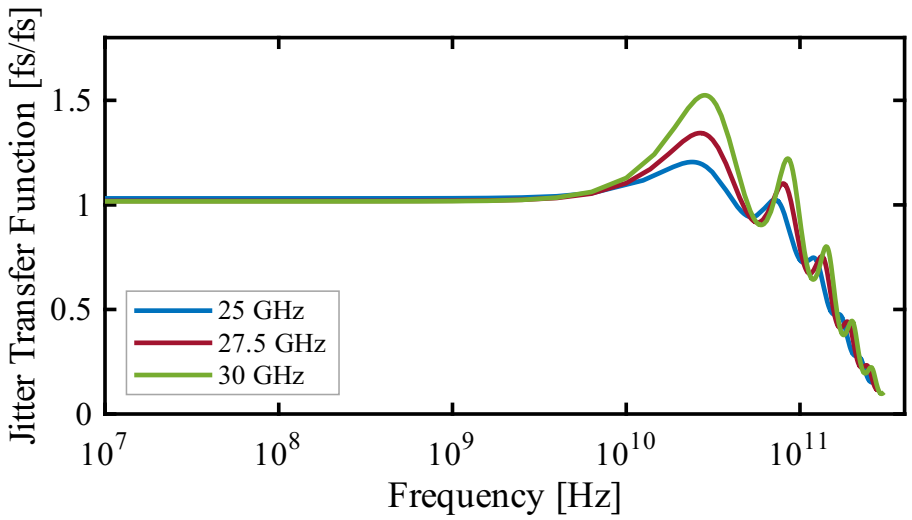


Figure 3.34: Jitter transfer functions for the four-segment 1 mm transmission-line implementation.

3.3.4 Comparison of line implementations

The comparison of the three implementations shows that segmenting a long transmission line can simplify its driving. By dividing the complete line into short sections, each driver sees a small equivalent load, reducing the effort required from the individual driving stages and allowing the clock waveform to be regenerated along the distribution path.

However, increasing the number of segments does not necessarily improve the overall performance. Indeed, each additional segment requires further inverter stages, whose noise and timing contributions accumulate along the clock path. The additional circuitry also

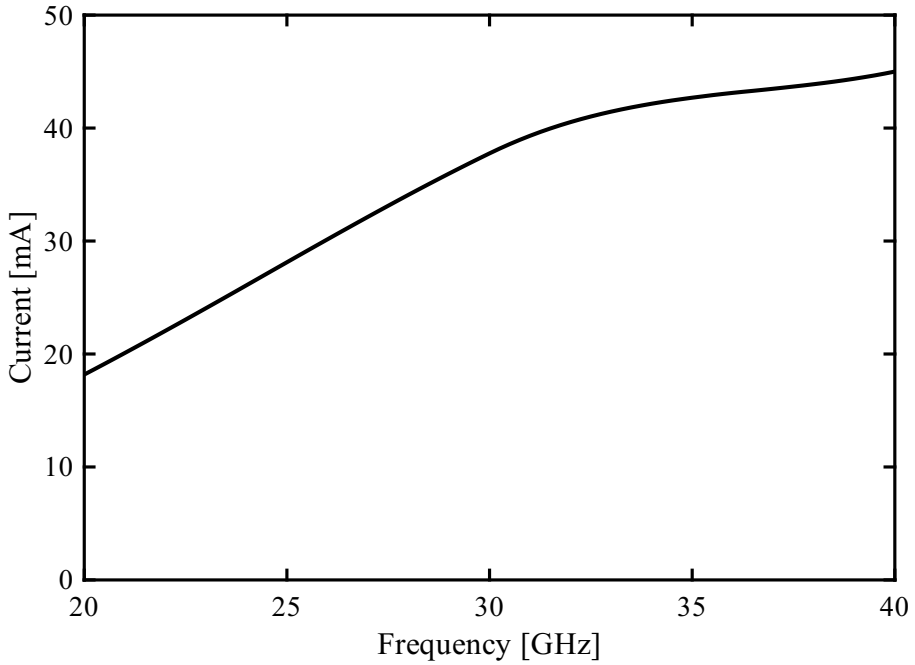


Figure 3.35: Average current consumption for the four-segment 1 mm transmission-line implementation.

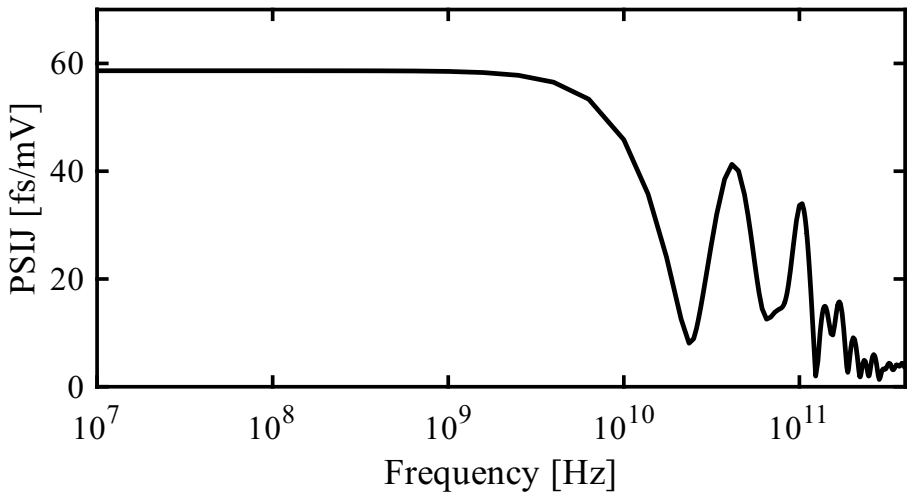


Figure 3.36: Power-supply-induced jitter transfer function for the four-segment 1 mm transmission-line implementation.

increases the switching activity and the sensitivity of the distribution network to supply disturbances. Therefore, excessive segmentation results in high random jitter, large jitter amplification, large average current consumption, and increased power-supply-induced jitter.

The simulated trends of the main performance metrics for the single-segment, two-segment, and four-segment implementations are compared in Figures. 3.37 and 3.38.

For the 1 mm transmission line considered in this work, the two-segment implementation provides the best compromise among the analyzed configurations. Indeed, compared with the single-segment implementation, it reduces the load driven by each stage and improves the jitter-amplification behavior, while introducing only one additional regeneration block. Conversely, the four-segment implementation introduces too many cascaded inverter stages, and the resulting penalties outweigh the benefit obtained from the shorter transmission-line sections.

These results show that an appropriate number of segments must be selected by balancing the easier driving of the individual transmission-line sections against the cumulative penalties introduced by the additional regeneration stages.

3.4 Effect of Back-to-Back Inverter Placement

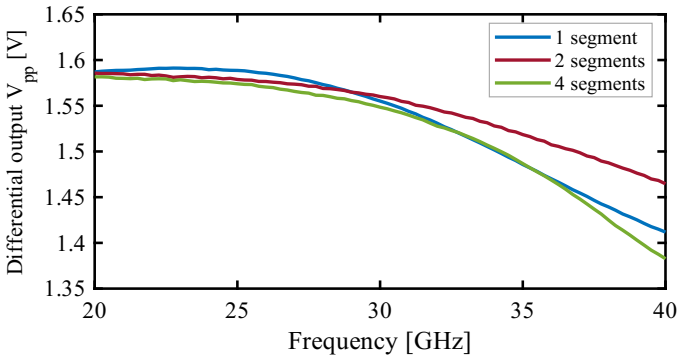
As an additional design variant, back-to-back inverter pairs are introduced at different positions along the 1 mm clock-distribution path. A back-to-back inverter pair consists of two cross-coupled CMOS inverters connected between the differential branches. The resulting local positive feedback reinforces the differential voltage during a transition, providing a regenerative action that can sharpen degraded clock edges and potentially improve their timing quality. This design strategy is therefore considered as a simple architectural solution for introducing local signal regeneration using the same CMOS inverter-based structure already adopted in the distribution network. Three configurations are considered: one pair at the transmission-line input, one pair at the line midpoint, and one pair at the line output, as shown in Figure 3.39. In each configuration, both inverters of the additional pair have multiplicity M_{BTB}. The purpose of this analysis is to evaluate how the position and size of the additional inverter pair affect the performance of the complete clock-distribution path.

To focus specifically on the effect of the back-to-back inverter pair, the analysis is performed at a single operating frequency of 27.5 GHz. This frequency is chosen as the midpoint of the target operating range from 25 to 30 GHz, providing a representative operating condition while avoiding the additional variability introduced by frequency. For each configuration, M_{BTB} is varied, and the maximum jitter amplification, output random jitter, rise time, average current consumption, and output voltage swing are evaluated. The comparison is primarily based on jitter amplification, since this represents the most

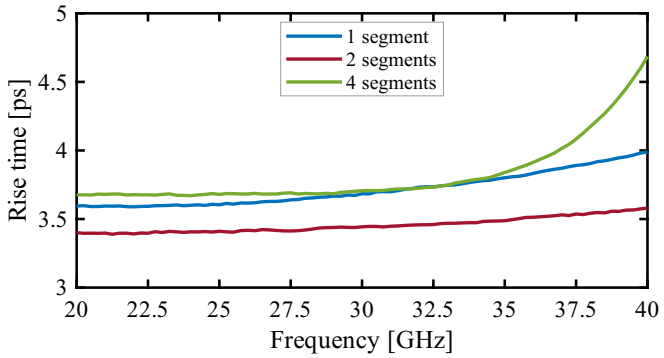
critical metric for the original architecture. The remaining quantities are examined to identify their trends and any potential trade-offs.

3.4.1 Placement at the transmission-line input

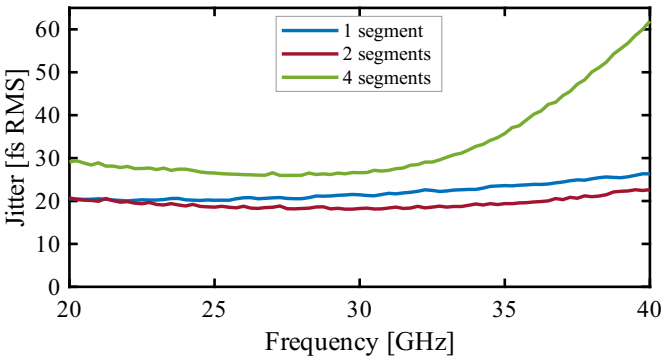
When the back-to-back inverter pair is placed at the transmission-line input, increasing M_{BTB} provides a progressive reduction in the output rise time. This behavior is consistent with the stronger local regenerative action provided by a large back-to-back pair, which reinforces the differential transition before the signal propagates along the transmission



(a)



(b)



(c)

Figure 3.37: Comparison of the single-segment, two-segment, and four-segment 1 mm transmission-line implementations in terms of (a) differential output peak-to-peak voltage, (b) output rise time, and (c) output random jitter.

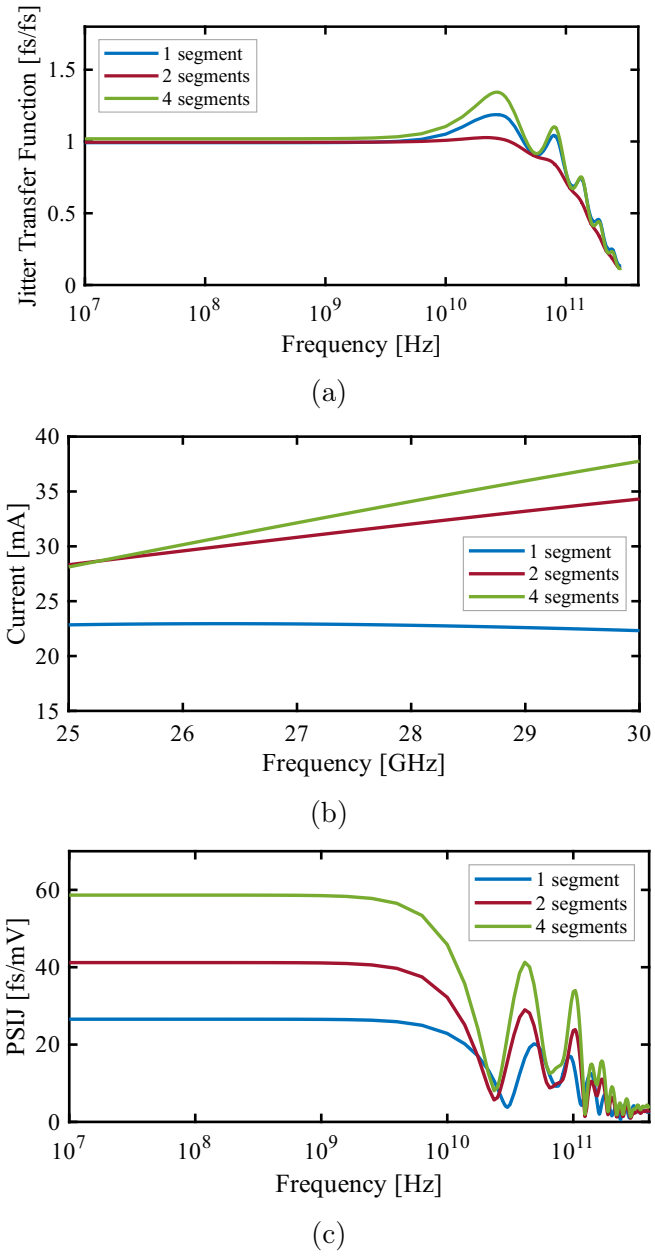


Figure 3.38: Comparison of the single-segment, two-segment, and four-segment 1 mm transmission-line implementations in terms of (a) jitter transfer function, (b) average current consumption, and (c) PSIJ. The jitter transfer function comparison is evaluated at a clock frequency of 27.5 GHz.

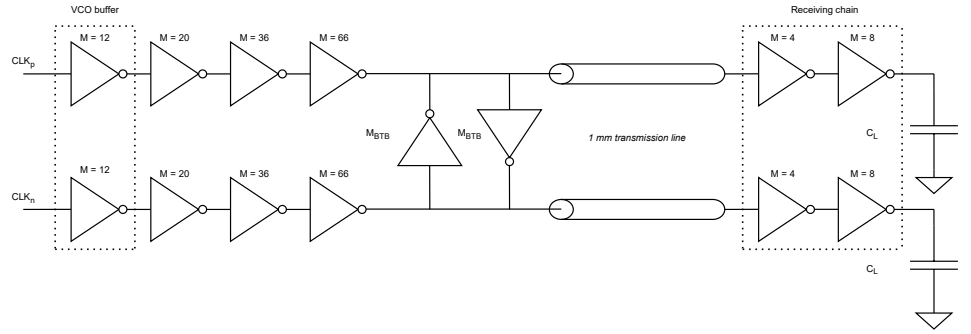
line. The corresponding trends in rise time, maximum jitter amplification, and average current consumption are reported in Figure 3.40.

Although the output rise time decreases as the input pair is enlarged, the maximum jitter amplification increases over the complete analyzed range. The average current consumption also increases almost linearly with M_{BTB}, as expected from a large switching stage size introduced before the line.

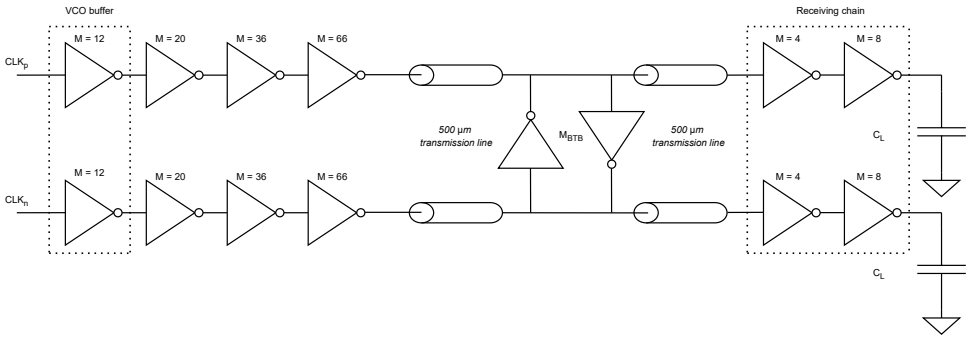
The output random jitter shows an overall increasing trend, particularly for large inverter multiplicities, whereas the output voltage swing varies only slightly and remains

close to its reference value over the relevant range. These quantities are therefore not reported separately.

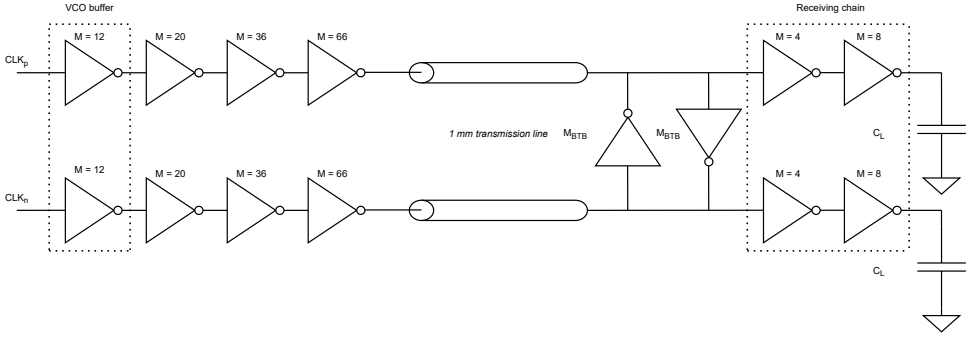
The input placement consequently provides a trade-off between transition speed and overall timing and power performance. A larger back-to-back pair improves the rise time, but simultaneously increases jitter amplification, output random jitter, and current consumption. It therefore does not provide an overall advantage for the considered distribution path.



(a)



(b)

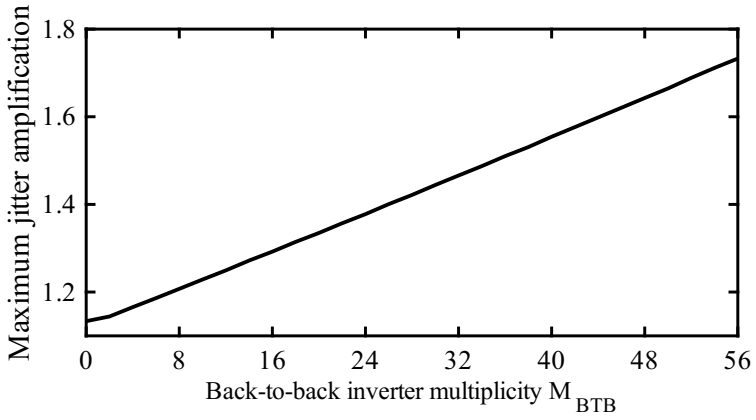


(c)

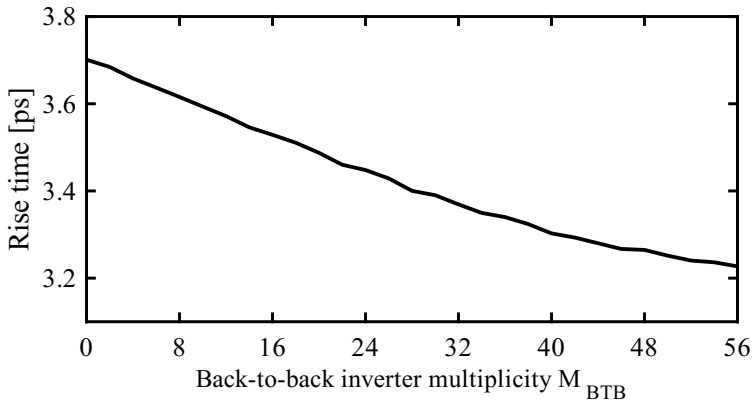
Figure 3.39: Back-to-back inverter-pair placements at the transmission-line input (a), midpoint (b), and output (c) along the 1 mm transmission-line path.

3.4.2 Placement at the transmission-line midpoint

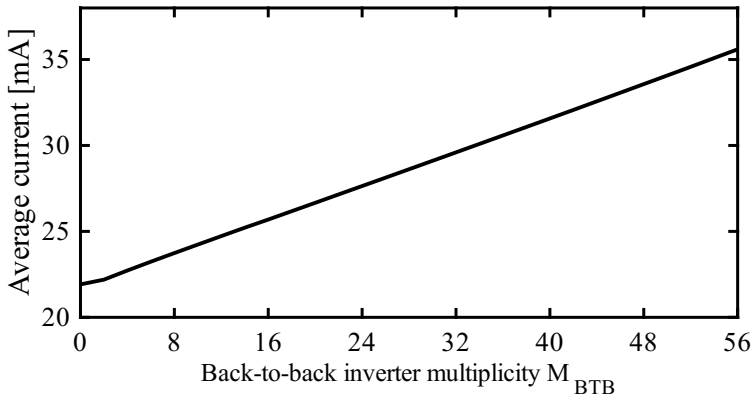
The back-to-back inverter pair is placed at the midpoint of the 1 mm transmission line, resulting in a symmetric configuration with two equal 500 μm sections. For this placement, Figure 3.41 shows the maximum jitter amplification as a function of the back-to-back inverter multiplicity.



(a)



(b)



(c)

Figure 3.40: Maximum jitter amplification (a), output rise time (b), and average current consumption (c) as functions of the back-to-back inverter multiplicity for the transmission-line input placement.

For small inverter multiplicities, the maximum jitter amplification remains close to the value obtained without the back-to-back inverter pair. As M_{BTB} is increased further, the maximum jitter amplification progressively increases.

The secondary metrics do not show a compensating advantage. Indeed, the output rise time increases as the inverter pair is enlarged, while the output random jitter remains close to the reference value only at the lowest multiplicities and subsequently increases. The average current consumption also rises with M_{BTB}.

The midpoint placement therefore does not provide a useful improvement. Small inverter multiplicities have only a marginal effect on the maximum jitter amplification, whereas larger multiplicities degrade the timing performance and increase the current consumption.

3.4.3 Placement at the transmission-line output

A different behavior is observed when the back-to-back inverter pair is placed at the transmission-line output. In this configuration, the pair acts on the signal after propagation through the complete interconnect, and no additional transmission-line section follows it. The corresponding maximum jitter amplification, output random jitter, and rise time are reported in Figure 3.42.

