

Uniform-Phase Uniform-Amplitude Resonant-Load Global Clock Distributions

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

Abstract—This paper presents a new approach to global clock distribution in which tree-driven grids are augmented with on-chip spiral inductors to resonate the clock capacitance. In this scheme, the energy of the fundamental frequency resonates between electric and magnetic forms, with the reduced admittance of the clock network allowing for significantly lower gain requirements in the buffering network. The substantial improvements in jitter and power resulting from this approach are presented using measurement results from two test chips, one fabricated in a 90-nm and the other in a 0.18- μm CMOS technology.

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

I. INTRODUCTION

CLOCKING large microprocessors with a single high-frequency global clock is becoming an increasingly difficult task. Spatial variation in clock arrival time (skew) and compression or expansion of the clock period from nominal (jitter) can limit the maximum operating frequency of a processor or cause functional failure. A 10-GHz global clock distribution in the 75-nm process node can be expected to have a latency of four to five cycles and a gain (as measured by the ratio of total clock capacitance to the capacitance seen at the output of the PLL) of 10^5 or more. A skew and jitter budget of 10% of cycle time translates into a delay variation of no more than 2% in the presence of process, voltage, and temperature variation, a realistically unachievable target (2% of five cycles of latency is 10% of cycle time).

Most global clock distributions today take the form of tree-driven grids. Clock trees are efficient in their use of wiring resources and in minimizing wiring capacitance, while clock grids provide phase averaging and reduce skew [1]. Unfortunately, because of the many levels of buffering required, tree-driven grid global clock distributions are having difficulty meeting skew and jitter requirements as clock frequencies approach 10 GHz because of power-supply noise and intrachip process and temperature variations. Although techniques such as power-supply filtering and deskewing [2] have been effective in mitigating

clock edge uncertainty, continued scaling results in increased clock latency (as measured by number of cycles) and increased clock loading. Both these trends worsen clock skew and jitter, and make global clock power consumption a growing concern.

In the search for alternatives, it has been long recognized that inductive reactance can be used to “cancel” capacitive reactance around a given resonance frequency [3]. In *LC* oscillators, this results in lower power and improved phase stability (skew and jitter). Resonant techniques have been applied to global clock distributions by exploiting the inductance of transmission lines [4]–[6]. Standing-wave, or salphasic, clock distributions have been proposed at both the board level [4] and at the chip-level [6]. These distributions are low-skew and low-jitter, but produce a clock amplitude which varies spatially across the network. Traveling-wave clock distributions [5] use coupled transmission line rings to generate low-skew and low-jitter clocks, but must contend with nonuniform phase across the distribution. Both traveling wave and standing wave clocks are implemented as distributed oscillators, with gain elements uniformly spaced throughout the clock network to overcome losses. Other traveling-wave distributed oscillators have been studied [7], [8], but these have not been explicitly exploited as a distribution. The scale, topology, and resonant frequency of these distributions exploiting transmission-line inductance are inherently linked, which leads to more complex scalability concerns. While the use of *LC*-generated clocks for energy recovery has been recognized [9]–[11], the true scope of power savings possible with resonant clock distributions has yet to be fully studied.

In this work, we design a uniform-phase, uniform-amplitude resonant-load global clock distribution that provides the same straightforward design scalability (in clock frequency and area of the distribution) enjoyed by traditional tree-driven grids. In our approach, a tree-driven grid is rendered resonant with a set of discrete on-chip spiral inductors distributed throughout the clock network. The large clock capacitance then resonates with this inductance, dramatically reducing the gain required in the distribution. This gain reduction saves power, reduces skew through a reduction in clock latency, and reduces power-supply-noise-induced jitter. In addition, less clock buffering can reduce on-chip power supply noise, while the bandpass filtering characteristic associated with the resonance can further reduce skew and jitter.

This paper is organized as follows. Section II describes the details of the resonant-load global clock distribution topology. In Section III, the two test chips used to prototype the proposed resonant clocking scheme are described. One chip was designed in a 90-nm 1.0-V ten-level Cu CMOS technology [12],

Manuscript received April 15, 2004; revised July 30, 2004. The work at Columbia University was supported in part by the MARCO/DARPA C2S2 Center, the National Science Foundation under Contract CCR-00-86007, the SRC, and gifts from IBM and Intel.

S. C. Chan and K. L. Shepard are with the Columbia Integrated Systems Laboratory, Department of Electrical Engineering, Columbia University, New York, NY 10027 USA (e-mail: schan@cisl.columbia.edu; shepard@cisl.columbia.edu).

