

# Performance analysis and optimisation of NCL self-timed rings

W. Kuang, J.S. Yuan, R.F. DeMara, M. Hagedorn and K. Fant

**Abstract:** A self-timed ring using NULL convention logic (NCL) is presented. An analytical method to evaluate the speed of NCL rings has been developed. The analytical predictions are verified by a Synopsys simulation and excellent agreement between the theoretical predictions and simulation results is obtained. Some important principles for ring optimisation are obtained. The analysis leads to the speed optimisation of a 24-bit NCL divider.

## 1 Introduction

In recent years, the implementation of floating-point division has received increasing attention. The computation time of a division operation is fairly long compared with other floating-point operations such as multiplication and addition. Moreover, applications such as three-dimensional graphics use division frequently. It has been shown that ignoring its implementation can result in significant system performance degradation for many applications [1].

A self-timed ring is an efficient approach to the implementation of division based on the SRT algorithm [2–5], which is the most commonly used algorithm in modern processors [1]. In this ring structure, the computation can progress as fast as the circuit configuration and fabrication technology allow [6], rather than being slowed down by the worst-case delay or clock skew margins found in synchronous counterparts. Furthermore, due to the identity of the iterative operations in division, a self-timed ring with a few stages can be used to perform a floating-point mantissa division operation without speed loss compared with the fully expended pipeline version. One of the early designs was proposed by Williams [3]. Since then, several modifications have been suggested in order to improve performance [4, 5]. However, all of these methods use differential cascode voltage switch logic (DCVSL) [6]. Some conventional CMOS design principles are not applicable to DCVSL. For example, the rise and fall times of output nodes in DCVSL circuits are generally different, due to an inherent asymmetry in NMOS tree and PMOS load. Careful sizing of transistors is required for DCVSL to avoid possible problems associated with races [8]. For DCVSL, a timing assumption appears during the pre-charge phase. It is assumed that, when request signal goes low, the input data are all reset. Yet, this constraint is not always checked for validity in practice [9].

The method for self-timed ring design presented in this paper is based on NULL convention logic (NCL), which is a new clockless methodology for digital system design [10, 11]. It offers a quasi delay-insensitive logic paradigm, where control is inherent with each datum, and follows the so-called ‘weak conditions’ of Seitz’s delay-insensitive signalling scheme [12]. The NCL paradigm assumes that forks in wires are isochronic [13]. Besides the common advantages of general asynchronous circuits, such as being speed-independent, the average-case delay, low power, low noise and electromagnetic interference (EMI), a formal and systematic method of design and optimisation has been developed to make the design of NCL circuits straightforward [14, 15]. A complete register-transfer level (RTL) design flow for NCL can be performed using commercial CAD tools [16], such as Synopsys and Modelsim VHDL simulator.

However, the speed of an NCL ring circuit depends not only on the latencies of individual stages in the data path, but also on the speeds of handshake circuits and the number of stages. To investigate how these factors affect the ring speed, an analytical method is developed using dependency graphs. Performance analysis of self-timed rings yields important implications for designing NCL rings with optimal speed performance or with an appropriate speed–area tradeoff. These implications are applied to the design of a 24-bit floating-point mantissa division.

## 2 Overview of null convention logic

NCL uses symbolic completeness of expression to achieve self-timed behaviour [10]. A symbolically complete expression is defined as an expression that only depends on the relationships of the symbols presented in the expression without a reference to the time of evaluation. In general, a multi-rail signal can be used to incorporate data and control information into one mixed signal path to eliminate the time reference, and therefore to form a symbolically complete expression. Typically, a dual-rail signal  $D$  consists of two wires,  $D_0$  and  $D_1$ , which represent a value from the set {Data0, Data1, Null}, shown in Table 1. The Data0 state ( $D_0 = 1, D_1 = 0$ ) corresponds to a Boolean logic 0 value, the Data1 state ( $D_0 = 0, D_1 = 1$ ) corresponds to a Boolean logic 1 value, and Null state ( $D_0 = 0, D_1 = 0$ ) corresponds to the empty set, meaning that the value of  $D$  is not yet available. The state ( $D_0 = 1, D_1 = 1$ ) is forbidden.

