

A Robust High Speed Serial PHY Architecture With Feed-Forward Correction Clock and Data Recovery

William Redman-White, *Senior Member, IEEE*, Martin Bugbee, *Member, IEEE*, Steve Dobbs, Xinyan Wu, Richard Balmford, Jonah Nuttgens, Umer Salim Kiani, Richard Clegg, and Gerrit W. den Besten

Abstract—This paper describes a robust architecture for high speed serial links for embedded SoC applications, implemented to satisfy the 1.5 Gb/s and 3 Gb/s Serial-ATA PHY standards. To meet the primary design requirements of a sub-system that is very tolerant of device variability and is easy to port to smaller nanometre CMOS technologies, a minimum of precision analog functions are used. All digital functions are implemented in rail-to-rail CMOS with maximum use of synthesized library cells. A single fixed frequency low-jitter PLL serves the transmit and receive paths in both modes so that tracking and lock time issues are eliminated. A new oversampling CDR with a simple feed-forward error correction scheme is proposed which relaxes the requirements for the analog front-end as well as for the received signal quality. Measurements show that the error corrector can almost double the tolerance to incoming jitter and to DC offsets in the analog front-end. The design occupies less than 0.4 mm² in 90 nm CMOS and consumes 75 mW.

Index Terms—Clocks, CMOS digital integrated circuits, data communications.

I. INTRODUCTION

RIVALING analog-to-digital conversion signal channels, embedded high-speed serial data interfaces are now an essential component of modern system-on-chip (SoC) designs. Interfaces are often needed between processor ICs, connections with display systems, disk drives, other memory functions, etc. A typical multimedia SoC may have many high speed data streams, and without the use of serial ports the pin count, package size and hence cost can be excessive. Power can also be saved with careful system partitioning. Typical data rates are presently in the Gb/s region, and are increasing over time.

Many such standards exist, some visible to the end user such as HDMI, USB2/3 and Ethernet, while others such as PCI-Express and Serial-ATA (SATA) are generally only for internal use [1]–[5]. It is common for these standards to share similar features, and architectural techniques can also thus be shared between designs. To the system designer, such interfaces are just another set of pins, and hence should not occupy a significant

area or consume large amounts of power, particularly if multiple parallel data lanes are required on the same die. Despite their mundane role, there is significant subtlety in realizing a robust and efficient serial data interface architecture.

Traditionally, there are a number of precision analog functions employed in such sub-systems, and porting these blocks for fabrication in different foundries has been challenging and time-consuming compared with the digital sections. Further, with demands for very high data rates, small devices must be used in the critical analog blocks to achieve sufficient bandwidth. As a consequence, mismatch effects are significant and can affect manufacturing yield if there is no strategy employed to calibrate or correct for these imperfections. An additional concern for such SoC designs is the ease of transfer from one generation of CMOS to the next. In these applications, market pressures demand that the core digital functionality is implemented in the most advanced CMOS, and so will be resynthesized as new technology becomes available for design. The ancillary mixed signal block must also be available at the same time or the exploitation of new scaled technology will be delayed. Architectures should therefore be developed where there is little reliance on specific transistor characteristics or a particular supply voltage, and where the functions can be expected to work in a new process with only moderate transistor-level optimization.

In this paper we describe an architecture for a low-power embedded high speed serial data physical layer capable of 1.5 Gb/s or 3 Gb/s operation, where the system and circuit functions have been optimized to minimize the sensitivity to the analog limitations of nanometre CMOS, and to allow porting to new and smaller technologies.

A. System Requirements

Many standards exist depending on the application and technology, (e.g., Ethernet, PCI-e, SATA, HDMI, DisplayPort) but most have strong electrical similarities, often using 50 Ω terminated, balanced lines with defined low signal swing. In the case of SATA [1] this is nominally 0.5 V_{p-p} differential. The data rate is well defined for this standard, but there can be some small degree of frequency modulation added to spread the spectrum of any electromagnetic radiation from the link. In this case, the receiver must be able to cope with 0.53% frequency deviation. The transmitted random and deterministic jitter are defined for a worst-case eye opening, thereby allowing the use of straightforward symbol recovery hardware. Short and long term jitter limits are often specified in the context of tracking PLL bandwidths. The SATA standard employs the commonly used 8b10b symbol encoding [6] which has a run length limit of 5, and hence

Manuscript received January 14, 2009; revised March 10, 2009. Current version published June 24, 2009.

W. Redman-White is with NXP Semiconductors, Southampton SO15 0DJ, U.K., and also with the University of Southampton, U.K. (e-mail: bill.redman-white@nxp.com).

M. Bugbee, S. Dobbs, X. Wu, R. Balmford, J. Nuttgens, U. S. Kiani, and R. Clegg are with NXP Semiconductors, Southampton SO17 1BJ, U.K.

G. W. den Besten is with NXP Semiconductors Research, 5656 AE Eindhoven, The Netherlands.

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/JSSC.2009.2020230

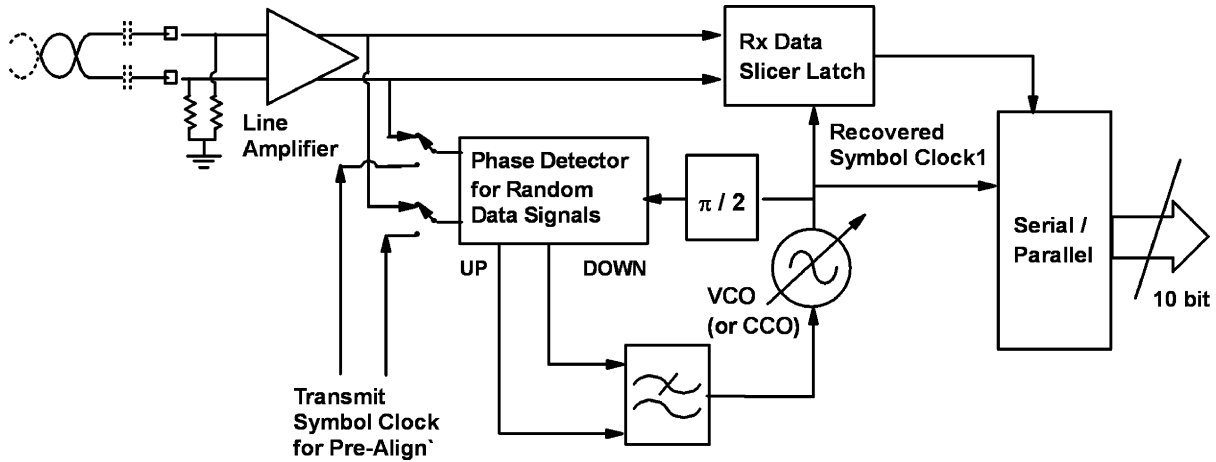


Fig. 1. Classical CDR architecture using tracking PLL.

the receiver must include both clock and data recovery (CDR) functions to allow the symbols (one received data bit) to be re-sampled accurately.

B. CDR Techniques

The CDR is one of the most critical aspects of the PHY architecture, determining how well the receiver can acquire and track the incoming data rate, and how well the symbol eye is sampled. The classical technique is to use a tracking PLL [7], [8] (Fig. 1). The incoming signal is fed to a phase detector able to handle the random transitions in the data. The VCO in the loop has outputs in phase and in quadrature with the input so that the quadrature edge is targeted to align with the center of the data eye for sampling the symbol values. In such an architecture, a PLL is needed with a precise settling time to acquire and track the frequency of data bursts, while the relative phase of the sampler is set by dead reckoning so that phase errors must be well controlled. With several precision analog cells in the system, porting this architecture to another technology requires significant effort. Double sampled bang-bang tracking CDRs provide a significant improvement on phase alignment [9]. However for any of the tracking CDR solutions, the response to jitter in the incoming signal is also affected by the PLL bandwidth and since the receive PLL tracks the incoming data frequency, it typically cannot be used to simultaneously generate the transmit signal.

Rather than try to track the incoming data continuously, an alternative strategy is to take many decision samples for each symbol period, and then use high speed logic to determine the positions of transitions in the data stream. This technique, referred to as blind oversampling [10], [11], has become more attractive with scaled technologies where the fast, dense logic can realize the required algorithms in a very small die area. A major advantage is that the PLL controlling the receive path sampling does not need to be exactly synchronous with the incoming data, eliminating start-up and tracking issues. As a result, the same PLL can also be used to control the transmitter at the same time. It is a development of this approach that is described in this paper.

