A Standard Cell Compatible Bidirectional Repeater with Thyristor Assist

Sudhir Satpathy, Dennis Sylvester, David Blaauw
University of Michigan, Ann Arbor

Abstract

A thyristor-assisted standard cell compatible self-timed bidirectional repeater with no configuration overhead enables 8mm interconnects to achieve 37% higher speed at 20% lower energy over conventional repeaters in 65nm CMOS at 1.0V. Bidirectional operation without the need for configuration logic removes the need for clocking, yielding up to 14× higher energy efficiency at low data switching activity.

Motivation and Proposed Approach

Bidirectional interconnects are an integral part of global communication networks in multiprocessor chips. They facilitate high bandwidth with low silicon overhead by eliminating the need for replicating unidirectional signal wires [1,2]. Conventional bidirectional links are based on duplication of unidirectional repeaters [3], one of which is selectively activated for signal propagation. This implementation incurs logic and interconnect overhead to configure the repeaters, degrading performance and energy efficiency. Additionally, a synchronizing signal in the form of a clock is needed to eliminate contention when reversing signal propagation direction. Further, a bidirectional link can be driven at multiple locations in many snoop-based signaling schemes, making signal propagation information challenging to obtain *a priori* [4]. Recently, a number of custom-designed repeater-less signaling techniques have achieved high speed with low energy dissipation based on reduced voltage swing [5,6,7,8,9]. Such techniques typically require careful custom design that is tailored to each specific interconnect situation and involves precise device matching, additional supply voltages, and wider wire thickness/pitch. Hence, they cannot be easily used in synthesis-based design flows or re-used in different interconnect situations in the same design.

Instead, we target a drop-in replacement for conventional repeaters within a standard cell based design flow. To this end, we present a thyristor-assisted bidirectional signaling (TABS) technique with the following key features. 1) The fully static circuit implementation allows TABS to be used as a standard cell in synthesis-based design flows; 2) The self-timed bidirectional repeater has no configuration overhead and enables 8mm interconnects to achieve 37% higher performance at 20% less energy per bit compared to conventional repeaters in 65nm CMOS at 1.0V; 3) Absence of configuration logic obviates the need for clocking, simplifying the design flow and providing up to 14× better energy efficiency for low data switching activity; 4) Robust operation across 0.6–1.2V supply targetting usage in voltage scaled SoCs. Besides improving performance, these features increase the optimal repeater insertion interval for TABS from 625mm to 1mm, yielding 38% fewer repeaters in the signal propagation path. At 1.5mm repeater spacing TABS energy efficiency is improved by 51% and repeater count is reduced by 58% while maintaining iso-performance with conventional optimally spaced repeaters. We also present a synchronous version of the repeater for use at synchronizing boundaries in place of conventional bidirectional flip-flops, saving 36% area.

Fig. 1 shows a TABS half-repeater with a pair of NMOS and PMOS transistors arranged in a CMOS "thyristor configuration" and used as a transition amplifier. The expanded circuit senses a falling transition on *IN_L* and/or a rising transition on *IN_R*, thereby causing the thyristor to switch and pull both interconnect nodes to opposite supply rails. Two such repeaters are used in parallel with their ports connected to opposite nodes of the link to form a TABS bidirectional repeater. Each repeater has 3 operating states: *High gain*, *low gain*, and *passive*. In the absence of any switching or when *reset* is asserted, the repeater is in *low gain* state with the thyristor nodes pre-charged/pre-discharged to voltage levels where the sensing transistors have their lowest gain ($V_{gs} = 0$). The sensing nodes (*S1/S2*) are connected to *IN_L* and *IN_R* through transmission gates (*T1/T2*) that are turned ON. *S1/S2* node voltages are held by keepers placed on *IN_L* and *IN_R* (not shown). When *IN_L* transitions from high to low, node *S1* initially follows it. The PMOS device in the thyristor gradually turns ON, raising *S2* which causes the NMOS thyristor device to turn ON. This regeneration mechanism causes *S1* and *S2* to switch rapidly. Inverter *d_n (d_p)* is skewed with a stronger NMOS (PMOS) to provide a faster response to *S1*/*S2* switching and quickly turn OFF transmission gates *T1 (T2)*. This decouples *S1 (S2)* from *IN_L (IN_R)* and allows for faster switching. Once the transistors in the thyristor transition, the repeater enters the *high gain* state where the sensing transistors have highest gain ($V_{gs} = V_{DD}/-V_{DD}$), causing the large internal drivers *(D_L/D_R)* to rapidly pull *IN_L* and *IN_R* towards supply rails. Once *IN_L* and *IN_R* have both transitioned, their delayed signals *in_l_d* and *in_r_d* enable the pre-charge signals *cut_n* and *cut_p* which reset and hold the thyristor to its precharge state. The repeater is now in the *passive* state with its thyristor disconnected from *IN_L* and *IN_R*. Finally, when *IN_L* and *IN_R* transition back, the repeater again enters the *low-gain* state.

