

Regenerative Feedback Repeaters for Programmable Interconnections

Ivo Dobbelaere, *Member, IEEE*, Mark Horowitz, *Senior Member, IEEE*, and Abbas El Gamal, *Senior Member, IEEE*

Abstract—The use of regenerative feedback repeaters to reduce the delay in programmable interconnections is described. A static, complementary regenerative feedback (CRF) repeater is proposed. This CRF repeater locally regenerates the new level for a fixed time after a transition has been detected. Design issues and limitations are discussed. It is shown that rising transitions can propagate faster than falling transitions through a chain of overdriven nMOS switches with CRF repeaters. Experimental results from a 1.2 μm CMOS implementation show that the loaded delay through 64 switches for static and dynamic repeaters can be reduced by a factor 1.4–2 over conventional repeaters.

I. INTRODUCTION

FIELD Programmable Gate Arrays (FPGA's) consist of an array of programmable logic cells interspersed by a programmable interconnection network as shown in Fig. 1 and provide a medium for prototyping, emulating, and implementing logic designs. The interconnection network consists of programmable switches that are organized in connection blocks and switch blocks. Several programming technologies are available such as SRAM cells, anti-fuses, and (E)EPROM's. Programmable switches may be implemented as MOS transfer gates controlled by a memory element or as anti-fuses. The performance of FPGA's is mainly limited by the delay of the programmable interconnection network. This delay increases quadratically with the number of series switches and linearly with the number of switches loading each node and is especially a problem when the programmable switches are implemented using MOS transistors since these have an appreciable resistance and capacitance.

Changing the size of the switches cannot reduce the delay much since RC remains relatively constant. The accumulation of quadratic delay can be limited by inserting repeaters as shown in Fig. 2. Conventional repeaters consist of pairs of unidirectional tristate buffers and memory cells. These repeaters have a high area penalty since after programming at least half of the buffers is not used. In addition, they contribute their own propagation delay to the signal delay, which limits the delay reduction that can be obtained.

In this paper, we describe an alternative repeater based on regenerative feedback, which offers a smaller area and delay penalty. In Section II, the principle is explained, and

Manuscript received May 3, 1995; revised August 17, 1995. This work was supported in part by a donation from Altera, by FBI Contract J-FBI-89-101 and by ARPA Contract DABT-63-94-C-0054. Software was provided by Meta-Software and by Mentor Graphics.

The authors are with the Department of Electrical Engineering, Stanford University, Stanford, CA 94305 USA.

IEEE Log Number 9415197.

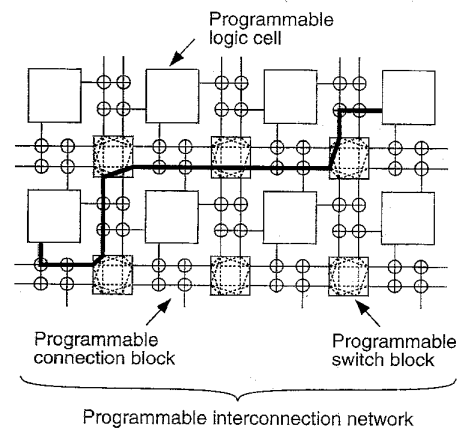


Fig. 1. FPGA architecture. An interconnection obtained after programming is indicated.




Fig. 2. Interconnection network using pairs of unidirectional tristate buffers to limit delay.

a number of CMOS implementations are described. In Section III, we discuss the design issues for high performance and show a delay comparison with conventional repeaters. In Section IV, the limitations of regenerative feedback repeaters are addressed. Section V details the design of a test chip, and Section VI contains measurement results.

II. REGENERATIVE FEEDBACK REPEATERS

The area and delay penalty of providing signal amplification in both signal directions along the interconnection paths can be alleviated by using regenerative feedback repeaters. Such repeaters have a single signal terminal on which the potential is monitored. When a level change past a threshold is sensed, the repeater locally regenerates the signal, shorting out the large series resistance that is driving the interconnection. Fig. 3 conceptually shows a network with regenerative feedback repeaters.

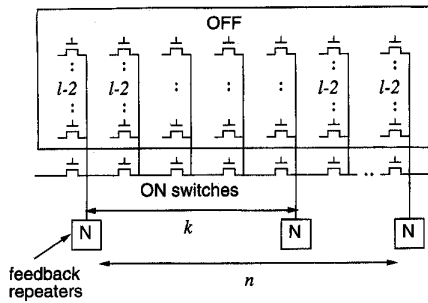


Fig. 3. Interconnection network using regenerative feedback repeaters.




