23.7 A 16Gb/s 1 IIR + 1 DT DFE Compensating 28dB Loss with Edge-Based Adaptation Converging in 5μs

Shayan Shahramian, Behzad Dehlaghi, Anthony Chan Carusone

University of Toronto, Toronto, ON, Canada

I/O receivers routinely equalize ISI over 10 or more post-cursor UI. IIR DFEs are a low-power technique for canceling long post-cursor ISI tails, and have been demonstrated compensating over 20dB loss at $f_{bit}/2$ up to 10Gb/s [1-5]. Equalizer adaptation is required to maintain signal integrity in time-varying channel and circuit conditions. Robust adaptation algorithms suitable for discrete-time (DT) DFEs are well-established, but there are few examples of adaptive algorithms for IIR DFEs [2,4], each exhibiting relatively slow convergence, additional high-bandwidth hardware and/or requiring the input data statistics to meet specific criteria. In this work, a 16Gb/s IIR DFE is integrated into a CDR, and the adaptation algorithm makes use of signals available in a regular binary phase detector (PD) to simultaneously adapt the IIR and DT taps. The novel algorithm provides faster and more robust convergence than has been previously demonstrated for IIR DFEs.

Figure 23.7.1 shows the receiver block diagram. A half-rate 1 IIR + 1 DT DFE is incorporated into a PD providing binary samples of the received data and edges. The half-rate outputs are demultiplexed 2:64 using custom high-speed demultiplexers (2:8) followed by synthesized logic (8:64), and then supplied to both the clock recovery unit (CRU) and DFE adaptation algorithm. The PD employs double-tail latches followed by SR latches, with DFE subtraction performed inside the latch [5]. A key challenge which has limited the speed of past IIR DFEs is their feedback loop delay, which includes a full-rate multiplexer. The IIR feedback filter input is here taken directly from the output of the regenerative latches, instead of after the SR latches, to reduce the feedback delay. Unlike [1,4,5], this work uses a 2:1 clockless multiplexer to reduce loading on the clock buffers. The multiplexer is, in fact, two SR latches in parallel, which alternately control its output, as shown in Fig. 23.7.1. When the even path of the DFE is evaluating, it sets the output of the multiplexer via "in1" while in the odd path both $in2_{p,n}$ are reset to zero. Similarly, during the alternate half-rate clock phase, the odd data latch evaluates setting the multiplexer output while $in1_{p,n}$ are reset to zero. Therefore, together the two SR latches function as a 2:1 clockless multiplexer.

The digital CRU uses the 64 demultiplexed data, a_k, and edge, E_k, samples to track the incoming data phase. Figure 23.7.2 shows a block diagram of the CRU and phase rotator. Early/late clock occurrences within each 64b block are counted by the PD logic, subtracted and passed through a proportional path with gain Kp and a parallel integral path with gain Ki. The proportional path tracks variations in the phase of the recovered clock relative to the data, while the integral path helps the CRU track frequency offsets between the incoming data and the receiver clock. The outputs are summed, integrated and truncated to 7b. The resulting clock phase code is converted into thermometer- and gray-coded signals for the phase rotator. The phase rotator consists of a multi-phase generator (MPG) and 8 single-ended phase interpolators (PIs), 4 data-sampling and 4 edge-sampling, followed by 2 pseudo-differential multiplexers. The MPG input differential clock is AC coupled to a two-stage ring oscillator to generate quadrature clocks. Each of the 4 data-sampling PIs is responsible for covering one quadrant of clock phases: 000-090, 090-180, etc. Identically, 4 PIs are used for edge-sampling. Each PI is comprised of 2 sets of 31 inverters driving the same node. The inverters are selectively (de)activated to provide a weighted combination of the input phases depending on the input phase code. To improve PI linearity, capacitor C1 reduces the swing at the inverter-bank output to approximately 400mVpp. AC coupling capacitor C2 and an inverter with resistive feedback follow to alleviate sensitivity to common-mode variations. Finally, the multiplexer outputs select between the different PI outputs depending on the quadrant of the selected phase. PIs corresponding to unused quadrants are dynamically powered up/down during rotation, saving 5mW out of 41.9mW total in the phase rotator at 16Gb/s.

