LC^2: Limited Contention Level Converter for Robust Wide-Range Voltage Conversion

Yejoong Kim, Dennis Sylvester, David Blaauw
University of Michigan, Ann Arbor, MI
yejoong@umich.edu, dennis@eecs.umich.edu, blaauw@umich.edu

Abstract

We propose a robust single-stage static level converter called LC^2 that uses a pulsed control strategy to avoid contention. It reliably converts from 0.3V to 2.5V across wide PVT ranges. Fabricated in 130nm CMOS, 80 measured converters have an average delay of 2.38FO4 and 229fJ switching energy, marking 3.0× and 7.4× improvements over the best prior work. It consumes 475pW static power.

Introduction

Low-voltage circuit design has been widely investigated for ultra-low power applications, reaching as low as 230mV in a recent multi-pipelined processor [1], and requiring wide-range level conversion (LC) for communication with IO pads and high-voltage circuit blocks. In addition, cores on a chip multiprocessor are increasingly voltage scaled independently [2], necessitating LC between core voltage domains in high performance applications. However, level conversion is challenging at reduced voltages since conventional approaches suffer from severe contention between weak pull-down devices and strong pull-up devices, making them vulnerable to PVT variations.

Fig. 1 shows the operation of a conventional DCVS approach. A Zero-V_{TH} device prevents oxide breakdown in the thin oxide devices, making it possible to use a fast SVT pull-down device [3]. The DCVS LC suffers from a two-sided constraint on the PMOS device: if the PMOS is too weak, the pull-up transition becomes slow and the node may not be kept high, giving rise to performance/robustness issues; if the PMOS is too strong, the NMOS cannot overcome it and the circuit fails. The current margin plots in Fig. 1 show that severe variations at the low voltages exacerbate this two-sided constraint. Although the circuit is designed such that $I_{NMOS} >> I_{PMOS}$ to discharge node *n1* or *n2*, as little as 2σ V_{TH} variation causes failure due to $I_{NMOS} < I_{PMOS}$. Increasing NMOS size by 3.5× guarantees 3σ robustness, but results in very large devices (W_{NMOS} = 105µm) with undesirable leakage (9nA). In addition, the increased diffusion capacitance slows the pull-up transition. This two-sided constraint severely limits DCVS LC robustness under PVT variation. Multiple LC stages can improve robustness but introduce overhead due to intermediate supplies and increased latency. Other static LCs [4][5] have similar two-sided constraints and require precise transistor sizing, and have lacked silicon measurements. A recently proposed dynamic LC [6] uses a high-voltage clock, which improves robustness but increases layout size and power consumption. Furthermore, none of the previous LCs has demonstrated robustness through comprehensive silicon measurements.

Proposed Solution

We propose a new approach called Limited Contention Level Converter (LCLC or LC^2) that eliminates the two-sided constraint without the use of high-voltage clocks. Fig. 2 shows the conceptual operation of LC^2. Before the rising transition, node *n1* is held high by the weak keeper, which is subthreshold-biased, while all other switches are off; hence $V_{n1}=V_{DDH}$ and $V_{n2}=0$. Once V_{IN} rises to V_{DDL}, the pull-down driver starts to discharge *n1* and easily overcomes the weak keeper. This transition on *n1* causes "Pull-Up Control" to activate both the weak keeper and the strong switch on the other side, which quickly charges up *n2*. "Pull-Down Control" is then triggered to directly connect *n1* to ground, rapidly discharging it and completing the transition. Finally, a delay element turns off all switches (except the appropriate keeper) after all transitions are finalized. The next transition can then proceed such that the only contention is with the weak keeper. The use of separate and different strength pull-up devices for holding state and charging/discharging *n1* and *n2* substantially improves design robustness and performance.

