Design and Analysis of Actively-Deskewed Resonant Clock Networks

Zheng Xu, *Member, IEEE*, and Kenneth L. Shepard, *Fellow, IEEE*

***Abstract*—Active deskewing is an important technique for managing variability in clock distributions but introduces latency and power-supply-noise sensitivity into the resulting networks. In this paper, an adaptively deskewed resonant clock network, based on an injection-locked distributed differential oscillator, is described, in which the delay lines required for deskewing are incorporated into the injection-lock source, dramatically improving jitter immunity. A power management system based on automatic amplitude control of the resonant grid further enhances energy efficiency. A prototype system operates at a nominal 2-GHz frequency in a 0.18 μm technology with on-chip jitter and skew measurement circuits and with more than 25 pF/mm^2 of clock loading.**

***Index Terms*—Clocking, resonant clocking, resonant, clock, power, skew, deskew, jitter, filtering.**

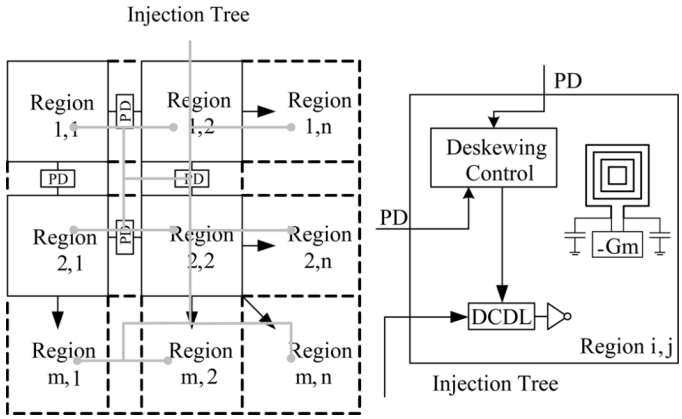

Fig. 1. System diagram of m-by-n system with deskewing.

I. Introduction

TRADITIONALLY, clock signals are distributed globally using either tree- or grid-based networks, requiring many stages of buffering and consuming a large percentage of system power [1]. Many levels of buffers also leave these systems sensitive to process, supply-voltage, and temperature (PVT) variability, both spatially and temporally. For a reasonably large design, cross-chip variations give rise to skew between different sections of the design. Active deskewing approaches [2] designed to mitigate these problems increase clock latency, making the system more sensitive to power supply noise and degrading jitter performance.

Resonant clocking techniques, in which the clock capacitance is rendered resonant around the target clock frequency by a set of on-chip inductors, addresses many of the challenges associated with standard tree- or grid-based networks. Power-supply-noise-induced jitter is significantly reduced and power savings can be realized in driving the global clock [3]. In the distributed-differential-oscillator (DDO) approach to resonant clock distributions, a differential global clock grid with a distribution of spiral inductors and loss-compensating negative transconductors forms a distributed LC oscillator, injection-locked to an external reference [4]. Despite the advantages demonstrated with these earlier resonant designs, resonant systems are still vulnerable to cross-chip variations and transient loading changes. With increasing system size, these differences in clock loading can lead to significant clock skew between different sections of the DDO clock network.

In this paper, we describe how a DDO resonant network can incorporate active deskewing without the jitter degradation associated with this approach in traditional tree-driven networks by incorporating the deskewing delays into the injection-lock network [5]. We also demonstrate active power management that can significantly improve the energy efficiency of resonant networks. Both control loops benefit from nearly all-digital implementations. In Section II, we describe the design of this actively deskewed DDO resonant clock network. Section III describes the properties of the digital deskewing control loop and the jitter-filtering properties of the injection-locked design. Specifics of the test chip implementation are described in Section IV with measurement results in Section V. Section VI concludes.

II. Actively Deskewed DDO Resonant Networks

A generalized DDO network incorporating active deskewing is shown in Fig. 1. The differential global grid, rendered resonant with symmetric inductors connecting the two phases and distributed throughout the grid, is divided into m-by-n regions which are deskewed with respect to each other. All the oscillating regions are connected together in shunt forming one distributed oscillator. Gain elements, to compensate for losses and sustain oscillation, are also distributed throughout the grid. (The test chip described in Section IV implements a two-by-two version of this more generalized network.) Clock load is determined by both wire loading and the gate capacitance of the local

Manuscript received November 26, 2007; revised September 11, 2008. Current version published January 27, 2009. This work was supported in part by the SRC FCRP C2S2 Center and the SRC GRC.

Z. Xu is with the Boston Design Center, Advanced Micro Devices, Boxborough MA 01719 USA (e-mail: zhengxu.xu@amd.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@cisl.columbia.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JSSC.2008.2010760

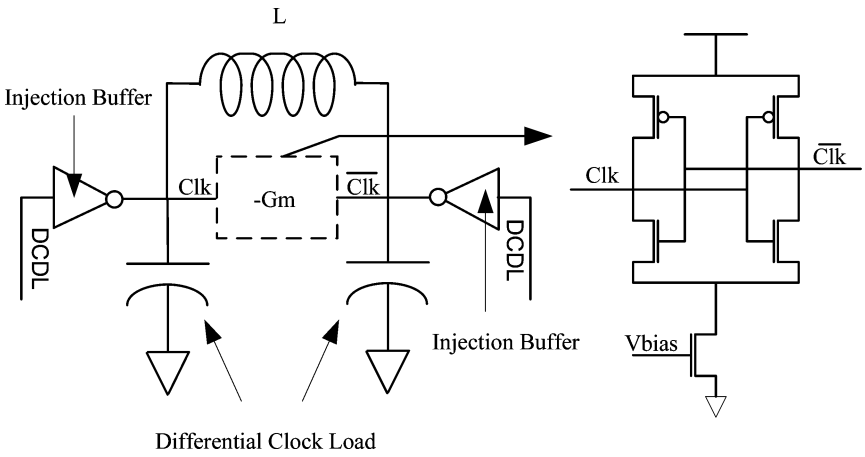

Fig. 2. Diagram of a distributed-differential-oscillator cell, and a transistor level implementation of the negative resistance element used.

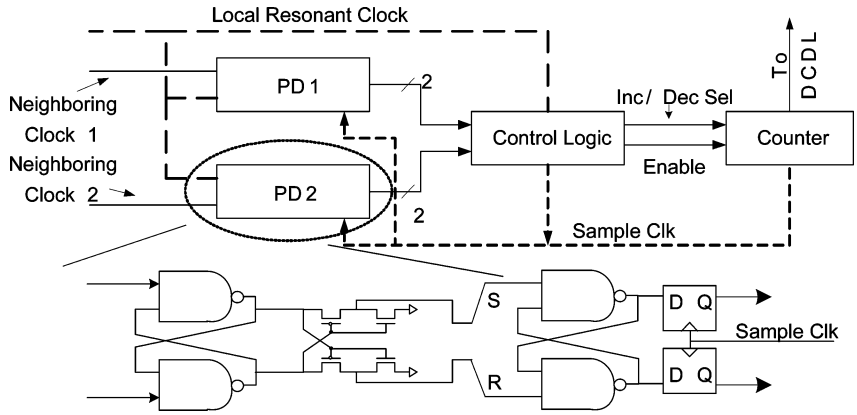

Fig. 3. System diagram of deskewing control.

clock buffers. The injection-lock source is constructed as a tree with each region of the grid associated with a leaf node of the tree. The delay of the injection-lock source to each region is independently controllable with digitally controlled delay lines (DCDLs). The output of each regional DCDL is then directly injected into the regional DDO cell as shown in Fig. 2. Phase detectors are placed between clock regions to detect skew with the information from the phase detectors passed onto the control circuitry in each clock region (see Fig. 3) determining the delay settings of the DCDLs. The deskewing system in each region is synchronized using a sample clock generated locally by the control logic and is a buffered version of the injection clock. One region is chosen as a reference with each region compared only to the neighbors that are closer to the reference region than itself. In this manner, closed loops can be avoided, preventing mode lock [6], [7]. While the system is primarily designed to distribute the resonant clocks to the leaf node clock buffers, it is also possible to bring the resonance directly to the flop-flops themselves [8], [9] for additional power savings.

