

Distributed Differential Oscillators for Global Clock Networks

Steven C. Chan, *Member, IEEE*, Kenneth L. Shepard, *Senior Member, IEEE*, and Phillip J. Restle, *Member, IEEE*

Abstract—This paper presents a distributed differential oscillator global clock network where the clock capacitance is rendered resonant with a set of on-chip spiral inductors. The clock amplitude and clock phase are both uniform across the entire global distribution, making this design scalable and compatible with existing local clocking methodologies. The resonant network, combined with phase averaging of the distributed oscillator, provides high immunity to process-, voltage-, and temperature-variation-induced timing uncertainty. Measurement results from a prototype design implemented in a 0.18- μm CMOS technology show almost an order of magnitude less jitter and power than a traditional tree-driven grid global clock distribution. On-chip measurement circuits are used to characterize the jitter on the test chip, while a simulation model is used to examine skew and higher-order resonances in the resonant clock network.

Index Terms—Clock distribution, inductance, jitter, resonant clocking, skew, timing circuits.

I. INTRODUCTION

THE global clock on a microprocessor has traditionally been distributed using a hierarchical approach, in which a tuned and balanced tree drives a grid [1]–[5]. Ensuring that this large tree-driven-grid global clock network is low-skew and low-jitter in the presence of process, voltage, and temperature (PVT) variation is a significant challenge. The increasing clock frequencies that accompany technology scaling exacerbate the challenge, making alternative global clocking solutions that offer better scalability than traditional tree-driven grids highly desirable.

Resonant clock distributions have been proposed as an alternative clock technology and have shown promise in reducing clock timing uncertainty and power dissipation. Standing-wave clock distributions have been implemented at both the board level [6] and chip level [7]. These designs achieve low-skew and low-jitter clocks and can save power due to the resonance between the clock capacitance and the clock wire inductance. Standing-wave clock distributions must, however, contend with nonuniform clock amplitude which may result in skew or make

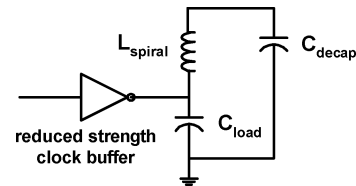


Fig. 1. Simple lumped model of the resonant-load global clock distribution described in [17].

local clock buffering more complex. Traveling-wave clock distributions [8] use coupled transmission line rings to generate a low-skew and low-jitter clock and also benefit from the power advantage of resonance. Non-uniform phase across the distribution makes integration with existing local clocking methodologies more difficult, however. Other resonant clock generation approaches include “power” clocks for adiabatic logic [9]–[11] and an LC oscillator local clock [12]. Neither of these two approaches, however, attempt to distribute a resonant clock across a large area.

Oscillator array clocks, another proposed alternative to tree-driven grid clocks, distribute clock generation, thereby reducing the distance and latency between clock source and clock load. Some oscillator array clocks use phase detectors to synchronize the array [13], [14], while others directly couple the oscillators together using interconnect [7], [8], [15], [16]. Both resonant [7], [8] and nonresonant [13]–[16] oscillator array clock distributions have been examined. Many of these distributions however, are complicated by nonuniform phase, nonuniform amplitude, or complex synchronization schemes.

The authors have previously proposed a resonant-load global clock distribution which augments traditional tree-driven grids with a set of spiral inductors, which resonates with the clock load capacitance [17]. This results in a clock amplitude and clock phase that are both uniform across the entire global distribution, making this design scalable and compatible with existing local clocking methodologies. A sizable portion of the jitter reduction and power savings results from reducing the strength of the clock buffers driving the resonant-load, exploiting the proclivity of the network to sustain a tone at the target resonant frequency, resulting in a nearly sinusoidal clock signal. Sinusoidal clocks are however generally undesirable because of slower signal transition times, which exacerbate timing uncertainty at single-ended local clock buffers in the presence of PVT variation.

Another drawback associated with the resonant-load global clock distribution in [17] is the requirement for large on-chip decoupling capacitance to serve as a charge reservoir. In the simple lumped model of Fig. 1, the spiral inductance L_{spiral} resonates

Manuscript received October 9, 2005; revised May 1, 2006. The work at Columbia University was supported by the MARCO/DARPA C2S2 Center, the National Science Foundation under Contract CCR-00-86007, the Semiconductor Research Corporation, and gifts from IBM and Intel.

S. C. Chan was with Columbia Integrated Systems Laboratory, Department of Electrical Engineering, Columbia University, New York, NY 10027 USA. He is now with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA. (e-mail: stchan@us.ibm.com).

K. L. Shepard is with the Columbia Integrated Systems Laboratory, Department of Electrical Engineering, Columbia University, New York, NY 10027 USA (e-mail: shepard@ee.columbia.edu).

P. J. Restle is with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA.

Digital Object Identifier 10.1109/JSSC.2006.880610

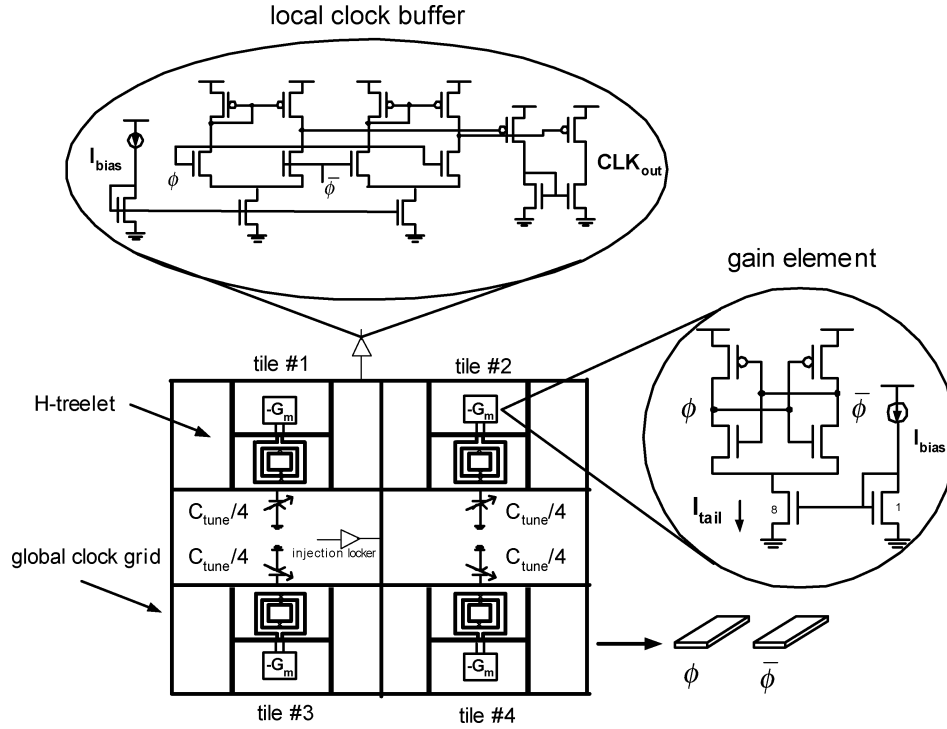


Fig. 2. Distributed differential oscillator (DDO) global clock network.

with the clock load capacitance C_{load} , driven by the clock buffer to a full-rail (V_{DD}) amplitude. The voltage across C_{decap} is approximately $V_{DD}/2$ and establishes the dc level around which the clock load oscillates. C_{decap} must usually be at least ten times as large as C_{load} .

In this paper, an improved resonant clock design is presented which addresses the concerns of the resonant-load design. A distributed differential oscillator (DDO) global clock network is presented which combines the idea of resonating a clock grid with a set of distributed spiral inductors with the low-latency advantages of an oscillator array distribution. The distribution is a free-running oscillator injection-locked to an external reference. Furthermore, the distribution is differential with the use of symmetric inductors placed between the two clock phases (ϕ and $\bar{\phi}$), eliminating the need to add large capacitors to the clock distribution as in the resonant single-ended distribution of [17]. However, as in [17], at resonance, the DDO global clock is uniform in both amplitude and phase across the distribution. Differential detection at the local clock buffers reduces skew and jitter due to power-supply noise, process variations, and other common-mode noise sources [18]–[20], mitigating the disadvantages of the more sinusoidal clocks characteristic of resonant distributions.

The paper begins with an overview in Section II on the DDO global clock network architecture. The design and implementation of a prototype test chip is discussed and comments on the interconnect design, resonant network analysis, jitter, injection locking, and on-chip jitter measurement circuits are presented. Measurement results from the test chip showing the benefits of this new resonant clocking scheme is presented in Section III. Section IV concludes.