Fig. 4. Precharged regenerative feedback circuits. (a) Domino. (b) improved Domino. (c) NORA.

The principle is similar to the signal propagation in myelinated nerve axons [2]. Due to the periodical amplification of the signal along its propagation path, the delay increases linearly rather than quadratically with the number of stages [3]. The deviation from the initial node level that is required to activate the repeater will be referred to as the “relative threshold.” The propagation delay decreases as the relative threshold is chosen smaller.

Since regenerative feedback repeaters are not inserted in the signal path, they do not contribute their own propagation delay but only add extra capacitance to the network. As a result, the interconnection delay keeps decreasing further as repeaters are placed more frequently, and the delays can be made smaller than when conventional repeaters are used. The regenerative feedback repeater only needs one buffer per repeater and no memory cells. This allows a substantial area reduction over the conventional repeater.

A. Dynamic Implementation

The simple implementation of regenerative feedback using cross-coupled static inverters does not function due to hard-latching, but this is avoided by using dynamic logic. Precharged regenerative feedback repeaters have been used to reduce the delay of memory word lines [4] and carry chains [5], [6]. Fig. 4(a) shows a Domino buffer [7] with its output connected back to its input. When the precharged node N is pulled below the threshold voltage of the inverter, the low level is enforced locally. When added to a chain of switches as in Fig. 3, these circuits rapidly propagate a falling transition. By using a gated clock, as shown in Fig. 4(b), only a single nMOS is needed in the pull-down path resulting in a lower output impedance and a shorter propagation delay. The delay




Fig. 5. Waveforms in a chain with regenerative feedback repeaters.

can be minimized by choosing the NOR gate threshold close to V_{DD} while maintaining a low input capacitance. By using a dynamic NOR gate, a small input capacitance can be combined with a relative threshold equal to the pMOS device threshold $|V_{TP}|$. This results in the NORA [8] feedback circuit of Fig. 4(c). The NORA implementation is the fastest but suffers from poor noise margin and poor noise performance.

In Fig. 3, the topology parameters used to describe the regenerative feedback repeater chain are indicated. The repeater interval k is defined as the number of switches between two nodes that contain repeaters, the chain length n is the total number of series switches, and the node load l is the total number of switches attached to each intermediate node. The corresponding parameters for conventional repeaters are indicated in Fig. 2.

Fig. 5 shows the waveforms at subsequent nodes in a chain of switches with Domino feedback repeaters as in Fig. 4(b), for a repeater interval $k = 4$, chain length $n = 16$, and node load $l = 16$. The input node (Node 0) is driven by a clocked driver. The potentials on the subsequent nodes show an exponentially decaying step response until the potential on the fourth node reaches the threshold of the NOR gate. This activates the local pull-down path and speeds up the falling transition on the further nodes. In turn, the repeaters at the eighth, twelfth, and sixteenth nodes are activated.

The precharged circuit is fast, but it is not compatible with most FPGA's since it requires clocking and monotonic signaling.

B. Static Implementation

To address FPGA's, we propose the self-timed, complementary regenerative feedback (CRF) circuit of Fig. 6. This is the complementary version of the circuit in Fig. 4(b), where the precharge transistor is removed and to which the complement of the pull-down circuitry is added, resulting in the NAND gate and pMOS transistor. The clock input to NOR and NAND gates, which determines when the feedback loop is turned off, is replaced by the delayed, inverted signal of node N .

In steady state, both drivers are off, so N is not actively driven. In the case of a falling transition on node N , both inputs to the NOR gate are temporarily low, and the nMOS driver is turned on for approximately the imposed delay or regeneration time t_d . Complementarily, when a rising transition is detected by the NAND gate, the pMOS driver turns on for a time t_d . Hence, this circuit can propagate both rising and falling transitions, does not require a clock, and, if




Fig. 6. Self-timed complementary regenerative feedback (CRF) repeater.

there is sufficient time between transitions, is not limited by ratioed logic constraints.

III. DESIGN FOR PERFORMANCE

In this section, the relation between design choices and performance is discussed. First, the switch choice is discussed. Since the rising and falling thresholds can be chosen independently in CRF repeaters, overdriven nMOS switches can propagate a rising transition as fast as a falling transition. The propagation speed further depends on the output impedance and on how fast the feedback turns on, which is determined by the relative threshold, the input capacitance, and the predriver strength. Finally, the delay of regenerative feedback repeaters is compared with that of conventional repeaters.

