



MICROELECTRONIC ENGINEERING

An International Journal of Semiconductor Manufacturing Technology

Including Nanoelectronic Engineering
www.elsevier.com/locate/mee

Including:

**Special Issue on VLSI Design and Test
(VLSI 2003)**

Guest Editors:

Laurence T. Yang,
Jose G. Delgado-Frias,
Yiming Li,
Mohammed Niamat,
Dimitrios Soudris,
Srinivasa R Vemuru

This article was originally published in a journal published by Elsevier, and the attached copy is provided by Elsevier for the author's benefit and for the benefit of the author's institution, for non-commercial research and educational use including without limitation use in instruction at your institution, sending it to specific colleagues that you know, and providing a copy to your institution's administrator.

All other uses, reproduction and distribution, including without limitation commercial reprints, selling or licensing copies or access, or posting on open internet sites, your personal or institution's website or repository, are prohibited. For exceptions, permission may be sought for such use through Elsevier's permissions site at:

<http://www.elsevier.com/locate/permissionusematerial>



ELSEVIER

Available online at www.sciencedirect.com

ScienceDirect

Microelectronic Engineering 84 (2007) 280–287

MICROELECTRONIC
ENGINEERING

www.elsevier.com/locate/mee

Design and characterization of NULL convention arithmetic logic units

Satish K. Bandapati, Scott C. Smith *

*University of Missouri – Rolla, Department of Electrical and Computer Engineering, 133 Emerson Electric Co. Hall,
1870 Miner Circle, Rolla, MO 65409, United States*

Available online 20 March 2006

Abstract

In this paper, a number of 4-bit, 8-operation arithmetic logic units (ALUs) are designed using the delay-insensitive NULL convention logic paradigm, and are characterized in terms of speed and area. Both dual-rail and quad-rail, pipelined and non-pipelined versions are developed, and the tradeoffs and design considerations for each are discussed. Comparing the various architectures shows that the fastest dual-rail and quad-rail ALUs achieve average speedups of 1.72 and 1.59, respectively, over their non-pipelined counterparts, while requiring 133% and 119% more area, respectively. Overall, the dual-rail designs are both faster and require less area than their respective quad-rail counterparts; however, the quad-rail versions are expected to consume less power.

© 2006 Elsevier B.V. All rights reserved.

Keywords: Asynchronous logic design; Delay-insensitive circuits; Self-timed circuits; Computer arithmetic; Arithmetic logic unit

1. Introduction

For the last two decades, digital design has focused primarily on synchronous, clocked architectures. However, because clock rates have significantly increased while feature size has decreased, clock skew has become a major problem. To achieve acceptable skew, high-performance chips must dedicate increasingly larger portions of their area to clock drivers, thus dissipating increasingly higher power, especially at the clock edge, when switching is the most prevalent.

As this trend continues, the clock is becoming more difficult to manage, causing renewed interest in asynchronous digital design. Researchers have demonstrated that correct-by-construction asynchronous paradigms, particularly NULL convention logic (NCL), require less power, generate less noise, produce less electromagnetic interference, and allow for easier reuse of components than their synchronous counterparts, without compromising performance [1]. Furthermore, these paradigms should allow much greater flexibility in the design of complex circuits

such as systems-on-a-chip (SoCs). Because these circuits are delay-insensitive, they should drastically reduce the effort required to ensure correct operation under all timing scenarios, compared to equivalent synchronous designs. Also, the self-timed nature of correct-by-construction SoCs should allow designers to reuse previously designed and verified functional blocks in subsequent designs, without significant modifications or retiming effort within a reused functional block. Such SoCs may also provide for simpler interfacing between the digital core and nontraditional functional blocks.

NCL technology is currently being utilized to design low-power, low-EMI ASICs for use in mobile electronics and smart cards, and is being used in the design of soft processor cores for use in SoCs. The initial version of the Motorola STAR08 processor using NCL technology shows a 40% reduction in power and a 10 dB reduction in noise over its clocked Boolean counterpart, while operating at a comparable frequency [1].

One of the first tasks necessary to help integrate NCL into the semiconductor design industry is to develop and characterize the key components of a reusable design library. Of fundamental importance are arithmetic circuits, including the ALUs described in this paper, as well as the

* Corresponding author. Tel.: +1 573 341 4232; fax: +1 573 341 4532.
E-mail address: smithsco@umr.edu (S.C. Smith).

multipliers [2], MACs [3], divider [4], and counters [5] described elsewhere. An overview of the NULL Convention Logic (NCL) paradigm is provided in the following section.