II. DDO ARCHITECTURE

In the prototype DDO global clock network designed here (see Fig. 2), four resonant oscillators (or “tiles”) are coupled together through interconnect to form a distributed differential oscillator global clock network. Each tile consists of a spiral inductor and negative differential transconductor (“gain element”) to compensate for loss and to maintain oscillation. The bias current in the gain elements, I_{tail} , determines the clock amplitude. “H-treelets” are used to route the clock from the inductors and gain elements to the global clock grid. The treelets distribute the drive of the inductors and gain elements, thereby improving skew. A small variable strength buffer positioned at the center of the clock network provides injection locking to a reference clock.

This DDO global clock network is fabricated in a $0.18\text{-}\mu\text{m}$ 1.8 V mixed-signal CMOS technology with six levels of aluminum wiring. The distribution, shown approximately to scale in Fig. 2, is $2\text{ mm} \times 2\text{ mm}$. Local clock buffers perform differential-to-single-ended conversion before jitter measurement. On-chip jitter histogram circuits allow for period jitter characterization of the clock network. Both ϕ and $\bar{\phi}$ have a capacitive load of 13 pF nominally (3.25 pF per tile); 9 pF (2.25 pF per tile) is from the clock wires themselves, while the remainder is from the local clock buffering and the on-chip test and measurement circuitry. Up to 5 pF (1.25 pF per tile) can be added to each clock phase through switchable MOS capacitors, C_{tune} , to provide tuning of the oscillation frequency from 1.6 to 1.1 GHz . Fig. 3 shows a die micrograph of the test chip. A differential nonresonant global clock network driven by a conventional clock H-tree and a clock grid was also designed on the same test chip for comparison purposes.

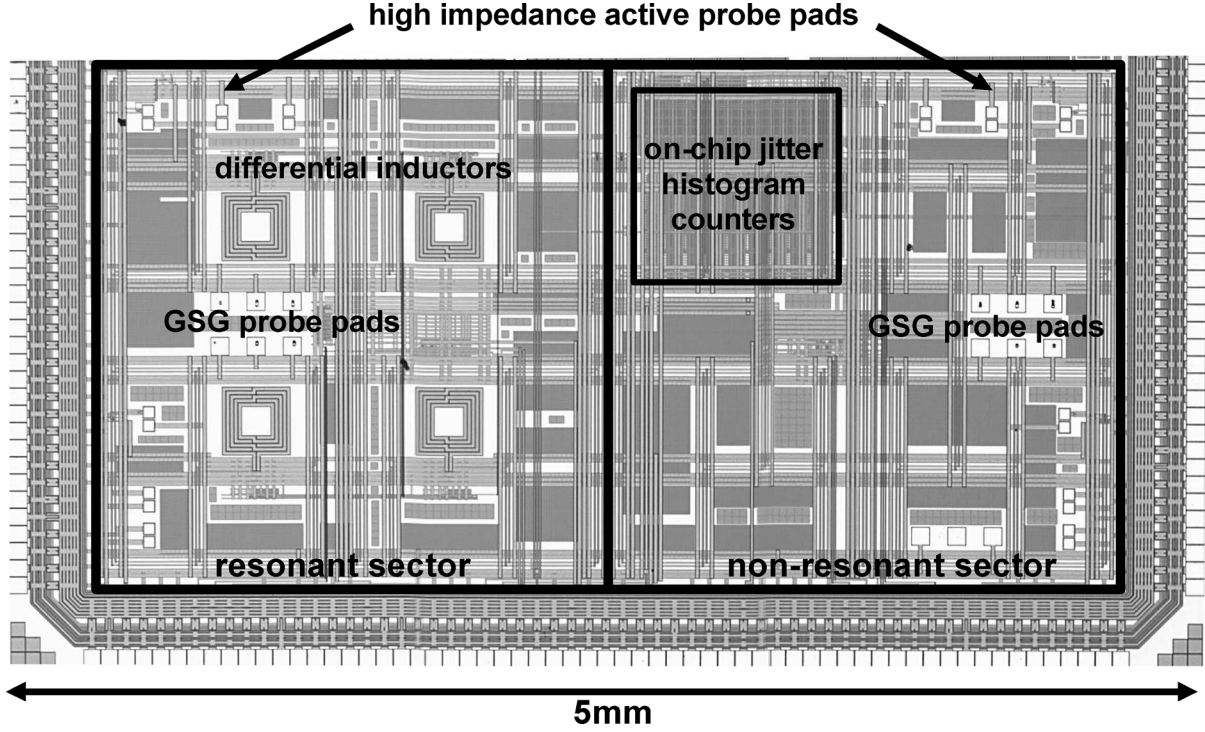


Fig. 3. Die photo of the test chip.

A. Interconnect Design and Resonant Network Analysis

No shielding of the clock grid wires, which are $6\ \mu\text{m}$ wide and spaced $8\ \mu\text{m}$ apart, is necessary to maintain a controlled clock grid inductance since the ϕ and $\bar{\phi}$ wires provide current return paths for each other. Two-dimensional cross-sectional simulations [21] show the clock grid wire inductance to be approximately $0.4\ \text{nH/mm}$, helping to create a low-inductance, low-resistance clock grid (i.e., one that is “rigid”) and ensuring a uniform-phase and uniform-amplitude clock at the target resonant frequency. Each of the differentially-driven symmetric on-chip spiral inductors [22] has a diameter of $280\ \mu\text{m}$, with the wires of the spiral inductor determined by the top two aluminum metal layers of the stack (M6 and M5) connected through vias, thereby reducing the series resistance.

Fig. 4(a) shows the results of three-dimensional field solver [23] simulations on the inductor from $100\ \text{MHz}$ to $4\ \text{GHz}$. The inductor Q , series inductance L_s , and series resistance R_s are extracted from simulated S-parameter data [24]. The equivalent half-circuit of the differentially-driven inductor is shown in Fig. 4(b). The inductor has a peak Q of 7.7 at $800\ \text{MHz}$ (slightly below the free-running DDO global clock frequency of between 1.1 and $1.6\ \text{GHz}$) and a self-resonant frequency of $2.5\ \text{GHz}$ due to the parasitic capacitance C_p . At $1.6\ \text{GHz}$, the extracted Q , L_s , and R_s are 4.3 , $6.2\ \text{nH}$ and $8.5\ \Omega$, respectively, while at $1.1\ \text{GHz}$ they are 6.8 , $6.4\ \text{nH}$, and $5.2\ \Omega$, respectively. C_p is approximately $325\ \text{fF}$.

It is useful to think of the entire differential clock network as a tank, the half-circuit of which can be similarly modeled as in Fig. 4(b). In this case, C_p is the capacitance of the (ϕ or $\bar{\phi}$) clock network. Expressions for the Q , resonance frequency, admittance, and clock amplitude that do not make the assumption

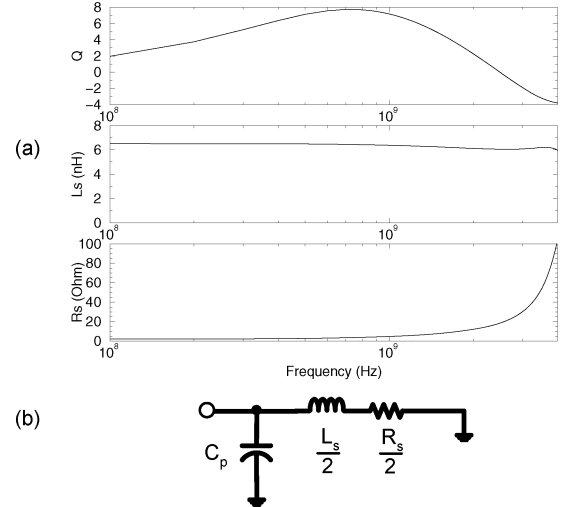


Fig. 4. (a) Simulated Q , series inductance L_s , and series resistance R_s of the differentially-driven on-chip spiral inductor used in the DDO global clock network. (b) The equivalent half-circuit of inductor using a simple three element frequency independent lumped RLC one-port, where C_p represents parasitic inductor capacitance. This half-circuit can also be used to represent the resonant clock network when C_p is the clock load capacitance.

of low loss follow from this three element RLC model. In particular, resonance occurs at

$$\omega_o = \frac{1}{\sqrt{(L_s/2)C_p}} \sqrt{1 - \frac{(R_s/2)^2 C_p}{(L_s/2)}} \quad (1)$$