P. J. Restle is with the IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA (e-mail: restle@us.ibm.com).

Digital Object Identifier 10.1109/JSSC.2004.838005




Fig. 1. Global clock distribution with a resonant load. Eight clock sectors, which form the basic building block of the distribution, are shown.

while the other chip was designed in a $0.18\text{-}\mu\text{m}$ 1.8-V six-level Al mixed-signal CMOS technology [13]. Measurement results from these chips are presented in Section IV. Comments on scaling these distributions to different frequencies and different capacitive loads are presented in Section V. Section VI concludes the paper.

II. RESONANT LOAD GLOBAL CLOCK DISTRIBUTION

Fig. 1 shows a resonant-load global clock distribution. The clock is distributed from a single synchronous source and is buffered through a tuned, balanced, global H-tree. The tree then drives a set of clock sectors, a basic unit of the distribution driven by the lowest buffer level of the global clock tree. For simplicity, only eight clock sectors are shown in Fig. 1, while on a real microprocessor, there might be several dozen clock sectors.

The sector buffer associated with each clock sector provides the gain needed to drive a local H-tree, a globally connected clock grid, and local clock buffers, as shown in Fig. 2. In our single-ended topology, four spiral inductors have one end attached to the clock tree and the other end attached to a large decoupling capacitance which establishes a mid-rail dc voltage around which the clock network swings. By attaching the inductor into the tree (as shown in Fig. 2) a small “treelet” distributes the current flowing into and out of the inductor, reducing skew in the distribution. The MOS decoupling capacitors are positioned adjacent to the spiral inductors. Local clock buffers tap into the global clock grid within a sector and provide additional gain needed to drive the latches and gates in the design.

Although a distributed model is needed for detailed understanding of the clock sector, a simple lumped circuit representation of the resonant distribution as seen by the clock sector buffers can be used to highlight some important features. In Fig. 3, C_{clock} is the capacitive load of the clock network, C_{decap} is the decoupling capacitance, L is the spiral inductance, while R_{eddy} , R_{ind} , and R_{clock} model losses in the network: R_{ind} models the series resistance in the spiral inductor wires, R_{clock} models losses in the clock wires and losses due to displacement current flow into neighboring conductors and the substrate, R_{eddy} is the loss from eddy currents due to induced EMFs in the surrounding metal lines and the substrate. It is worth noting that the substrate is not likely to have a significant




Fig. 2. Components and topology of a resonant clock sector.




Fig. 3. Simple lumped circuit model of the resonant clock sector.

effect on either R_{clock} or R_{eddy} in the metal-dense environment of a microprocessor. The driving point admittance of the simplified lumped model is given by

$$Y(s) = s(C_{\text{decap}} + C_{\text{clock}}) \frac{1 + sR_s C_p + s^2 L C_p}{1 + sR_s C_{\text{decap}} + s^2 L C_{\text{decap}}}$$

where $C_p = C_{\text{decap}} C_{\text{clock}} / (C_{\text{decap}} + C_{\text{clock}})$ and R_s is the total effective series resistance of the inductor, which includes wire loss and eddy current loss. C_{decap} must be chosen large enough to ensure that the poles associated with the resonance at approximately $f_{\text{decap}} \cong 1/2\pi\sqrt{LC_{\text{decap}}}$ do not interfere with the zeros at the desired clock resonance frequency, $f_{\text{clock}} \cong 1/2\pi\sqrt{LC_p} \cong 1/2\pi\sqrt{LC_{\text{clock}}}$. At f_{clock} , the capacitive reactance of clock load is cancelled by the inductive reactance of the spirals. The amount of ripple on node B in Fig. 3 depends on the ratio $f_{\text{clock}}/f_{\text{decap}}$, which we choose to be three in our design. This corresponds to decoupling capacitance that is approximately ten times greater than the amount of clock capacitance. Such a quantity of decoupling capacitance is no more than what is typically required on the power-ground network in

the case of nonresonant clocking to prevent more than a 10% collapse of the supply rails during clock switching.

The driver Q -loads the network but the benefits of resonant clocking come from reducing the strength of this driver (and its Q -loading) to be just strong enough to sustain oscillation.

It is expected that higher-order resonances exist in the clock network, characterized, for example, by standing wave patterns in the clock grid or clock tree or the self-resonance frequency of the spiral inductors. The frequencies of these resonances must be kept sufficiently high so as not to interfere with the engineered f_{clock} uniform-phase eigenmode of the network. This can be achieved by ensuring that the clock grid has low inductance and is driven from a sufficient number of points and “suspended” by a sufficient number of inductors. This will be discussed in more detail in Section V.