In order to detect skew between different clock regions, a binary phase detector is used, implemented using an SR detector latch, metastability filter, and sampling latches [10] as shown in Fig. 3. To avoid wire delay mismatches in sampling the clock waveforms, phase detectors are physically placed at mid-points between the clock regions that they are designed to compare. Clocks are sampled at the center of each region and are connected to the phase detector with matched wiring.[1] The two clocks are compared at the S and R inputs of a SR latch, implemented using two NAND gates. The output of the SR latch is connected to a metastability filter, the output of which is buffered and feeds another SR latch. The output of the second SR latch is then sampled using two D-type flip-flops. For the cases in which two input clocks are very close in phase, the metastabiliy filter ensures only one of the SR latch's complementary output is asserted *HIGH* at any time.

A digital filter (contained within the "control logic" block of Fig. 3) is implemented to help reduce steady-state dithering. The filter is built as a cascade of n identical filter cells as shown in Fig. 4. When cascaded, the filter stages form an n-bit binary counter, only producing an inc or dec signal at the output when 2^n inc or dec signals are received at the input. While the filter helps to reduce steady-state phase dithering by a factor-of-two per filter stages, it also increases the locking time by the same factor. Despite the additional latency of the filter, locking transients generally involve a simple monotonic adjustment of DCDL through the action of the inc or dec signal.

The proper choice of granularity (region size) for the DDO network depends on many factors, such as the density of clock loading and the strength of the clock grid. When the clock grid

[1]While practical constraints might prevent the detectors from being placed at these optimal locations, it is possible to layout the wiring to ensure both inputs into the phase comparator experience comparable delays.

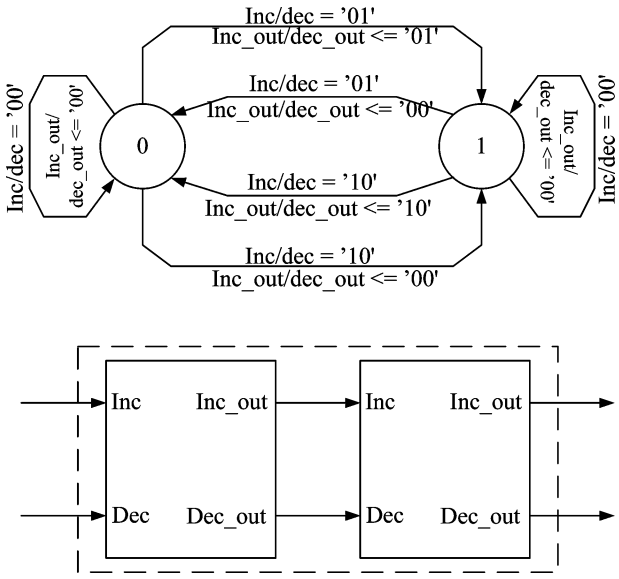

Fig. 4. State transition graph for a digital filter cell (top), the digital filter shown as a cascade of filter cells (bottom).

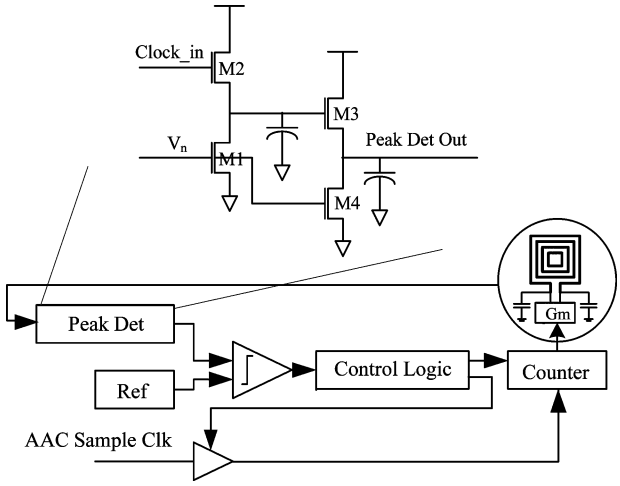

Fig. 5. System diagram of the automatic amplitude control system. A two stage cascaded source follow used as a simple peak detector, the current sinks are biased with very low current to capture the peak input voltage.

is dense and lightly loaded, a smaller number of larger regions can be employed. Generally, one would want to make size of a region as large as possible, while achieving tolerable skew limits. There is also a stability concern; if a dense clock grid is divided into too many regions, it will be more vulnerable to the positive feedback effect discussed in Section III-A.

It is also important to make the region size and loading uniform. If different regions have very different local natural frequencies, the systems could fail to lock or there could be enough skew that the deskewing system would be unable to compensate. Making the clock grid coupling stronger averages out variances in the natural frequencies, reducing skew, but due to the positive feedback effect discussed in Section III-A, making the clock grid between regions too strong reduces the system's ability to correct for the remaining skew.

Depending on the Q of the resonator and the strength of the negative resistance elements, when operating full-rail, it is possible to supply more power to the resonator than what is needed to sustain oscillation, resulting in wasted energy and giving power-supply noise greater influence on overall phase noise. Fig. 5 shows the block diagram of the automatic amplitude control (AAC) system designed to achieve optimal biasing of the gain elements, consisting of a peak detector, clocked comparator, counter, and control logic. On power up (or reset), the AAC starts with the lowest possible amplitude setting and increments the control counter until a desired amplitude is reached. After achieving this desired amplitude, the system goes into a standby mode in which it consumes negligible power. The AAC sample clock must be generated independently from the resonant clock to avoid start-up problems since the AAC system controls the resonant clock amplitude.

III. Properties of The Deskewing Control Loop

In this section, we consider the phase error and locking transients of actively deskewed DDO networks as well as the jitter-filtering properties of injection locking, which allows active deskewing to be introduced without degrading the jitter immunity of the network.

A. Properties of Actively Deskewed DDO Networks

Consider the generalized m-by-n network shown in Fig. 1. Once injection locking of the DDO network is achieved, the frequency of the resonant clock remains constant and the action of the deskewing control loop changes the phase of each region with respect to the phase of the injected clock. Region (1,1) denotes the reference region with the non-reference regions locking to the reference in order of their proximity. The worst case locking time is determined by Region (m, n). Each path from Region (1,1) to Region (m, n) determines a locking time bound of

$$T_{\text{lock}} = A_g \cdot \sum_{\text{path}} (\phi_{i+1} - \phi_i) \tag{1}$$

where ϕ_i is the clock phase in the ith region along the path and A_g is the rate at which phase error can be corrected. Since each region only takes reference from regions that are physically closer to Region (1,1), all paths from Region (1,1) to Region (m,n) have length $m + n$, and the path with the largest phase offset will determine the maximum locking time. Typical locking times will actually be shorter, since each region begins to match phase with its preceding region before the preceding region has locked to its own predecessor.