As M_{BTB} is increased from small values, the maximum jitter amplification progressively decreases and approaches unity. The reported quantity is the maximum of the jitter-transfer function over the complete analyzed frequency range. Although the amplification occurring at some frequencies is reduced, the jitter-transfer magnitude remains approximately equal to one at the clock frequency, where the timing perturbation is transferred together with the clock signal. This unity-gain component therefore sets the plateau observed at intermediate multiplicities. For low to moderate values of M_{BTB}, the output

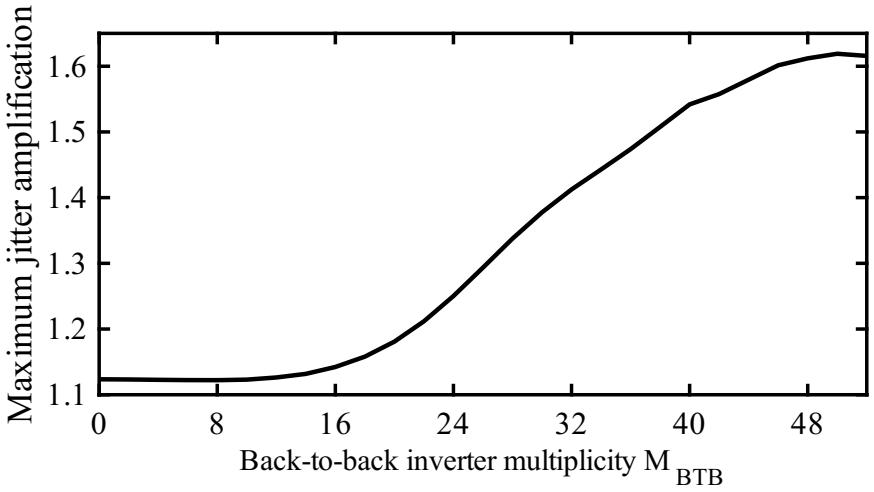
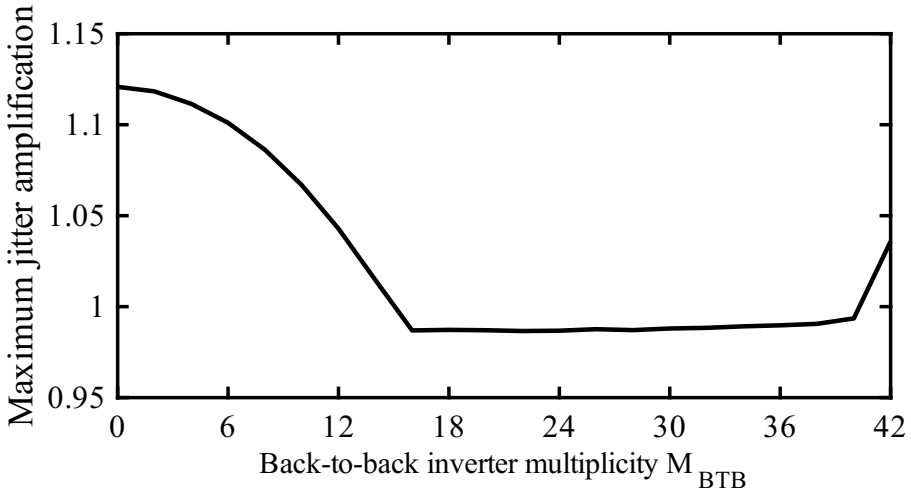
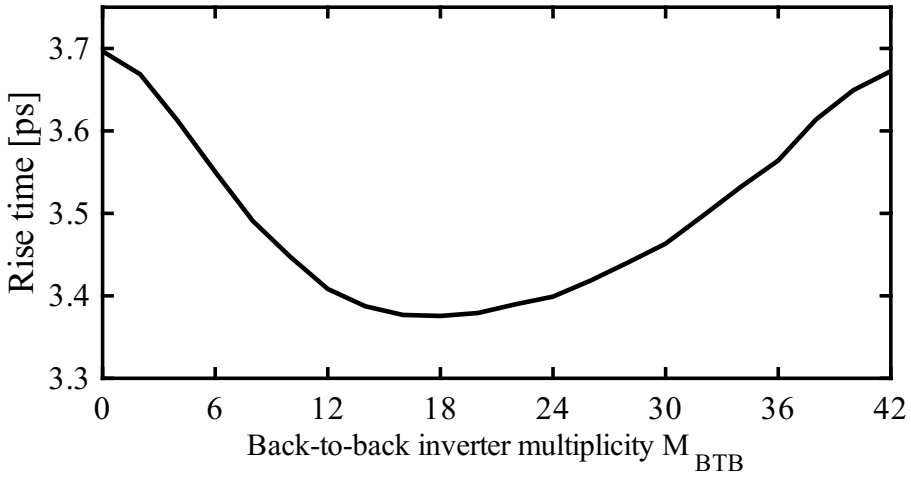


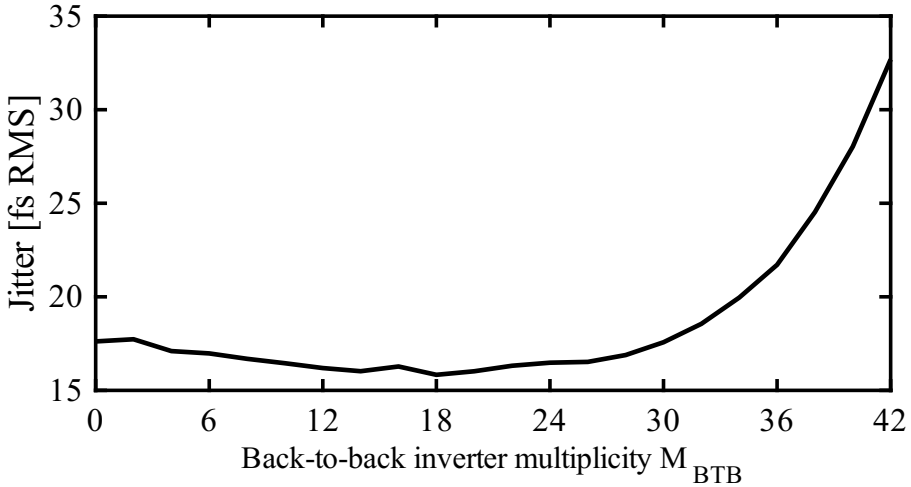
Figure 3.41: Maximum jitter amplification as a function of the back-to-back inverter multiplicity when the inverter pair is placed at the transmission-line midpoint.



(a)



(b)



(c)

Figure 3.42: Maximum jitter amplification (a), output rise time (b), and output random jitter (c) as functions of the back-to-back inverter multiplicity for the transmission-line output placement.

random jitter and rise time also decrease, reaching their minimum values at moderate inverter multiplicities. As M_{BTB} is increased further, the additional capacitive loading and the latch-like behavior of the cross-coupled pair oppose subsequent differential transitions, leading to an increase in maximum jitter amplification, output random jitter, and rise time.

The average current consumption remains close to the reference value over the range in which the timing metrics improve. The output voltage swing decreases gradually, but its variation remains limited over the same range and does not determine the selection of the inverter multiplicity.

The output placement therefore exhibits a useful operating range at moderate values of M_{BTB}. Within this range, the back-to-back pair reduces the maximum jitter amplification, output random jitter, and rise time without introducing a significant current penalty. For larger inverter multiplicities, these improvements progressively diminish.

3.4.4 Summary of the Back-to-Back configurations

The comparison of the three configurations shows that the effect of the back-to-back inverter pair strongly depends on its position along the transmission line. At the line input, increasing the pair size improves the output transition speed but simultaneously increases jitter amplification, output random jitter, and current consumption. The reduction in rise time therefore does not translate into an overall timing improvement.

When the pair is placed at the midpoint, only marginal variations in maximum jitter amplification are observed for the smallest multiplicities. Increasing M_{BTB} further degrades the jitter amplification, output random jitter, and rise time, while also increasing the average current consumption. The midpoint placement therefore does not provide a useful operating range.

A clear overall improvement is obtained only when the pair is placed at the transmission-line output. Within a moderate multiplicity range, this configuration reduces the maximum jitter amplification, output random jitter, and rise time while maintaining an average current consumption close to that of the reference implementation. For larger multiplicities, the timing metrics progressively deteriorate and the advantage is lost.

The different behavior of the three placements can be interpreted in terms of the amount of signal regeneration required at each position. At the transmission-line input, the clock waveform has just been generated by the line driver and already presents a well-defined transition. The regenerative benefit provided by the back-to-back pair is therefore limited, while the additional pair introduces extra loading. As the signal propagates along the transmission line, its transition is progressively affected by the interconnect. The midpoint represents an intermediate condition, whereas at the line output the back-to-

back pair acts after the complete propagation path, where signal regeneration is more beneficial. In this position, a moderately sized pair can sharpen the final differential transition and improve the associated timing performance. For larger M_{BTB}, however, the increasing loading progressively offsets this regenerative benefit.

The results therefore indicate that the most effective configuration is a moderately sized back-to-back inverter pair placed directly at the transmission-line output.

Chapter 4

Resonant clock distribution

4.1 Motivation for resonant clock distribution

The non-resonant clock-distribution architecture investigated in Chapter 3 allows the clock signal to be transmitted over the considered interconnect while satisfying the required voltage-swing, rise-time, and random-jitter specifications. However, the jitter-transfer behavior represents one of its main limitations. In particular, for the 1 mm transmission-line implementation, the maximum jitter amplification is close to the corresponding specification limit and therefore represents the main constraint of the distribution implementation.

Driver upsizing can partially mitigate this limitation. Indeed, large driver size reduces the transition time and allows the signal to fully settle within each half-cycle, thereby reducing the dependence of each threshold crossing on the signal state left by previous cycles. As already shown in Chapter 3, increasing the driver size progressively reduces the maximum jitter-amplification magnitude toward unity, without providing significant attenuation below unity in the investigated non-resonant configurations. At the same time, the average current consumption continues to increase. The non-resonant architecture therefore exhibits a direct trade-off between jitter-transfer performance and power consumption.

The impedance-transformation analysis developed in Chapter 3 suggests an alternative direction. The capacitive load connected at the receiving end is transformed by the transmission line, and the resulting input impedance therefore depends on the electrical length of the line. As the electrical length increases, the input reactance approaches zero and subsequently changes sign, suggesting that an appropriately selected inductive element can be used to establish a resonant condition at the desired operating frequency.

Based on this observation, an inductive element is introduced into the clock-distribution network to compensate the capacitive behavior of the loaded transmission line and force the system to resonate at a desired operating frequency. The following analysis investigates whether this resonant condition can overcome the main limitations

of the non-resonant implementation, particularly as regards jitter-transfer performance and power consumption trade-off, while preserving the required voltage-swing, rise-time, and random-jitter specifications.

4.2 Resonant architectures and operating principle

The general principle of resonant clock distribution is to introduce an inductive element into the distribution network to let the system resonate at a desired operating frequency. In this condition, at resonance, the inductive and capacitive contributions compensate each other, reducing the net reactive loading presented to the active driver. The energy stored in the capacitive elements is periodically transferred to the inductive element and subsequently returned to the network, while the driver mainly supplies the energy dissipated by the circuit losses. This mechanism is expected to reduce the power required from the active driver and improve the transition speed, allowing the clock signal to settle within the available half-cycle. Since improved settling reduces the dependence of each threshold crossing on the waveform state left by the previous cycles, an improvement in the jitter-transfer behavior is also expected.

The impedance seen by the inductive element depends on its position along the transmission line and therefore corresponds to a different effective reactive load at the operating frequency. Three representative placements are considered, as illustrated in Figure 4.1: at the near end, at the midpoint, and at the far end of the line. The resonance condition and the procedure used to determine the inductance are first introduced for the near-end configuration and then applied to the other two architectures by considering the equivalent capacitance at each inductor connection node. The comparison of these configurations allows the effect of the inductive-element placement on the clock-distribution performance to be evaluated.

Following the approach adopted for the non-resonant clock distribution in Chapter 3, a preliminary study is first carried out to determine whether introducing resonance can provide any performance advantage over the corresponding non-resonant implementation. For this initial analysis, a total transmission-line length of 500 μm is considered.

The study is focused on the nominal operating frequency of 27.5 GHz. For each configuration, the inductance is determined to satisfy the resonance condition at this frequency. The main objective at this stage is to assess the effect of resonant operation itself and determine if the resulting improvements justify pursuing the resonant architecture. The inductor quality factor is then varied as a secondary check to verify that the observed trends are not limited to an ideal-inductor assumption. The complete frequency-dependent behavior over the target operating range is subsequently evaluated in Section 4.3 using the final 1 mm implementations.

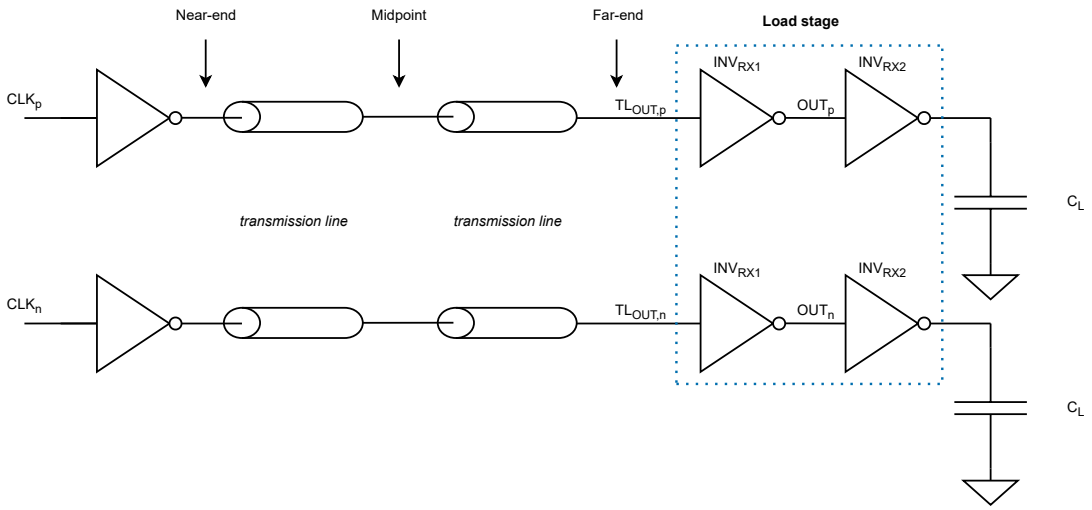


Figure 4.1: General clock-distribution architecture indicating the three possible placements of the shunt resonant inductor: near-end, midpoint, and far-end.

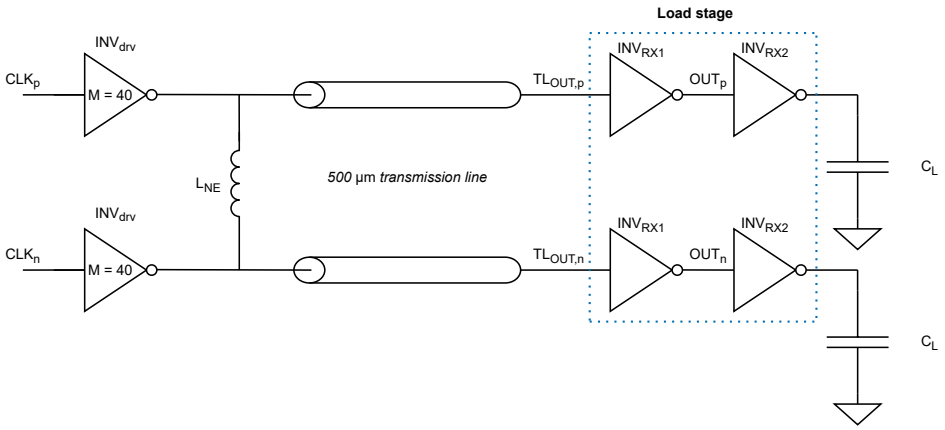


Figure 4.2: Near-end resonant configuration, with the shunt inductive element connected at the input of the transmission line.

4.2.1 Near-end resonant configuration

The first architecture considered is the near-end resonant configuration, shown in Figure 4.2, in which a shunt inductive element is connected at the input of the transmission line. As indicated in Figure 4.1, the near-end position corresponds to the outputs of the driving inverters INV_{DRV}. Each node is loaded by the output capacitance $C_{\mathrm{out,drv}}$ of the corresponding driving inverter INV_{DRV} and by the input impedance $Z_{\mathrm{in,line}}$ presented by the transmission line.

The node and inverter nomenclature introduced in Figure 3.1 is adopted throughout this chapter. INV_{RX1} and INV_{RX2} denote the first and second inverters of the receiving stage, respectively. Then, $TL_{\mathrm{OUT},p}/TL_{\mathrm{OUT},n}$ identify the differential transmission-line output, whereas OUT_p/OUT_n corresponds to the differential nodes between INV_{RX1} and INV_{RX2} used for performance evaluation.

When the loaded transmission line exhibits a capacitive input behavior, the reactive part of its input admittance can be represented by the equivalent input capacitance $C_{\text{in,line}}(\omega)$. At the near-end node, this contribution appears in parallel with the output capacitance of the final CMOS driver, INV_{DRV}. Therefore, the total equivalent capacitance at this node is equal to

$$C_{\text{eq,NE}}(\omega) = C_{\text{in,line}}(\omega) + C_{\text{out,drv}}. \tag{4.1}$$

The shunt inductance L_{NE} is connected in parallel with this equivalent capacitive contribution $C_{\text{eq,NE}}$. Using (4.1), the imaginary part of the near-end admittance can therefore be expressed as

$$\text{Im}\{Y_{\text{NE}}(\omega)\} = \omega C_{\text{eq,NE}}(\omega) - \frac{1}{\omega L_{\text{NE}}}. \tag{4.2}$$

At the target angular frequency ω_0, resonance is obtained when the capacitive and inductive susceptances are equal in magnitude and opposite in sign, so that the total susceptance at the near-end node is zero. From Eq. (4.2), the resonance condition is therefore obtained when

$$\omega_0 C_{\text{eq,NE}}(\omega_0) = \frac{1}{\omega_0 L_{\text{NE}}}. \tag{4.3}$$

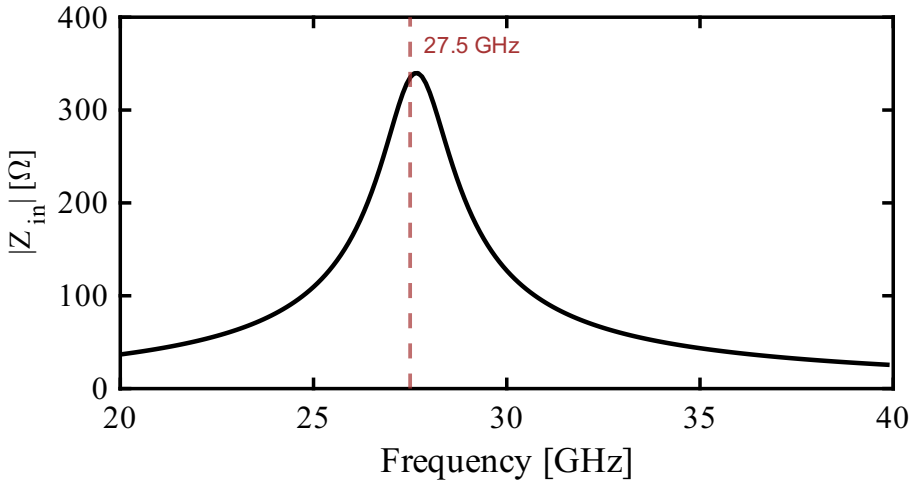
Since $C_{\text{eq,NE}}(\omega)$ depends on frequency through the impedance transformation of the loaded transmission line, the resonance condition is frequency dependent. In the present analysis, however, the target resonant frequency is set to 27.5 GHz, corresponding to the midpoint of the 25–30 GHz target operating range. The required near-end inductance can therefore be computed by evaluating the equivalent capacitance at ω_0. In this way, the resulting inductance required to establish resonance at the desired operating frequency is:

$$L_{\text{NE}} = \frac{1}{\omega_0^2 C_{\text{eq,NE}}(\omega_0)}. \tag{4.4}$$

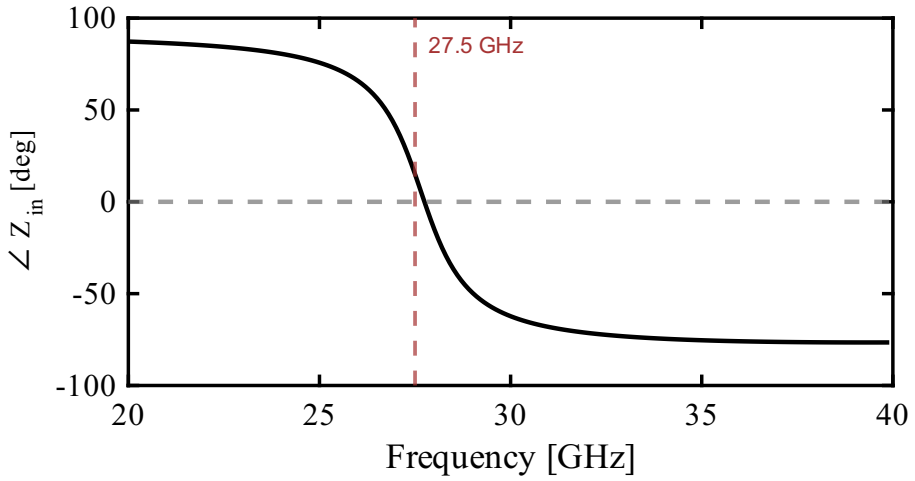
To verify that the inductance defined in (4.4) establishes resonance near the target frequency, the resulting impedance at the near-end node is studied as a function of frequency.

Figure 4.3 shows the evolution of the magnitude and phase of the total impedance at the near-end node as a function of frequency. Below the resonant region, the impedance exhibits a predominantly inductive behavior. As the frequency increases, the phase approaches zero while the impedance magnitude rises to a pronounced maximum. The coincidence of the impedance peak with the zero-phase crossing identifies the parallel resonance of the network. Above resonance, the impedance becomes predominantly capacitive. L_{NE} has been selected to establish resonance near the target frequency of 27.5 GHz.

The near-end resonant configuration is then evaluated in terms of average driver cur-



(a)



(b)

Figure 4.3: Magnitude (a) and phase (b) of the total impedance at the near-end node as functions of frequency.

rent consumption. To account for the losses of the resonant inductor, L_{NE} is modeled as an ideal inductance in series with a resistance R_s. Therefore, for a fixed operating frequency and inductance, the corresponding inductor quality factor is given by

$$Q_L = \frac{\omega_0 L_{\mathrm{NE}}}{R_s}. \tag{4.5}$$

The effect of inductor losses is investigated by varying R_s, and consequently Q_L, according to (4.5), while keeping the remaining circuit parameters unchanged. The resulting performance of the resonant configuration is compared with that of the corresponding non-resonant configuration, assuming a driver multiplicity $M_{\mathrm{drv}} = 40$.

