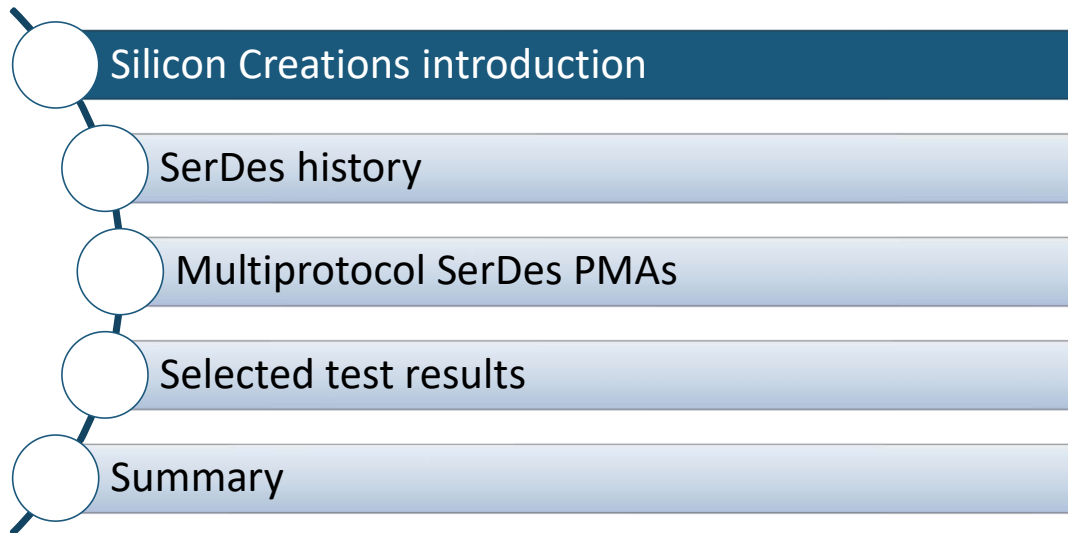


Versatile, wide range Multiprotocol SerDes solutions for challenging standards and applications including DisplayPort and PCIe5

Andrew Cole, CC Chen & Jeff Galloway




In this talk I'll provide a little history of SerDes interfaces and a couple of interesting comparisons to parallel data interfaces. After this I'll introduce you to our multiprotocol SerDes PMA product line highlighting some differentiating features. Following this I'll show a selection of Silicon test results. Before doing this, I'll begin with three slides reminding you about our company and main product lines.

Silicon Creations Overview



- IP provider: PLLs, Oscillators and SerDes
- Design: Atlanta, USA and Krakow, Poland
- TSMC 2017, 2018, 2019 and 2020 “Analog Mixed-Signal Partner of the Year”
- SMIC awards 2013 – 2016 including “Best Customer Support”
- ISO9001 certified
- GF22FDXcelerator member
- Mass production from 5nm to 180nm
- 3nm PLL ready for design starts
- >250 customers, >100 in China



© Silicon Creations, 2020

Galloway, Chen & Cole – Multiprotocol SerDes

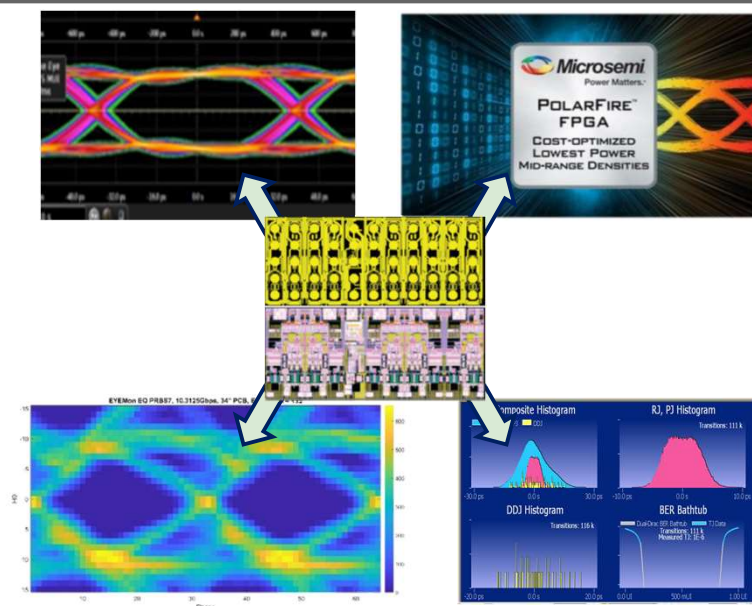
3

We have been providing clocking IPs to our fabless IC customers since 2006 and have been growing consistently since 2010 as you can see from the graph of customers by region and year. We have development centers in Atlanta, USA and Krakow Poland our customers chips using our IPs are in volume production from 180nm down to 5nm FF, and we can support 3nm design starts today. Our design engineers understand one of their main missions is to help our customers get to market quickly and our quality and support has been recognized with best-of awards from TSMC and SMIC. We’ve also recently achieved ISO9001 certification which is appreciated by our customers making automotive chips.

SerDes from Silicon Creations



- Robust and in production from 12nm to 180nm and from <100Mbps to >25Gbps
- Multiprotocol in TSMC, GF, UMC 6nm, 12nm, 16nm, 28nm, 40nm
- Targeted protocols including SGMII, XAUI, RapidIO, V-by-1 HS/US, DP, FPDLink, OIF-CEI, JESD204, CPRI, PCIe1-5, 10G-KR, ...



© Silicon Creations, 2020

Galloway, Chen & Cole – Multiprotocol SerDes


4

SerDes is one of our key product lines, and we focus primarily on the PMA – the high-speed analog part including the serializer, line driver, line receiver, equalization, CDR and deserialization. Our PLLs have enabled extremely good performance in the Microsemi FPGA, and also very low power. Over the past few years we've ported and proven this multiprotocol PMA to a number of mainline process nodes. In addition to our multiprotocol PMA, we have protocol-specific PMAs built for V-by-one, Display Port, JESD204, FPD-link and several custom interfaces.