The skew-correction control loop has first-order dynamics. If the phase of a sector is different from that of the reference then the phase will be incremented or decremented by a fixed amount, α, in the direction of the difference, where α is the phase change produced by a single correction step in the DCDL of the injection source. Because of steady-state dithering (a limit-cycle oscillation) after lock, there will be some residual phase error (skew) in the network. For Region (i, j), this phase error is given by $\Delta\phi_{(i,j)} = \phi_{(i,j)} - \phi_{(1,1)}$, which can be *approximated* as the uniform sum distribution of $i + j - 2$ random numbers uniformed distributed between $-\alpha$ and α with zero mean. $\Delta\phi_{(i,j)}$ is then bounded by

$$-\alpha * (i + j - 2) < \Delta\phi_{(i,j)} < \alpha(i + j - 2) \tag{2}$$

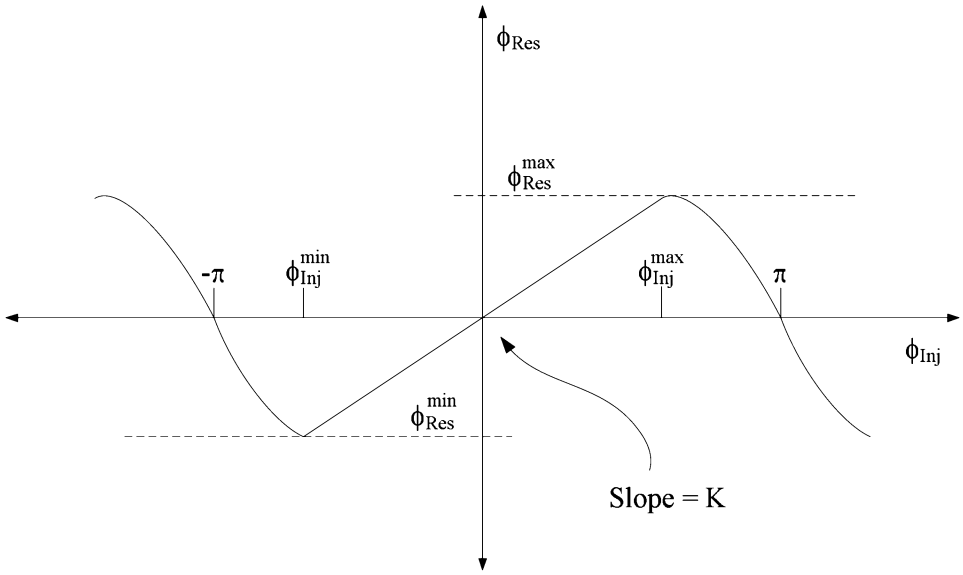

Fig. 6. Injection to resonant phase transfer function.

assuming a skew of α is added for each unit distance between the reference region and Region (i,j) with a total distance of $i+j-2$. The probability distribution function (PDF) for $\Delta\phi_{(i,j)}$ is then given by [11]

$$P(\Delta\phi) = \frac{1}{2\alpha(i+j-3)!}\sum_{k=0}^{i+j-2}(-1)^k \times \binom{i+j-2}{k}\left(\frac{\Delta\phi}{2\alpha}+\frac{i+j-2}{2}-k\right)^{i+j-3} \times \operatorname{sgn}\left(\frac{\Delta\phi}{2\alpha}+\frac{i+j-2}{2}-k\right). \quad (3)$$

Even though the residual phase error of Region (i,j) grows with increasing system size, it increases at a exponentially decreasing rate. Furthermore, the local dithering amplitude for any given region is bounded by $\pm\alpha T_i/2\pi$, where T_i is the period of the injected clock reference. Since the phase detector always try to match each region with its neighbor, the residue skew between neighboring regions after lock will never exceed one phase step, α, in this analysis.

More accurate estimation comes from detailed time-domain simulations of the injection-locked grid, which includes the effects of coupling through the grid and the latency of injection. These further influence the magnitude of the steady-state dithering and the resulting residual skew. We have noted several observations from these simulations:

- *Delay in injection locking*. It takes 3 to 20 cycles after a sudden change in the injection phase before the corresponding change in the resonant clock phase can fully stabilize. The number of cycles required depend on the injection strength and the Q of the slave oscillator. Higher injection strengths and lower Qs reduce the number of cycles needed—usually three to four cycles is sufficient. If the clock phase is sampled before its phase has fully stabilized, the resulting incorrect phase information can cause an incorrect decision in the phase correction circuitry and lead to additional dithering.
- *Reduced injection efficiency due to coupling through the grid*. Changes in the injection phase in each clock region also affect neighboring regions by the action of the clock grid. As a result, in order to achieve a phase shift $\phi_{\rm res}$ in the clock network, a phase change of $\phi_{\rm inj}$ is required in the injection clock with the nonmonotonic transfer function of Fig. 6. In the stable range of injection phases (the region of positive slope), the transfer function is given by

$$\Delta\phi_{\rm Res} = K\Delta\phi_{\rm inj} \quad (4)$$

  with $K < 1$.[2] The values of $\phi_{\rm Res}^{\rm max}, \phi_{\rm Res}^{\rm min}$ and K increase with decreasing clock grid density, since lower clock grid density increases isolation between different local regions, giving the local injection clock signal greater control. The deskewing control loop becomes unstable if the delay in the injection clock is large enough to reach the negative slope regions.
- *"Positive feedback" due to coupling through the grid*. Another consequence of coupling through the grid is that a correction step in a given region causes the clock phase of the region to which it is locked to also change in the same direction, potentially increasing the amplitude of limit-cycle oscillations.[3]

B. Jitter Filtering Properties of Injection-Locked DDO Networks

One of the major advantages of active deskewing in the context of injection-locked DDO networks is that the required delay lines are introduced into the injection-lock source. Despite the extra latency that this produces in the injection path, the filtering

[2]This transfer function also causes the phase difference across the final (inverting) injection buffer to become less than π, reducing the effective injection strength and resulting in static power dissipation in this driver. Therefore, large K values should be avoided.

[3]An alternative is to decouple all the clock regions and have each running independently [12]. Unfortunately, decoupling the clock regions eliminates the variability-averaging effect provided by the shunting common grid.