As shown in Figure 4.4, the near-end resonant configuration exhibits a lower average driver current consumption than the corresponding non-resonant reference over the entire investigated range of Q_L. This confirms the benefit of resonant operation at the design

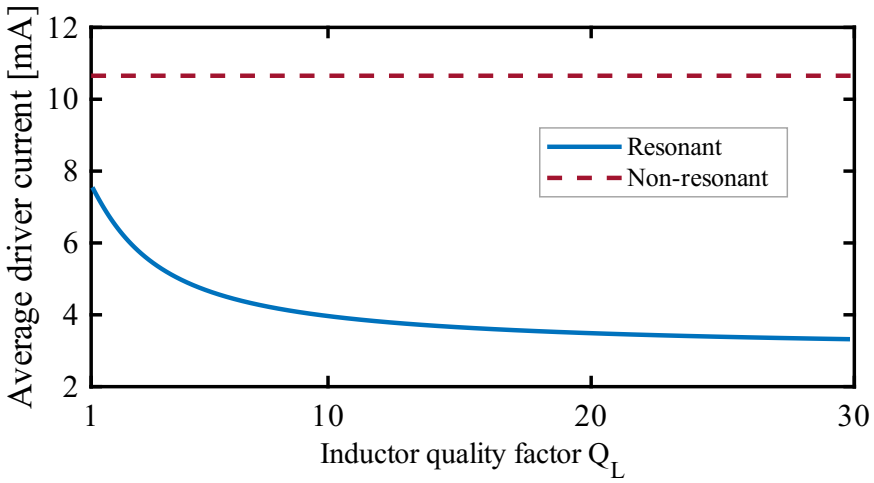


Figure 4.4: Average driver current as a function of the inductor quality factor for the near-end resonant configuration at 27.5 GHz. The non-resonant configuration simulation data with $M_{\text{drv}} = 40$ is included as a reference.

frequency. Part of the energy associated with the near-end equivalent capacitance C_{eq} is periodically exchanged with the inductive element rather than being supplied entirely by the active driver during each transition. Consequently, the required current is reduced and, for the same supply voltage, so is the average power consumption.

The extent of this reduction depends on the inductor quality factor. A low Q_L corresponds to a large series resistance R_s, which increases the losses in the inductive branch. Part of the energy involved in the resonant exchange is therefore dissipated and must be restored by the driver, reducing the benefit with respect to the non-resonant case. In the limiting case of $Q_L \to 0$, R_s tends to infinity, effectively disconnecting the inductive branch and recovering the current consumption of the non-resonant reference. As Q_L increases, these losses decrease and the current progressively approaches a limiting value, indicating that the contribution of the inductor losses becomes small compared with the remaining dissipative mechanisms of the clock-distribution circuit.

The near-end resonant configuration is also evaluated in terms of output voltage swing and rise time. As shown in Figure 4.5, the resonant configuration exhibits a slightly larger output voltage swing than the corresponding non-resonant case. As discussed in the near-end impedance analysis, the compensation of the capacitive loading at resonance allows the driver to apply a wide voltage swing to the transmission line. This contributes to the slight increase in output swing observed after regeneration by INV_{RX1}. At the same time, the peak-to-peak voltage remains nearly constant as Q_L varies, indicating that the inductor losses only have a limited influence on the final signal amplitude.

The rise time of the differential signal between OUT_p and OUT_n is then used to investigate the signal transitions. As shown in Figure 4.6, the resonant configuration provides a marked reduction in rise time with respect to the corresponding non-resonant implementation over the entire investigated range. The influence of the inductor quality

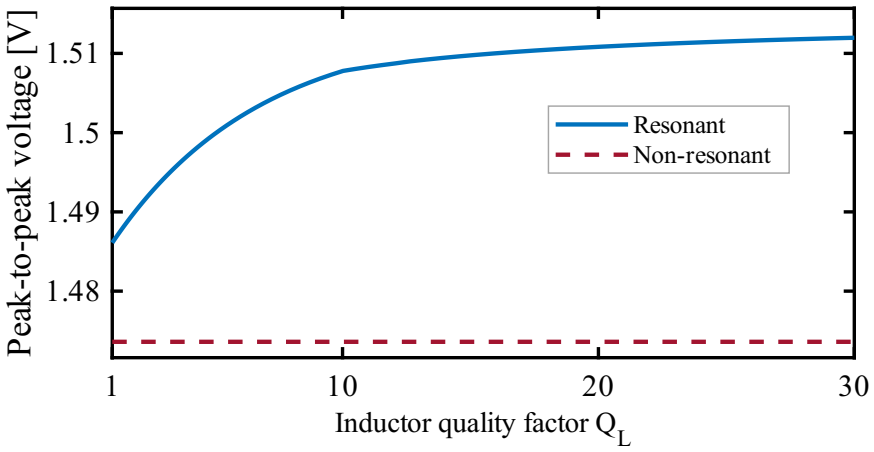


Figure 4.5: Differential output peak-to-peak voltage evaluated between OUT_p and OUT_n as a function of the inductor quality factor for the near-end resonant configuration at 27.5 GHz. The corresponding non-resonant configuration is included as a reference.

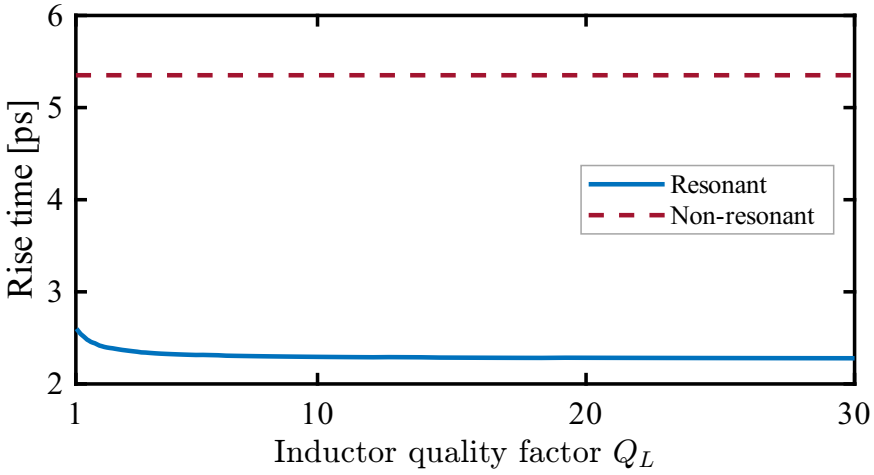


Figure 4.6: Output rise time evaluated between OUT_p and OUT_n as a function of the inductor quality factor for the near-end resonant configuration at 27.5 GHz. The corresponding non-resonant configuration is included as a reference.

factor on the rise time is limited. Most of the variation occurs at low values of Q_L, whereas at higher Q_L the rise time changes only slightly. As Q_L increases, the curve progressively approaches an almost constant value, indicating that further reductions in the inductor losses provide only marginal improvements. Increasing Q_L reduces the inductor losses but leaves the transmission-line losses unchanged, making further reductions in rise time progressively small. The resonant configuration therefore provides a marked improvement over the non-resonant reference even at low Q_L, while further reductions in the inductor losses have little effect on the output rise time at high Q_L.

The output random jitter is shown in Figure 4.7. The near-end resonant configuration exhibits lower jitter than the corresponding non-resonant reference, confirming the benefit of resonant operation also in terms of timing uncertainty. The short rise time produces a steep signal slope at the switching threshold, reducing the timing variation caused by

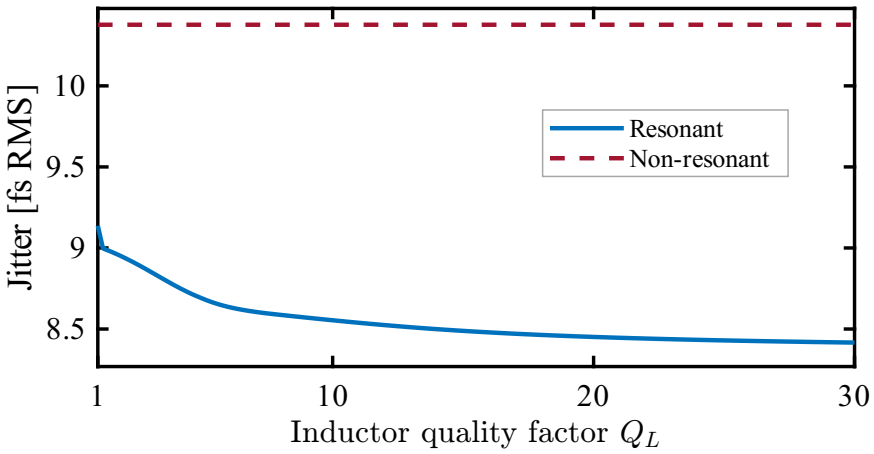


Figure 4.7: Output random jitter evaluated between OUT_p and OUT_n as a function of the inductor quality factor for the near-end resonant configuration at 27.5 GHz. The corresponding non-resonant configuration is included as a reference.

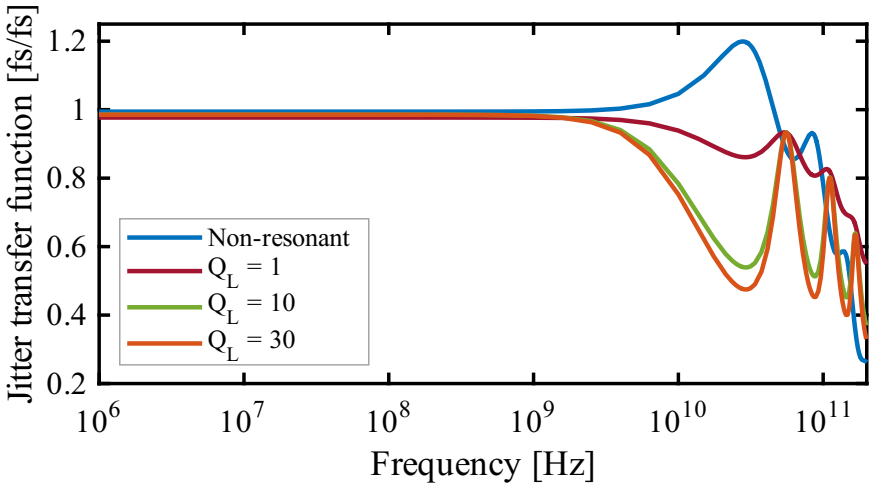


Figure 4.8: Jitter transfer function magnitude of the near-end resonant configuration as a function of the perturbation frequency for three representative values of the inductor quality factor. The corresponding non-resonant configuration is included as a reference.

voltage fluctuations. A reduction in voltage noise also contributes to the lower jitter. The effect of resonant operation on voltage noise is investigated in Section 4.3.1 through the power spectral density analysis of the final 1 mm near-end implementation.

The dependence on the inductor quality factor is weak. The jitter decreases only slightly as Q_L increases, with most of the variation confined to the low-Q_L region. As discussed for the rise time, at high Q_L the unchanged transmission-line losses limit the benefit of reducing the inductor series resistance. Consequently, the signal slope at the switching threshold changes little, limiting further reductions in the timing uncertainty caused by voltage noise.

Figure 4.8 compares the jitter transfer function response of the near-end resonant architecture with the corresponding non-resonant configuration for three representative values of the inductor quality factor. The main difference appears in the perturbation-

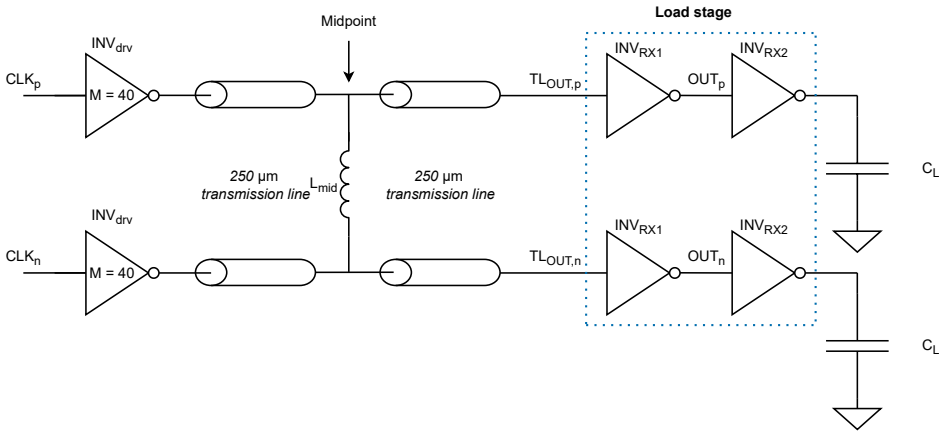


Figure 4.9: Midpoint resonant configuration, with the shunt inductive element connected between two 250 μm transmission-line sections.

frequency region where the non-resonant response exceeds unity and reaches its amplification peak. In the same region, all three resonant cases remain below unity and develop an attenuation valley. Resonant operation therefore changes the behavior of the distribution path from jitter amplification to jitter attenuation in the region where the non-resonant implementation is most critical.

The reduction in the output rise time shown in Figure 4.6 allows the signal to approach its logic levels earlier within each clock half-cycle. Consequently, a timing perturbation affecting one edge has less influence on the signal voltage at the beginning of the subsequent transition, limiting the propagation of timing errors across successive clock edges. The attenuation observed over specific perturbation-frequency ranges is instead associated with the frequency-selective response of the resonant network.

The impact of the inductor quality factor on performance is also more evident here than for the previously analyzed performance metrics. Increasing Q_L reduces the damping and makes the resonant response more selective, producing a clear deepening of the attenuation valley. The overall frequency dependence remains similar for the three resonant cases, while the main variation is concentrated around the attenuation region. Away from this region, the curves progressively approach unity and the differences among the considered quality factors are small.

4.2.2 Midpoint resonant configuration

The second architecture considered is the midpoint resonant configuration, shown in Figure 4.9. In this implementation, the shunt inductive element is connected at the center of the 500 μm transmission line. The interconnect is therefore divided into two 250 μm sections.

At the node labeled *Midpoint* in Figure 4.9, the two halves of the transmission line provide two frequency-dependent reactive contributions. Looking from this node toward

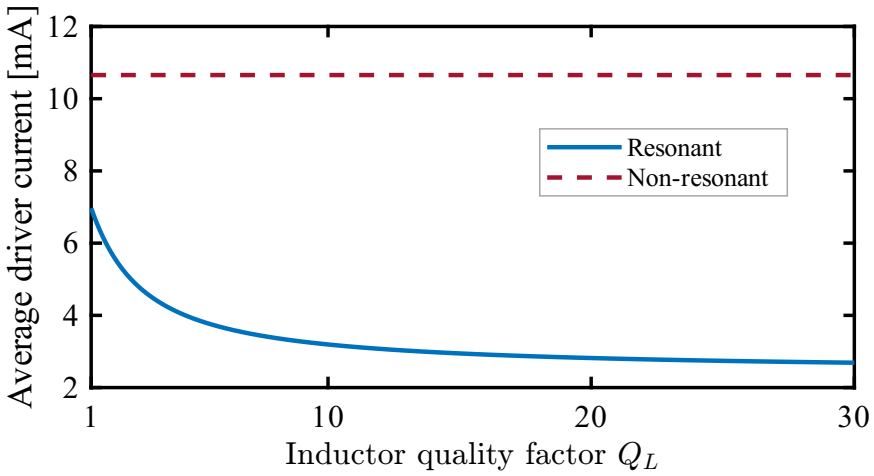


Figure 4.10: Average driver current as a function of the inductor quality factor for the midpoint resonant configuration at 27.5 GHz. The corresponding non-resonant configuration is included as a reference.

the receiving stage, the receiver input capacitance is transformed through the right-hand 250 μm transmission-line section. Looking from the same node toward the driver, the output network of the final CMOS stage is transformed through the left-hand 250 μm section. Applying the same transmission-line impedance transformation used for the near-end configuration gives the two equivalent capacitive contributions $C_{\mathrm{mid,R}}(\omega)$ and $C_{\mathrm{mid,D}}(\omega)$. Since these contributions appear in parallel at the node labeled *Midpoint*, the total equivalent capacitance is

$$C_{\mathrm{eq,mid}}(\omega) = C_{\mathrm{mid,R}}(\omega) + C_{\mathrm{mid,D}}(\omega). \tag{4.6}$$

Using the same parallel-resonance condition introduced for the near-end configuration in (4.4), the required midpoint inductance is given by

$$L_{\mathrm{mid}} = \frac{1}{\omega_0^2 C_{\mathrm{eq,mid}}(\omega_0)}. \tag{4.7}$$

Following the same approach adopted for the near-end configuration, the midpoint architecture is evaluated in terms of average driver current, output voltage swing, rise time, random jitter, and jitter-amplification behavior. The corresponding non-resonant implementation is included as a reference, while the inductor quality factor is varied to assess the impact of the resonant-branch losses over system performance.

As shown in Figure 4.10, the midpoint resonant configuration exhibits a lower average driver current than the non-resonant reference over the entire investigated range of Q_L, showing that the reduction in current consumption is preserved also for this resonance strategy. More specifically, the current decreases as Q_L increases, with most of the variation occurring at low quality factors, and progressively approaches a limiting value as the inductor losses become small. This trend is consistent with that already observed for the

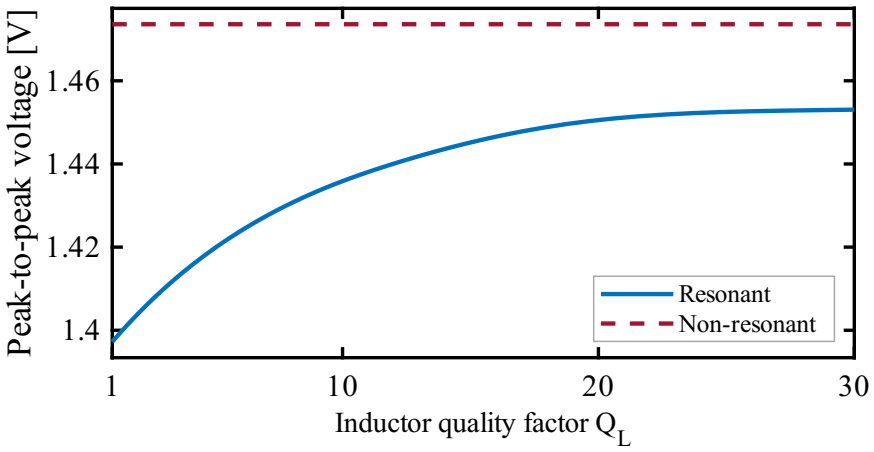


Figure 4.11: Differential peak-to-peak voltage evaluated between OUT_p and OUT_n as a function of the inductor quality factor for the midpoint resonant configuration at 27.5 GHz. The corresponding non-resonant configuration is included as a reference.

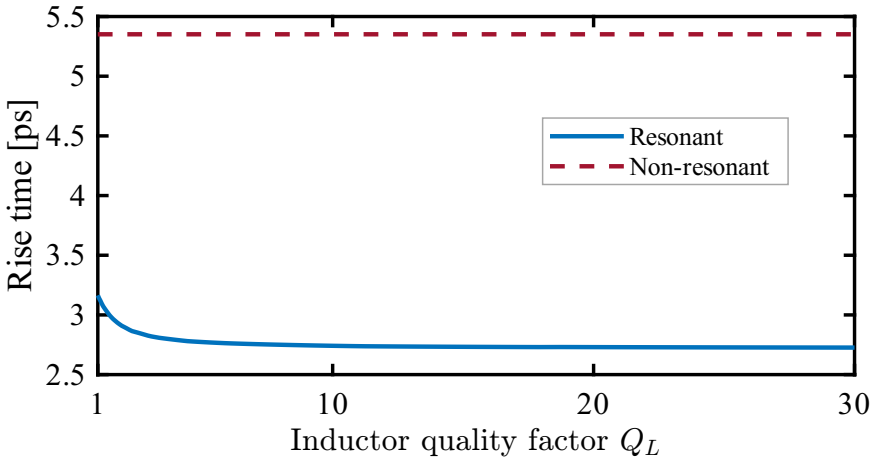


Figure 4.12: Output rise time as a function of the inductor quality factor for the midpoint resonant configuration at 27.5 GHz. The corresponding non-resonant configuration is included as a reference.

near-end configuration.

The output peak-to-peak voltage shows only a limited dependence on the inductor quality factor. Most of the variation is confined to the low-Q_L region, while the voltage swing rapidly approaches an almost constant value as Q_L increases. The resonant configuration exhibits a slightly lower output swing than the corresponding non-resonant reference, but the overall variation with the inductor losses remains small.

A more significant advantage is observed in the output rise time. The midpoint resonant configuration provides faster transitions than the corresponding non-resonant case over the investigated range. Similarly to the near-end configuration, the dependence on Q_L is relatively weak.

The random jitter evaluated between OUT_p and OUT_n is shown in Figure 4.13. The midpoint resonant configuration remains below the corresponding non-resonant reference

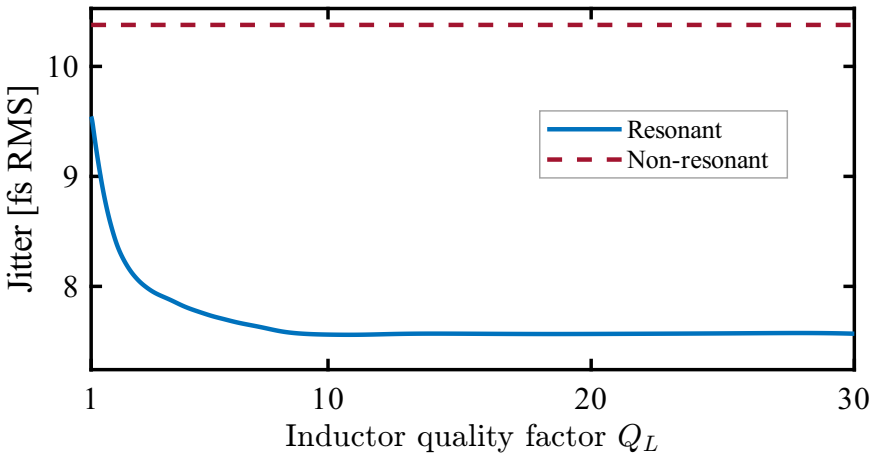


Figure 4.13: Output random jitter evaluated between OUT_p and OUT_n as a function of the inductor quality factor for the midpoint resonant configuration at 27.5 GHz. The corresponding non-resonant configuration is included as a reference.

over the entire investigated range, showing that the jitter reduction obtained with resonant operation is preserved also for the midpoint placement.

