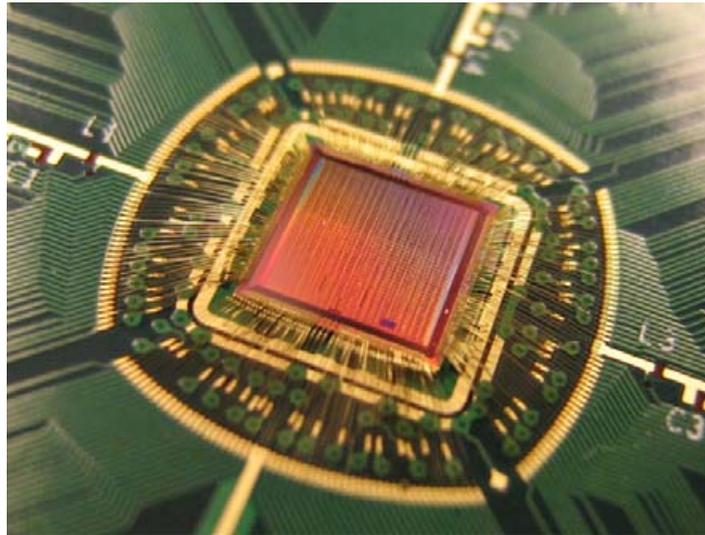




## Addressing the Power-Performance IC Design Conundrum



*A Novel Clock Design Technique to  
Reduce Power and Increase Performance*

June 1, 2012

## The Need for Speed

We're witnessing the emergence of a whole new generation of adults who expect to always have immediate and ever-present access to on-line information, interpersonal communications, social networks, and financial transactions. These are the teenagers who woke on Christmas morning to find a new smartphone under their tree, or as they prepared to return to school had a new tablet PC in their backpacks, or who on their way to pick up a first time date punched the address into a street navigator instead of unfolding an unruly map. The generation that is entering the workforce today and who will be the future leaders of our companies, communities, and governments operate in a world of endless texting, emails, apps, videos, and postings.

The enormous spending capability of this new generation of mobile users warrants intense scrutiny of their needs by semiconductor vendors who have scarily dwindled in numbers during the last global recession. Winning a socket in an Apple iPad or Samsung smartphone can be a game changer for a semiconductor startup or a Fortune 500 vendor due to the tremendous and rapid ramp up of IC production and revenue.

So what are their needs? Quite simply – the best performance and battery life possible.



Mobile Internet devices (MIDs) are however only as useful as the networking infrastructure that supports them. Equally important to semiconductor vendors are the needs of the data center vendors that are currently investing in upgrading all of their hardware to accommodate not only the large global increase in the number of mobile users, but also the large demand in throughput brought on by the video revolution. Vendors are racing to increase data center efficiency and throughput by implementing 10GBase-T technology that enables increased bandwidth and speed over the existing standard, structured copper infrastructure.

## The Need to Be Green

But with the soaring cost of energy, these same performance-hungry data center vendors are desperately looking to decrease energy consumption and lower their environmental/cooling costs. In fact, one semiconductor vendor has even trademarked "GrEEEN" that stands for "Energy Efficient Ethernet" in the same style as

environmentalists promote energy saving appliances and buildings via the U.S. government's EnergyStar program ([www.energystar.org](http://www.energystar.org)).



PLX Technology/Teranetics 10GBase-T PHYs

Consider the “power consumption cost” of a simple Google search made from a MID. A fair amount of processing has to be performed locally on the MID and then the MID must transmit the request to its cellular network, which consumes a small portion of the MID's battery life, which is not really a big deal. But what about the power consumption cost of the data centers that need to respond to the search request from a mobile user? First the request has to be processed by multiple servers as it skips across the Internet until it finally reaches the Google data center. An average Google search is reported to require ~ 0.3 watt-hours, about the same amount of power that it takes for a 100 watt light bulb to be lit for 10 seconds. Then the results of the search must be communicated back to the MID, bouncing back across multiple servers on the Internet. I doubt most Internet users have even contemplated how much energy is being spent every time they perform a Google search – and I'm quite sure Google doesn't want them to know!

“Google operates huge data centres around the world that consume a great deal of power,” said Alex Wissner-Gross, a Harvard University physicist whose research on the environmental impact of computing is due out soon. “A Google search has a definite environmental impact.”

