
School of Electrical, Electronic & Computer Engineering



A Signature-based Single-Wire On-Chip Communication Scheme

Crescenzo D' Alessandro

Technical Report Series

NCL-EECE-MSD-TR-2007-121

August 2007

Contact:

Crescenzo.DAlessandro@ncl.ac.uk

Supported by EPSRC grant EP/C512812/1 and EP/C512804/1

NCL-EECE-MSD-TR-2007-121

Copyright © 2007 University of Newcastle upon Tyne

School of Electrical, Electronic & Computer Engineering,
Merz Court,
University of Newcastle upon Tyne,
Newcastle upon Tyne, NE1 7RU, UK

<http://async.org.uk/>

A Signature-based Single-Wire On-Chip Communication Scheme

Crescenzo D'Alessandro

August 2007

Abstract

This paper describes the use of two types of noncoherent on-chip communication over a single wire. The first, baseband binary phase-shift keying (BPSK), carry the information over the phase relationship between two delayed versions of the same carrier; the other, pulse-width modulation (PWM) exploits different pulse widths to carry the data. We show that these techniques allow high bitrate using single-wire serial links. The system proposed is also robust to process variations; this robustness is achieved employing redundancy at the receiver to filter out metastable latches. Concept and implementation examples are proposed and discussed.

1 Introduction

The requirement for high-speed on-chip interconnects is becoming a crucial aspect of chip design, so much that communication-centric and throughput-centric design is becoming an important topic of research. Serial communication for on-chip interconnects has been identified as a solution to the problem of fast, reliable and, importantly, low-power interconnects [4, 6, 8]. In its simplest form, a serial link consists of a serialiser at the sender side and a deserialiser at the receiver's (SERialiser/DESerialiser systems or SERDES); these are traditionally implemented as simple shift registers. However, the increased latency introduced by the interconnects makes it impractical to transmit a single bit over a long distance on-chip [4]; the serialisation has been therefore implemented at various degrees. Additionally, synchronous and asynchronous techniques have been proposed for long-distance interconnects, although the asynchronous paradigms respond to the needs expressed by the ITRS [7] of 40% asynchronous signalling by 2020.

Bainbridge in [1] proposes the use of 1-of-4 links. In this technique a pair of bits is sent over 4 lines, of which only one switches for every two bits. The system is inherently asynchronous and is of easy implementation; however, it does not reduce the number of physical wires necessary for the communication: in fact it doubles the requirement, as every wire carries 1/2 a bit. The same capacity is achieved employing differential links. In these, a bit is transmitted when the voltage difference over two lines exceeds a given value. This technique can be either synchronous or asynchronous and is both robust and suitable for low-swing application, allowing low-power consumption. An example is proposed in [6], where a low-swing technique achieves a power consumption of $<1\text{pJ/bit}$ in $0.18\mu\text{m}$ technology at 1 GHz.

Dual-rail applications are similar to the differential link paradigm, but are conceptually different, being a special case of the $m - of - n$ scheme (1-of-4 is also a special case of this technique). Dual-rail

schemes are inherently asynchronous, but in general require an acknowledge signal back from the receiver, which imposes long latency. A dual-rail system which reduces the number of transitions on the link, hence reducing the overall power consumption is Level-Encoded Dual-Rail or LEDR [9]. In this system, an implementation of the phased logic paradigm, only one line toggle to indicate a 0 or a 1 according to which line had switched. This is effectively a non-return-to-zero (NZR) dual-rail protocol.

Various techniques have been proposed, based on the idea of wave pipelining [2, 11].

In [5] the authors propose the use of LEDR for long-distance serial interconnects. In this case one line represents the 'phase' of the signal, while the other represents the 'state'. The 'state' line always represents the data stream; an event on the 'phase' line indicates that two bits of the same value have been transmitted. This scheme is particularly efficient in terms of transitions per bit, as a single transition per bit is used, with evident advantages in terms of power consumption. However, the link requires high matching between the two lines, as well as being affected by cross-talk and environmental noise; also, the metastability issue of the latches used in the receiver is not well analysed. To counteract the effect of cross-talk on the link, the authors propose the use of current-mode differential signalling for each line, coupled with wire placement to improve the shielding of each wire; the link requires 6 wires including the shielding for a single downstream link.

