29 November 1996

Comparison of Completion-Detection Techniques for Asynchronous Circuits

V.A. Bartlett, E. Grass

School of Electronic and Manufacturing Systems Engineering
University of Westminster
115 New Cavendish Street
London W1M 8JS

1. Introduction

Whereas in synchronous systems the clock frequency governs the speed of the whole circuit, which in turn is determined by its critical path, asynchronous circuits use Completion-Detection (CD) methods to determine when a Combinational Logic block (CL) has completed its operation. The signal generated by the completion-detection circuitry, indicating validity of the CL's outputs, is used by the handshake protocol to control the transfer of data tokens between consecutive stages. This paper focuses on the methodology of completion-detection. Naturally the CD mechanism has some repercussions on the CL itself and in some cases becomes an integral part of it. Several methods of completion-detection are known and will be considered in Section 2 whilst Section 3 deals with a relatively new class of single-rail completion-detection techniques which will be considered in more detail and compared with each other.

2. Review of Known Completion-Detection Techniques

Double-Rail Coding

Single rail logic itself has no capacity to carry information on the sequence or validity of data. Since a sequence of data values need not contain any transitions, there is no mechanism that allows generation of a completion signal by monitoring only the data outputs of a combinational logic block. A possible solution to this problem is to introduce an additional signal state (space) and to ensure that this state occurs between all signal transitions. Furthermore, in order to ensure the unambiguous identification of processed data, the logic has to be hazard-free. Under these constraints, the availability of a new data token at the output of a CL is signalled by a transition from 'space' (00) to 'data-1' (10) or 'data-0' (01).

These double-rail coding techniques are currently widely used but come at some cost in terms of hardware overhead, speed and energy consumption. Double-rail circuits are often implemented in Differential Cascode Voltage Switch Logic (DCVSL). DCVSL has the benefit that in the evaluation phase, signals are established in both logic polarities thereby enabling a single NAND gate to be used to generate a completion signal for each signal pair [1].

Ternary Logic

The principle behind the use of ternary logic is similar to that of double-rail coding in which invalid data is indicated by means of a third state. With ternary logic however, the three signal states, 'one' ,'zero' and 'space' are represented by three different voltage levels. In common with the double-rail technique, the same transition restrictions apply and again the logic must operate hazard free. An implementation of the ternary logic approach is investigated in [2]. Its disadvantages include a considerable silicon overhead and the requirement for two supply rails.

Bounded Delay

The bounded delay approach is the easiest method of completion-detection to implement (although strictly speaking, this approach is not a completion-detection method at all since the validity of data is solely dependent on timing assumptions). This method requires a priori knowledge of the CL's critical path and uses a delay circuit whose input is the START signal and whose output indicates completion. The delay is tuned such that completion is indicated after the critical path delay of the CL has expired. This technique can be used to advantage for CLs with fairly short critical paths.

Auxiliary Function

This approach to completion-detection is based on the idea of modelling the data-dependent delay of the CL with an additional circuit which generates the RDY signal. Often internal signals within the CL are used to support the operation of this auxiliary circuit [3]. The additional silicon area required is often relatively high.

Current-Sensing Completion-Detection

Current-Sensing Completion-Detection (CSCD) is an approach to self-timing which exploits the fact that the operating current of CMOS circuits is several magnitudes higher than its quiescent current [4] [5]. Hence by monitoring the supply current of a CL, a decision can be made as to whether the circuit is in a transient or steady state.

Unfortunately, the proposed circuits in both [4] and [5] have the major disadvantage that their quiescent currents are substantially increased in comparison with synchronous circuits as a result of the high power consumed by the CSC itself. Several variants of this technique have been reported but most suffer from high power consumption or require an extra supply voltage. However the disadvantages can be largely overcome by application to BiCMOS circuits if the inherently high processing costs can be tolerated.

3. Novel Completion-Detection Techniques

3.1 Activity-Monitoring Completion-Detection (AMCD)

AMCD is a single rail method which exploits the fact that if no transitions on internal signals occur within a certain period the circuit must have reached a steady state. Activity Monitors (AMs) are connected to internal signals of the CL [6]. Their Open-Drain outputs are connected together and, when fitted with a pull-up device, form a “wired-OR” circuit whose output, ACT_L, remains low while the CL is in a transient state.