***Revealed: the environmental impact of Google searches***

**The Sunday Times**

**[http://technology.timesonline.co.uk/tol/news/tech\\_and\\_web/article5489134.ece](http://technology.timesonline.co.uk/tol/news/tech_and_web/article5489134.ece)**

**January 11, 2009**

So the world of computing must change. Not only in reducing power consumption, but also in how massive amounts of computing power are provided to users. Thankfully, a new reality is emerging for organizations of every size from every part of the planet. It's called the cloud—a profound evolution of IT with revolutionary implications for business and society. Sometime soon, users won't need to purchase and maintain local

computers. Computing will happen magically via the clouds of mega-data centers that are being built today all across our planet. Users will simply have a terminal or MID that provides them instant, convenient, and maintenance-free access to these mega-data centers via the Internet. The adoption of cloud computing will create hockey-stick like growth for those semiconductor vendors who can successfully provision energy efficient data centers, delivering content and computing to the new generation of consumers with high-throughput, low-cost, and extreme reliability.

Again, their needs are quite simple. High-performance and throughput are absolutely necessary but not completely sufficient. Also critically important to facilitate the adoption of cloud computing in the near future is a significant and meaningful reduction in data center power consumption.



The IBM "Green" Mega-Data center

## All That's Been Done is Still Not Enough

IC design engineers and electronic design automation (EDA) vendors have been diligently addressing the need to reduce power in IC designs for nearly a decade now. Designers have successfully used power reduction techniques like clock and power gating, multi-VT, and multiple voltage domains for many years. EDA vendors have delivered synthesis and power analysis tools that implement a myriad of advanced techniques for reducing power such as logic factorization, don't care optimization, path balancing, technology mapping, state encoding, finite-state machine decomposition, and retiming. But it hasn't been enough.

Systems engineers architecting new hardware products must now have awareness of their software applications behavior to take advantage of opportunities to power-down or slow-down functional blocks when they are not in use. Dynamic frequency scaling (also known as CPU throttling) is a technique whereby the performance of a processor can be automatically adjusted "on-the-fly" to conserve power. MIDs today are smart enough to know that when a user is reading email, for example, the performance

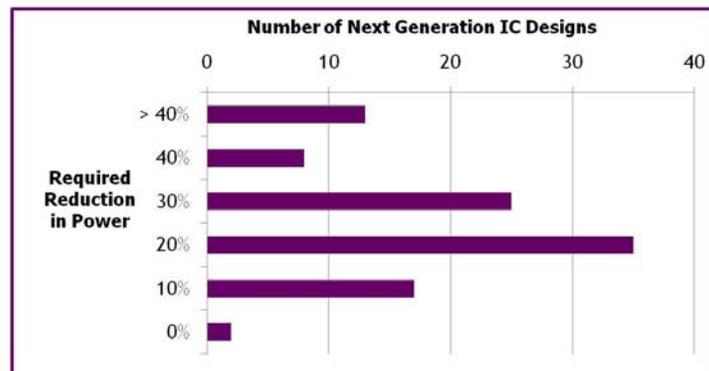
requirements of the processor are much less than when downloading and viewing a video or playing an interactive video game. Dynamic voltage scaling is another power management technique where the voltage supplied to a processor is decreased to reduce power consumption when high-performance is not required. These ways of reducing power consumption based on the performance required by the software application are not trivial to implement but have shown to be very effective by Apple and other MID vendors in extending battery life. But it's still not enough.

**“It’s still possible to save lots of power in devices. It’s just harder to do—and more expensive. It will require new materials, new techniques, and in some cases a new way of thinking.”**

*The Easy Stuff is Done*  
Chip Design Magazine  
January 13, 2011

Semiconductor process engineers have been developing new materials for lowering static power consumption in advanced process nodes where leakage currents are known to increase exponentially. Fully-Depleted Silicon on Insulator [FDSOI] materials have been shown to be very effective in eliminating leakage power and FDSOI is on a path to be implemented in mainstream silicon processing with the adoption of the 22 nm process node. Over two-thirds of vendors developing semiconductors for next generation MIDs have indicated they plan to use FDSOI for their next semiconductor devices. New materials will certainly help reduce power, but it is still not enough.

In a 2010 survey conducted by the Global Semiconductor Alliance entitled “Technology Needs for Mobile Internet Devices”, the critical need to reduce power consumption in next generation ICs for MIDs is clearly apparent. The chart from the survey is shown below and the results are alarming. Over 80% of the respondents indicated at least a 10% reduction in power is required, with over 20% of the vendors needing a whopping 40% reduction. It’s crystal clear – much more work needs to be done to reduce power. And indeed ... it does require new techniques and new ways of thinking.



## RETHINK Clocks

One new way of thinking attacks the "beating heart" of every chip - the clock. Like the heart in a human body, the clock inside an integrated circuit is what makes it tick. Virtually every digital semiconductor relies on a highly accurate, highly stable, cyclical clock to orchestrate all the data movement and processing being done by the various functional blocks on the chip. In today's IC designs, clock frequencies above a GHz are fairly common and that means the clocks inside these chips are oscillating a mind boggling billion times per second.