While one repeater is in the *low gain* state, the other half-repeater is in the *passive* state. A transition on the interconnect causes the sensing repeater to switch from *low gain* to *passive* state via the *high gain* state while the other repeater transitions from *passive* to *low gain* and begins sensing the interconnect to detect the next transition (Fig. 2).

Using the same technique, a bidirectional latch is also designed for data hand-off at synchronizing boundaries as shown at bottom right in Fig. 2. The transition direction is stored in a 6T SRAM based on the state of the interconnect at the falling edge of *clock,* which is later used in the positive phase for sensing the appropriate transition.

TABS offers 3 key advantages over inverter-based repeaters: 1) The regeneration allowed by the decoupling mechanism makes the latency less dependent on slow slew rates on long interconnect as commonly seen in global wires; 2) In a conventional repeater a transition is performed only by the driver, as opposed to TABS where the receiver aids the driver after detection; 3) The internal feedback and state tracking mechanism eliminates the need for a global synchronizing signal.

Measurement Results

Fig. 3 shows simulated energy / delay curves for an 8mm link with varying repeater sizes. Optimal insertion interval is 1000µm for TABS and 625µm for conventional repeaters reducing the number needed repeaters. A test prototype was fabricated in 65nm bulk CMOS where TABS was treated as a standard cell replacement for traditional repeaters in the design flow. Both links were implemented using M4 and M5 with 100nm wire width and 200nm spacing (2x min.) and worst-case aggressor switching. In silicon measurements, the TABS link operates at 732MHz, which is 37% faster than the conventional link at 1.0V while dissipating 20% less energy (Fig. 4.) Due to the absence of any configuration overhead TABS improves energy efficiency by 27% to 47% with increasing bidirectional traffic at iso-throughput of 602Mb/s per link (Fig. 5). The absence of clock facilitates seamless signaling across different frequency domains and improves TABS energy efficiency by up to 14× at a low, 0.005 data switching activity over conventionally bidirectional repeaters (Fig. 6.) To exploit the ability of TABS to operate with longer insertion intervals, we also implemented TABS with a 1.5mm insertion interval (measured results shown in Fig. 7). TABS energy efficiency increases to 51% over conventional repeaters at iso-performance while using 58% fewer repeaters. This makes TABS highly suitable for links that run over caches and other IP blocks that prohibit frequent repeater insertion. TABS is fully functional across a temperature range of -20ºC to 90ºC as shown in Fig. 7 at 1.0V.

References

[1] C.Park *et al*, ISSC 2010, pp. 181-182
[2] S.Satpathy *et al*, SoVC 2011, pp. 138-139
[3] B.Stackhouse *et al,* JSSC, 2009, pp. 18-31, Vol. 44

[4] M.Ghoneima *et al,* TVLSI, 2008, pp. 1904-1910
[5] B.Kim *et al,* ISSCC, 2009, pp. 66-67
[6] D.Schinkel *et al,* TVLSI, 2009, pp. 12-21
[7] R.Ho *et al,* ISSCC, 2010, pp. 412-413
[8] E.Mensink *et al,* ISSCC, 2007, pp. 414-415
[9] M.Nazari *et al,* ISSCC, 2011, pp. 446-447