2. Overview of NCL

NCL offers a self-timed logic paradigm where control is inherent with each datum. NCL follows the so-called “weak conditions” of Seitz’s delay-insensitive signaling scheme [6]. As with other self-timed logic methods discussed herein, the NCL paradigm assumes that forks in wires are isochronic [7], meaning that the wire delays are much less than the logic element delays *within a component*, which is a valid assumption even in future nanometer technologies. Wires connecting components do not have to adhere to the isochronic fork assumption. The origins of various aspects of the paradigm, including the NULL (or spacer) logic state from which NCL derives its name, can be traced back to Muller’s work on speed-independent circuits in the 1950s and 1960s [8].

2.1. Delay-insensitivity

NCL uses symbolic completeness of expression [9] to achieve delay-insensitive behavior. A symbolically complete expression is defined as an expression that only depends on the relationships of the symbols present in the expression without a reference to their time of evaluation. In particular, dual-rail signals, quad-rail signals, or other *Mutually Exclusive Assertion Groups* (MEAGs) can be used to incorporate data and control information into one mixed signal path to eliminate time reference [10]. A dual-rail signal, D , consists of two wires, D^0 and D^1 , which may assume any value from the set {DATA0, DATA1, NULL}. The DATA0 state ($D^0 = 1, D^1 = 0$) corresponds to a Boolean logic 0, the DATA1 state ($D^0 = 0, D^1 = 1$) corresponds to a Boolean logic 1, and the NULL state ($D^0 = 0, D^1 = 0$) corresponds to the empty set meaning that the value of D is not yet available. The two rails are mutually exclusive, such that both rails can never be asserted simultaneously; this state is defined as an illegal state. A quad-rail signal, Q , consists of four wires, $Q^0, Q^1, Q^2,$ and Q^3 , which may assume any value from the set {DATA0, DATA1, DATA2, DATA3, NULL}. The DATA0 state ($Q^0 = 1, Q^1 = 0, Q^2 = 0, Q^3 = 0$) corresponds to two Boolean logic signals, X and Y , where $X = 0$ and $Y = 0$. The DATA1 state ($Q^0 = 0, Q^1 = 1, Q^2 = 0, Q^3 = 0$) corresponds to $X = 0$ and $Y = 1$. The DATA2 state ($Q^0 = 0, Q^1 = 0, Q^2 = 1, Q^3 = 0$) corresponds to $X = 1$ and $Y = 0$. The DATA3 state ($Q^0 = 0, Q^1 = 0, Q^2 = 0, Q^3 = 1$) corresponds to $X = 1$ and $Y = 1$, and the NULL state ($Q^0 = 0, Q^1 = 0, Q^2 = 0, Q^3 = 0$) corresponds to the empty set meaning that the result is not yet available. The four rails of a quad-rail NCL signal are mutually exclusive, such that no two rails can ever be asserted simultaneously; these states are defined as illegal states. Both dual-rail and quad-rail signals are space

optimal 1-out-of- N delay-insensitive codes, requiring two wires per bit. Other higher order MEAGs are not typically wire count optimal; however, they can be more power efficient due to the decreased number of transitions per cycle.

Most multi-rail delay-insensitive systems, including NCL, have at least two register stages, one at both the input and at the output. Two adjacent register stages interact through their request and acknowledge lines, K_i and K_o , respectively, to prevent the current DATA wavefront from overwriting the previous DATA wavefront, by ensuring that the two DATA wavefronts are always separated by a NULL wavefront.

2.2. Logic gates

NCL differs from many other delay-insensitive paradigms in that these other paradigms only utilize one type of state-holding gate, the *C-element* [8]. A C-element behaves as follows: when all inputs assume the same value then the output assumes this value, otherwise the output does not change. On the other hand, all NCL gates are state-holding. Thus, NCL optimization methods can be considered as a subclass of the techniques for developing delay-insensitive circuits using a pre-defined set of more complex components, with built-in *hysteresis* behavior.

NCL uses *threshold gates* for its basic logic elements [11]. The primary type of threshold gate is the TH_{mn} gate, where $1 \leq m \leq n$, as depicted in Fig. 1. TH_{mn} gates have n -inputs. At least m of the n -inputs must be asserted before the output will become asserted. Because NCL threshold gates are designed with hysteresis, all asserted inputs must be de-asserted before the output will be de-asserted. Hysteresis ensures a complete transition of inputs back to NULL before asserting the output associated with the next wavefront of input data. Therefore, a TH_{mn} gate is equivalent to an n -input C-element and a TH_{1n} gate is equivalent to an n -input OR gate. In a TH_{mn} gate, each of the n -inputs is connected to the rounded portion of the gate; the output emanates from the pointed end of the gate; and the gate’s threshold value, m , is written inside of the gate. NCL threshold gates may also include a *reset* input to initialize the output. Resettable gates are denoted by either a D or an N appearing inside the gate, along with the gate’s threshold, referring to the gate being reset to logic 1 or logic 0, respectively.