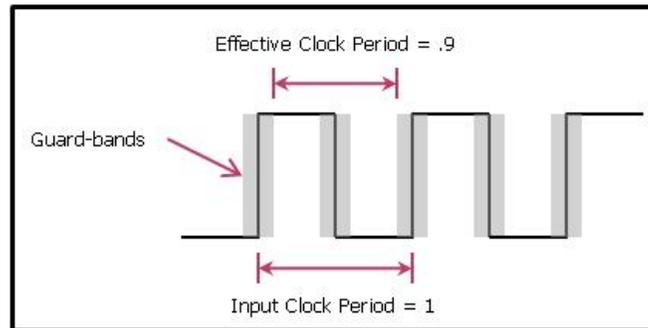
Since the beginning of semiconductor devices, designers have traditionally used clock trees to evenly distribute the clock signals inside a chip. An external oscillator tuned to the desired operating frequency is fed into a pin on the chip where it is buffered into many "branches" that are then connected to all the "leaves" of the clock tree. That doesn't sound all that difficult until you try to comprehend the incredible accuracy that is required for SoCs that operate at a GHz or above. While IC designers work harder and harder to increase performance, the accuracy requirement of the clock is sabotaged by the phenomena of clock skew.

Clock skew refers to the difference in time when each clock signal arrives at the various "leaves" (or clock loads) on the clock tree branches. In a perfect world, there would be no clock skew – but in reality, it's a major design headache. And worse, any late design changes due to engineering change orders (ECOs) can cause the clock branches to become unbalanced and thus require a completely new clock tree to be designed very late in the project schedule.

The most common clock tree design methodology used today is to utilize clock tree synthesis (CTS) and optimization (CTO) tools available from many EDA vendors. It's not unusual for a chip design project to have several engineers working diligently to design and analyze their clock trees. Engineers meticulously place clock buffers and route clock signals across the entire chip, and then exhaustively simulate the timing of all clock signals over all temperatures, voltages, and process variations. Care must be taken at excruciating levels of detail during clock tree buffer and interconnect design and layout to ensure clock skew is minimized and balanced across all the branches of the clock tree. Why do IC designers pay such meticulous attention to clock skew? Because clock skew reduces performance.

Clock skew requires that circuit designers "guard-band" their chip-level timing, in essence eliminating a portion of the available clock period. For example, it's not unusual to find design specifications for high-performance chips requiring that 5 to 10% of a given clock cycle be "off limits" due to the uncertainty caused by clock skew, which

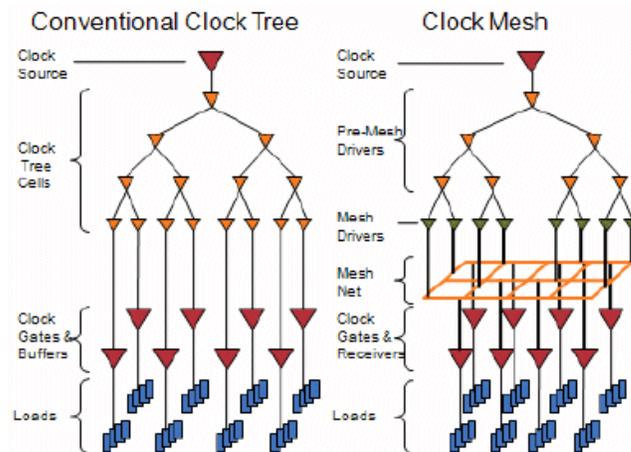
thereby reduces overall performance by the same amount. The diagram below illustrates the reduction in available clock periods, and therefore a reduction in performance, due to the guard-banding requirements of clock tree signals.



Performance Reduction Due to Clock Skew/Uncertainty

For many years ultra-high performance microprocessor design engineers have utilized clock meshes instead of clock trees to improve performance. A clock mesh is a uniform metal grid that is laid across the chip to distribute all of the clock signals, with local connections to clock loads available via a metal strap at every horizontal row or vertical column. The mesh differs from a clock tree in the way that it is driven by the incoming clock signal. In a clock tree, all signals radiate out in branch-like fashion from the clock source (tree root). This is why clock skew exists in clock trees - each branch of the clock tree is made up of a unique chain of buffers and intermediate loads, and therefore exhibits delays and timing characteristics that may have little in common with those of the other branches of the clock tree. Designers have to design, route, and analyze each clock branch individually, so that all the branch delays magically "line up" at the end.