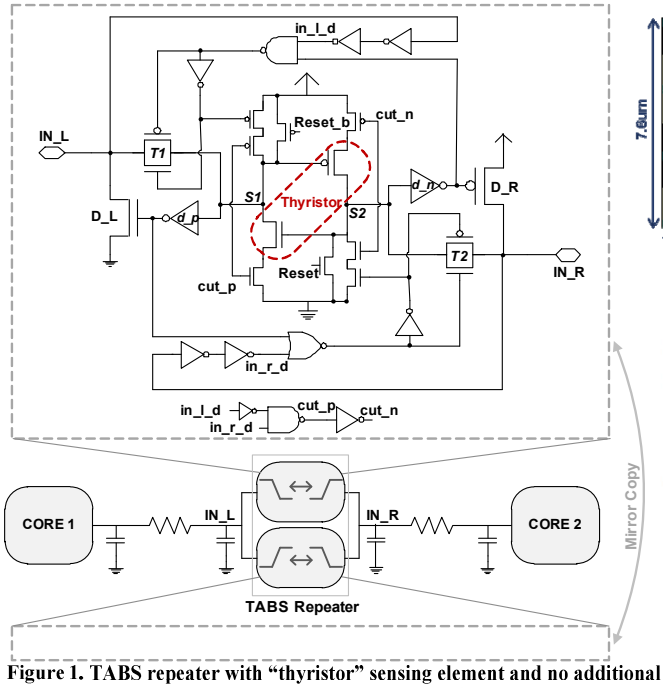

Figure 1. TABS repeater with "thyristor" sensing element and no additional configuration logic.

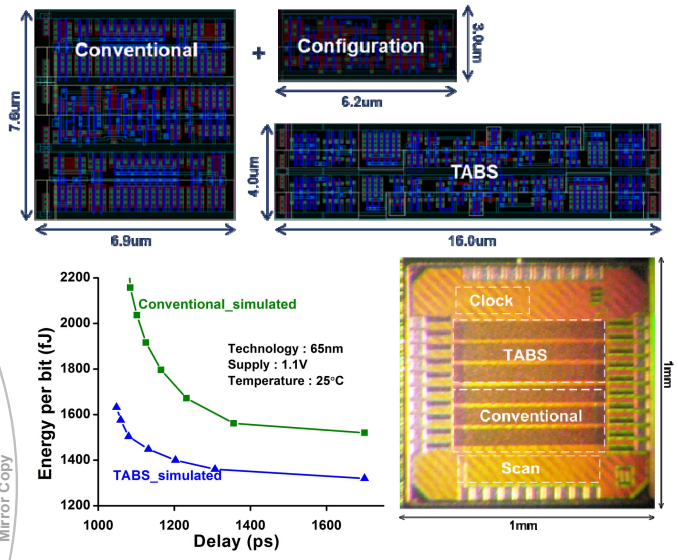

Figure 3. Simulated energy versus delay curves for TABS (inserted every 1mm) and conventional bi-directional repeaters (inserted every 625μm) in 65nm for 8mm link. TABS spans 11% less area than conventional repeaters with configuration overhead.

