An exploration of synchronization solutions for parallel short-range optical interconnect in mesochronous systems

Harald Devos, Joni Dambre, Wim Meeus, Dirk Stroobandt and Jan Van Campenhout

Department of Electronics and Information Systems, Ghent University
Sint-Pietersnieuwstraat 41, B-9000 Ghent, Belgium

ABSTRACT

As a result of the increasing complexity of electronic chips, the bandwidths required for inter- and intra-chip communication are rapidly increasing. As optoelectronics provides high-bandwidth and high-density interconnection it is considered as a candidate for short-range interconnection. For such interconnections, situated at a low level in the systems hierarchy, the interconnect latency is extremely critical for the systems performance. This paper describes some methods for mesochronous synchronization, needed for such interconnections. It will be shown that it can be beneficial to use an additional optical link to transfer a synchronization signal. Such a reference signal can be used efficiently for phase detection, provided that the data skew is sufficiently small, and result in a decrease of the cost-per-link.

Keywords: Synchronization, mesochronous, latency, skew, optical interconnect, short-range, parallel

1. INTRODUCTION

As a result of the technological progress towards bigger chips, smaller feature sizes and larger clock frequencies, the bandwidths required for inter- and intra-chip communication are rapidly increasing. When trying to meet these requirements, the bandwidth limitation for traditional electrical interconnects has become one of the main worries of digital system engineers. Hence, over the years, an increasing amount of research effort has been spent on alternative interconnect technologies for short distance (inter- or intra-chip) communication. Optical interconnect technology, which is not plagued by the bandwidth limitation problem, has effectively solved this problem where it first occurred: for high data-rate, long- and medium-distance communication (e.g. networking and telecom applications). Since optoelectronics can provide flexible, high-bandwidth and high-density interconnection, it is also considered one of the most promising candidates for short-range (inter- and even intra-chip) interconnection.

However, for short-range interconnections, which are typically situated at much lower levels in the system hierarchy, the overall (electrical-to-electrical) interconnect latency, relative to the system clock period, is extremely critical for system performance.¹ Furthermore, especially if they are to be used in consumer-applications, they should be low-cost and often low-power. In contrast to the situation with long-haul interconnections the latency is not dominated by the time-of-flight (TOF), but spread over all parts of the link. To minimize latency, careful design and exploration are required for each of the elements constituting the electrical-to-electrical path. One of these elements is the receiver, an electronic circuit that converts the detector output current into a locally synchronous digital signal. For high-density interconnect, all skew between parallel channels, caused by unavoidable differences in the physical properties of those channels, must be eliminated by the receiver circuit.

In this paper, we address the design exploration issues in high-speed, high-density short-range interconnect. We focus on optimizing synchronization circuitry for mesochronous communication, i.e. communication where the frequency of the incoming signal is equal to the local clock frequency, but its phase is unknown and possibly variable in time. This situation is typical for systems with short-distance interconnects where both ends of the link are fed by the same clock. The delays in the clock distribution and the optical link prevent it from being a synchronous system. As a result synchronization is needed at the receiver side.

Further author information: E-mail: hdevos@elis.rug.ac.be

The exploration is performed by means of detailed design and quantitative analysis of various receiver circuits. Our results show that it can be beneficial to use one optical link to transfer a synchronization signal. Indeed, such a reference signal can be used very efficiently for phase detection, provided that the skew between data channels is sufficiently small (relative to the clock frequency). Hence, the investment of an additional link can result in a significant decrease of the cost-per-link, especially when a lot of links are used in parallel.

2. OPTICAL LINKS AND THE NEED FOR SYNCHRONIZATION

In this paper an optical link is considered as the entire electrical-to-electrical path consisting of driver, LED or laser, optical channel, detector and receiver. Each of these elements can be chosen from a lot of varieties, and the choice made for one can affect the behavior of the others. Also the used encoding can have an influence. A lot of receivers need a DC-balanced input.

Receivers can be clocked or unclocked. Unclocked receivers will change its digital output whenever the optical signal at its input changes while clocked receivers will only do this at certain (clocked) moments. A clocked receiver can be modelled as an unclocked one followed by a sample flip-flop. This allows us to solve the synchronization problem for an unclocked receiver and only consider clocked receivers in the end.

We will assume that the system is mesochronous. This means that the digital circuits at the sender and at the receiver side of the link are fed by the same clock (exactly the same frequency) but with a different phase due to delay in the clock lines. Even if an ideal clock distribution prevents phase differences, synchronization will be necessary because of delay, also called latency, in the optical data lines. This situation is typical for short range links, situated at a low level in the system hierarchy. As a result the receiver signal has to be synchronized. In an ideal situation this can be done by a flip-flop that samples the signal each period of the local clock (at the receiver side).

