Why Synthesizable-digital PLLs Are No Substitute for Hardened Mixed-signal PLLs

John G. Maneatis, Ph. D.
President, True Circuits, Inc.
May 22, 2019

I reviewed John Cooley's Deepchip article titled 'Movellus beats out True Circuits PLL/DLL IPs as #9 "Best of 2018"' (http://www.deepchip.com/items/dac18-09.html) and found it quite entertaining. The article suggests that Movellus has us all beat (e.g. TCI, AB, SC). In fact, it is quite the opposite. I cannot speak for AB or SC, but I can certainly speak for TCI.

We were in the PLL licensing business at the beginning, before people even knew the meaning of IP in an SOC, in the days of 0.25um silicon. Over this time period, we have developed a lot of innovative PLL technology. With each technology generation, the preferred mix of analog and digital circuitry has changed and we have addressed the challenges with a variety of architectures. We currently support all of our PLL and DLL designs in over **100 IC processes**. We have a lot of experience and understand the problem very well.

The question is not analog versus (synthesizable) digital -- it is the delivery model and how the responsibility is split between the customer and the supplier. Movellus does not have a lock on all digital and synthesizable PLL, DLL, LDO designs. We have developed similar technology and some of the IP in our portfolio is fully synthesizable, but offering it as soft IP is another issue. At TCI, we always try to deliver the best solution to address our customer's needs. **It is inaccurate to label TCI as a strictly analog PLL/DLL supplier.**

Let's consider the technology trade-offs in a PLL. First consider the analog versus digital split. Making the loop filter digital instead of analog has the potential for large area savings in low-bandwidth PLLs (commonly those that multiply from low clock frequencies like 32KHz). Each added bit in the accumulator is like doubling the loop filter capacitor sizes, which often dominate the area of a PLL (assuming the charge pump current is already set to a minimum value). Unfortunately, the power of the digital loop filter is another matter. If no TDC or DTC is used, then the power will be limited mostly to the accumulator running at the reference clock frequency. If a TDC or DTC is used to improve the loop dynamics or support fractional-N, the power will be much higher. In comparison, an analog loop filter uses very little power and is pretty much limited to a small multiple of the tiny charge pump current.

PLL Sub-Function Analog Versus Digital Comparison		
Item	Analog	Digital
phase detector	linear PWM output	binary output without TDC (or DTC)
loop filter	lower power	higher power
	can dominate area	small area
VCO	lower power	higher power (especially for low frequencies)
	lower jitter with PSRR	higher jitter, needs linear regulator
	wide continuous tuning	narrow tuning, needs many bands
	linear tuning	potentially non-linear tuning
	analog control	digital control with limited resolution
biasing	built-in supply regulation	uses core voltage with no regulation

Now consider the VCO. The VCO is all about jitter performance, low power, linearity, and tuning range. Analog VCOs can be made to have very good jitter performance with integrated linear regulators on the local supplies or in the biasing. Analog VCOs might be composed of "digital" inverters, but they have some analog biasing to control their local supplies. Synthesized digital VCOs will have to operate on the core supply with no regulation and thus poor jitter performance, or require an external linear regulator to improve jitter performance (an all-digital LDO cannot easily be used because of the quantized load steps and slower response time). As such, better jitter performance will always be possible in an analog VCO and it can be achieved using the core supply voltage. The power consumed by the VCO is C*Vs*Vdd*F (where C is capacitance in VCO, Vs is VCO swing, Vdd is the supply voltage, and F is the frequency). Synthesized digital VCOs have Vs=Vdd, but analog VCOs use Vs equal to Vdd/2 or Vdd/4, so analog VCOs have the potential to run at 1/4 the power of synthesized digital VCOs. For low frequencies, the difference is much more dramatic, since synthesized digital VCOs cannot increase R (where F=1/(R*C)), but analog VCOs can, so such synthesized digital VCOs require a larger C (from load capacitance or long inverter chains, or indirectly from output division) and thus have much higher power dissipation. Common trends may suggest that that synthesized digital VCOs are lower power, but it is only because analog VCOs typically use larger devices (with larger C) to improve jitter performance.

The VCO tuning range is also an important issue. To make a PLL robust and able to deal with spread-spectrum input clocks and large temperature and voltage changes, a wide continuous tuning range is needed. Analog VCOs with linear tuning ranges that cover several orders of magnitude have been

demonstrated (by me and others) -- a guaranteed 10:1 or 5:1 range over PVT is no problem. Synthesized digital VCOs ideally would have similar wide linear continuous tuning ranges with high resolution on the clock frequency. In fact, our digital VCOs have this property. Unfortunately, most designs proposed do not. They have a large number of frequency bands with highly non-linear frequency control. What this means for your PLL is that once locked, where a locking assist circuit has found the center of one of the bands, very little change in input frequency, voltage, or temperature can be tolerated without the design producing large amounts of jitter and glitching. In a synthesis flow, it is also difficult to guarantee a worst-case maximum step size and the step monotonicity, which can lead to excessive jitter at random operating frequencies that change with PVT. In fact, you would need to characterize all VCO codes possible for the PLL in PRODUCTION TESTING to guarantee that the PLL will behave as desired and not produce too much jitter.