The dependence on the inductor quality factor is very weak, except at very low values of Q_L, where the jitter is noticeably higher than at moderate and high quality factor values. As Q_L increases from this region, the jitter rapidly decreases and then becomes nearly constant. This trend is expected because increasing Q_L initially reduces the contribution of the inductor losses to the damping of the resonant network. Once the transmission-line losses become dominant, however, they limit the effective quality factor of the overall network. Further increases in Q_L therefore have little effect on the resonant response and, consequently, on the output random jitter.

The jitter-amplification response of the midpoint configuration is shown in Figure 4.14. As already observed for the near-end architecture, the non-resonant response exhibits several peaks, one of which exceeds unity. Around the frequency of this amplification peak, the resonant cases remain below unity and develop an attenuation valley. Increasing Q_L makes the attenuation more pronounced, while the differences among the resonant curves become smaller away from this region.

Overall, the midpoint configuration preserves the main advantages introduced by resonant operation. Compared with the non-resonant reference, it provides lower average current consumption, faster output transitions, lower random jitter, and reduced jitter amplification. These results confirm that the benefits of resonant operation are preserved when the resonant element is placed at the midpoint of the transmission line. A direct comparison of the Near-End, Midpoint, and Far-End resonant configurations is presented in Section 4.2.4, where the influence of the inductor placement on the different performance metrics is discussed.

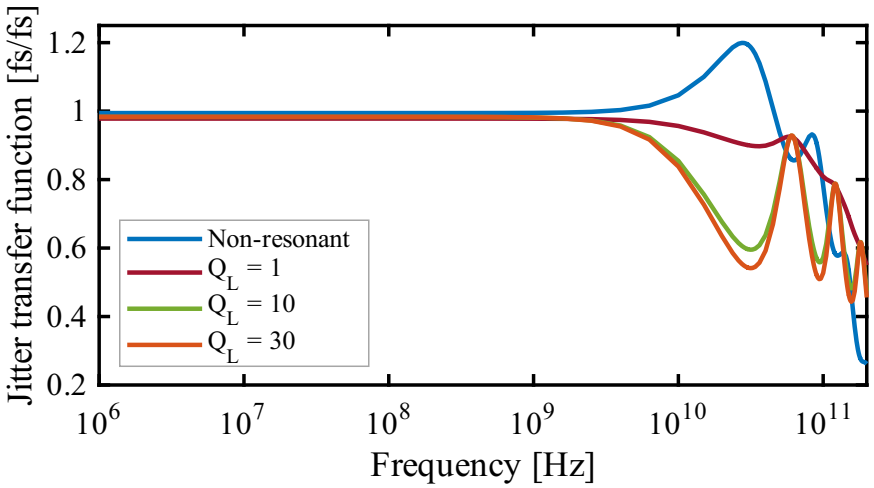


Figure 4.14: Jitter transfer function of the midpoint resonant configuration as a function of the input jitter frequency, for three representative values of the inductor quality factor at a clock frequency of 27.5 GHz. The corresponding non-resonant configuration is included as a reference.

4.2.3 Far-end resonant configuration

The third architecture considered is the far-end resonant configuration, shown in Figure 4.15, in which the shunt inductive element is connected at the output of the 500 μm transmission line, directly at the input node of the receiving stage.

At the far-end node, the receiver input capacitance contributes directly to the local reactive load, while the output network of the final CMOS driver is observed through the complete transmission line. Applying the same transmission-line impedance transformation used for the previous configurations gives the equivalent capacitive contribution $C_{\mathrm{far,D}}(\omega)$ associated with the driver side. Denoting the local receiver input capacitance by $C_{\mathrm{far,R}}$, the total equivalent capacitance at the far-end node is

$$C_{\mathrm{eq,far}}(\omega) = C_{\mathrm{far,R}} + C_{\mathrm{far,D}}(\omega). \tag{4.8}$$

Using the same parallel-resonance condition introduced for the near-end configuration, the required far-end inductance is given by

$$L_{\mathrm{far}} = \frac{1}{\omega_0^2 C_{\mathrm{eq,far}}(\omega_0)}. \tag{4.9}$$

The far-end configuration completes the preliminary analysis of the three possible inductor placements. Its performance is assessed through the average driver current, output voltage swing, rise time, random jitter, and jitter transfer function, using the corresponding non-resonant configuration as a reference. The inductor quality factor is varied to determine how the losses of the resonant branch affect the results for this placement.

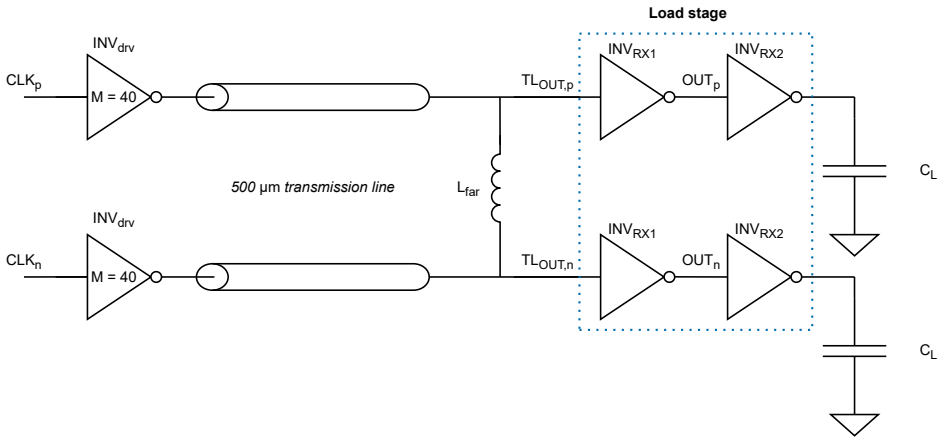


Figure 4.15: Far-end resonant configuration, with the shunt inductive element connected at the output of the 500 μm transmission line and in parallel with the input of the receiving stage.

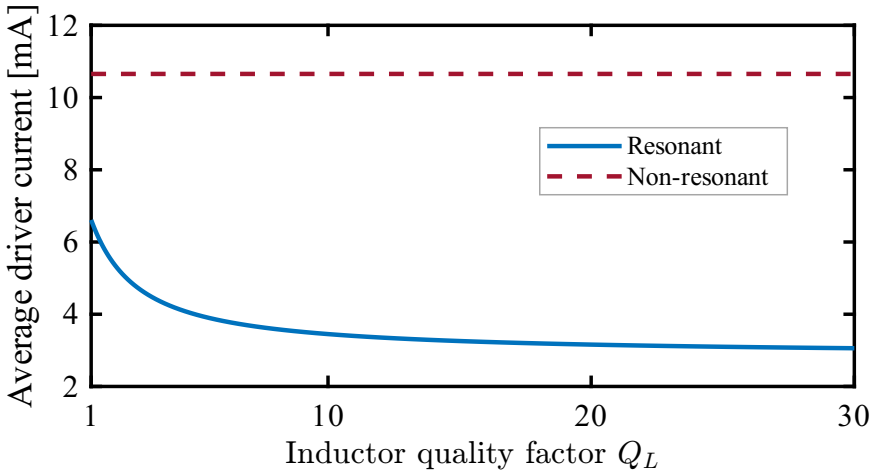


Figure 4.16: Average driver current as a function of the inductor quality factor for the far-end resonant configuration at 27.5 GHz. The corresponding non-resonant configuration is included as a reference.

As shown in Figure 4.16, the far-end resonant configuration exhibits a lower average driver current than the non-resonant reference over the entire investigated range of Q_L. The current decreases as Q_L increases, with most of the variation occurring at low quality factors, and progressively approaches a limiting value. This behavior is consistent with that observed for the previous resonant configurations.

The output peak-to-peak voltage exhibits a strong dependence on Q_L at low quality factors. Starting from a value close to the non-resonant reference at very low Q_L, the voltage swing decreases rapidly and then approaches an almost constant value. Therefore, at moderate and high quality factors, any further variations in Q_L would only have a limited influence on the output amplitude.

The far-end resonant configuration also provides a lower rise time between OUT_p and OUT_n than the corresponding non-resonant case over the entire investigated range.

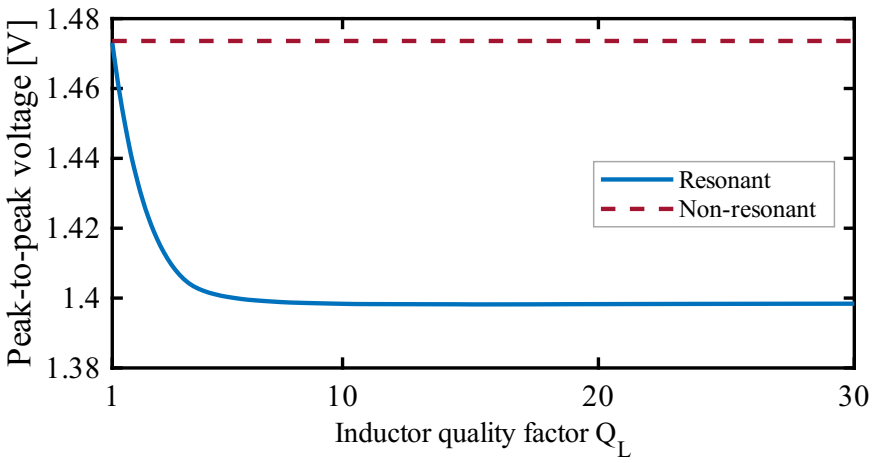


Figure 4.17: Differential output peak-to-peak voltage evaluated between OUT_p and OUT_n as a function of the inductor quality factor for the far-end resonant configuration at 27.5 GHz. The corresponding non-resonant configuration is included as a reference.

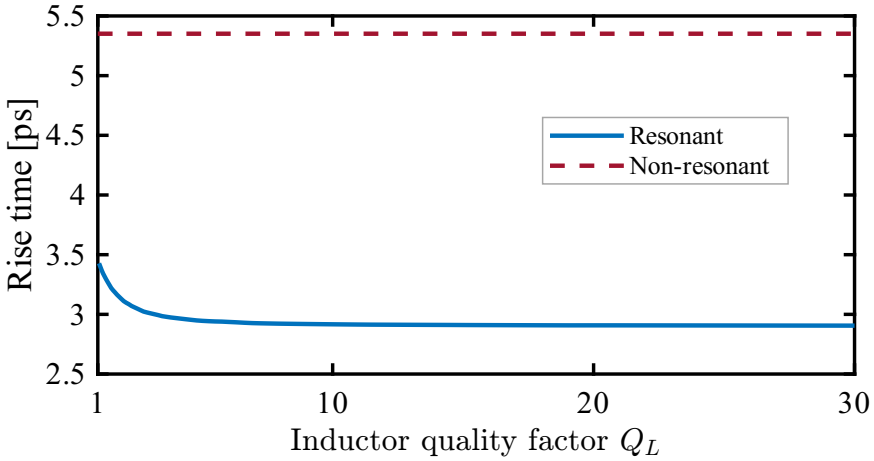


Figure 4.18: Rise time evaluated between OUT_p and OUT_n as a function of the inductor quality factor for the far-end resonant configuration. The corresponding non-resonant configuration is included as a reference.

The rise time decreases rapidly as Q_L increases from very low values and subsequently approaches an almost constant value, showing that the influence of the inductor losses becomes weak at moderate and high quality factors.

The output random jitter of the far-end configuration is shown in Figure 4.19. At very low values of Q_L, the resonant configuration exhibits higher jitter than the corresponding non-resonant approach, despite the reduction in rise time observed in Figure 4.18. As Q_L increases, the jitter rapidly decreases, falls below the non-resonant level, and then approaches an approximately constant value at moderate and high quality factors. As in the previous configurations, the plateau is due to the transmission-line losses limiting the effective quality factor of the resonant network.

Figure 4.20 shows the far-end jitter-transfer response. In the frequency interval where the non-resonant transfer magnitude exceeds unity, all resonant responses remain below

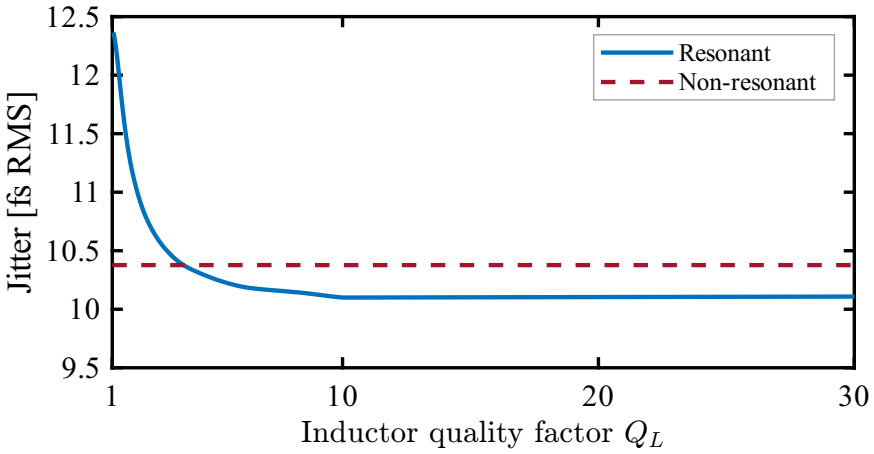


Figure 4.19: Random jitter evaluated between OUT_p and OUT_n as a function of the inductor quality factor for the far-end resonant configuration. The corresponding non-resonant configuration is included as a reference.

unity, so the input timing perturbation is attenuated rather than amplified. Increasing Q_L lowers the minimum transfer magnitude in this interval, thereby increasing the attenuation depth.

Overall, the far-end configuration preserves the main advantages of resonant operation once very low quality factors are excluded. The average current and rise time are reduced with respect to the non-resonant implementation, while the critical jitter-amplification region is replaced by attenuation. The strongest sensitivity to the inductor losses is observed at very low Q_L, particularly in the output voltage swing and random jitter, whereas the performance becomes only weakly dependent on Q_L at moderate and high quality factor values.

4.2.4 Comparison of the resonant placements

Table 4.1 includes the main performance metrics of the three resonant configurations at the nominal operating frequency of 27.5 GHz for $Q_L = 1$, 10, and 30. These values cover a wide range of inductor losses, from the strongly lossy case represented by $Q_L = 1$ to the nearly ideal-inductor limit represented by $Q_L = 30$. The relevant non-resonant implementation is reported for the sake of completeness.

The comparison shows a substantial reduction in average current for all three resonant placements, with the Midpoint configuration providing the lowest consumption at moderate and high quality factors. The rise time is also markedly reduced with respect to the non-resonant reference for all three placements. As discussed in Chapter 3, the capacitive loading presented at the transmission-line input makes the transition at the output of INV_{DRV} particularly critical. In the Near-End configuration, the inductive branch compensates this loading directly at the driver output. As the inductor is placed farther from this node, its compensation of the load seen by INV_{DRV} becomes less direct, consistently

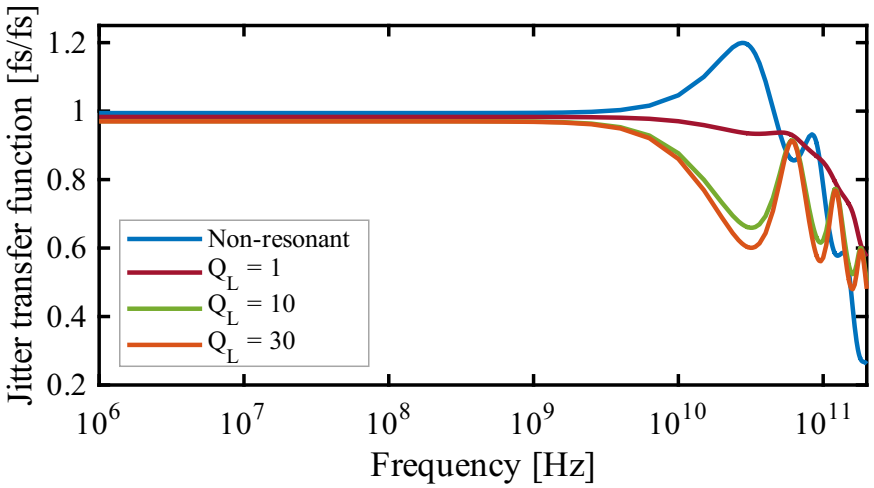


Figure 4.20: Jitter transfer function of the far-end resonant configuration as a function of the input jitter frequency, for three representative values of the inductor quality factor at a clock frequency of 27.5 GHz. The corresponding non-resonant configuration is included as a reference.

with the progressive increase in output rise time from the Near-End to the Midpoint and Far-End configurations.

The differential peak-to-peak voltage changes only moderately among the considered implementations. This is expected since the signal at nodes OUT_p and OUT_n is already well established, leaving limited room for substantial variations in voltage swing. Larger differences are instead visible in random jitter. The Near-End and Midpoint configurations remain below the non-resonant reference, with the Midpoint providing the lowest jitter at moderate and high quality factors.

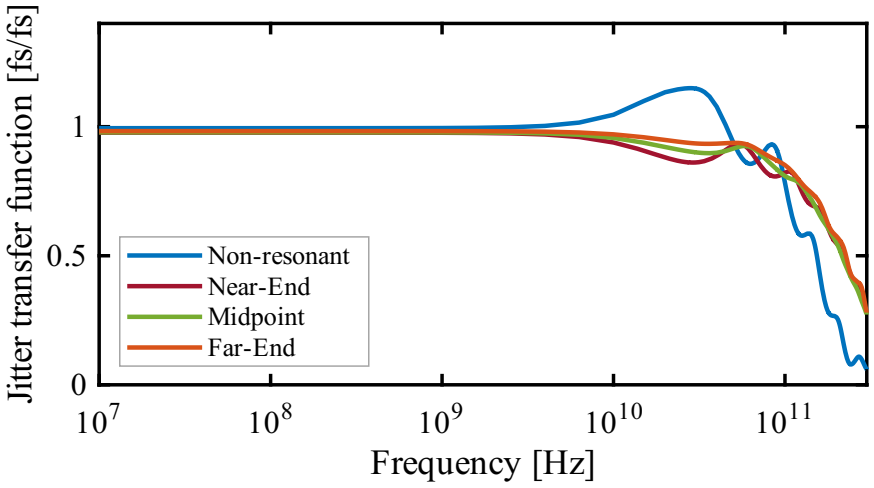
The jitter-amplification comparison is reported in Figure 4.21. For all three resonant placements, the amplification peak exhibited by the non-resonant reference is removed and replaced by an attenuation region. This directly addresses the main limitation identified for the non-resonant implementation, namely its tendency to amplify input timing perturbations over the critical frequency range. Increasing the inductor quality factor makes the attenuation more pronounced, although the additional improvement becomes progressively smaller as the inductor losses are reduced. Among the considered placements, the Near-End configuration provides the strongest attenuation, followed by the Midpoint and Far-End cases.

Overall, the preliminary study confirms that resonant operation provides a clear improvement over the non-resonant architecture. The main benefit is associated with the introduction of resonance itself, while the dependence on the inductor quality factor is comparatively limited for most of the considered metrics. Together with the reduction in current consumption and rise time and the suppression of the critical jitter-amplification peak, these results support the resonant approach and motivate its evaluation in the complete 1 mm implementations over the target operating-frequency range. No single

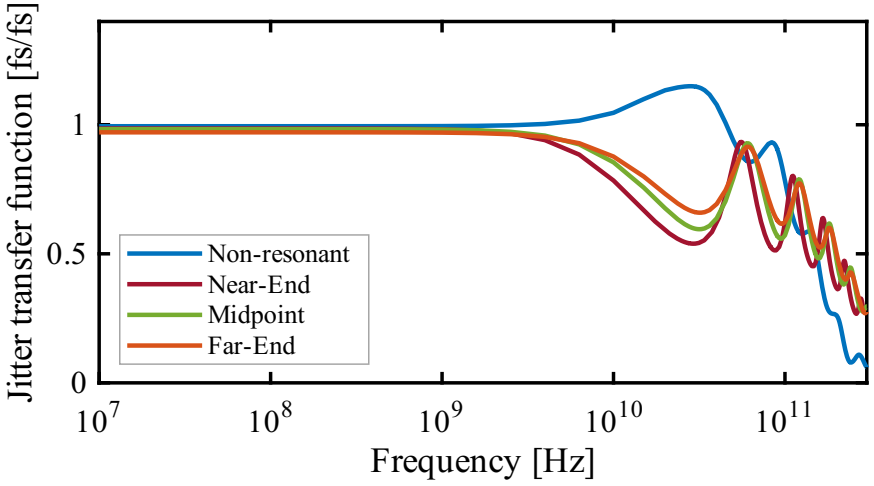
Table 4.1: Comparison of the preliminary 500 μm resonant configurations at 27.5 GHz for $Q_L = 1, 10$, and 30. The corresponding non-resonant implementation is included as a common reference.

Metric	Q_L	Non-resonant	Near-End	Midpoint	Far-End
Resonant inductance [pH]	–	–	272	300	265
Average current [mA]	1	10.65	7.61	6.96	6.61
	10	10.65	3.96	3.19	3.45
	30	10.65	3.07	2.50	2.92
V_{pp} [V]	1	1.474	1.486	1.397	1.473
	10	1.474	1.508	1.436	1.398
	30	1.474	1.514	1.455	1.399
Rise time [ps]	1	5.35	2.60	3.16	3.46
	10	5.35	2.29	2.74	2.92
	30	5.35	2.27	2.72	2.90
Jitter [fs RMS]	1	10.38	9.14	9.54	12.36
	10	10.38	8.55	7.56	10.10
	30	10.38	8.37	7.39	9.99

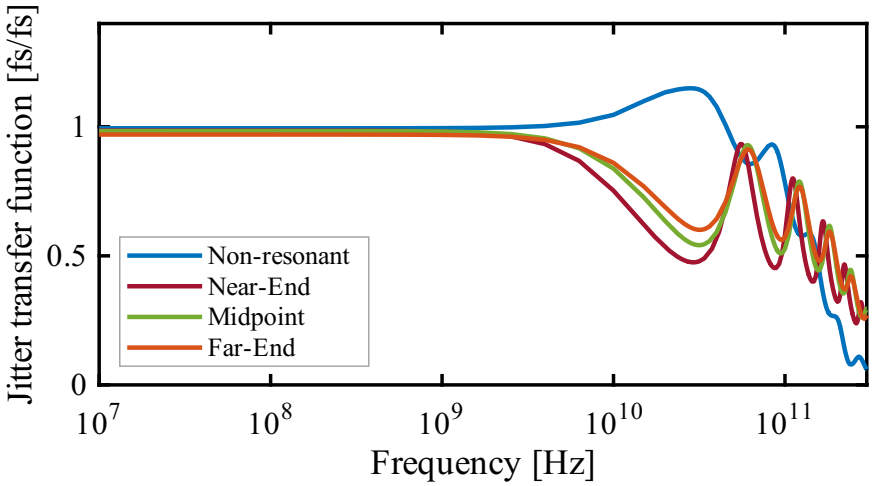
resonant placement provides the most favorable results for all the considered metrics. The Near-End configuration exhibits the lowest rise time and the strongest jitter attenuation, whereas the Midpoint configuration provides the lowest average current and random jitter at moderate and high Q_L. These results identify Near-End and Midpoint as the most promising placements, whose performance will be further assessed in the complete 1 mm implementations.