By employing threshold gates for each logic rail, NCL is able to determine the output status without referencing time. Inputs are partitioned into two separate wavefronts,




Fig. 1. TH_{mn} threshold gate.

the NULL wavefront and the DATA wavefront. The NULL wavefront consists of all inputs to a circuit being NULL, while the DATA wavefront refers to all inputs being DATA, some combination of DATA0 and DATA1. Initially, all circuit elements are reset to the NULL state. First, a DATA wavefront is presented to the circuit. Once all of the outputs of the circuit transition to DATA, the NULL wavefront is presented to the circuit. Once all of the outputs of the circuit transition to NULL, the next DATA wavefront is presented to the circuit. This DATA/NULL cycle continues repeatedly. As soon as all outputs of the circuit are DATA, the circuit's result is valid. The NULL wavefront then transitions all of these DATA outputs back to NULL. When they transition back to DATA again, the next output is available. This period is referred to as the DATA-to-DATA cycle time, denoted as T_{DD} , and has an analogous role to the clock period in a synchronous system.

2.3. Completeness of input

The completeness of input criterion [9], which NCL combinational circuits and circuits developed from other delay-insensitive paradigms must maintain in order to be delay-insensitive, requires that:

1. all the outputs of a combinational circuit may not transition from NULL to DATA until all inputs have transitioned from NULL to DATA, and
2. all the outputs of a combinational circuit may not transition from DATA to NULL until all inputs have transitioned from DATA to NULL.

In circuits with multiple outputs, it is acceptable, according to Seitz's weak conditions [6], for some of the outputs to transition without having a complete input set present, as long as all outputs cannot transition before all inputs arrive.

2.4. Observability

There is one more condition that must be met to ensure delay-insensitivity for NCL circuits and other delay-insensitive circuits. No *orphans* may propagate through a gate [12]. An orphan is defined as a wire that transitions during the current DATA wavefront, but is not used in the determination of the output. Orphans are caused by wire forks and can be neglected through the isochronic fork assumption [7], as long as they are not allowed to cross a gate boundary. This *observability* condition, also referred to as indicatability or stability, ensures that every gate transition is observable at the output, which means that every gate that transitions is necessary to transition at least one of the outputs.

3. Dual-rail ALU

A block diagram of the dual-rail 4-bit ALU is shown in Fig. 2. S selects which operation is to be performed on the




Fig. 2. Dual-rail ALU block diagram.

4-bit inputs, A and B , and the carry/borrow, C_{in}/B_{in} , determined by the function table in Fig. 3. F is the 4-bit output and C_{out}/B_{out} is the carry/borrow output. This circuit, as do all NCL systems, contains a complete request/acknowledge interface, including K_o for requesting the inputs, and K_i for acknowledging the outputs, and a *reset* input to initialize the NCL registers to NULL. Notice that C_{in}/B_{in} is not used in operations 0–3, and that B is not used in operations 3–5. However, to ensure delay-insensitivity the ALU must still be input-complete with respect to these inputs, even for the operations where the C_{in}/B_{in} and/or B input(s) are not used. This ensures that the unused inputs are received before the output can transition.

3.1. Non-pipelined version

The logic diagram of the non-pipelined dual-rail ALU is shown in Fig. 4. It consists of dual-rail registers [9], completion components, denoted as COMP [9], a Convert to MEAG function, a Demultiplexer, NCL OR, AND, XOR, invert, shift right, and shift left functions, a ripple-carry subtractor and adder, two Multiplexers, and Carry Logic.

The Convert to MEAG function is comprised of eight TH33 gates that convert S , which consists of three dual-rail signals, into an 8-rail MEAG. The OR, AND, and XOR function are all input-complete versions [5], consisting of two gates and one gate delay. The invert, shift right, and shift left function are performed by renaming signals, hence they have no logic delay. The ripple-carry subtractor and adder consist of four full-adders [13,14], while the

S_2	S_1	S_0	F	C_{out}/B_{out}
0	0	0	A OR B	0
0	0	1	A AND B	0
0	1	0	A XOR B	0
0	1	1	NOT A	0
1	0	0	SHR C_{in} , A_{A_0}	
1	0	1	SHL A, C_{in} A_3	
1	1	0	A-B-1+ B_{in}	B_{out}
1	1	1	A+B+ C_{in}	C_{out}

Fig. 3. ALU function table.