II. RECEIVER FRONT-END ARCHITECTURE

The basic architecture used is shown in Fig. 2. The first parameter to be fixed in the design is the number of decision samples to be taken for each symbol period, the oversampling ratio (OSR). A lower OSR requires less samplers and less hardware, and may be necessary if operating close to the limits of the silicon technology [12]. However, there is less redundancy in the sampling process, and increases difficulty in extracting the data transitions if signal quality is poor. A higher OSR enables more advanced clock and data recovery, but requires great care in the generation of many sampling phases with very small time separations, and can increase the scale of the logic considerably along with the digital dynamic power consumption [13].

A. Signal Input Path

The received signal is amplified before the digital processing by the input amplifier (Fig. 3). Adaptive line terminations are present on the die, and the operating value is adjusted to compensate for fabrication tolerances by comparing the value of a matched on-chip resistor with an external reference resistor. The amplifier uses a differential grounded gate structure, with outputs at a signal level sufficiently large for direct sampling by a fast rail-to-rail CMOS latch. To achieve a high bandwidth in the amplifier and in the succeeding fast latch, the designs necessarily use very small MOS transistors. As a consequence, the variance in the input referred offset is not negligible. To ensure that this does not lead to unacceptable errors in the received data, some strategy is needed to calibrate the error, or to compensate for imperfections. In this design the latter strategy is adopted in the CDR structure.

B. PLL

In this architecture $5\times$ oversampling is used, giving a good compromise between complexity and robustness. The single fixed-frequency PLL runs at 1.5 GHz with its current-controlled oscillator (CCO) delivering 10 output phases in both 1.5 Gb/s and 3 Gb/s modes. The PLL sampling phases are separated by only 67 ps, and the p-p jitter should thus be significantly less than this. The PLL uses a 25 MHz reference with the loop bandwidth is made as high as possible to minimize jitter

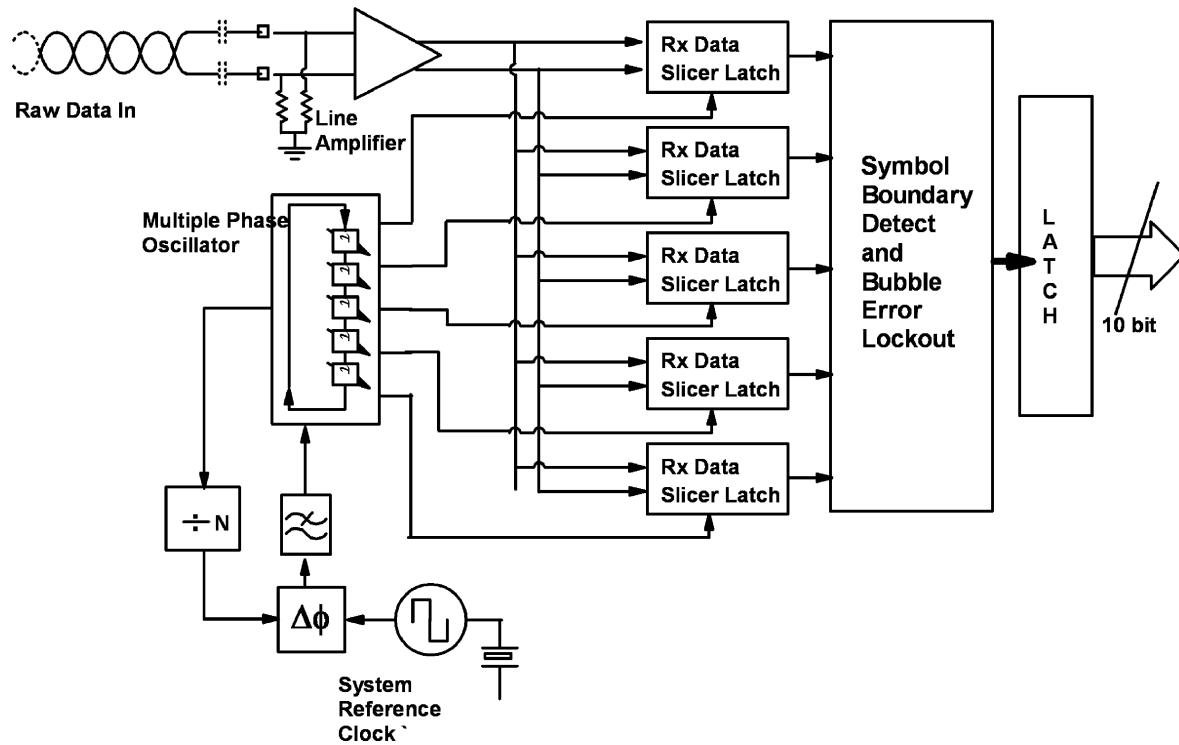


Fig. 2. Blind oversampling CDR architecture.

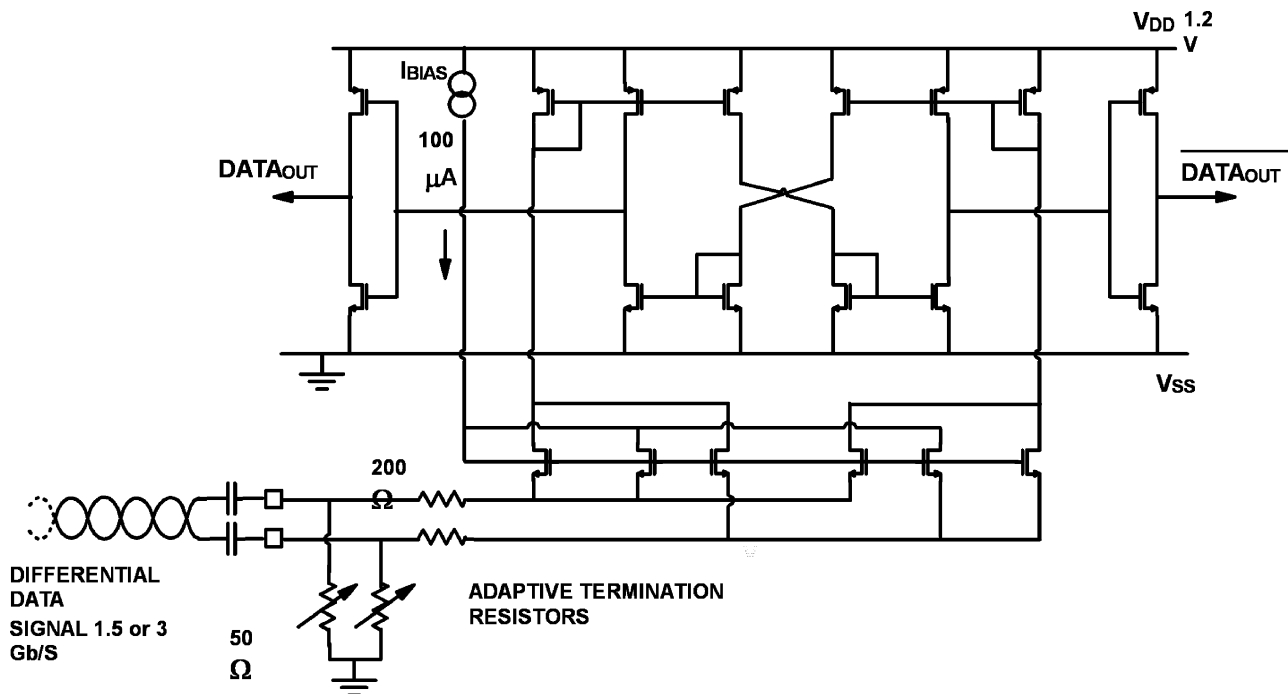


Fig. 3. Receive front-end line amplifier.

from the CCO (Fig. 4). The time resolution between stages is a function of transistor matching parameters, and hence transistor sizes are significantly larger than minimum dimensions; this in turn implies an increased CCO operating current for a given frequency and jitter budget. Much attention is also paid to the CCO layout to avoid systematic errors in the temporal spacing, particularly in the output phase routing to ensure balance in the parasitic loads on each node. The outputs from the CCO are

level shifted to an internal low noise supply and then again to the digital core supply to ensure the fastest possible edge speed before entering this noisy supply domain.