A clock mesh however is driven evenly by multiple drivers spread out across an interconnected "grid" of metal wires. The outputs of all the drivers are thus shorted together by the metal grid, ensuring that the entire clock mesh oscillates as a single entity instead of as multiple, independent branches as in a clock tree. The result is that clock skew is virtually eliminated using clock meshes and higher performance can be achieved.



Clock Tree versus Clock Mesh  
Illustration: Synopsys, Inc.

Another key advantage of clock meshes is their immunity to on-chip process variations (OCV). As manufacturing processes shrink, the likelihood that the electrical characteristics of transistors across the chip will vary increases substantially. For example, a small variance in the amount of doping atoms in a diffused source or drain, or in the thickness of gate oxide layers, or in the critical dimensions of transistors can cause clock buffers on the same chip to have slightly different electrical properties – and that requires further guard-banding of the clock branches to ensure proper operation due to the potential of clock skew from OCV.

Clock meshes are immune to OCV because the clock mesh is composed of nothing but metal wires - no transistors in the clock mesh means no transistor characteristics can cause variation. Clock meshes therefore improve manufacturing yield due to their immunity to OCV. Higher yield means lower prices, along with higher performance, and IC designers have long considered switching to clock meshes for just these reasons.

However, until recently, clock mesh design was an entirely manual, difficult to analyze technology that was only utilized by very experienced and well resourced design teams. New advances in clock mesh design automation from EDA vendors, however, now allow clock mesh design to be adopted by mainstream IC design teams. So then, why aren't clock meshes being used extensively in SoC designs today?

Because of one simple problem – *clock meshes increase power consumption.*

Pretty dramatically too! Instead of each clock branch being uniquely optimized to minimize loading, a clock mesh consists of a big chain link fence-like grid covering the entire chip. And that big grid of metal means there is lots of capacitance to drive by the clock buffers. The clock buffers thus consume a tremendous amount of power driving the large capacitance of a clock mesh a billion times or more every second.

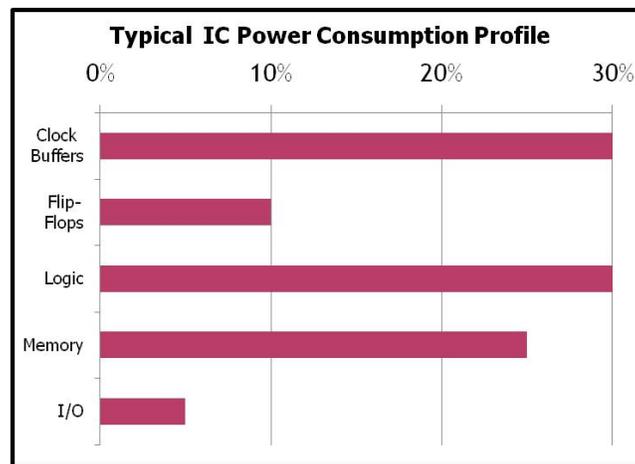
So, despite their higher performance, reduced design engineering requirements, and immunity to OCV, clock meshes have not replaced clock trees because semiconductor vendors are desperate to REDUCE power while increasing performance in SoCs. That's the power-performance conundrum ... designers today must sacrifice one or the other. Any engineering manager however who decides to implement a new design technology that increases power consumption most likely won't be around for the company's next design project. And in many cases, unfortunately, neither will their company!

Is the power consumed by the clock tree really that meaningful? It seems intuitive to believe that 100<sup>+</sup> million transistors on a chip would draw a lot more power than just a simple clock tree. You might be questioning the relevance of this whole discussion and

wondering if we're making a mountain out of a mole hill. Well, let's take a look at the power consumption profile of a typical high performance SoC so you can decide if indeed it is worthwhile to consider reducing the power consumed by clock trees.

## Where's the heat?

Surprise! The power consumed by the clock tree is meaningful ... it's on the order of 30% of the total power consumed by a chip! Think back to the GSA survey results on page 5 – the majority of designers of devices for MIDs believe they have to reduce power consumption in their next generation chip by at least 20 to 40%. If the power consumed by the clock buffers could be substantially reduced or eliminated, the power-performance conundrum for next generation MID devices would be solved!