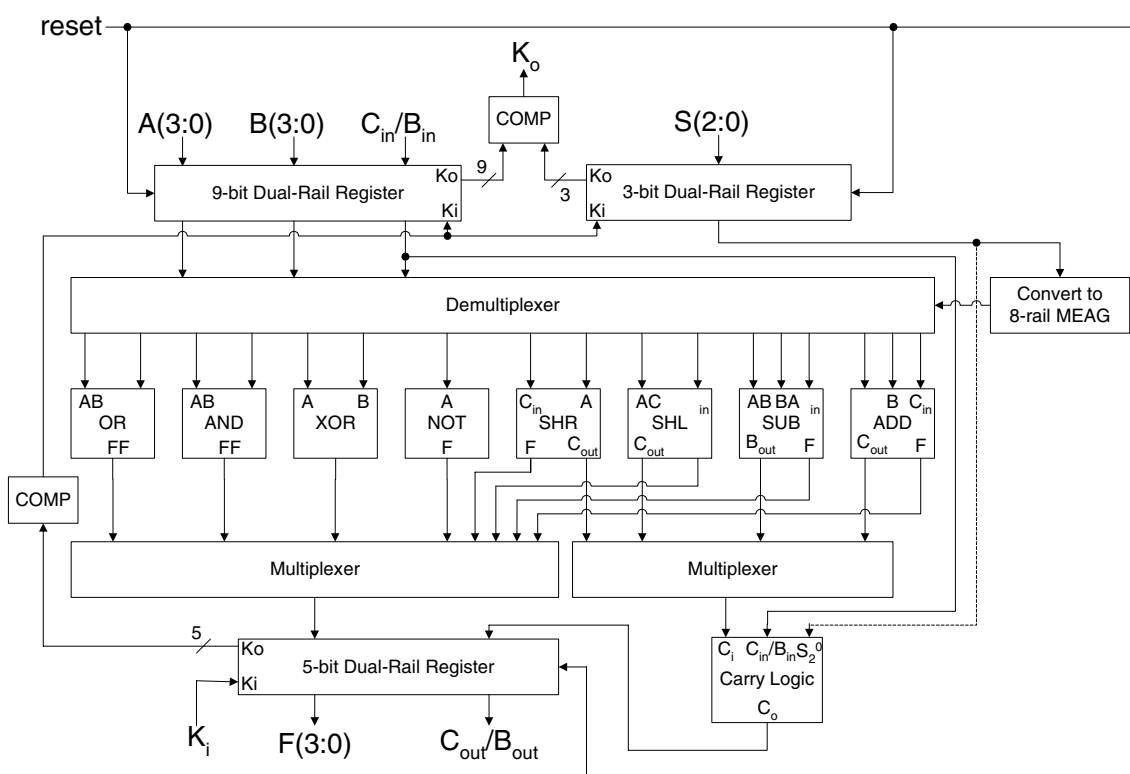


Fig. 4. Logic diagram of dual-rail non-pipelined ALU.

Multiplexers consist of TH14 and TH12 gates, which OR each rail of the demultiplexed results together to form a single result. The Carry Logic is required to generate a C_{out} of logic 0, when operation 0–3 is chosen, and to ensure input-completeness with respect to C_{in} in these cases, in order to maintain delay-insensitivity, as shown in Fig. 5.

The Demultiplexer consists of TH22 gates, which pass the A input, and the B and C_{in}/B_{in} inputs, when necessary, to the correct function, determined by the select MEAG. For the functions in which B is not required, the Demultiplexer consists of TH34 gates to pass A , while maintaining input-completeness with respect to B . Input-completeness of C_{in}/B_{in} is ensured in the carry logic. The alternative is to always pass B to every function, whether it is needed or not, and then ensure input-completeness of B within functions 3–5. However, this would require an extra gate delay for these three functions, which reduced average throughput by 3% and required an additional 24 gates; hence this alternative was not chosen.




Fig. 5. Carry logic.

3.2. Pipelined version

The interface for the pipelined dual-rail ALU is the same as for the non-pipelined version, shown in Fig. 2; the logic diagram is shown in Fig. 6. To pipeline the ALU, registration stages must be added within the combinational logic of the non-pipelined version, without violating the completeness of input criterion [9]. To minimize both delay and area, many embedded registration stages were used, where the combinational logic and registration are combined together (i.e. the Convert to MEAG Register, the Carry MEAG Register, the Demultiplexer Register, the Select (Sel) Registers, and the Multiplexer Register). Fig. 7 shows a 2-to-1 multiplexer with registered output; while Fig. 8 shows a multiplexer having the same functionality but utilizing embedded registration, where the registration functionality becomes part of the multiplexer’s combinational logic. Note that utilizing embedded registration in this case results in a decrease in 1 gate delay for the generation of the output, F , and a reduction of 2 gate delays for the generation of the completion signal, K_o . It also requires 2 fewer gates, but because of the larger gate sizes, it requires 4 additional transistors (100 vs. 96).