The other parts of the PLL are quite conventional with dead-band elimination in the phase-frequency detector and low-noise locally-generated supplies for the digital divider as well as the charge-pump. Since the multi-phase oscillator is current controlled, the loop filter voltage output is converted to a current by

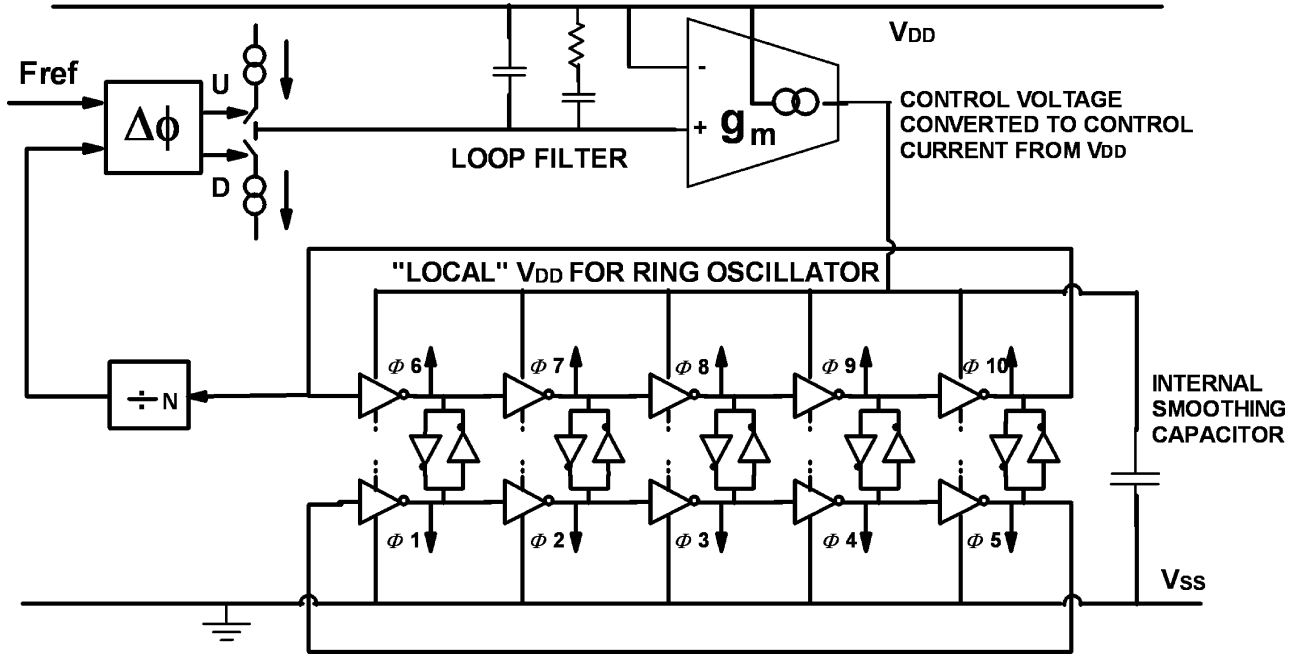


Fig. 4. PLL architecture.

means of a transconductor with a high output impedance which provides good power supply rejection. There is also some local decoupling of the CCO to remove high frequency noise components; this creates an additional pole in the loop, so care must be taken to avoid stability problems. The whole design is made in a triple well process which allows the use of deep n-well for isolation purposes and the main analog supply for the PLL comes from a dedicated device pin.

C. Sampler and Serial to Parallel Conversion

Each phase from the CCO controls the sampling of the amplified signal into one of 10 latches, each connected to the output of the input amplifier. The outputs of these latches are stable for most of one complete period of the CCO, but all outputs cannot be resampled at one instant. Hence, the data are first realigned in two 5-bit blocks by two opposite phases of the CCO, and then latched with another single CCO phase. These 10-bit blocks are then pipelined through a short shift register to allow four consecutive blocks of 10 samples to be assembled for processing in the CDR. For the 3 Gb/s mode, the raw samples are taken directly in 40-bit blocks at a clock rate of 375 MHz (Fig. 5). In the 1.5 Gb/s mode the raw sample stream in this shift register is decimated by 2 and the 40-bit output blocks are taken at 187.5 MHz (Fig. 6). This decimation is almost the only mode switching required in the receive path. Normal CMOS library logic can handle both modes, and so the PLL is not required to switch frequency or to track incoming signal variations, simplifying the design and giving more freedom to optimize for the jitter target.

III. CDR ARCHITECTURE

In previous versions designed by the authors the CDR had been quite simple where the design was not intended for significant reuse. The objective in the design described here is to ensure that analog imperfections due to manufacturing tolerances

such as amplifier offset and internal as well as external jitter should be allowed for in the digital algorithms with the goal of greater robustness and higher yield. The strategy is to use simple, pragmatic error tolerant algorithms with low hardware overhead.

A. Synchronization and Symbol Extraction Strategies

The main tasks in a blind oversampling CDR are to retrieve the transmitted bit values from the stream of raw samples by using the position of signal transitions, and furthermore to determine where the symbol boundaries are located. From this information, the data payload can be recovered. Since the sampling and the incoming data are not synchronous, the definition of the symbol boundaries is only approximate, and some elasticity must be built into the data recovery. The simplest method of extracting bit transitions from the raw samples is to use an EXOR function on adjacent samples, and look for non-zero outputs. Over a defined sample block length, the EXOR '1' values are expected to be present in positions at multiples of 5 samples, but the starting reference position is unknown. For each of 5 possible reference sample positions the number of EXOR '1' values appearing every 5th sample are counted; there should be a very clear winner when the totals are compared. This gives the CDR logic the positions of the symbol boundaries in this block of samples, and hence the data may be recovered (as shown in Fig. 7, top).

This approach works well if the received signal and receiver sampling function are fairly ideal, but can be significantly less reliable if there are imperfections in the incoming signal (e.g., jitter) or in the receiver hardware (e.g., offset).

B. Window Algorithm

If the receiver is subject to jitter in the received signal and sampler timing, input noise or offset due to small, poorly