2.1. Metastability

However in a real situation this is not sufficient. If the phase of the incoming signal is not known data transitions can occur during the sampling time. This can cause metastability: the time needed by the output of the flip-flop to reach a logically defined state is not limited.² Normal circuit simulators do not allow to simulate metastable behavior what hinders the study of this problem. Data transitions should be avoided at sampling moments or better: sampling should happen when the data is stable.

2.2. Jitter

Another difficulty is caused by jitter (J). The delay of an optical link is not a constant. Fluctuations may occur at the sender and the receiver (e.g. Ref. 3,4). It is defined as the maximum deviation of the transition moment from the average transition moment. Using this definition the separation between two eyes in an eye diagram is two times the jitter (Fig.1). Here it is assumed that each link in a parallel interconnect has the same jitter. If this is not the case J is the maximum jitter of all the links (= worst case taken for each link). Jitter can be pattern dependent, so encoding can have an influence on it.(e.g. Ref. 5)

As a result of the jitter, the transitions at the output of the receiver might occur before the sampling of the clock one time and after it another time. Fig. 2 illustrates this. $Clk1$ is in the region where transitions can occur. Only the samples made by $Clk2$ are always correct. Sampling must happen within the eye.

If one has the freedom to change the phase of the local clock these problems can be solved. Usually this is not the case, and even if it is, problems may arise when several or bidirectional links are used. The phase can not be adjusted to all links at the same time.

Another solution is to change the phase of the incoming data (to shift the eye-diagram). This is done in a synchronization circuit. Different solutions will be explained in Sect. 3. Synchronization includes two actions:

1. Determining the phase of the data in comparison with the local clock.
2. Adjusting (delaying) the phase of the data so that the data will be sampled in the eye of the eye-diagram.

The phase of the data is not a constant. Slow variations in the latency (low frequent jitter) should be compensated by the synchronizer.

Figure 1. Eye-diagram of receiver signal (1 link) and schematic representation (2 links)

Figure 2. Jitter may cause uncorrect sample values.

2.3. Synchronization Latency

Synchronization may introduce an extra delay on the data lines. We will call this delay synchronization latency. This latency has to be kept as low as possible (as all latency) to optimize the behavior of the system. Theoretically it can be kept lower than one clock period.

2.4. Skew

In a parallel interconnection there will be a difference between the latency of the various links. This phenomenon, called skew, is caused by unavoidable differences in the physical properties of the different links. The received optical power for example influences the delay caused by the receiver.⁶ To reduce this skew a high uniformity of parallel links is needed. The skew (S) of a bundle of parallel links is defined as the maximum difference in latency between two links.(Fig. 1)

If the skew is small enough, knowing the phase of one link suffices for synchronizing them all. Else, each link has to be synchronized independently. The main purpose of this paper is to investigate how synchronization of a parallel interconnect with minimal synchronization latency can be optimized when the skew is sufficiently low and to determine what sufficiently low means.

3. SYNCHRONIZERS

3.1. Changing the phase

We will now look at some circuits to shift a signal in time before it is sampled by the local clock. Most of these circuits are described in Ref. 7.

Figure 3. Delay-line synchronizer.

3.1.1. Delay-line synchronizer

In this case the signal is really delayed by a delay line consisting of several delay elements. There are two possibilities to change the delay. In Fig. 3.a an analog voltage is used to adjust the delay of the elements. In Fig. 3.b each element has a constant delay. A multiplexer is used to vary the delay in discrete steps by selecting the number of elements included in the signal path.

The synchronization latency is kept minimal if the signal is delayed until the sampling time is in the middle of the eye. If the delay is varied in discrete steps (δ) the middle can only be approximated (in the worst case the deviation is $\delta/2$, supposing that the eye-diagram is known exactly).

The delay should be able to be varied from 0 to one clock period T . This limits the frequencies at which the circuit can be used. The minimum frequency is reached when the maximum delay corresponds with one clock period. In Fig. 3.b δ is a constant and the resolution is δ/T . This gives an upper limit to the frequency.

In a parallel interconnection each link needs its own delay-line. This is a big cost in chip area and power consumption when the number of links increases. (Ref. 8 gives an idea: one delay-element with a delay of 500 ps, 0.4 μm technology, 10 $\mu\text{m} \times 40 \mu\text{m}$, $P = 0.33 \text{ mW}$ at 100 MHz.)