Fig. 3 shows the schematic of LC^2 with detailed timing waveforms. At the beginning of a rising transition, $V_{n1}=V_{n3}=V_{DDH}$ and $V_{n2}=V_{n4}=0$, hence M6 and M11 are off and M1 contends only with the weak keeper Mx. Once M1 and M3 start to discharge *n1*, positive feedback from M10 and M7 boosts transition speed by pulling the gate of M7 to V_{DDH}. Thus, M10 can be sized for fast rising transitions on *n2* (using a min length device). In contrast, this transistor must remain weak in the conventional approach to minimize the contention, making it slower and less robust. Once the transition completes, M5 and M12 are turned off after an inverter chain delay to prepare for the next transition. Devices M5-M12 use minimum width, and the inverter chains simply require sufficient delay to fully charge *n1* or *n2*, simplifying device sizing. Although the pull-down drivers (M1 and M2) and keepers should be carefully sized, keeper size can be easily determined using known techniques [7], after determining M1 and M2 sizes based on the desired speed-power trade-off. A simple diode chain is used to generate the keeper voltage (V_{KEEPER}), setting the current supplied by the keeper. The current margin plot in Fig. 4 shows that this design is robust to $>3\sigma$ variation in simulation. Simulation results in Fig. 5 indicate that DCVS is highly vulnerable to V_{TH} shifts, while LC^2 functions correctly within the entire process corner without significant delay change.

Measurements

We measured 40 dies in 130nm CMOS; each die has two LC^2s and two DCVS LCs designed for 0.3V to 2.5V conversions (V_{DDL}=0.3V, V_{DDH}=2.5V) with a minimum-sized inverter as an output load. Fig. 6 shows measured delay across temperature. LC^2 is 3.2× faster than DCVS with 2.38FO4 delay at 25°C (FO4 measured at V_{DDL} supply and corresponding temperature). In addition, DCVS shows a 10.4× delay change across 10~100°C, while LC^2 changes by only 4.3×. Normalizing to FO4 delays, LC^2 delay increases 18% from 10 to 100°C while DCVS worsens by 104%. This is due to the much reduced contention in LC^2.

Fig. 7 shows measured power consumption across temperature. While DCVS consumes 7.15nW static power, LC^2 consumes 15× less (475pW) at 25°C, mainly due to the smaller pull-down device (1.5µm). It consumes 2.29nW active power at 25°C which is 4.9× less than DCVS (11.21nW), as well as nearly constant active power over a wide temperature range. Due to the lack of contention, its active energy is dominated by charging of capacitances rather than short-circuit current as in DCVS, making it temperature insensitive. Active power changes only 2% (from 2.27nW to 2.32nW) in the 10~100°C range while DCVS shows a 7.7× change (from 4.15nW to 31.88nW) and high power consumption at low temperature. Unlike LC^2, not all 80 DCVS LCs function below 10°C since the low temperature increases V_{TH}, weakening the NMOS exponentially and the PMOS linearly, exacerbating contention.

To show the impact of process variations, Fig. 8 displays measured delay distributions for the LCs at 25°C. LC^2 shows 6× smaller standard deviation than DCVS. For voltage variations, Fig. 9 shows performance degradations across voltage drop. While DCVS delay increases by 7.7× with 10% V_{DDL} drop, LC^2 slows by only 6% (normalized to FO4 delays at the corresponding voltages), indicating that the keeper sizing strategy is sufficiently robust to handle expected voltage variations. Fig. 10 shows the number of operating LCs at 1MHz across temperature. DCVS was designed to operate as fast as 20MHz at 25°C, and the 1MHz clock allows 20× delay degradation. While all LC^2s operate reliably in the -20~100°C range, the first DCVS fails at 20°C, and only 5 of 80 work at -20°C, showing the robustness of LC^2 to PVT variations.

Fig. 11 shows the die photo and comparisons to recent work. Despite having more transistors than DCVS, LC^2 is smaller than DCVS in layout even including the extra diode chain, which can be shared among multiple LC^2s. The static nature of LC^2 does not require clocks or complex synchronizing schemes, enabling 1093× smaller area and 7.4× lower energy per transition compared to recent work in 130nm [6], as well as 3× faster speed.

References

[1] H. Kaul *et al*., ISSCC, 2009, pp. 260-261.
[2] J. Howard *et al*., ISSCC, 2010, pp. 108-109.
[3] W. Wang *et al*., Symp. VLSI-TSA, 2001, pp. 307-310.
[4] H. Shao *et al*., ESSCIRC, 2007, pp. 312-315.
[5] I. Chang *et al*., ISLPED, 2006, pp. 14-19.
[6] I. Chang *et al*., Trans. VLSI, Aug. 2010.
[7] M. Seok *et al*., CICC, 2008, pp. 423-426.

