Brief Announcement: TRIX: Low-Skew Pulse Propagation for Fault-Tolerant Hardware

Christoph Lenzen* and Ben Wiederhake**

Max Planck Institute for Informatics, Saarland Informatics Campus; and Saarbrücken Graduate School of Computer Science, Germany

Abstract. We present a simple grid structure to use in a fault-tolerant clock propagation method and study it by means of simulation experiments. A key question is how well neighboring grid nodes are synchronized, even without faults. Our statistical approach provides substantial evidence that this system performs surprisingly well. In a grid of height H, the standard deviation of the delay seems to be $O(H^{1/4})$ (≈ 2.7 link delay uncertainties for $H = 2000$) and the standard deviation of the skew to be $o(\log\log H)$ (≈ 0.77 link delay uncertainties for $H = 2000$).

1 Introduction

Traditionally, clocking of synchronous systems is performed by clock trees or other structures that cannot sustain faulty components [12]. This imposes limits on scalability on the physical size of clock domains. To the best of our knowledge, work on fault-tolerant clocking schemes started in earnest in the last decade, with an upsurge of interest in single event upsets of the clocking subsystem [1, 2, 8, 10]. Larger systems and smaller components require going beyond these techniques.

There is a significant body of work on fault-tolerant synchronization from the area of distributed systems considering Byzantine faults [9, 11]. A line of works culminating in [6] additionally consider *self-stabilization*, the ability of a system to recover from an unbounded number of transient faults. These highly desirable properties come at a high price, usually in the form of high connectivity [4].

A suitable relaxation of requirements is proposed in [3], requiring that Byzantine faults are distributed across the system not in a worst-case fashion, but more "spread out". Distributing a clock signal through a grid-like network called HEX is proposed, which tolerates one out of each node's four in-neighbors being faulty. Unfortunately, HEX has poor synchronization performance: a crashed node causes a "detour" resulting in a clock skew between neighbors of at least one maximum node-to-node communication delay d. This is much larger than the *uncertainty* u in the node-to-node delay, which is engineered to be small ($u \ll d$).

We propose a novel clock distribution topology that overcomes the above shortcoming of HEX, in particular the high skew between neighboring nodes. Similar to HEX, the clock signal is propagated through layers, but for each node,

* clenzen@mpi-inf.mpg.de

** bwiederh@mpi-inf.mpg.de

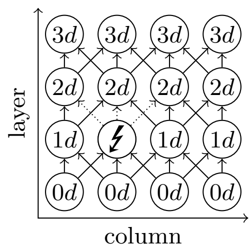

Fig. 1: Zoom-in on a part of a larger TRIX grid with one crashing node. Observe that the fault causes no significant additional skew.

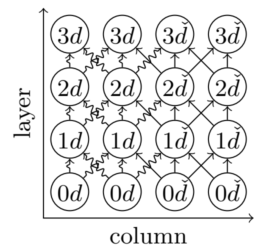

Fig. 2: Worst-case assignment of wire delays causing large skew for TRIX. Squiggly lines indicate slow wires, straight lines indicate fast wires. $\check{d} \coloneqq d - u$.

all of its three in-neighbors are on the preceding layer. If at most one in-neighbor is faulty, each node still has two correct in-neighbors on the preceding layer, as demonstrated in Figure 1. Hence, we can now focus on fault-free executions, because single isolated faults only introduce an additional uncertainty of at most $u \ll d$. Predictions in the fault-free model are therefore still meaningful for systems with rare and non-malicious faults.

The TRIX topology is acyclic, which conveniently means that self-stabilization is trivial to achieve, as any incorrect state is "flushed out" from the system.

Despite its apparent attractiveness and even greater simplicity, we note that this choice of topology should not be obvious. The fact that nodes do not check in with their neighbors on the same layer implies that the worst-case clock skew between neighbors grows as uH, where H is the number of layers and (for the sake of simplicity) we assume that the skew on the first layer (which can be seen as the "clock input") is 0, see Figure 2. However, reaching the skew of d between neighbors on the same layer, which is necessary to give purpose to any link between them, takes many layers, at least $d/u \gg 1$ many. This is in contrast to HEX, where the worst-case skew is bounded, but more easily attained.

While the worst-case behavior is easy to understand, it originates from a very unlikely configuration, where one side of the grid is entirely slow and the other is fast, see Figure 2. In contrast, correlated but gradual changes will also result in spreading out clock skews. Any change that affects an entire region in the same way will not affect local timing differences at all. This motivates to study the extreme case of independent noise on each link in the TRIX grid. Moreover, we assume "perfect" input, i.e., each node on the initial layer signals a clock pulse at time 0, and that the grid is infinitely wide. We argue that this simplistic abstraction captures the essence of (independent) noise on the channels.

We provide evidence that TRIX behaves better than conventional concentration bounds might suggest. The full version [7] argues in-depth that these results are not just artifacts of the simulation, the model, or due to various biases.