Microprocessor clock distributions must at times be able to operate in a low-frequency test mode in which the advantages of resonant clocking are not important. At frequencies significantly lower than f_{clock} , the resonant clock network presents more loading (higher admittance) than C_{clock} would present alone. If the resonant sector clock buffer, which would be sized much smaller than a comparable nonresonant sector clock buffer, is too weak to operate at the chosen test frequency, larger buffers may have to be switched in to help during test mode.

III. TWO TEST CHIP DESIGNS

In order to quantify the power savings and jitter reduction possible with resonant-load global clocking, two test chips were designed. The first chip, fabricated in an aggressive 90-nm CMOS technology, has a target resonance of 3.7 GHz. However, because this chip was part of a much larger experiment, there is limited controllability and observability, and accurately measuring jitter or viewing clock waveforms is not possible. Hence, a second chip was designed and fabricated in a 0.18- μm CMOS technology. Because of its more extensive test and measurement infrastructure, this chip is able to confirm not only power savings, but jitter reduction as well. On both chips, we designed a resonant clock sector, as shown in Fig. 2, and a nonresonant “control” clock sector, allowing us to quantify the power savings and jitter reduction that the proposed scheme enables, without the need to design an entire global clock distribution. A clock sector occupies an area of $2500\ \mu\text{m} \times 2500\ \mu\text{m}$ in both chips.

In the 0.18- μm test chip, the clock tree and grid occupy the top two Al metal layers, M6 and M5. In order to minimize resistance and inductance in the clock lines, each wire in the distribution is shielded, and each line is split into multiple fingers. Fig. 4 shows the shielding and fingering of the clock tree wires on M6 and M5.¹ We use 16- μm -wide segments, spaced 4 μm apart. For the M6 and M5 clock grid wires, a single 16- μm -wide segment surrounded by two 8- μm ground shields spaced 4 μm apart is




Fig. 4. Fingering and shielding of clock tree wires on M6 and M5 on the 0.18- μm test chip.

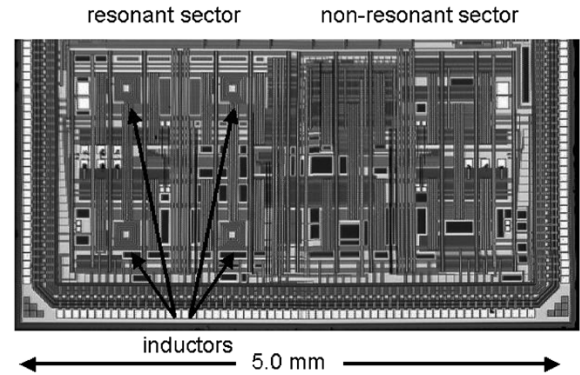


Fig. 5. Die photo of the 0.18- μm test chip. The spiral inductors in the resonant clock sector are visible.

used.² The capacitive loading of the clock wires in the sector is approximately 9.5 pF, while the drain capacitance of the tunable clock sector buffer contributes approximately 10 pF, giving a total $C_{\text{clock}} = 19.5$ pF. Each of the four decoupling capacitors contributes $C_{\text{decap}} = 60$ pF, and each spiral inductor is approximately $L = 6$ nH. Since the four decoupling capacitors and four spiral inductors act in parallel, the total effective decoupling capacitance is 240 pF and the total effective inductance is 1.5 nH. This results in an f_{clock} resonance of approximately 930 MHz. Each of the four thin-oxide inversion-mode decoupling capacitors consume about $150\ \mu\text{m} \times 150\ \mu\text{m}$ of chip area, while each 5.5 turn spiral inductor has a diameter of $250\ \mu\text{m}$ and is routed on M6 using 15- μm -wide Al traces, spaced 1.5 μm apart. Fig. 5 shows a die photo of the chip, including the two clock sectors (resonant and nonresonant) and the spiral inductors of the resonant sector.

In the 90-nm test chip, $C_{\text{clock}} = 7.5$ pF, $C_{\text{decap}} = 20$ pF, and $L = 1$ nH. This results in an effective decoupling capacitance of 80 pF, an effective inductance of 0.25 nH, and an f_{clock} resonance of 3.7 GHz. Because of the more advanced process technology, the decoupling capacitors on this chip were implemented as thick-oxide accumulation-mode capacitors, each consuming about $80\ \mu\text{m} \times 80\ \mu\text{m}$ of chip area. The three-turn spiral inductors have a diameter of $90\ \mu\text{m}$, and use 6- μm -wide Cu traces, spaced 13 μm apart. An unusually large turn spacing was used here because a very dense power grid was adjacent to the each inductor wire segment. Because of layout image requirements, the vertical segments of the inductor are on M10, while the horizontal segments are on M9.