PLLs from Silicon Creations



- Highest volume analog IPs
... Many millions of wafers 5nm to 180nm
... robust design and good QA are essential
- PLL products include general purpose, fractional, low jitter AFE, μ W IoT, Automotive



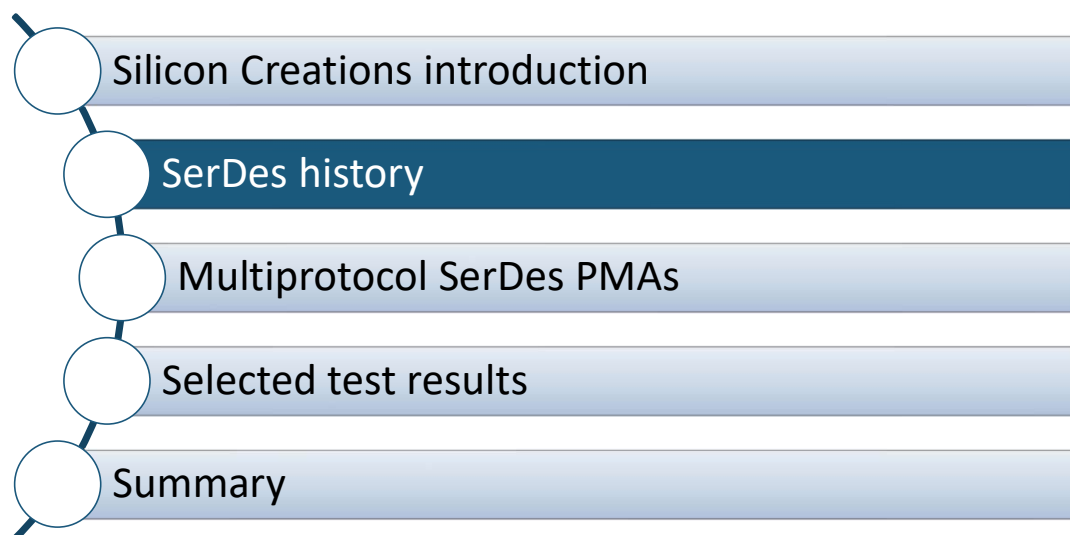
© Silicon Creations, 2020

Galloway, Chen & Cole – Multiprotocol SerDes

5

Slightly less than half our revenue now comes from PLLs. The core of this business is general purpose Fractional-synthesizer which so widely programmable it's been used in hundreds of chips and a wide variety of markets. This has resulted in some truly gigantic volumes. For example in TSMC 28nm over 140 production chips use this IP and in 12/16nm over 50 chips use it. This has resulted in many millions wafers delivered to customers with our PLL IP. These kinds of volumes are a testament to our robust, high-yield design and customer support. We've leveraged this high-volume design to make PLL variants that are optimized for low jitter, extremely low power and Deskew/clean-up PLLs for the Cadence DDR PHYs.

Outline



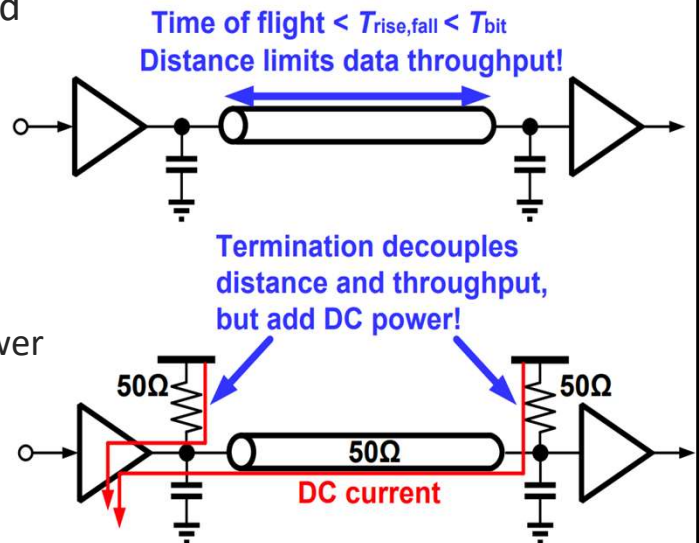
Why SerDes?

- Before SerDes, chip-to-chip data used TTL, LVCMOS, ...

- ✓ Simple
- ✗ Bit rate limited by time of flight
- ✗ Power = $k \cdot \text{rate} \cdot \text{PCB capacitance}$

- 50Ω interfaces

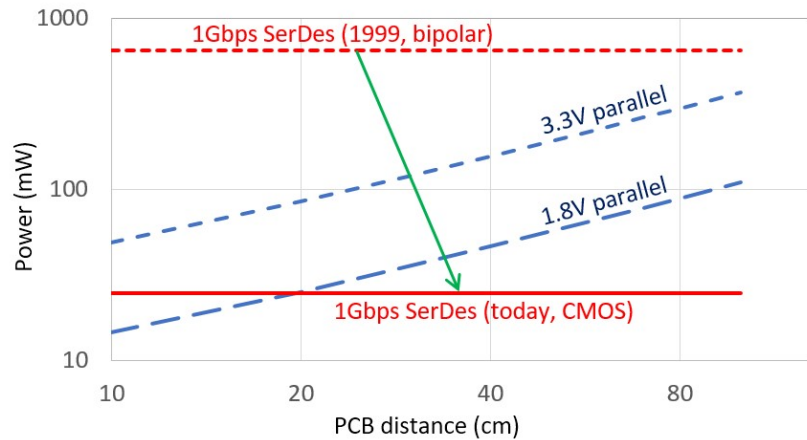
- ✓ Power = $k \cdot V_{\text{swing}} / R_{\text{term}}$
- ✓ Distance weakly impacts bit rate, power
- ✓ Serialization lowers pin count, PCB traces, energy/bit
- ✗ Complex (✓ good IP fixes that)



Let me start the technical portion of the presentation with a reminder why we make SerDes at all. Before SerDes we used bunches of wires with single-ended rail-rail signaling and either a shared system clock, or the data sender also generated a clock. For example, PCI initially used 16 data wires at 33Gbps. Over the years this evolved – adding more signal wires and running them faster until we got to PCIx with 64 wires at 133Mbps each. If you multiply this out, it's not that bad. PCIx has a bandwidth of 8.5Gbps. But there are some issues. 64 pins is a lot of wires on the PCB, and when you account for the supply pins on the package you probably needed about 128 pins on the package for a PCIx interface – about 66Mbps/pin. Going faster proved impossible and keeping the wire lengths matched and limited wire length became a limit on the number of cards we could plug into the PC. The lower graphic shows the solution telecom providers had found much earlier – treating the wires on the PC as matched transmission lines rather than lumped capacitances. At the end of the 90's we saw PCIe1 at 2.5Gbps begin rolling out. For this we get about 250Mbps/pin (accounting for the unidirectional nature of most SerDes) and could make the backplanes or board-board connections much longer. The higher data density allowed processors to communicate with multiple chips simultaneously.