Figure 23.7.3 illustrates the 1 IIR + 1 DT DFE adaptation algorithm. The same edge, E_k, and data, a_k, samples required by the CRU are used to inform the adaptation. The correlations between early-late PD outputs, E_{k-1}, and the four preceding bits a_{k-2} … a_{k-5} are proportional to the post-cursor edge ISI terms $h_{1.5}$, $h_{2.5}$, $h_{3.5}$, and $h_{4.5}$. The algorithm updates DFE coefficients iteratively moving the observed correlations ($a_{-k} \times E_{-1}$) towards zero, thereby minimizing post-cursor ISI in the channel pulse response edge samples $h_{1.5}$, $h_{2.5}$, $h_{3.5}$, and $h_{4.5}$. The binary product ($a_{-k} \times E_{-1}$) requires simply a logical XOR. Using this approach, no training pattern or lengthy BER measurements are required to perform adaptation as in, for example, [4]. No additional high-speed comparator is required for adaptation, avoiding the associated extra power, loading on a critical node in the DFE and phase-adjustment circuitry.

Fundamentally, to infer complete information about a channel response, spectrally rich data patterns are required. Other edge-based adaptation algorithms wait for specific patterns to arrive before updating the equalizer [2]. By contrast, this algorithm updates the equalizer upon receiving any 64b demultiplexed word containing at least 10 different 6b sequences (a_{-5}, a_{-4}, … a_0) having transitions ($a_0 \neq a_{-1}$). This criteria is easy to implement in digital logic, and prevents instability of the adaptation algorithm in the presence of patterns with insufficient spectral diversity. Yet the criteria also provides much faster convergence than previous approaches that await a specific pattern [2]. A second challenge is how to independently adapt the IIR DFE gain, B, and time constant, τ, both of which contribute to the cancellation of all post-cursor ISI terms, along with the DT tap weight, G. In [2] only $h_{2.5}$ information is used to adapt the IIR time constant which may not result in a good fit to a long tail in the channel pulse response. Moreover, [2] does not include a discrete-time tap for the DFE which leaves its performance sensitive to any process or voltage variations in the DFE feedback delay. In this work, the DT tap weight is iteratively updated to drive the correlation ($a_{-2} \times E_{-1}$) towards zero, minimizing ISI at $h_{1.5}$. The product ($a_{-3} \times E_{-1}$) is used to guide the IIR gain coefficient, B, towards zero ISI at $h_{2.5}$. Finally, the IIR time constant, τ, is guided by both the products ($a_{-4} \times E_{-1}$) and ($a_{-5} \times E_{-1}$), thereby adjusting τ to remove ISI at $h_{3.5}$ and $h_{4.5}$. To ensure the IIR gain has time to respond to changes in IIR time constant, τ is updated at $1/3^{rd}$ the rate of B. All equalizer coefficient updates are calculated at the demultiplexed clock rate, $f_{bit}/64$.

Figure 23.7.4 shows jitter tolerance (JT) with PRBS7 input at 16Gb/s with 2.7dB of setup loss. Measurements are shown for both mesochronous, and plesiochronous half-rate receiver input clocks. Both show similar low-frequency JT, demonstrating proper phase rotation as plotted in Fig. 23.7.4 (left). Figure 23.7.5A plots insertion loss for 3 channels having 15.7dB, 22dB, and 28dB loss at 8GHz. Figure 23.7.5B,C illustrates measured equalizer adaptation curves for channels 1 and 2 with PRBS7 input. Initial convergence is achieved within 80,000UI, over an order of magnitude faster than in [2], after which the BER is $<10^{-12}$. Figure 23.7.5D shows measured adaptation curves for channel 1 when repeating patterns are inserted. It is evident that the equalizer coefficients are not updated when the repeating patterns are present. Deactivating this feature, the coefficients diverge in Fig. 23.7.5D and the BER increases when the repeating patterns arise. Figure 23.7.6 shows measured bathtub curves for all three channels; all coefficients were adapted, except for channel 3 where DT tap (G) was fixed and IIR coefficients (B,τ) adapted. Figure 23.7.7 shows a die photo and area breakdown of the chip.

In conclusion, a 16Gb/s 1 IIR + 1 DT DFE was demonstrated in 28nm FD-SOI CMOS with integrated clock recovery and adaptation. The edge-based adaptation algorithm reuses the high-speed circuitry and signals required for clock recovery, is robust in the presence of ill-conditioned data statistics, and yet converges over an order-of-magnitude faster than previous techniques.

Acknowledgements:
Huawei and Semtech for financial support, ST Microelectronics for IC fabrication, and CMC Microsystems for test equipment.