3.1.2. Multi-register synchronizer

Another possibility is to sample the signal on different equidistant moments and select, with a multiplexer, the sample taken closest to the center of the eye in the eye-diagram. This sample is resampled by the local clock to put the signal in the local clock domain. (Fig. 4.a)

Let n be the number of equidistant clock phases (each shifted over T/n). n should be large enough to guarantee that there is always one phase situated in the eye. On the other hand the number of phases is limited by the aperture and propagation time of the flip-flops. The output signal of the first flip-flop should be stable when it is sampled by the second (lower bound for T/n).

3.1.3. Synchronizer with auxiliary clock

In Fig. 4.b the multiplexer selects one clock phase, henceforth called the auxiliary clock (H), instead of one sample. This results in a reduction of the circuit. We now have the receiver followed by a flip-flop. This means a clocked receiver can be used here.

If the auxiliary clock is used for all parallel links, the cost per link is only 2 flip-flops (probably the lowest cost possible). The cost of the generation of the clock phases is shared by all the parallel links. It relatively diminishes as the number of links grows.

Fig. 5 illustrates the operation with four clock phases. In this case both clock phases 4 and 1 are close to the center of the eye. Phase 4 is selected as auxiliary clock and the synchronization latency is minimal. If

Figure 4. (a) multi-register synchronization and (b) use of an auxiliary-clock.

Figure 5. Synchronization with an auxiliary clock.

Figure 6. Change of the auxiliary clock (H) from C4 to C1.

phase 1 was chosen this latency would have been more than one clock period. If the choice of the auxiliary clock is altered*, e.g. as a result of low frequent jitter, between these two clock phases, or in general between the first and the last clock phase, the difference in synchronization latency must be compensated. Else a bit will be transmitted twice or get lost as shown in Fig. 6 and Fig. 7. This problem — we call it the date line problem by analogy with people who lose or win a day by crossing the date line — can be solved by adding an extra flip-flop. By inserting and removing this flip-flop in the data path a delay of one clock period can be added and removed to avoid the large jumps in synchronization latency. This flip-flop stores the bit that is in danger of being lost, is used as a one bit shift-register, causing the one period delay, and is finally removed when the auxiliary clock returns to the last clock phase (no bit is transmitted twice). This problem also arises in delay-line and multi-register synchronizers.

The same auxiliary clock can be used for all links on condition that the skew is not to large. There should be a region that is common to all the eyes ($T - 2J - S > 0$, see Fig. 8) and there should be at least one clock phase in that region. With T/n the distinction between the clock phases this results in the condition $T - 2J - S > T/n$. This condition only is sufficient when the choice of the auxiliary clock is based on the phases of all links. If it is based on only one link, e.g. H = clock phase closest to the center of the eye of Chan.1 on Fig. 8, the condition becomes more stringent: $T/2 - J - S > T/2n$ or

$$J + S < \frac{(n-1)T}{n} \frac{1}{2}. \quad (1)$$

The condition becomes weaker as the number of phases becomes higher but $J + S$ must be smaller than $T/2$ anyway. The formula can be seen as a limit to the skew and jitter given a certain frequency but it can also be seen as a limit to the used frequency, given the properties of the optical link:

$$f = \frac{1}{T} < \frac{(n-1)}{2n} \frac{1}{J+S}. \quad (2)$$

The aperture and propagation time of the flip-flops are neglected to keep these formula's simple.

Several measurements of the OIIC-system demonstrator,⁹⁻¹¹ an experimental system consisting of optically interconnected FPGA's,¹² are listed in table 1. Application of Eq. 2 results in table 2. It is clear that these results have to be improved.

*The choice of the auxiliary clock is considered to be synchronous with the local clock (C1)

Figure 7. Change of the auxiliary clock (H) from C1 to C4.

Figure 8. Maximum jitter and skew.

Table 1. Measurements on the OIIC-system demonstrator

Measurement (ns)	9 cm POF	23 cm POF
Jitter	1.45	3.25
Skew	2.4	4.3

Table 2. Maximum frequency in MHz, n = number of clock phases

n	9 cm POF	23 cm POF
2	65	33
4	97	50
8	114	58
∞	130	66

3.1.4. FIFO-synchronizer

In this case an extra clock signal, $xclk$, is sent in parallel with the data signal x , that samples x in turn with the upper and with the lower flip-flop (Fig. 9). $x0$ and $x1$ are synchronous with $xclk$ but change only every other period. xm is swapped between these two such that it is stable until the next edge of the local clock clk , and then resampled to obtain xs . The synchronization latency lies between $3T/2$ and $5T/2$, and is not minimal. To reduce this loss some combinatorial logic can be inserted between xm and xn if its delay is sufficiently less than a period.