Power comparison – 1Gbps

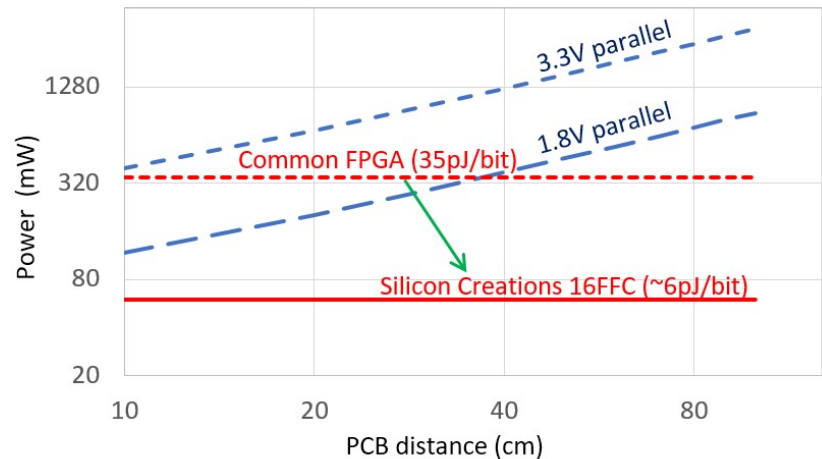
- Between 1999 and 2019
 - Power for a 1Gbps link dropped from ~600mW to 25mW
 - = 600pJ/bit to 25pJ/bit (compared to 3pJ/bit for a 32Gbps interface)
- For a 1Gbps link, parallel interfaces are competitive for short links



It's interesting, however, to compare parallel interfaces to SerDes. Because the serial busses have transmission lines their power is independent of the chip-chip distance. On the right, the red lines show how the power of a typical 1Gbps interface has evolved with improvements in design concepts and available process technology – from typically 600mW to perhaps 25mW today in an older CMOS process. Those of us making the first gigabit interfaces in the bipolar technologies needed at the time to run this fast knew they were not extremely efficient. The blue dashed lines indicate the power consumed by parallel interfaces running at 3.3V, and 1.8V. As the bus length increases the bus has more capacitance and most Gigabit serial interfaces today use less power than a parallel bus. But even today, parallel interfaces are more power efficient for gigabit chip-chip distances below 20cm or so.

Power comparison – 10Gbps


- Comparing SerDes solutions for 10Gbps backplane with parallel
- FPGA SerDes #1 uses > 5x power of Silicon Creations solution in FPGA #2
- Parallel power beats FPGA #1 for short distances!
- Silicon Creations beats Parallel over any distance



Of course, on the first technical slide I pointed out that serial interfaces only have a weak dependence of power on speed. In fact, the energy per bit sent drops a lot as we design the interface to run faster. The lower red line on this slide represents our TSMC 12FFC SerDes PMA running at 10Gbps with amplitude and equalization set for a 40" to 80" backplane. 6pJ/bit translates to 60mW. The upper red line compared this to the interface power in a tier-1 manufacturer's FPGA and shows how SerDes IPs in ASICs can be a lot more efficient than FPGAs. The same blue dashed lines are shown here for 10Gbps now. It's interesting to extrapolate to the left and draw the conclusion that parallel interfaces will still win on power for really short distances and remain attractive when the pin count doesn't matter much. This is one of the driving forces behind 2.5d packaging. After all, parallel interfaces will win every day for interface latency. Not having data encoders/decoders and the like simplifies the data path.

CMOS for everything

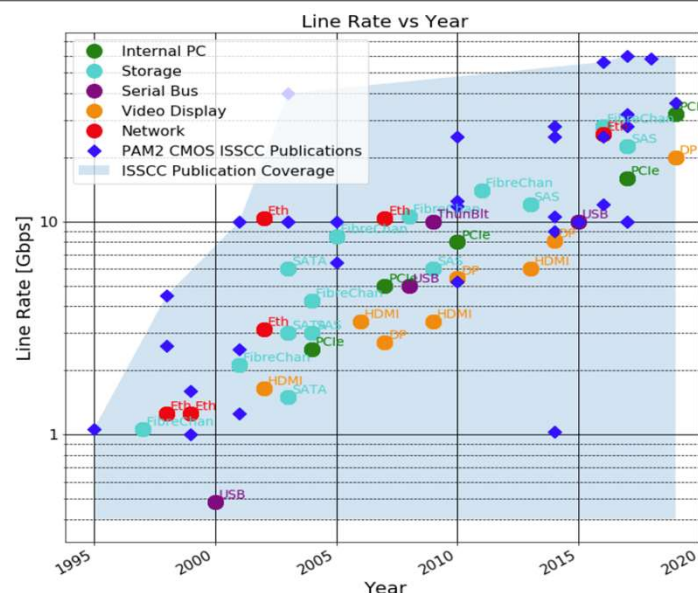
- International Solid-State Circuits Conference (ISSCC) publications
- < 2000: mostly Bipolar/BiCMOS
 - Process needed for high rates ... SerDes standards limited to (expensive) networking/optical
- > 2005: mostly CMOS
 - Allowed high levels of integration at low cost (PCIe1 @2.5Gbps)



The previous slides showed how SerDes became more efficient as integrated circuit process technology developed. This slide maps the speed of new SerDes interfaces proven in silicon since 1995. The source data is papers published at IEEE's ISSCC conference. The green dots indicating bipolar or BiCMOS technology show that by 2005 or so, we'd worked out how to design all the circuits in a SerDes using pure CMOS which is significantly cheaper and enabled higher levels of integration. These advantages are so compelling it's extremely rare to see SerDes designed in anything else. The trend line shows the common interface speed increasing about 6x per decade.