The primitive element for NCL is an  $M$ -of- $N$  threshold gate with hysteresis. This gate has  $N$  input signals and one

© IEE, 2003

IEE Proceedings online no. 20030343

doi:10.1049/jip-cds:20030343

Paper first received 2nd January and in revised form 1st October 2002

W. Kuang, J.S. Yuan and R.F. DeMara are with the School of Electrical Engineering and Computer Science, University of Central Florida, Orlando, FL, 32816, USA

M. Hagedorn is with Theseus Logic, Inc., Maitland, FL, 32751, USA

K. Fant is with Theseus Research, San Jose, CA, 95054, USA


**Table 1: Dual-rail encoding**

Logic value	Encoding	
	$D_0$	$D_1$
Data1	0	1
Data0	1	0
Null	0	0
Invalid	1	1

output, shown in Fig. 1a. The signal values are referred to as DATA and NULL for a high voltage level and a low voltage level, respectively. This gate has two important properties, threshold behaviour and hysteresis behaviour. The threshold behaviour requires that the output becomes DATA if at least  $M$  of the  $N$  inputs have become DATA. The hysteresis behaviour requires that the output only changes after a sufficiently complete set of input values has been established. In the case of a transition to DATA, the output remains at NULL until at least  $M$  of the  $N$  inputs become DATA. In the case of a transition to NULL, the output remains at DATA until all  $N$  inputs become NULL. As special examples, an  $N$ -of- $N$  gate is an  $N$ -input Muller C-element while a 1-of- $N$  gate corresponds to an  $N$ -input OR gate. Transistor-level design for  $M$ -of- $N$  threshold gates is described in [17]. As an example, a 2-of-3 gate having inputs A, B and C is shown in Fig. 1b. It includes four blocks (go to NULL, go to DATA, hold NULL, hold DATA), one inverter and two feedback transistors (pMOS for hold NULL, nMOS for hold DATA).


Threshold gates can be used to build combinational circuits. An NCL half-adder is shown in Fig. 2, where  $(x_0, x_1)$  and  $(y_0, y_1)$  denote the dual-rail encoded input

adds.  $(c_0, c_1)$  and  $(s_0, s_1)$  denote the dual-rail encoded carry and sum outputs, respectively.

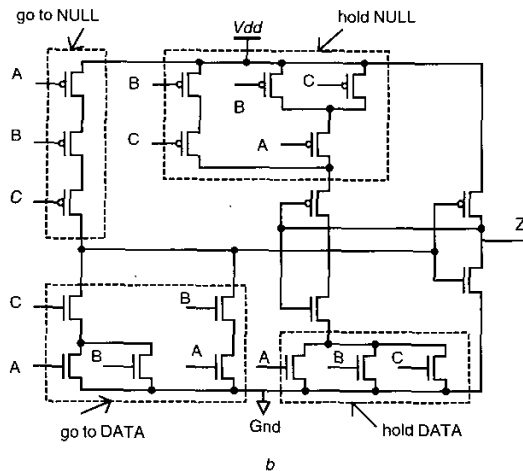


**Fig. 2** NCL half-adder

A typical NCL circuit structure forms a pipeline, in which each stage consists of a register, completion detection circuit and a combinational circuit [10]. An NCL ring is formed when the output of a pipeline is connected to its input as shown in Fig. 3. For simplicity, a half-adder is used as the combinational circuit for each stage in the ring. Note that, although one of the three bits in the data bus is not involved in the half-adder computation, it puts a load on the completion detection block. All of the registers are initialised to Null outputs except that the outputs of register  $R_N$  are initialised to Data0 or Data1. Fig. 4 illustrates the gate-level structures of register and completion detection circuit used in the ring. Th22nx0 (Th22dx0) is a 2-of-2 threshold gate that is initialised to NULL (DATA). Th12bx0 is a 1-of-2 threshold gate with an inverter following. Th22x0 is a 2-of-2 threshold gate without initialisation control.