(a)



(b)



(c)

Figure 4.21: Jitter-amplification comparison of the non-resonant, near-end, midpoint, and far-end configurations for (a) $Q_L = 1$, (b) $Q_L = 10$, and (c) $Q_L = 30$.

4.3 Case study: 1 mm transmission line

The previous analysis showed that resonant operation can provide a performance advantage with respect to the corresponding non-resonant implementation. However, the resonant configurations were mainly characterized at the nominal operating frequency of 27.5 GHz, where the inductance was selected to establish the resonant condition.

The final analysis is therefore performed starting directly from the 1 mm single-segment non-resonant implementation introduced in Chapter 3. The CMOS driving and receiving stages, the load conditions, and the total transmission-line length are kept unchanged. The resonant implementations are obtained by introducing in the circuit described in Chapter 3 an appropriately sized shunt inductance at the near end, midpoint, or far end of the same distribution path. In this way, the effect of resonance can be evaluated while keeping the other circuit conditions unchanged.

For each configuration, the inductance is selected to establish resonance around the nominal operating frequency of 27.5 GHz. The resulting implementations are then characterized over the complete target range from 25 to 30 GHz and compared with the same non-resonant reference. This analysis makes it possible to verify whether the advantages observed at resonance are preserved over the required operating range and to evaluate the impact of the resonant-element placement on the overall clock-distribution performance.

4.3.1 Near-end resonant implementation

Following the same approach adopted in the previous analysis, the near-end configuration is considered first. The complete 1 mm resonant implementation is shown in Figure 4.22. The inductance value is estimated using (4.4), with the equivalent capacitance evaluated at the near-end node. This results in a nominal inductance of 110 pH.

As observed in the near-end analysis presented in Section 4.2.1, the resonant network can increase the signal amplitude delivered to the receiving stage. INV_{RX1} regenerates the signal delivered by the transmission line. A large differential swing at its input improves the regenerated transition at OUT_p and OUT_n, even though the final voltage swing is already close to its full value. The inductance is therefore varied at the nominal operating frequency of 27.5 GHz, and the differential peak-to-peak voltage between $TL_{\mathrm{OUT},p}$ and $TL_{\mathrm{OUT},n}$ is evaluated as a function of the inductance. The resulting sweep confirms that the nominal value of 110 pH lies in the peak-voltage region and is therefore selected for the subsequent analysis.

Using the selected inductance of 110 pH, the near-end resonant implementation is characterized over the complete target frequency range and compared with the corresponding non-resonant reference. The differential peak-to-peak voltage evaluated between OUT_p and OUT_n, shown in Figure 4.24, exhibits only a limited variation between the resonant and non-resonant implementations over the target frequency range. This behavior is

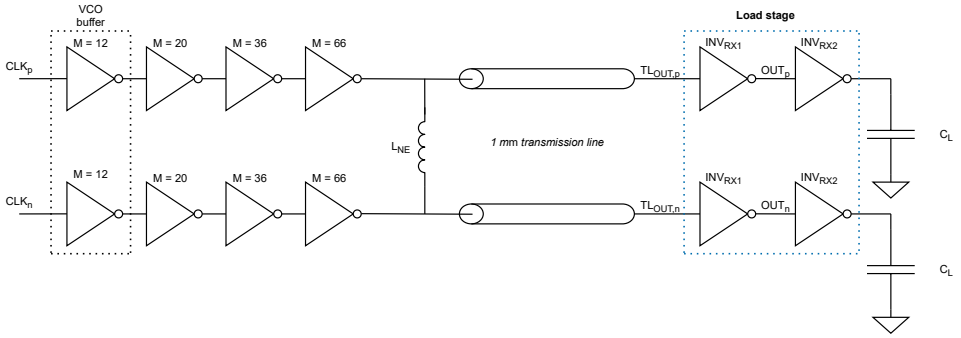


Figure 4.22: Near-end resonant implementation for the 1 mm transmission line.

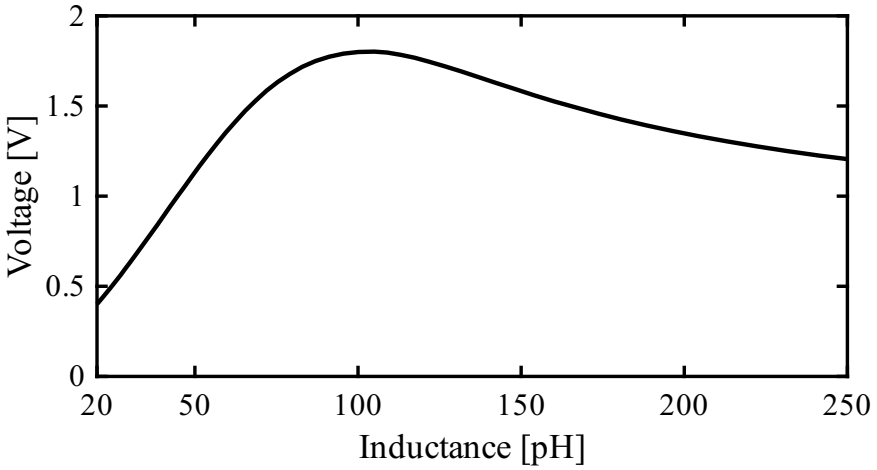


Figure 4.23: Differential peak-to-peak voltage evaluated between $TL_{\mathrm{OUT},p}$ and $TL_{\mathrm{OUT},n}$ as a function of the inductance for the near-end resonant implementation at 27.5 GHz.

mainly related to the regenerative action of the receiving chain. In the non-resonant implementation, the waveform is already restored close to its full differential voltage swing, leaving only limited margin for a further increase in V_{pp}. Therefore, the benefit introduced by the resonant network is not primarily reflected in the voltage swing evaluated between OUT_p and OUT_n.

An improvement is observed in the rise time evaluated between OUT_p and OUT_n, as shown in Figure 4.25. Within the 25–30 GHz target range, the resonant implementation provides faster transitions than the corresponding non-resonant configuration. The preliminary 500 μm study showed a pronounced separation between the resonant and non-resonant rise times. In the final 1 mm implementation, the driving chain was dimensioned to reduce jitter, and the non-resonant configuration already exhibits a short rise time. This limits the additional reduction achievable through resonance, which can be associated with the partial compensation of the reactive loading around resonance.

Moreover, the origin of this improvement can be further understood by observing the differential signal directly between $TL_{\mathrm{OUT},p}$ and $TL_{\mathrm{OUT},n}$, as shown in Figure 4.26. At these nodes, the effect of the resonant network is more pronounced, and a larger voltage swing is obtained around the target operating region. Although the receiving chain largely

masks this amplitude difference at OUT_p and OUT_n by regenerating the waveform toward its logic levels, the strong signal available at $TL_{\mathrm{OUT},p}$ and $TL_{\mathrm{OUT},n}$ provides favorable conditions for a rapid transition to take place. This behavior is consistent with the reduction in the rise time observed at OUT_p and OUT_n, while producing only a limited variation in the corresponding final V_{pp}.

The output random jitter is then evaluated over the complete operating-frequency range, as shown in Figure 4.27. Within the target operating range, the resonant implementation exhibits lower random jitter than the corresponding non-resonant one. This reduction is consistent with the faster transitions previously observed. To further inves-

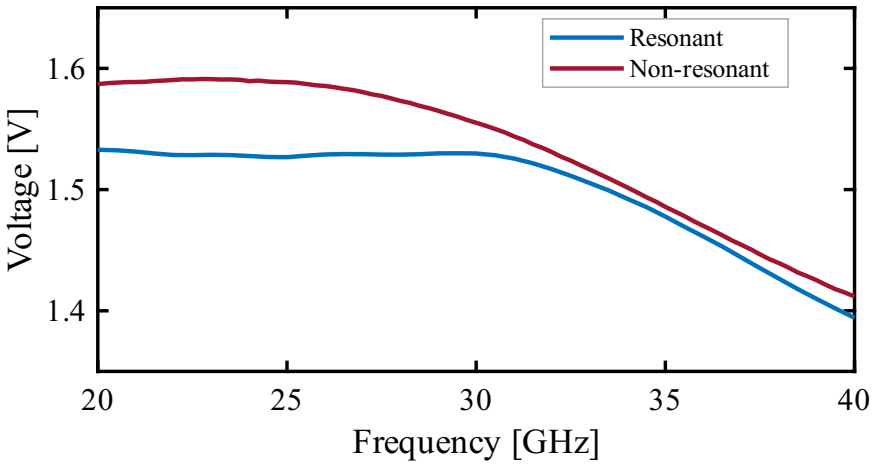


Figure 4.24: Differential peak-to-peak voltage evaluated between OUT_p and OUT_n as a function of the operating frequency for the near-end resonant and non-resonant implementations.

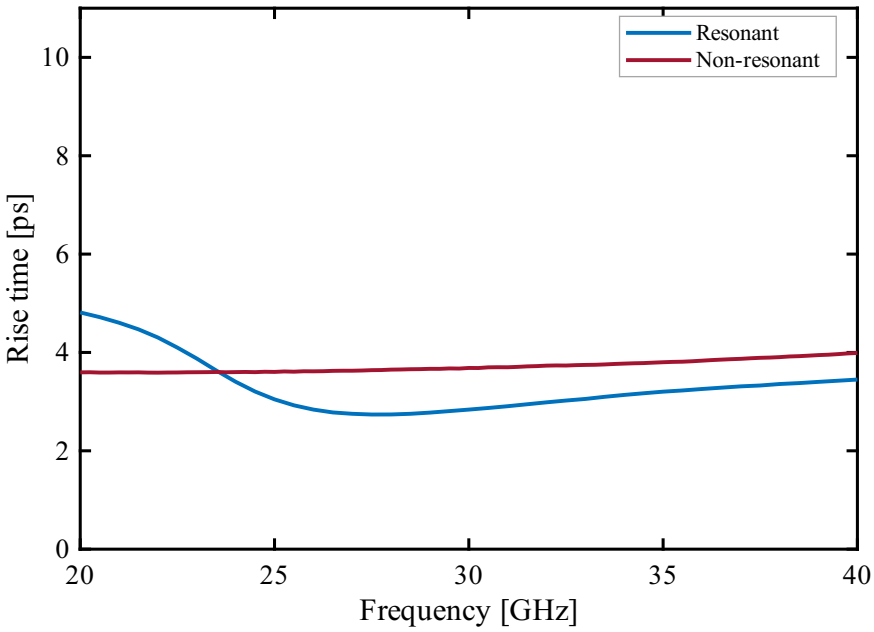


Figure 4.25: Rise time evaluated between OUT_p and OUT_n as a function of the operating frequency for the near-end resonant and non-resonant implementations.

tigate the origin of the random-jitter reduction, the voltage noise is also simulated at the differential transmission-line output $TL_{\mathrm{OUT},p}/TL_{\mathrm{OUT},n}$ and at the differential output OUT_p/OUT_n. This allows us to check whether resonant operation also reduces voltage noise, rather than improving jitter solely by enabling fast transitions. The corresponding voltage-noise spectra are shown in Figure 4.28.

At low frequencies, the shunt inductor presents a low impedance and attenuates noise propagating from the driver. At higher frequencies, the magnitude of the inductor impedance increases, reducing the shunting effect of the inductive branch; consequently, the noise spectrum progressively approaches that of the non-resonant implementation. At OUT_p/OUT_n, the resonant implementation still exhibits a slightly lower noise level, although the spectral shaping is no longer clearly visible. This behavior is consistent with the additional noise introduced by INV_{RX1}, which raises the noise floor at the output.

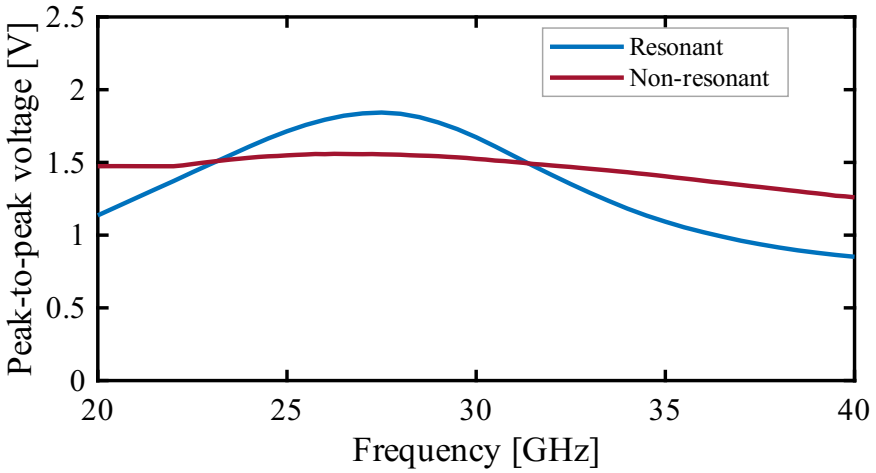


Figure 4.26: Differential peak-to-peak voltage evaluated between $TL_{\mathrm{OUT},p}$ and $TL_{\mathrm{OUT},n}$ as a function of the operating frequency for the near-end resonant and non-resonant implementations.

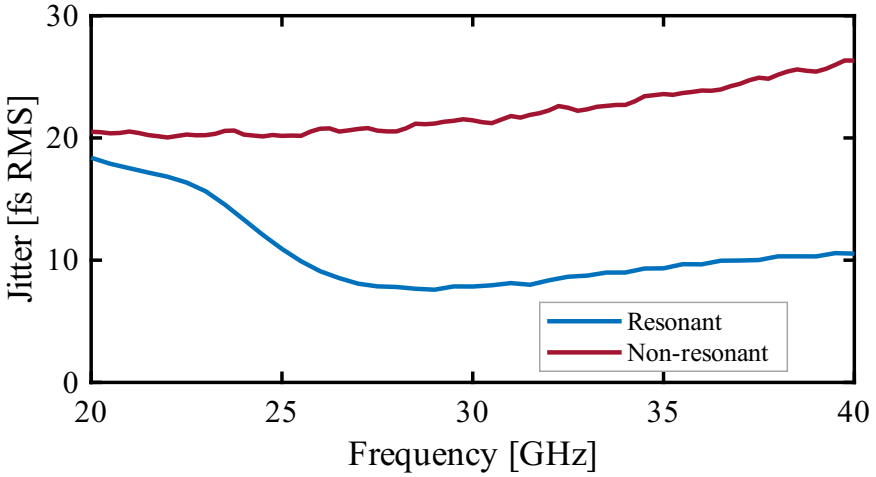
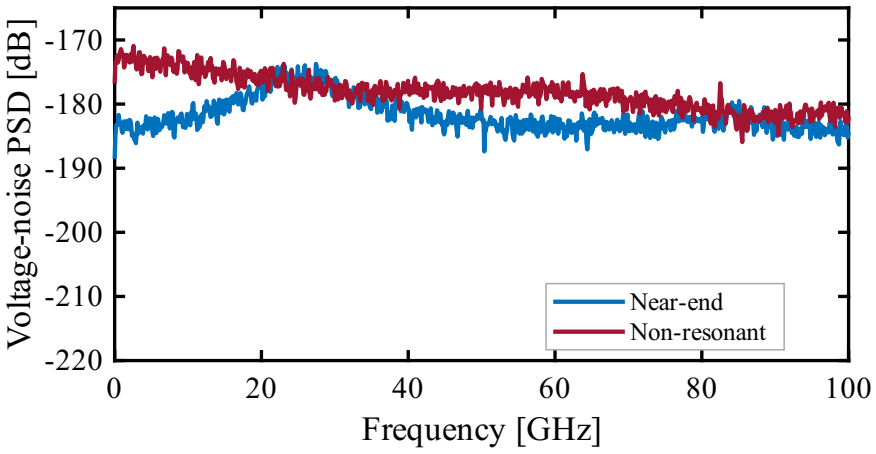


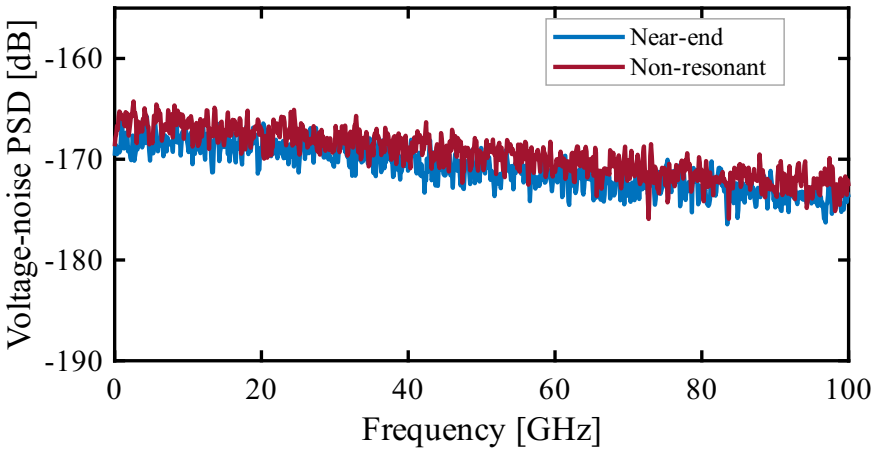
Figure 4.27: Output random jitter evaluated between OUT_p and OUT_n as a function of the operating frequency for the near-end resonant and non-resonant implementations.

The jitter transfer function is then considered in Figure 4.29. The non-resonant implementation exhibits peaks above unity over specific ranges of perturbation frequency, indicating that the timing perturbation is amplified through the distribution path. Conversely, the resonant implementation provides jitter attenuation over the corresponding frequency regions. This behavior can be related to the frequency-selective nature of the resonant network. Indeed, the clock signal is efficiently transferred around the intended operating condition, whereas perturbation components away from this region are less effectively propagated. Consequently, the resonant network reduces the transfer of these timing perturbations and suppresses the amplification observed in the non-resonant implementation.

The PSIJ response is then considered, as shown in Figure 4.30. This metric reflects the sensitivity of the complete distribution path to supply-voltage disturbances, since the



(a)



(b)

Figure 4.28: Voltage-noise power spectral density for the near-end resonant and non-resonant implementations, evaluated (a) at the differential transmission-line output $TL_{\mathrm{OUT},p}/TL_{\mathrm{OUT},n}$ and (b) at the differential regenerated output OUT_p/OUT_n. For readability, the voltage-noise power spectral density is reported in dB, while the corresponding linear quantity is expressed in V^2/Hz.

perturbation affects the switching behavior and propagation delay of the different CMOS stages along the chain. Its frequency dependence therefore results from the combined response of the active stages and the resonant network rather than from a single circuit element.

As shown in Figure 4.30, the resonant implementation exhibits a clear reduction in PSIJ at low perturbation frequencies. This behavior can be associated with the reduced dependence of the line-driving transition on the instantaneous current supplied by the active stage, since part of the reactive energy is exchanged within the resonant network. The frequency-selective behavior introduced by the resonant branch can also limit the transfer of disturbance components through the distribution path. As the perturbation frequency increases, the impedance of the shunt inductance increases and its influence on the disturbance transfer becomes progressively weak. Consequently, the resonant and non-resonant responses progressively become similar at high perturbation frequencies.

Finally, the average current consumption is compared over the operating-frequency range, as shown in Figure 4.31. Around the nominal operating region, the resonant implementation exhibits a lower average current than the corresponding non-resonant reference. At low frequencies, the inductive reactance decreases, providing a low-impedance current path between the two differential driver outputs and increasing current consumption. This behavior is consistent with the near-end analysis presented in Section 4.2.1, where the resonant branch was shown to partially compensate the reactive loading seen by the line driver. As a result, part of the energy associated with the capacitive load is exchanged with the inductive element rather than being supplied entirely by the active driver during each transition.

The reduction is less pronounced than in the previous near-end analysis. In the present 1 mm implementation, the transmission line is driven through a tapered CMOS chain. Resonant compensation mainly reduces the reactive burden seen by the final stage directly connected to the transmission line, whereas the preceding stages of the tapered chain continue to switch and contribute to the overall current consumption. The benefit obtained at the line-driving stage is therefore only partially reflected in the current of the complete driving structure. Moreover, toward the upper end of the target frequency range, the compensation becomes less effective and the current advantage with respect to the non-resonant implementation progressively decreases. Outside the target range, this advantage is eventually lost as the operating frequency moves further away from resonance.

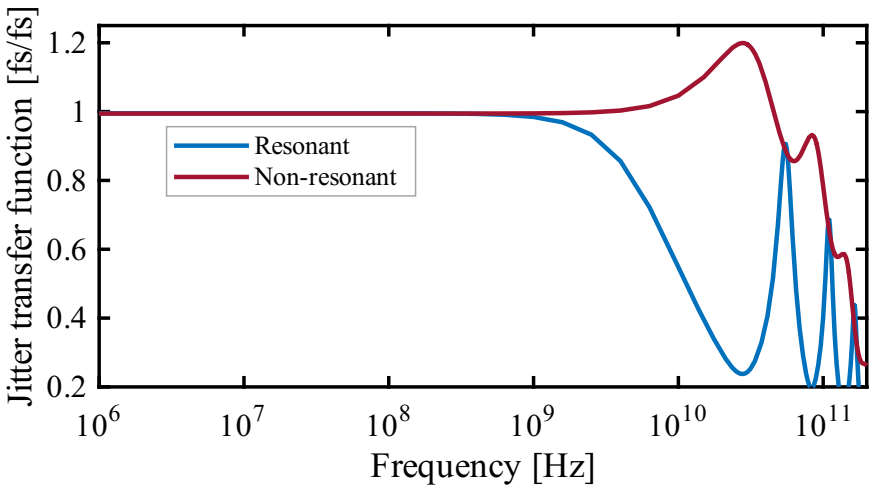


Figure 4.29: Jitter-amplification magnitude as a function of the signal frequency for the near-end resonant and non-resonant implementations.

