40LP Digital PLL (DPLL_40LP)

Granite SemiCom Inc.

High-Speed, Wide-Range Digital PLL in TSMC 40LP

Granite Semicom has just shipped GDS-II of its totally-integrated (no off-chip components) Clock-Driver and Clock-Multiplying-Unit IP block including fractional-N in TSMC's 40LP process; silicon is expected Q2 2013. In addition to porting its TSMC 40G version, a number of minor changes were incorporated. The oscillator was improved to have higher speed and less jitter, for a given power. Some critical current sources were modified to give an expected 40% reduction in long-term accumulated jitter. In addition, some feedback dividers where changed to give lower power. This block is intended for applications such as the CMU in a SERDES PHY, and for clock-driver applications. This IP block is a digital phased-lock-loop, plus an integrated voltage and current reference, a variety of input amplifiers (single-ended and differential, dc and ac-coupled), a number of programmable dividers, a serial interface for programming, and a high-speed 50 ohm driver capable of driving off-chip at full-speed. The specification is for the Digital PLL (DPLL_40LP) to operate between 0.5GHz and 7.5GHz over process corners between -40 and +125 degrees celsius; the power dissipation is less than 30 mw (for a 5GHz output), and the IP requires a 0.11 mm^2 area (not including pads and the output driver). Long term accumulated jitter is specified at less than 1ps rms accumulated over 260 periods.

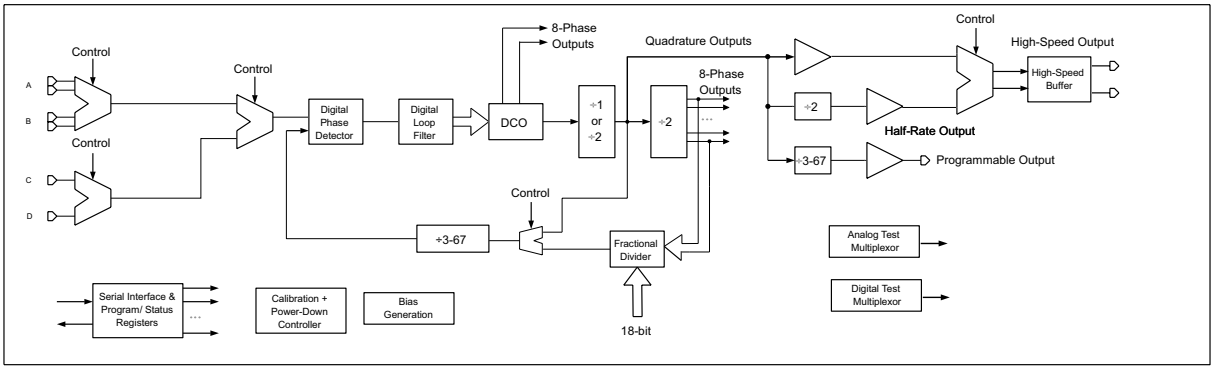

Fig. 1. Block diagram of DPLL_40LP digital PLL and clock-driver IP.

A Digital PLL?

A *Digital PLL* is a PLL that is mostly digital, that does not have an analog loop filter, and that also does not have analog charge-pump circuits. These sub-components are replaced with digital equivalents in a ***Digital PLL.***

Why a Digital PLL?

1. **Wide Range and Programmability:** A digital PLL is much easier to design to be used over a wide range of reference frequencies and output frequencies; the dynamics of the loop filter automatically track with the reference frequency. This should be compared with an analog approach where the analog loop filter must be *'digitally trimmed'* for different bandwidths and reference frequencies, and the tuning range is limited. In addition, a highly-programmable digital approach can be digitally adjusted for many different application requirements; this allows for a ***One Size Fits All,*** minimizing both costs and support complexities. The programmability of the IP block includes output impedance, loop bandwidth, output division ratio, VCO load capacitance, input reference amplifier type, current bias levels, etc.
2. **Predictability:** the PLL is primarily composed of digital logic circuits which are highly predictable especially in sub-micron processes. ***What you simulate is what you get.*** Digital circuits are not affected (to the same degree) by analog issues such as leakage currents through loop capacitors, offset voltages due to transistor mismatches, substrate noise, reference feedthrough, etc (effects minimized, not eliminated). Predictability is critical as one goes to 40nm and smaller technologies; a second iteration caused by a poorly working PLL can incur millions of dollars of costs and loss of revenue and months lost in time to market.
3. **Porting:** porting a PLL to a new technology is much faster and involves much less risk for a Digital PLL; most of the analog-like circuits are realized using only I/O transistors, and these seldom need to be changed during the porting process; most of the porting process when going to a smaller technology is to ensure layout design rules are not violated.
4. **Size:** because the loop-filter and analog charge pumps have been eliminated, the size ($0.11 mm^2$) is similar to or smaller compared to analog approaches; this is especially true as one goes to smaller technologies.
5. **Jitter:** a major source of jitter for analog PLL's (often the dominant source of jitter) comes from the loop filter, charge-pumps, and V/I conversion circuits. This is especially the case when non-calibrated charge-pump offsets, and power-supply noise are considered. These jitter noise sources can be significantly minimized with a properly designed digital approach.
6. **Testing:** the HSDPLL has a number of built-in testing features that can assessed through the serial SPI bus; these include access to the calibration values to characterize the speed of the process, an "IN-LOCK" output, and a built-in accumulated jitter measurement that gives the highest confidence all is working as expected.