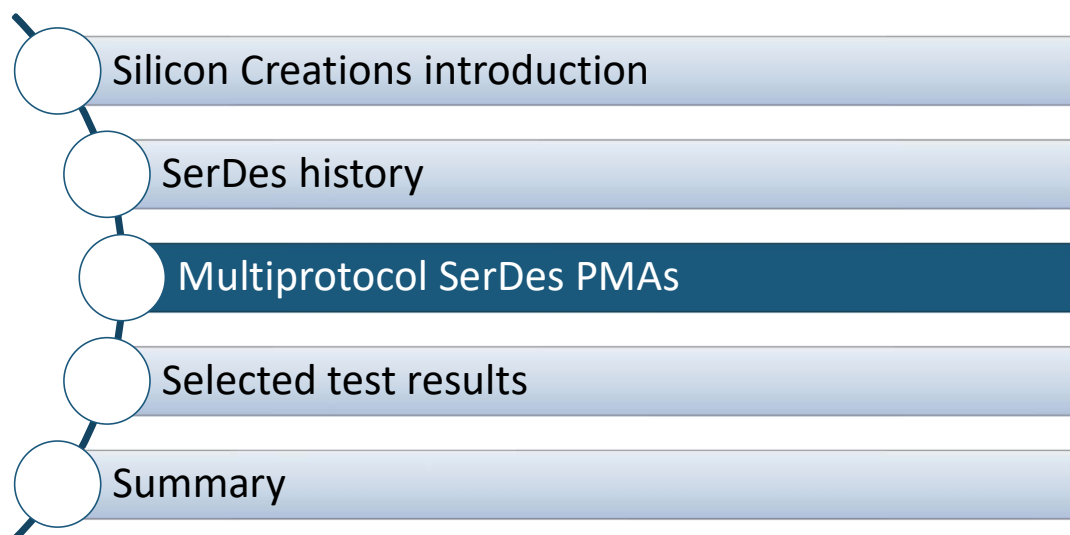
SerDes Standards by year

- Today, links such as PCI Express, HDMI, USB are ubiquitous.
- But it wasn't that way 20 years ago.
- The last 20 years have seen an explosion in the number of serial link applications.
- The path to higher data rates has been foreshadowed by research (such as presented at ISSCC) followed by Networking, Storage, PC/PCIe, and video standards



The last way of looking at the evolution of SerDes is to show how the Serdes standards for various markets have evolved in line rate. The standards for all applications using SerDes have unsurprisingly evolved to follow the 6x/decade trend enabled by the technology. Video displays using cheap cables have been cheaper while storage and networking which are hungry for data have consistently driven to the highest rates.

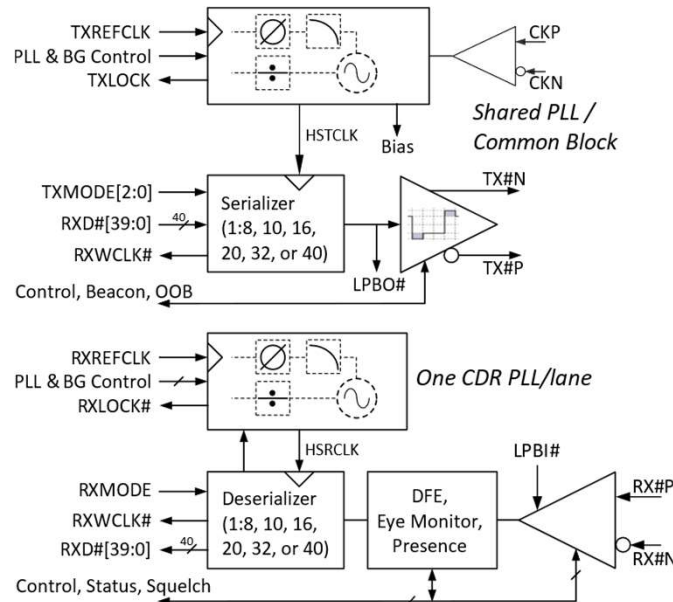
Outline




Now I'll present the structure of our multiprotocol SerDes PMA along with a few key features and a list of the protocols supported.

Overview of PMA

- From IP we made for an FPGA (many protocols)
- Uses ring PLLs → small, low power, flexible
- Xtal clock good enough for PCIe (no clock chip)
- Fast CDR
 - JTOL passes Display Port with margin
 - SDI Pathological pattern
 - Burst mode (OTU/Sync-E)
- Jitter cleaner & SS modulator integrated with Tx PLL