a



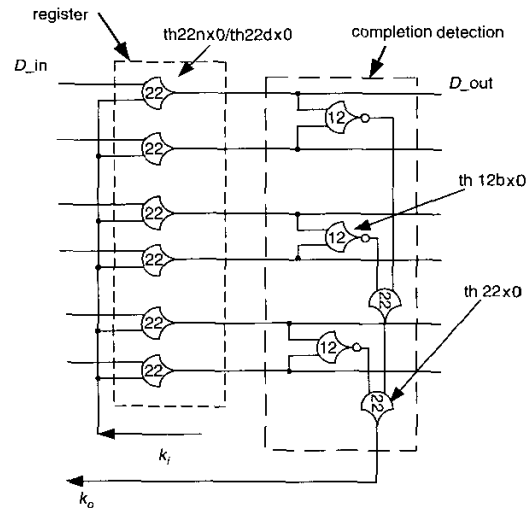
b

**Fig. 1** Threshold gate with hysteresis

- a Symbol of  $M$ -of- $N$  gate
- b Schematic of 2-of-3 gate



**Fig. 3** NCL ring each stage of which is a half-adder



**Fig. 4** 3-bit register and completion detection

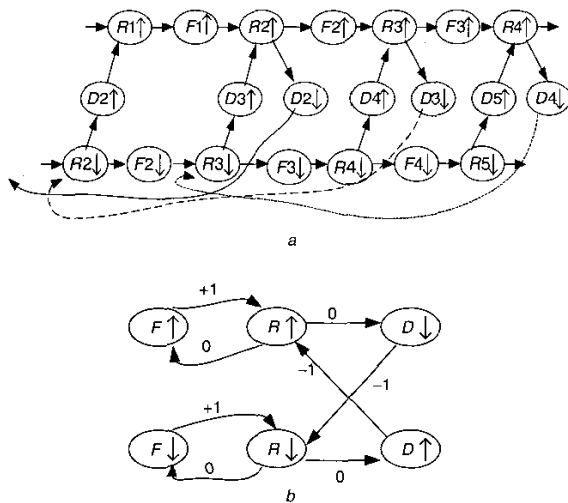
### 3 Analysis of rings using dependency graphs

An efficient analysis method for ring speed performance is based on the use of a dependency graph, which was proposed by Williams [2, 18]. Williams' theory on rings shows how to analyse a ring performance in terms of speed and area. This theory will be used below to develop an analytical model for evaluating self-timed rings implemented in NCL.

First, several parameters need to be defined. Local parameters include the forward latency ( $L_f$ ), the reverse latency ( $L_r$ ), and the local cycle time ( $P$ ). The forward latency specifies the delay from valid data outputs at one stage to valid data outputs at the following stage without waiting for request signal. The reverse latency specifies the delay from the request of a stage output to the request of its predecessor's output without waiting for Data or Null in the data path. The local cycle time specifies the delays of all the transitions necessary for a stage to pass a token, and become enabled again for the next token. This limits the maximum throughput of a ring. A key global parameter is the total latency  $\lambda$ , that is the delay between the introduction of a new data token into the ring and the removal of the corresponding processed token after the token has passed through stages necessary for solution of an iterative problem. Since the number of iterations for a given problem is constant, and each iteration corresponds to one stage, the total latency is proportional to the average time for a token to pass through one stage. The average propagation delay of a single stage will be used as a metric of ring speed.

#### 3.1 Folded dependency graph

A dependency graph of a pipeline can be constructed from its structure. Fig. 5a shows the dependency graph of a pipeline corresponding to Fig. 3. The nodes of the graph correspond to specific rising (up arrow) or falling (down arrow) transitions of circuit components ( $R$  refers to register,  $F$  refers to combinational function block for one stage,  $D$  refers to completion detection), and the edges represent the dependencies of each transition on the outputs of other components. The delay of each transition is represented by a value attached to the node in the graph.