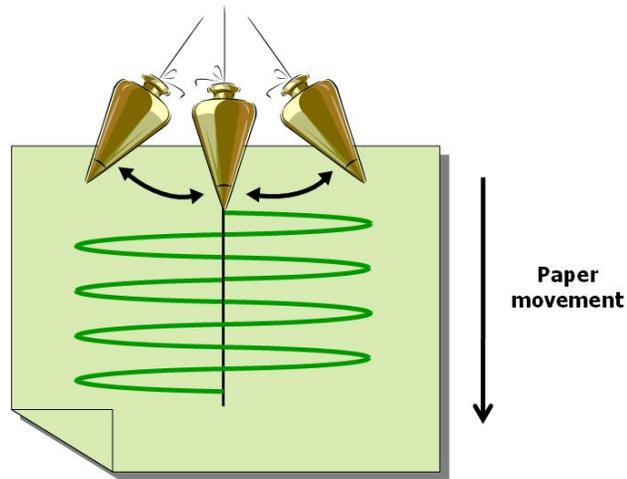
## An Electric Pendulum

Sometimes the simplest solutions are right in front of us. Consider the motion of a pendulum ... once someone provides the initial potential energy by extending its end to one side, it will oscillate back and forth in a steady, cyclical manner. And if it wasn't for friction, it would continue forever. Of course, every real system has friction so if we want the pendulum to continue swinging, we have to occasionally "nudge" it to replace the lost energy. But that's no big deal is it? We get a lot of work out of the pendulum at a very low cost – only an occasional nudge is required to keep it moving ... back and forth forever.

Simple, elegant, and reliable. All things IC design engineers greatly appreciate.

What if the end of the pendulum was a pencil point and had a slowly moving piece of paper attached below it that was synchronized with the swing of the pendulum? The

image the pencil draws would look like an ideal clock signal as the pendulum swings back and forth and the paper slowly moves forward.



RETHINK clocks. If it were possible to create an on-chip electric pendulum, wouldn't that be a really simple and clever way to generate a clock signal on an SoC? Like a mechanical pendulum, the only power required to keep it oscillating forever would be an occasional nudge ... else ... tick tock, tick tock, it just keeps swinging back and forth, providing us with a very useful clock signal that consumes very little energy. And that's the beautiful thing about a pendulum. Very little energy is required once the oscillations begin. But is it really possible to build an electric pendulum on a chip?

Of course it is ... all it takes is a simple inductor (L) and capacitor (C) connected in parallel to a power source. It's commonly called a "tank circuit" ... and Cyclos Semiconductor with its partners has proven it works on high-performance chips.

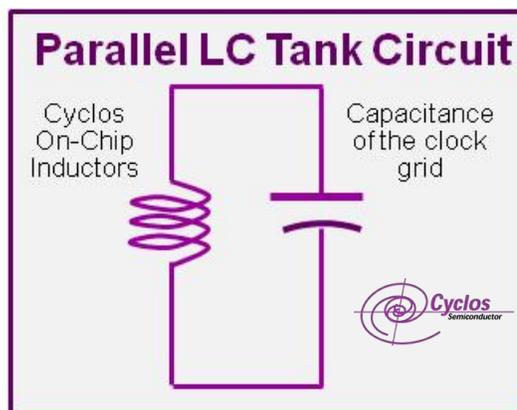
## Resonant Clock Meshes

Resonant clock meshes exhibit almost zero-skew and thus allow SoCs to operate 5 to 10% faster. And, a resonant clock mesh design process plugs right into an existing IC design flow as Cyclos is working with all the major EDA vendors tools. But most importantly, resonant clock meshes can reduce total chip power consumption by up to 30%. Here's how it works.

A tank circuit exhibits "simple harmonic motion", or resonance, almost exactly like a mechanical pendulum. When a charged capacitor is connected in parallel with an inductor, charge begins to flow through the inductor. The current flowing through the inductor in turn causes a magnetic field to form around the inductor and reduces

the voltage across the capacitor. Eventually all the charge will be gone and the voltage across the capacitor will reach zero.

However, the current will continue to flow because inductors resist changes in current. The energy to keep the current flowing comes from the magnetic field that formed around the inductor, which will begin to decline as the capacitor charges. When the magnetic field is completely dissipated, the current will stop and the charge will have been again stored in the capacitor, albeit with opposite polarity as when it began the cycle.



The cycle begins again with the current flowing in the opposite direction through the inductor. The result is a circuit where the energy continues to flow back and forth from the capacitor and inductor, creating an electric pendulum that provides a nifty clock signal. If perpetual motion was more than a pipe dream, no energy would be required to keep the tank circuit oscillating. In reality, a small "nudge" is required, like with a mechanical pendulum, to keep the tank circuit in continual motion due to resistance ... but it's much less energy compared to what it takes to drive a clock tree or clock mesh. Measurements in silicon have shown clock power can be reduced significantly using resonant clock meshes, which reduces total power by up to 30%.