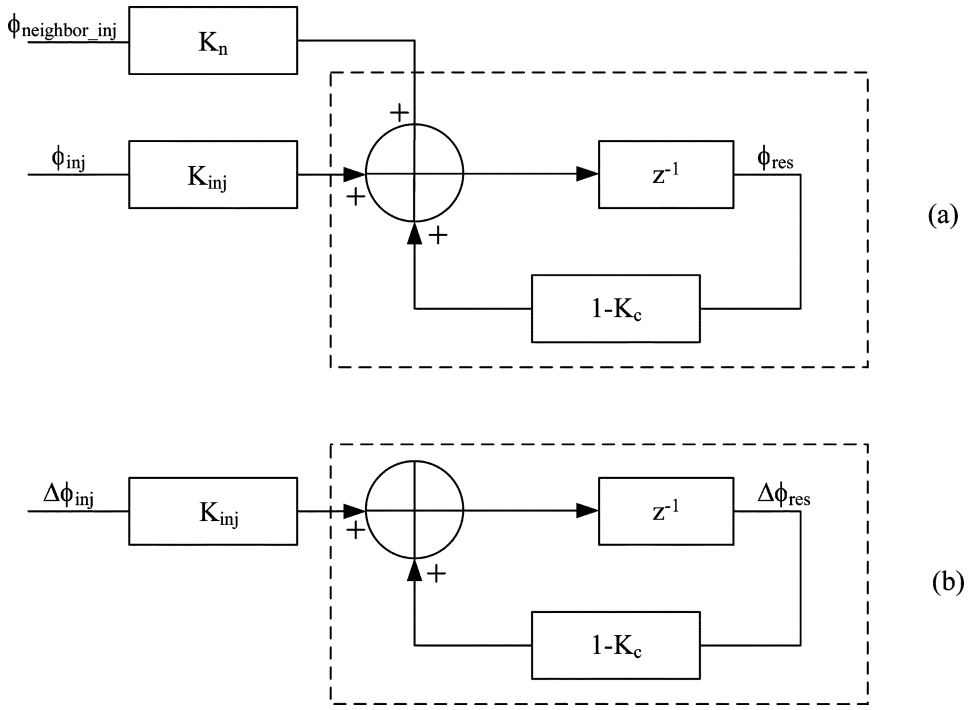

Fig. 7. System model for injection locking with multiple sources: (a) the case of system being influenced by the neighboring regions as well as the injection source. (b) the case in which the phase of the neighboring region is held constant, which model the change in resonant clock phase due to changes in the injection source.

properties of injection locking ensure that the overall system jitter performance is not compromised.

In injection locking, or entrainment, an oscillatory system is synchronized to an external frequency through the introduction of an external stimulus. The injection locking of electrical oscillation has been thoroughly studied beginning with the seminal work of Adler [13]. Fig. 7(a) shows a simple discrete-time model for the phase of a particular region of the injection-locked DDO system. The summer represents the effect of injecting an external signal into the system. The z^{-1} element represents the delay require for the oscillator to adjust to changes in the injection source. While this model was originally used to characterize a ring oscillator system [14], all the same elements apply to LC-based resonators. $\phi_{\rm res}$ is the resonant clock phase of a given region, and $\phi_{\rm inj}$ is the injected clock phase for that region. $\phi_{\rm neighbor\_inj}$ models the phase of one of the neighboring regions which influences of the phase of the given region through the coupling grid. Other coupling regions can be modelled by additional $\phi_{\rm neighbor\_inj}$ factors with their own characteristic K_n values. $K_{\rm inj}$ can be expressed as the ratio of peak injected current to peak current in the resonant region [14], [15], [3]:

$$K_c = \frac{I_{\rm inj}}{I_{\rm inj} + I_{\rm res} + I_{\rm neighbor}} \tag{5}$$

where $I_{\rm inj}$ is the peak injected current, $I_{\rm res}$ is the peak amplitude of the current in the region and $I_{\rm neighbor}$ is the net current entering the resonator for neighboring regions. K_n is determined by how strongly the clock regions are connected together in the resonant grid (reflecting either dense or sparse wiring) and can be expressed as the ratio of the component of the resonant current from neighboring resonators to the total resonator current:

$$K_n = \frac{I_{\rm neighbor}}{I_{\rm res} + I_{\rm neighbor} + I_{\rm inj}} \tag{6}$$

where $I_{\rm neighbor}$ is the net current entering the resonator from neighboring regions. The combined injection coefficient, K_c, can then be expressed as

$$K_c = \frac{I_{\rm inj} + I_{\rm neighbor}}{I_{\rm res} + I_{\rm neighbor} + I_{\rm inj}} = K_{\rm inj} + K_n \tag{7}$$

with $K_{\rm inj}, K_n$ and K_c are all positive constants less than one.

Assuming $\phi_{\rm neighbor\_inj}$ remains constant, the change in $\phi_{\rm res}$ due to a change in $\phi_{\rm inj}$ can be modeled by Fig. 7(b). We can derive the injection-to-resonant jitter transfer function as

$$\frac{\Delta\phi_{\rm res}}{\Delta\phi_{\rm inj}} = \frac{K_{\rm inj}}{z - (1 - K_c)}. \tag{8}$$

The transfer function of (8) clearly exhibits a low-pass characteristic with a cut-off frequency ω_c given by

$$\omega_c = \frac{ln(1 - K_c)}{T_i} \tag{9}$$

where T_i is the injection clock period.

In addition to the low-pass characteristic of injection locking which helps to isolate the resonant clock from the noise in the injection source, other characteristic of injection locking also help jitter performance. Given a relatively clean injection clock, injection locking helps to reduce the phase noise (or jitter) in the slave oscillator [16]. This filtering effect is most effective when

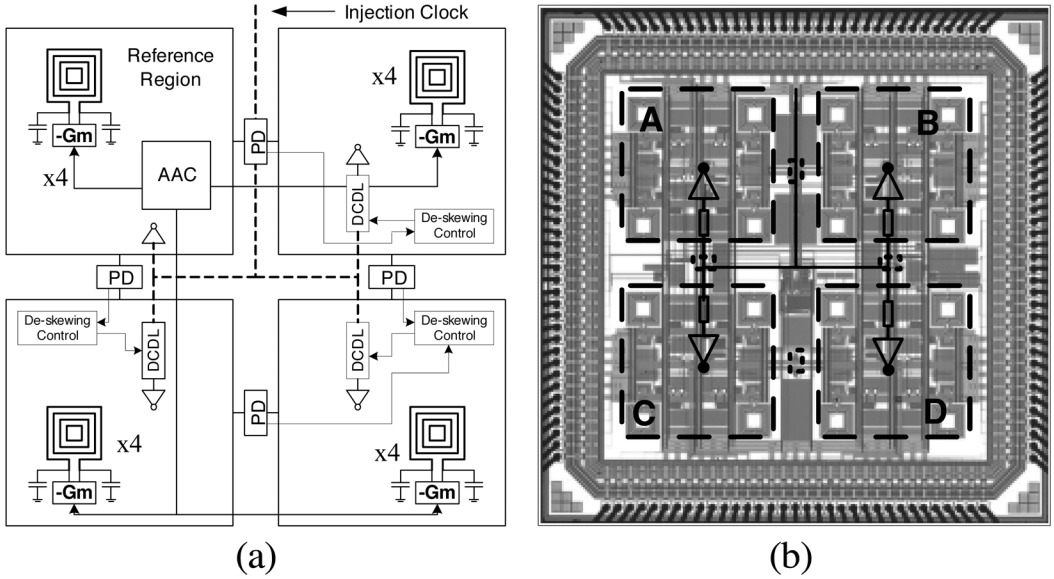

Fig. 8. (a) System diagram of the test chip, the components are labeled. (b) Die photo of the test chip, with key components labeled.

the injection frequency is near the natural resonance of the slave oscillator. [17], [9]. In the same way, as the injection clock get closer to the edge of the injection-lock range, the low-pass filtering of noisy injection sources becomes more effective since the injection clock's influence over the resonant clock network weakens. This effect is manifest through a reduced injection strength, K_c, moving the pole in (9) closer to zero.

IV. Test Chip and Measurement Circuitry

A test chip is implemented using a six-level-metal TSMC 0.18 μm technology. This technology has a FO4 delay of approximately 70 ps and the clock tree and grid are constructed from metal layers with a sheet resistance of 0.078 $\Omega/\square$. The prototype clocking system implemented on the test chip allows us to measure and confirm the predicted properties of the deskewing and amplitude control features. The test chip runs at a nominal injection frequency of 2 GHz, with an injection range of around 300 MHz.