We point out the open problem of analyzing the stochastic process we use as an abstraction for TRIX. Understanding of the underlying cause would allow making qualitative and quantitative predictions beyond the considered setting.

2 Model

The network topology is a grid of height H and width W. To simplify, we choose $W = \infty$, because we aim to focus on the behavior in large systems. We refer to the grid nodes by integer coordinates (x, y), where $x \in \mathbb{Z}$ and $y \in \mathbb{N}_0$. Layer $0 \leq \ell \leq H$ consists of the nodes (x, ℓ), $x \in \mathbb{Z}$.

Nodes in layer 0 represent the clock source. Note that for the purposes of this paper, we assume that the problem of fault-tolerant clock signal *generation* has already been sufficiently addressed (e.g. using [5]), but the signal still needs to be *distributed*. All other nodes (x, ℓ) for $\ell > 0$ are TRIX nodes. Each TRIX node propagates the clock signal to the three nodes "above" it, i.e., the vertices $(x+c, y+1)$, $c \in \{-1, 0, +1\}$. Each of the wire delays is modeled as i.i.d. random variables $w_c^{x,y}$ (or w_c for short) that are fair coin flips, i.e., attain the values 0 or 1 with probability 1/2 each. This reflects that any absolute delay does not matter, as the number of wires is the same for any path from layer 0 (the clock generation layer) to layer $\ell > 0$; also, this normalizes the uncertainty from u to 1.

Let $d(x, y)$ be the time at which node (x, y) fires. Clock generation provides us with $d(x, 0) = 0$. Each TRIX node fires when receiving the second signal from its predecessors: Define $t_c := d(x - c, y) + w_c^{x-c,y}$ as the time at which node $(x, y+1)$ receives each clock pulse. Then node $(x, y+1)$ fires a clock pulse at the median time $t := \mathrm{median}\{t_{-1}, t_0, t_{+1}\}$.

We concentrate on two important metrics to analyze this system: absolute delay and relative skew. Our main interests are the random variables $d(H) := d(0, H)$, i.e. the total delay at the top, and $s(H) := d(1, H) - d(0, H)$, i.e. the relative skew between neighboring nodes.

3 Delay is Tightly Concentrated

We examine $d(2000)$, the delay at layer 2000. The estimated probability mass function of $d(2000)$ looks like a binomial distribution. The empiric standard deviation is only 2.741, i.e. less than three delay uncertainties. The full version [7] explains the statistic methods in detail, contains more figures, and proves all following lemmas. The peak of the probability mass function falls in the middle of the support $[0, H]$:

Lemma 1. $\mathbb{E}[d(H)] = H/2$.

The behavior at $H = 2000$ is similar for other heights and changes slowly with increasing H.

The empiric standard deviation for various values of H can be seen in the data plotted in Figure 3 as a log-log plot. This suggests a polynomial relationship between standard deviation σ and grid height H. The slope of the line is close to 1/4, which suggests $\sigma \sim H^{\beta}$ with $\beta \approx 1/4$. This is a quadratic improvement over standard concentration bounds, which would predict $\beta \approx 1/2$.

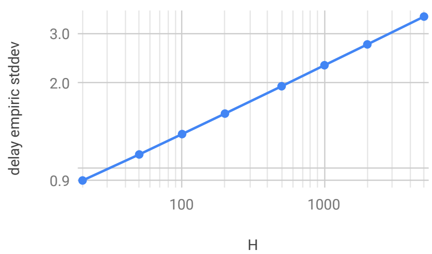

Fig. 3: Log-log plot of the empiric standard deviation of $d(H)$.

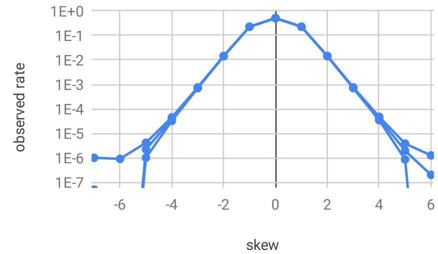

Fig. 4: Estimated probability mass function for $s(2000)$ with logarithmic Y-axis. Error bounds are only visible at fringes.

4 Skew is Tightly Concentrated

We examine $s(2000)$, the skew at layer 2000 between neighboring nodes. As expected, we see a high concentration around 0 in Figure 4, with roughly half of the probability mass at 0.

Observe that the skew does not follow a normal distribution at all: The probability mass seems to drop off exponentially like $e^{-\lambda|x|}$ for $\lambda \approx 2.9$ (where x is the skew), and not quadratic-exponentially like $e^{-x^2/(2\sigma^2)}$, as it would happen in the normal distribution. The probability mass for 0 is a notable exception, not matching this behavior.

We observe that the skew seems to be symmetric with mean 0.

Corollary 1. *$s(H)$ is symmetric with $\mathbb{E}[s(H)] = 0$.*

Furthermore, the worst-case skew on layer H is indeed H, c.f. Figure 2.