A different approach, also based on wave pipelining is proposed by Lee in [8]: the WaFT system replaces the shift registers in the sender and receiver with delay elements and multiplexers. The system is ideally suited for mesochronous systems, although the timing assumption and the effects of PVT variations are not discussed. The problem of synchronisation between two communicating blocks is also not discussed.

From the point of view of reliability, the shift towards sub-micron technology nodes imposes strong limits to the traditional design of long-distance interconnects. Systematic timing variations increase, due to the process variability imposed by the manufacturing technology available. The high sensitivity of the smaller transistors also results in random transient variations, due to voltage or temperature fluctuations. These appear as jitter on the communication lines, which can cause failure of the link. Worm *et al.* in [12] propose the use of two-phase communication, in the sense that the data is encoded in two different ways alternatively; the receiver keeps track of the "phase" (or state) of the system and decodes the data according to the phase. If the decoding fails it means that a massive timing variation has occurred and the system recovers. This approach is used to implement a voltage-scaled link: the operating voltage is determined dynamically by the presence or absence of errors.

D'Alessandro in [3] initially proposes the use of phase relationship to encode a bit of data, following on the tradition of phase-modulated links. The implementation described in the paper refers to a dual-rail link where each of the lines is differentially phase-modulated (or delay-modulated if the link is not driven by a periodic clock), so that the order of arrival of events on the lines encodes the item of data being sent.

This paper proposes a refinement of the work described in [3]. Instead of two wires, a single wire is used; the data modulates the phase of a clock of given frequency, so that every pulse is either in-phase or out-of-phase, indicating the presence of a 0 or a 1. An alternative system is also presented where a PWM implementation is used. We show that the implementation allows energy requirements to be reduced to around 3 pJ/bit at 1.8V power supply for a 0.18 μm technology, with a link speed of 1 GHz, referring to a bitrate of 1 Gb/s.

2 Concept

The work proposed in this paper employs the time difference between two events *on the same wire* to encode/decode the data. This can be done using a TDC (Time to Digital Converter), a device which, given two events at a time difference between them, produces a numerical output proportional to the time separation. The most important parameters of a TDC are the resolution ρ (minimum time difference between the events to produce different results), the capture range κ (the maximum time between the edges above which the TDC saturates and produces the same output, minus the minimum time below which the TDC again saturates) and the latency λ (the time required to produce an output). Of course, area and power consumption are implied parameters which have to be taken into account; however, from the time measurement point of view, the mentioned parameters will discriminate between the performance of different styles of TDCs.

The simplest TDC can be seen as a bank of latches; a “start” signal will be connected to a delay line where each stage has a given delay τ and each stage of the delay line is connected to the D pin of the latch. A “stop” signal will be instead applied to the $clock$ pin of the latch. If the “start” signal is issued before the “stop” signal, the two signals will meet along the delay line and a number of latches will latch a 1 and the rest will latch a 0. This will provide a “temperature coded” number which encodes the time between the two edges. An example of this device is shown in Figure 1.




Figure 1: Delay line-based TDC

This design, although very simple, has a number of drawbacks:

- The resolution time is effectively τ – the delay of each stage
- Where the two edges overlap there might be the risk of metastability
- The capture range is proportional to the length of the delay chain, imposing large area overheads

The first problem can be solved using a Vernier delay line, where the “start” and the “stop” signals are fed to two different delay lines, the first with elements of delay τ_1 , the second with elements of delay τ_2 . This will improve the resolution time to $\tau_1 - \tau_2$. However, in this case there also is the problem of metastability where the two edges overlap, and the area overhead is actually slightly worse than in the previous case,

as twice as many delay elements are required (although having the same number of stages with smaller resolution results in shorter capture range).