In these designs, the spiral inductors exist in the metal-rich environment of a microprocessor, quite different from that of

¹We chose grounded nets for the shields because ground was available throughout the chip, while for purposes of monitoring power in this testchip, the V_{dd} network was fractured into domains. In general, shielding with both ground and V_{dd} nets would be preferred since it creates balanced loading on both networks.

²Since the loss in the resonant clock network is dominated by the loss in the spiral inductors, the actual widths of the clock tree and grid wires are not critical, so long as the network is sufficiently gridded so that resistive loss is reduced. Shielding insures that the parasitic inductance in the clock wires is well controlled, and is typically an order of magnitude less than the spiral inductance added.




Fig. 6. Spiral inductors on the 90-nm test chip. Vias in the power-ground network near the inductor are dropped and small cuts are made along the two axis shown in order to reduce eddy currents in power-ground network beneath and adjacent to the inductor.




Fig. 7. Worst case simulated capacitive and inductive coupling noise on a signal line which encircles one of the spiral inductors on the 90-nm test chip. (a) Signal line response due to a step input on the clock. (b) Signal line response due to continuous full-rail switching of the clock at 3.7 GHz.

typical RF applications. Careful attention must be paid to limit eddy current losses in neighboring wires; this is important both to prevent Q degradation and to prevent inductive noise in the power-ground distribution and in neighboring signal lines. Because the spiral inductors are much larger than the power grid pitch, simulations using a full-wave PEEC time-domain simulator [14] show that most of the potential deleterious couplings are to the underlying power grid adjacent and beneath the inductors. To impede eddy current formation, the vias in the grid are dropped and small cuts are made in the wires down to M2, analogous to the ground plane laminations used for spiral inductors in RF circuits [15]. Fig. 6 shows the inductor layout on the 90-nm test chip, along with the cuts in the nearby power grid.

An important consideration in incorporating large spiral inductors is that they do not represent a significant area overhead; that is, active circuitry can use the area under the inductors. To consider the magnitude of the induced voltages on underlying

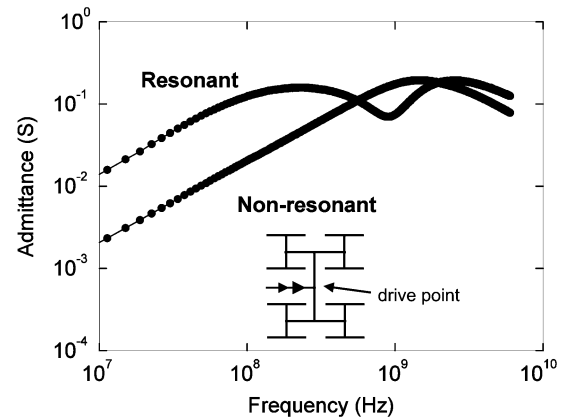


Fig. 8. Measured driving point admittance for the resonant and nonresonant clock sectors.

interconnect, Fig. 7 shows the noise, from PEEC simulation, on a hypothetical signal line which encircles one of the spiral inductors in the 90-nm test chip. A minimum width, $350\ \mu\text{m}$ unshielded trace on M5 is loaded with $10\ \text{fF}$ at the far end, and at the near end, a driver with a linearized resistance of $50\ \Omega$ is used to tie the line to V_{dd} . A 44-mA peak current is flowing through each inductor, producing a peak magnetic flux of approximately $4.4 \times 10^{-11}\ \text{Wb}$. The peak noise induced is less than $10\ \text{mV}$ for a step input on the clock, and less than $30\ \text{mV}$ for continuous full-rail switching of the clock at $3.7\ \text{GHz}$. Simulations (not shown) also show that capacitive and inductor noise coupling onto the power-ground network is small as well (less than $5\ \text{mV}$) because of the extremely low impedance of these nets.