Lemma 2. *There is an assignment for all c_w such that $s(H) = H$.*

We conjecture that the probability mass of high-skew assignments is very low.

Again, the behavior at $H = 2000$ is similar for other heights and changes extremely slowly with increasing H. Figure 5 shows that the skew remains small even for large values of H. Note that the X-axis is doubly logarithmic. This suggests that the standard-deviation of $s(H)$ grows strongly sub-logarithmically, possibly even converges to a finite value. In fact, the plot suggests that $s(H) \in O(\log \log H)$.

Note that if we pretended that adjacent nodes exhibit independent delays, the skew would have the same concentration as the delay. In contrast, we see that adjacent nodes are tightly synchronized; this is ideal for clock propagation.

So far, we have limited our attention to the skew between neighboring nodes. In the other extreme end, at horizontal distances $\delta \geq 2H$, node delays are independent, as they do not share any wires on any path to any clock generator. In Figure 6, we see that the skew grows steadily with increasing $\delta \ll 2H$. The plot suggests that the standard deviation increases roughly proportional to δ^γ for $\gamma \approx 1/3$. This is noticeably less steep than the naive guess $\gamma \approx 1/2$ for small

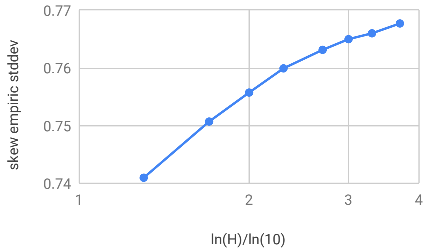

Fig. 5: Empiric standard deviation of $s(H)$ as a function of H, as a loglog-lin plot.

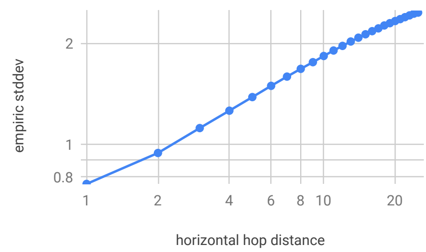

Fig. 6: Empiric standard deviation of $d(\delta, 500) - d(0, 500)$ as a function of horizontal distance δ in a log-log plot.

δ. It is not surprising that the slope falls off towards larger values, as it must become constant for $\delta \geq 2H = 1000$.

References

1. Chipana, R., Kastensmidt, F.L.: SET Susceptibility Analysis of Clock Tree and Clock Mesh Topologies. In: ISVLSI. pp. 559–564 (2014)
2. Chipana, R., Kastensmidt, F.L., Tonfat, J., Reis, R., Guthaus, M.: SET susceptibility analysis in buffered tree clock distribution networks. In: RADECS. pp. 256–261 (2011)
3. Dolev, D., Függer, M., Lenzen, C., Perner, M., Schmid, U.: HEX: Scaling honeycombs is easier than scaling clock trees. J. Comput. Syst. Sci. **82**(5), 929–956 (2016)
4. Dolev, D., Halpern, J.Y., Strong, H.R.: On the Possibility and Impossibility of Achieving Clock Synchronization. J. Comput. Syst. Sci. **32**(2), 230–250 (1986)
5. Dolev, D., Függer, M., Lenzen, C., Schmid, U.: Fault-tolerant Algorithms for Tick-generation in Asynchronous Logic. J. ACM **61**(5), 30:1–30:74 (2014)
6. Lenzen, C., Rybicki, J.: Self-Stabilising Byzantine Clock Synchronisation Is Almost as Easy as Consensus. J. ACM **66**(5), 32:1–32:56 (2019)
7. Lenzen, C., Wiederhake, B.: TRIX: Low-skew pulse propagation for fault-tolerant hardware (2020), https://arxiv.org/abs/2010.01415
8. Malherbe, V., Gasiot, G., Clerc, S., Abouzeid, F., Autran, J.L., Roche, P.: Investigating the single-event-transient sensitivity of 65 nm clock trees with heavy ion irradiation and Monte-Carlo simulation. In: IRPS. pp. SE–3–1–SE–3–5 (2016)
9. Srikanth, T.K., Toueg, S.: Optimal Clock Synchronization. J. ACM **34**(3), 626–645 (1987)
10. Wang, H.B., Mahatme, N., Chen, L., Newton, M., Li, Y.Q., Liu, R., Chen, M., Bhuva, B.L., Lilja, K., Wen, S.J., Wong, R., Fung, R., Baeg, S.: Single-Event Transient Sensitivity Evaluation of Clock Networks at 28-nm CMOS Technology. IEEE Trans. Nucl. Sci. **63**(1), 385–391 (2016)
11. Welch, J.L., Lynch, N.A.: A New Fault-Tolerant Algorithm for Clock Synchronization. Information and Computation **77**(1), 1–36 (1988)
12. Xanthopoulos, T. (ed.): Clocking in Modern VLSI Systems. Springer (2009)