Other solutions have been proposed to address these problems. However, in general, the solutions allow for metastability to occur. This is because in general the time measurement accuracy takes the last unit into account. Also, in many cases a number of trials are possible; as the metastability of one node causes the output to be seen as either 1 or 0, statistical techniques can be used to extrapolate the correct time measurement within the accepted margin.

If the time information encodes data, though, this is not acceptable. The decision unit, in fact, will need a threshold below which and above which the results mean different things. Thus, if the bit which represents the threshold is metastable, the output might be incorrect 50% of the times. Obviously one solution would be to send an item of data an odd number of times in succession, so that a voting system can be employed. This, however, will impose a strain on the system in terms both of area (the extra area required for the circuitry) and time (the time redundancy introduced by the additional pulses).

Instead, we propose the use of signature-based communication. The sender first produces a “calibration” pulse; the result of this first measurement is stored as a “signature”. Every time another pulse arrives, the time measurement is compared to the stored signature, and the data can be recovered checking whether a match is found or not. This method makes the communication channel more adaptive, as PVT variations are taken into account at run time, at the expense of an additional pulse to calibrate the system and additional circuitry.

3 Proposed Schemes

3.1 BPSK System

Binary Phase Shift Keying (BPSK) is a well known and widely used communication paradigm. The system consists of a sender which is able to change the phase of a transmitted waveform according to the item of data being transmitted. The frequency of the carrier is normally much higher than that of the data stream for RF communication. In this work the data stream is at the same frequency as the carrier (baseband communication).




Figure 2: Block diagram of BPSK sender

Figure 2 shows the block diagram of a BPSK sender. The strobe line is modulated onto two lines, each of which with a different phase. The two phases are then multiplexed over the link line by the data; this scheme approximates the traditional multiplying mixer used in RF communication. The data line forces a toggle flip-flop to change the output when the data to send is a 1; the toggle is operated by the strobe signal.




Figure 3: Block diagram of receiver

At the falling edge of the strobe signal, that is, when a pulse has been sent, the toggle output is updated and the system is ready for a new item.

Figure 3 shows the receiver block for the same system. The waveform is delayed by a one-bit delay; the time of arrival of the rising edge is then compared to the time of arrival of the rising edge of the undelayed waveform using a phase detector; the logic then decides whether a 1 or a 0 was received. This system also approximates the correlator used in standard BPSK communication systems.

A standard BPSK receiver also employs a discriminator (normally an op-amp) which, after the correlator has completed the integration, decides which item of data has been received. For the baseband case, this introduces a fundamental problem: the sensitivity of the receiver needs to be high to allow fast communication, but at the same time it needs enough resolution to avoid loss of precision. We propose a novel approach to the problem of designing a robust discriminator logic. The logic operates as follows: a “pilot” pulse is initially sent by the sender; the logic ignores the output of the TDC for the first pulse. An additional pulse, at the same phase as the previous pulse, is now sent. This is used to calibrate the system: the TDC in fact will output a “signature” which refer to the same-phase case; the logic stores this signature into a register.

From this point onwards the pulses sent over the link will contain data. When the output of the TDC corresponds to the signature, the last two pulses were in phase with each other: referring to the sender of Figure 2, a 0 is therefore received; otherwise, if the signature does not corresponds to the stored signature, the two pulses were not in phase and a 1 is received.

3.2 PWM System

A natural evolution of the system just described in the previous section consists in removing the one-bit delay and use the pulse width to identify whether a 1 or a 0 has been received.




Figure 4: Block diagram of PWM sender

In this case the sender needs to output a rising transition at time $kT + d_k \delta$, where $d_k \in \{0, 1\}$ is the data item to be transmitted, $k = 0, 1, 2, \dots, T$ is the period of the sender strobe signal and δ is the time difference between a pulse indicating a 1 and one indicating a 0. The falling transitions, instead, are generated at times $kT + \frac{T}{2}$. Figure 4 shows an implementation of such a sender, based on the Transition Sequence Encoder




Figure 5: Block diagram of PWM receiver

(TSE) described in [10]. Note that the choice of the association of a long pulse with a 1 or a short pulse with a 0 or viceversa is trivially arbitrary: in fact in the equation above a 1 will produce a short pulse, while in Figure 4 the opposite is true.