A. Switch Selection

Due to the higher current driving capability of nMOS transistors, a single nMOS transistor with gate voltage $V_{GG} = V_{DD}$ has a smaller equivalent RC for propagating a falling transition than a similarly sized CMOS transfer gate [9]. By using an overdriven gate voltage $V_{GG} \geq V_{DD} + V_{TN}$ with V_{TN} , the nMOS device threshold, the switch resistance can be further reduced, and the threshold voltage drop for a rising transition can be avoided. Such voltages above V_{DD} can be generated on-chip using a charge pump [10]. With an inverter threshold at 50% V_{DD} , the rising transition propagates roughly 1.5 times slower than the falling transition in an overdriven switch with $V_{GG} = V_{DD} + V_{TN}$. In order to obtain equal rising and falling propagation delays, the transistor sizes should be adjusted such that the gate threshold is around 40% V_{DD} .

When using CRF repeaters, the rising and falling thresholds V_R and V_F may be chosen independently. This allows a rising transition to propagate faster than a falling transition through a chain with overdriven nMOS transistors as will be shown next.

The current through an nMOS transistor with an overdriven voltage of $V_{GG} = V_{DD} + V_{TN}$ can be approximated by neglecting body and short-channel effects

$$I = B \left\{ V_{GS} - V_{TN} - \frac{1}{2} V_{DS} \right\} V_{DS} \quad (1)$$

where $B = \mu C_{ox} W/L$, in which μ is the electron surface mobility, C_{ox} is the gate oxide capacitance per unit area, and W and L are the width and length of the channel, and where V_{GS} and V_{DS} are the gate-source and drain-source potentials, respectively. Let the source and drain potentials be V_1 and V_2 with $V_1 \leq V_2$. This yields

$$I \approx B \left\{ V_{DD} - \frac{1}{2} (V_1 + V_2) \right\} (V_2 - V_1). \quad (2)$$




Fig. 7. Lumped and distributed RC networks with nonlinear resistance R_{nl} and linear capacitance C .

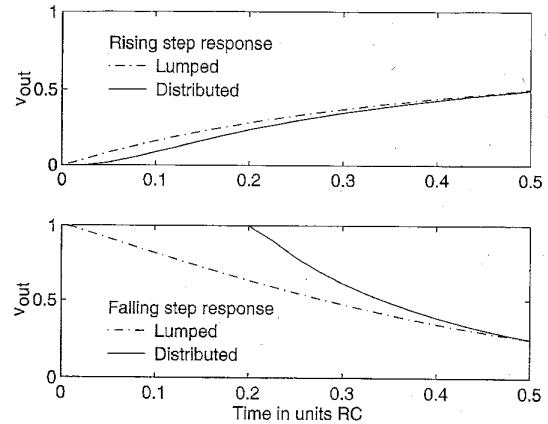


Fig. 8. Rising and falling step response through a lumped and a distributed nonlinear RC network modeling overdriven nMOS transistors.

We define the equivalent resistance R as the resistance at $V_1 = V_2 = V_{DD}/2$ or

$$R = \frac{2}{BV_{DD}}. \quad (3)$$

An overdriven nMOS switch can in first order be modeled as a nonlinear resistance

$$R_{nl}(V_1, V_2) \approx \frac{R}{\left(2 - \frac{V_1 + V_2}{V_{DD}}\right)}. \quad (4)$$

We now consider an n -stage chain consisting of nonlinear series switches as above with equivalent resistance R/n and linear shunt capacitances C/n for large values of n as shown in Fig. 7. For comparison, we also consider a lumped, one-stage network. Fig. 8 shows the step response at the output of the lumped and the distributed network. For the lumped response, the falling transition is always faster than the rising transition for equal relative thresholds, i.e., when $V_R = V_{DD} - V_F$. However, for the distributed network, the falling response stalls initially, while the rising response follows quickly. This can be explained by the fact that in the case of a rising transition, the switches are initially in the region where the conductance is largest, while for the falling transitions the initial switch conductance is zero. For finite switch chains, the response falls in between the lumped and the distributed response after normalization, and for length $n > 1$, the initial rising response is faster than the initial falling response. As a result, a rising transition propagates faster than a falling transition when the relative threshold is chosen below about 30% V_{DD} .