The major differences between the pipelined and non-pipelined designs are the input-completeness of the B and C_{in}/B_{in} inputs, and the subtractor and adder circuits. Since a registration stage was to be added between the Demultiplexer and the functions, an extra level of logic would have been required for this registration if input-completeness of

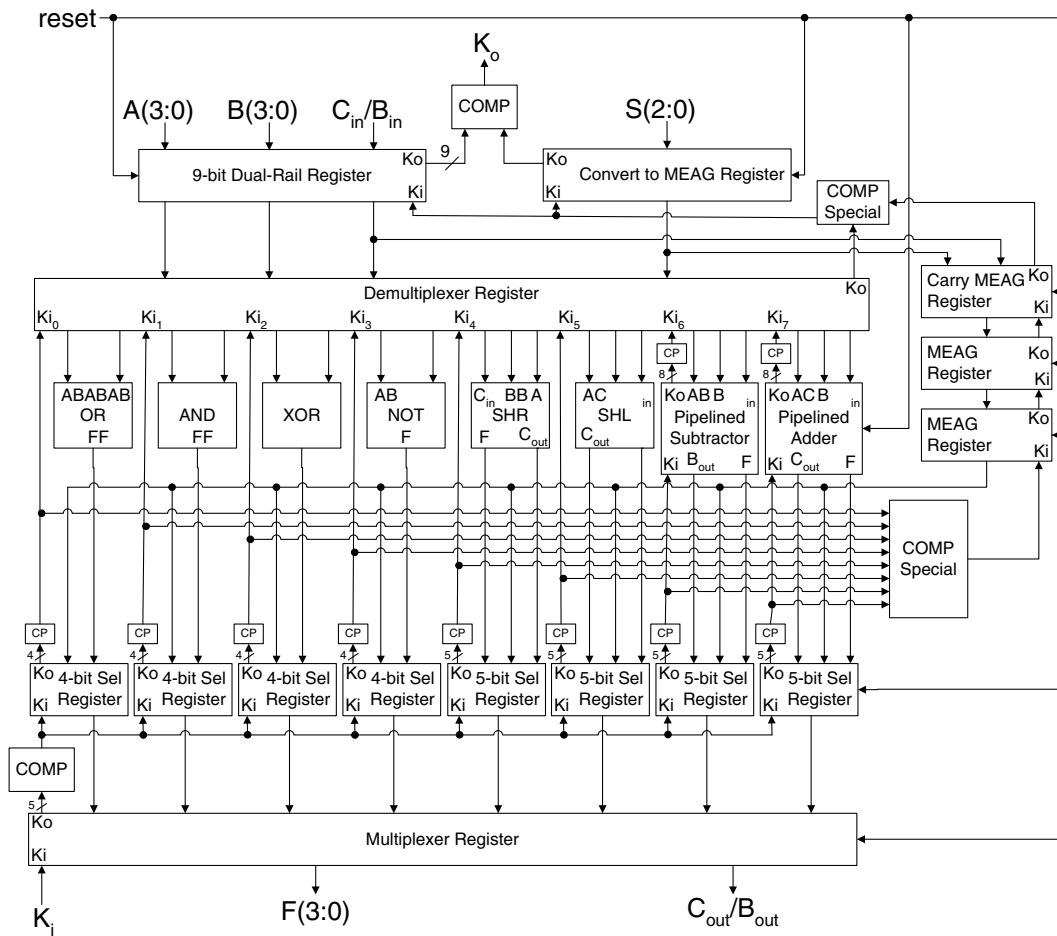


Fig. 6. Logic diagram of dual-rail pipelined ALU.

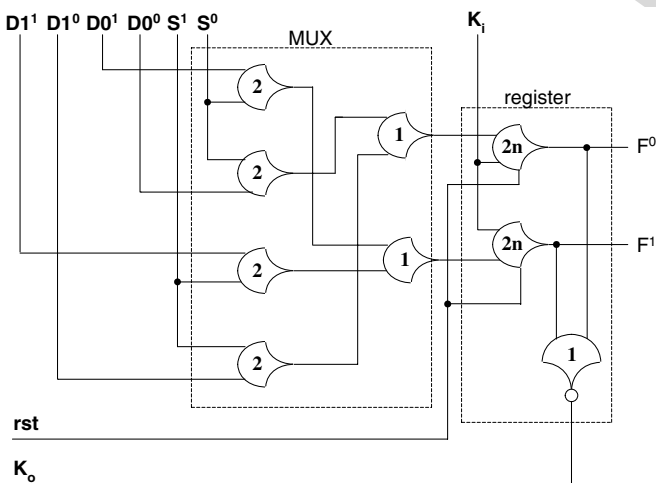


Fig. 7. Multiplexer without embedded registration.