A. Test Chip Features and Specifications

The chip is 3-mm-by-3-mm and is divided into four clock regions in a two-by-two configuration. The die photo and corresponding system diagram are shown in Fig. 8. In each clock region there are four inductors and four negative transconductors distributed evenly on a differential clock grid (distributing clock and $\overline{\text{clock}}$). The grid in each region is composed of vertical and horizontal metal lines with a spacing of 200 μm and width of 6 μm. The vertical strips are made using level-four metal (M4), and the horizontal strips are level-five metal (M5). Both vertical and horizontal routes have comparable resistances of 0.013 Ω/μm.

The four clock regions are connected together forming a large distributed resonant oscillator. The connections between the four clocks regions are made using relatively thin interconnect, 1 μm compared to the 6 μm used in the grids. These weak connections are designed to emulate a longer grid, increasing the effective distance between different regions by a factor of six. With four sets of inductor and negative resistance element in each clock region the entire test chip is composed of sixteen resonator "units" running in parallel.

The total clock loading in each one-mm-by-one-mm region is approximately 25 pF which includes fixed MOS clock load of 10.5 pF, a variable MOS load of 0.6 pF to 5.4 pF, 2 pF of wiring capacitance, and additional parasitic capacitances associated with the gain element and injection buffers. The variable MOS loads are implemented as pass-transistor-controlled switchable loads, configured digitally using scan registers. The clock loads distributed evenly throughout the clock grids with four banks variable cap per clock region. The inductance values are set so that the natural resonance frequency of the distributed resonator is approximately 2 GHz. Each of the sixteen inductors on the test chip is sized to around 3 nH.

The clock grid is synchronized to an external clock signal through an injection tree, which is laid out in an H-tree fashion, distributing the clock signal to all four clock regions. The DCDLs are implemented as chains of CMOS inverters, each with a variable load controlled by a counter as shown in Fig. 9. The DCDLs have a programmable delay range of between 800 ps and 1.1 ns on steps of 4.5 ps. The counter combines binary and thermometer codes to achieve both high linearity and high dynamic range. The relatively poor power-supply-noise rejection of this implementation [18] is mitigated by the filtering properties of injection locking. The injection clock is distributed as a single-ended signal and is converted to differential form before the final injection buffers. The injection buffers are composed of a set of tristate buffers with programmable, three-bit-binary-coded injection strength.

In order to characterize clock jitter performance and its sensitivity to power supply noise, power supply noise generators are included in the test chip. Each generator is composed of a single NMOS transistor with the drain connected to power, the source connected to ground, and the gate connected to a chip-wide digital noise input clock. Toggling the noise clock introduces noise of the same frequency onto the power supply node. There are

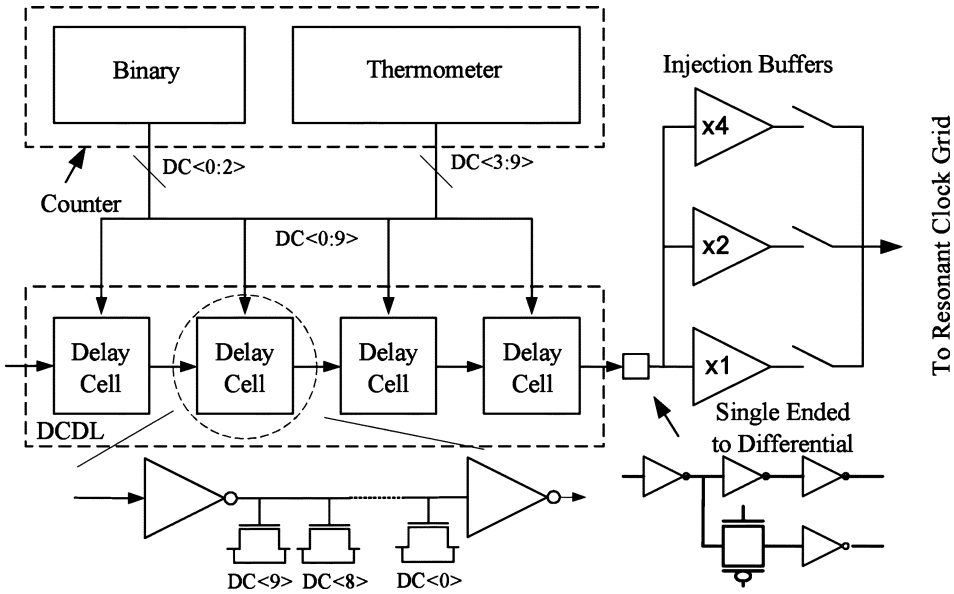

Fig. 9. Counter, delay line and injection buffer.

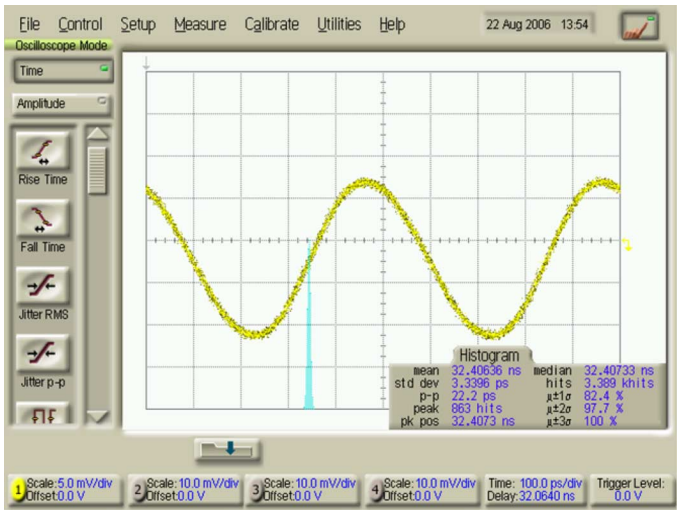

Fig. 10. The resonant clock waveform measured through an open drain driver on the test chip. Since the open drain driver is used, the absolute amplitude information is lost but the shape of the waveform is maintained.

four noise generators for each clock region, equidistant from the center.

Normally, it is not possible to operate a resonant clock network at the low frequencies necessary for certain test modes [19], [20], because the inductor becomes a short at these frequencies. This test chip, however, supports a mode in which the clock network can operate at frequencies below 200 MHz as a conventional tree-driven grid. In this mode, the negative resistance elements are disabled and the injection clock tree directly drives both the clock and $\overline{\text{clock}}$ grids together in a single-ended manner. The injection buffers have sufficient strength at these frequencies to generate full-rail clock signals on the grid.

B. On-Chip Measurements

On-chip skew and jitter measurement circuits constitute an important part of the test chip design, avoiding the jitter and skew introduced by off-chip drivers. In an effort to reduce offset, the on-chip measurement block is placed in the middle of the test chip, ensuring it is equidistant to all four clock regions. Both on-chip jitter and skew measurements are done using the same circuitry which can be configured to do either measurement. The measurement circuits run on their own power supply and are not affected by the noise generators.

Period jitter is measured as shown in Fig. 11, in which two delay lines nominally differing in delay by a clock period drive the clock and data inputs of a differential sense-amplifier flip-flop [21]. A reference counter counts every clock edge, while the test counter counts every clock edge for which the output of the flip-flop is *HIGH*. Once the reference counter reached its maximum value, both counters are stopped. If the delay of the variable delay line is swept from a delay slighter less than a period to a delay slightly more than a period, the ratio of the two counters gives the cumulative distribution function (CDF), $f(t)$, the derivative of which, $f'(t)$, represents the jitter distribution. The Agilent 81133A signal generator used to generate the reference clock has jitter of less than 0.1 ps.