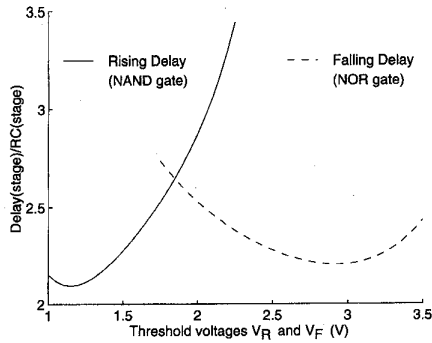


Fig. 9. Delay per stage versus sense gate threshold voltage.

B. Rising and Falling Thresholds

Due to the asymmetry in the falling and rising step responses, a larger delay reduction is obtained by making the rising threshold V_R of the NAND gate close to GND than by making the falling threshold V_F of the NOR gate close to V_{DD} . In addition, it is difficult to design a NOR gate in CMOS with high threshold without having to resort to long-channel transistors with high input capacitance, while it is relatively easy to design a NAND gate with low threshold and low input capacitance. Both of these factors favor the propagation of the rising transition.

Fig. 9 shows how the rising and falling delay in a CRF chain with overdriven nMOS switches depends on the rising and falling threshold, respectively, for repeater interval $k = 2$. The trade-off between input capacitance and threshold voltage is apparent for the NOR gate. This comparison was obtained by changing the thresholds in a delay-optimized CRF repeater chain while keeping the driver and predriver driving strengths constant. The input capacitance becomes a less important factor as the repeater interval k increases.

C. Driver and Predriver Size

The delay further depends on the size of the driver and predriver transistors. The delay optimum occurs for large transistor sizes due to the large node capacitances of the interconnection network, and optimization under area constraint is required.

D. Delay Comparison

Figs. 2 and 3 show the topologies that were used for the conventional repeaters and for the regenerative feedback repeaters, respectively, for delay comparison. The topology for the conventional repeaters is less general than the topology for the regenerative feedback repeaters and than the one we considered before in [11], but it is appropriate in a comparison directed toward FPGA's. Fig. 10 shows the circuit schematic of the tristate buffer used in the conventional repeater. The buffer threshold was 2.0 V.

Fig. 11 compares delays per stage, expressed in units of equivalent RC per stage, for simulations of the propagation delay in four chains with variable repeater interval k . The first curve shows the average propagation delay for a chain that uses conventional repeaters, which were optimized for




Fig. 10. Tristate buffer for conventional repeater.

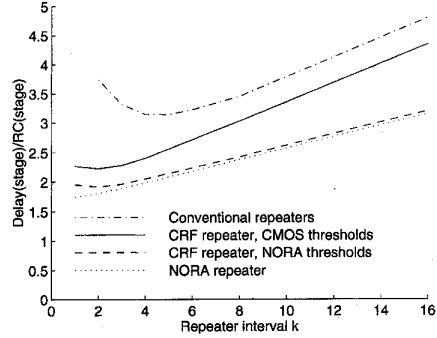


Fig. 11. Delay through 64 loaded stages versus repeater interval.

minimum delay, with equal rising and falling delay, within a given area. The second curve shows the average propagation delay for a chain that uses CRF repeaters with CMOS NAND and NOR gates. These CRF repeaters were optimized for minimum worst case delay for small repeater interval k . The sum of driver and predriver area per buffer was kept equal for both repeater types. The rising threshold was (see Fig. 9) $V_R = 1.3$ V, and the falling threshold was $V_F = 2.6$ V. For large values of the repeater interval k , the delay can be reduced further by changing the threshold of the NAND and NOR gates since the larger input capacitance has a smaller influence on the delay. A lower bound on the delay is obtained by replacing the CMOS NAND and NOR gates by NORA-like gates with relative thresholds equal to the pMOS and nMOS device thresholds $|V_{TP}|$ and V_{TN} , respectively. This yields the third curve of Fig. 11. The fourth curve shows the falling delay for a chain that used dynamic, NORA regenerative feedback repeaters, with a threshold of $|V_{TP}|$. This delay is smaller than that of the CRF repeater due to the smaller capacitance contributed by the repeater. While the noise margin of a NORA implementation is probably too small to be used on an FPGA, the ideal gate will be one with skewed thresholds in the sense gates. In this experiment, the NORA chain was 1.8 times faster than the conventional approach. The CRF chain with NORA-like sense gates was 1.6 times faster, and the CRF chain with CMOS sense gates was 1.4 times faster. A sparser placement of the smaller CRF repeaters can still yield a smaller delay than the conventional approach. Other topologies were considered as well, and speed-ups of over 2 were obtained.