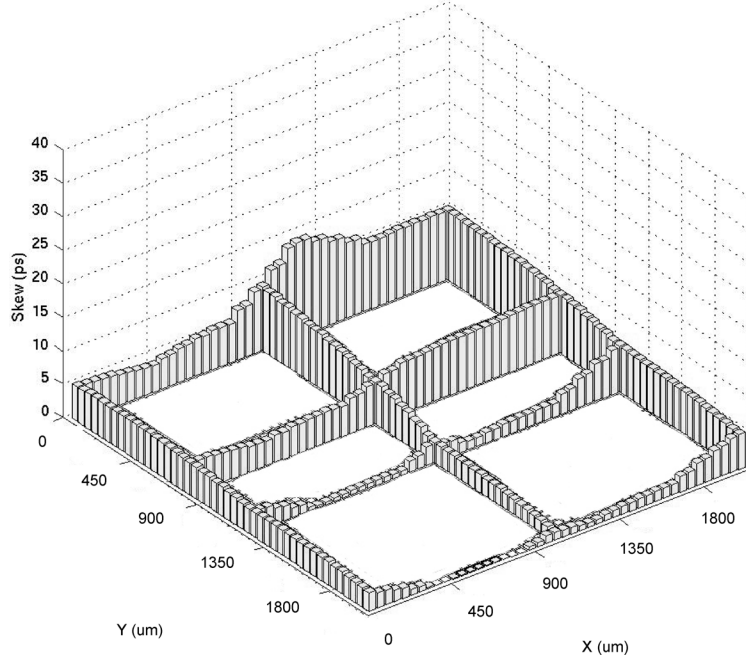


Fig. 5. Simulated skew on the resonant clock network when 3.25 pF of additional clock loading, approximately 18% of the total clock capacitance, is lumped at one point on the global clock grid. The worst case skew is 13.4 ps, or 1.5% of the clock period, with $C_{\text{tune}} = 5$ pF at 1.1 GHz.

with a Q of

$$Q = \frac{\omega_o L_s}{R_s} = \frac{\sqrt{(L_s/2)/C_p}}{(R_s/2)} \sqrt{1 - \frac{(R_s/2)^2 C_p}{(L_s/2)}}. \quad (2)$$

We use R_s to model all the losses in the network, including series resistance losses, eddy current losses in the substrate and neighboring conductors, and displacement current losses. At resonance, the magnitude of the admittance is given by

$$|Y(\omega_o)| = \frac{R_s C_p}{L_s}. \quad (3)$$

Assuming a square wave excitation of the tank by the gain elements, the clock amplitude is approximated by

$$V_{\text{clock}} \cong \frac{4}{\pi} \frac{I_{\text{tail}}}{|Y(\omega_o)|} \quad (4)$$

where I_{tail} is bias current in the gain elements (see Fig. 2).

As in traditional clock distributions, the presence of a grid helps to amortize capacitance nonuniformity—the denser the grid, the more effective this amortization. The same applies to the resonant network designed here. To demonstrate the skew caused by capacitance nonuniformity, 3.25 pF of additional clock loading, approximately 18% of the total clock capacitance, is lumped at one point on the global clock grid and simulations are used to examine its effect. The worst case skew (with the skew distribution shown in Fig. 5) is 13.4 ps, or 1.5% of the clock period, with C_{tune} adjusted to 5 pF for a 1.1 GHz resonant frequency. The nonresonant clock network,

which shares the same grid design, has a comparable amount of worst case skew (17.3 ps with 3.25 pF of additional clock loading lumped at a single point). Without the load imbalance, the nominal skew is less than 3 ps for both networks.

While this is a small prototype design, it is straightforward to scale the DDO global clock network to more advanced process technologies, higher frequencies, and larger areas. In a more advanced CMOS technology node, thick upper-level copper interconnect can be used for the inductors and clock distribution wires, resulting in substantially higher inductor and tank Q compared to the present prototype design for the same routing utilization. The robustness of DDO global clock networks to variability compared with tree-driven-grid global clocks should also make them more attractive with technology scaling. In contrast, the noise sensitivity of traditional tree-driven grids can only be expected to get worse with technology scaling as the latency of clock distributions (as measured in number of clock cycles) increases.

The DDO clock network easily scales to larger designs (for the same clock capacitance per unit area) by adding more tiles to the grid. Scaling to higher clock frequencies (or larger capacitance per unit area) can be achieved by reducing the size of the inductors or by adding more inductors to the grid (smaller tile). Adding more inductors (rather than scaling inductors to be smaller) consumes more wiring resources but pushes unwanted standing wave resonances in the grid to higher frequencies (as described in Section III-A).

B. Jitter Analysis and Injection Locking

The *period jitter*¹ is the most important metric determining the quality of the clock in high-performance digital designs and

¹Period jitter is often (incorrectly) used interchangeably with the term *cycle-to-cycle jitter*. However, the latter is given by $\sigma_{\text{cycle-to-cycle}} = \lim_{N \rightarrow \infty} \sqrt{(1/N) \sum_{k=1}^N (t_{k+1} - t_k)^2}$ and is not the same as period jitter.

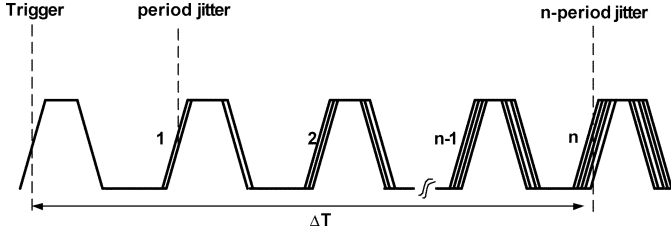


Fig. 6. Relationship between the reference edge and the clock edge for which jitter is being measured.

specifies the variation of the actual clock period, t_k , from the ideal period, t_{cycle} . The rms period jitter is given by

$$\sigma_{\text{period}} = \lim_{N \rightarrow \infty} \sqrt{\frac{1}{N} \sum_{k=1}^N (t_k - t_{\text{cycle}})^2}. \quad (5)$$

RMS n -period jitter, used to quantify jitter accumulation in oscillators, is given by

$$\sigma_{n\text{-period}}(\Delta T) = \lim_{N \rightarrow \infty} \sqrt{\frac{1}{N} \sum_{k=1}^N (\Delta T - nt_{\text{cycle}})^2} \quad (6)$$

where ΔT , as shown in Fig. 6, is the time to the n th active clock edge after the reference edge [25].

The jitter performance of these networks is largely limited by power-supply noise, which is highly correlated period-to-period. As a result, jitter accumulation proceeds through the addition of standard deviations rather than variances, and one finds that the rms n -period jitter varies linearly with the measurement interval ΔT [26]

$$\sigma_{n\text{-period}} = \eta \Delta T \quad (7)$$

which is confirmed in the measurement results presented in Section III. The constant of proportionality, η , can be estimated from the power-supply sensitivity of the network

$$\eta \cong \frac{S_{\text{ps}} \sigma_{\text{ps}}}{t_{\text{cycle}}} \quad (8)$$

where S_{ps} is the sensitivity of t_{cycle} to small-signal changes in power-supply voltage and σ_{ps} is the standard deviation of the power-supply noise. Fig. 7 shows simulated values of S_{ps} for the prototype clock network of Fig. 2 at three different values of I_{bias} :² 0.8, 2.0, and 5.0 mA. The gain elements are more sensitive to power-supply-noise-induced jitter at larger values of I_{bias} .

The injection locking of the clock network to a external reference limits jitter accumulation by providing phase filtering. The

² I_{bias} is nominally eight times smaller than I_{tail} , the bias current flowing through the gain elements.

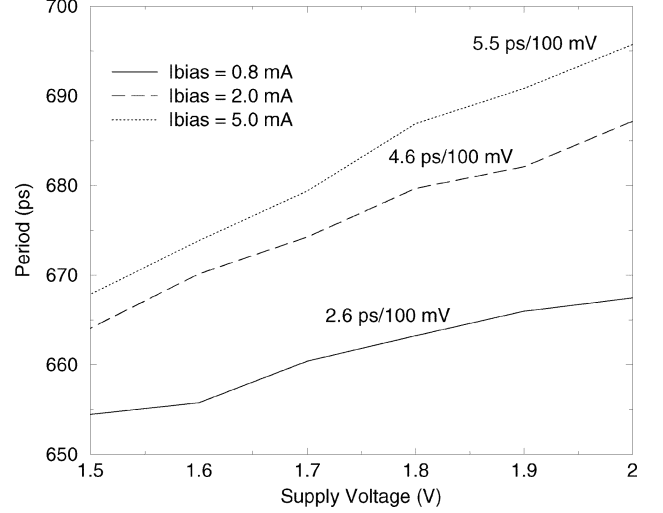


Fig. 7. Simulated values of S_{ps} , the variation in clock period as a function of power-supply voltage, for three different values of I_{bias} .

jitter transfer function from the external reference to the clock network is given by [27]

$$\frac{\sigma_{\text{clock}}}{\sigma_{\text{external}}} = \frac{S_c}{z - (1 - S_c)} \quad (9)$$

where S_c is the injection coupling strength given approximately by

$$S_c = \frac{I_{\text{inj}}}{I_{\text{ind}} + I_{\text{inj}}} \quad (10)$$

where I_{inj} is the peak current supplied to the clock network by the injection locking buffer and I_{ind} is the peak current flowing through the spiral inductors. This defines a time constant

$$\tau = -\frac{t_{\text{cycle}}}{\ln(1 - S_c)} \quad (11)$$

which determines how much the jitter is allowed to accumulate in the oscillator before the phase filtering action of the injection locker becomes significant.