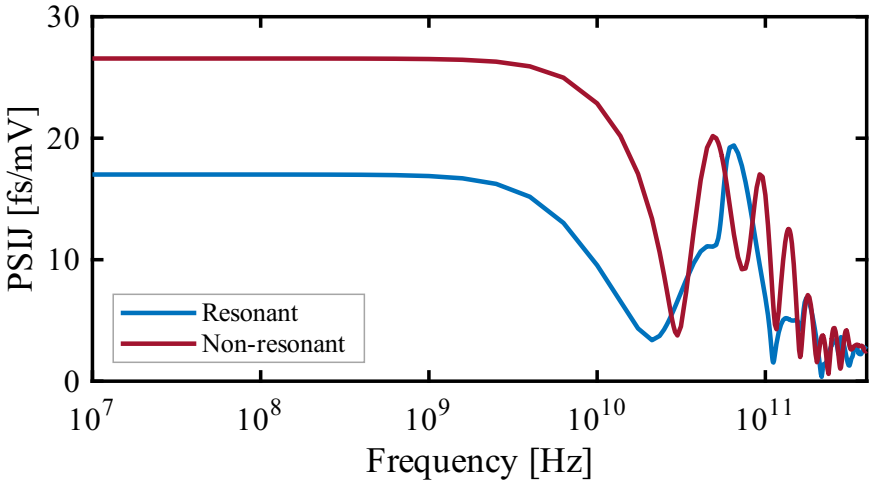


Figure 4.30: PSIJ as a function of the perturbation frequency for the near-end resonant and non-resonant implementations.

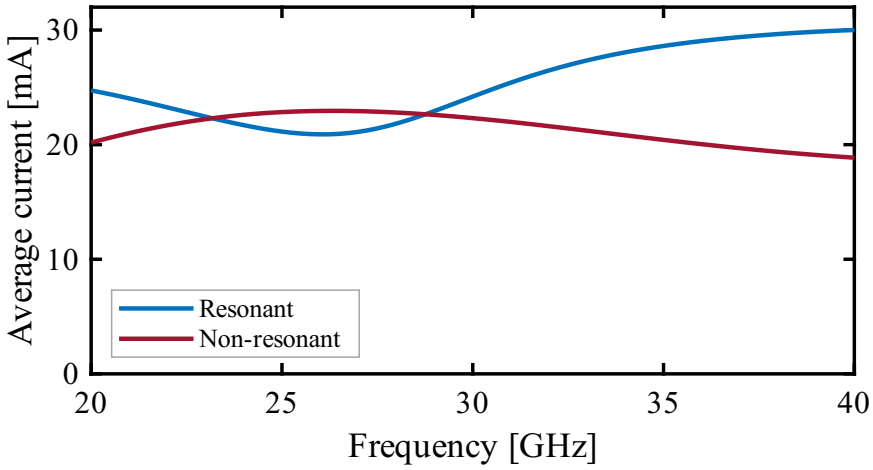


Figure 4.31: Average current consumption as a function of the operating frequency for the near-end resonant and non-resonant implementations.

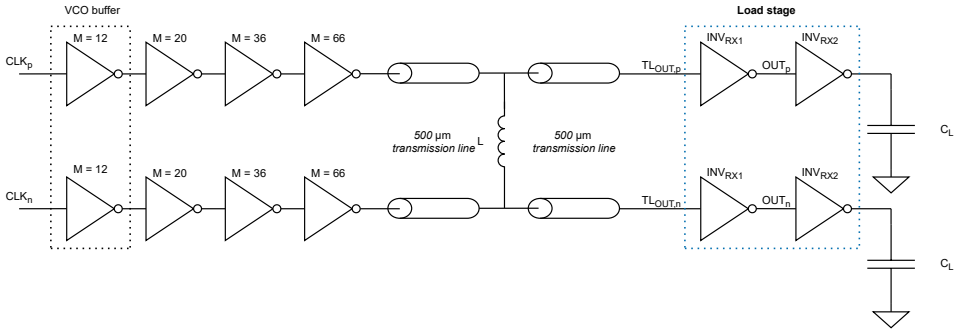


Figure 4.32: Midpoint resonant implementation for the 1 mm transmission line.

4.3.2 Midpoint resonant implementation

In this configuration, the shunt inductance is connected at the midpoint of the 1 mm transmission line, as shown in Figure 4.32. The CMOS driving and receiving stages, as well as the load conditions, are kept unchanged with respect to the corresponding non-resonant implementation.

The midpoint inductance is determined using (4.7), based on the equivalent reactive loading seen at the intermediate node. Evaluating this condition at the nominal operating frequency of 27.5 GHz leads to an inductance value close to 145 pH, which is adopted for the midpoint implementation. The configuration is then characterized over the complete target frequency range and compared with the same non-resonant reference used for the near-end analysis.

The differential peak-to-peak voltage evaluated between OUT_p and OUT_n shows only limited differences between the midpoint resonant and non-resonant implementations over the target frequency range, as shown in Figure 4.33. Therefore, similarly to the near-end case, the resonant configuration does not provide a significant advantage in terms of final regenerated voltage swing.

An improvement is observed in the output rise time, shown in Figure 4.34. Within the 25–30 GHz target range, the midpoint resonant implementation consistently provides faster transitions than the corresponding non-resonant reference. This result is consistent with the behavior already observed for the resonant configurations in the previous analysis, where resonance was shown to reduce the output rise time with respect to the non-resonant case.

The output random jitter remains approximately constant over the analyzed frequency range and is significantly lower than the corresponding non-resonant values, as shown in Figure 4.35. This behavior is consistent with the relatively limited variation observed in the output voltage swing and rise time over the same frequency range, as shown in Figure 4.33 and Figure 4.34.

As in the Near-End analysis, the voltage noise is examined to check whether resonant operation also reduces noise, in addition to improving the transition speed. The

voltage-noise spectra evaluated at $TL_{\mathrm{OUT},p}/TL_{\mathrm{OUT},n}$ and OUT_p/OUT_n are shown in Figure 4.36. As observed for the near-end implementation, the midpoint resonant configuration reshapes the noise spectrum at $TL_{\mathrm{OUT},p}/TL_{\mathrm{OUT},n}$, with a marked reduction of the low-frequency noise contribution with respect to the non-resonant case.

At OUT_p/OUT_n, the resonant implementation still exhibits a slightly lower noise level than the non-resonant implementation, while the spectral shaping is less evident, consistently with the behavior observed for the near-end implementation.

The jitter-amplification response is shown in Figure 4.37. As already observed for the near-end configuration, the non-resonant implementation exhibits an amplification peak

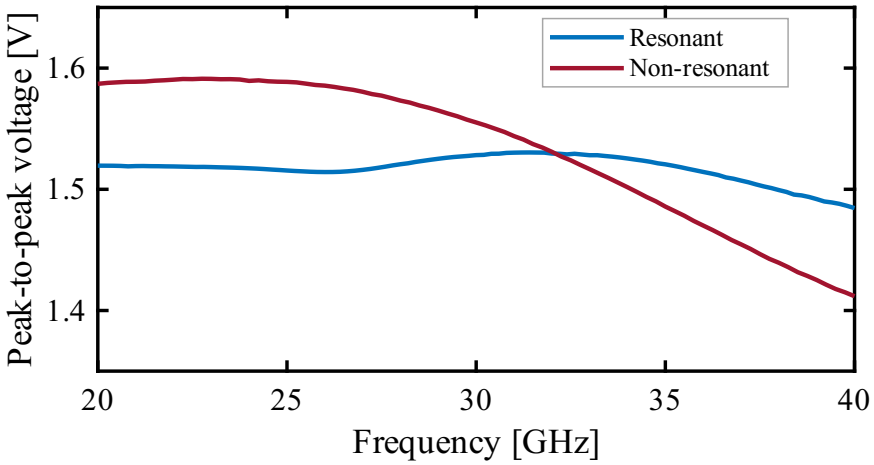


Figure 4.33: Differential peak-to-peak voltage evaluated between OUT_p and OUT_n as a function of the operating frequency for the midpoint resonant and non-resonant implementations.

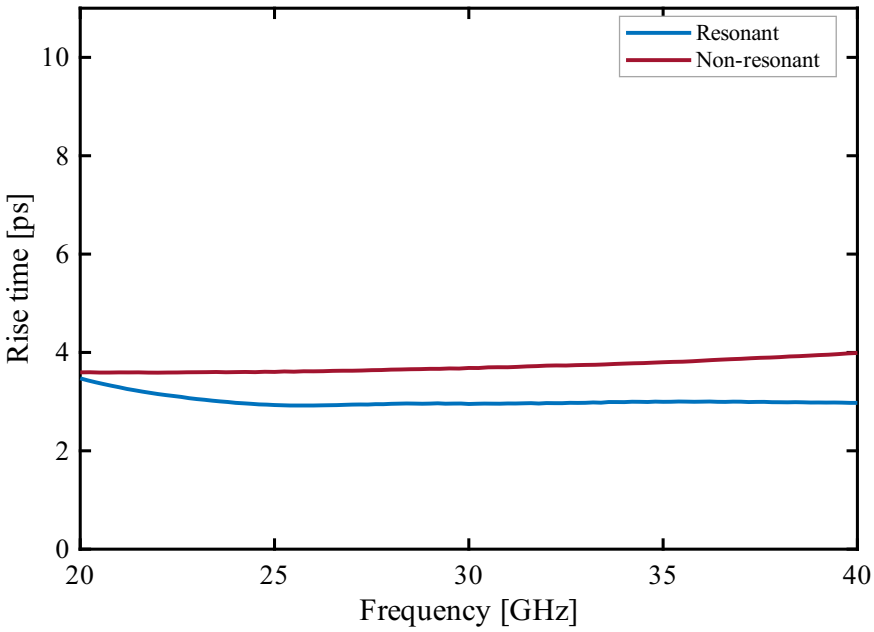


Figure 4.34: Rise time evaluated between OUT_p and OUT_n as a function of the operating frequency for the midpoint resonant and non-resonant implementations.

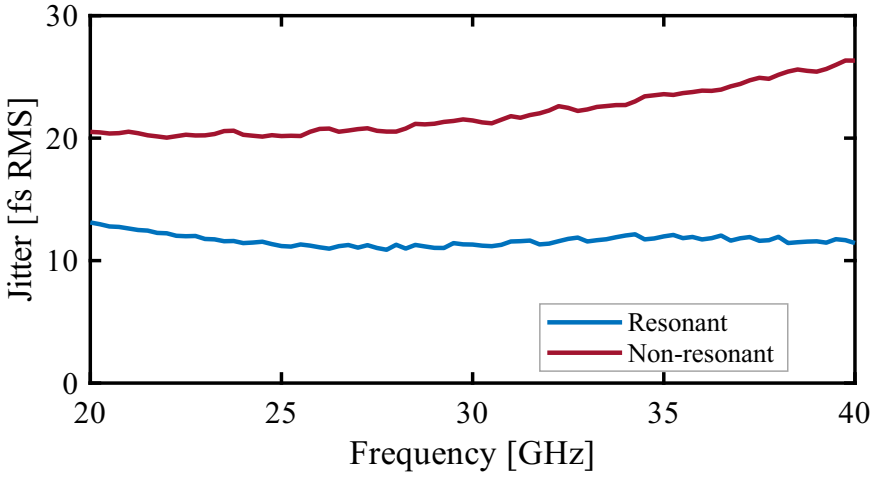
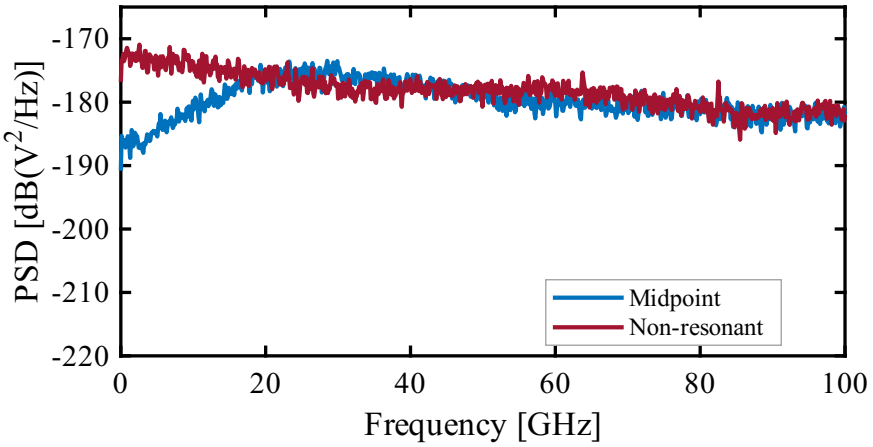
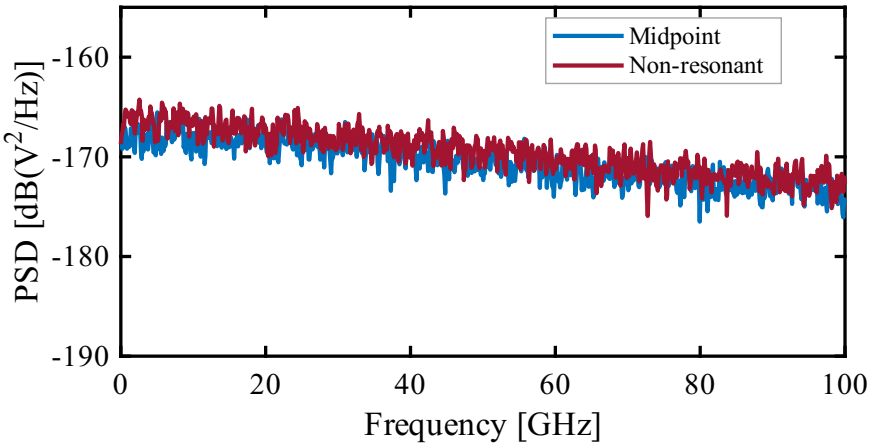


Figure 4.35: Output random jitter evaluated between OUT_p and OUT_n as a function of the operating frequency for the midpoint resonant and non-resonant implementations.



(a)



(b)

Figure 4.36: Voltage-noise power spectral density for the midpoint resonant and non-resonant implementations, evaluated (a) at $TL_{\mathrm{OUT},p}/TL_{\mathrm{OUT},n}$ and (b) at OUT_p/OUT_n. For readability, the voltage-noise power spectral density is reported in dB, while the corresponding linear quantity is expressed in V^2/Hz.

above unity, whereas the midpoint resonant implementation suppresses this peak and remains below unity over the corresponding frequency region. The midpoint placement therefore preserves the reduction in jitter amplification previously observed for resonant operation.

The PSIJ response is then evaluated for the midpoint configuration, as shown in Figure 4.38. Compared with the non-resonant reference, the midpoint resonant implementation exhibits a lower PSIJ over the low- and intermediate-frequency portions of the investigated perturbation-frequency range. At high perturbation frequencies, the difference between the two responses becomes small and the curves approach each other. The midpoint placement therefore preserves the reduction in sensitivity to supply disturbances observed for resonant operation.

Finally, the average current consumption is compared over the operating-frequency range, as shown in Figure 4.39. Within the 25–30 GHz target range, the midpoint resonant implementation requires less average current than the corresponding non-resonant reference. The difference progressively decreases toward the upper end of the target range. Above the target frequency range, the current advantage progressively decreases and is eventually lost as the operating frequency moves further away from resonance.

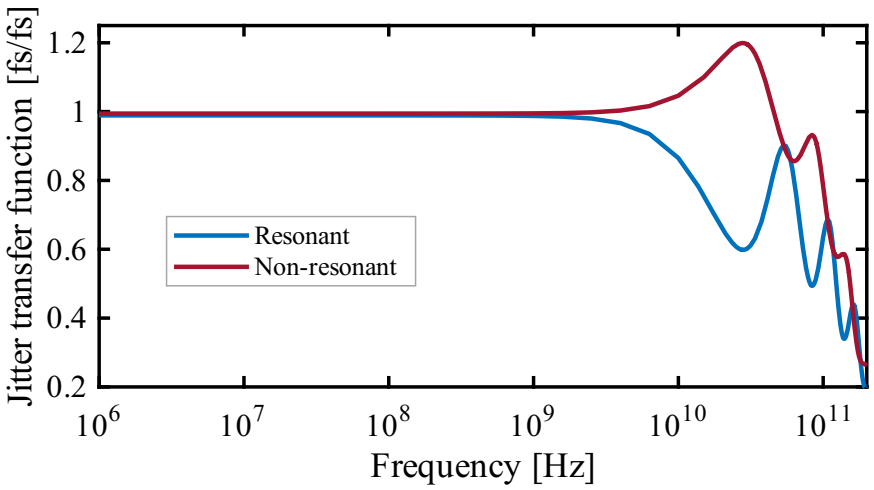


Figure 4.37: Jitter-amplification magnitude as a function of the perturbation frequency for the midpoint resonant and non-resonant implementations.

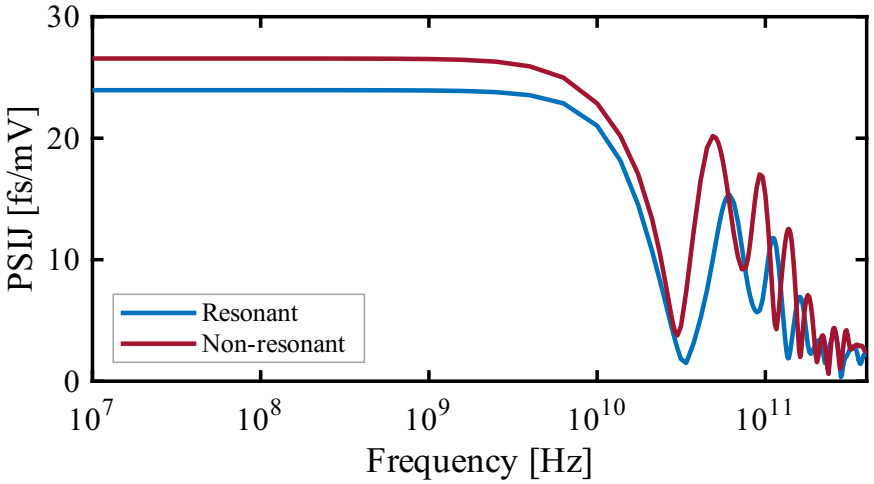


Figure 4.38: PSIJ as a function of the perturbation frequency for the midpoint resonant and non-resonant implementations.

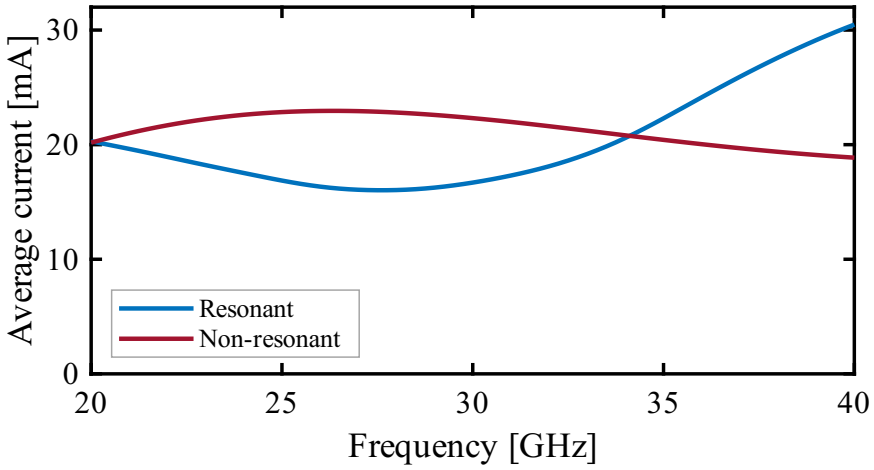


Figure 4.39: Average current consumption as a function of the operating frequency for the midpoint resonant and non-resonant implementations.

4.3.3 Far-end resonant implementation

The far-end resonant implementation is considered last. In this configuration, the shunt inductance is connected at the differential transmission-line output $TL_{\mathrm{OUT},p}/TL_{\mathrm{OUT},n}$, before INV_{RX1}. The CMOS driving and receiving stages, as well as the load conditions, are kept unchanged with respect to the corresponding non-resonant implementation. The complete 1 mm implementation is shown in Figure 4.40.

The inductance is determined using (4.9), with the equivalent capacitance evaluated at the transmission-line output. A value of 100 pH is adopted for the final implementation. The resulting configuration is then characterized over the 25–30 GHz target frequency range and compared with the same non-resonant reference used for the previous implementations.

The differential peak-to-peak voltage evaluated between OUT_p and OUT_n is shown in Figure 4.41. The far-end resonant implementation exhibits a lower output voltage swing than the corresponding non-resonant reference over the target frequency range. Nevertheless, the difference remains limited, and both responses remain relatively stable with frequency.

The output rise time is shown in Figure 4.42. Over the target frequency range, the far-end resonant implementation exhibits rise times comparable to the non-resonant reference, with a slight improvement around the nominal operating frequency and toward 30 GHz.

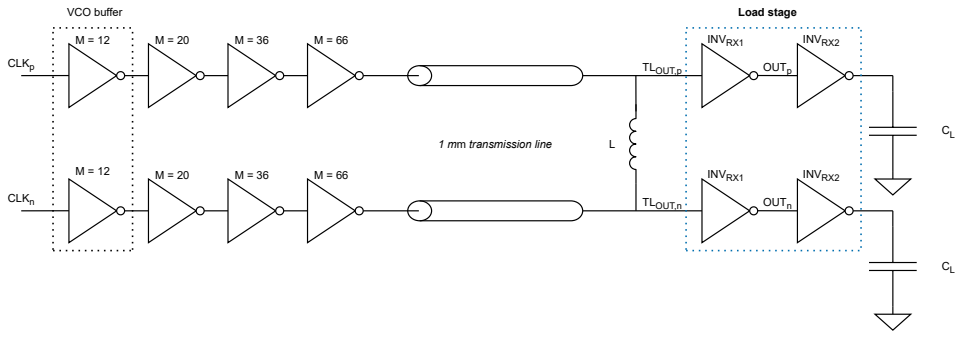


Figure 4.40: Far-end resonant implementation for the 1 mm transmission line, with the shunt inductance connected at $TL_{\mathrm{OUT},p}/TL_{\mathrm{OUT},n}$.