IV. MEASUREMENT RESULTS

A. $0.18\text{-}\mu\text{m}$ Test Chip

1) *Driving Point Admittance*: The drive points of the resonant and nonresonant clock sector buffers on the $0.18\text{-}\mu\text{m}$ test chip include ground-signal-ground probe pads that allow for network analyzer sweeps to be made to obtain S-parameters. Fig. 8 shows the magnitude of the measured driving point admittance as a function of frequency around zero bias. At low frequencies, the resonant sector has a larger admittance than the nonresonant sector because of the additional loading from the large decoupling capacitors and the poles of the driving point admittance associated with f_{decap} . At higher frequencies however, the inductors shield the decoupling capacitors and the engineered f_{clock} parallel resonance between the spirals inductors and the clock capacitance lowers the admittance in the resonant sector to be below that of the nonresonant sector over the frequency range from $600\ \text{MHz}$ to $1.2\ \text{GHz}$. The Q of this resonance is about 1.6 . While a higher Q would be advantageous for reducing power and improving skew and jitter, the low Q network provides the benefit of a wide operating frequency range without tuning. The measured resonant frequency is $970\ \text{MHz}$ at zero bias. If the dc bias is increased to $1.8\ \text{V}$, the resonance frequency decreases slightly to approximately $859\ \text{MHz}$ because of an increase in the nonlinear drain capacitance (which constitutes approximately half the clock load) with bias.




Fig. 9. Noise generators are switched on and off at 100 MHz to produce power-supply noise near the sector clock buffers on the 0.18- μm test chip.

2) *Jitter*: Since power supply noise in the clock buffers is the main contributor to clock jitter, we are most interested in quantifying the jitter in the resonant and nonresonant clock sectors in the presence of power supply noise. Fig. 9 shows the noise generators used to induce power supply noise on the test chip, simple MOS switches which periodically short V_{dd} to ground near the sector clock buffers. Two different power-supply switch sizes are used which correspond to measured noise amplitudes of 200 and 365 mV. The noise generators are switched on and off at 100 MHz to introduce noise at one-tenth the nominal clock frequency, typical of the most problematic low-frequency resonances determined by package inductance and die capacitance [16]. Fig. 9 shows the measured power-supply noise between the 1.8 V V_{dd} supply pin and local ground.

To measure the jitter, open-drain drivers attached to a corner of the clock grid buffer the clock off chip to the 50- Ω input of a sampling scope. Figs. 10 and 11 show the measured characteristics at a 900-MHz operating frequency for both the resonant and nonresonant clock sectors. Jitter is plotted as a function of sector buffer strength (the x -axis indicates the size of the nFET in the final stage inverter of the sector buffer: 44, 88, 175, 350, 700, or 1400 μm) for three different conditions: no added supply noise, 200 mV of added supply noise, and 365 mV of added supply noise. The figures show better than 60% reduction in jitter in the resonant sector at all buffer strengths and at all noise levels.

In the nonresonant clock sector, the two smallest sector buffers, the 44- μm and 88- μm buffers, are too weak to drive the clock network at 900 MHz and as a result, no jitter is recorded for these buffers. In fact, as we will show, only the strongest buffer is able to deliver a full-rail clock. This is in contrast to the resonant clock sector, where only the weakest buffer is unable to drive the network full-rail. Fig. 10 shows that as the sector buffers are strengthened in the nonresonant clock sector, jitter is reduced. This occurs as the clock edges are sharpened and rendered less sensitive to power supply noise. If we were able

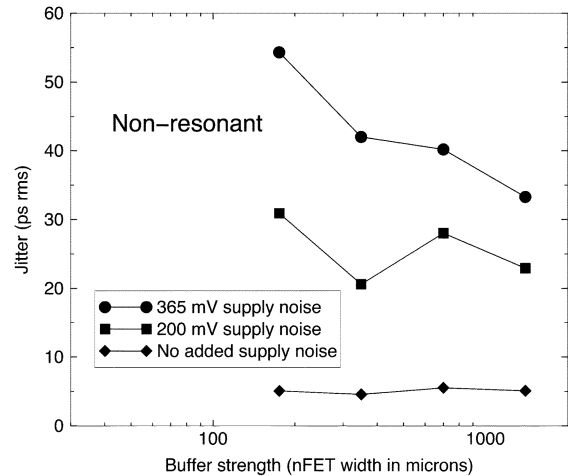


Fig. 10. Jitter as a function of sector buffer strength as measured in the nonresonant clock sector of the 0.18- μm test chip. The jitter is shown for the case of no added supply noise and for two different amounts of added power supply noise.

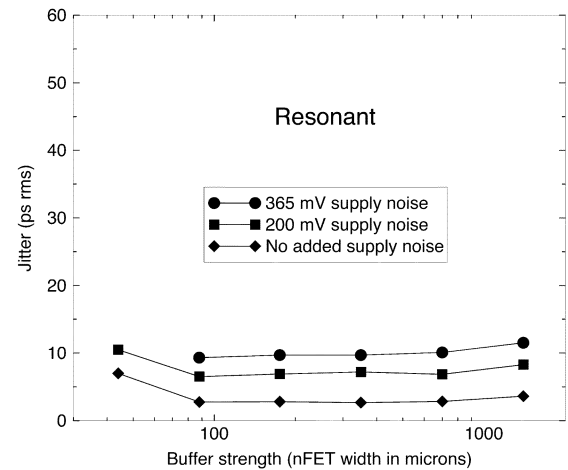


Fig. 11. Jitter as a function of sector buffer strength as measured in the resonant clock sector of the 0.18- μm test chip. The jitter is shown for the case of no added supply noise and for two different amounts of added power supply noise.

to increase the buffer strength beyond 1400 μm , further jitter reduction would have been possible. Eventually, however, the jitter would increase with increasing buffer strength as larger buffers are more able to convert power supply noise into jitter.