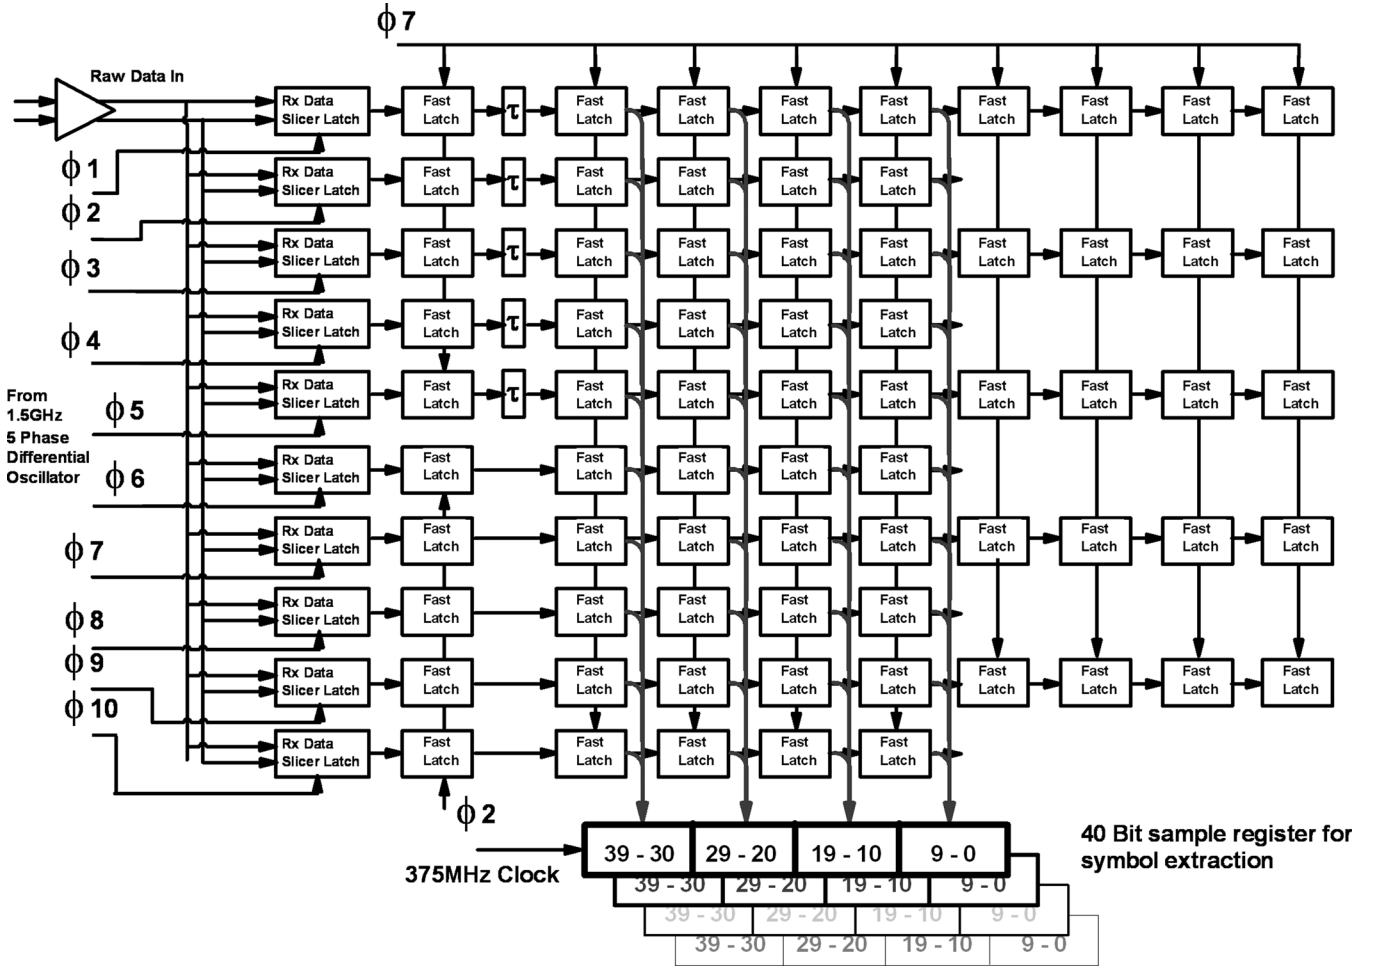


Fig. 5. Serial-to-parallel sample rate reduction, 3 Gb/s mode.

matched transistors in the pre-amplifier and samplers, sample errors can arise, making it difficult to determine the actual transition moments and therefore the bit values. Amplifier offset and jitter could, in the worst case, corrupt two consecutive samples, one either side of the ideal transition boundary instant. We can model these effects simplistically by a variation in the decision threshold (Fig. 7, center and bottom traces). However, if the signal is not completely lost, in a system with a $5 \times$ OSR the three center samples are generally reliable. We can thus define a simple window function to look for differences between valid symbols separated by two samples, implying finding strings of three samples having the same sign, followed by strings of three samples of the opposite sign. This is implemented using a simple AND-OR function. As with the EXOR approach, it is then necessary to locate the positions of the symbol boundaries at multiples of 5 samples from the occurrences of the transitions in the sample stream.

Note that with this strategy, an ideal signal with perfect sampling leads to results which are ambiguous, showing where in the sample stream the transitions could possibly be, but not identifying the positions exactly (Fig. 8). However, if the raw sample data are affected by jitter and offsets, the averaged window function results readily converges to the correct position. This algorithm is robust against in the presence of bubble errors due to noise (Fig. 9) and errors due to DC offsets (Fig. 10). From

the foregoing it is a reasonable inference that a combination of these schemes could be beneficial in recovering data from raw samples of varying quality. Switching between the detection modes is cumbersome, but a voting system which combines results from both can be readily implemented to achieve more robust data recovery.

C. Transition Detection and Data Recovery

To allow for the slippage due to non-synchronous sampling, as well as for jitter, offset and noise, the bit transition timing must be estimated from a sample buffer long enough such that there are always sufficient transitions in the samples to make a reliable decision. In this design a buffer of effectively 200 raw samples is used, guaranteeing that at least 8 transitions are present. Results from all 200 window edge detections are multiplied by a weighting factor and combined with the 200 EXOR results, also multiplied by a second weighting factor, before then being summed in five groups (since there is $5 \times$ oversampling, see Fig. 11). The group with the largest vote sum is deemed to be the sample index modulo 5 which represents the best estimate of the symbol edges (Fig. 12). It is now possible to assume that the samples in between the transitions represent valid data values. However, to ensure that the transition timing estimate is only used for a block of samples in which there cannot have been

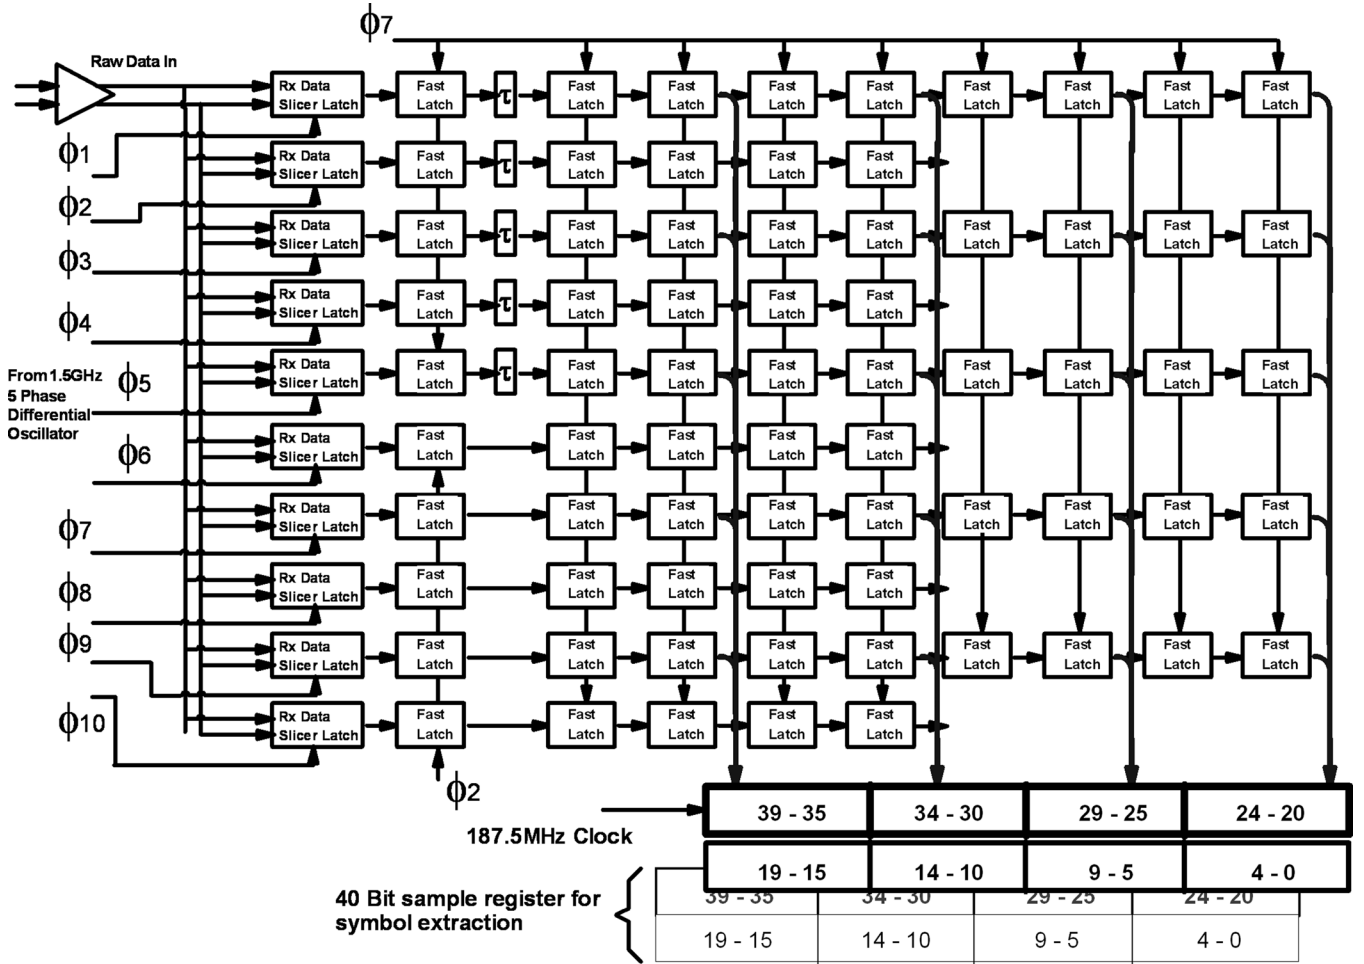


Fig. 6. Decimated serial-to-parallel sample rate reduction, 1.5 Gb/s mode.

significant slippage in the true clock, only the central 40 samples (Fig. 11) of the 200 tested are used. (In reality, a few extra samples are also tested to allow for slippage between transmit and receive clock frequencies). The preceding and succeeding 80 samples are only used as run-in and run-out data for the algorithm. At the end of this evaluation, a new block of 40 raw samples is loaded to one end of the 200 sample buffer and the oldest 40 samples are discarded. The central three samples from each group of five are used to determine the symbol value by another majority vote evaluation.