Figure 5 shows the receiver part of a PWM system. This is very similar to the BPSK receiver, with some important differences. First, the absence of the one-bit delay makes the system lighter, removing the need for long delay lines. Also, the TDC now is started and stopped by the same link signal: it is started when a rising transition is present on the link and is stopped when a falling transition is seen. The logic operates in a manner similar to the BPSK logic, with the difference that the calibration pulse is unnecessary: the first pulse indicates the length of a 0 or a 1 (according to the designer's decision); the "signature" of this pulse is recorded. If the subsequent pulse lengths are equal to the stored signature, then the appropriate item of data has been received.

3.3 Notes on Link Speed

The design of the TDC imposes constraints on the minimum size of the pulses, which in turn will affect the speed of the system. This depends on the type of gates used for the detection; we devise an example of the required analysis in the case of latches used to determine the width of the pulse.

First, the pulses should respect the minimum width for the high and low stages of the clock signal specified in the data sheets for the particular technology. This imposes an upper bound to the speed of the system. Second, the time separation difference between a data item 1 and a 0 should be recognisable, and therefore greater than one TDC stage delay. Third, the stages should be designed so that the latches don't enter metastability, or so that at most only one latch enter metastability. This implies that the TDC stages should be spaced by at least $t_{setup} + t_{hold}$, so that no two latches can both enter metastability at the same time. In fact, if this is not the case, two (or more if the spacing is very small) latches could enter metastability and produce outputs which violate the thermometer coding. Thus, the minimum difference in pulse width between a 1 and a 0 must be at least $t_{setup} + t_{hold}$. To look at a proper case, consider the specifications of a major vendor's $0.18\mu m$ technology standard cell library. The latches are transparent when the clock signal is high and retain the value of D just before the clock goes low as the clock goes low. This is exploited to simplify the design: the link in input is sent both to the delay line and directly to the clock signals; when the link goes high the latches are transparent and the pulse is propagated through the delay line. As the pulse goes low, the latches retain the value and the time measurement can be performed. The minimum pulse width specification for the latches is around $500ps$ for both the high and low state of the clock. The setup time, however, is $-1ps$ if the data goes high while the clock goes low, which would be the case for the PWM system. Finally, the hold time is $11ps$ for the same case. Using these conditions as a

guideline, the figure of $500ps$ can be used to indicate a 0 and $500ps + t_{setup} + t_{hold} = 510ps$ to indicate a 1. This can be made more resilient by increasing the difference in pulse width, so the 1 is indicated by a pulse of width $500ps + 2(t_{setup} + t_{hold}) = 520ps$. This makes the link achieve a bandwidth between $1/1ns=1GHz$ and $1/1.02ns=0.98GHz$, or an average bitrate of $0.99Gb/s$. In practice, the delay introduced by a buffer is around $100ps$, so each stage is separated by a value well greater than the required $10ps$. This limits the speed of the link if a simple delay line is used to an average bitrate of $0.916 Gb/s$.

However, the reader must take into consideration the fact that the link needs a pilot and a calibration pulse, or a single pilot/calibration pulse in the case of PWM, for correct operation. In this case the effective link speed in terms of data items is of course reduced by the presence of these pulses. The reduction in data rate depends on the frequency of occurrence of these pulses; this is improved if the calibration is performed at idle times or if the packet length is large. In the case of BPSK the link will have an effective data rate of $bitrate \cdot \frac{packet\ length}{packet\ length+2}$; for a bitrate of $0.916 Gb/s$ and a packet length of 8 bit the effective data rate will be $732 Mb/s$. PWM is better because only a single calibration pulse is required, so the data rate is $bitrate \cdot \frac{packet\ length}{packet\ length+1}$; for the same conditions the effective data rate will be $814 Mb/s$. Increasing the packet length improves the data rate: using a packet length of 32 bits already increases the data rate to $862 Mb/s$ and $888 Mb/s$ for BPSK and PWM respectively. If the calibration is performed every n packets the data rate is also increased.