Fig. 11 shows the jitter performance of the resonant clock sector. For the weakest buffer (44 μm) and the maximum added power supply noise (365 mV), the clock waveform does not have the integrity to perform a jitter characterization. Overall, however, the resonant sector has a much better ability to reject jitter from power-supply noise. It should be noted that the resonant clock sector shows similar jitter reduction when the 44- μm buffer is strengthened to 88 μm . Further increases in sector buffer strength does not improve jitter significantly because the resonant load is already filtering much of the jitter.

3) *Waveforms*: In addition to the open drain drivers, the test chip includes picoprobe sites to enable probing of the clock grid as well as a buffered clock driven from the grid. Because of the use of passive probes, we are not able to probe the clock grid directly without being invasive. As a result, Fig. 12 only shows




Fig. 12. Measured clock waveforms, as locally buffered from the clock grid on the $0.18\text{-}\mu\text{m}$ test chip. (a) Non-resonant sector. (b) Resonant sector. Two different sector clock buffer sizes are used.

the waveforms as measured from a buffered clock using a passive $1\text{-k}\Omega$ S-G Z_0 -probe (adjusted for the factor-of-20 attenuation of the probe). The “strong driver” curves correspond to the $1400\text{-}\mu\text{m}$ clock sector buffer, while the “weak driver” curves correspond to the $88\text{-}\mu\text{m}$ clock sector buffer. Fig. 12(a) shows the waveforms for the nonresonant clock sector. In the weak driver case, there is insufficient amplitude on the clock grid to buffer out a full-rail clock. Fig. 12(b) shows that for the resonant clock sector, both weak and strong sector clock buffers drive the grid with sufficient amplitude to buffer out a full-rail clock. The weak driver case shows a slight duty-cycle reduction, most likely due to the fact that the clock grid waveform is nearly sinusoidal at this sector buffer strength, and the buffer driving the picoprobe pad has a slightly skewed switch point.

4) *Power*: Fig. 13 shows the measured power dissipation of the resonant and nonresonant clock sectors. Average sector buffer current is plotted as a function of frequency, with each curve in the graph representing a different sector buffer strength.

In the nonresonant clock sector (dashed curves), except for the largest buffer, the sub-linear dependence of the average current on frequency is because the clock network is not being driven full-rail. In addition, the $44\text{-}\mu\text{m}$ and $88\text{-}\mu\text{m}$ buffers do not have sufficient strength to drive the network beyond 400 MHz and 700 MHz , respectively. Fig. 13 also shows that for equivalent sector buffer strengths, the resonant clock sector (solid curves) consumes approximately 20% less power than the nonresonant clock sector. It is worth noting that the Q of the resonant network, 1.6, measured in Fig. 8, is consistent with approximately 20% of the clock energy being recovered each cycle. A more interesting comparison, however, can be made by comparing sector buffer power consumption based not on size, but on ability to drive the clock network full-rail and minimize jitter. Figs. 10–12 show that for the nonresonant clock sector, the $1400\text{-}\mu\text{m}$ buffer is the only buffer which is able to drive the clock network full-rail. It is also the buffer which minimizes jitter. For the resonant clock sector, the $88\text{-}\mu\text{m}$ sector buffer is the buffer which minimizes jitter and is able to drive the clock network full-rail. Comparing the power consumption of the nonresonant $1400\text{-}\mu\text{m}$ buffer and the resonant $88\text{-}\mu\text{m}$ buffer in




Fig. 13. Resonant (solid curves) and nonresonant (dashed curves) sector buffer current as a function of frequency, as measured on the $0.18\text{-}\mu\text{m}$ test chip. Each curve represents a different sector buffer strength, shown as numbers next to the curves.




Fig. 14. Sector buffer currents as a function of frequency on the 90-nm test chip, as measured (solid line) and as simulated (dashed line).

Fig. 13 shows that the power savings achieved by the resonant clocking scheme is more than 77%.