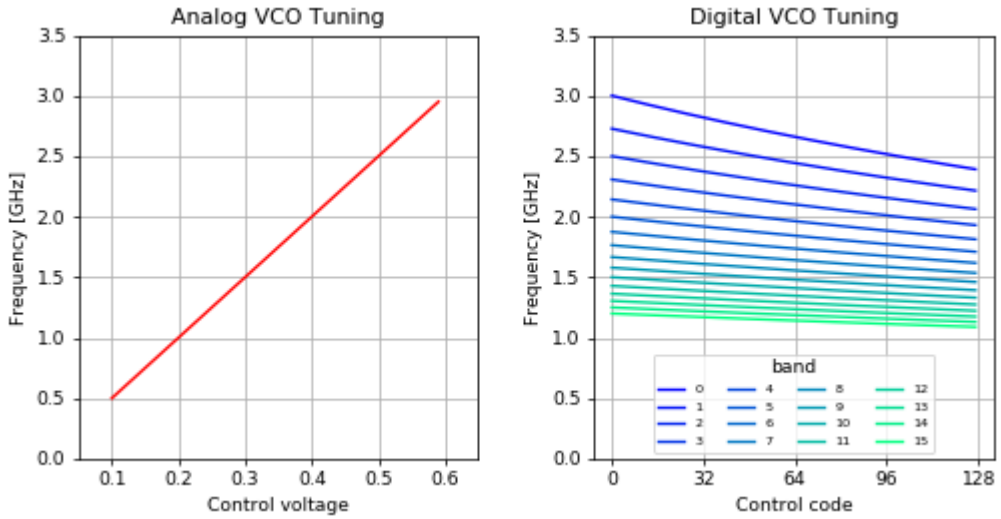

In summary, given these technological issues, we are only considering shipping fully-synthesizable PLLs for low-end low-risk applications. For high-performance applications like SERDES clock generation, we will only ship our hardened mixed-signal PLLs, which may employ a full spectrum of analog and fully-digital components.

I should point out that the trade-offs can be lessened to some extent when custom standard cells are used in a fully-synthesized design, but using them will further complicate the customer delivery model. Also, my above discussion focused on only PLLs. DLLs have many similar issues (especially jitter), but DLLs are fundamentally simpler designs. In contrast, LDOs tend to be digital designs, and except for the header transistors and some needed capacitors, can be fully synthesized.

As part of our design flow, TCI performs a lot of synthesis and place-and-route for our PLL and DLL designs, but we currently do not ship them as soft (customer synthesizable) IP. Instead, to maximize design efficiency, porting and reuse, TCI employs a unique and proprietary CAD flow, with a broad set of tools designed for analog and mixed signal design that you will not find anywhere else. This automation gives us a unique capability to design mixed-signal circuits quickly, easily support them

over a wide range of foundries/processes, and obtain optimal performance. The result of this automation is a large set of modular designs which can be easily ported forward or backward or across foundries, are "correct by construction", and are delivered hard not soft. When a customer needs a PLL in "the other foundry", or a different metal stack or with particular oxides or Vts, we can support those requests in hours or days, not weeks or months. You may remember the IP/Design Track presentation I gave at DAC in 2017 titled "The Secret to Building IP at the Cutting Edge", where I went into detail about how and why we develop and maintain our own highly automated CAD environment.

Finally, let's consider key question of the delivery model and how the responsibility is split between the customer and the supplier. It may seem at the surface that fully-synthesizable design makes life easy for the customer. They just drop it into their RTL, synthesize, place and route, and they are done, right? In fact, it is not that simple. A lot of issues can occur in synthesis, such as introduction of non-linearity or non-monotonicity in frequency or delay, that will lead to poor jitter performance. Dealing with these issues is tricky. There is no guarantee that one synthesized design will work as well as the same design synthesized for a different chip in the same process. This directly or indirectly places a lot of the responsibility on the customer, but they do not want it. They just want a silicon proven hard macro that they can drop into their chip. They want it designed by experts and want the experts to consider all issues so they do not have to consider them. While it may look that the chip is built around these hard macros, it happens for free as the place and route logic is happy to fill all available space.

PLL Hard Versus Soft Delivery Model Comparison		
Issue	Hard macro	Soft macro
integration	included in RTL but not synthesized	impacts synthesis
floorplan	must plan around block	no constraint
RTL config	limited	flexible
pin programmability	flexible	flexible
production testing	test at extremes	must test VCO at all points
performance	low jitter common	low jitter requires linear reg
predictability	predicted by test chips	can be different on customer chip
customer responsibility	only requires correct usage	requires correct synthesis

You may think that SoC customers want the ability to tweak their PLLs to get exactly what they need. Our 20+ years of experience says they really **do not**, and **probably should not even try**. Instead, we have found that **they would rather have a pin programmable solution** that supports a wide range of operating requirements without the need for tweaking. At TCI, we take the time to understand our customer's requirements, direct them to the correct solution and deliver and support a hard macro that best fits their needs, one they can count on for silicon success.