Skew measurements (for example, between clock_1 and clock_2 of two different regions of the grid) are performed using the same circuitry but with different clocks feeding the two delay lines. A reference clock, derived from the injection clock source, feeds the bottom delay line, while clock_1 is fed into the top delay line. With this input selection, the delay of the second delay line is swept and the corresponding counter values are stored. The same is done for clock_2. The skew between clock_1 and clock_2 is then determined by the temporal distance between the two CDFs. The two corresponding CDFs are shown as $f'(t)$ and $g'(t)$ in Fig. 11. The on-chip measurement is calibrated with an external reference, through which a delay-to-control-voltage calibration function is performed.

V. Measurement Results

Resonant Clock Waveform: The resonant clock waveform on the test chip is observed by means of an open-drain driver and

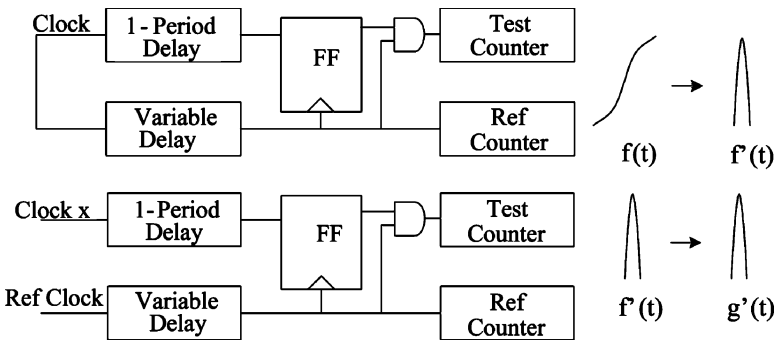

Fig. 11. System diagram of the jitter and skew measurement circuits.

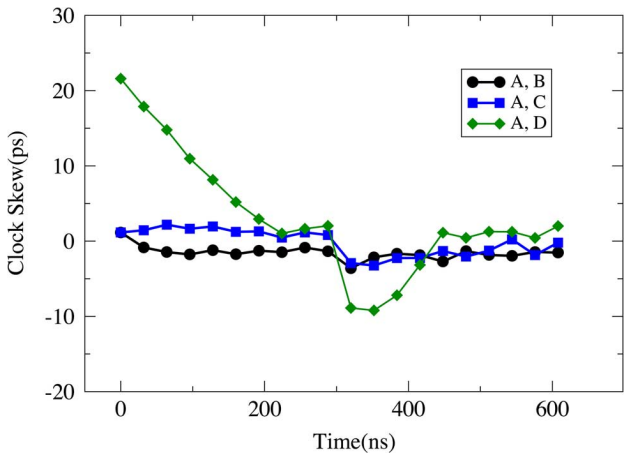

Fig. 12. Measured skew correction behavior of the test chip. The each curve shows the skew of in each region from with respect to the reference region.

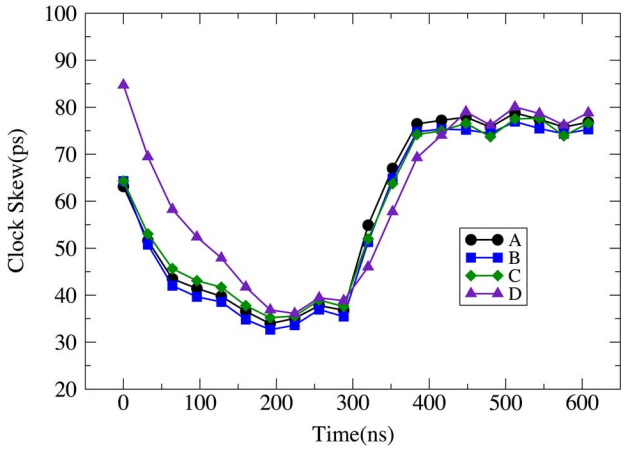

Fig. 13. Measured skew between each clock region and a fixed reference clock (derived from the injection clock). As shown in this figure, the absolute value of the phase in each region changes significantly as Region D attempts to correct its phase.

is displayed on an oscilloscope as shown in Fig. 10. The clock waveform is sinusoidal, a typical shape for a resonant clock output. A similar shape waveform is observed when the resonant grid is probed directly with an active high-bandwidth picoprobe.

Skew: Fig. 12 shows the measured dynamics of skew correction. During this measurement, the clock used in the deskewing control circuitry is brought from off chip. The time scale of the correction in Fig. 12 is determined by the 500-MHz on-chip sample clock. By providing Region D with an initial capacitance of 2 pF/mm^2 over that of Regions A, B, and C, Region D has a skew of 22 ps when driven from a balanced injection lock source with an injection coupling strength of $K_{\text{inj}} = 0.08$. The skew correction loop corrects this skew by $t = 300$ ns, at which point another capacitance offset of 1 pF/mm^2 is introduced, which is further corrected. For $K_{\text{inj}} = 0.08$, approximately 75 ps of delay in the injection lock source is required to compensate for each 1 pF/mm^2 offset.

To provide more insight into how the skew is being corrected, Fig. 13 shows the phase of each region with respect to a fixed external reference. As Region D is decreasing its phase in an effort to match the other regions, the phases of the other regions are also changing in the same direction. This reflects the parasitic feedback described in Section III. This four-region test chip should have been given a larger "effective" area by further weakening the grid routes between regions. Large amounts of feedback lead to increased locking time, reduced effective injection strength, and additional steady-state dithering. These problems are minimized because only four regions are present on the test chip. The absolute clock skew does settle once the active deskewing system has performed the skew correction.

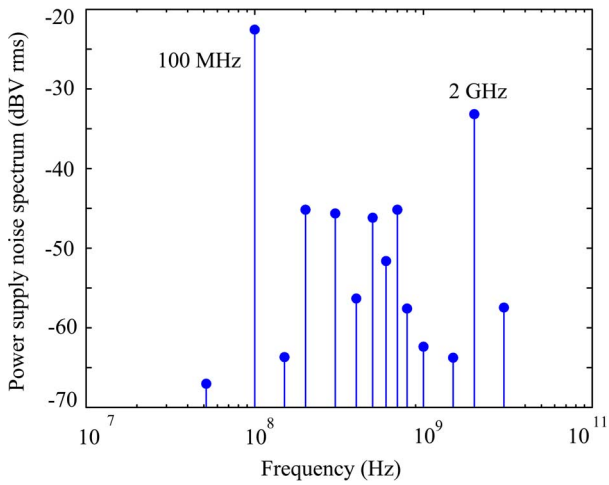

Fig. 14. Power supply noise spectrum (0 dB is 1 V rms) measured through active picoprobing with a 100-MHz-square-wave stimulation of all of the on-chip noise generators.