D. Asynchronous Clock Slippage

Because the receiver clock is fixed, the algorithms must handle variations in the received signal. There is an allowed tolerance in the nominal clock frequencies, as well as the spread spectrum deviation. Altogether the differences can amount to as much as 0.53%. If the single PLL is used for a transmit path with spread spectrum capability enabled, then the receiver must also work with this additional frequency difference. This clock slip is handled by extracting more data bits than are normally needed at each step of the CDR algorithm. Because the starting point of the data extraction varies, the number of samples in the center of the 200-samples buffer that must have the symbol values determined is actually 50, corresponding to 9 data bits. When the transmit baud rate is exactly 1/5 of the receiver

sample rate, the 9th bit is not needed, and only the first 8 are output. The 9th bit is effectively overwritten in the evaluation of the next 40 sample buffer. In the case that the transmitted baud rate is higher than 1/5 of the receiver sample rate, the symbol boundary index (the calculated sample position where a bit transition occurs) gradually advances through the buffer until it wraps around and an extra bit is periodically generated, so that 9 bits are output (Fig. 13). Alternatively, if the transmitted baud rate is lower than 1/5 of the receiver sample rate, the symbol boundary gradually moves back through the buffer until it wraps around. In this case, one of the bits is effectively recovered twice, so that one bit must be discarded and only 7 are output. A buffer with flag signals controls the transfer rate of these data to the link layer with sufficient elasticity in the buffers to allow for the clock slip budget.

E. CDR Hardware

Power and area are important attributes in this design, and so considerable effort has been invested in achieving an efficient implementation. A preliminary design used a direct mapping of the algorithm into standard library logic using several multiplier cells for the weighting process, but the area and operating speed were unsatisfactory. Two main strategies were employed to improve the design. Firstly, all multipliers were removed and simple left and right shift operations used, nearly halving the




Fig. 7. Use of EXOR for signal sign transition detection and clock and data recovery. Top: ideal signal and receiver. Centre, effect of DC offset in receiver front-end. Bottom: effect of jitter or noise in receiver front end.




Fig. 8. Use of Window transition detection approach on ideal signal.

area of the combinatorial logic, and reducing the logic depth in the critical paths of the design. As a consequence of this, the range of vote weights was constrained into powers of 2, and the final summation of the weighted votes is not normalized in any way. However, these are not serious limitations as the maximum vote summation value obtained is still valid. In the implemented logic the voting weights could range from 1:0, 8:1, 4:1, to 1:8 and 0:1.

The second change to the design was to introduce extensive pipelining. Only 40 samples are actually processed at one time (out of a possible 50 to allow for clock slippage) to give the

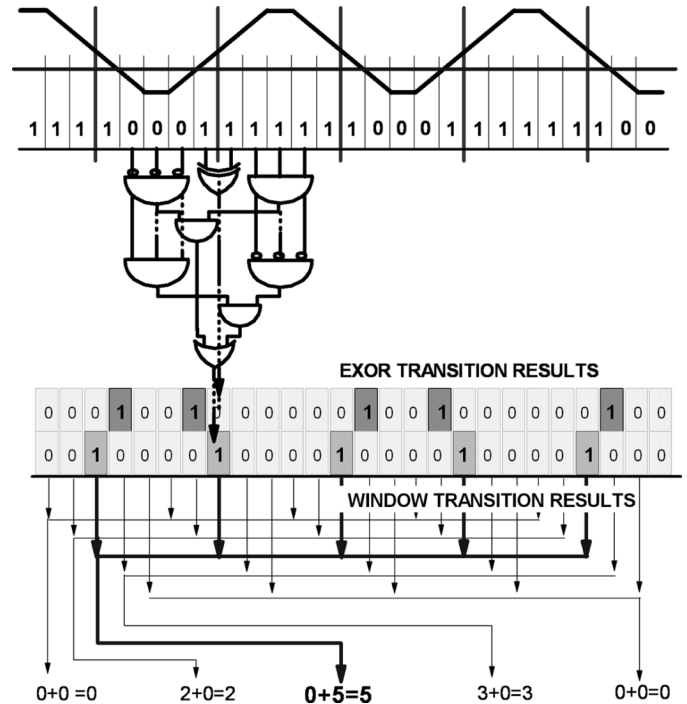


Fig. 9. EXOR and Window transition detection combined for case of DC offset in receiver front end.

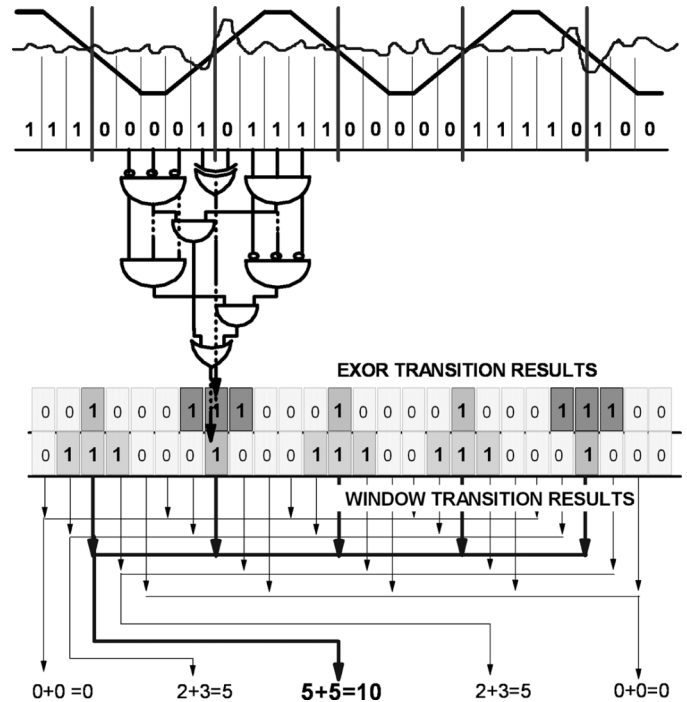


Fig. 10. EXOR and Window transition detection combined for case of jitter or noise in receiver front end.

EXOR and Window votes, and the values are held until the votes derived from all five 40 sample blocks are available, whence the summation and evaluation can be performed. At this point the 40 raw samples corresponding to the central block (plus the extra 10 to allow for clock slippage) are still present in the pipeline and the symbol data values can be extracted from them.

