

The Fanout-of-4 Inverter Delay Metric

David Harris, Ron Ho, Gu-Yeon Wei, and Mark Horowitz
Stanford University, Stanford, CA 94305

Introduction

Digital circuit delays vary with feature size, process corner, operating voltage, and junction temperature. Delays are steadily decreasing with advances in process technology, so comparing results reported in nanoseconds between process generations is difficult. This paper proposes using the delay of a fanout-of-4 inverter (FO4) to normalize process and operating condition variations and quantifies how well this normalization works.

A novel application of this correlation is a power-reduction technique. Power supply and operating frequency can be regulated on the fly to minimize power while a chip is performing non-critical operations while allowing full-speed operation when necessary. Proposed implementations [1,2,3] rely on a good correlation between ring-oscillator frequency and critical path latency. The tracking of chip delays with FO4 delay determines the necessary extra margin for functionality over process and environmental variation.

Fanout-of-4 Inverter Delays

We select the fanout-of-4 inverter as a representative delay element because such a fanout is typically used in tapered buffers driving large loads [4] (the optimal fanout is technology dependent, but delay is generally within 5% of minimum over a fanout range of 2.7 to 5.3). FO4 delays are also useful when thinking about circuits; e.g. a control signal driving a 64-bit datapath requires about $\log_4 64 = 3$ FO4 delays of buffering to drive the heavy load. Simulation shows that other fanout inverters track very well with FO4 delays.

Fig. 1 shows the simulation setup for determining the FO4 delay for a given process and environment. The first inverter shapes the input waveform to have realistic rise and fall times. The last inverter slows the switching time of the third inverter, preventing excessive Miller multiplication of the third inverter's C_{gd} . The same temperature and voltage should be used to measure FO4 delay as to measure other path delays. We chose nominal process and voltage at 70 degrees. The figure also shows actual delays of a variety of processes at different operating voltages.

Simulation Results

We ran simulations to measure how gate delay tracks with process, process corner, temperature, and voltage. Our baseline is the MOSIS CMOS14B 0.6 μm process running in the TT corner at 70 degrees and 3.3 volts. Fig. 2 shows the delay of various fanout-of-4 gates across process measured in FO4 delays. Across technology from 1.2 μm down to 0.35 μm , the maximum deviation is only 11% except in the case

of domino gates which track to within 14%. Of particular interest, the “adder” gate is a complete 64-bit adder self-bypass path [5] including domino logic, transmission gates, and interconnect. Figs. 3 and 4 show how gate delay expressed in FO4 delays varies with voltage and temperature. Maximum variation in FO4 delay of gates over voltage (2-3.5 v), temperature (0-125 °C), and process corner is summarized in Fig. 5. Finally Fig. 6 plots delay vs. voltage for a buffer driving a 0-5 mm wire plus a FO4 load.

Measured Results

We measured the access time of a fabricated low-power SRAM relative to the delay of an on-chip ring oscillator, and also normalized the delay of a fabricated color subband decoder to FO4 delays [6]. Fig. 7 shows these results for voltages ranging from 150% of V_{th} to the process's maximum power supply voltage.

Limitations

Simple static circuits track well with FO4 delay. Worse mismatches occur for paths dominated by a single type of transistor, paths operating at the extremes of operating conditions, paths dominated by diffusion and not gate capacitance (the ratio of diffusion to gate capacitance changes with process), and paths with significant wire RC delay. In the real paths examined, the larger variation of some elements is balanced by small variation of other elements (like inverters and simple gates), resulting in less severe variation than predicted by looking at isolated gates.

Conclusions

The delay of a gate-dominated path tracks well with the delay of a fanout-of-4 inverter. Reporting the FO4 delay of a process along with any circuit performance results (both taken at the same operating conditions) will facilitate comparing the circuit to alternative implementations in other processes. Over a wide range of processes, the delay measured in FO4 inverters changes by less than 15% for most domino circuits and 11% for static circuits. Over a wide range of process and environment, domino gates in the 0.6 μm process vary up to 30% while static gates only vary 20% in FO4 delay. This margin is required to “dead reckon” cycle time with a ring oscillator. Therefore a chip clocked with a ring oscillator may allow better typical performance as well as reduced power during non-critical computation.

Acknowledgements

The authors gratefully acknowledge B. Amrutur and B. Gordon for providing measurement data from their fabricated chips.

This research was partially supported by NSF and ARPA contracts #DABT63-95-C-0089 and #J-FBI-92-194.

Notes:

- Underlining added by B. Baas for EEC 116
- "TT corner" means "typical NMOS" and "typical PMOS" process conditions

References

- [1] P. Maken, M. Degrauwe, M. Van Paemel, H. Oguey, "A Voltage Reduction Technique for Digital Systems," *ISSCC*, February 1990.
- [2] V. Gutnik and A. Chandrakasan, "An Efficient Controller for Variable Supply-Voltage Low Power Processing," *Symposium on VLSI Circuits*, June 1996, pp. 158-159.
- [3] G. Wei and M. Horowitz, "A Low Power Switching Power Supply for Self-Clocked Systems," *SLPED*, August 1996, pp. 313-318.
- [4] L. Gal, "Reply to 'Comments on the optimum CMOS tapered buffer problem,'" *JSSC*, vol. 29, pp. 158-159, February 1994.
- [5] D. Harris and M. Horowitz, "Skew-Tolerant Domino Circuits," *ISSCC*, February 1997.
- [6] B. Gordon, T. Meng, N. Chaddha, "A 1.2mW Video-Rate 2D Color Subband Decoder," *ISSCC*, February 1995.