C. On-Chip Jitter Measurement Circuits

Fig. 8 shows the schematic of the on-chip jitter measurement circuit, used to characterize the period jitter of both the resonant and nonresonant clock networks on the test chip [28]. A DLL is first used to divide a reference clock into 14 equally spaced clock phases. Two levels of calibrated interpolation are used to increase the number of clock phases by a factor of four to 56. Each of these 56 clock phases is then sent to the data input of a sense-amplifier-based flip-flop [29] (see inset in Fig. 8); each flip-flop thus receives at its input a unique clock phase, with the edges of each adjacent clock phase spaced by 1/56 of a clock period. Since adjacent latches receive adjacent clock phases, any difference in the output levels of two adjacent latches indicates

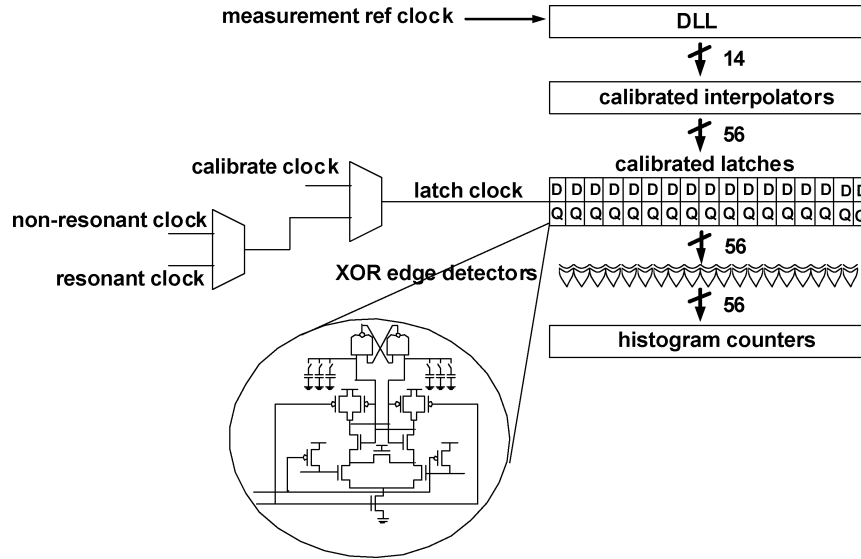


Fig. 8. On-chip jitter measurement circuit used to obtain a more precise jitter characterization of the clock networks on the test chip.

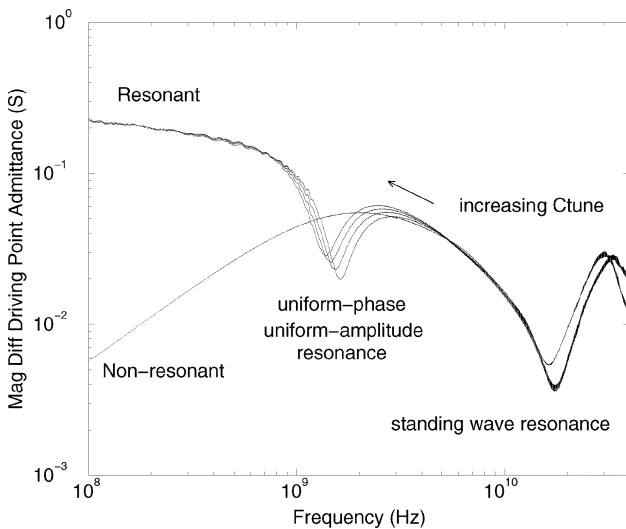


Fig. 9. The magnitude of the differential driving point admittance for the resonant and nonresonant clock networks as measured around zero common-mode bias on the test chip. For the resonant network, C_{tune} is swept from 0 to 5 pF in steps of 1.25 pF. A standing-wave resonance at 17 GHz is present in both networks.

the arrival of a clock edge in the associated interval. By XORing the outputs of adjacent latches, a “count” or “don’t count” signal is generated to trigger a set of counters to produce a jitter histogram.

Both the latch offset and the interpolation ratio are calibrated. As shown in Fig. 8, the latches can be triggered by either the resonant or the nonresonant clock, or an external calibrate clock. The calibration clock is an asynchronous clock that results in uniform distribution of clock edges in the histogram. The interpolation ratio can be adjusted to eliminate any phase spacing errors which result in nonuniform histogram distributions from the asynchronous calibration clock. As shown the inset of Fig. 8, the sense-amplifier latch is loaded by six switchable 3 fF MOS capacitors, allowing for the removal of static offsets.

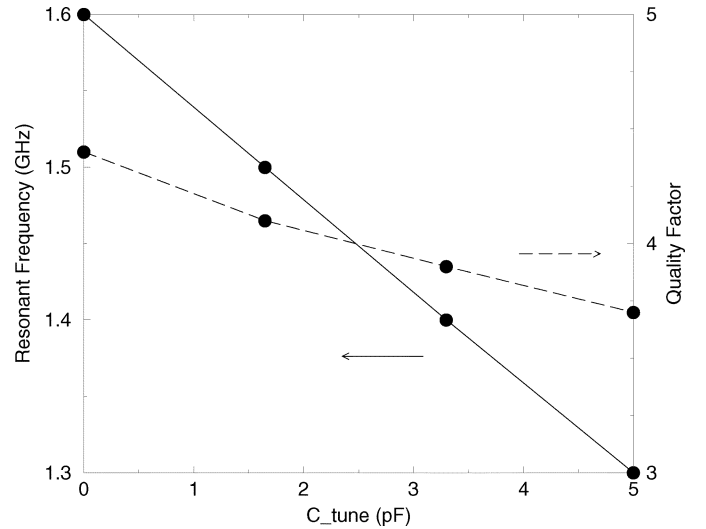


Fig. 10. Measured variation in resonant frequency and Q as a function of C_{tune} .

III. MEASUREMENT RESULTS

We perform frequency-domain characterization of the passive clock network as well as time-domain analysis of the operating DDO.

A. Frequency-Domain Behavior of the Passive Clock Network

Fig. 9 plots the magnitude of the differential driving-point admittance ($|Y_{\text{dd}}|$) of the resonant and nonresonant clock networks (around zero common-mode bias) as measured by probing at the center of clock grids (see Fig. 3) with GSG probes and an Agilent E8363B PNA Network Analyzer. For the resonant network, the family of curves correspond to C_{tune} from 0 to 5 pF in steps of 1.25 pF. At low frequencies, $|Y_{\text{dd}}|$ for the resonant network exceeds the nonresonant control, reflecting the difficulty in driving the resonant clock network differentially because the on-chip inductors short the two clock phases.³ As

³A low-frequency test mode would, therefore, be possible if the ϕ and $\bar{\phi}$ network are driven *together* at the same phase.