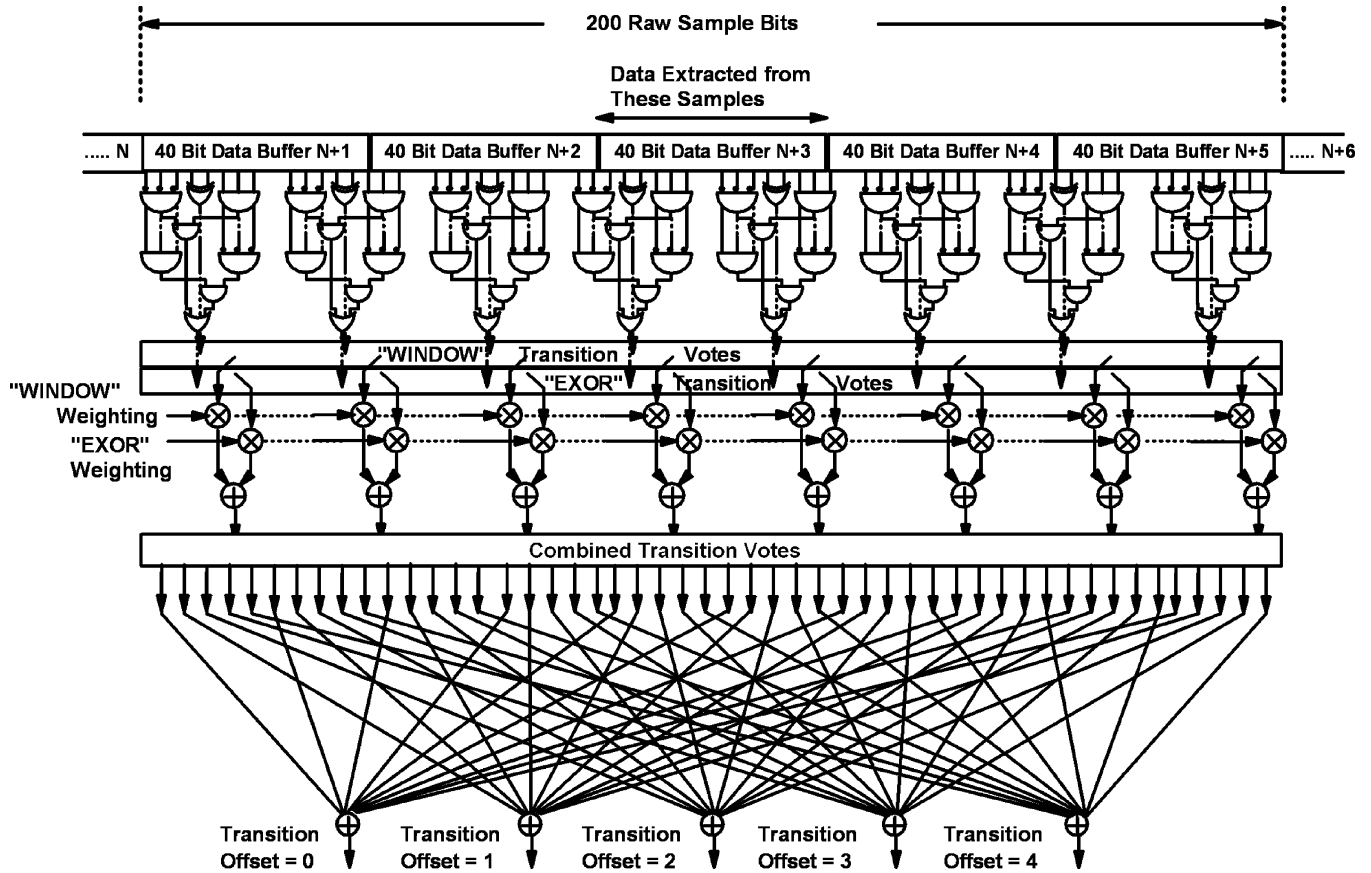


Fig. 11. Weighting of transition detector outputs from 200 sample buffer to derive transition reference position in central block of samples.

IV. TRANSMIT ARCHITECTURE

The top level architecture follows broadly conventional structure. Eight-bit-wide parallel data are encoded into 10 bits and delivered from the link layer to the PHY transmit section at a moderate clock speed [at 75 MHz, for 150 MHz DDR (Gen1)]. The low jitter receiver PLL oscillator clock signal is reused to drive the parallel to serial conversion function and the transition timing in the line driver. EVEN and ODD data bits are parallel loaded into a shift register and then clocked out serially at half the data rate. The serial data are then interleaved and retimed using clock edges fed directly from the CCO (Fig. 14).

The transmit line driver uses a simple differential current steering scheme. The output current is derived from a current source referred to an external close tolerance resistor. The line driver differential pair is fed from a current source with a programmable replica bias scheme that allows the output amplitude to be varied by the configuration software, while compensating for variations in the individual transistors' operating conditions over temperature etc. The gate drive to the differential pair is configured to give make-before-break switching, thereby ensuring that the tail current is never turned off and the common mode voltage remains constant. Slew-rate control is also included to meet the Gen. 1 (1.5 Gb/s) and Gen. 2 (3 Gb/s) targets [1] and to ensure precise and symmetrical eye crossing points [1]. As in the receive path, the output terminations are adaptively set with an external resistor and a replica circuit (Fig. 15).

V. IMPLEMENTATION AND RESULTS

The layout of the complete PHY is shown in Fig. 16. The total area is less than 0.4 mm² in a standard triple-well 90 nm CMOS, including a significant decoupling capacitance. The PHY characteristics are summarized in Table I. Note that the circuit is embedded in a large multimedia IC for which Gen 1 operation and compliance has been verified; hence the measurements are taken in this environment, not as an isolated test chip. The design is also fully functional at 3 Gs/s (Gen 2), but the product has not been fully qualified at this rate at the time of writing. In the present application the MAC is locked in Gen 1 mode, and some of the Gen 2 features of the PHY are not accessible. Some basic parameters of the PHY can be observed in a test mode whereby data can be directed through to the transmit buffer for testing the jitter of the PLL, but the final retiming logic required for Gen 2 operation as well as the slew rate options are bypassed in this mode. Nonetheless, some low-level testing is possible in Gen 2 and the results are also presented.

A. Basic Performance

The PHY performance has been verified using a TDS6808B (32M) oscilloscope running the TDSRT-EYE (Serial ATA) compliance package [14].

The measured transmit and PLL performance is summarized in Table II. A special test mode allows direct measurement of the PLL behavior at the pins with measured 1- σ jitter less than 3 ps for one clock period unit interval (UI) and less than 8.3 ps 1- σ

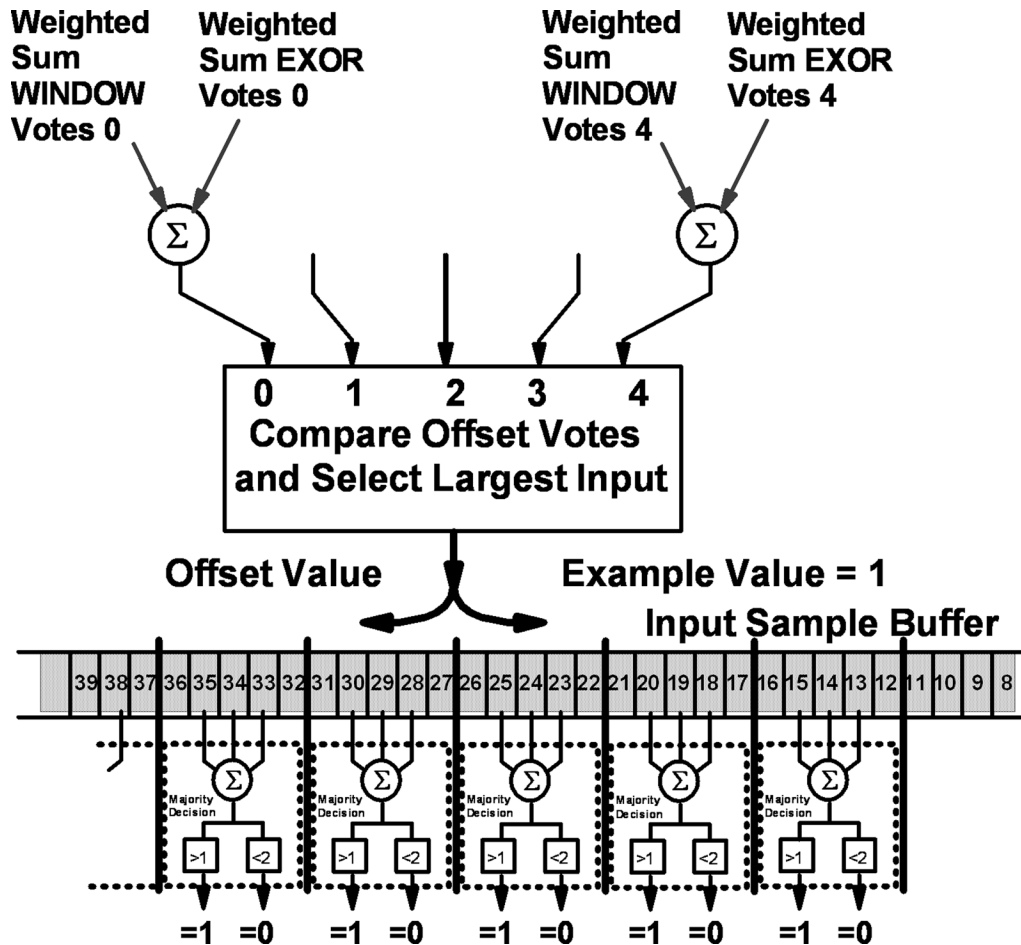


Fig. 12. Evaluation of maximum vote to set transition reference position in buffer.