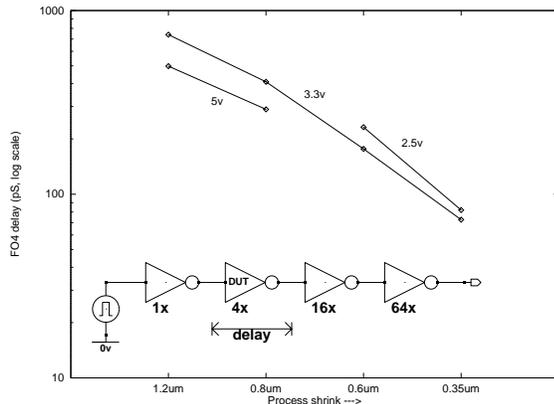


Figure 1: Fanout-of-4 inverter delays

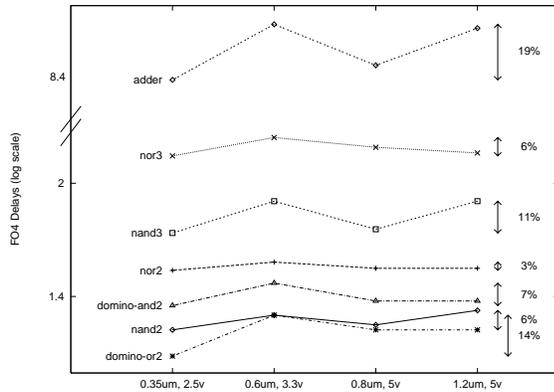


Figure 2: Gate delay vs. process

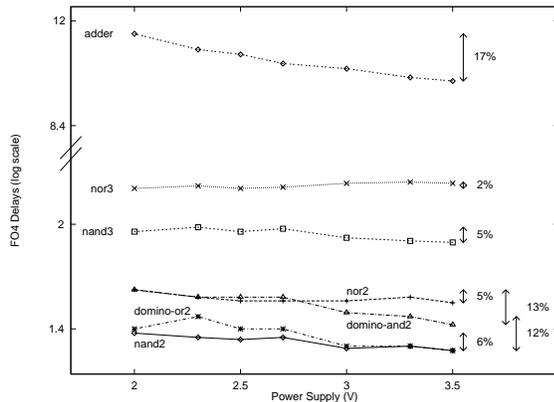


Figure 3: Gate delay vs. voltage

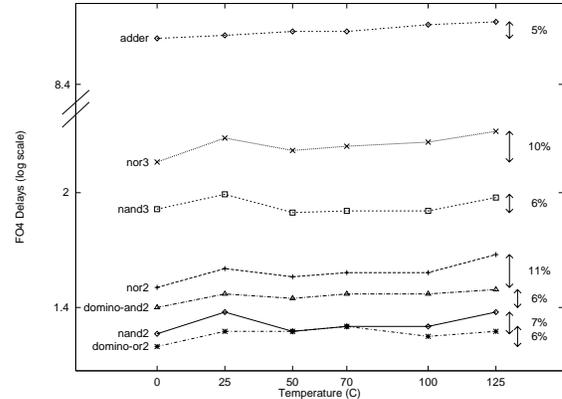


Figure 4: Gate delay vs. temperature

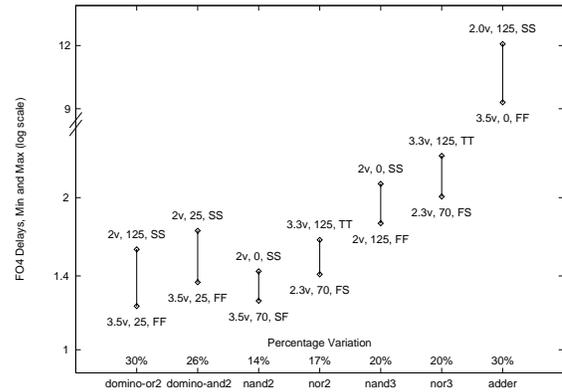


Figure 5: Maximum variation over voltage, temperature, corner

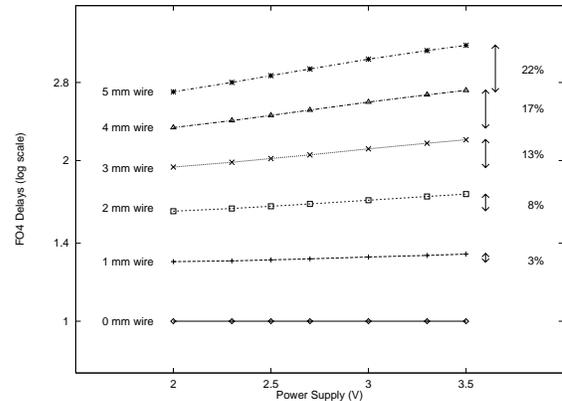


Figure 6: Driver + Interconnect delay vs. voltage

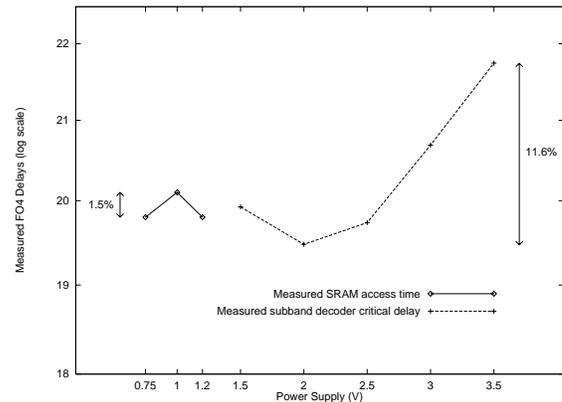


Figure 7: Measured delay vs. voltage

Notice how relatively constant (flat) the delays are in terms of FO4 delays when the process, supply voltage, temperature are varied! This means that FO4 is a very stable way to measure delays.