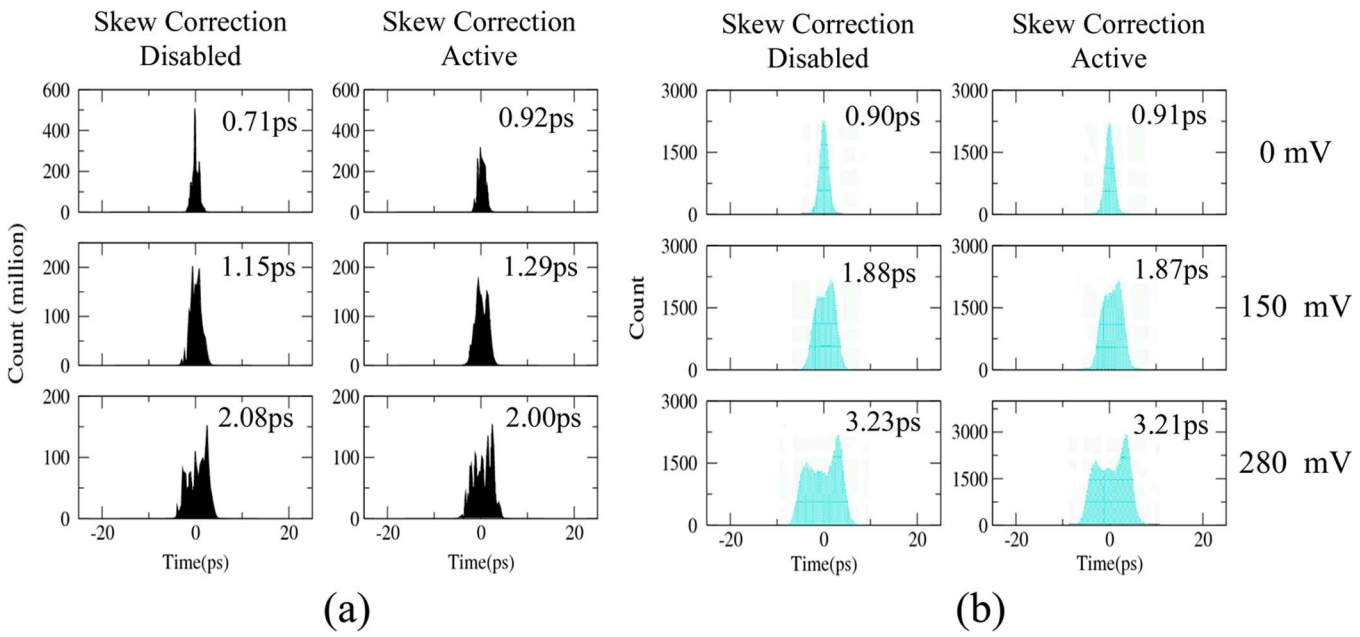

Fig. 15. Jitter histogram from on-chip and off-chip measurements. The graph shows RMS period jitter at three different levels of power supply noise. (a) on-chip results (b) off-chip results, the two are fairly close.

Jitter: Power-supply noise is the primary source of jitter in most clock distributions. The power supply noise is created by the noise generator described in Section IV and is measured using active probing sites on the chip. Fig. 14 shows the noise spectrum (in dBV rms) in the case that all the noise generators are activated by a 100 MHz square wave, producing a noise amplitude of 300 mV. In addition to features at 100 MHz and its harmonics, a significant component at 2 GHz is also observed. Fig. 15 shows the jitter histogram with no added power supply noise and with 150 mV and 280 mV of added supply noise using a square wave on the noise generator at 50 MHz. The clock jitter, as expected, increases with increasing power-supply noise. The jitter that is measured off-chip is slighting larger then the on-chip measurement, mainly due to noise introduced in buffering the clock off chip and in the off-chip measurement setup. The activation of the deskewing system has little impact on the jitter performance.

In order to verify the jitter filtering properties of injection locking, the period jitter is measured as a function of jitter frequency in the injection source. Fig. 16 shows the ratio of measured jitter in the resonant clock to the measured jitter in the injection source as a function of the noise frequency of the injection source at different injection strengths. While the injection jitter (generated by a voltage-delay modulator in the the Agilent 81133A waveform generator) is measured at the source before it goes onto the test chip, it should be close to the actual jitter in the injection tree. As predicted by (8), the system has a low-pass response. The frequency cut-off, ω_c, increases with increasing injection strength as predicted by (9). With $K_{\text{inj}} = 0.094$, the predicted cut-off frequency of $\omega_c = 197$ Mrad/s or 31 MHz is close to the measured value. Stronger injection strengths give the injection source more control over the resonator. When noise in the injection source dominates, stronger injection strength results in greater noise in the resonator.

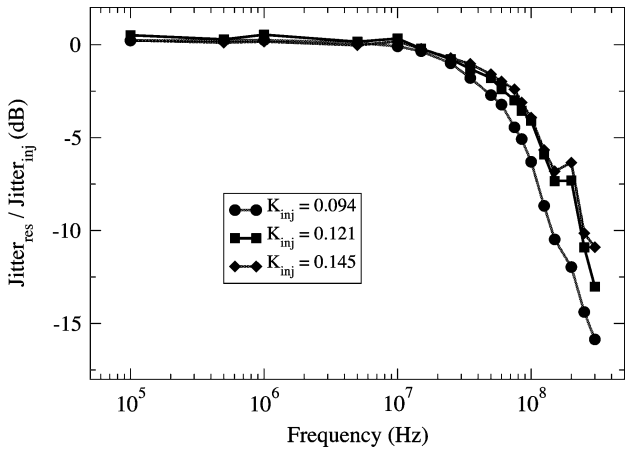

Fig. 16. Jitter transfer between the injection source and the resonant clock is plotted as a function of jitter frequency.

Automatic Amplitude Control: The operation of the automatic amplitude control (AAC) system is also measured. This is done by overriding the AAC sampling clock with a independently controlled signal supplied by an off-chip generator.[4] For each period of the sample clock, the resonant clock amplitude is measured using active probe sites on the test chip. The clock amplitude and the corresponding resonator current consumption are plotted as a function of correction step as shown in Fig. 17. The x-axis is mapped to the time scale determined by the sampling clock at 1 kHz. The system behaves as expected; the amplitude starts from its lowest level and increases until it reaches the desired amplitude. After exceeding the desired amplitude bound, the system backtracks one step, and remains steady at this amplitude. As seen in the Fig. 17, the power consumption

[4]It is also possible to use a simple ring oscillator on-chip to perform the same function.

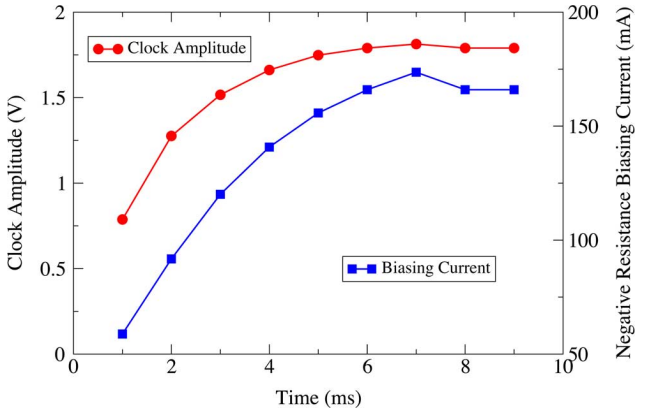

Fig. 17. On-chip clock amplitude as AAC is seeking the correct levels, measured with active pico-probes. Both the clock amplitude and the corresponding biasing currents are shown.