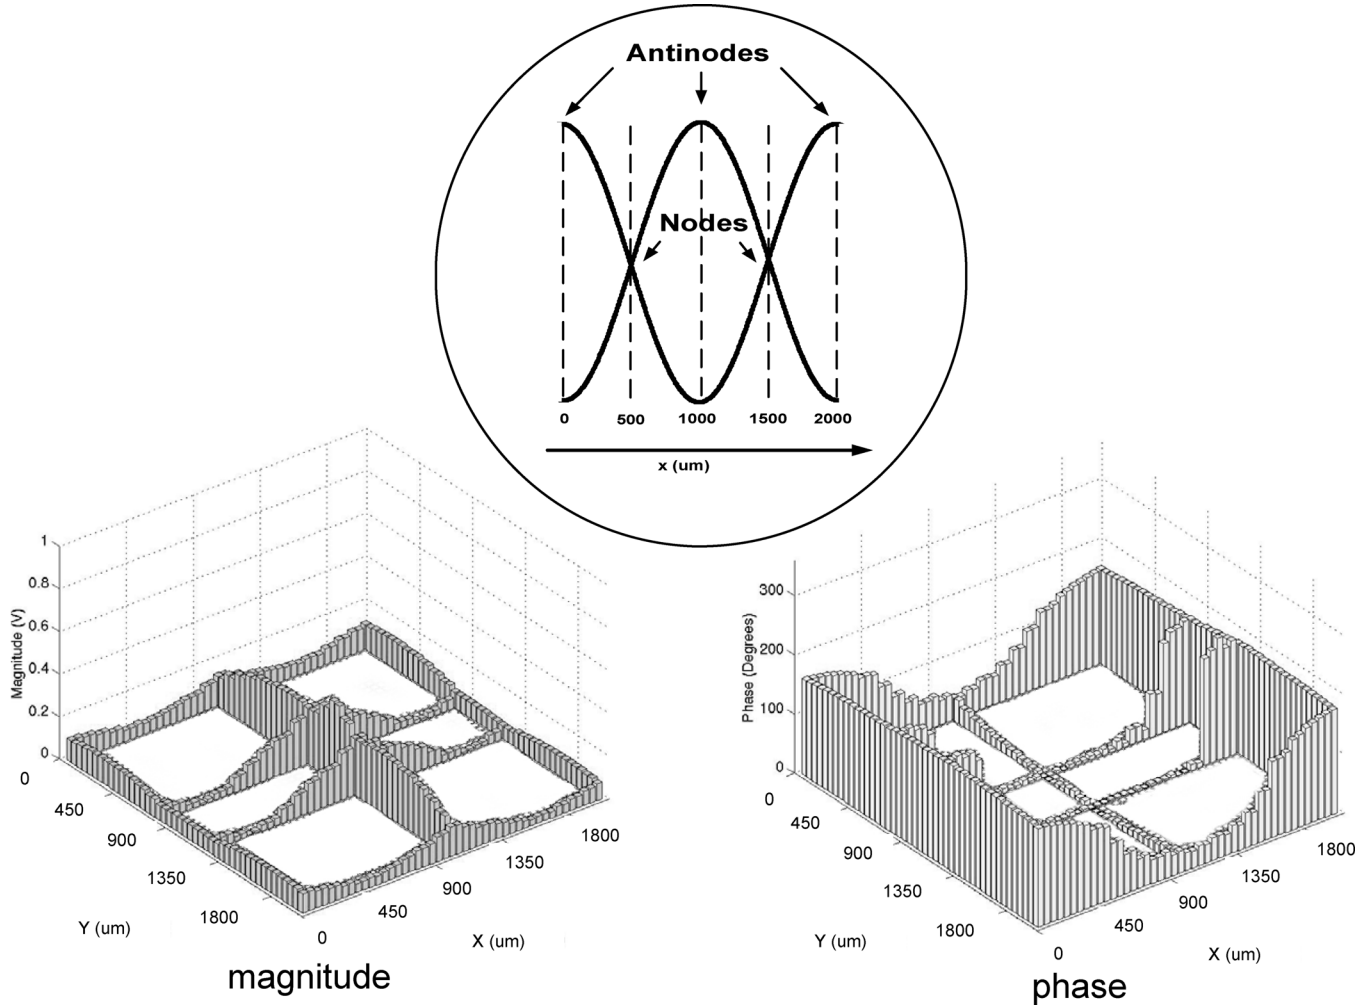


Fig. 11. Simulated magnitude and phase response of the resonant clock network around zero common-mode bias for $C_{\text{tune}} = 0$ pF when driven by a 50Ω source at 17 GHz. The magnitude varies sinusoidally across the distribution while the phase is uniform, with discrete 180° changes in phase at every $\lambda/2$ wavelength change in position. The inset shows schematically this standing-wave resonance along the x axis of the clock grid with two nodes and three antinodes.

the resonant frequency is approached, $|Y_{\text{dd}}|$ for the resonant network becomes less than the nonresonant control, indicating that the resonant network is “easier” to drive differentially.

Fig. 10 shows the variation in resonance frequency and Q as C_{tune} is varied. The measured Q (as determined from the 3-dB bandwidth at resonance) corresponds to $R_s \cong 7.3 \Omega$. The resonant frequencies obtained from the minimums in the admittance measurements differ slightly from the free-running oscillation frequencies of the clock network because of nonlinear device capacitance in the clock load. When $C_{\text{tune}} = 5$ pF, nonlinear gate capacitance is half the total tank capacitance and the resonant frequency varies by up to 9% with common-mode bias.

At 17 GHz, Fig. 9 shows a higher order resonance common to both the resonant and nonresonant networks. This standing-wave resonance is characterized by uniform-phase, but sinusoidally varying clock amplitude across the distribution. Fig. 11 shows the simulated magnitude and phase response of the resonant clock network at 17 GHz, demonstrating a standing-wave resonance along the x axis of the clock grid as shown schematically in the inset of Fig. 11 (with two nodes and three antinodes).

The network has salphasic behavior [6], [7] characterized by discrete 180° changes (at positions $x = 500 \mu\text{m}$ and $x = 1500 \mu\text{m}$) in an otherwise uniform phase distribution in the x direction. The slight deviations from constant 0° and constant 180° phase and the nondiscontinuous nature of the phase changes are due to losses in the network [7]. One full wavelength of the standing-wave is seen to span the entire 2 mm length of the resonant clock grid along the x -axis. At 17 GHz, this corresponds to a propagation velocity of approximately $0.1c_0$, where c_0 is the speed of light in vacuum, significantly slower than the expected $0.5c_0$ propagation velocity in silicon dioxide. This slow-wave effect [30] comes from the fact that the clock is a grid, the orthogonal wiring of which adds significant additional per unit length capacitance without contributing to inductive current return.

While this standing-wave resonance exists in the network when probed from the center of the clock grid, this resonance is unlikely to be excited in the actual network because of the action of the gain elements which force the network at these locations to be antinodes. Resonances at frequencies higher than 40 GHz (the highest frequency measured) are likely to exist. However,

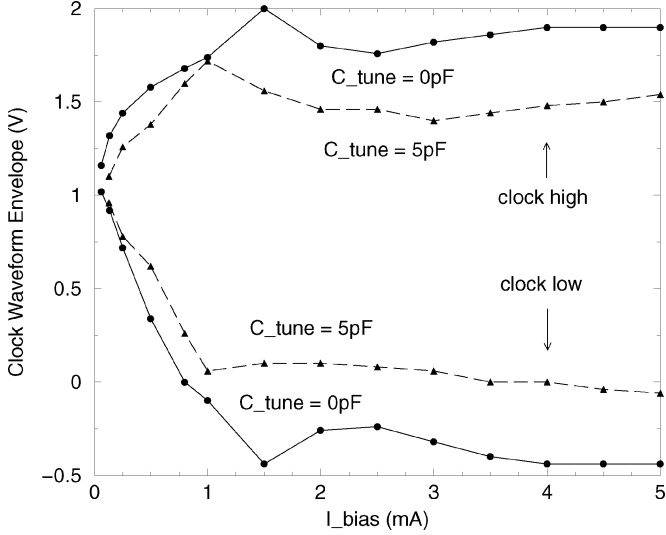


Fig. 12. Measured resonant clock waveform oscillation envelope (the high and low voltage levels of the clock) as a function of I_{bias} for $C_{\text{tune}} = 0$ and 5 pF on a 1.8 V supply.

since the interconnect loss at these frequencies is large, the impact of these resonances on skew and jitter at the nominal clock frequency is not expected to be significant. If the clock network is scaled to higher clock frequencies only by reducing inductor values, high order resonances can be expected to become a limitation as the inductance of the discrete inductors becomes comparable to the inductance of the grid and it becomes difficult to excite the desired fundamental resonance without also exciting standing-wave resonances.

B. Clock Amplitude

The minimum bias current, I_{bias} , needed to initiate oscillation of the DDO global clock network is measured to be 0.06 mA for $C_{\text{tune}} = 0$ and 0.13 mA for $C_{\text{tune}} = 5$ pF. At these minimum bias currents, the clock amplitude is measured to be only 140 mV.

Fig. 12 shows the envelope of the resonant clock oscillations (the high and low voltage levels of the clock) as measured using a high impedance active probe (GGB Model 34A with a 3 dB bandwidth of 3 GHz). The envelope voltage is plotted as a function of I_{bias} for $C_{\text{tune}} = 0$ and 5 pF. The clock amplitude increases and the envelope widens as I_{bias} increases beyond the minimum biases needed for startup.

The amplitude measurements in Fig. 12 and the admittance measurements in Fig. 9 can be used to confirm that (4) correctly relates the amplitude of the clock (V_{clock}) to the gain element bias current (I_{tail}) and the driving point admittance of the clock network (Y). Fig. 9 shows that the admittance Y per tile (as seen by a single gain element⁴) is 5.0 mS at 1.6 GHz for $C_{\text{tune}} = 0$ and 7.1 mS at 1.1 GHz for $C_{\text{tune}} = 5$ pF. Equation (4) therefore predicts that the clock will reach a full rail 1.8 V swing when $I_{\text{tail}} \cong 7.1$ mA ($I_{\text{bias}} \cong 0.88$ mA) for $C_{\text{tune}} = 0$ and $I_{\text{tail}} \cong 10$ mA ($I_{\text{bias}} \cong 1.3$ mA) for $C_{\text{tune}} = 5$ pF. Fig. 12

⁴The admittance in Fig. 9 must be divided by four in order to obtain the approximate admittance per tile.