**Fig. 5** Dependency graphs for NCL pipeline  
a Extended dependency graph  
b Folded dependency graph

Dependency graphs can be used to determine both the stage latencies (forward and reverse) and the local cycle time [2]. When the stages are identical, the folded dependency graph can be used to calculate the local parameters more conveniently as shown in Fig. 5b. In the folded dependency graph, it is not necessary to subscript the nodes with a particular stage index since each node represents that same transition in all stages. For each edge, there is a stage index offset attached, giving the offset in stage index to which the dependency refers.

The local parameters can be calculated by the following equations [18]:

$$L_f = \max \left( \frac{\sum_{i \in path} t_i}{\sum_{i \in path} w_i} \right) \text{ over all non-repeating cyclic paths}$$

$$\text{with } 0 < \sum_{i \in path} w_i, w_i \geq 0$$

$$L_r = \max \left( \frac{\sum_{i \in path} t_i}{-\sum_{i \in path} w_i} \right) \text{ over all non-repeating cyclic paths}$$

$$\text{with } 0 > \sum_{i \in path} w_i, w_i \leq 0$$

$$L_r = \max \left( \sum_{i \in path} t_i \right) \text{ over all non-repeating cyclic paths}$$

$$\text{with } 0 = \sum_{i \in path} w_i$$

where  $t_i$  is the  $i$ th node transition delay in the cyclic path chosen and  $w_i$  is the corresponding stage index offset.

#### 3.2 Local analysis

As a particular example, the latencies and the local cycle time of NCL pipeline can be analysed, based on folded dependency graph. The forward latency is given by

$$L_f = \max(F \uparrow + R \uparrow, F \downarrow + R \downarrow) \quad (1)$$

The reverse latency is given by

$$L_r = \frac{1}{2}(R \uparrow + D \downarrow + R \downarrow + D \uparrow) \quad (2)$$

To determine the local cycle time, the longest cycle with zero offset needs to be found. There are three possibilities for the longest zero offset cycle in Fig. 5b as follows:

$$\begin{aligned} P1 &= F \uparrow, R \uparrow, D \downarrow, R \downarrow, D \uparrow, R \uparrow, F \uparrow, R \uparrow \\ P2 &= F \uparrow, R \uparrow, D \downarrow, R \downarrow, F \downarrow, R \downarrow, D \uparrow, R \uparrow \\ P3 &= F \downarrow, R \downarrow, D \uparrow, R \uparrow, D \downarrow, R \downarrow, F \downarrow, R \downarrow \end{aligned}$$

The lengths of the three paths are

$$P1 = 2F \uparrow + 2R \uparrow + R \uparrow + R \downarrow + D \uparrow + D \downarrow \quad (3a)$$

$$P2 = F \uparrow + F \downarrow + R \uparrow + R \downarrow + R \uparrow + R \downarrow + D \uparrow + D \downarrow \quad (3b)$$

$$P3 = 2F \downarrow + 2R \downarrow + R \uparrow + R \downarrow + D \uparrow + D \downarrow \quad (3c)$$

Since the sum of the first four terms of  $P2$  is the middle point of the sum of the first two terms of  $P1$  and the sum of the first two terms of  $P3$ , the local cycle time is:

$$P = \max(P1, P2, P3) = \max(P2, P3) \geq P2 \quad (4)$$

From (1), (2), and (4), the local cycle time is obtained in terms of the forward and reverse latencies:

$$P = 2(L_f + L_r) \quad (5)$$

### 3.3 Ring performance analysis

The local cycle time of an NCL ring satisfies  $P = 2(L_f + L_r)$ , which means that the token flow rate is limited either by the forward latency (data-limited) or by the reverse latency (bubble-limited) [18]. We consider the case that there is only one token in an  $N$ -stage ring. There are  $(N-2)$  bubbles in the  $N$ -stage ring. When  $N$  is larger, there are so many bubbles in the ring that the token circulation period will fall in the data-limited region, and

$$P_{ring} = N \cdot L_f \quad (6)$$

When  $N$  is smaller, there are so few bubbles that the token circulation period will fall in the bubble-limited region [18], and

$$P_{ring} = \frac{2N \cdot L_r}{N-2} \quad (7)$$

In the data-limited region, the delay of the completion detection is removed from the critical path. If two local parameters are given, the optimal number of stages is defined as the minimum  $N$  in the data-limited region. It is achieved when (6) and (7) are equal to

$$N_{optimal} = 2 \left( 1 + \frac{L_r}{L_f} \right) \quad (8)$$

The final expression of the token circulation period is given by

$$P_{ring} = \begin{cases} \frac{2N \cdot L_r}{N-2}, & N \leq 2 \left( 1 + \frac{L_r}{L_f} \right) \\ N \cdot L_f, & N > 2 \left( 1 + \frac{L_r}{L_f} \right) \end{cases} \quad (9)$$