Layout delivered with ESD & RDL



© Silicon Creations, 2020

Galloway, Chen & Cole – Multiprotocol SerDes

14

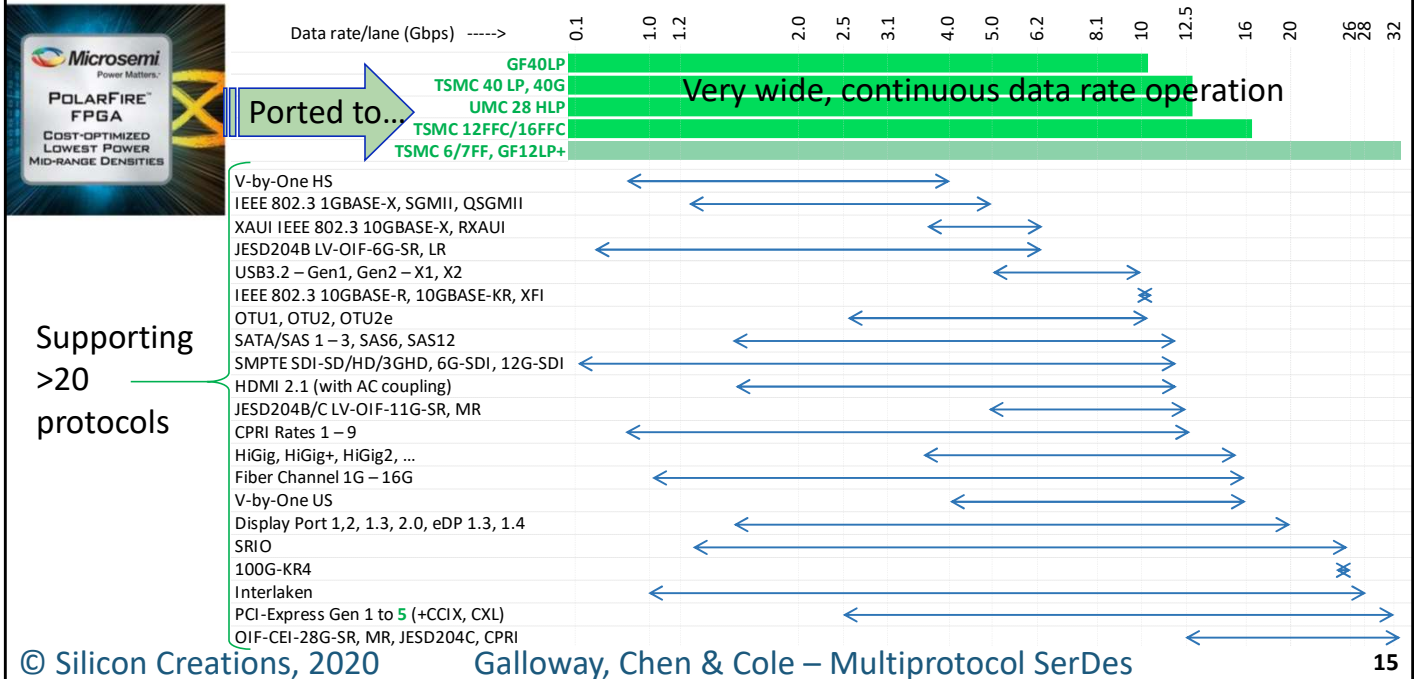
In our architecture we leverage our wide range and efficient ring PLLs.

In the transmitter we share one PLL over multiple transmit lanes. This Tx PLL is good enough to pass PCIe's jitter requirements using just a crystal oscillator which has reduced system cost for some of our customers, and it can also be configured as a jitter cleaner and generate spread spectrum representing even higher system cost savings.

On the receive side we use a ring PLL in a bang-bang loop for CDR.

The PMA includes all the features you'd expect for an enterprise SerDes product such as hot plug presence detection and adaptive equalization for PCIe. The CDR is quite fast but stable and this enables it to pass the difficult DisplayPort Jitter Tolerance tests as well as supporting the SDI pathological pattern as well as burst mode recovery for optical networking. The layout on the right shows the absence of inductors which are not needed in ring PLLs enabling a compact, bump limited layout. Integration of our PMAs is easy because they are self contained, including all supplies PLLs and ESD.

One IP, many protocols



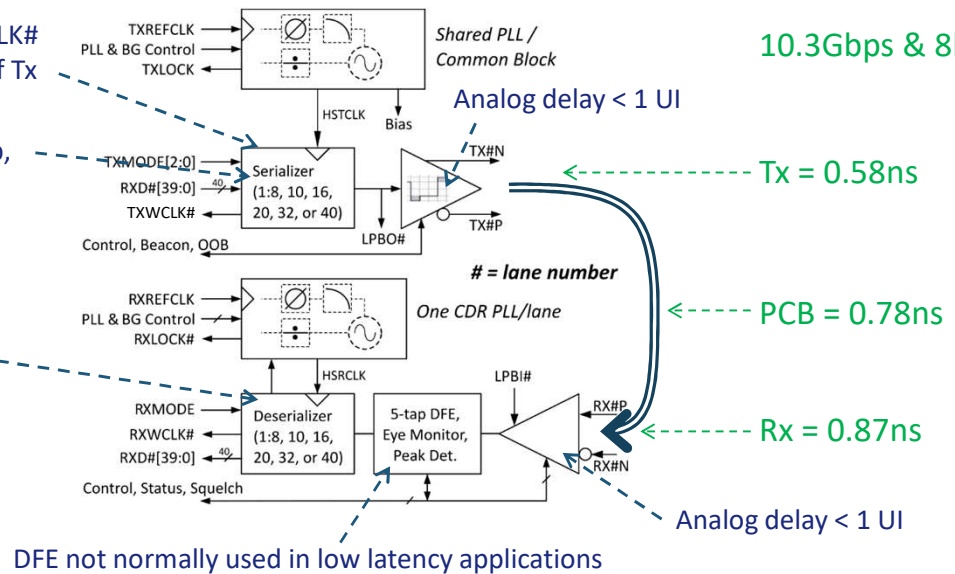
There's a lot of information on this slide. On the horizontal axis of the graph is the data rate supported per lane. At the top of the graph we indicate the process nodes we support with a multiprotocol PMA, and the green bars represent the speed range supported in each process node. Most of these IPs are built for flip-chip operation but with careful package design it's possible to support 10Gbps per lane, and sometimes it makes economic sense. With their heritage in a design we made for an FPGA, our multiprotocol PMAs are known to support a lot of protocols and the lower half lists the main protocols supported by our PMAs along with blue bars to indicate the speed ranges for each protocol. The 40nm IPs support up to 12.5Gbps, so include Ethernet 10G-KR, SAS12 and PCIe3 while the more recent FinFet IPs we've made support up to PCIe4 at 16Gbps or PCIe5 at 32Gbps. Our first PCIe5 PMA is in GLOBALFOUNDRIES 12 LP+ and we have started porting this to TSMC 6FF. We're usually not the first company in each process, but our repeat customers have come to value our comprehensive support and Power, Performance and Area as good as the best in the market.