4 Fuzzy Receiver

The design of Time-to-Digital Converters (TDCs) requires the use of a reliable time reference, which is then used in different ways according to the system requirements. The simplest TDC is composed of a high-frequency clock and a counter, which is incremented by the high-frequency clock between the occurrence of the two events whose time difference is being measured: when the first event occurs, the counter starts being incremented. The value of the counter at the occurrence of the second event is proportional to the time between the events. This approach, although simple and accurate, has the obvious drawback that can only be used when the time between two events is of orders of magnitude greater than the period of the counting clock. This is not acceptable when measuring small time differences.

A different approach consists in using a delay line and “tapping” the line at different points; a number of arbiters is then used to arbitrate between one input signal and the various taps of the other.

We propose a different approach to the use of TDCs for BPSK. The philosophy behind this approach is based on the assumption that the absolute time difference between the two events is immaterial: only the relative phase is important. Therefore it is not necessary to calculate with accuracy this time separation. Additionally, PVT variations can affect the delay line employed, degrading the performance of the circuit.

The system proposed instead first produces a “pilot” pulse; after this pulse, an additional pulse is sent, at the same phase of the pilot pulse. This second pulse is used to calibrate the system: the result of the TDC is stored into a register as a signature. The subsequent pulses will produce the same signature if the phase between two pulses is the same or a different result if the phase is different.

This system can be made more robust by noticing that the highest arbiter can go metastable: in this case the output of this arbiter could be either 1 or 0, so the phase difference may appear different even if it is in fact the same. To avoid this problem it is necessary to store not only the signature, but two copies of the same signature, one shifted left and one right: when a pulse arrives, the signature of the current pulse is compared to the other two and if the incoming signature matches any of the stored ones, the phase of the

two pulses was the same.

The “catch” is that the time separation necessary to indicate a change in phase must be greater than two delay units used in the delay line. In fact, if the delay is less than this value, the signature might match one of the stored values, introducing an error.

Figure 6 shows a block diagram of the decoder. The result of the arbitration for the pilot/calibrate pulse pair is stored in the signature register, against which all subsequent arbitrations will be compared. This comparison is performed simply using XOR gates or equivalent.

However the result of a measurement is also compared against the signature shifted one bit to the left and then to the right. Instead of storing three different signatures, the system is wired in such a way that the comparison is performed immediately. The result of the three comparisons are then “merged” and the decoder outputs the item of data.

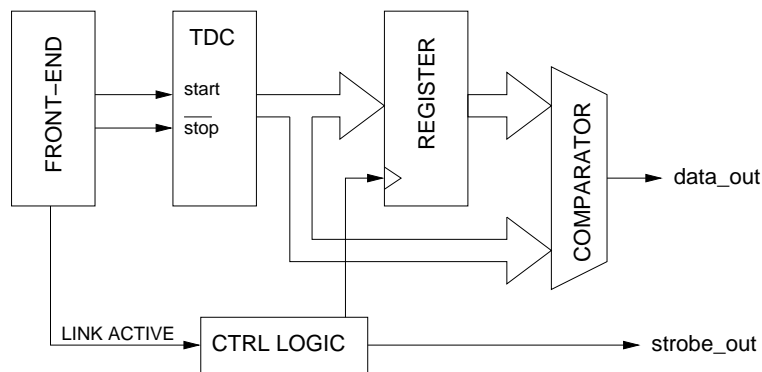


Figure 6: Block diagram of decoder

5 Implementation and Results

An implementation example of the system was designed for a major vendor’s $0.18\mu\text{m}$ technology and simulated using SPECTRE. The link is able to reliably communicate across a 10-mm wire using small drivers at a speed of 1 GHz, or 1Gb/s, as one bit only is sent for every pulse. Preliminary results show an overall power consumption of 3 pJ/bit using BPSK; using pulse-width the energy per bit is 1.6pJ.

The next set of figures shows the dynamic behaviour of the proposed systems through the SPECTRE simulations.