Therefore, the average time for a token to pass through one stage is

$$t_{stage} = \begin{cases} \frac{2 \cdot L_r}{N-2}, & N \leq 2 \left( 1 + \frac{L_r}{L_f} \right) \\ L_f, & N > 2 \left( 1 + \frac{L_r}{L_f} \right) \end{cases} \quad (10)$$

We see from (8) that, when  $L_r$  is decreased by speeding up the register and/or completion detection circuit, the number of stages required for optimal speed will be reduced. On the other hand, speeding up the forward data path will lead to an increase of the number of stages required for a higher optimal speed. It is also implied from (10) that a ring with too few stages will fall in the bubble-limited region, where the ring speed can be improved by speeding up the reverse latency and/or increasing the number of states until the optimal speed is achieved. A ring with more stages will fall in the data-limited region, where the ring speed is the optimal speed that depends only on the forward latency  $L_f$ .

### 3.4 Ring initialisation

A token introduction, as well as the characteristics of the threshold gates, requires a correct initialisation for NCL rings. Since the threshold gates used in combinational circuits have no reset input for initialisation, the validity of the initial output for each gate is guaranteed only by the completeness of input signals. However, dual-rail code (1,1) is forbidden in NCL. Therefore, it is convenient that all inputs of a combinational circuit are initialised as Null so that the circuit initial outputs are validly all Null. To realise this idea for NCL ring initialisation, an additional register with initialised Null output is inserted between the register with initialised Data output and the following combinational circuit, shown in Fig. 6.




Fig. 6 Ring with initialising stage ( $R_i, D_i$ )

For a ring initialised in this way, it is not necessarily the case that each stage performs identically. Based on the analysis in subsection 3.3, the additional register contributes a delay to the token circulation period in the data-limited region, while an additional bubble will modify the token circulation period in the bubble-limited region. The token circulation period for a ring with an initialising stage is given by

$$P'_{ring} = \begin{cases} \frac{2(N+1) \cdot L_r}{N-1}, & N \leq N'_{optimal} \\ N \cdot L_f + R, & N > N'_{optimal} \end{cases} \quad (11)$$

where  $R$  is the maximum of  $R \uparrow$  and  $R \downarrow$ . The optimal number of stages for a ring with an initialising stage is similarly defined as the value corresponding to the intersection of the bubble-limited region and data-limited region, which is given by

$$N'_{optimal} = \frac{L_f + 2L_r - R + \sqrt{(L_f + 2L_r - R)^2 + 4L_f(R + 2L_r)}}{2L_f} \quad (12)$$

Note that  $N'_{optimal}$  is slightly smaller than  $N_{optimal}$  in (8). The corresponding average time for a token to pass through one stage is given by

$$t'_{stage} = \begin{cases} \frac{2 \cdot \left( 1 + \frac{1}{N} \right) L_r}{N-1}, & N \leq N'_{optimal} \\ L_f + \frac{R}{N}, & N > N'_{optimal} \end{cases} \quad (13)$$

Noting the fact that  $R$  is usually small, (13) shows that the introduction of the initialising register improves the ring speed in bubble-limited region while leading to a negligible degradation of the ring speed in a data-limited region.