Specifications:

As is the case for all PLL's, the specifications are dependent on the power, oscillation frequency, reference frequency, loop bandwidth, and loop damping factor. All specification numbers are based upon either measurements or careful simulations over process and corners.

An example set of specifications for a loop optimized for high-frequency (a 5GHz oscillation frequency for a 10Gb/s half-rate SERDES application) 100MHz reference frequency, a 2MHz loop bandwidth (with less than 0.1dB transfer function peaking) and a 1.8V high-voltage option, are:

1. **Power:** less than 30 mW excluding output driver, for a 5GHz output frequency; much less at lower output frequencies.
2. **Temperature:** -40° C to 125° C.
3. **Jitter:** accumulated total jitter using a 1667 golden-filter (equivalent to accumulating jitter over 260 periods) is specified to be less than 1ps rms depending on bandwidth and reference frequency. The numbers go up about 10-15% when fractional-N division is enabled.
4. **Lock Times**: less than 10 us for complete calibration at start-up and lock (nominally 5-6 us); and less than 1us for lock from power-down when calibration has been remembered from previous power-up. When starting up without calibration, the loop starts in frequency lock and their is no cycle-slipping.
5. **Off-Chip Components:** none.
6. **PSRR:** better than 46 dB (voltage to free-running oscillation frequency) at lower frequencies, peaking less than 20 dB at all frequencies across process and temperature.
7. **Fractional/N Capability:** included; this can be removed to save space when not needed. When included, it can be enabled or disabled. When enabled, it increases jitter by 10%-15%.
8. **Temperature Stability:** all process variations are calibrated out at start-up. The integral path eliminates any variations over the complete temperature range when in lock. Since all loop parameters are proportional to the free-running frequency, which is calibrated, effectively all loop parameters are accurate to around 1%-2%.
9. **Programmability:** there is a divide-by one or two, followed by a multi-output-phase (8) divide by 1, or 2 divider right after the oscillator inside the feedback loop followed by a 6-bit programmable feed-back divider. These give a very wide range on the ratio of effective oscillator frequency (the output of the first dividers) to reference frequency (the output of the second divider). In addition, outside the loop at the output, one can take the full-rate output, a half-rate output, or the output from a 6-bit programmable output. The loop bandwidth, damping factor, free-running frequency, Frac/N, and bias currents are all programmable (through an SPI serial interface - this can easily be replaced by an I^2C interface). There is also an option to access all registers using a parallel bus instead of the SPI bus. All programming registers have default values, so often programming is not necessary. A Graphic-User-Interlace (GUI) (used in both simulation and design) is available for programming the registers, to guarantee defaults are all set correctly, and to ease the integration into an SOC.
10. **Optional Custom Features:** a number of optional features (not included in the current implementation) have been designed and can be included in “custom” realizations: These include: an output divide-by-1, divide-by-2, or divide-by-4 block, that can be user programmed “on-the-fly” without runt pulses or missing pulses. In addition, the 6-bit feedback divider can be swapped with a 9-bit feedback divider, that also can be programmed “on-the-fly” without glitches. This allows for the output frequency to be wide-band dynamically programmed to lower frequencies during times when the highest speed is not necessary, thereby saving system power. The output frequency “ramps over time” to the new desired frequency without the loop losing lock. Finally, many different variations on the input, feedback, and output dividers, as well as output multiplexor options are available for custom applications.