The AM implementation shown in Figure 1 has been found to be an optimal solution with respect to symmetrical operation, speed and silicon-area and allows adjustment of the sensitivity of the circuit to H-L-H glitches and L-H-L glitches independently. If such independent control (effected by altering the inverter thresholds appropriately) is not required then the AM in [6] can be used which presents a lower capacitive load to the signal being monitored.

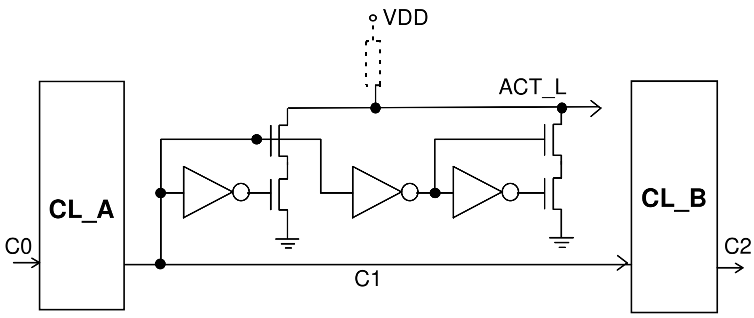

Figure 1 Circuit of an Activity Monitor attached to signal C1

Monitoring every internal signal of a CL can result in a substantial increase in circuit area and normally AMs are attached to certain exposed signals only.

For a ripple-carry adder an appropriate place to attach AMs is to each carry signal. For reliable operation, the pulse width (t_p) generated by each AM must cover the critical path delay of a one-bit full adder plus the switch-on delay of the subsequent AM. In practice, as with synchronous circuits, a safety margin of up to 100% may need to be added to the nominal value of t_p. The AM itself requires approximately half the area of one full-adder.

A crucial issue with circuits based on AMCD is their response to glitches on signals within the CL. This problem is discussed in more detail in the following section. A number of simulations have been carried out to ascertain whether glitches that might be propagated to the primary outputs of the CL are safely detected by the AMs; for certain classes of circuits the above strategy is sufficiently safe.

3.2 Transition-Monitoring Completion-Detection (TMCD)

The AMCD techniques described above as well as the CSCD technique discussed previously have a potential problem in the presence of glitches on the inputs to sub-circuits. It is essential that any glitch of sufficient magnitude to be propagated within the CL is also detected by the Completion-Detection Circuit (CDC) otherwise completion will be detected erroneously. Therefore the CDC has to be more sensitive to glitches than the actual CL. However, for some classes of circuits, it is conceivable that glitches occur at a number of inputs simultaneously and whilst no single glitch is sufficient to trigger a CDC, the (dynamic or static) combination of glitches causes a hazard within the CL. This hazard might then be propagated to the primary outputs of the circuit and could cause a malfunction [7] [8]. (The situation described cannot occur in ripple-carry adders since every sub-circuit has only one (carry) input on which glitches can be present). A possible, but somewhat limited, solution to this potential problem requires either increasing the sensitivity of the CDC (making its design relatively difficult and incurring extra penalties in terms of delays and energy consumption) and/or decreasing the sensitivity of the downstream sub-circuit, inevitably causing extra delays.

An alternative solution called Transition-Monitoring Completion-Detection (TMCD) has been developed by the authors. This is a variant of AMCD incorporating combined glitch detection and suppression circuitry. TMCD is a strategy for monitoring transitions on a signal which guarantees that no glitch can be passed, undetected, to the primary outputs of the CL. This is achieved by use of an 'interlock' circuit whose function is to inhibit the propagation of new data values until their presence has been detected. Since the only data that can be propagated is that which has already been detected, the possibility of undetected glitches propagating is eliminated. An implementation of such a Transition Monitor (TM) is shown in Figure 2.