## 4 Simulation results and discussion

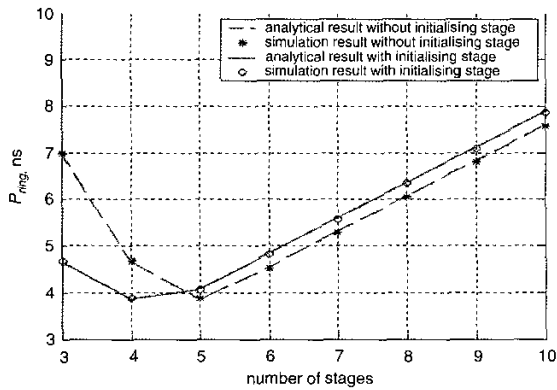
To verify the effectiveness of the above analysis, the ring circuits are simulated by VHDL code in Synopsys. The delays of threshold gates specified in the VHDL library are listed in Table 2. The three local parameters  $L_f$ ,  $L_r$ ,  $P$  based on the gate delays are also calculated in Table 2.

First, the ring performance can be evaluated by equations in Section 3 without running the VHDL simulation. The token circulation periods as a function of the number of stages are plotted in Fig. 7. The average times of one stage as a function of the number of stages are plotted in Fig. 8. In these Figures, the dashed lines correspond to (9) and (10), while solid lines correspond to (11) and (13).

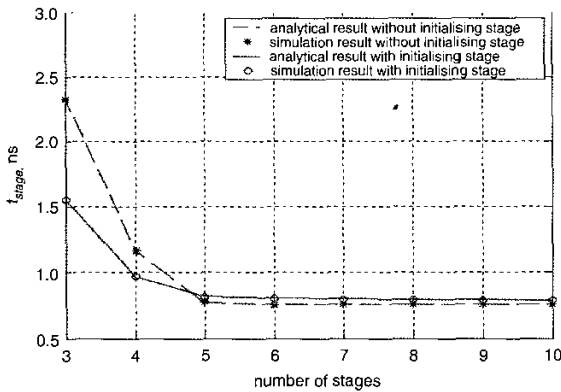
Second, by writing VHDL codes for ring circuits and running simulation, the token circulation period can be measured based on the output waveform and thus the ring performance can be obtained. The corresponding results are plotted in Figs. 7 and 8.

**Table 2: Gate delays and local parameters**

Gate delays/local parameters	Value, ns
Th22nx0↑ (R↑)	0.29
Th22nx0↓ (R↓)	0.16
Th12bx0↑ (D1↑)	0.45
Th12bx0↓ (D1↓)	0.22
Th22x0↑ (D2↑)	0.40
Th22x0↓ (D2↓)	0.20
Half adder	0.47
$L_f = (\text{half adder}) + R↑$	0.76
$L_r = 0.5(R↑ + R↓ + D1↑ + D1↓ + 2D2↑ + 2D2↓)$	1.16
$P = 2(L_f + L_r)$	3.84



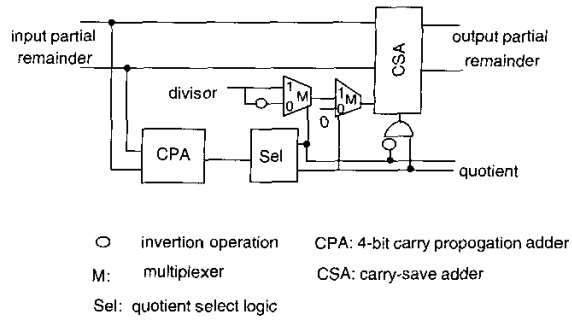
**Fig. 7** Circulation period against number of stages