References:
[1] B. Kim et al., "A 10-Gb/s Compact Low-Power Serial I/O With DFE-IIR Equalization in 65-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 44, no. 12, Dec. 2009.
[2] Y. Huang et al., "A 6Gb/s receiver with 32.7dB adaptive DFE-IIR equalization," *ISSCC Dig. Tech. Papers*, Feb. 2011.
[3] O. Elhadidy et al., "A 10 Gb/s 2-IIR-tap DFE receiver with 35 dB loss compensation in 65-nm CMOS," *IEEE Symp. VLSI Circuits*, June 2013.
[4] S. Son et al., "A 2.3-mW, 5-Gb/s Low-Power DFE Receiver Front-End and its Two-Step, Minimum Bit-Error-Rate Adaptation Algorithm," *IEEE J. Solid-State Circuits*, vol. 48, no. 11, Nov. 2013.
[5] S. Shahramian, A. Chan Carusone, "A 0.41 pJ/Bit 10 Gb/s Hybrid 2 IIR and 1 Discrete-Time DFE Tap in 28 nm-LP CMOS," *IEEE J. Solid-State Circuits*, vol. 50, no. 7, July 2015.

 978-1-4673-9467-3/16/$31.00

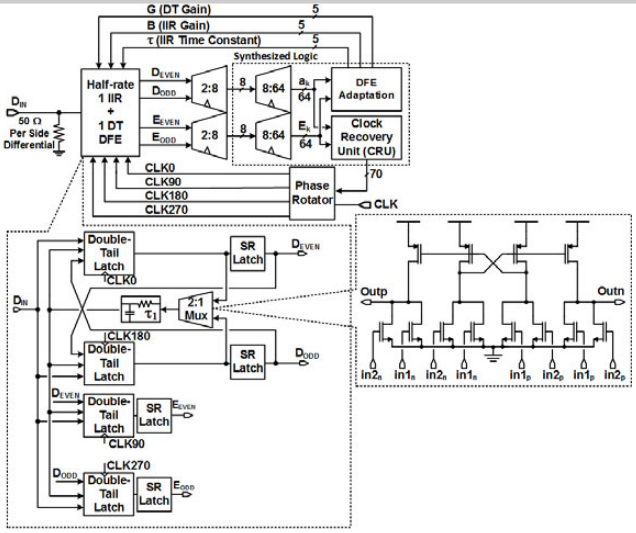

Figure 23.7.1: Complete receiver block diagram showing half-rate IIR DFE details, including the 2:1 clockless mux (inset).

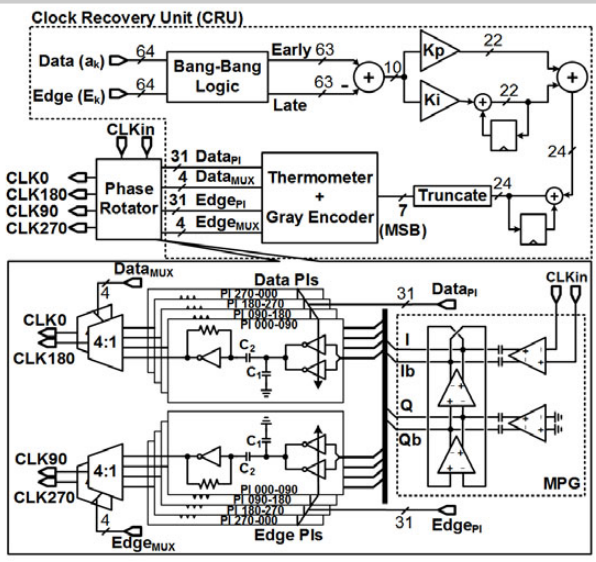

Figure 23.7.2: Clock recovery unit and phase rotator details.

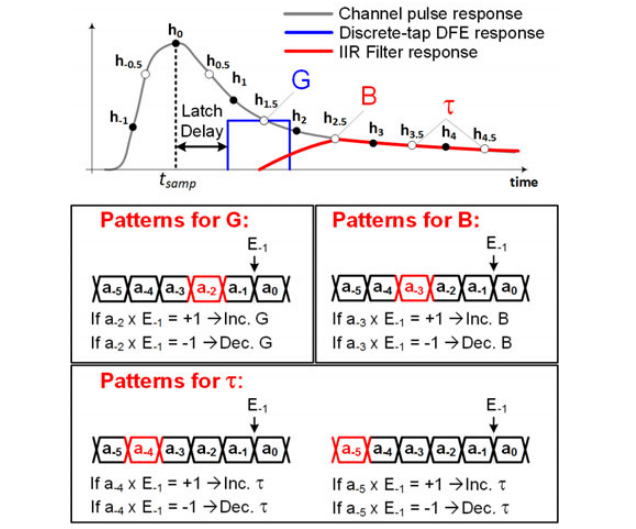

Figure 23.7.3: Proposed edge-based adaptation algorithm using all patterns with a transition ($a_0 \neq a_{-1}$) to update the DFE coefficients.