B. 90-nm Test Chip

Fig. 14 shows a plot of sector buffer current as a function of frequency for the resonant and nonresonant sectors on the 90-nm test chip, as measured (solid line) and as simulated (dashed line). At low frequencies, the clock driver in the resonant sector, which is sized to be the same as the driver in the nonresonant sector,³ draws approximately ten times more current than the driver in the nonresonant sector. This is because the resonant driver sees not only the clock capacitance, but also the decoupling capacitors. As f_{clock} is approached, the buffer in the resonant sector begins to draw less current since the clock load’s capacitive reactance is canceled by the spiral’s inductive reactance, and the load of the decoupling capacitors is more effectively shielded by the inductors. The buffer in the nonresonant sector continues to draw linearly more current as frequency is increased up to 3 GHz . A measurable skew in the

³On this chip, there was no ability to vary the strength of the sector buffer.

duty cycle of the test chip results in a reduced voltage swing and sub-linear current increase above 3 GHz. The hardware shows 35% less current in the resonant sector at 4.6 GHz, 15% of this reduction is due to a reduced load capacitance on the resonant sector. This suggests that approximately 20% of the clock power is being recycled, which is comparable to the power savings observed on the 0.18- μm test chip when comparing equivalent buffer strengths. Had we had the ability to reduce the strength of the sector buffer in the resonant sector on this chip, the power savings would be much greater, and would likely approach the 80% observed in the 0.18- μm test chip.

Fig. 14 shows that at frequencies below resonance, the resonant clock network draws more current than the nonresonant clock network because of the additional loading presented by the decoupling capacitors. It is thus important to operate the resonant network at frequencies at or above resonance in order to avoid driving this extra load. In designs with many different frequencies of operation, the target resonance should be at the lowest such frequency. In situations where this is not possible, some mechanism to turn-off the resonance by switching out the spiral inductors would be desirable.

V. SCALING

The approach presented here allows easy scalability to different load capacitance and different operating frequencies. For example, one can scale to higher clock frequencies for a given clock load by the addition of more inductors to the network, reducing the effective L in Fig. 3. Adding more spirals to the grid is preferable to reducing the inductance of each spiral because the addition of more “attach” points help to suppress high order resonance in the clock network and preserve the uniform phase and uniform amplitude across the clock distribution. In fact, by designing the clock tree and grid to be low inductance, we deliberately push other resonances associated with the network to high frequencies so that they do not interfere with the dominant engineered resonance.

The low inductance clock grid also helps to reduce skew due to mismatch in the spiral inductors, similar to the way a low inductance clock grid helps to reduce skew due to clock buffer mismatch in nonresonant clocking. Simulation of the resonant network in the 0.18- μm test chip shows that a 100% variation in the inductance value of one of the four spiral inductors does not degrade the skew by more than 15% across the clock sector. This immunity to inductor mismatch induced skew is due to the low inductance clock grid phase averaging the low Q clock network over the relatively small area of single clock sector.

It is commonly argued that power reduction approaches applied to the global clock distribution will have only a small effect on overall clock power requirements because much of the clock distribution is local, that is, gain stages driven from the global distribution to the latches. Similar arguments can be made about jitter and skew introduced in the local clock distribution. The approach presented here, because of its straightforward scalability to different load capacitance and because of the dramatically reduced gain requirements on the driving network,

allows a general collapsing of the buffering network, both global and local. More capacitance can be transferred up to the global level and fewer levels of local clock buffering can be employed. In this way, the jitter, power, and skew advantages of resonant clocking can be extended closer to the latches. Given that the bulk of the capacitance is in the “leaves” of the tree, the largest power advantage will come in extending resonance down to the latches. This will require more detailed understanding of the skew and jitter implications of temporal and spatial variations in clock capacitance, more consideration of how clock gating and other “local” clock manipulations will be accomplished, and how latches will perform with the more “sinusoidal” clocks characteristic of resonant networks.

VI. CONCLUSION

In this paper, we have presented a new approach to global clock distribution in which traditional tree-driven grids are augmented with on-chip inductors to resonate the clock capacitance at the fundamental frequency of the clock node. The resulting distribution delivers a uniform-phase, uniform-amplitude clock and allows straightforward scalability to different clock load capacitances and different operating frequencies. With the low Q s typically achieved, approximately 20% of the energy of the fundamental is recovered and reused each cycle. Further power reductions come about because of the ability to significantly scale back the required buffering in the global clock distribution, allowing total power saving approaching 80%. This decrease in buffering, along with the natural band-pass characteristics of the resonant network, results in over 60% reduction in jitter as well.