Low latency option

5 UI from rising edge of TXWCLK#
(parallel load) to first bit out of Tx

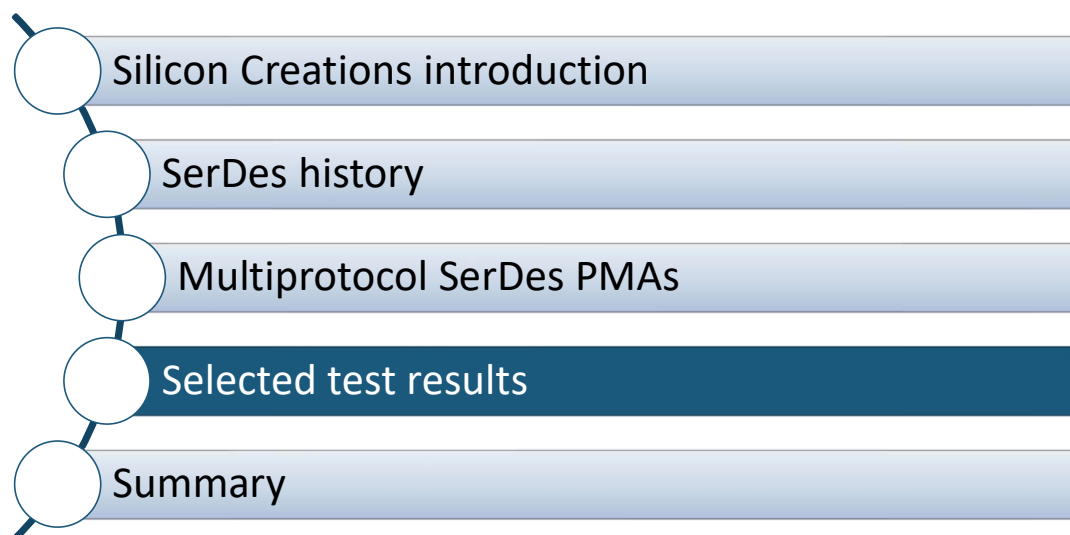
Ser/Des width can be made 8b,
22b, to match FIFO...

8 UI from last bit in to
rising edge of RXWDCLK#
(parallel data ready)



In several process nodes we've also developed a variant of the PMA with extremely low-latency. This slide shows the loop latency from loading a data word to sending it out the Tx pins, getting it all back through the line receiver and deserializing it can be as low as 23b for an 8b wide word. At 10.3Gbps this is less than 2ns in the PMA. This latency is very interesting for high-speed financial trading and critical die to die interfaces with functions like cache coherency.

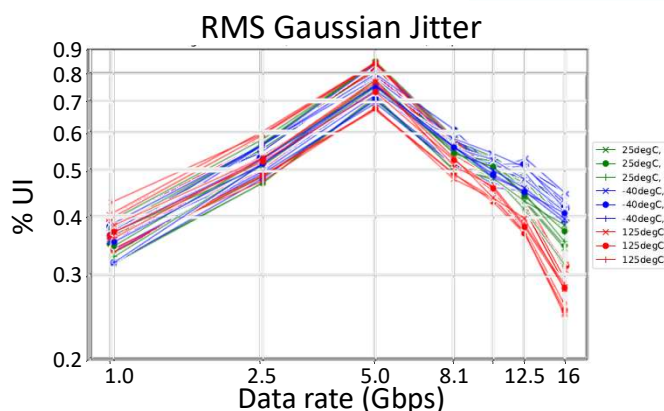
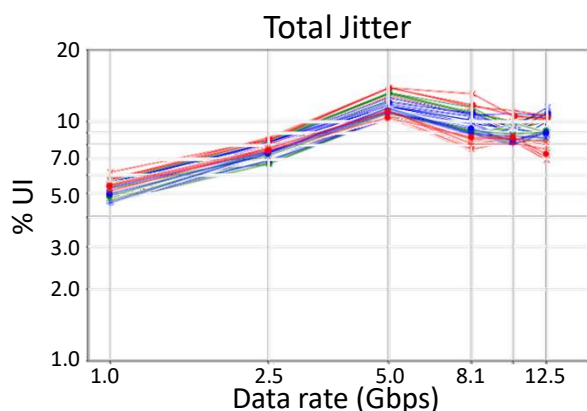

Outline



Now I've introduced you to the structure of the IP, I'd like to show some of the performance measured on our multiprotocol silicon.,

Tx performance (16FFC)

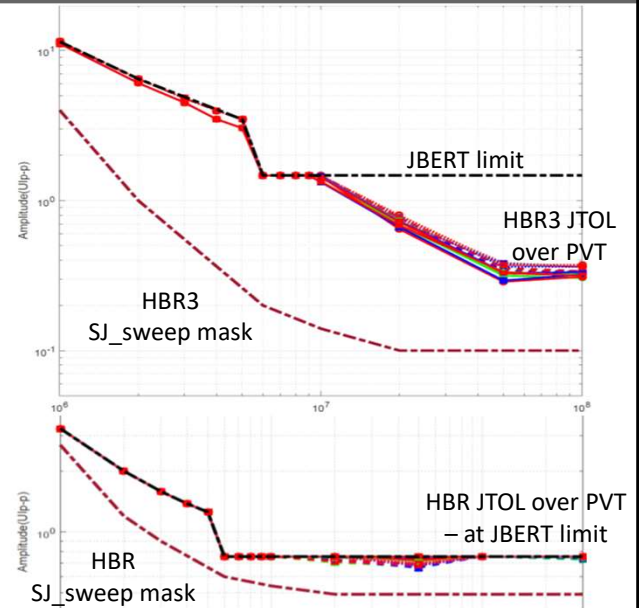
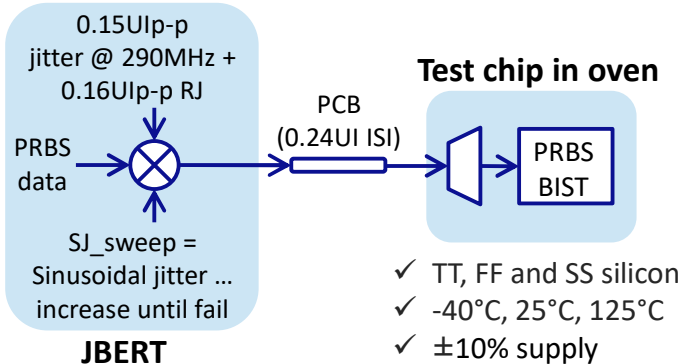
- Overlays of Tx output jitter shows excellent performance from ring-based Tx PLL, very consistent over PVT (5-process corners, $\pm 10\%$ supplies, -40°C to 125°C)
- Package + test socket causes slow rise-time, increased ISI