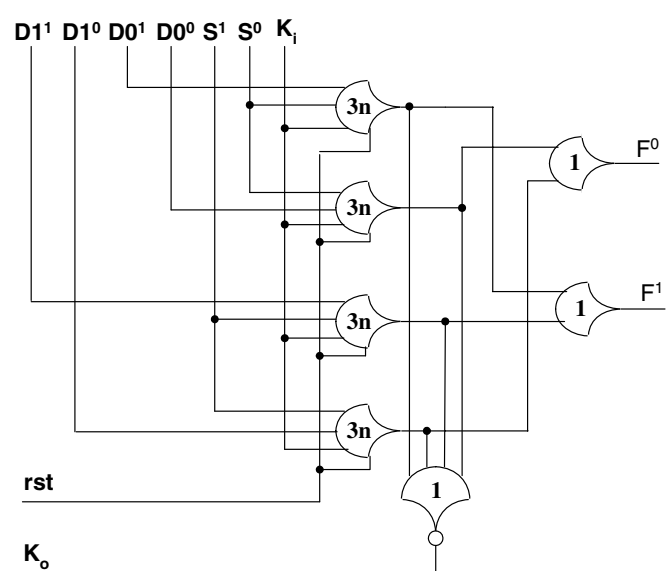


Fig. 8. Multiplexer with embedded registration.

B was performed within the demultiplexer, as was done for the non-pipelined design, since this required maximum 4-input, TH34 gates; thus excluding the possibility of embedded registration due to the needed additional request input and the 4-input limitation. Therefore, the alternative to pass *B* to functions 3–5 and then ensure input-completeness

of *B* within these functions was chosen, since this option only required 2-input, TH22 gates, for the demultiplexer, thus enabling the extra request input to be added, changing

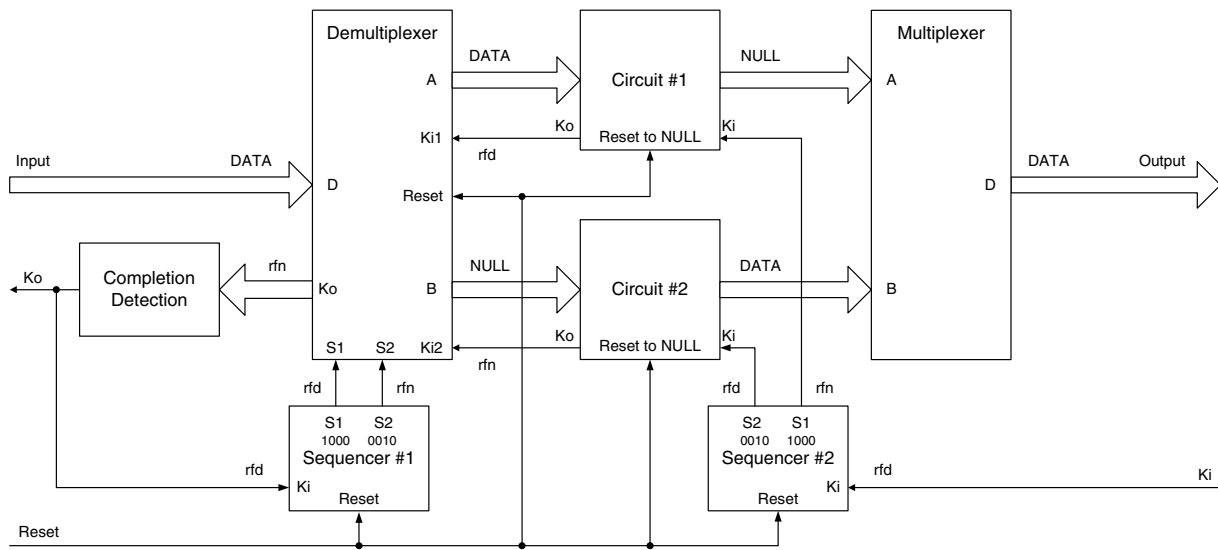


Fig. 9. NCR architecture.

the TH22 gates to TH33 gates, in order to embed the registration stage within the combinational logic of the demultiplexer. Also, both the Demultiplexer Register and the Select Register require special completion components, since the inputs are only sent to one out of the eight output sets, whereas normally all outputs transition every DATA/NULL cycle.