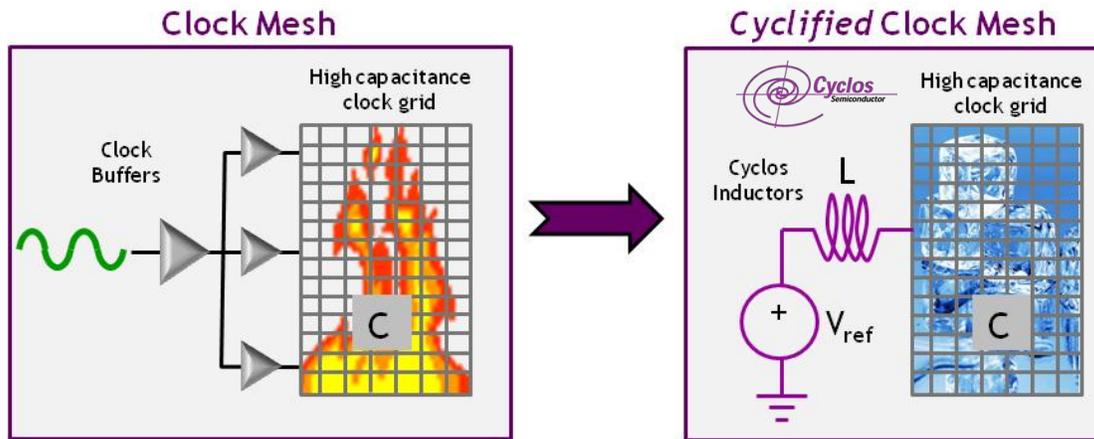
Recall that the natural resonant frequency of a tank circuit is easily calculated to be:

$$f_{\text{resonant}} = \frac{1}{2\pi\sqrt{LC}}$$

If the resonant frequency is our desired operating frequency, then all we need to determine are the values for L and C on our SoC to produce a low-power clock signal.

Now recall the benefits of a clock mesh versus a clock tree design. Designers love the simplicity of design and elimination of clock skew provided by the clock mesh, but couldn't tolerate the large increase in power consumption of the clock buffers required to drive the large capacitance of the all-metal grid. So let's start then with just a clock mesh, without any of those power-hungry clock buffers. All we need for a resonant clock mesh is the all-metal grid laid out on an optimal pitch across the chip's floorplan. Next we'll use a parasitic extraction tool to determine the actual capacitance of the all-metal grid, and since we know the desired resonant frequency, we can easily solve for the amount of inductance required to create an electric pendulum.

Resonant clocking marries the benefits of a clock mesh with the ultra-low power consumption of a tank circuit, which we call a resonant clock mesh. The inductors that are built on the chip, in parallel with the capacitance of the clock grid, produce a natural oscillator to drive the clock grid ... with hardly any power consumption.



No clock skew means much less engineering time and resources are required and the SoC will perform faster. Clock meshes are much more immune to OCV so vendors benefit from higher manufacturing yields (and hence lower parts cost). Also, last minute engineering changes will be much easier to accommodate without having to rip up and redesign a balanced clock tree. Best of all, however, is that the resonant clock meshes eliminate the large power consumption of traditional clock meshes. So IC designers get:

- ✓ Higher performance
- ✓ Reduced engineering costs
- ✓ Reduced schedule risk
- ✓ Higher manufacturing yield
- ✓ Faster time-to-market
- ✓ Up to 30% less power consumption

Cyclos resonant clock meshes solve the power-performance conundrum of next generation SoC design by eliminating clock trees and their power-hungry buffers.

## The Results

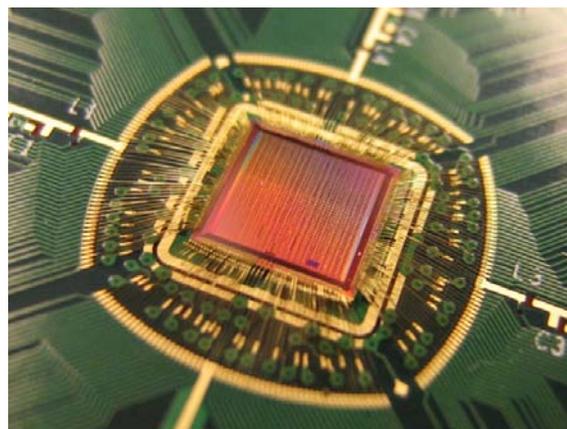
Cyclos has been working with several major semiconductor vendors for many years testing the commercial viability of resonant clock meshes and our customers are now implementing resonant clock meshes in their next production SoC designs.

Cyclos worked with ARM, for example, to perform a side-by-side comparison ... two ARM cores were fabricated next to each other on the same test die, one with the traditional ARM clock tree and the other with a Cyclos resonant clock mesh. What happened? Depending on what software application the cores were executing, total power consumption was reduced from 25 to 35% using resonant clock meshes

compared to the original ARM core. And considering that most new SoC designs at 1 GHz or higher implement multiple processor cores, the power savings from resonant clock meshes in embedded processor cores is quite compelling.