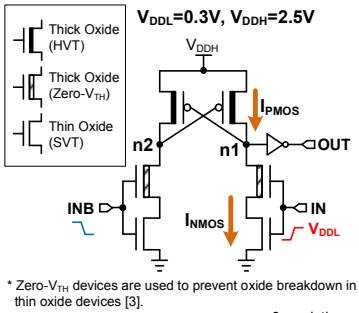

* Zero-V_{TH} devices are used to prevent oxide breakdown in thin oxide devices [3].

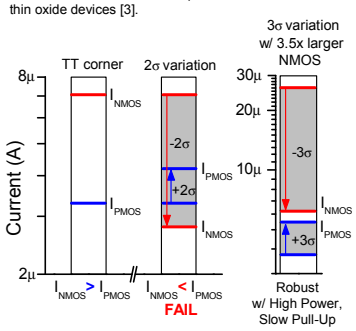

Figure 1. DCVS LC and its current margin plots indicate that only 2σ-variation causes functional failure.

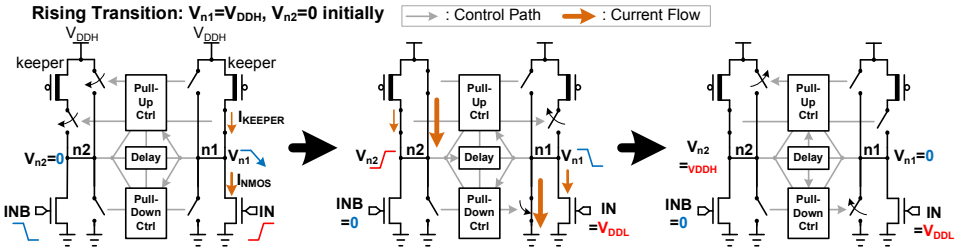

Figure 2. LC^2 eliminates the strong contention by contending only with the weak keeper at the start of the transition. Once the pull-down device overcomes the weak keeper, "Pull-Up/Pull-Down Ctrl" boosts the transition speed, and "Delay" turns off all the main switches after the transition is finalized.

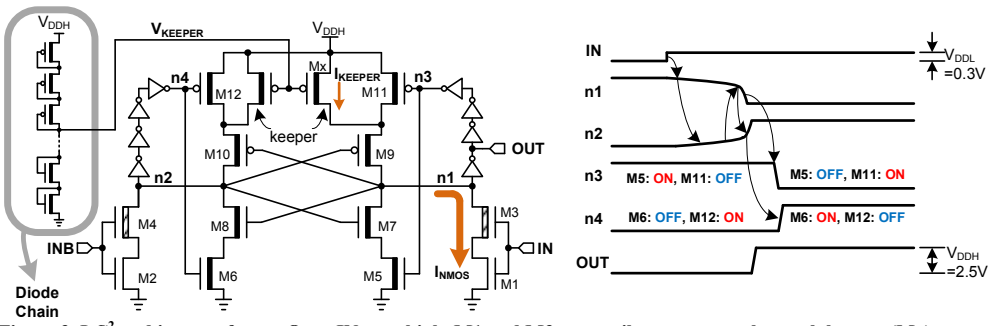

Figure 3. LC^2 and its waveforms. Once IN goes high, M1 and M3 can easily overpower the weak keeper (Mx), discharging n1. M7~M10 boosts up the speed of the transition since M6 and M11 remain turned off. The inverter chain turns off M5 and M12 after the transition in order to prepare for the next transition. A simple diode chain is used for generating V_{KEEPER}.

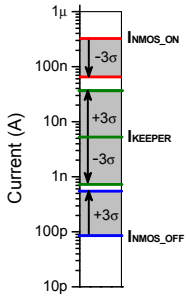

Figure 4. The current margin plot of LC^2 shows that it is robust to 3σ-variations.