This slide shows the jitter measured on the output of our TSMC 16FFC test chip. The graphs overlay measurements for 5 process corners, min and max supply and -40°C to 125°C . On the left we present the total jitter combining deterministic jitter like ISI and DCD with peak-peak random jitter for a channel BER of $1\text{E}-12$. For the low range of the PLL the jitter rises but for all rates stays below 15% of the bit width. We stopped at 12.5Gbps because the inexpensive package and socket used did not support 16Gbps very well. But we could measure the random jitter over the whole native rate range of the PLL. As with total jitter this is presented as RMS jitter as a percentage of the bit width. Staying well below 1% UI is a good result and shows the IP supports the most demanding protocols. In both graphs the tight grouping of the lines shows the design is very consistent over PVT which means you can expect to not have problems in the field once you've qualified your chip with our SerDes inside.

DisplayPort JTOL

- Jitter Tolerance tested per DP 1.4a (3.5.2.7.1)
- Margin > 0.2UI for HBR (2.7Gbps), HBR2 (5.4Gbps) & HBR3 (8.1Gbps)

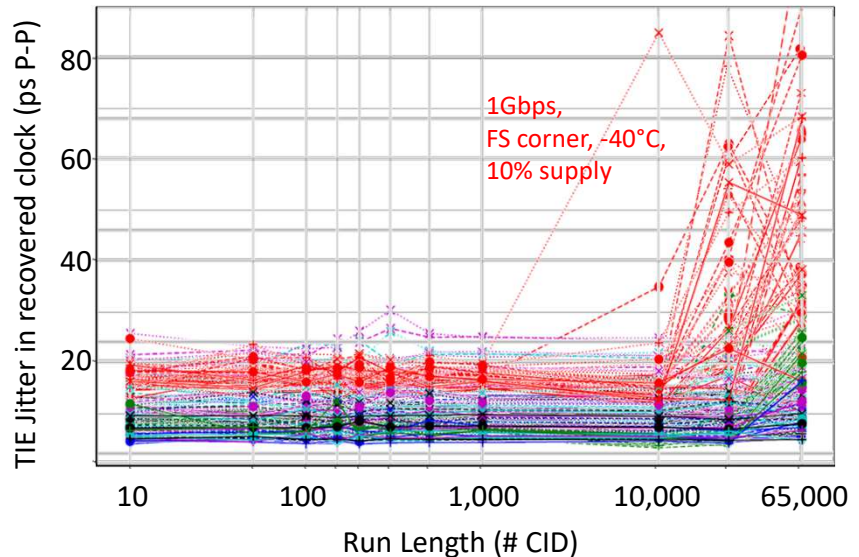


I mentioned DisplayPort earlier. It's difficult to make a CDR that is nimble enough to follow the various deterministic jitter disturbances expected in some DisplayPort systems. The graphic shows how this measurement is defined in the standard, and the overlay graphs on the right show the performance we've measured. For much of the frequency range of sinusoidal jitter added to the data, our CDR tracks up to the modulation limit of the Bit Error Rate Tester. This shows a consistent eye-width margin of about 0.2UI. We've heard that some of the leading SerDes providers can't pass this test at all.

Robust for long run length

Tolerance for very long run lengths (over 5,000 CID):

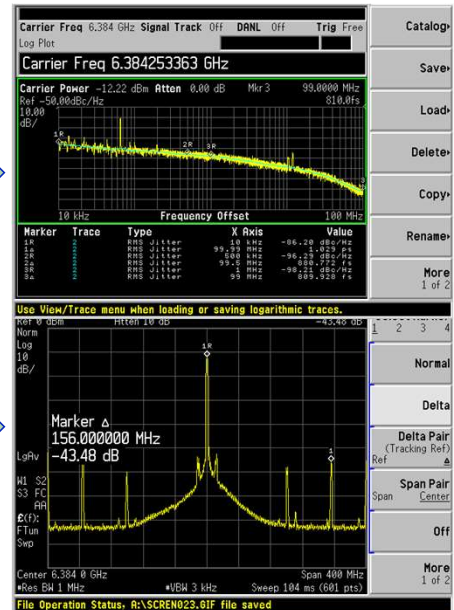

1Gbps to 16Gbps,
Process corners,
-40°C to 125°C,
-10% to +10% supply



Having just shown you how fast our CDR is able to respond, you might be asking if this comes at the expense of tolerance for data patterns with long run lengths. Actually, in our circuit architecture we manage the leakage from our filter capacitors very well. This graph shows the measurement of our CDR's ability to track very long sequences of repeating 0's or 1's in the data. 8b/10b encoding has a maximum run length of 5 CID and even aggressive encoding schemes like PCIe and USB have run lengths at least 100x shorter than our IP can support. With this huge margin our IP is even able to faithfully track data in the SDI pathological pattern.