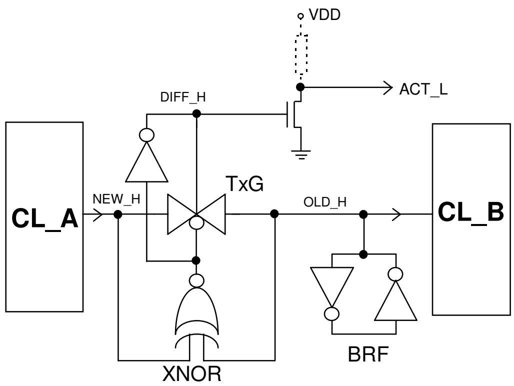

Figure 2 Transition Monitor (TM) with glitch suppression

Data propagation is inhibited by insertion of transmission gate TxG into the signal path on which activity is to be monitored. Since the input, OLD_H, to the downstream logic, CL_B, is under certain circumstances, now undriven a Bus-Retention Flip-flop (BRF), is added to retain the previous signal value. New data appearing on the output of the combinational logic block CL_A is compared with the old value by means of an exclusive-NOR gate, XNOR. If the two values are different TxG turns on allowing the new data to propagate. Conversely, if the two values are the same then TxG is switched off.

In the interests of subsystem speed, the propagation of data must be as fast as possible which requires that the H-L propagation delay of XNOR must also be fast. The actual ACT pulse duration can be controlled by adjusting a number of circuit parameters.

The Transition Monitor has been laid out using a 1.0 micron n-well CMOS process and its area is comparable to that of one full-adder circuit. Simulations after parasitic extraction show the delay added to the signal path between circuits is approximately 1.2ns and 1.8ns for H-L and L-H transitions respectively.

Extensive simulations of the circuit's response to dynamic glitches (H-L-H-L and L-H-L-H transitions) have shown safe operation under all conditions. Due to its feedback structure the circuit can be easily designed such that it is more sensitive to glitches than the downstream combinational logic. However, its disadvantages are a relative high additional silicon area and an increased energy consumption.

3.3 Performance Summary

The performance of the AMCD and TMCD approaches was measured using a 16-bit ripple-carry adder as a benchmark and a synchronous circuit, using the same full-adder cells, for comparison. Average case performance was measured assuming random input data and the average delays show the benefits of data dependent operation of asynchronous circuits.

The AMCD version incurs an area overhead of around 20% and degrades the worst case delay by a similar amount. Average speeds, however, show significant improvement (a factor of more than 3 when allowing a 100% safety margin) over the synchronous circuit albeit at some cost in terms of energy consumption.

The more glitch-robust TMCD version, incurs an area overhead of around 50% and suffers from worst case delays of roughly double those of the synchronous circuit. This is mainly because, in contrast to AMCD, TMs interrupt the adder's critical carry path. Average case delays, however, are shorter by a factor of about 3 but at a significant cost in terms of energy consumption.

A technique combining AMCD with minimal use of TMCD offers the potential of glitch-robust operation without the costs incurred with pure TMCD and is currently under investigation.

REFERENCES

[1] T.H.-Y. Meng, R.W. Broderson, D.G. Messerschmidt — 'Automatic Synthesis of Asynchronous Circuits from High-Level Specifications'. IEEE Transactions on Computer-Aided Design, Vol. 8, Nov. 1989, pp1185-1205

[2] Chung-Yu Wu, Hong-Yi Huang, — 'Design and Application of Pipelined Dynamic CMOS Ternary Logic and Simple Ternary Differential Logic', IEEE Journal of Solid-State Circuits, Vol. 28, No.8, Aug. 1993, pp.895-906.

[3] V. Varshavsky, V. Marakhovsy, M. Tsukisaka, — 'Data-Controlled Delays in the Asynchronous Design', Proc. of the 1996 IEEE International Symposium on Circuits and Systems (ISCAS'96), Atlanta, USA, May 1996, pp. 153-155.

[4] Dean, M. E. Dill, D. L. Horowitz, M. — 'Self-Timed Logic Using Current-Sensing Completion Detection (CSCD)', Journal of VLSI Signal Processing, 1994, Vol. 7, (1), pp. 7-16

[5] Izosimov, O. A. Shagurin, I. I. Tsylyov, V. V. — "Physical approach to CMOS Self-Timing", *Electronic Letters*, Vol. 26, pp. 1835-1836,1990.

[6] E. Grass, R.C.S. Morling, I. Kale — "Activity-Monitoring Completion-Detection (AMCD): A new single rail approach to achieve self-timing", Proc. of the IEEE 2nd international symposium on Advanced Research in Asynchronous Circuits and Systems, Aizu (Japan), March 18-21, 1996, pp. 143-149

[7] S. Nowick — Private Communication, March 1996.

[8] S. H. Unger — 'The Essence of Logic Circuits' Prentice-Hall 1989.