### ARM926 test-chip with Cyclos resonant clock mesh

- Working first silicon
- 25% to 35% lower power (application dependent)
- Identical performance to original ARM core
  - ✓ 250MHz+ on the same silicon test die
  - ✓ Faster design and timing closure
  - ✓ Standard RTL IC design flow
- No special libraries or processes required



After this encouraging demonstration of resonant clock mesh technology, Cyclos extended the technology to work in multiple GHz DSP and embedded processor designs with equally impressive results. And while many optimizations are possible in the resonant clock mesh design process, Cyclos has already addressed many of the finer details that are required for implementation of resonant clock meshes. That's why resonant clock meshes are being designed today into next generation SoC designs by semiconductor vendors. It's the silicon proven, low-risk, and affordable solution to the power-performance conundrum of today's IC design process.

### Too Good to Be True?

Resonant clock meshes have been silicon proven to be very effective in reducing power consumption and increasing performance, but as with most new or alternative design technologies, there are tradeoffs to be carefully considered before taking a leap of faith.

- The power reduction benefit of resonant clock meshes is significantly reduced at low-speed operation. For SoCs operating below a GHz, IC designers usually continue to use their tried and true clock tree methodologies ... after all, why risk changing anything that is known to work unless the benefit of a new design technique is meaningful? At a GHz, resonant clock meshes provide meaningful power reduction. Above 2 GHz? Resonant clock meshes are a saving grace.

- Additional area is required to implement the on-chip inductors. Our results show, on average, a resonant clock mesh will increase die area by ~ 4-5%. Bear in mind, many SoCs designed today are pad limited and therefore much empty die area already exists. And it's important to consider that the area increase due to inductors will not reduce manufacturing yield as significantly as an area increase due to adding logic, memory, and interconnect. Why? Because the defect density of all-metal inductors is orders of magnitude less than for logic and memory blocks where transistors are minimally sized and packed together. An area increase due to resonant clock meshes will however decrease the number of candidates per wafer so it does need to be considered during product planning. Typical inductor values required for GHz<sup>+</sup> operation are in the .75nH to 1.25nH range and require ~ 100 x 100 μm<sup>2</sup>.
- Thick metal layers make better inductors and thus older process nodes may not be suitable for resonant clock meshes. Most process nodes at and below 65 nm have adequate metallization for highly efficient inductors – and most GHz<sup>+</sup> SoCs are being manufactured in these processes so it's not really an issue, but should be discussed with your foundry when considering resonant clock meshes.
- Most modern SoC designs have multiple clock domains. Resonant clock meshes require no changes to the existing methods for managing signals that cross clock domains. The fact that the clock signal is being derived from a resonant clock mesh versus a clock tree makes no difference to the RTL designers who must deal with clock domain crossing.
- With any new IC design technique there is always risk. Most SoC designs today are so highly complex and have such intense time-to-market requirements that IC designers tend to avoid change as much as possible to avoid risk. Every once in a while though, the benefits of a change are so advantageous and proven that IC designers will commit themselves to adopting a new IC design technique.

Ask yourself ... what if IC designers had resisted change in the 90's and not adopted RTL design flows. Change is difficult but when the benefit is large, it's a worthwhile undertaking. Rethink clocks.

### Been there, done that?

As was mentioned earlier, clock meshes have been researched and analyzed by IC design groups for many years due to their huge success in high-performance microprocessor designs. But the initial results were less than impressive for two

reasons. First, without the resonant clocking technology, the large increase in power consumption due to the clock mesh, that was not a problem for microprocessor designs, was completely unacceptable for SoC designers who have to meticulously count every milliwatt of power consumed. And, without an investment from the EDA vendors in clock mesh design tools and flows, it was typically too clumsy and immature for SoC designers to consider adopting because SoC design teams seldom have the large CAD support group that a microprocessor team is fortunate to have.

So some SoC design groups that set off eagerly to analyze clock meshes with great expectations now feel the same way about clock meshes as someone who just touched a hot stove or licked a frozen flag pole. They typically are very reluctant to do it again!

Fear not. With the availability of silicon proven IP, expert design services, and mature design utilities, Cyclos is making resonant clock mesh design adoption minimally painful and maximally beneficial. Ask yourself if you can afford not to start adopting resonant clock meshes today... will your competitor leap ahead of you in the power-performance curve? And the important word to fear is LEAP because resonant clock meshes provide substantial power reduction and performance improvement as proven in silicon. It's real. And it's reasonable to expect, like with the rapid adoption of logic synthesis in the early 1990s, that 10 years from now resonant clock meshes will be the standard, not the exception. RETHINK clock meshes – this time with resonance!