TABLE I
PHY SUMMARY

Power	Total	Analogue	Digital
	75mW	65mW	10mW
Die Area	0.4 mm ²		
Tech.	90nm CMOS		
Power Supplies	3.3V and 1.2V		

jitter for 250 UI. The Gen1 (1.5 Gb/s) transmit eye pattern shows less than 170 ps peak-to-peak jitter (TJ) at 250 UI (Fig. 17) and the Gen2i (3.0 Gb/s) transmit eye pattern shows less than 100 ps peak to peak jitter at 500 UI (Fig. 18). All eye diagrams were generated using the composite pattern as defined in the SATA specification [1]. Both Gen1 and Gen2 operation show considerably lower deterministic jitter (DJ) than the specification limit.

The receiver front end shows excellent sensitivity, being able to recover clean signals down to 100 mV p-p and can handle 100 mV p-p noise on a 200 mV p-p signal. There are no lock time issues, and the system easily follows far more than the 0.53% frequency differences specified, such that no special spread spectrum tracking is needed. Full compliance with SATA Gen1 requirements has been verified.

B. Error Tolerance of Voting System

Tests were also performed to establish the improvements due to the new EXOR/Window CDR algorithm.

The resilience of the system in the presence of analog offsets in the input amplifier was tested by adding a differential DC voltage to the amplifier inputs via external resistors. These resistors were made sufficiently high that the static impact on the line termination was not significant. The value of the offset was measured with no signal applied. As a rigorous bit error rate measurement was not possible in the SoC, a signal was then applied from a hard disk drive (325 mV p-p as measured at the connector with 2M long cables) and the system condition monitored to establish the offset value at which loss of synchronization occurred. The weighting factors of the EXOR and Window detector paths were then changed and the test repeated. Fig. 19 shows the tolerance to the offset as a function of the detector weighting factors. When only the EXOR detector is operating, as in a conventional oversampling CDR, the link fails with around 70 mV applied DC offset. As the contribution of the Window detector output is increased beyond about 67%, the offset tolerance increases by a factor of 2 to around 140 mV. Increasing the Window detector contribution further to the point where there is no EXOR contribution will eventually lead to the link failing if the signal is ideal, as predicted by MATLAB simulations. However, if there is a significant DC offset present, the link will still function.

A similar test was undertaken to establish the tolerance to input jitter. IDLE and SYNC/ALIGN patterns were sent from

Normal Buffer Shift Index Change Without Wrapping




Fig. 13. Data extraction strategy with clock slippage.




Fig. 14. Transmit path parallel to serial conversion architecture.

a Tektronics 5334 Data Timing Generator, with an amplitude of 340 mV p-p. Gaussian jitter was applied to both edges of the data, and the system monitored as before to establish the jitter level required for loss of sync (the fail condition was determined from the average of several fixed duration tests) at each of the possible weighting factors for the EXOR and Window detectors. Fig. 20 shows how the jitter tolerance varies. With only the EXOR detector operating, the system can just tolerate 0.24 UI jitter. As the contribution of the Window detector is increased there is again a sharp improve-

ment by nearly a factor of 2 when the Window detector has more than 67% weighting. The jitter measurements show the same trends as with the DC offset tests, except that with the Window detector contributing 100% to the CDR, the system fails completely.

These results confirm the choice of the default weighting factors, derived from MATLAB simulations as being EXOR:Window at 1:2. The tolerance to offsets and jitter are shown to be nearly doubled by the use of the EXOR and Window detectors with a weighted voting system.

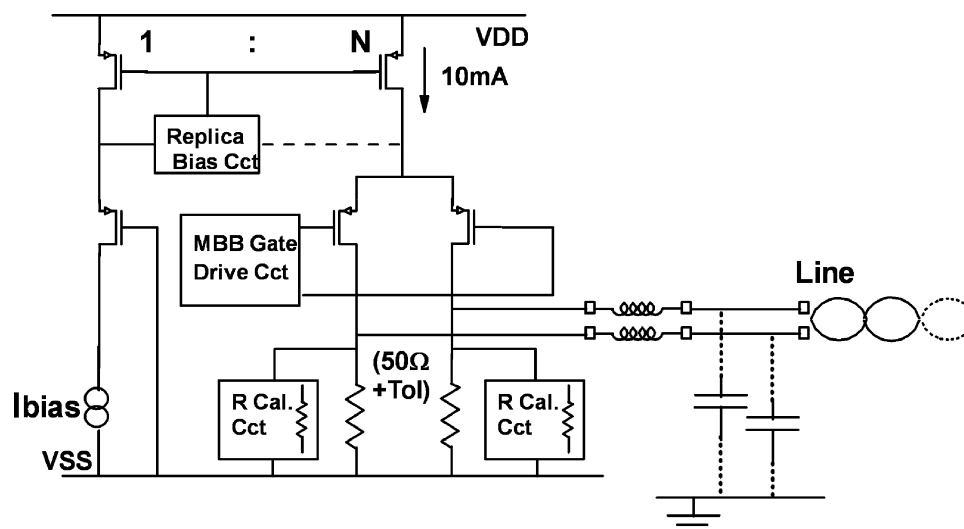


Fig. 15. Transmit buffer with calibrated terminations.

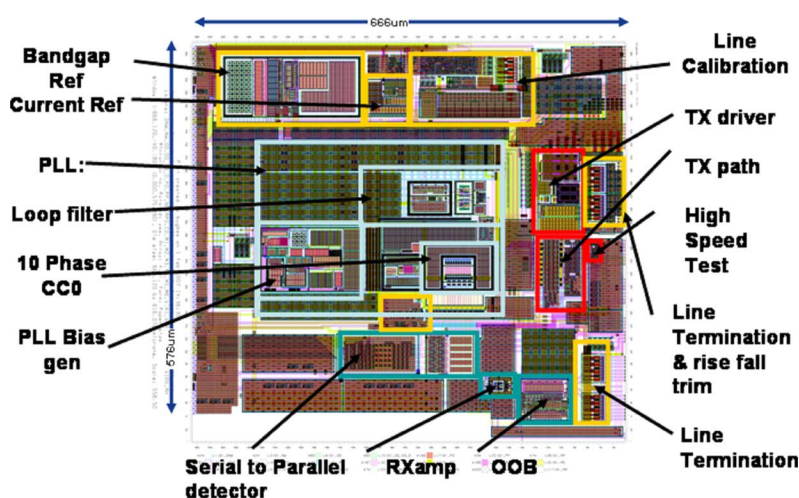


Fig. 16. PHY die layout.

TABLE II
JITTER RESULTS

	Measured		Specification	
	DJ	TJ	DJ	TJ
GEN1, Transmit				
5 UI	< 35ps	< 100ps	117ps	237ps
250 UI	< 35ps	< 170ps	147ps	313ps
GEN2, Transmit				
10 UI	< 30ps	< 65ps	57ps	100ps
500 UI	< 30ps	< 100ps	63ps	123ps
PLL, 1.5GHz				
	1UI	5UI	250UI	
σ	< 3.2ps	< 4.3ps	< 8.0ps	
p-p	< 27.5ps	< 33.9ps	< 76.8ps	

VI. CONCLUSION

A robust high-speed serial data PHY has been developed for the SATA Gen 1/2 specifications, with features applicable to a wider range of similar standards. The architecture uses a minimum of precision analog blocks for yield and process portability. A single fixed frequency low-jitter PLL is used for both transmit and receive paths in both modes, saving power




Fig. 17. Transmit 1.5 Gb/s jitter eye plot at 250 UI.

and eliminating locking problems. Optimization of conventional CMOS digital circuitry and extensive pipelining




Fig. 18. Transmit 3.0 Gb/s jitter eye plot at 500 UI.




Fig. 19. Receiver behavior as a function of CDR vote setting with externally applied DC offset at input to line amplifier. Curves show offset value at loss of SYNC with programming of different vote weights in CDR.





Fig. 20. Receiver behavior as a function of CDR vote setting with Gaussian jitter added to signal. Curves show p-p jitter values at threshold of loss of SYNC with programming of different vote weights in CDR.

is used to achieve small die area and low power consumption. A new CDR architecture has been demonstrated with enhanced tolerance to imperfections in the system. The use of weighted voting to combine the results of EXOR and Window transition detectors shows that the immunity to DC offsets and jitter is improved by almost a factor of two with little overhead in hardware and power.