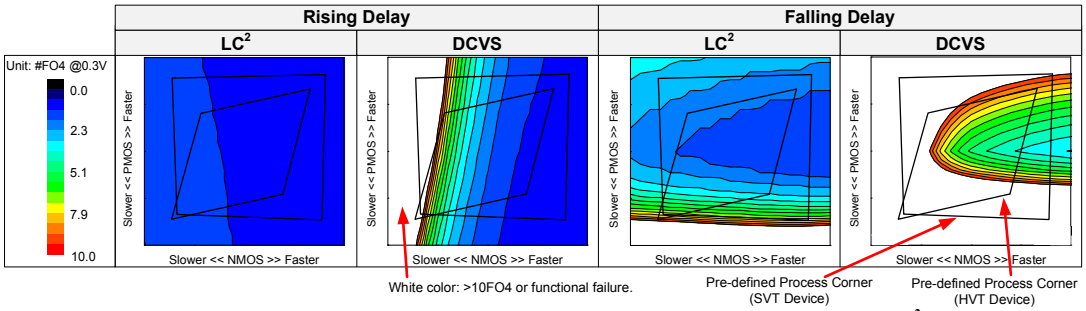

Figure 5. The simulation results show that DCVS is vulnerable to V_{TH} shifts (process variations), while LC^2 works correctly within the entire process corner without significant delay change. Note that the vertices of polygons represent the pre-defined process corners (FF, FS, SF, SS) of the specified devices above. White regions indicate a delay larger than 10 FO4 or functional failure.

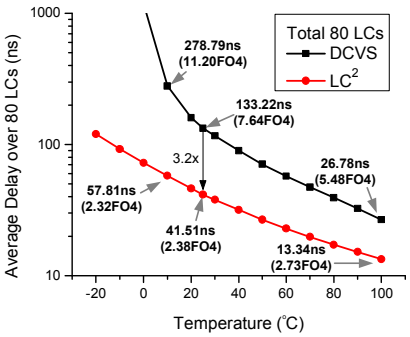

Figure 6. Measured delay compared to DCVS

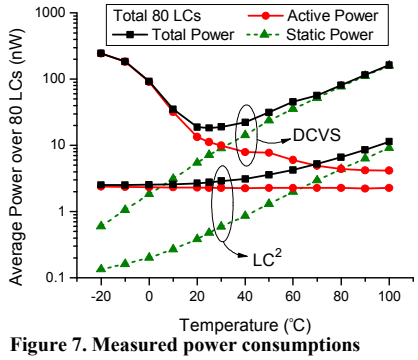

Figure 7. Measured power consumptions (freq=5kHz, α=2)

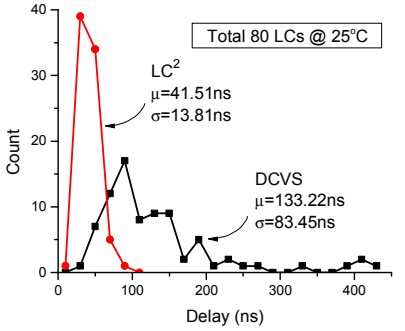

Figure 8. Measured delay variations

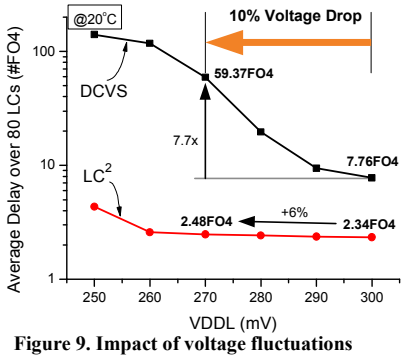

Figure 9. Impact of voltage fluctuations

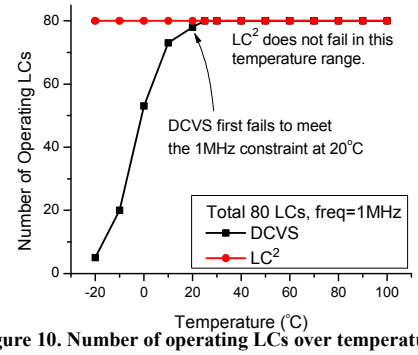

Figure 10. Number of operating LCs over temperature

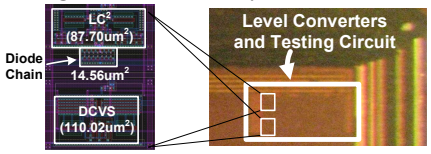

Figure 11. Die photo and a comparison table

25°C	LC^2	[6]	[4]
Technology	130nm	130nm	180nm
Conversion	0.3V to 2.5V	0.3V to 2.5V	0.25V to 1.8V
Type	Static	Dynamic (w/ 2.5V clock)	Static
Delay	41.51ns	125ns	~190ns
Static Power	475pW	N/A	N/A
Energy/ Transition	229fJ	1.7pJ	~5pJ
Area	102.26μm^2 (including the diode chain)	0.1118mm^2 (1093x larger than LC^2)	Silicon measurement not reported