## What's the Payback?

Every engineer and manager has to justify spending money. And usually the bean counters in the Finance department want the justification in terms of return-on-investment (ROI). With resonant clock meshes, ROIs are easily estimated based simply on asking yourself "what's the incremental revenue that can be expected from an SoC that has 5-10% more performance AND 20-30% less power consumption than any competitor?".

One smartphone vendor estimated that they could charge an extra \$5 to \$10 per phone if battery life was 20% longer than their competitors. Extra performance was just icing on the cake. Predictions are that production of the iPhone will reach 8 million units in the third quarter of 2011. If battery life was extended through the use of resonant clock meshes in the iPhone processors, one would expect then an additional \$40M to \$80M in revenue – per quarter! With the estimated cost of implementing a resonant clock mesh being ~ \$1M, that's an ROI that easily exceeds 40x in the first quarter alone! Those are numbers every bean counter will love to see!

## The Future

Resonant clock meshes have been silicon proven to be very effective in reducing power consumption, increasing performance, and are scalable for advanced IC design implementation and manufacturing. As mobile consumers demand more performance and longer battery life from their handheld devices, resonant clock meshes provide the ideal solution and thus are being deployed by semiconductor vendors in 2011. As cloud computing becomes more pervasive, data center vendors will also be taking actions to reduce their power consumption and cooling costs. Once again, resonant clock meshes are an ideal solution and they will soon be deployed in networking devices that enable the Internet. Other design applications for resonant clock meshes are plentiful too ... there is no shortage of IC designers who need to reduce power consumption and increase performance. With a novel IC design solution to the power-performance conundrum provided by Cyclos, SoC vendors are going to benefit greatly.

Typically, however, a new IC design technique with such incredible and proven benefits as resonant clock meshes always seem to have some sort of gotchas ... something that takes the wind out of your sails, bursts your balloon, or leaves you standing all alone at the altar. But fortunately with resonant clock meshes, it just isn't the case. In fact, resonant clock meshes, in addition to reducing power and increasing performance, also:

- ✓ Are easier to design than traditional clock trees = reduced schedule risk
- ✓ Readily accommodate ECOs late in the design cycle = reduced schedule risk
- ✓ Require fewer engineering resources to implement = reduced engineering costs
- ✓ Are less sensitive to on-chip process variations = increased manufacturing yield and thus reduced part costs and/or higher profit margins
- ✓ Easily integrate with existing IC design tools and physical design flows = no special engineering or IT/CAD support required.
- ✓ Require no special processes or libraries = no increase in manufacturing costs

Cyclos provides electronic design automation tools, intellectual property, and consulting services to IC design engineers. If you would like to find out more about how resonant clock meshes can solve your power-performance IC design conundrum, contact Cyclos at [info@cyclos-semi.com](mailto:info@cyclos-semi.com). We're ready to assist you in achieving GHz design success.

## RETHINK Clocks

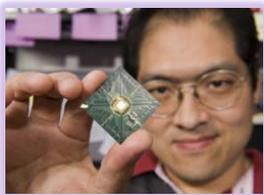
## Background Information

Cyclos Semiconductor was established in 2006 via the Technology Transfer Office of the University of Michigan in Ann Arbor, Michigan. The company currently is located in Berkeley, CA and the team consists of the following members:



- Marios Papaefthymiou (PhD MIT '93), Co-Founder, President

- ✓ Leading researcher in commercial application of resonant-clocking technologies
- ✓ Tenured Full Professor of EECS, Univ. of Michigan



- Alexander Ishii (PhD MIT '92), Co-Founder, VP of Engineering

- ✓ 15 years of experience across entire semiconductor value chain (NEC USA, several semiconductor startups)
- ✓ Extensive experience and expertise in EDA flows, ASIC and custom VLSI design

- Cyclos Design Engineering team



- ✓ Experienced, diverse team with expertise in all aspects of advanced IC design
- ✓ 100+ years of IC design experience with over 50 tape-outs
- ✓ Previous designs include: HDTV chips, processors, graphics chips, memories, standard cell library development, Sony Playstation Floating Point Unit

Investors in Cyclos currently consist of ARM, Siemens, and High-Tech Venture Capital. The company intends to execute another round of funding in 2012 to facilitate faster commercial growth via an expanded global sales and marketing channel. Funds will also be used to staff an EDA software development team to develop design automation tools that will substantially improve engineering productivity in resonant clock mesh design flows. Interested investors are urged to contact Cyclos at [info@cyclos-semi.com](mailto:info@cyclos-semi.com) for more detailed investment information.