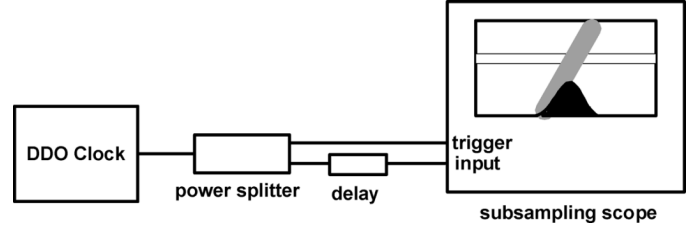


Fig. 13. Setup for measuring jitter on the test chip.

shows the clock reaching full rail by 0.9 mA and 1.0 mA of I_{bias} for $C_{\text{tune}} = 0$ and 5 pF, respectively. There is a slight discrepancy in using (4) for $C_{\text{tune}} = 5$ pF because nonlinear gate capacitance results in Y having a strong dependence on the common-mode voltage level of the clock network. A similar analysis using (4) and Figs. 9 and 12 when the clock is low swing (when the variation in Y with common-mode voltage is negligible and the assumption of square wave excitation by the gain elements holds) yields more accurate results. In particular, at start-up, when the voltage swing is measured to be 140 mV, (4) using measured Y predicts that $I_{\text{tail}} \cong 0.57$ mA ($I_{\text{bias}} \cong 0.071$ mA) for $C_{\text{tune}} = 0$ and $I_{\text{tail}} \cong 1.1$ mA ($I_{\text{bias}} \cong 0.13$ mA) for $C_{\text{tune}} = 5$ pF. These values of I_{bias} are in close agreement with the measured I_{bias} of 0.06 mA and 0.13 mA in Fig. 12.

C. Jitter

The jitter on the test chip is measured by converting the differential global clock into a full-rail single-ended clock using the local clock buffer shown in the inset of Fig. 2 [31]. This signal is analyzed with the on-chip jitter measurement circuits and by buffering the clock off-chip with an open-drain buffer into the input of a subsampling oscilloscope (Agilent 86100A) for external measurement. Fig. 13 shows the external period jitter measurement setup. Once off-chip, the clock is split to drive both the oscilloscope's trigger and input channel. The 24 -ns delay line in the input channel path compensates for ΔT , the latency internal to the oscilloscope between the trigger and the sampling event. ΔT on the oscilloscope can be varied between 24 ns and 2 μ s.

Fig. 14 shows the measured rms n -period jitter on the DDO global clock network, for $C_{\text{tune}} = 5$ pF at 1.1 GHz, as a function of ΔT . The jitter due to the measurement setup is removed by subtracting the rms jitter of the triggering edge ($\Delta T = 0$) from the rms jitter at a delay ΔT after the triggering edge [26]

$$\sigma_{\text{eff}} = \sqrt{\sigma_{\Delta T}^2 - \sigma_{\Delta T=0}^2}. \quad (12)$$

The case where the clock network operates as a free-running oscillator ($S_c = 0$) is shown as a solid line in Fig. 14, while the cases in which the clock network is injection-locked to an external reference is shown with a long dashed line ($S_c \cong 0.043$) and a short dashed line ($S_c \cong 0.083$). In all cases, $I_{\text{bias}} = 0.8$ mA. The linear relationship between σ_{eff} and ΔT for $\Delta T > 100$ cycles reflects the effect of highly correlated power-supply noise sources. Fig. 14 shows the injection-locked oscillator jitter

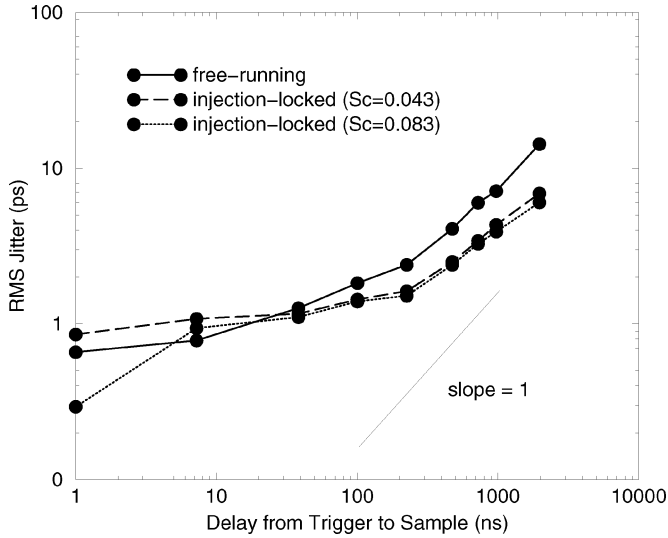


Fig. 14. Measured n -period jitter as a function of delay ΔT for the DDO global clock operating as a free-running oscillator and as an injection-locked oscillator.

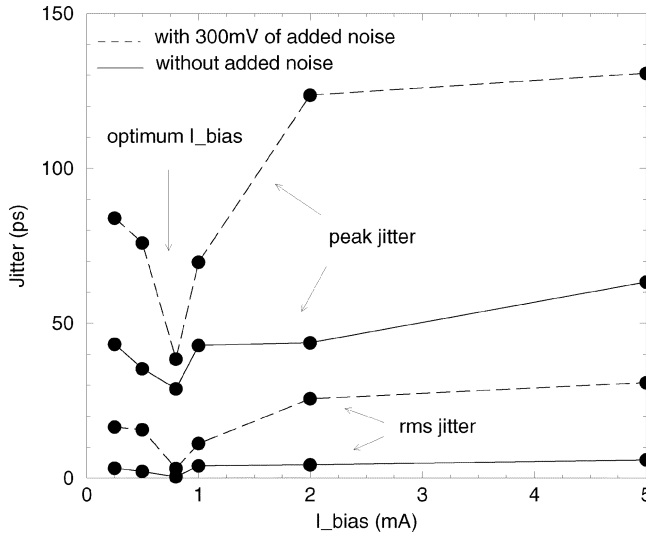


Fig. 15. Measured peak and rms period jitter in the resonant clock network as a function of I_{bias} with and without added power-supply noise for $C_{\text{tune}} = 5$ pF at 1.1 GHz.

deviating from the free-running oscillator jitter at around ΔT of over 40 cycles, consistent with the τ of (11).

Fig. 15 shows the measured period jitter (which is the most important timing uncertainty metric for the DDO as a global clock network) as a function of I_{bias} for $C_{\text{tune}} = 5$ pF at 1.1 GHz. Peak and rms jitter are plotted as a function of I_{bias} for no added power-supply noise and 300 mV of added power-supply noise. The added power-supply noise is introduced through MOS switches which periodically short the power supply to ground at approximately one-tenth the nominal clock frequency, 100 MHz, which is typical of the most problematic low-frequency resonances determined by package inductance and die capacitance [32]. At low values of I_{bias} , the clock has large jitter due to reduced clock signal swing. As I_{bias} is increased and clock signal swing increases, jitter decreases. Jitter is minimized by setting I_{bias} to be just large enough to ensure a close to full-rail clock (approximately 0.8 mA). Larger values of I_{bias} , which do not result in additional signal power,

degrade jitter because of the increased sensitivity of the gain elements to power-supply noise (see Fig. 7) at the same time that larger current demands on the power distribution network introduce additional power-supply noise.⁵

Figs. 16 and 17 shows the period jitter histograms as measured using both the on-chip and off-chip measurement circuits for the resonant and nonresonant clocks, respectively, for three different cases: (a) no added power-supply noise, (b) 150 mV of added power-supply noise, and (c) 300 mV of added power-supply noise. For the resonant clock network, $C_{\text{tune}} = 5$ pF, $I_{\text{bias}} = 0.8$ mA, $f_{\text{clock}} = 1.1$ GHz, and $S_c = 0.045$, while for the nonresonant clock network $f_{\text{clock}} = 1.1$ GHz. An external pulse generator (Agilent 81133A) at 1.1 GHz synchronizes all the clocks in the system (the measurement circuit reference clock, the reference clock feeding the injection locking buffer in the resonant network, and the reference clock driving the nonresonant network). The histogram bin size is 16 ps for the on-chip measurement which limits its resolution.

The shapes of the jitter histogram for the resonant clock is the same for all three cases, a single reasonably Gaussian distribution. The jitter histogram for the nonresonant clock is similar for case (a). However, for cases (b) and (c), the histogram has pronounced side peaks, which are most evident for case (c) with two large peaks centered around a smaller central peak. The histogram shape results from the introduction of low frequency power-supply noise. The clock buffers have periods in which their voltage is above the nominal supply voltage (resulting in transiently shorter clock cycles), and periods in which their voltage is below the nominal supply voltage (resulting in transiently longer clock cycles). These shorter and longer clock cycles give rise to the two peaks in the jitter histogram centered around the nominal clock period.

The measured jitter characteristics of the resonant and nonresonant clock networks are included in the summary Table I on the test chip. There is approximately an order of magnitude more jitter in the nonresonant clock network than in the resonant clock network, which shows the significance of the improved phase stability that this resonant clocking scheme enables. It is also clear from Table I that buffering the clock off-chip introduces considerable additional jitter.