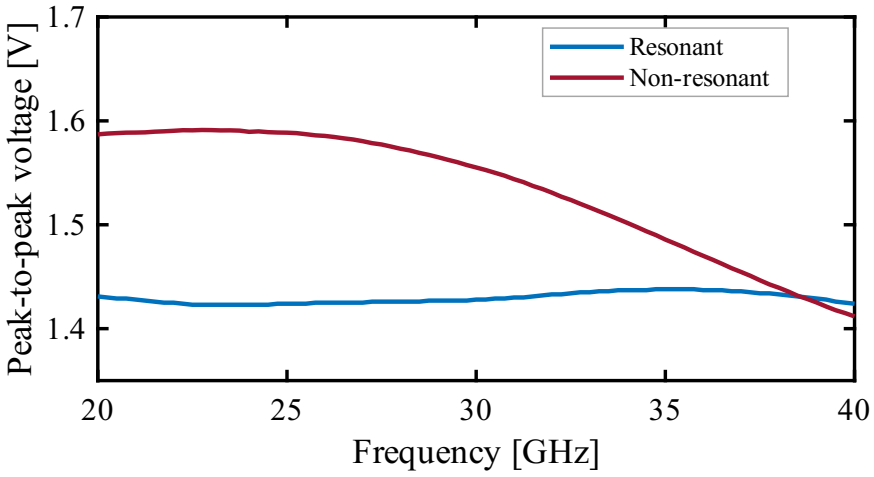


Figure 4.41: Differential peak-to-peak voltage evaluated between OUT_p and OUT_n as a function of the operating frequency for the far-end resonant and non-resonant implementations.

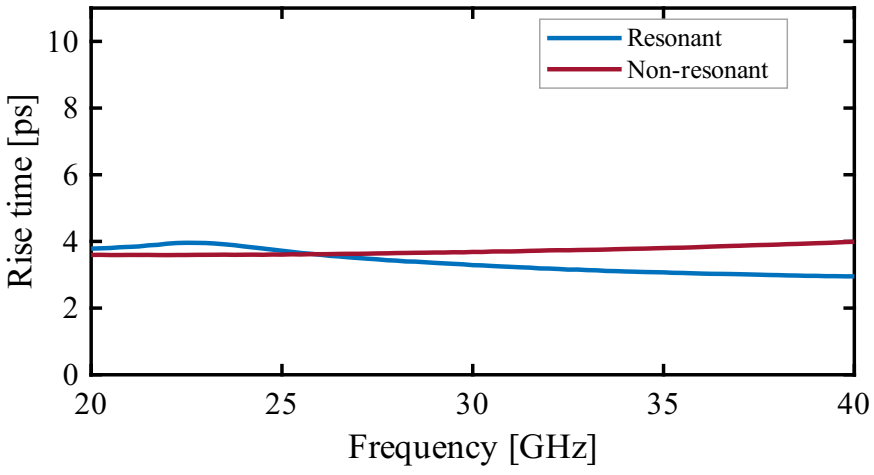


Figure 4.42: Rise time evaluated between OUT_p and OUT_n as a function of the operating frequency for the far-end resonant and non-resonant implementations.

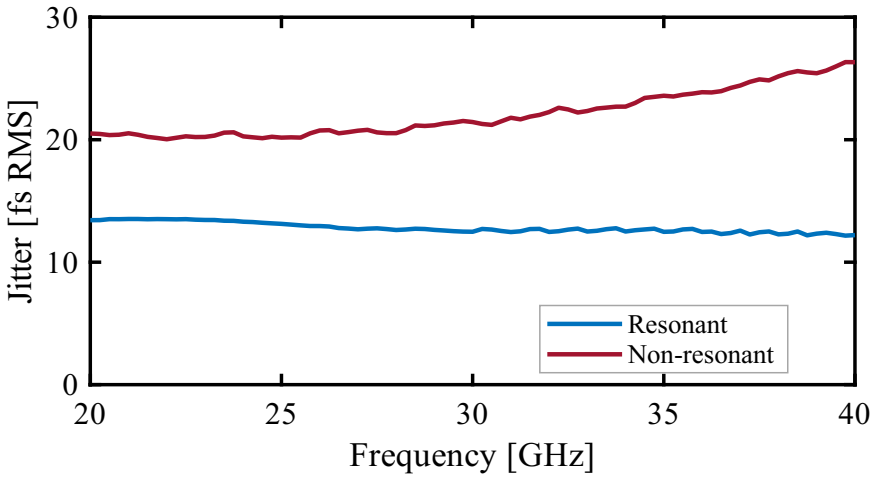


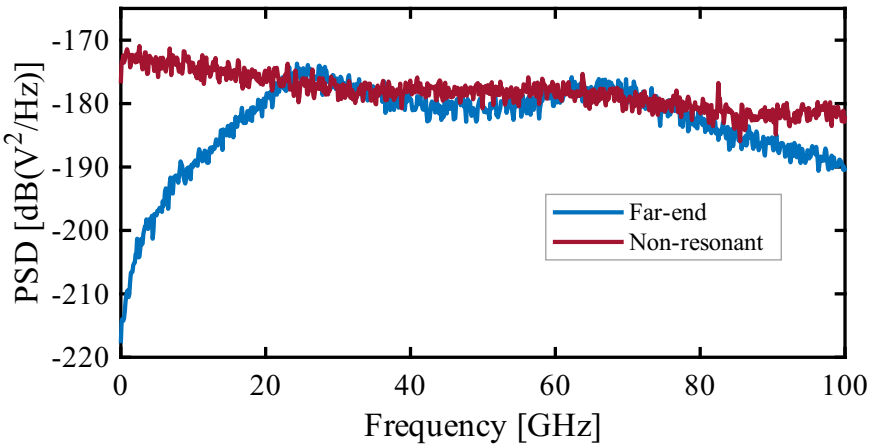
Figure 4.43: Output random jitter evaluated between OUT_p and OUT_n as a function of the operating frequency for the far-end resonant and non-resonant implementations.

The output random jitter is shown in Figure 4.43. The far-end resonant implementation exhibits lower random jitter than the corresponding non-resonant implementation throughout the target frequency range, while remaining relatively stable with frequency.

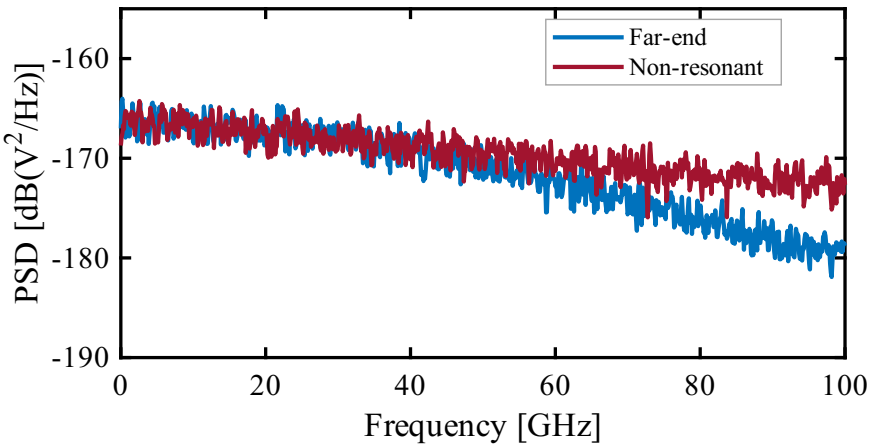
The noise behavior is further characterized through the voltage-noise power spectral density at $TL_{\mathrm{OUT},p}/TL_{\mathrm{OUT},n}$ and OUT_p/OUT_n, shown in Figure 4.44. The Far-End configuration produces a pronounced reduction in voltage noise at the transmission-line output. Unlike the other placements, the inductor is connected after the entire transmission line, directly at the input of INV_{RX1}. The inductive branch can therefore attenuate the differential noise reaching this node from the full transmission-line path. Its low impedance at low frequencies makes this effect particularly evident in that portion of the noise spectrum.

The jitter-amplification response is shown in Figure 4.45. Consistently with the other resonant strategies, the far-end implementation suppresses the amplification region observed in the non-resonant reference and provides attenuation over the corresponding perturbation-frequency range.

The PSIJ response is shown in Figure 4.46. The far-end resonant implementation exhibits lower PSIJ over the low- and intermediate-frequency range, while the difference with respect to the non-resonant reference becomes small at high perturbation frequencies.



(a)



(b)

Figure 4.44: Voltage-noise power spectral density for the far-end resonant and non-resonant implementations, evaluated (a) at the differential transmission-line output $TL_{\mathrm{OUT},p}/TL_{\mathrm{OUT},n}$ and (b) at the differential regenerated output OUT_p/OUT_n. For readability, the voltage-noise power spectral density is reported in dB, while the corresponding linear quantity is expressed in V^2/Hz.

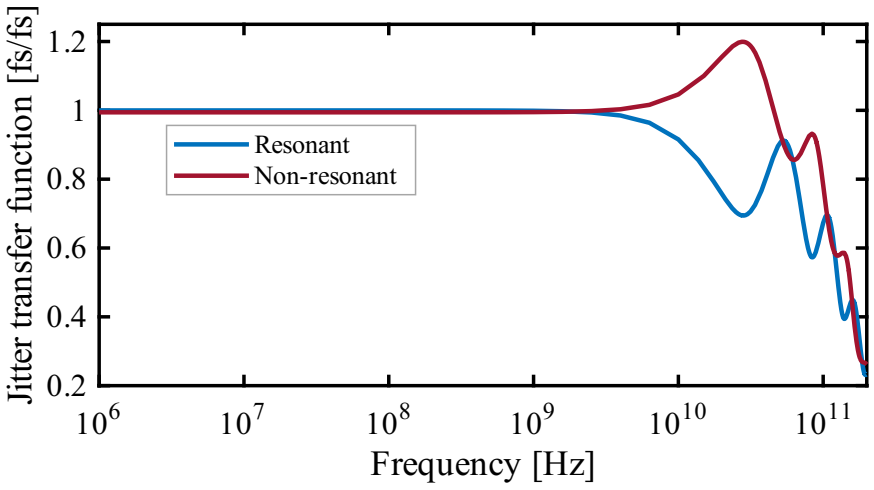


Figure 4.45: Jitter-amplification magnitude as a function of the signal frequency for the far-end resonant and non-resonant implementations.

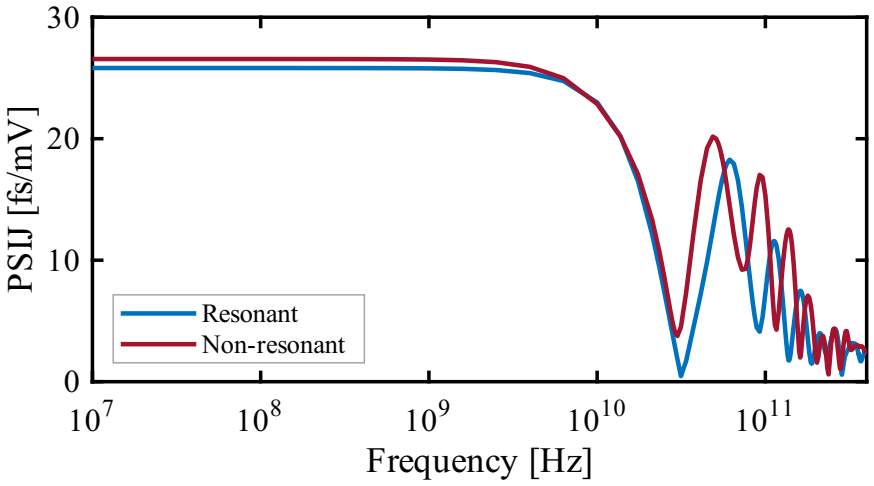


Figure 4.46: PSIJ as a function of the signal frequency for the far-end resonant and non-resonant implementations.

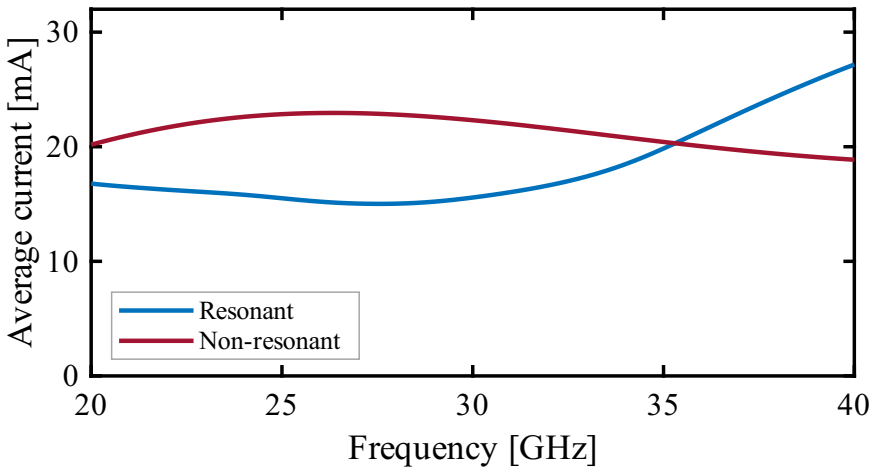


Figure 4.47: Average current consumption as a function of the operating frequency for the far-end resonant and non-resonant implementations.

Finally, the average current consumption is compared as a function of the operating frequency in Figure 4.47. Within the 25–30 GHz target range, the far-end resonant implementation requires less average current than the corresponding non-resonant reference. The difference decreases toward 30 GHz, while above the target range the resonant current eventually exceeds the non-resonant value.

4.3.4 Comparison of the three resonant approaches

Following the individual analysis of the three 1 mm implementations, their performance is compared to assess the differences among the three resonant approaches.

The clearest distinction concerns the jitter transfer function. Although all three configurations suppress the amplification peak of the non-resonant implementation, Near-End provides substantially stronger attenuation than Midpoint and Far-End in the correspond-

ing frequency region. Midpoint and Far-End, on the other hand, exhibit similar average current consumption, both lower than Near-End throughout the 25–30 GHz operating range. Around the design frequency, the Midpoint implementation exhibits a shorter output rise time and lower random jitter than the Far-End implementation. Near-End achieves the lowest random jitter of the three implementations around the design frequency, but its current advantage over the non-resonant reference decreases toward the upper end of the operating range. These results distinguish Near-End and Midpoint as the two approaches offering the most relevant benefits for the 1 mm implementation. Near-End provides the strongest suppression of jitter amplification, whereas Midpoint combines reduced current consumption with improved output rise time and random jitter. Far-End, despite drawing a current comparable to Midpoint, exhibits less favorable output timing and does not achieve the same jitter-transfer attenuation as Near-End.

4.4 Summary

This chapter investigated resonant clock distribution as an alternative to the non-resonant architecture, with the main objective of improving jitter-transfer behavior without relying on increased driver strength and the associated current penalty. Three possible positions of the resonant element have been discussed: Near-End, Midpoint, and Far-End.

The preliminary analysis at the nominal operating frequency showed that resonant operation can reduce the current required by the driver and provide jitter attenuation in the frequency region where the non-resonant implementation exhibits jitter amplification. The influence of the inductor quality factor was found to be limited for moderate and high values of Q_L.

The final 1 mm implementations confirm the suppression of jitter amplification and show that the benefits of resonant operation depend on both the selected approach and the operating frequency. These benefits are concentrated around the intended resonant condition, highlighting the narrowband nature of resonant operation in contrast to the broader operating behavior of the non-resonant architecture.

A further practical consideration follows directly from the resonance condition. The required inductance depends on the operating frequency and on the equivalent impedance at the resonant node. For the frequency range and interconnect lengths considered in this work, the resulting inductance values remain practical. However, this is not guaranteed in general, since different operating conditions may lead to inductance values that are difficult to implement.

Chapter 5

Comparison and conclusions

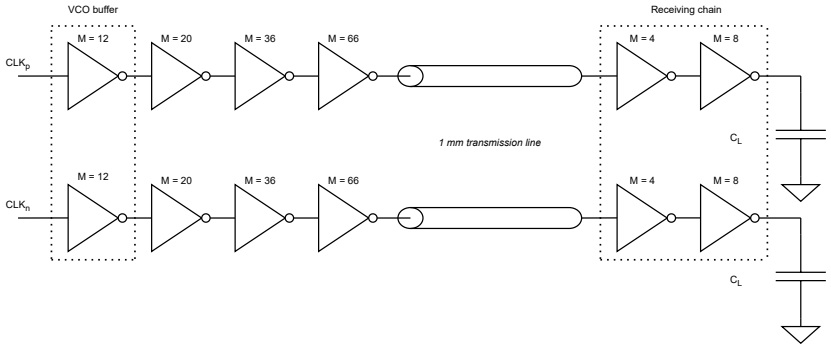
5.1 Comparison between non-resonant and resonant approaches

This chapter compares the final 1 mm resonant implementations with the non-resonant clock-distribution architectures investigated in Chapter 3. The single-segment non-resonant implementation is included as the direct reference, since the resonant circuits were obtained from it while keeping the active stages and load conditions unchanged. The two-segment non-resonant implementation is also considered to evaluate segmentation as an alternative to resonance. Among the three resonant approaches analyzed in Chapter 4, the comparison focuses on Near-End and Midpoint, which emerged as the two configurations of greatest interest for jitter-transfer attenuation and the combination of current consumption and output timing performance, respectively.

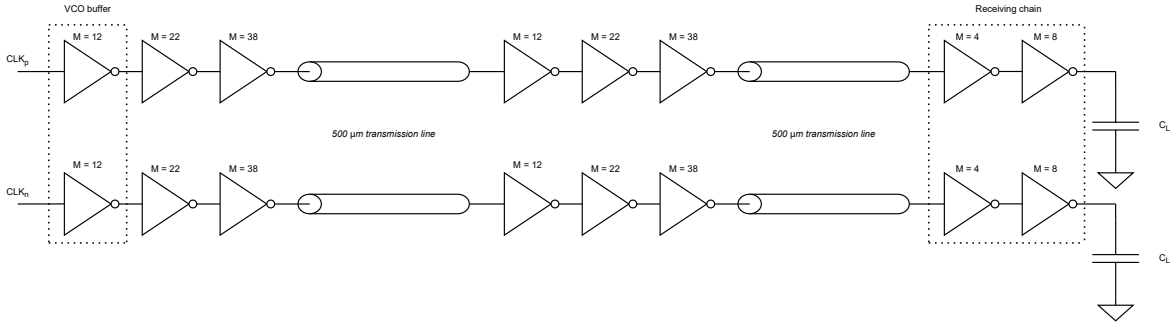
The four clock-distribution implementations considered in the following comparison are shown in Figure 5.1.

The comparison first considers the average current consumed by the four implementations, as shown in Figure 5.2. In the non-resonant architecture, segmentation is accompanied by a substantial increase in current consumption. In contrast, both resonant configurations require significantly less current than the two-segment non-resonant solution over the target frequency range. Both resonant configurations also require less current than the single-segment non-resonant approach over most of the target frequency range. This reduction is consistent with the resonant exchange of reactive energy between the inductive and capacitive elements, which reduces the amount of energy that must be supplied by the active driver during each transition.

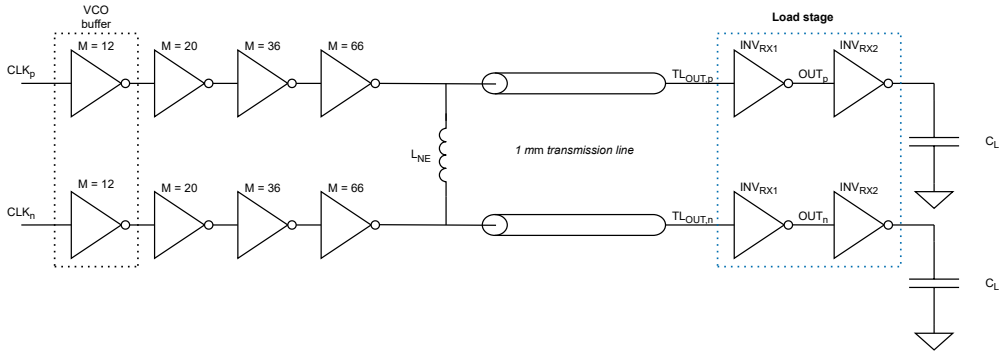
No significant variation in the differential output voltage swing is observed among the considered implementations, since the receiving stages regenerate the clock signal close to its full swing. More noticeable differences emerge in the output rise time, as shown in Figure 5.3. Segmentation reduces the rise time of the non-resonant architecture compared



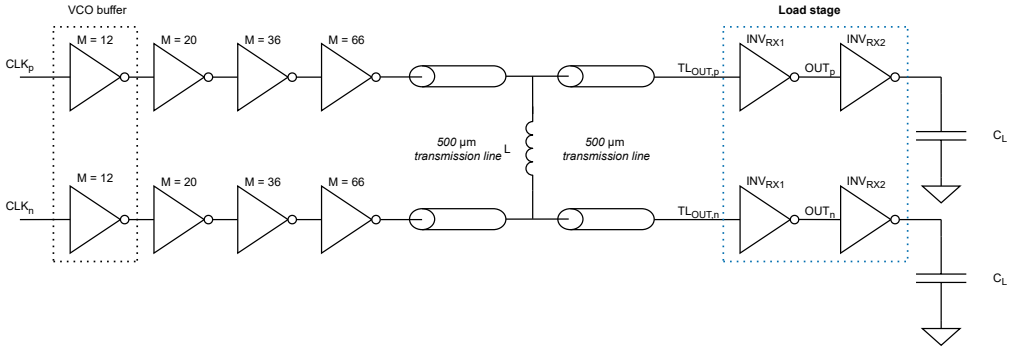
(a)



(b)



(c)



(d)

Figure 5.1: Clock-distribution implementations considered in the final comparison: (a) single-segment non-resonant, (b) two-segment non-resonant, (c) Near-End resonant, and (d) Midpoint resonant.

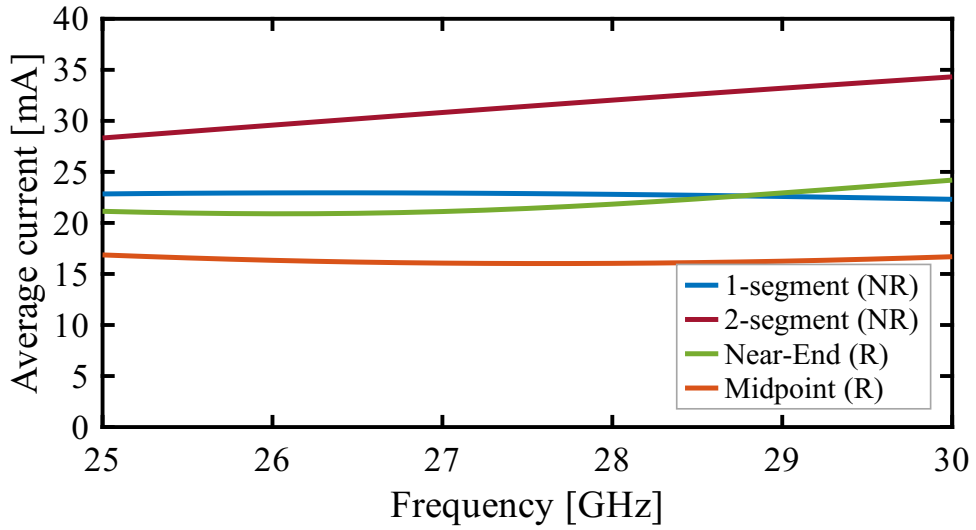


Figure 5.2: Average-current comparison between the non-resonant and resonant implementations.

with the single-segment implementation. Introducing resonance into the single-segment architecture also improves the output transitions, with both Near-End and Midpoint generally providing shorter rise times than the two-segment non-resonant solution over the target frequency range.