Jitter Cleaner operation

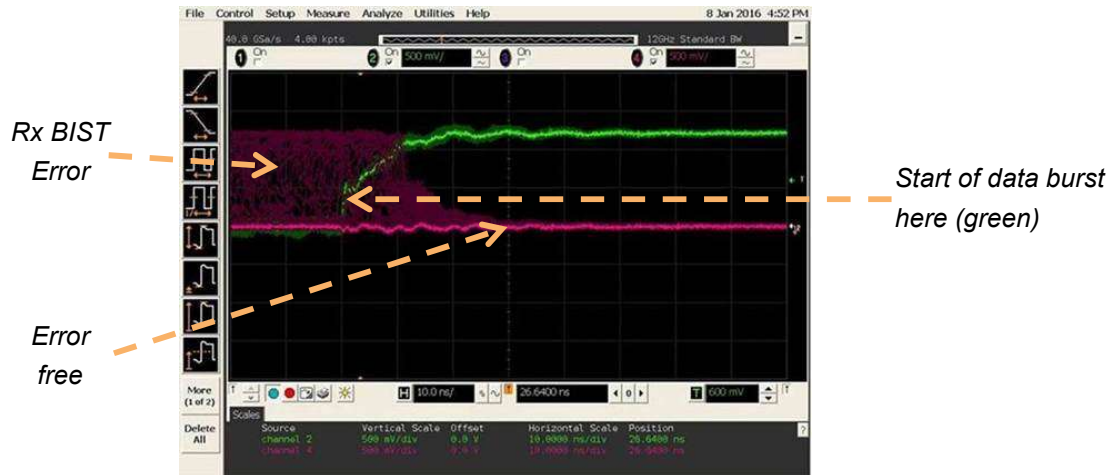
- Tx PLL can generate SSMOD and has a jitter cleaner function (PON, Sync-E)
- Input jitter = gapped clock (~6.4ns jitter), output jitter < 1ps RMS
- Reference clock = 156.25MHz
- Measure eye diagram for Tx
- Spurs from the DPLL are all < 43dBc (<1ps DJ)
- TIE jitter, integrated from 500kHz = 660fs RMS
- Integ. PN from 500kHz = 880fs RMS ($\approx \sqrt{2} * \text{TIE}$)
- Jitter same as normal FRAC mode, good enough for Ethernet 10G_R



I also mentioned that our Tx PLL can be configured as a jitter cleaner. This slide shows a measurement of the performance. We programmed the Tx BIST to send out a repeating 1010 pattern known as D10.2 and use the Tx PLL together with our DPLL programmable loop filter to generate the average of a gapped reference clock that has effectively 6.4ns or so of jitter. With the jitter cleaner bandwidth programmed to about 100Hz this gapped clock which might come from a Synchronous Ethernet framer is cleaned to have comfortably less than 1ps RMS of jitter – good enough for 10G_R Ethernet. Our customers have used this function for optical networking, eCPRI in base-station infrastructure, Synchronous Ethernet, and spread spectrum clock cleaning in Video interfaces.

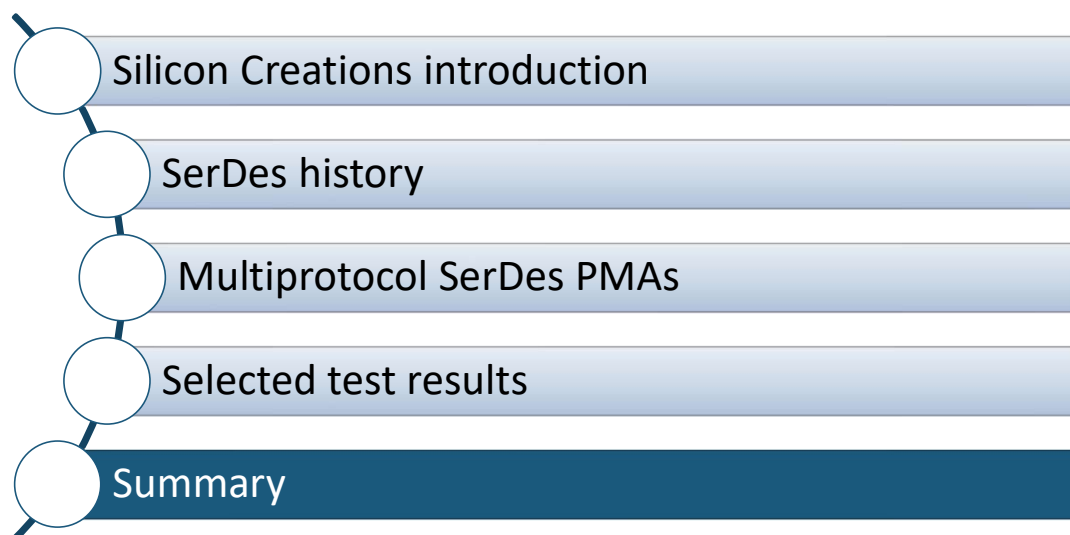
Burst mode CDR (PON / OLT)

- Measured TCDR < 92UI @ 10Gbps & < 48UI @ 2.488Gbps



In passive optical networks the Optical Line Terminal receives serial bit streams from multiple ONU end points. The number of clients and net bandwidth that can be serviced by a single piece of exchange equipment depends on how fast the CDR can switch between uncorrelated ONUs. Our CDR has an extremely fast burst mode. On this slide the rise of the green signal shows the beginning of the data burst and the difficult to see purple the end of the BIST errors. CDR locking is chaotic so many cases are much faster. The longest time measured for XGPON is 92UI and for 2.5Gbps it is only 48UI.

Outline



Thank you for listening to our presentation

Summary



- Silicon Creations has been providing reliable, high performance clocking and SerDes solutions since 2006
- PLLs in production from 180nm to 5nm ... millions of wafers == Low Risk
- Multiprotocol SerDes PMA made for PolarFire FPGA in production has been ported to UMC 28 HLP, TSMC 40LP, 40G, GF40LP and TSMC 12FFC/16FFC (PCIe4)
- Porting now to GF12LP+ & TSMC 6FF (32Gbps)
- Comprehensive feature set means one mixed signal IP is programmable to support many protocols – typical of FPGA
- Ring PLLs used provide excellent performance, continuous data rate range, low area, power

In the short introduction we explained how our low risk offerings of PLLs and SerDes PMAs has been used by well over 200 customers to bring their chips quickly to market in technologies down to 5nm. If you search for our IP on the TSMC website you'll see some of the truly impressive production numbers.

In the past few years our presence in the SerDes space has been growing quickly since we made a multiprotocol PMA for a leading FPGA and ported it to many different process nodes. We highlighted some of the interesting features and great performance from this PMA and its very likely that you will find your target SerDes protocols in the list this PMA supports. We look forward to meeting you and helping you reduce your development risk. Thank you for your time.