**Fig. 8** Per-stage time against number of stages

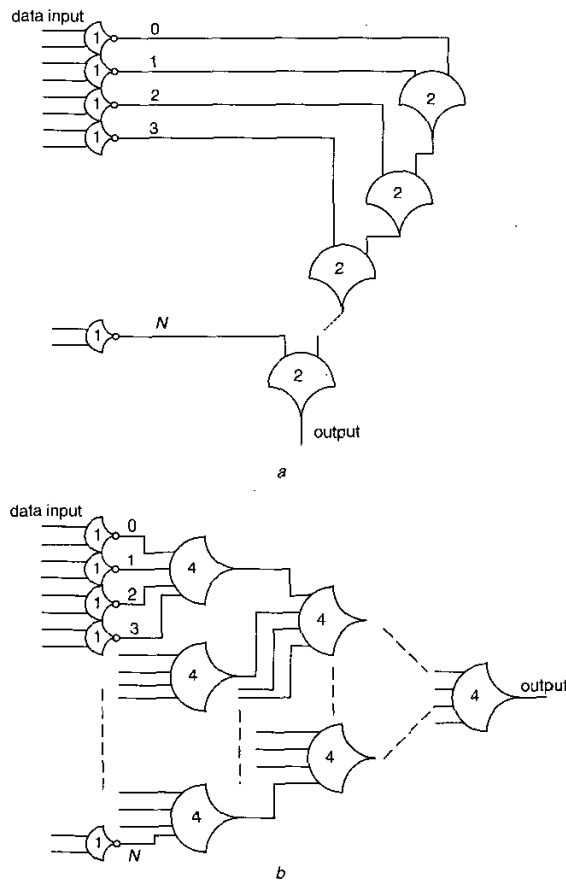
The analytical and simulation curves agree with each other very well. Although a discrepancy occurs due to data-dependent delay and the assumption that the half-adder has the same delay for rising and falling transition, it is negligible.

The results in the previous Section can be used as a guideline to Optimisation for SRT division circuit. A radix-2 SRT division algorithm can be implemented by designing the combinational circuits in the ring structure. Fig. 9 shows the combinational circuit for one stage of a 24-bit division. When the data wavefront arrives, the combinational circuit produces one bit of the quotient, and the partial remainder that will be used in the next stage. To speed up the calculation of partial remainder, the partial remainder is

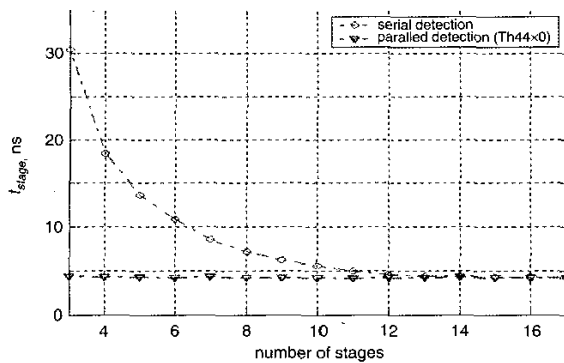


**Fig. 9** Combinational logic for one stage of radix-2 SRT division

usually represented by a carry-save redundant binary format [19]. The forward latency of the division ring is approximately 4.19 ns calculated from the delays of components. Each register in the ring passes 75-bit data, including 24-bit divisor, 25-bit carry and 24-bit sum for partial remainder, 1-bit for logic '1', 1-bit for logic '0'. If the completion detection circuit is implemented in series as shown in Fig. 10a, the reverse latency of 22.5 ns can be determined by the method in Table 2. From (12) and (13), a 13-stage ring will achieve the optimal speed 4.21 ns. A fast completion detection circuit can significantly reduce the number of stages for the optimal speed. A parallel implementation of the completion detection circuit, shown in Fig. 10b, can reduce the reverse latency to 1.865 ns, and



**Fig. 10** Completion detection circuits  
 a Serial architecture  
 b Parallel architecture



**Fig. 11** Simulation results for 24-bit divisions with different completion detection circuit

thus result in a 3-stage ring for the optimal speed. VHDL simulation results are plotted in Fig. 11.

## 5 Conclusions

From the analysis, as well as simulation, some important principles for optimal ring design are suggested. When the information of  $R$ ,  $F$ , and  $D$  delays is known, the optimal number of stages can be determined by (4), which means that further increase of the number of stages does not improve speed, and that a ring with fewer stages will become slower. The optimised speed only depends on the forward latency  $L_f$ , independent of the delay of completion detection block.

As for an optimised NCL ring, further speed improvements of components  $R$ ,  $F$ , and  $D$  provide different opportunities to further improve the ring overall performance. Improving register speed directly leads to increasing the overall ring speed, while the number of stages almost doesn't need to change. Speeding up the combinational circuit does not improve the ring overall speed unless the number of stages is increased. Speeding up the completion detector does not improve the ring speed, but it allows the number of stages to be reduced without speed loss.