The output random jitter, evaluated between OUT_p and OUT_n, is compared in Figure 5.4. Unlike the rise time, random jitter shows only a limited improvement with segmentation, since the additional regeneration stage also introduces noise, as discussed in Chapter 3. Both resonant approaches provide lower random jitter than the two non-resonant implementations throughout the target operating range. Over most of the operating range, Near-End exhibits lower random jitter than Midpoint; this trend is consistent with the rise-time comparison in Figure 5.3, where Near-End also exhibits a shorter rise time than Midpoint.

The jitter-transfer comparison highlights a further advantage of resonant solutions, as shown in Figure 5.5. The single-segment non-resonant implementation exhibits jitter amplification in the frequency range of interest, whereas segmentation reduces the jitter-transfer magnitude and brings its maximum close to unity. This improvement in jitter-transfer behavior is accompanied by a substantial increase in current consumption, as shown in Figure 5.2. As shown in Chapter 3, local regeneration through a back-to-back inverter pair placed at the transmission-line output can further reduce jitter amplification, bringing the maximum jitter-transfer magnitude approximately to unity. Conversely, both resonant implementations provide a reduction of the jitter-transfer magnitude below unity in the frequency range of interest. Among them, the Near-End configuration provides a more pronounced jitter attenuation than the Midpoint configuration. Resonant operation therefore provides improved jitter-transfer performance without the current penalty associated with segmentation, highlighting a key advantage of the resonant approach.

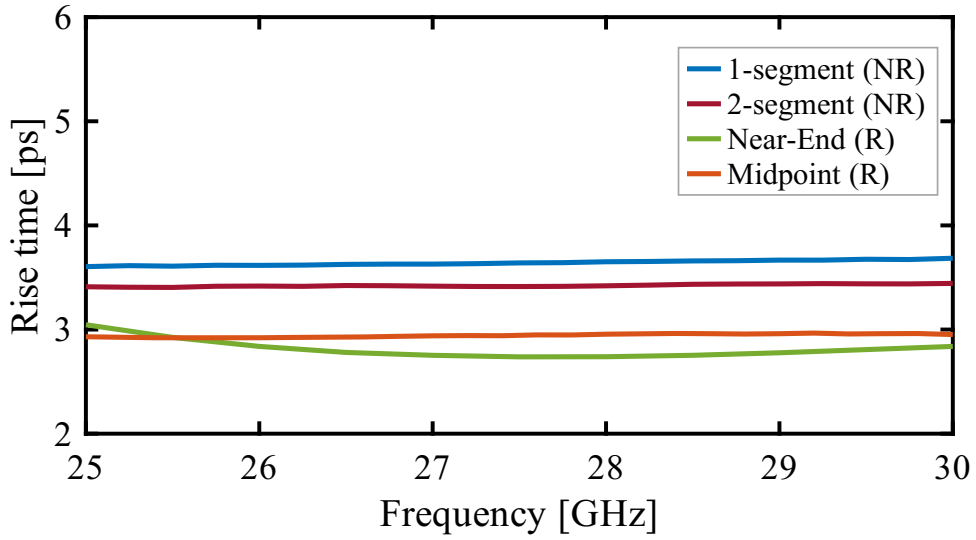


Figure 5.3: Output rise-time comparison between the non-resonant (NR) and resonant (R) implementations, evaluated between OUT_p and OUT_n.

As observed in Chapter 3, the single-segment non-resonant implementation provides better PSIJ performance than the two-segment solution. The resonant configurations further modify this behavior, as shown in Figure 5.6. At low perturbation frequencies, where the PSIJ responses are approximately constant, the Near-End configuration provides the largest reduction with respect to the single-segment non-resonant reference. The Midpoint configuration also improves PSIJ with respect to the two-segment non-resonant solution, while its response is similar to the single-segment non-resonant case.

The results show that the main difference between the two approaches lies in the trade-off between jitter-transfer performance, current consumption, and frequency selectivity. In the non-resonant architecture, segmentation improves the jitter-transfer behavior and brings the corresponding magnitude close to unity, but this improvement is obtained at the cost of a significant increase in average current. This current penalty represents the main drawback of segmentation. Resonant operation changes this trade-off substantially. Where the non-resonant implementation exhibits jitter amplification, the resonant configurations can instead provide jitter attenuation while requiring less current than the two-segment non-resonant solution. At the same time, the output voltage swing remains substantially comparable, while transition time, random jitter, and sensitivity to supply disturbances are generally improved. Resonance therefore does not simply reproduce the improvement obtained through segmentation with a different implementation, but provides a favorable combination of timing, jitter-transfer, and current-consumption performance within the target operating range.

The resonant approach nevertheless introduces additional design constraints. Its benefits rely on a tuned resonant condition and are therefore inherently frequency selective, whereas the non-resonant architecture does not require frequency tuning and can operate

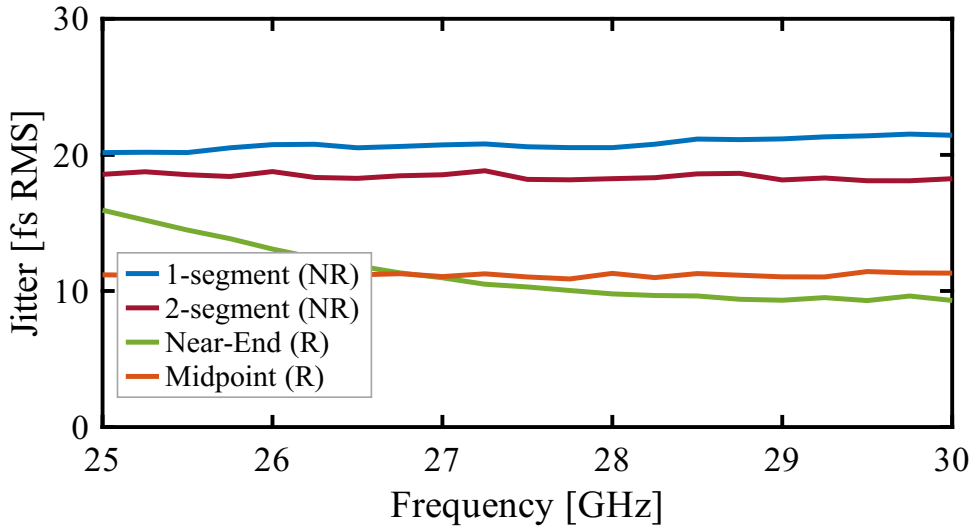


Figure 5.4: Output random-jitter comparison evaluated between OUT_p and OUT_n for the non-resonant (NR) and resonant (R) implementations.

over a broad frequency range. The design of the resonant network also requires the inductance value and its placement to be selected according to the impedance provided at the resonant node. In addition, the losses associated with the resonant network must be taken into account. The analysis of the inductor quality factor showed that low-Q_L values reduce the benefits of resonance. As shown in Chapter 3, increasing the transmission-line length requires a larger non-resonant driver to maintain the required jitter-transfer performance, with a corresponding increase in current consumption. As long as the electrical load presented by the transmission line can be driven by the non-resonant CMOS distribution while maintaining a jitter-transfer magnitude close to unity without a large current penalty, the introduction of an inductor may not be justified. As the line becomes more demanding to drive, the required drive strength and current consumption increase, making the resonant approach more attractive. The required inductance also depends on the operating frequency, line length, and resonant-node impedance. Although the values obtained in this work are practical, different operating conditions may require inductors with excessive silicon area or parasitic capacitance, potentially limiting their usable frequency range.

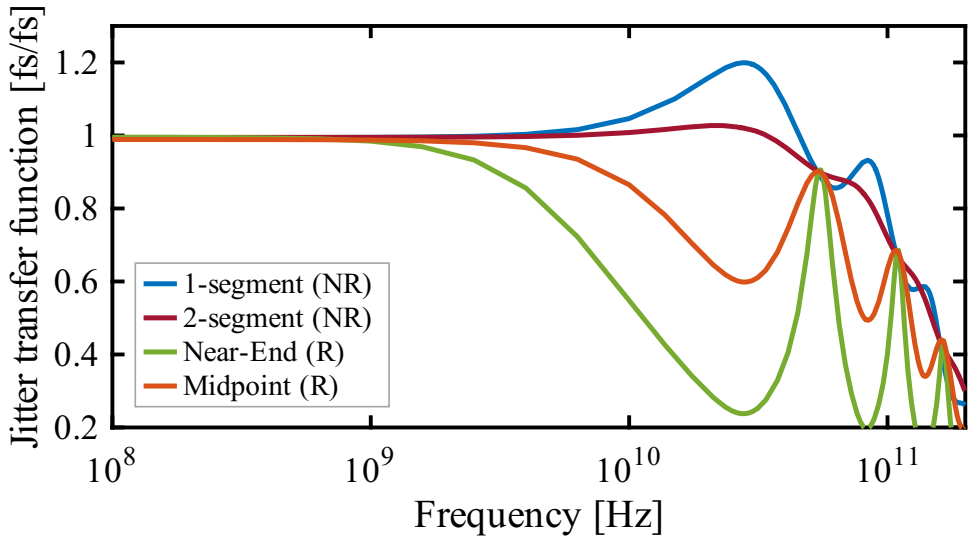


Figure 5.5: Jitter transfer function comparison between the non-resonant (NR) and resonant (R) implementations.

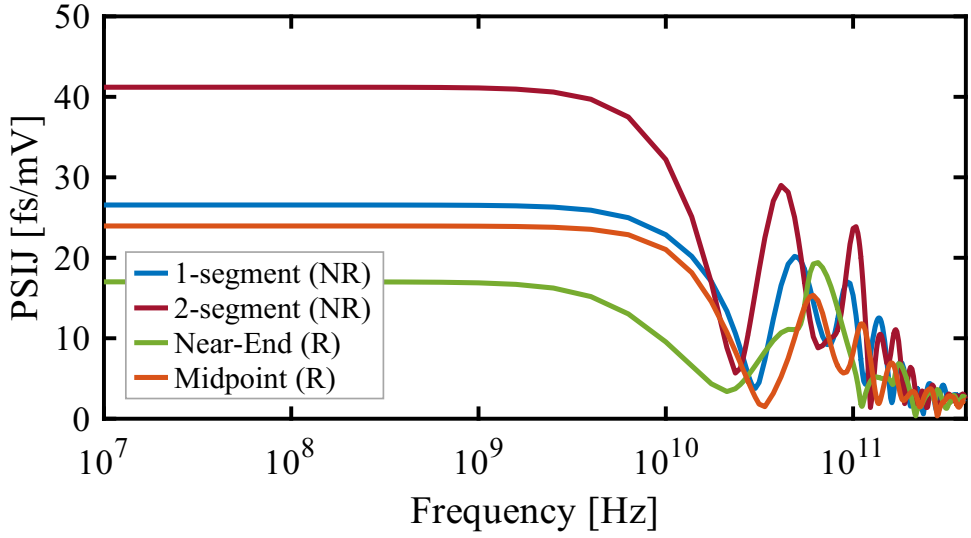


Figure 5.6: Power-supply-induced jitter comparison between the non-resonant and resonant implementations.

5.2 Future developments

Several extensions of the present work can be considered. A first direction concerns the design of the transmission line itself. In the analysis presented in this thesis, the clock-distribution circuits were designed and optimized starting from a given interconnect implementation. However, the electrical characteristics of the transmission line depend strongly on its physical realization, including conductor geometry, spacing, metal layer, and coupling. A more complete design methodology could therefore include the transmission line among the design variables rather than treating it as a fixed element. This would allow the interconnect geometry to be optimized together with the clock-distribution circuitry, with the objective of identifying the line implementation that best supports the required performance for a given distribution length, operating frequency, and load.

A second direction concerns the optimization of the resonant implementation itself. In this work, the resonant architectures were derived from the corresponding non-resonant solution while keeping the active circuitry substantially unchanged, so that the effect introduced by resonance could be evaluated independently. The non-resonant driver, however, was sized also to limit jitter amplification, which represents one of the main constraints of the conventional architecture. Once resonance is introduced, the resonant network itself improves the jitter-transfer behavior and can provide attenuation below unity. Consequently, the driver inherited from the non-resonant design may be larger than necessary for the resonant implementation. A dedicated optimization could therefore reduce the driver size and current consumption while accepting a controlled degradation of the timing performance, provided that the required jitter-transfer behavior is still maintained. In this case, the driver size and the inductive-element value should be optimized jointly, since changing the active-stage sizing also modifies the impedance seen at the resonant node.

Another possible development concerns the limited operating bandwidth associated with the resonant architecture. In the implementations investigated in this work, the resonant condition is established around a fixed operating frequency. The introduction of tunable elements, such as switched or variable capacitances, could allow the resonant frequency to be adjusted and therefore extend the usable frequency range of the architecture. In addition, alternative resonant configurations could be investigated with the objective of obtaining a broad operating region. Hybrid architecture could also be considered, for example by combining resonance with segmentation or with the back-to-back regeneration stages investigated for the non-resonant implementation.

An extension of this work would be the investigation of current-mode logic (CML) clock distribution. The architectures considered in this thesis are based on CMOS inverters and therefore distribute full-swing clock signals, whereas CML relies on differential current steering and reduced voltage swing. Indeed, CML clock buffers can provide improved

rejection of supply disturbances, which can be particularly attractive when power-supply-induced jitter is a critical metric. At the same time, CML drivers introduce static power consumption, they require bias circuitry, and a CML-to-CMOS interface is needed when the distributed clock must be delivered to CMOS circuitry.

A relevant continuation of this work would therefore be to design a CML-based distribution network under the same operating conditions as considered in this thesis and evaluate it using the same set of metrics. Such a comparison could include power consumption, random jitter, jitter transfer, and sensitivity to supply disturbances, while also accounting for the CML-to-CMOS interface required at the receiving side. This would make it possible to directly assess the trade-offs between CMOS and CML clock distribution and to identify the operating conditions and design requirements for which one approach becomes more advantageous than the other.

5.3 Conclusions

The objective of this work was to investigate and compare non-resonant and resonant CMOS clock-distribution architectures for high-frequency operation, focusing on their operating principles and on the main trade-offs introduced by the two approaches. The analysis was carried out under the same reference conditions to understand how the two architectures behave, which mechanisms determine their performance, and under which conditions one solution may lead to better performance than the other.

The analysis of the non-resonant architecture first provided a methodology for sizing the CMOS driving network according to the electrical loading introduced by the transmission line and to the settling requirements imposed by high-frequency operation. The investigation of the complete 1 mm distribution showed that jitter amplification represents the most critical metric for the implementation considered. In this respect, segmentation can improve the jitter-transfer behavior by reducing the loading of the individual driving stages, although the additional regeneration stages also introduce additional noise, increase current consumption, and degrade supply sensitivity. Among the investigated segmented solutions, the two-segment implementation provided the most favorable non-resonant compromise. The analysis of additional back-to-back inverter pairs also showed that local regeneration can further reduce jitter amplification when appropriately designed, with the most effective improvement obtained when the pair is located at the transmission-line output. Overall, the non-resonant analysis highlighted a fundamental trade-off between power consumption and jitter performance: reducing jitter amplification and improving the timing behavior generally requires additional active circuitry, thus increasing the current consumption.

The resonant approach was investigated starting from this limitation. By introducing an inductive element appropriately sized according to the impedance presented at the

selected resonant node, the reactive loading of the distribution network can be partially compensated. The results obtained for the Near-End, Midpoint, and Far-End configurations show that resonance improves the trade-off between power consumption and jitter performance identified in the non-resonant case. The resonant implementations provide jitter attenuation rather than simply limiting jitter amplification, while requiring less current than the segmented non-resonant solution and maintaining comparable clock-signal quality. At the same time, these advantages are obtained at the cost of a frequency-selective behavior and of additional constraints associated with the value, losses, and placement of the inductive element. Overall, none of the investigated resonant configurations provides the best performance across all the considered metrics.

The comparison therefore shows that the choice between non-resonant and resonant clock distribution depends on the requirements of the specific application. The non-resonant approach provides a broadband solution and does not require a tuned inductive element, but its jitter performance depends on the amount of active circuitry and, consequently, on power consumption. As long as the transmission-line load can be driven while maintaining a jitter-transfer magnitude close to unity without a large current penalty, resonance may not be necessary if additional jitter attenuation is not required. For clock-distribution paths in which the non-resonant implementation does not provide sufficient jitter-transfer performance, the resonant approach can instead provide a more favorable trade-off between power consumption and jitter performance within the frequency range for which the network is designed, at the expense of increased frequency selectivity and additional resonant-design constraints.

Thus, the appropriate choice of clock-distribution strategy depends on the distribution length and on the required balance between jitter performance, power consumption, and operating bandwidth.

Bibliography

[1] T. Chan Carusone, T. O. Dickson, S. Palermo, S. Shekhar, and M. Mansuri, "Modern wireline transceivers," *IEEE Journal of Solid-State Circuits*, vol. 61, no. 2, pp. 395–422, Feb. 2026.

[2] F. Giunco, M. Sosio, C. Nani, I. Fabiano, T. Lovitt, V. Karam, D. Albano, C. Asero, N. Codega, M. Garampazzi, N. Ghittori, L. D. Herbas Burgos, S. S. K. Ho, E. Monaco, B. Reyes, P. Rossi, E. Temporiti, P. Pascale, F. De Bernardinis, S. Scouten, and S. Jantzi, "An eight-lane 800-gb/s transceiver for pam-4 optical direct-detection applications in 5-nm finfet process," *IEEE Journal of Solid-State Circuits*, vol. 60, no. 4, pp. 1277–1288, Apr. 2025.

[3] Statista, "Volume of data/information created, captured, copied, and consumed worldwide from 2010 to 2020, with forecasts from 2021 to 2025," Jun. 2021, based on International Data Corporation data. [Online]. Available: https://www.statista.com/statistics/871513/worldwide-data-created/

[4] B. Razavi, "Design techniques for high-speed wireline transmitters," *IEEE Open Journal of the Solid-State Circuits Society*, vol. 1, pp. 53–66, Sep. 2021.

[5] X. Mo, J. Wu, N. Wary, and T. Chan Carusone, "Design methodologies for low-jitter cmos clock distribution," *IEEE Open Journal of the Solid-State Circuits Society*, vol. 1, pp. 94–103, Oct. 2021.

[6] D. Abts and B. Felderman, "A guided tour of data-center networking," *Communications of the ACM*, vol. 55, no. 6, pp. 44–51, Jun. 2012.

[7] S. Palermo, "High-speed serial i/o design for channel-limited and power-constrained systems," in *CMOS Nanoelectronics: Analog and RF VLSI Circuits*. McGraw-Hill, 2011, ch. 9.

[8] A. Boletti, D. Giacomuzzi, G. Parladori, P. Boffi, M. Ferrario, and M. Martinelli, "Performance comparison between electrical copper-based and optical fiber-based backplanes," *Optics Express*, vol. 21, no. 16, pp. 19 202–19 208, Aug. 2013.

[9] D. A. B. Miller, "Rationale and challenges for optical interconnects to electronic chips," *Proceedings of the IEEE*, vol. 88, no. 6, pp. 728–749, Jun. 2000.

[10] G. T. Reed, G. Mashanovich, F. Y. Gardes, and D. J. Thomson, "Silicon optical modulators," *Nature Photonics*, vol. 4, no. 8, pp. 518–526, Aug. 2010.

[11] H. Tong, Y. Wang, S. Tian, C. Tong, Z. Wei, and L. Wang, "Parasitic parameters extraction of high-speed vertical-cavity surface-emitting lasers," *Applied Sciences*, vol. 12, no. 12, p. 6035, Jun. 2022.

[12] Q. Liao, Y. Zhang, S. Ma, L. Wang, L. Li, G. Li, Z. Zhang, J. Liu, N. Wu, L. Liu, Y. Chen, X. Xiao, and N. Qi, "A 50-gb/s pam-4 silicon-photonic transmitter incorporating lumped-segment mzm, distributed cmos driver, and integrated cdr," *IEEE Journal of Solid-State Circuits*, vol. 57, no. 3, pp. 767–778, Mar. 2022.

[13] H. W. Johnson and M. Graham, *High-Speed Signal Propagation: Advanced Black Magic*. Upper Saddle River, NJ, USA: Prentice Hall, 2003.

[14] B. Razavi, *RF Microelectronics*, 2nd ed. Upper Saddle River, NJ, USA: Prentice Hall, 2012.

[15] ——, *Design of Analog CMOS Integrated Circuits*, 2nd ed. New York, NY, USA: McGraw-Hill Education, 2017.

[16] R. J. Baker, *CMOS: Circuit Design, Layout, and Simulation*, 4th ed. Hoboken, NJ, USA: Wiley-IEEE Press, 2019.

[17] *IEEE Standard for Transitions, Pulses, and Related Waveforms*, IEEE, New York, NY, USA, Sep. 2011.

[18] B. Casper and F. O'Mahony, "Clocking analysis, implementation and measurement techniques for high-speed data links—a tutorial," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 56, no. 1, pp. 17–39, Jan. 2009.

[19] Tektronix, "Understanding and characterizing timing jitter," Application note, document 55W-16146-6. [Online]. Available: https://download.tek.com/document/Understanding-and-Characterizing-Timing-Jitter_55W_16146_6.pdf

[20] W. T. Beyene, "Modeling and analysis techniques of jitter enhancement across high-speed interconnect systems," in *2007 IEEE Electrical Performance of Electronic Packaging*, Oct. 2007, pp. 29–32.

[21] F. Rao and S. Hindi, "Mechanism of jitter amplification in clock channels," in *DesignCon 2014*. Agilent Technologies, Jan. 2014.

[22] D. M. Pozar, *Microwave Engineering*, 4th ed. Hoboken, NJ, USA: John Wiley & Sons, 2012.