IV. LIMITATIONS

In this section, we discuss the limitations and constraints that must be addressed when designing interconnections with CRF repeaters. The most obvious constraint is that the regeneration




Fig. 12. Minimum delay t_d , measured at the typical process, that still ensures proper operation over all process corners.

time t_d must be chosen longer than the propagation delay between subsequent repeaters. This regeneration time also limits the maximum rate at which a signal can change. The minimum pulse width that is propagated without any disturbance is the sum of the regeneration time and the propagation delay between opposite levels. Shorter pulses are stretched out to this minimum pulse width. Sufficiently short pulses are filtered out. There is a range of pulse widths at the boundary of propagation for which the repeaters are left in a metastable state. This metastability is the main limitation of CRF repeaters because it prevents the use in an environment with glitches. The adverse effect of metastability can be limited by ensuring that the positive feedback is turned off after the regeneration time t_d .

A. Minimum Regeneration Time t_d

The regeneration time t_d of each repeater must be long enough such that the repeater can activate the next repeater before the local current path is turned off. Otherwise, stored charge may initiate the reverse transition, resulting in oscillations. This requirement must be satisfied over all process corners and operating conditions. The worst situation occurs when there is a large discrepancy between the interconnect network delay, which we implemented using only nMOS transistors, and the delay of the delay element, for which we used both nMOS and pMOS transistors. Fig. 12 shows the delay of the delay element, measured at the typical process and conditions, that ensures sufficiently long delay over all processes and conditions, for both rising and falling transitions. The delay was implemented as an odd chain of inverters and as a single inverter for $k \leq 2$. The constraint on t_d can be relieved by using a Schmitt-trigger to replace the first inverter in the delay element. In that case, some of the dependency of the driver impedance on process and operating parameters can be absorbed by the fact that the delay t_d only starts counting when the node potential reaches a level close to the new level.

B. Minimum Pulse Width

After a CRF repeater detects a transition, the corresponding new level is locally regenerated for a regeneration time t_d . Hence, a propagating rising transition is followed by a propagating high level regenerative phase of duration t_d , and




Fig. 13. Minimum number of stages between the trailing edge of a rising regeneration and an oncoming falling transition.




Fig. 14. The minimum period T_{\min} is determined by the minimum regeneration time t_d and the time to reach V_R or V_F .

a propagating falling transition is followed by a propagating low level regenerative phase of duration t_d . A minimum phase distance p_F exists between the trailing edge of a high level regenerative phase and a subsequent falling transition. This number of stages p_F is determined by the ratioed logic constraint, as shown in Fig. 13, for which the potential on the next repeater node V_{nn} can be pulled below the falling threshold V_F by the first repeater that regenerates the oncoming falling transition against the pull-up from the last repeater that regenerates the previous rising transition. The minimum phase distance between two opposite regenerative phases for rising and falling transitions p_R and p_F may be different due to the asymmetry of the switches. By proper choice of the design parameters, we can ensure this phase distance is minimum: $p_F = p_R = 2k$. The minimum period of the input signal T_{\min} can then be estimated as (Fig. 14)

$$T_{\min} \approx 2t_d + t_R(k) + t_F(k) \quad (5)$$

where $t_R(k)$ and $t_F(k)$ are the rising and falling propagation delay, respectively, through k switches in the chain of length $2k$ between two opposite regenerative phases.

C. Glitches

The minimum signal pulse width t_{\min} for undisturbed propagation is approximately $T_{\min}/2$. Pulses that are shorter than the minimum pulse width t_{\min} are either propagated and stretched out to duration t_{\min} , or, if they are sufficiently short, are not propagated. There is a range of pulse widths around the boundary of propagation for which the repeaters are left in a metastable state as shown in Fig. 15. For example, immediately after a glitch to the low level, the circuit consists of two cross-coupled inverters. The first inverter, from V_y to V_x , consists of the pMOS driver MP4 with a load of nMOS transistors MN3-MN2-MN1 since the gate of MN1 is kept high after the glitch. The second inverter consists of the NAND gate, which




Fig. 15. Metastable state may be induced by a glitch.