D. Power

At 1.1 GHz, an average DC current of 28 mA is drawn from the supply of the resonant clock network for $C_{\text{tune}} = 5$ pF with $I_{\text{bias}} = 0.8$ mA, while the nonresonant clock network draws 81 mA. The nonresonant clock network has more than two times less load capacitance of the resonant clock network (7 pF compared with 18 pF), yet still consumes three times more power. The resonant clock network dissipates 2.8 mW/pF of load capacitance, while the nonresonant network dissipates almost an order of magnitude more, 21 mW/pF.

The clock distribution on a typical high-performance microprocessor may be up to ten times larger than the distribution designed on the prototype test chip (4 mm² on the test chip compared to 389 mm² on the IBM POWER5 microprocessor

⁵This trend is similar to that observed in LC oscillators in which thermal noise dominates. The least phase noise is observed for tail currents just large enough to drive the amplitude to its maximum value. [33], [34]

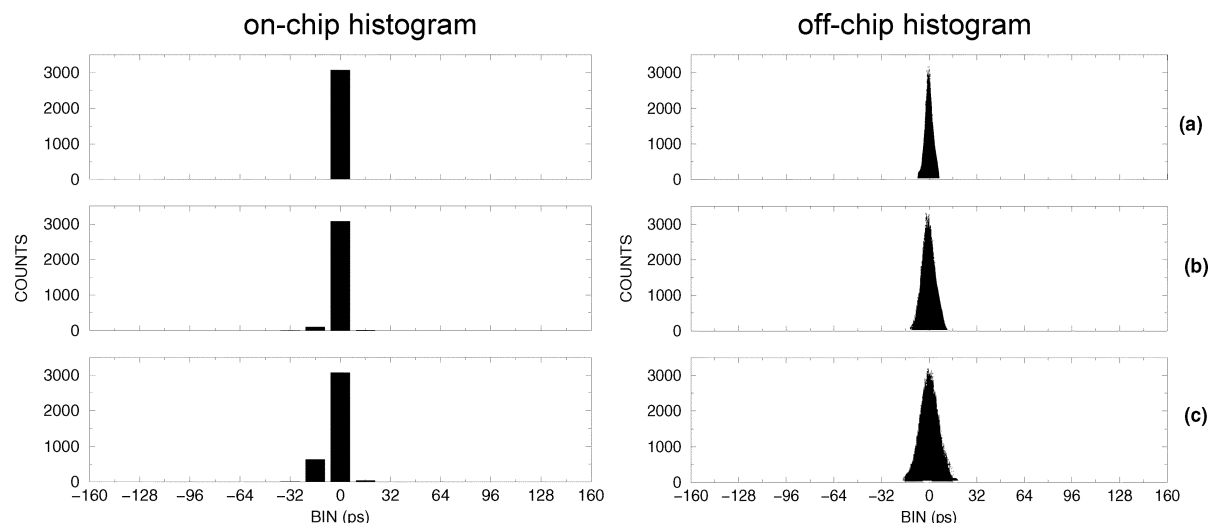


Fig. 16. Jitter histogram of the resonant clock operating at 1.1 GHz as measured using the on-chip jitter measurement circuits and as measured by buffering the clock off-chip: (a) no added power-supply noise; (b) 150 mV of added power-supply noise; (c) 300 mV of added power-supply noise.

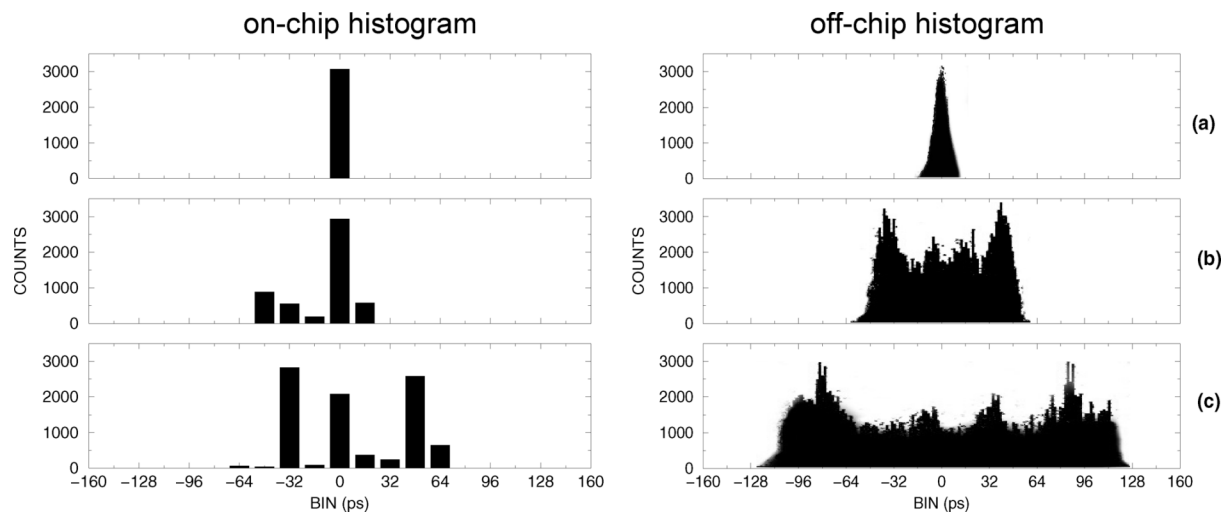


Fig. 17. Jitter histogram of the nonresonant clock operating at 1.1 GHz as measured using the on-chip jitter measurement circuits and as measured by buffering the clock off-chip: (a) no added power-supply noise; (b) 150 mV of added power-supply noise; (c) 300 mV of added power-supply noise.

TABLE I
MEASURED PERIOD JITTER AND POWER FOR $f_{\text{clk}} = 1.1$ GHz AND $C_{\text{tune}} = 5$ pF

| | resonant | | | non-resonant | | |
|---|-----------|------|--------------|--------------|------|--------------|
| | rms | peak | peak on-chip | rms | peak | peak on-chip |
| jitter – no added supply noise (ps) | 0.51 | 34 | 16 | 2.7 | 40 | 16 |
| jitter – 150 mV added supply noise (ps) | 1.8 | 40 | 64 | 31 | 122 | 80 |
| jitter – 300 mV added supply noise (ps) | 3.1 | 42 | 64 | 70 | 239 | 144 |
| average power (per unit load) | 2.8 mW/pF | | | 21 mW/pF | | |

[35] and 374 mm² on the third generation Intel Itanium microprocessor [36]). Scaling the DDO global clock network by a factor of ten in area would only result in approximately ten times more power being consumed because of the distributed oscillator topology. In a nonresonant tree-driven grid distribution or in the resonant-load distribution in [17], the power would increase by more than a factor of ten because a larger predriver

network (which is nonresonant) would be needed to distribute the clock across the larger chip.

IV. CONCLUSION

In this paper, we have presented a distributed differential oscillator (DDO) global clock network. The global clock capacitance resonates with a set of on-chip spiral inductors,

resulting in a clock which has uniform-amplitude and uniform-phase across the entire distribution. A prototype test chip implemented in a 0.18- μm CMOS technology shows strong immunity to power-supply-noise-induced jitter and low power dissipation, achieving approximately an order of magnitude less jitter and power than a conventional differential nonresonant tree-driven-grid global clock distribution.