REFERENCES

- [1] Serial ATA 2.6 Specification. Feb. 2007 [Online]. Available: www.sata-io.org.
- [2] USB 2.0 Standard, 2000 [Online]. Available: www.usb.org
- [3] PCI Express specifications. PCI-SIG [Online]. Available: www.pcisig.com/specifications/pciexpress
- [4] High-Definition Multimedia Interface (HDMI) specification. [Online]. Available: www.hdmi.org
- [5] DisplayPort, [Online]. Available: www.displayport.org
- [6] A. X. Widmer and P. A. Franaszek, "A DC-balanced, partitioned-block, 8B/10B transmission code," *IBM J. Res. Devel.*, vol. 27, no. 5, pp. 440–451, Sep. 1983.
- [7] R. Walker *et al.*, "A 2.488 Gb/s Si-bipolar clock and data recovery IC with robust loss of signal detection," in *IEEE ISSCC Dig. Tech. Papers*, 1997, pp. 246–247.
- [8] Y. M. Greshishchev and P. Schvan, "SiGe clock and data recovery IC with linear-type PLL for 10-Gb/s SONET application," *IEEE J. Solid-State Circuits*, vol. 35, no. 9, pp. 1353–1359, Sep. 2000.
- [9] A. Fiedler *et al.*, "A 1.0625 Gbps transceiver with $2 \times$ oversampling and transmit signal pre-emphasis," in *IEEE ISSCC Dig. Tech. Papers*, 1997, pp. 238–239.
- [10] K. Lee *et al.*, "A CMOS serial link for fully duplexed data communication," *IEEE J. Solid-State Circuits*, vol. 30, no. 4, pp. 353–364, Apr. 1995.
- [11] C.-K. K. Yang, R. Farjad-Rad, and M. A. Horowitz, "A 0.5- μ m CMOS 4.0-Gbit/s serial link transceiver with data recovery using oversampling," *IEEE J. Solid-State Circuits*, vol. 33, no. 5, pp. 713–722, May 1998.
- [12] G. W. den Besten, "The USB 2.0 physical layer: Standard and implementation," in *Analog Circuit Design*. Boston, MA: Kluwer, 2003, pt. III, pp. 359–378.
- [13] G. den Besten, F. Gerfers, J. Conder, A. Kollmann, and P. Petkov, "A 200 Mb/s-2 Gb/s oversampling RX with digitally self-adapting equalizer in 0.18 μ m CMOS technology," in *Symp. VLSI Circuits Dig. Tech. Papers*, 2006, pp. 196–197.
- [14] Tektronix, Inc., Beaverton, OR, *TDSRT-Eye v2.0 Serial Data Compliance and Analysis Software*. 2007 [Online]. Available: www.tek.com



William Redman-White (M'83–SM'08) has been with NXP (formerly Philips Semiconductors) since 1990, presently as a Fellow in Southampton, U.K. He has also worked in San Jose, CA, and Caen, France, on optical storage, WLAN, cellular radio, Bluetooth, digital audio, TV, satellite baseband, high-speed serial links and car security. He was previously with Motorola, Geneva, GEC-Marconi Research London, and Post Office Telecommunications, London. Concurrently with his industrial activities, he has also had a faculty position in Southampton University, U.K., since 1983, currently as a full Professor of integrated circuit design. His research and teaching is centered on analog and RF IC design, and design issues in SOI CMOS technology.

Dr. Redman-White has published 120 papers and has had more than 12 patents granted with several pending. He served as an associate editor for the *IEEE JOURNAL OF SOLID-STATE CIRCUITS* from 1996 to 2002. He is currently Analog Sub-Committee chair for the IEEE International Solid State Circuits Conference, has twice been Technical Programme Chair of the European Solid State Circuits Conference (1997 and 2008), and is a member of the steering committee for the European Solid State Circuits and Devices Conference series (ESSCIRC/ESSDERC).



Martin Bugbee (M'99) received the B.Eng (Hons) degree in electrical engineering with medical electronics from University College London, U.K., in 1994 and the Ph.D. degree from University College London, U.K., in 2000.

From 1998 to 1999, he worked as a Research Assistant at University College London developing implantable nerve stimulators and from 1999 to 2000 as a Research Assistant at Aalborg University, Denmark. During 2000 to 2001, he returned to University College London. In 2001, he joined Philips Semi-

conductors developing analog subsystems for CD/DVD and TV products including high-speed serial interfaces. He is currently with NXP Semiconductors, Southampton, U.K., working on very low-power circuits for the car access and immobilization market.



Steve Dobbs was born in 1966. He received the B.A (Hons) in Engineering from the University of Cambridge, Cambridge, U.K., in 1987.

He subsequently joined Philips Components, initially working on a DSP chip for audio applications. He then moved disciplines to focus on the design of transistor-level analog circuits for the NICAM and CD/DVD markets. In 1995, he left Philips to join Acapella Ltd., a small U.K.-based design house providing design effort on a contractual basis. In 1998, after Acapella was acquired by Semtech Corp.,

he focused on the design and development of high-speed PLLs for the newly formed communications division, specifically for optical networking products (SONET). In 2002, he left Semtech and rejoined Philips Semiconductors, working on various analog subsystems for CD/DVD and TV products including high-speed serial interfaces. He is currently with NXP Semiconductors (founded by Philips) working on very low-power circuits for the car access and immobilization market.



Xinyan Wu received the M.A.Sc. and Ph.D. degrees in electrical engineering from Xian Jiaotong University, Xian, China, in 1989 and 1994, respectively.

From 1994 to 1995 and from 1997 to 1999, he worked in Xian Jiaotong University, China. From 1995 to 1997 and from 1999 to 2000, he was with University of Southampton, U.K. From 2000, he joined Philips Semiconductors (now NXP Semiconductors), his R&D interests now are in analog design, modelling and verification.



Richard Balmford received the M.Eng. degree in electrical and electronic engineering in 1991 from Imperial College of Science and Technology, University of London, London, U.K., followed by research into analog circuit performance in the Microelectronics Department, University of Southampton, Southampton, U.K.

Since 1996, he has been with Philips Semiconductors, Southampton, U.K., developing analog subsystems for TV and optical disc CMOS ICs; Semtech Inc. Advanced Communications division; and NXP

Semiconductors, Southampton, developing BiCMOS front-end ICs for DVD systems, CMOS high-speed serial interfaces, and RF contactless ICs for car access and immobilization.



Jonah Nuttgens received the M.Eng. degree in electronic engineering from the University of Southampton, Southampton, U.K., in 1999.

He joined Philips Semiconductors in 1999 where he worked on digital design of ASICs for TV and set-top-box applications. In 2005, he moved into analog design, staying with the company which was later relaunched as NXP Semiconductors. In 2008, he submitted a patent for a novel low-power MOS threshold-voltage extraction circuit. He has been involved in a variety of projects including high-speed serial data links, RFID technology and complex SoCs for Digital TV.



Umer Salim Kiani received the B.Sc. degree in computer systems engineering degree from National University of Sciences and Technologies (NUST), Rawalpindi, Pakistan, in 1999, and the M.Sc. degree in communications and digital signal processing from Imperial College London, London, U.K., in 2001.

He worked as a Design Engineer at Avaz Networks in Pakistan from 1999 to 2000 where he was involved in developing VoIP related hardware accelerators. Since 2001, he has been with Philips/NXP

Semiconductors where his work has mainly focused on implementing communication interfaces on digital media storage drives and digital TVs. His main research interests include finding optimum hardware implementations for communication and digital signal processing algorithms and in optimal hardware and software partitioning in a system design.

Richard Clegg, photograph and biography not available at the time of publication.



Gerrit W. den Besten was born in 1971 in Heinen-Boeicop, The Netherlands. He started his study of electrical engineering at the University of Twente, Enschede, The Netherlands, in 1989. His graduation project involved the design of a high-speed low-jitter phase-locked loop, which was performed in cooperation with Philips Electronics. He graduated *cum laude* in 1994, and accepted a job in Philips Research at the NatLab in Eindhoven.

Initially, he worked on integrated reference circuits, on-chip power regulators, clock-generation and clock-multiplication. Since 1996, he has been active in the field of high-speed data communication. Besides finding new proprietary solutions, he has also been driving new interface standards. He has been a member of the working group which standardized USB2 and he is a contributor to the MIPI Alliance to define new interface industry standards for mobile applications. Furthermore, he has been involved in circuit implementations of several high-speed interfaces standards. With the disentanglement of the semiconductors division of Philips in 2006, he joined NXP Semiconductors, where he works now as a Senior Principal Architect in their research department. His current interests are in high-speed power-efficient interfaces and networks.