Fig. 16. Arrival time of rising and falling glitch edges versus glitch duration for a glitch to the low level.

acts as an inverter from V_x to V_y as long as its input from the delay element remains high. A critical glitch duration exists, for which the amount of charge placed on the network exactly puts V_x and V_y at the metastable point of these two inverters.

Fig. 16 shows the arrival times of the falling and rising edge of a low level glitch, at the output of a chain of length $n = 8$ with repeater interval $k = 2$, versus the glitch duration. Both the arrival times and the glitch duration are measured in units regeneration time t_d . Glitches shorter than $t_{\min} \approx 2t_d$ are propagated but are stretched out to t_{\min} . Glitches of $\approx 0.5t_d$ induce a metastable state with a long resolve time [12].

Such metastability is a fundamental property of complementary circuits that use regenerative feedback. As a result, a small range of pulse widths exists that cannot be propagated. This pulse width range can be avoided by requiring glitch-free signaling. Unfortunately, CRF repeaters are not readily compatible with conventional FPGA architectures since the output of a static programmable logic cell can produce a glitch when two inputs change within a short time interval. Glitch-free signaling can be obtained, e.g., by using self-timed design throughout the FPGA or by using clocking at the interconnect drivers. Such architectures are more complex than conventional FPGA's, and their complexity must be compared with that of using simple precharged repeaters (Fig. 4) in a dynamic FPGA.

Finally, by using separate delay elements for NAND and NOR gate, as shown in Fig. 17, with thresholds $V_{de,R} < V_R$ and $V_{de,F} > V_F$, respectively, it can be guaranteed that the positive feedback is turned off, even when in the metastable state, after delay t_d . The worst case propagation delay is now determined by a glitch that moves each repeater in turn into the metastable state for about t_d . This circuit limits the adverse effect of metastability but may leave some of the repeaters unable to fire after signals with long transition times.




Fig. 17. CRF repeater with limited metastable state duration.




Fig. 18. Chip block diagram.

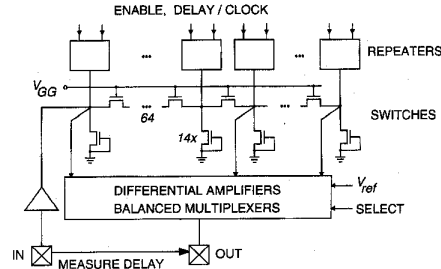


Fig. 19. Programmable chain.

V. TEST CHIP

In order to evaluate regenerative feedback repeaters, a test chip was fabricated in a MOSIS $1.2 \mu\text{m}$ n-well CMOS technology. Fig. 18 shows a block diagram of the chip. We implemented several chains each consisting of 64 MOS switches in series. Each intermediate node is loaded by 14 switches in the off state, representing unused switches in a typical programmable interconnection of an FPGA. Each chain employs a different combination of feedback circuits and switch types. The switches are either nMOS transistors with an external gate voltage V_{GG} (Fig. 19) or CMOS transfer gates.

The voltage V_{GG} can be set above V_{DD} for lower resistance and to avoid voltage drop. In order to evaluate different repeater intervals k , NOR and NAND gates were added at the clock input or in the delay loop of the repeaters such that the feedback can be enabled externally.

The delay t_d of the CRF repeaters is controlled by an external supply. At certain nodes, the signal is observed by a measurement circuit (Fig. 20) consisting of two pMOS and two nMOS differential pairs for waveform measurements in the 0–2.5 V and the 2.5–5 V ranges, respectively. The outputs are connected to a balanced network of multiplexers




Fig. 20. Each measurement circuit consists of a two pMOS and two nMOS differential pairs for waveform measurements in the 0–2.5 V and 2.5–5 V ranges, respectively. The reference voltages are externally supplied.




Fig. 21. Chip micrograph.

for measurement of relative delays and transition times. All reported results are for chain length $n = 64$, with 14 extra switches per node, $V_{DD} = 5$ V, and nMOS switches with $V_{GG} = 6.5$ V or 5 V.

Fig. 21 shows a chip micrograph. The two dynamic chains used the Domino circuit of Fig. 4(b) with 2.5 V threshold or the NORA circuit of Fig. 4(c). The three static chains used a CRF repeater with $V_F = 2.5$ V and $V_R = 1.3$ V. One static chain used nMOS transfer gates, and two static chains used CMOS transfer gates.