REFERENCES

- [1] P. J. Restle *et al.*, "A clock distribution network for microprocessors," *IEEE J. Solid-State Circuits*, vol. 36, no. 5, pp. 792–799, May 2001.
- [2] T. Xanthopoulos, D. W. Bailey, A. K. Gangwar, M. K. Gowan, A. K. Jain, and B. K. Prewitt, "The design and analysis of the clock distribution network for a 1.2 GHz alpha microprocessor," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2001, pp. 402–403.
- [3] S. Tam, S. Rusu, U. N. Desai, R. Kim, J. Zhang, and I. Young, "Clock generation and distribution for the first IA-64 microprocessor," *IEEE J. Solid-State Circuits*, vol. 35, no. 11, pp. 1545–1552, Nov. 2000.
- [4] D. Pham *et al.*, "The design and implementation of a first-generation CELL processor," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2005, pp. 184–185.
- [5] J. Hart *et al.*, "Implementation of a 4th-generation 1.8 GHz dual-core SPARC V9 microprocessor," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2005, pp. 186–187.
- [6] V. L. Chi, "Salphasic distribution of clock signals for synchronous systems," *IEEE Trans. Comput.*, vol. 43, no. 5, pp. 597–602, May 1994.
- [7] F. O'Mahony, C. P. Yue, M. A. Horowitz, and S. S. Wong, "A 10-GHz global clock distribution using coupled standing-wave oscillators," *IEEE J. Solid-State Circuits*, vol. 38, no. 11, pp. 1813–1820, Nov. 2003.
- [8] J. Wood, T. C. Edwards, and S. Lipa, "Rotary traveling-wave oscillator arrays: A new clock technology," *IEEE J. Solid-State Circuits*, vol. 36, no. 11, pp. 1654–1665, Nov. 2001.
- [9] W. C. Athas, L. J. Svensson, J. G. Koller, N. Tzartzanis, and Y. Chou, "Low-power digital systems based on adiabatic-switching principles," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 2, no. 4, pp. 398–407, Dec. 1994.
- [10] W. C. Athas, L. J. Svensson, and N. Tzartzanis, "A resonant signal driver for two-phase, almost-nonoverlapping clocks," in *Proc. IEEE Int. Symp. Circuits and Systems (ISCAS)*, 1996, pp. 129–132.
- [11] W. Athas, N. Tzartzanis, L. J. Svensson, L. Peterson, H. Li, X. Jiang, P. Wang, and W.-C. Liu, "AC-1: A clock-power microprocessor," in *Proc. IEEE Int. Symp. Low Power Electronics and Design (ISLPED)*, 1997, pp. 328–333.
- [12] A. J. Drake, K. J. Nowka, T. Y. Nguyen, J. L. Burns, and R. B. Brown, "Resonant clocking using distributed parasitic capacitance," *IEEE J. Solid-State Circuits*, vol. 39, no. 9, pp. 1520–1528, Sep. 2004.
- [13] V. Gutnik and A. Chandrakasan, "Active GHz clock network using distributed PLLs," *IEEE J. Solid-State Circuits*, vol. 35, no. 11, pp. 1553–1560, Nov. 2000.
- [14] M. Saint-Laurent and M. Swaminathan, "A multi-PLL clock distribution architecture for gigascale integration," in *Proc. Int. Conf. Computer Design*, 2001, pp. 214–220.
- [15] L. Hall, M. Clements, W. Liu, and G. Bilbro, "Clock distribution using cooperative ring oscillators," in *Proc. 17th Conf. Advanced Research in VLSI*, 1997, pp. 62–75.
- [16] H. Mizuno and K. Ishibashi, "A noise-immune GHz-clock distribution scheme using synchronous distributed oscillators," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 1998, pp. 404–405.
- [17] S. C. Chan, K. L. Shepard, and P. J. Restle, "Uniform-phase, uniform-amplitude, resonant-load global clock distributions," *IEEE J. Solid-State Circuits*, vol. 40, no. 1, pp. 102–109, Jan. 2005.
- [18] F. E. Anderson, J. S. Wells, and E. Z. Berta, "The core clock system on the next generation Itanium microprocessor," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2002, pp. 146–147.
- [19] S. Tam, R. D. Limaye, and U. N. Desai, "Clock generation and distribution for the 130-nm Itanium 2 processor with 6-MB on-die L3 cache," *IEEE J. Solid-State Circuits*, vol. 39, no. 4, pp. 636–642, Apr. 2004.
- [20] P. Mahoney, E. Fetzer, B. Doyle, and S. Naffziger, "Clock distribution on a dual-core, multi-threaded Itanium-family processor," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2005, pp. 292–293.
- [21] Maxwell 2-D Field Solver. Ansoft Corp. [Online]. Available: <http://www.ansoft.com/products/em/max2d/>
- [22] M. Danesh and J. R. Long, "Differentially driven symmetric microstrip inductors," *IEEE J. Solid-State Circuits*, vol. 50, no. 1, pp. 332–341, Jan. 2002.
- [23] HFSS: High Frequency Structure Simulator 3-D Field Solver. Ansoft Corp. [Online]. Available: <http://www.ansoft.com/products/hf/hfss/>
- [24] C. P. Yue and S. S. Wong, "On-chip spiral inductors with patterned ground shields for Si-based RF ICs," *IEEE J. Solid-State Circuits*, vol. 33, no. 5, pp. 743–752, May 1998.
- [25] J. A. McNeill, "Jitter in ring oscillators," *IEEE J. Solid-State Circuits*, vol. 32, no. 6, pp. 870–879, Jun. 1997.
- [26] A. Hajimiri, S. Limotyrakis, and T. H. Lee, "Jitter and phase noise in ring oscillators," *IEEE J. Solid-State Circuits*, vol. 34, no. 6, pp. 790–804, Jun. 1999.
- [27] H.-T. Ng *et al.*, "A second-order semidigital clock recovery circuit based on injection locking," *IEEE J. Solid-State Circuits*, vol. 38, no. 12, pp. 2101–2110, Dec. 2003.
- [28] D. Weinlader, R. Ho, C.-K. K. Yang, and M. Horowitz, "An eight channel 36 GSAMPLE/s CMOS timing analyzer," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2000, pp. 170–171.
- [29] J. Montanaro *et al.*, "A 160-MHz, 32-b, 0.5-W CMOS RISC microprocessor," *IEEE J. Solid-State Circuits*, vol. 31, no. 11, pp. 1703–1714, Nov. 1996.
- [30] H. Hasegawa, M. Furukawa, and H. Yanai, "Properties of microstrip line on Si-SiO₂ system," *IEEE Trans. Microw. Theory Tech.*, vol. MTT-19, no. 11, pp. 869–881, Nov. 1971.
- [31] J. G. Maneatis, "Low-jitter process-independent DLL and PLL based on self-bias techniques," *IEEE J. Solid-State Circuits*, vol. 31, no. 11, pp. 1723–1732, Nov. 1996.
- [32] A. Muhtaroglu, G. Taylor, and T. Rahal-Arabi, "On-die droop detector for analog sensing of power supply noise," *IEEE J. Solid-State Circuits*, vol. 39, no. 4, pp. 651–660, Apr. 2004.
- [33] A. Hajimiri and T. H. Lee, "Design issues in CMOS differential LC oscillators," *IEEE J. Solid-State Circuits*, vol. 34, no. 5, pp. 717–724, May 1999.
- [34] J. J. Rael and A. A. Abidi, "Physical processes of phase noise in differential LC oscillators," in *Proc. IEEE Custom Integrated Circuits Conf. (CICC)*, 2000, pp. 569–572.
- [35] J. Clabes *et al.*, "Design and implementation of the POWER5 microprocessor," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2004, pp. 56–57.
- [36] J. Stinson and S. Rusu, "A 1.5 GHz third generation Itanium processor," in *IEEE Int. Solid-State Circuits Conf. (ISSCC) Dig. Tech. Papers*, 2003, pp. 252–253.



Steven Chan received the B.S. and M.S. degrees in electrical engineering and computer science from the University of California at Berkeley in 1996 and 1998, respectively, and the Ph.D. degree in electrical engineering from Columbia University, New York, in 2005.

From 1997 until 2001, he worked in electronic design automation on signal integrity and noise analysis of digital circuits. He is currently a Research Staff Member in the VLSI Design Department at the IBM T. J. Watson Research Center, Yorktown Heights, NY. His research interests include clocking for high-performance microprocessors, interconnect analysis, and on-chip test and measurement circuits.



Kenneth L. Shepard (SM'85–M'92–SM'03) received the B.S.E. degree from Princeton University, Princeton, NY, in 1987 and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 1988 and 1992, respectively.

From 1992 to 1997, he was a Research Staff Member and Manager in the VLSI Design Department at the IBM T. J. Watson Research Center, Yorktown Heights, NY, where he was responsible for the design methodology for IBM's G4 S/390 microprocessors. Since 1997, he has been at Columbia University, New York, where he is now Associate Professor. He also served as Chief Technology Officer of CadMOS Design Technology, San Jose, CA, until its acquisition by Cadence Design Systems in 2001. His current research interests include design tools for advanced CMOS technology, on-chip test and measurement circuitry, low-power design techniques for digital signal processing, low-power intrachip communications, and CMOS imaging applied to biological applications.

Dr. Shepard received the Fannie and John Hertz Foundation Doctoral Thesis Prize in 1992. At IBM, he received Research Division Awards in 1995 and 1997. He was also the recipient of an NSF CAREER Award in 1998 and IBM University Partnership Awards from 1998 through 2002. He was also awarded the 1999 Distinguished Faculty Teaching Award from the Columbia Engineering School Alumni Association. He has been an Associate Editor of IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION (VLSI) SYSTEMS and was the technical program chair and general chair for the 2002 and 2003 International Conference on Computer Design, respectively. He has served on the program committees for ICCAD, ISCAS, ISQED, GLS-VLSI, TAU, and ICCD.



Phillip J. Restle (M'87) received the Ph.D. degree in physics from the University of Illinois at Urbana-Champaign in 1986.

He is with IBM Research, Yorktown Heights, NY, where he has worked on CMOS modeling, package test, DRAM variable retention time, and on-chip transmission-line modeling. For the past decade, he has concentrated on tools and designs for high-speed clock distribution networks, and has contributed to more than a dozen high-performance microprocessors including IBM mainframes as well as the Power4, Power5, Power6, and the CELL processor.