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] N. A. Kurd, J. S. Barkatullah, R. O. Dizon, T. D. Fletcher, and P. D. Madland, “A multigigahertz clocking scheme for the Pentium 4 microprocessor,” *IEEE J. Solid-State Circuits*, vol. 36, no. 11, pp. 1647–1653, Nov. 2001.
- [3] M. I. Pupin, “Electrical Transmission by Resonance Circuits,” U.S. Patent no. 640 516, 1900.
- [4] V. L. Chi, “Salphasic distribution of clock signals for synchronous systems,” *IEEE Trans. Comput.*, vol. 43, no. 5, pp. 597–602, May 1994.
- [5] 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.
- [6] 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.
- [7] B. Kleveland *et al.*, “Monolithic CMOS distributed amplifier and oscillator,” in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, vol. 36, Feb. 1999, pp. 70–71.
- [8] H. Wu and A. Hajimiri, “Silicon-based distributed voltage-controlled oscillators,” *IEEE J. Solid-State Circuits*, vol. 36, no. 3, pp. 493–502, Mar. 2001.
- [9] W. C. Athas, L. J. Svensson, and N. Tzartanis, “A resonant signal driver for two-phase, almost-nonoverlapping clocks,” in *Int. Symp. Circuits Syst.*, 1996, pp. 129–132.
- [10] T. Ye and K. Roy, “QSERL: Quasistatic energy recovery logic,” *IEEE J. Solid-State Circuits*, vol. 36, no. 2, pp. 239–248, Feb. 2001.
- [11] A. J. Drake, K. J. Nowka, T. Y. Nguyen, J. L. Burns, and R. B. Brown, “Resonant clocking using distributed parasitic capacitance,” in *Proc. Int. Custom Integrated Circuits Conf.*, 2003, pp. 647–650.

- [12] S. C. Chan, P. J. Restle, K. L. Shepard, N. K. James, and R. L. Franch, "A 4.6 GHz resonant global clock distribution network," in *IEEE Int. Solid-State Circuits Conf. Dig. Tech. Papers*, 2004, pp. 342–343.
- [13] S. C. Chan, K. L. Shepard, and P. J. Restle, "Design of resonant global clock distributions," in *Proc. Int. Conf. Comput. Design*, 2003, pp. 238–243.
- [14] P. J. Restle, A. E. Ruehli, S. G. Walker, and G. Papadopoulos, "Full-wave PEEC time-domain method for the modeling of on-chip interconnects," *IEEE Trans. Computer Aided Design*, vol. 20, no. 7, pp. 877–887, Jul. 2001.
- [15] C. P. Yue and S. S. Wong, "On-chip spiral inductors with patterned ground shields for Si-based RF IC's," *IEEE J. Solid-State Circuits*, vol. 33, no. 5, pp. 743–752, May 1998.
- [16] 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.



Steven C. Chan received the B.S. degree with high honors and the M.S. degree in electrical engineering and computer science from the University of California at Berkeley in 1996 and 1998, respectively. His research at Berkeley focused on design methodologies for complex systems.

From 1997 to 2001, he was with CadMOS Design Technology, San Jose, CA, where he helped develop the first commercial transistor-level static noise analysis tool. Since 2001, he has been a doctoral student at Columbia University, New York,

working on clock distribution for high-performance microprocessors. His research at Columbia was funded in part by a fellowship from IBM and a fellowship from the SRC. Mr. Chan has held summer internship positions at IBM's T. J. Watson Research Center, where he designed and tested global clock distributions, and at Intel, Santa Clara, CA, where he developed a transistor-level power consumption estimator for the IA-64 microprocessor. His research interests include clock distribution, mixed-signal circuit design, on-chip test and measurement, and interconnect analysis.



Kenneth L. Shepard 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 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 an 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 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 received the B.A. degree in physics from Oberlin College, Oberlin, OH, in 1979, and the Ph.D. degree in physics from the University of Illinois at Urbana in 1986.

He then joined the IBM T. J. Watson Research Center as a Research Staff Member, where he initially worked on CMOS parametric test and modeling, CMOS oxide-trap noise, package testing, and DRAM variable retention time. Since 1993, he has concentrated on tools and designs for VLSI clock distribution networks contributing to 12 IBM microprocessors, as well as high-performance ASIC designs. He holds six patents, has written 21 papers, and has given keynotes, invited talks, and tutorials on clock distribution, high-frequency on-chip interconnects, and technical visualizations in VLSI design.

Dr. Restle has received IBM awards for the Mainframe G4, G5, and G6 microprocessors, for the Power4 and Power5 microprocessors, for the PowerPC 970 used in the Apple G5 machine, as well as an IBM corporate award for VLSI clock distribution design and methodology.