VI. MEASUREMENTS

Fig. 22 compares the measured delay for Domino and NORA dynamic regenerative feedback repeaters. The difference in threshold between Domino and NORA results in a delay reduction that relatively increases with the repeater interval k . In Fig. 23, the delay of the rising transition for a static CRF repeater chain with nMOS transistors is compared for nMOS gate voltage V_{GG} at 5 V and at 6.5 V. By using overdriven switches, the rising delay is reduced by about 35%. Since in the test chip the disabled repeaters still contribute capacitance to the chain, the measured delay is linear with

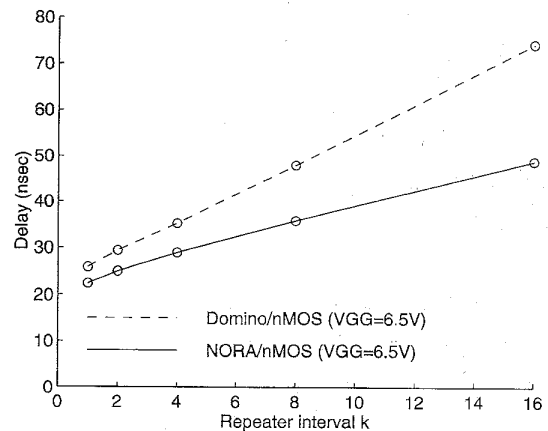


Fig. 22. Measured delay of dynamic regenerative feedback repeaters.

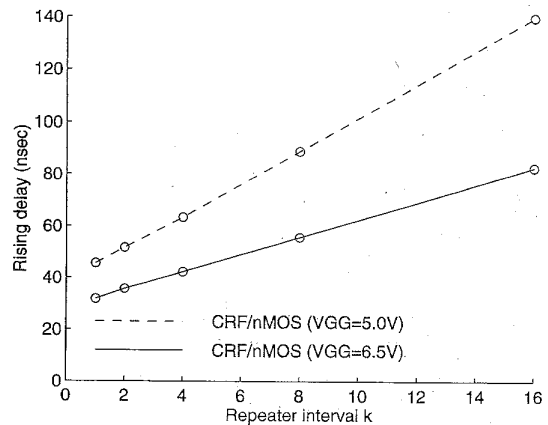


Fig. 23. Measured delay of CRF repeaters.

the repeater interval k and does not show a minimum as the curves in Fig. 11.

As expected from the higher RC values, the chains with CMOS switches had longer delays. The rising transition propagated slightly faster than the falling transition in the chain with overdriven nMOS switches in line with the discussion in Section III.

Table I shows the chip characteristics. The area of the repeaters included enable circuitry. For comparison with conventional repeaters; we added the estimated layout area from layouts that did not include enable circuitry. The higher current consumption of the CRF repeater over the dynamic repeater is mainly due to higher capacitive loading.

VII. CONCLUSION

We described the use of regenerative feedback repeaters in programmable interconnections. A static, complementary regenerative feedback (CRF) repeater was proposed, which regenerates the signal locally for a fixed time, after a transition has been detected. Experimental results from a 1.2- μm CMOS implementation show that the loaded delay through 64 switches for static and dynamic repeaters can be reduced by a factor 1.4–2. The main limitation of CRF repeaters, and of all complementary circuits that use regenerative feedback, is the

TABLE I
CHIP CHARACTERISTICS

Process	n-well 1.2 μ m CMOS	
Size	3.2mm \times 3.3mm	
Transistor count	21K	
Cell size (μ^2)	w/enable	[optimized]
Domino repeater	1200	[600]
CRF repeater (k=2)	3750	[2000]
Conventional repeater		[3600]
Current per chain:		
Domino rep. (k=4)	0.9mA @ 10MHz	
CRF rep. (k=4), nMOS sw.	1.5mA @ 10MHz	
I_{DD} per chain	0.1nA	

inability to propagate signals with arbitrary pulse width due to the presence of metastability. As a result, the use of CRF repeaters is limited to FPGA architectures that use self-timing or clocking to eliminate glitches.

ACKNOWLEDGMENT

The authors thank B. Fowler, M. Godfrey, B. Amrutur, and F. Klass for suggestions on design and testing, and D. Wingard for valuable comments.