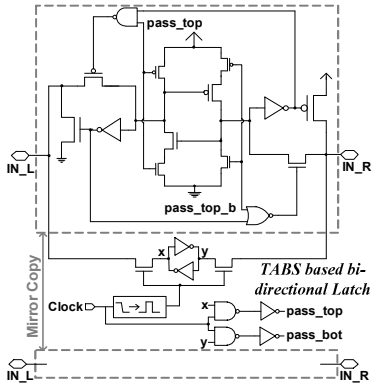

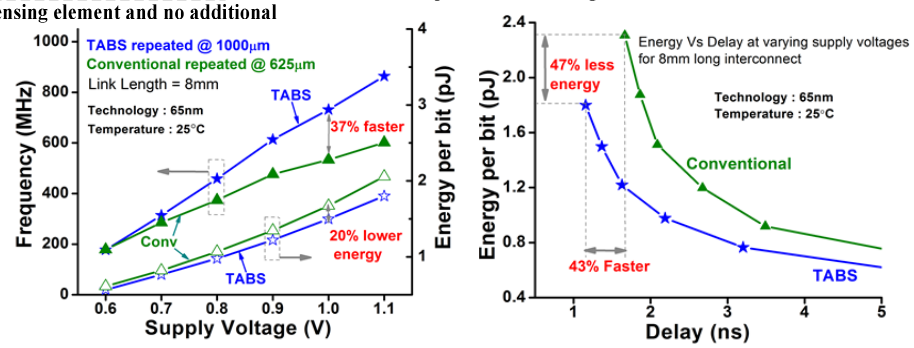

Figure 4. Measured performance and energy for 8mm long interconnect with TABS (1mm insertion interval) and conventional repeater (625μm insertion interval) at 50% bi-directional traffic. TABS is 43% faster and dissipates 47% lower energy at 1.1V and 100% bi-directional traffic

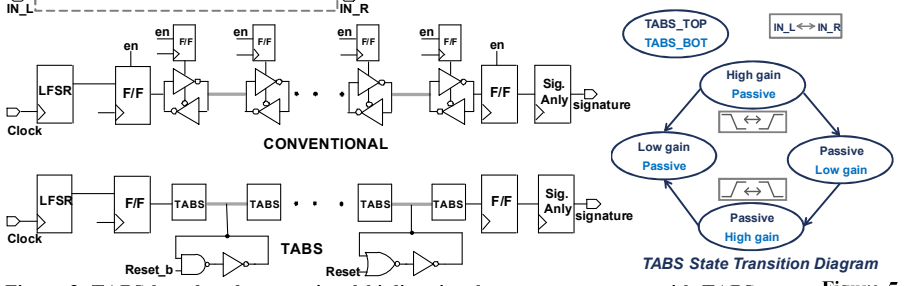

Figure 2. TABS-based and conventional bi-directional repeater structure with TABS-based bi-directional latch. Thyristor is used for transition amplification as well as a memory element for tracking state, eliminating the need for additional control logic.

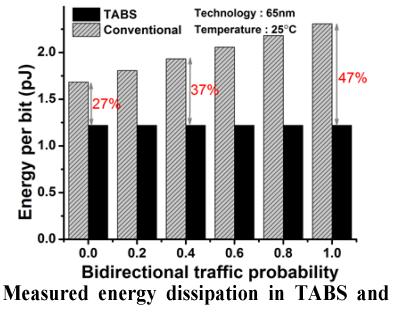

Figure 5. Measured energy dissipation in TABS and conventional repeaters with varying percentage of bi-directional traffic at iso-performance at 1.1V

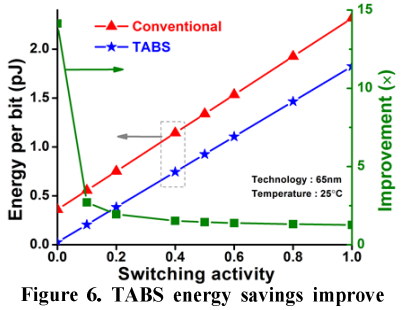

Figure 6. TABS energy savings improve from 27% to 14× as switching activity reduces.

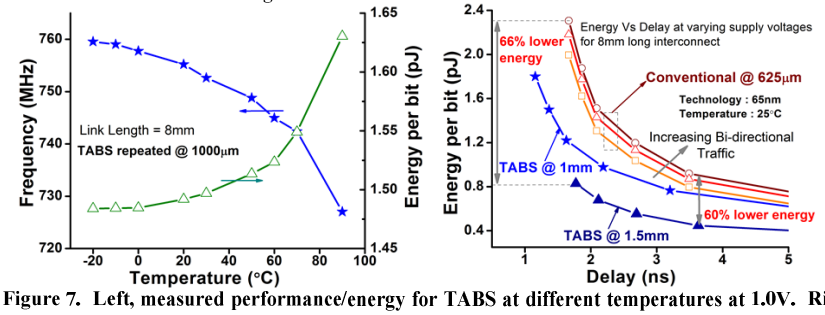

Figure 7. Left, measured performance/energy for TABS at different temperatures at 1.0V. Right, TABS energy efficiency improves by 51% (measured) over conventional repeater at iso-performance at 1.0V with 58% fewer repeaters by increasing insertion interval to 1.5mm.