This circuit will only work if $S + J < T/2$, so $xclk$ can sample x , and xp and rp are initialized well with respect to each other.

3.1.5. Generation of clock phases

The most common circuit used to generate clock phases is a DLL (Delay Locked Loop), a delay line of which the delay is adjusted to one or one half clock period, with taps that generate the different phases. This circuit uses more power and chip area than a delay-line synchronizer (e.g. Ref. 8: 16 mW, 21000 μm^2 , 0.4 μm technology), but it is shared by the synchronizers of all links. Using delay-line synchronizers becomes less beneficial when the number of parallel links increases.

3.2. Determination of the phase

The phase of a signal can only be observed by looking at its edges. This means that the signal needs transitions, also in the worst case that a long series of identical bits are transmitted. Encoding (e.g. with a DC-balanced

Figure 9. FIFO-synchronizer.

code) can solve this problem but uses extra chip area and introduces latency.[†] Another possibility, if it is permitted by the skew, is to use a reference signal, with transitions, transmitted by an additional link. For a serial link this is unacceptable as it doubles the used optical hardware whereas for parallel links it is acceptable or even necessary when clocked receivers are used. Some phase detectors need a well defined signal while others just need a signal with transitions, which can also be useful for other purposes, e.g. to transmit a PRBS (pseudo random binary sequence) when data scrambling is used.⁴ If it is certain that a data signal contains enough transitions for synchronization purposes, it can be used as a reference signal for parallel links.

A few methods to determine the phase of a reference link are described in Sect.4.

Oversampling vs. tracking

Until now the data was sampled only once per period. The edges on a data or reference signal were observed and followed to get the right sampling moment. This is called tracking. One wrong sample meant the loss of a bit.

Another possibility is oversampling. Each data signal is sampled at least three times per period at fixed moments. By comparing the samples of different bits the phase and the data can be reconstructed, e.g. by majority voting or by choosing the bit in the middle between the transitions.¹³ If a sample is wrong the other samples will help to find the correct value and no bit is lost. As a result this method is less sensitive to high frequent jitter. However, oversampling uses more power and die area and introduces quantization jitter.^{8, 14}

The method used in the designs and simulations (Sect. 4) is somewhere in between. A reference signal is oversampled to determine the phase and the data is sampled once per period.

4. DESIGNS AND SIMULATIONS

To control if the theoretical results mentioned above can be realized and to get an indication of the cost, designs have to be made. Each circuit was designed using the standard libraries of the AMS 0.6 technology. No layout was drawn. The mentioned die area (table 3) is the sum of the area of the used standard cells without interconnection. This will only give an idea of the used area. In the same way the mentioned power consumption only gives a rough idea of the power consumption of the synchronizers.

4.1. Using Artificial Jitter Injection (AJTI)

Two designs were made: the first was found in Ref. 15 and extended from two to four clock phases with the second design as a result.

The reference signal toggles each period. When it is sampled by a clock phase, the samples will alternate between 0 and 1. If the pulses of the reference signal are made wider or narrower (artificial jitter injection) the samples will remain the same except for the clock phases that are too close to the edges of the reference signal (Fig. 2). Those clock phases are not suited to sample the data and another clock phase will be chosen as auxiliary clock.

The AJTI introduces a delay in the reference signal. To obtain correct phase information the data signals are also delayed (delay match).

4.2. Using flip-flop phase detectors

In this design the reference signal can be arbitrary, as long as it contains *enough* transitions. The flip-flop phase detector on Fig. 11 compares the two last samples of the reference signal taken on the negative edge of clock phase *Aclk*. *T* will become high if a transition took place and *Y* indicates if the last transition took place before or after the rising edge of *Aclk*.

Two phase detectors are used: one with *C1* and the other with *C2* as clock input (*Aclk*). This way it is known if a transition lies before (zone 3 and 4 on Fig. 10) or after (zone 1 and 2) the rising edge of *C1* and before (zone 4 and 1) or after (zone 2 and 3) the rising edge of *C2*. This way it is determined in which zone

[†]This may influence the complexity and behavior of the receiver, what can reduce these disadvantages.

Figure 10. Synchronization with flip-flop phase detectors.

Figure 11. A flip-flop phase detector.

the transition lies (zone 2 on the figure) and which phases can be used to sample the data (C1 and C4, the phase closest to the former auxiliary clock will be chosen as the new auxiliary clock). The distance between the auxiliary clock and the reference edge will be at least $T/4$ ($T/2 - T/n$ in general) so the distance from the center of the eye will be less than $T/4$, T/n in general instead of the $T/2n$ used to deduce Eqs. 2 and 1.