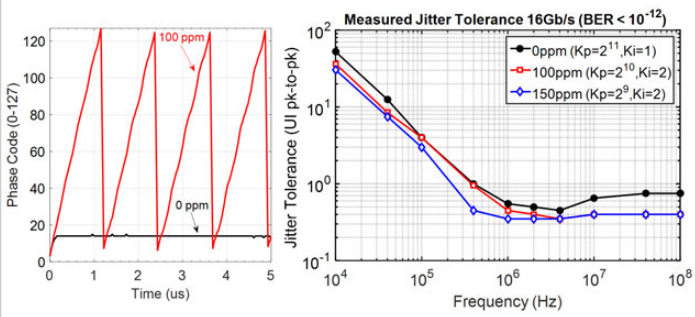

Figure 23.7.4: (left) Measured phase code vs time for 0ppm and 100ppm frequency error. (right) Measured jitter tolerance for 0ppm, 100ppm and 150ppm frequency offset.

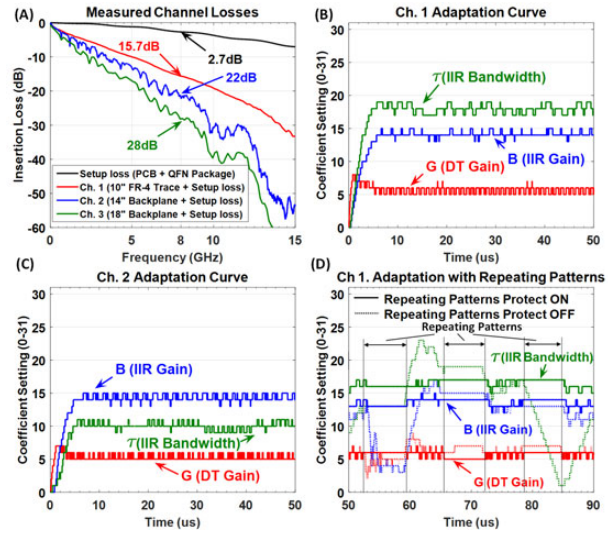

Figure 23.7.5: (A) Channel losses. Measured coefficient adaptation (B/C) Ch. 1/2 with PRBS7. (D) With 5µs intervals of repeating patterns.

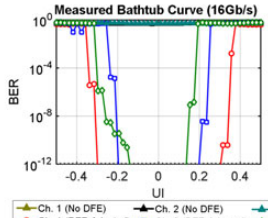

Block	Power Consumption
DFE (Data + Edge Latches + 2:1 Mux)	27 mW
Demux (2:8)	14.4 mW
Digital Logic (Adaptation, CRU, 8:64 demux)	24.3 mW
MPG	20.2 mW
PIs	21.7 mW
Clock Buffers	33.5 mW

	[1]	[2]	[3]	[4]	[5]	This work
Data Rate (Gb/s)	10	6	10	5	10	16
Architecture	1 IIR + 1 DT DFE	1 IIR DFE	2 IIR DFE	1IIR + 1DT + Pass. EQ	2 IIR + 1 DT	1 IIR + 1 DT DFE
Loss @ half bitrate	27 dB	32.7 dB	35dB	15 dB	24 dB	28dB
Technology	65nm	90nm	65nm	65nm (LP)	28nm (LP)	28nm FDSOI
Supply (V)	1	NA	1	1.2	1	1
DFE Power (mW) [Data path only]	3.5	4	9.9	2.3	4.1	15.8
mW / Gbps	0.35	0.67	0.99	0.46	0.41	0.99
Area (µm²)	17,250	89,000	30,400[a]	23,321	8,760[a]	8,100[a]
Adaptive Equalization	NO	YES	NO	YES	NO	YES
Adaptation Time	----	250 us (1.5 million UI)	----	Off Chip (1 hour 25 min) 2.55x10^13 UI	----	5 us (80,000 UI)

[a]Area of DFE core only.

Figure 23.7.6: Measured bathtub curves, power breakdown and performance comparison with previous IIR DFEs.

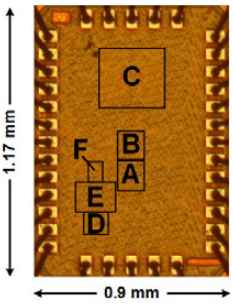

Block		Area (μm²)
A	DFE (Data + Edge Latches + 2:1 Mux)	11,900
B	Demux (2:8)	8,000
C	Digital Logic (Adaptation, CRU, 8:64 demux)	41,000
D	MPG	5,700
E	PIs	13,000
F	Clock Buffers	5,800

Figure 23.7.7: Die photo and area breakdown.

978-1-4673-9467-3/16/$31.00