REFERENCES

- [1] J. Rose, A. El Gamal, and A. Sangiovanni-Vincentelli, "Architecture of field-programmable gate arrays," *Proc. IEEE*, vol. 81, no. 7, pp. 1013–1029, July 1993.
- [2] J. Jack, D. Noble, and R. Tsien, *Electric Current Flow in Excitable Cells*. London, England: Oxford Univ. Press, 1985, ch. 10.
- [3] I. Richer, "The switch-line: A simple lumped transmission line that can support unattenuated propagation," *IEEE Trans. Circuit Theory*, vol. CT-13, no. 4, pp. 388–392, Dec. 1966.
- [4] L. Childs and R. Hirose, "An 18 ns 4 K \times 4 CMOS SRAM," *IEEE J. Solid-State Circuits*, vol. SC-19, no. 5, pp. 545–551, Oct. 1984.
- [5] L. Glasser and D. Dobberpuhl, *The Design and Analysis of VLSI Circuits*. Reading, MA: Addison-Wesley, 1985, p. 420.
- [6] J. Schutz, "A CMOS 100 MHz microprocessor," in *ISSCC Dig. Tech. Papers*, Feb. 1991, pp. 90–91.
- [7] R. Krambeck, C. Lee, and H.-F. Law, "High-speed compact circuits with CMOS," *IEEE J. Solid-State Circuits*, vol. SC-17, no. 3, pp. 614–619, June 1982.
- [8] N. Gonçalves and H. De Man, "NORA: A racefree dynamic CMOS technique for pipelined logic structures," *IEEE J. Solid-State Circuits*, vol. SC-18, no. 3, pp. 261–266, June 1983.
- [9] M. Shoji, *CMOS Digital Circuit Technology*. Englewood Cliffs, NJ: Prentice Hall, 1988, p. 168.
- [10] R. Guo *et al.*, "A 1024 pin universal interconnect array with routing architecture," in *Proc. CICC*, May 1992, pp. 4.5.1–4.5.4.
- [11] I. Dobbelaere, M. Horowitz, and A. El Gamal, "Regenerative feedback repeaters for programmable interconnections," in *ISSCC Dig. Tech. Papers*, Feb. 1995, pp. 116–117.
- [12] C. Portmann and T. Meng, "Metastability in CMOS library elements in reduced supply and technology scaled applications," *IEEE J. Solid-State Circuits*, vol. 30, pp. 39–46, Jan. 1995.



Ivo Dobbelaere (S'92–M'95) was born in Ninove, Belgium on June 13, 1966. He received the B.S. degree in electrical engineering from the Katholieke Universiteit Leuven, Belgium in 1989 and the M.S. and Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA in 1990 and 1995, respectively.


His fields of interest are in circuits and architectures for reconfigurable VLSI systems and in high-performance and low-power logic design.



Mark Horowitz (S'77–M'78–SM'95) received the B.S. and M.S. degrees in electrical engineering from the Massachusetts Institute of Technology, Cambridge, MA in 1978 and the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA in 1984.

Since September 1984, he has been working in the Computer Systems Laboratory at Stanford where he is currently an Associate Professor of Electrical Engineering. His research area is digital system design. He has led a number of processor design projects at Stanford including MIPS-X, one of the first processors to include an on-chip instruction cache, and TORCH, a statically-scheduled, superscalar processor. He has also worked in a number of other chip design areas including high-speed memory design, high-bandwidth interfaces, and fast floating point. In 1990, he took a leave from Stanford to help start Rambus, Inc., a company designing high-bandwidth memory interface technology. His current research includes multiprocessor design, low-power circuits, memory design, and processor architecture.

Dr. Horowitz is the recipient of the 1985 Presidential Young Investigator Award, an IBM Faculty Development Award, as well as the 1993 Best Paper Award at the 1994 International Solid-State Circuits Conference.



Abbas El Gamal (S'71–M'73–SM'83) received the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, in 1978.

He is currently an Associate Professor of Electrical Engineering at Stanford University. From 1978 to 1980, he was an Assistant Professor of Electrical Engineering at the University of Southern California. From 1981 to 1984, he was an Assistant Professor of Electrical Engineering at Stanford. He was on leave from Stanford from 1984 to 1987—first as Director of LSI Logic Research Lab, then as cofounder and Chief Scientist of Actel Corporation. In 1990, he cofounded SiArc, which is now part of Synopsys Inc. His research interests include VLSI circuits, architectures, and synthesis, FPGA's and mask programmable gate arrays, smart sensors, image compression, error correction, and information theory. He has authored or coauthored more than 60 papers and holds 20 patents in these areas.