6 Conclusions

An attractive method for single-wire on-chip communication has been proposed. The link allows speed approaching 1Gb/s in $0.18\mu\text{m}$ technology. The main attractiveness comes from the adaptability of the system to different operating conditions: the link can adjust its operation on the fly, or it can be calibrated at given intervals, or else the calibration can be issued by a higher-level management in idle times.

The bit rates achievable get close to 1 Gb/s; considering the calibration and pilot pulses, the effective data rate can reach values of 900 Mb/s.




Figure 7: Waveform of BPSK link




Figure 8: Waveform of PWM link

This type of link is ideally suited for point-to-point communication between proprietary IP blocks on a single chip; otherwise it could be used as the underlying physical interconnect fabric for a NoC.

Improvements on the scheme consist in developing a complete infrastructure for the scheme, in particular taking into account the handshaking between sender and receiver.

References

- [1] J. Bainbridge and S. Furber. Delay insensitive system-on-chip interconnect using 1-of-4 data encoding. In *Proc. 7th ASYNC*, pages 118–126, 2001.
- [2] L. Cotten. Maximum rate pipelined systems. In *Proc. AFIPS Spring Joint Computer Conference*, 1969.
- [3] C. D'Alessandro, D. Shang, A. Bystrov, and A. Yakovlev. PSK Signalling on SoC Buses. In *Proc. PATMOS 2005*. Springer, 2005.

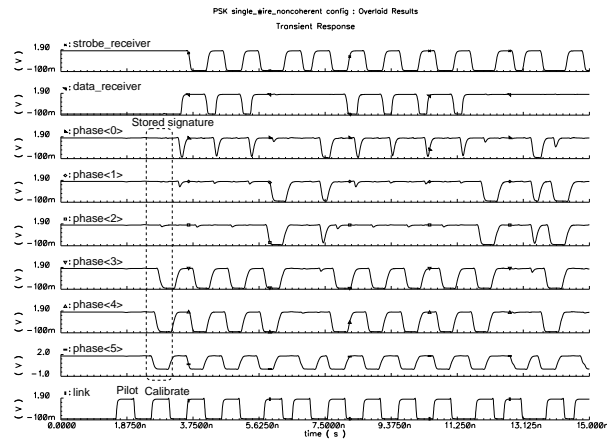


Figure 9: BPSK receiver waveforms




Figure 10: PWM receiver waveforms

- [4] Vinita Vasant Deodhar. *Throughput-Centric Wave-Pipelined Interconnect Circuits for Gigascale Integration*. PhD thesis, Dept. of Electrical and Computer Engineering, Georgia Institute of Technology, USA, 2005.
- [5] R. Dobkin, R. Ginosar, and A. Kolodny. Fast asynchronous shift register for bit-serial communication. In *Proc. 12th ASYNC*, pages 117–126, March 2006.
- [6] R. Ho, K. Mai, and M. Horowitz. Efficient on-chip global interconnects. In *Symposium on VLSI Circuits*, 2003.
- [7] International Technology Roadmap for Semiconductors (ITRS-2005), 2005.
- [8] Se-Joong Lee, Kwanho Kim, Hyejung Kim, Namjun Cho, and Hoi-Jun Yoo. Adaptive network-on-chip with wave-front train serialization scheme. In *Symposium on VLSI Circuits*, 2005.

- [9] Daniel H. Linder and James C. Harden. Phased logic: Supporting the synchronous design paradigm with delay-insensitive circuitry. *IEEE Transactions on Computers*, 45(9):1031–1044, September 1996.
- [10] A. Mokhov and A. Yakovlev. Transition Sequence Encoder. Technical Report NCL-EECE-MSD-TR-2006-117, Newcastle University (UK), 2006.
- [11] D. C. Wong, G. De Micheli, and M. J. Flynn. Designing high-performance digital circuits using wave pipelining: Algorithms and practical experiences. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 12(1):25–46, jan 1993.
- [12] F. Worm, P. Ienne, P. Thiran, and G. De Micheli. On-chip self-calibrating communication techniques robust to electrical parameter variations. *IEEE Design & Test of Computers*, 21:524–535, November-December 2004.