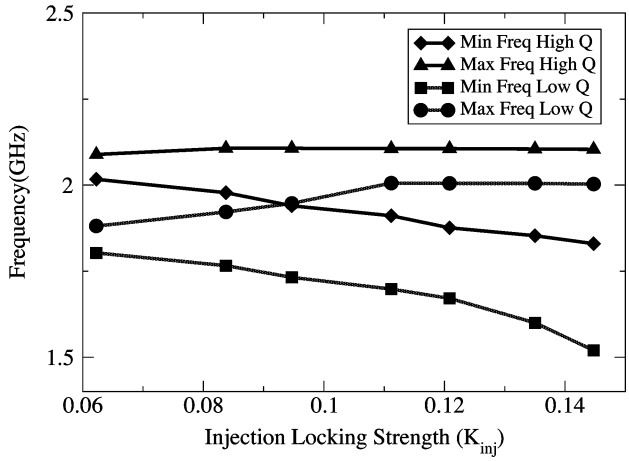

Fig. 18. Locking range of the DDO network as a function of injection strength, Kinj, for both the mininum (total clock capacitance of 92 pF) and maximum (total clock capacitance of 112 pF) available tunable capacitance for the grip.

increases sharply as clock amplitude reaches saturation. As a result, even a small reduction in clock amplitude can conserve significant power.

Injection-Lock Range: The operating frequency of the resonant clock network can be tuned by varying the tunable load capacitances described in Section II and through pulling from the injection-lock source, which entrains the clock system to the desired external frequency. The injection-lock range is governed by [13], [16]

$$\Delta f \cong \frac{f_0}{Q} K_c \tag{10}$$

where f_0 is the natural frequency and Q is the quality factor of the resonant tank. Adding capacitance reduces the Q of the tank, increasing the injection-lock range. At the highest injection-lock strength and utilizing all the available tunable capacitance, the resonant clock frequency can be varied from 1.5 GHz to 2.1 GHz. Fig. 18 shows the locking range of the DDO network as a function of injection strength, K_{inj}, for both the minimum (total clock capacitance of 92 pF) and maximum (total clock capacitance of 112 pF) available tunable capacitance for the grid. The injection strength values used in Figs. 16 and 18 are simulated (from seven selectable levels on the test chip).

Power: When operating at 2 GHz, the prototype clock network (driving a total clock capacitance of 100 pF, as seen by the finals stage of injection buffers) consumes an average power of 500 mW (5.4 mW/pF)–290 mW from the gain elements, 70 mW from the last-stage injection-lock buffers (at $K_c = 0.08$), 70 mW in the rest of the injection-lock tree (including the DCDLs), and 70 mW in the remainder of the AAC, deskewing circuitry, phase detectors, scannable configuration registers, and open-drain drivers used for measurement. For a conventional differentially tree-driven grid, more than 1 W of power, would be required for the same leaf level clock load (according to circuit simulation). This power includes CV^2f power of the clock load, the clock tree and short-circuit switching currents.

VI. Conclusion

In this paper, we have described the design of an adaptive distributed resonant clocking network architecture that increases system scalability beyond what is possible with current resonant-based techniques. A digital deskewing control loop varies the delay of the injection signal. As a result, any jitter introduced by the required delay lines is attenuated by the low-pass nature of the injection-locking process. The power management takes the form of automatic amplitude control, guaranteeing minimum energy for full-rail clock amplitudes while allowing for the possibility for low-swing operation.

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] 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.
[3] 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.
[4] S. C. Chan, K. L. Shepard, and P. J. Restle, "Distributed differential oscillators for global clock networks," *IEEE J. Solid-State Circuits*, vol. 41, no. 9, pp. 2083–2094, Sep. 2006.
[5] Z. Xu and K. L. Shepard, "Low-jitter active deskewing through injection-locked resonant clocking," in *Proc. IEEE Int. Custom Integrated Circuits Conf. (CICC)*, 2007, pp. 9–12.
[6] 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.
[7] G. A. Pratt and J. Nguyen, "Distributed synchronous clocking," *IEEE Trans. Parallel Distrib. Syst.*, vol. 6, no. 3, pp. 314–328, Mar. 1995.
[8] M. Hansson, B. Mesgarzadeh, and A. Alvandpour, "A 1.56 GHz on-chip resonant clocking in 130 nm CMOS," in *Proc. IEEE Int. Custom Integrated Circuits Conf. (CICC)*, 2003, pp. 241–244.
[9] B. M. *et al.*, "Jitter characteristic in charge recovery resonant clock distribuation," *IEEE J. Solid-State Circuits*, vol. 42, no. 7, pp. 1618–1625, Jul. 2007.
[10] W. J. Dally and J. W. Poulton, *Digital Systems Engineering*. Cambridge, U.K.: Cambridge Univ. Press, 1998, ch. Synchronization.
[11] E. W. Weisstein, From MathWorld—A Wolfram Web Resource. [Online]. Available: http://mathworld.wolfram.com/UniformSumDistribution.html
[12] L. Zhang, B. Ciftcioglu, M. Huang, and H. Wu, "Injection-locked clocking: A new GHz clock distribution scheme," in *Proc. IEEE Int. Custom Integrated Circuits Conf. (CICC)*, 2006, pp. 785–788.
[13] R. Adler, "A study of locking phenomena in oscillators," *Proc. IEEE*, vol. 61, no. 10, pp. 1380–1385, Oct. 1973.
[14] E. Lee, W. J. Dally, T. Greer, H.-T. Ng, R. Farjad-Rad, J. Poulton, and R. Senthinathan, "Jitter transfer characteristics of delay-locked loops-theories and design techniques," *IEEE J. Solid-State Circuits*, vol. 38, no. 4, pp. 614–621, Apr. 2003.

[15] 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.
[16] B. Razavi, "A study of injection locking and pulling in oscillators," *IEEE J. Solid-State Circuits*, vol. 39, no. 9, pp. 1415–1424, Sep. 2004.
[17] B. Mesgarzadeh and A. Alvandpour, "First-harmonic injection-locking ring oscillators," in *Proc. IEEE Int. Custom Integrated Circuits Conf. (CICC)*, 2006, pp. 733–736.
[18] W. J. Dally and J. W. Poulton, *Digital Systems Engineering*. Cambridge, U.K.: Cambridge Univ. Press, 1998, ch. Noise in Digital Systems.
[19] G. J. V. Rootselaar and B. Vermeulen, "Silicon debug: Scan chains alone are not enough," in *Proc. Int. Test Conf. (ITC)*, 1999, pp. 892–902.
[20] 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.
[21] K. A. Jenkins, A. P. Jose, and D. F. Heidel, "An on-chip jitter measurement circuit with sub-picosecond resolution," in *Proc. European Solid-State Circuits Conf. (ESSCIRC)*, 2005, pp. 157–160.

Zheng Xu (M'07) received the Bachelor degree in electrical engineering from Cooper Union for Advancement of Art and Sciences in 2002, and the M.S. and Ph.D. degrees in electrical engineering from Columbia University, New York, in 2004 and 2007, respectively, doing research focusing on alternative clock distributions and resonant clocking.

In 2007, he joined Advanced Micro Devices at its Boston Design Center designing a high-performance DDR memory interface.

Kenneth L. Shepard (M'91–SM'03–F'08) received the B.S.E. degree from Princeton University, Princeton, NJ, 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 with the VLSI Design Department, 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 with Columbia University, New York, where he is now Professor. He also was 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 mixed-signal design for biological applications.

Dr. Shepard was 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 ISSCC, VLSI Symposium, ICCAD, DAC, ISCAS, ISQED, GLS-VLSI, TAU, and ICCD. He received the Fannie and John Hertz Foundation Doctoral Thesis Prize in 1992, a National Science Foundation CAREER Award in 1998, and 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.