## 6 References

- Oberman, S.F., and Flynn, M.J.: 'Division algorithms and implementations', *IEEE Trans. Comput.* 1997, **46**, (8), pp. 833-854
- Williams, T.E.: 'Performance of iterative computation in self-timed rings', *J. VLSI Signal Process.*, 1994, **7**, (1-2), pp. 17-31
- Williams, T.E., and Horowitz, M.A.: 'A zero overhead self-timed 160-ns 54-b CMOS divider', *IEEE J. Solid-State Circuits*, 1991, **26**, (11), pp. 1651-1661
- Matsubara, G., Ide, N., Tago, H., Suzuki, S., and Goto, N.: '30-ns 55-b shared radix division and square root using a self-timed circuit'. Proceedings of the 12th symposium on Computer arithmetic. 1995, pp. 98-105
- Matsubara, G., and Ide, N.: 'A low power zero-overhead self-timed division and square root unit combining a single-rail static circuit with a dual-rail dynamic circuit'. Proceedings of the 3rd International Symposium on Advanced research in asynchronous circuits and systems. 1997, pp. 198-209
- Molnar, C.E., Jones, I.W., Coates, W.S., Lexau, J.K., Fairbanks, S.M., and Sutherland, I.E.: 'Two FIFO ring performance experiments', *Proc. IEEE*, 1999, **87**, (2), pp. 297-307
- Heller, L.G., Griffin, W.R., Davis, J.W., and Thoma, N.G.: 'Cascade voltage switch logic: a differential CMOS logic family'. Proceedings of the IEEE International Solid State Circuits Conference, San Francisco, CA. 1984, pp. 16-17
- Chu, K.M., and Pulfrey, D.L.: 'A comparison of CMOS circuit techniques: differential cascode voltage switch logic versus conventional logic', *IEEE J. Solid-State Circuits*, 1987, **22**, (4), pp. 528-532
- Renaudin, M., Hassan, B.E., and Guyot, A.: 'A new asynchronous pipeline scheme application to the design of a self-timed ring divider', *IEEE J. Solid-State Circuits*, 1996, **31**, (7), pp. 1001-1013
- Fant, K.M., and Brandt, S.A.: 'NULL convention logic: a complete and consistent logic for asynchronous digital circuit synthesis'. Proceedings of International Conference on Application-specific systems, architectures and processors, 1996, pp. 261-273
- Fant, K.M., and Brandt, S.A.: 'Null convention logic system', US Patent 5,305,463, 19 April, 1994
- Seitz, C.L.: 'System timing', in: 'Introduction to VLSI Systems' (Addison-Wesley, 1980), pp. 218-262
- Berkel, K.V.: 'Beware the isochronic fork', *Integr. VLSI J.*, 1992, **13**, (2), pp. 103-128
- Smith, S.C.: 'Gate and throughput optimizations for null convention self-timed digital circuits'. PhD dissertation, Department of Electrical and Computer Engineering, University of Central Florida, 2001
- Smith, S.C., DeMara, R.F., Yuan, J.S., Hagedorn, M., and Ferguson, D.: 'Delay-insensitive gate-level pipelining', *Integr. VLSI J.*, 2001, **30**, (2), pp. 103-131
- Lighthart, M., Fant, K., Smith, R., Taubin, A., and Kondratyev, A.: 'Asynchronous design using commercial HDL synthesis tools'. Proceedings of the 16th International Symposium on Advanced research in asynchronous circuits and systems, 2000, pp. 114-125
- Sobelman, G.E., and Fant, K.: 'CMOS circuit design of threshold gates with hysteresis'. Proceedings of the 1998 IEEE International Symposium on Circuits and systems, 1998, Vol. 2, pp. 61-64
- Williams, T.E.: 'Self-timed rings and their application to division'. PhD dissertation, Stanford University, 1991
- Parhami, B.: 'Computer arithmetic: algorithms and hardware designs' (Oxford University Press, New York, 2000)