Input-completeness of the C_{in}/B_{in} input was ensured within the Carry MEAG register, instead of at the end of the design, like in the non-pipelined version, since this would have required unnecessarily pipelining C_{in}/B_{in} throughout the design. This allowed for the carry output of zero to be generated from the F_0 output of functions 0–3; hence only one multiplexer component is shown in the diagram. Also, the adder and subtractor circuits have themselves been pipelined [13,14], by inserting two registration stages between the four full-adders and utilizing bit-wise completion [13,14], where the completion signal from bit b in register _{i} is only sent back to the bits in register _{$i-1$} that took part in the calculation of bit b . A third registration stage could have been inserted, however this did not increase throughput. Finally, the select MEAG was pipelined by adding two additional registration stages, after the Carry MEAG Register, such that the ALU could store the select signals associated with the concurrent DATA/NULL wavefronts being processed within the ALU. Both more and less MEAG pipelining stages decreased throughput. These pipelining optimizations resulted in a 32% increase in average throughput, with a 151% increase in area.

3.3. NULL cycle reduced version

Since pipelining the ALU resulted in more than twice the area of the non-pipelined design, with a speedup of only 1.32, a viable alternative was to use the NULL cycle reduction (NCR) technique [13,15] on the non-pipelined design

to increase its throughput. As shown in Fig. 9, NCR demultiplexes successive input wavefronts such that one circuit processes a DATA wavefront, while its duplicate processes a NULL wavefront. The first DATA/NULL cycle flows through the original circuit, while the next DATA/NULL cycle flows through the duplicate circuit. The outputs of the two circuits are then multiplexed to form a single output stream. The application of NCR to the non-pipelined design resulted in a 63% increase in average throughput, with only a 136% increase in area, thus the NCR version was both faster than the pipelined version and it required less area.

To further increase throughput and decrease area, embedded registration was applied to the non-pipelined design. This resulted in the Convert to MEAG function becoming an embedded register, replacing the 3-bit input register, and the multiplexer for F and the Carry Logic function becoming embedded registers, replacing the 5-bit output register. This optimization resulted in a speedup of 1.18 over the original non-pipelined design, and a 2% decrease in area. Now NCR could be applied to the non-pipelined design utilizing embedded registration, resulting in a speedup of 1.72 with only a 133% increase in area, versus the original non-pipelined design.

4. Quad-rail ALU

The interface for the quad-rail ALU is similar to that of the dual-rail ALU, shown in Fig. 2, except that A , B , and F now consist of two quad-rail signals, instead of four dual-rail signals, and the select input, S , consists of one dual-rail signal, S_2 , and one quad-rail signal, which represents $S(1:0)$. The logic diagram for the quad-rail ALU is also very similar to that of the dual-rail version, shown in Fig. 4. It consists of both quad-rail and dual-rail registers [13], completion components, a Convert to MEAG function, a Demultiplexer, NCL OR, AND, XOR, invert, shift right,

and shift left functions, a ripple-carry subtractor and adder, two Multiplexers, and Carry Logic. The main difference is that each component and each function was designed using quad-rail logic instead of dual-rail logic. This results in ripple-carry subtractor and adder circuits consisting of only two quad-rail adders (two quad-rail signals plus a dual-rail carry input, yielding a quad-rail sum and a dual-rail carry output) [2]. The other major difference was that B had to be passed to all functions through the demultiplexer and input-completeness of B had to be ensured within functions 3–5, due to the quad-rail logic and 4-input limitation for gates. Input-completeness of C_{in}/B_{in} was ensured in the same manner as for the dual-rail design.

Pipelining the quad-rail design followed in much the same manner as the dual-rail design, resulting in a 34% increase in average throughput, with a 116% increase in area. Again, NCR was applied to the non-pipelined design, resulting in a throughput increase of 53% with a 119% increase in area. Like the dual-rail design, embedded registration was applied to the non-pipelined design, resulting in a 10% increase in throughput with a minimal increase in area; and NCR was applied to the non-pipelined design with embedded registration, resulting in a speedup of 1.59 with an area increase of 119%, versus the original non-pipelined design.

5. Simulation results

To assess the performance of alternate designs, Mentor Graphics' ModelSim tool was used to simulate structural VHDL implementations of the circuits to generate their timing characteristics. The average rising and falling time for each NCL gate was determined through Cadence simulation of a 0.5 μm CMOS process operating at 3.3 V; and this information was then stored in a VHDL library to be used in the simulations. Table 1 summarizes the characterization of the various ALUs discussed herein, in terms of both speed and area. Gate count can be used as one measure of area for comparison purposes; however since the NCL gates vary greatly in size (i.e. from 2 transistors

for an inverter to 26 transistors for a TH24 gate), transistor count provides a better means of comparison, especially when embedded registration is used, since this increases transistor count without increasing gate count. Also, since NCL circuits are delay-insensitive, speed is data dependent; therefore average cycle time, T_{DD} , is calculated and used for comparison. T_{DD} is calculated as the arithmetic mean of the cycle times corresponding to all 4096 possible pairs of input operands.