4.3. Comparison of the designs.

The occupied die area is split in a basis, independent of the number of parallel links, and a part used for each link. The area per link of the AJTI is a little bit larger because of the delay match. Multi-register synchronization instead of an auxiliary clock would use $6230 \mu m^2$ per link. The power consumption is similar for both techniques and proportional to the frequency.

The AJTI-synchronizer needs a well defined reference signal while the phase detector synchronizer permits much more flexibility. The injected artificial jitter is independent of the used frequency (the ratio J/T is variable). As a result the circuit is designed for a limited frequency region. The phase detector synchronizer does not have this limitation.

Table 3. Area of components used in synchronization circuits and power consumption.

Area. (μm^2)	basis	per link	P (50 MHz)	P(100 MHz)
AJTI (2 phases)	9181	2194	1 mW	2 mW
AJTI (4 phases)	17469	2194	2.5 mW	4.2 mW
Phase detectors (4 phases)	20569	1573	1.8 mW	4.8 mW

5. CONCLUSION

In this paper several synchronization methods for parallel mesochronous optical interconnections were described. Which of the options is optimal depends on the properties of the optical links. The choice made for a type of link is dependent on the choice made for a certain synchronizer and vice versa.

If the skew is sufficiently small an additional link, used to send a reference signal, can simplify the synchronization circuits. Therefore it would be desirable that (research) efforts were made to improve the uniformity between parallel links.

REFERENCES

1. J. M. Van Campenhout, "Computing structures and optical interconnect: friends or foes?," in *proceedings of SPIE*, **4109**, pp. 206–216, 2000.
2. T. J. Chaney and C. E. Molnar, "Anomalous behavior of synchronizer and arbiter circuits," *IEEE Transactions on computers* **C-22**, pp. 421–422, apr 1973.
3. D. M. Cutrer and K. Y. Lau, "Ultralow threshold laser - how low a threshold is low enough?," *IEEE Photonics Technology Letters* **7**, Jan 1995.
4. J. Yang, J. Choi, D. M. Kuchta, K. G. Stawiasz, P. Pepeljugoski, and H. A. Ainspan, "3.3 V 500 Mb/s/ch parallel optical receiver," *IEEE Journal of Solid-state circuits* **33**, pp. 2197–2204, Dec 1998.
5. P. Pepeljugoski, D. Kuchta, and J. Crow, "Effect of bit-rate, bias and threshold currents on the turn-on timing jitter in lasers with uncoded and coded waveforms," *IEEE Photonics Technology Letters* **8**, Mar 1996.
6. H. Neefs, *Latency control in processors and implications on the opportunity of optical interconnects (in Dutch: Latentiebeheersing in Processors en implicaties op de Opportuniteit van Optische Interconnecties)*. PhD thesis, University of Ghent, 1999-2000.
7. W. J. Dally and J. W. Poulton, *Digital systems engineering*, Cambridge university press, 1998.
8. T. Aytur, J. Gebis, J. Golbus, B. Gribstad, C.-W. Lee, and J. Tuan, "The design of a high speed serial link for IRAM," aug 1997.
9. IMEC-ELIS, "Generic approach to manufacturable optoelectronic interconnects for vlsi circuits. OIIC ES-PRIT MEL-ARI project No.22641 work package 8 deliverable 8.13: Report on tested system demonstrator." commercial in confidence.
10. H. Neefs, "Optoelectronic interconnects for integrated circuits. Achievements 1996-2000," tech. rep., RUG, june 2000. Advanced research initiative in microelectronics MEL-ARI OPTO.
11. M. A. R. Initiative, "Technology roadmap optoelectronic interconnects for integrated circuits.," September 1999. MEL-ARI OPTO.
12. J. Van Campenhout, H. Van Marck, J. Depreitere, and J. Dambre, "Optoelectronic FPGA's," *IEEE Journal of Selected Topics in Quantum Electronics* **5**, pp. 306–315, March/April 1999.
13. S. Kim, K. Lee, Y. Moon, D.-K. Jeong, Y. Choi, and H. K. Lim, "A 960-Mb/s/pin interface for skew-tolerant bus using low jitter PLL.," *IEEE Journal of Solid-state circuits* **32**, pp. 691–700, May 1997.
14. J. Golbus, "Design of a 160 mW, 1 Gigabit/second, serial I/O link," nov 1998. URL = cite-seer.nj.nec.com/259300.html.
15. F. Mu and C. Svensson, "Self-tested self-synchronization circuit for mesochronous clocking," *IEEE Transactions on circuits and systems II: analog and digital signal processing* **48**, pp. 129–140, Feb 2001.