6. Conclusions

Comparing the various architectures shows that the dual-rail versions of all designs outperform their quad-rail counterparts in terms of both speed and area. However, the quad-rail designs are expected to consume less power, due to the fact that only one quad-rail signal transitions for every two dual-rail signals that transition. The reason that pipelining the designs did not further increase throughput is due to the long completion delays in the special completion components required for both the Demultiplexer Register and Select Registers, described in Section 3.2. However, the application of embedded registration to the non-pipelined design, followed by applying NCR yielded a significant additional increase in throughput over the original non-pipelined design, versus the throughput increase achieved by pipelining the design for both dual-rail and quad-rail architectures. Furthermore, this NCR approach required less area than pipelining for the dual-rail architecture, and only slightly more area than pipelining for the quad-rail architecture.

References

- [1] J. McCardle, D. Chester, Measuring an asynchronous processor's power and noise, in: Proceedings of the Synopsys User Group Conference, Boston, 2001.
- [2] S.K. Bandapati, S.C. Smith, M. Choi, Design and characterization of NULL convention self-timed multipliers, *IEEE Design and Test of Computers* 30/6 (November–December) (2003) 26–36 (Special Issue on Clockless VLSI Design).
- [3] S.C. Smith, R.F. DeMara, J.S. Yuan, M. Hagedorn, D. Ferguson, NULL convention multiply and accumulate unit with conditional rounding, scaling, and saturation, *Journal of Systems Architecture* 47 (12) (2002) 977–998.
- [4] S.C. Smith, Design of a NULL convention self-timed divider, in: Proceedings of the 2004 International Conference on VLSI, June 2004, pp. 447–453.
- [5] S.C. Smith, R.F. DeMara, J.S. Yuan, D. Ferguson, D. Lamb, Optimization of NULL convention self-timed circuits, *Integration, The VLSI Journal* 37 (3) (2004) 135–165.
- [6] C.L. Seitz, in: *System Timing in Introduction to VLSI Systems*, Addison-Wesley, Reading, MA, 1980, pp. 218–262.
- [7] C.H. (Kees) van Berkel, M. Rem, R. Saeijs, "VLSI Programming," in: Proceedings of the 1988 IEEE International Conference on Computer Design: VLSI in Computers and Processors, 1998, pp. 152–156.
- [8] D.E. Muller, Asynchronous logics and application to information processing, in: *Switching Theory in Space technology*, Stanford University Press, Stanford, CA, 1963, pp. 289–297.
- [9] K.M. Fant, S.A. Brandt, NULL convention logic: a complete and consistent logic for asynchronous digital circuit synthesis, in:

Table 1
ALU characterization

ALU architecture	Gate count	Transistor count	T_{DD} (ns)
Dual-rail non-pipelined	260	3757	7.75
Dual-rail pipelined	620	9926	5.89
Dual-rail NCR	542	8878	4.76
Dual-rail non-pipelined with embedded registration	245	3688	6.57
Dual-rail NCR with embedded registration	512	8740	4.51
Quad-rail non-pipelined	522	7179	8.85
Quad-rail pipelined	1027	15478	6.62
Quad-rail NCR	1137	15726	5.80
Quad-rail non-pipelined with embedded registration	514	7182	8.05
Quad-rail NCR with embedded registration	1121	15732	5.57

- Proceedings of the International Conference on Application Specific Systems, Architectures, and Processors, 1996, pp. 261–273.
- [10] T. Verhoff, Delay-Insensitive Codes-An Overview, *Distributed Computing* 3 (1988) 1–8.
- [11] G.E. Sobelman, K.M. Fant, CMOS circuit design of threshold gates with hysteresis, in: *Proceedings of the IEEE International Symposium on Circuits and Systems (II)*, 1998, pp. 61–65.
- [12] A. Kondratyev, L. Neukom, O. Roig, A. Taubin, K. Fant, Checking delay-insensitivity: 10^4 gates and beyond, in: *Proceedings of the Eighth International Symposium on Asynchronous Circuits and Systems*, 2002, pp. 137–145.
- [13] S.C. Smith, Gate, Throughput Optimizations for NULL Convention Self-Timed Digital Circuits, Ph.D. Dissertation, School of Electrical Engineering and Computer Science, University of Central Florida, May 2001.
- [14] S.C. Smith, R.F. DeMara, J.S. Yuan, M. Hagedorn, D. Ferguson, Delay-insensitive gate-level pipelining, *Integration, The VLSI Journal* 30 (2) (2001) 103–131.
- [15] S.C. Smith, R.F. DeMara, J.S. Yuan, M. Hagedorn, D. Ferguson, Speedup of delay-insensitive digital systems using NULL cycle reduction, in: *Proceedings of the 10th International Workshop on Logic and Synthesis*, June 2001, pp. 185–189.

Author's personal copy