Characteristics of Reversible Circuits for Error Detection

Lukas Burgholzer*   Robert Wille†‡   Richard Kueng*

*Institute for Integrated Circuits, Johannes Kepler University Linz, Austria
†Chair for Design Automation, Technical University of Munich, Germany
‡Software Competence Center Hagenberg GmbH (SCCH), Austria

lukas.burgholzer@jku.at   robert.wille@tum.de   richard.kueng@jku.at

https://www.cda.cit.tum.de/research/quantum/

Abstract

In this work, we consider error detection via simulation for reversible circuit architectures. We rigorously prove that reversibility augments the performance of this simple error detection protocol to a considerable degree. A single randomly generated input is guaranteed to unveil a single reversible error with a probability that only depends on the size of the error, not the size of the circuit itself. Empirical studies confirm that this behavior typically extends to multiple errors as well. In conclusion, reversible circuits offer characteristics that reduce masking effects – a desirable feature that is in stark contrast to irreversible circuit architectures.

Index Terms

emerging technologies, reversible logic, error detection, simulation, quantum computing

I. INTRODUCTION

The detection of errors is a fundamental problem in electrical engineering and computer science. Given a circuit C_1 with n inputs and m outputs (the *Golden Specification*), the task is to decide whether a given circuit realization C_2 (the *Design Under Verification*) describes the same functionality on the logical level.

Many approaches exist that address this important and challenging problem. In this work, we focus on error detection protocols that only require simulation runs of the two circuits—as opposed to formal verification techniques which explicitly utilize structural knowledge about both circuits [1]–[6]. This is a severe restriction, but simulations alone are—in principle—sufficient to solve this task. If the two circuits are equivalent, they have the same input-output behavior. Conversely, suppose that they are functionally distinct. Then, there exists at least one input string for which the two circuits produce *distinct* outputs. In formulas:

$$\exists \vec{x} \in \{0,1\}^n \quad \text{such that} \quad C_1(\vec{x}) \neq C_2(\vec{x}). \tag{1}$$

Such an input successfully detects the discrepancy between C_1 and C_2 and serves as a counterexample for the equivalence of the circuits.

The problem, however, is how to find counterexamples (1). If we only allow simulations of both circuits, i.e., we consider them as black boxes, we do not have actionable advice on how to choose promising input strings and we may as well generate inputs uniformly at random: $\vec{x} \sim \text{Unif}(\{0,1\}^n)$, i.e., we flip an unbiased coin for each input value ($\vec{x} = (x_n, \ldots, x_1)$, where $x_n, \ldots, x_1 \sim x$ and $\Pr[x = 0] = \Pr[x = 1] = 1/2$). Subsequently, we simulate both circuits with this input and check whether they produce the same output: $C_1(\vec{x}) \stackrel{?}{=} C_2(\vec{x})$. If the outputs are distinct, we have found a counterexample. The circuits cannot be equivalent. But if the outputs are the same, the test is inconclusive. In this case, we must repeat it with new (randomly generated) inputs until we either find a counterexample (non-equivalence) or have exhausted all 2^n possible inputs (equivalence). The latter, unfortunately, can be a very real possibility. The two circuits C_1 and C_2 may differ on a single input only and it is extremely unlikely to quickly find this input by (random) chance.

To make matters worse, classical circuits can mask even "small" errors very effectively. For $n = 8$, this is illustrated in Fig. 1. A cascade of logical AND gates, realizing the functionality $y = x_n \cdot \ldots \cdot x_1$ (ideal circuit C_1), is affected by a single bit-flip error (erroneous implementation C_2) in the second layer. It is easy to check that only 4 out of all $2^8 = 256$ input strings can detect this discrepancy.

Masking is a serious issue for error detection using simulation techniques. No malicious intent is required to fool randomly generated inputs. The circuit may do it all by itself. Needless to say, this issue has been well-known for decades. Error detection based on random inputs (alone) often pales in comparison to other more sophisticated techniques. Today's state of the art is governed by constrained-based stimuli generation techniques [7]–[11], fuzzing [12], etc. But on the positive side, error detection using randomly-chosen inputs is based on minimal assumptions, namely the possibility to simulate two circuits as black boxes. Moreover, it is intuitive and individual simulation runs are easy and fast to execute.
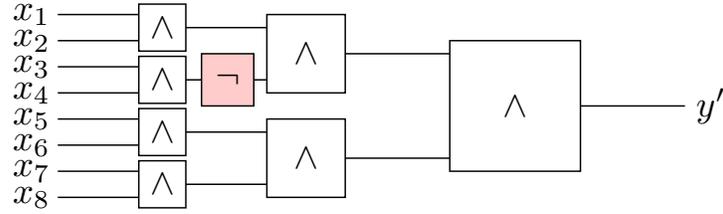
Fig. 1. *Error detection in classical circuits is hard:* Suppose that a cascade of logical AND gates, realizing the Boolean function $y = x_8 \cdot \ldots \cdot x_1$, is affected by a single bit-flip error (red) in the second layer. Only 4 out of the $2^8 = 256$ possible input strings can detect this error.
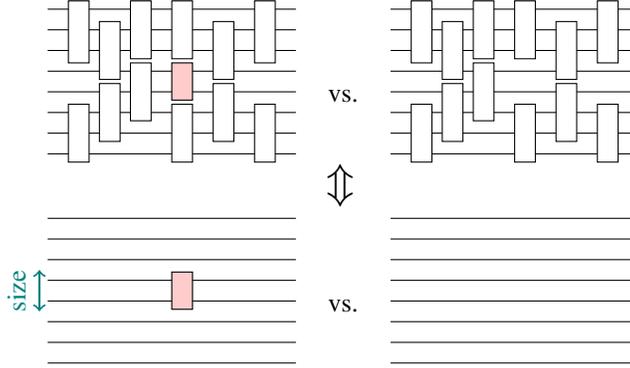

Fig. 2. *Illustration of main rigorous contributions:* Simulations with uniformly random inputs completely expose any single reversible error in a given reversible circuit. The two scenarios are *exactly* equivalent ("no masking"). In the lower scenario, the probability of correct distinction is governed by the size k of the error, not the total number of lines.

II. Summary of results: Error detection in reversible circuits

We have seen that, in general, simulation with (uniformly) random inputs is not a viable strategy for detecting errors in classical circuits. Already a single "small" error can be exceedingly difficult to detect (masking). Perhaps surprisingly, this dark picture lightens up considerably if we consider *reversible* implementations of logical functionalities. As the name suggests, reversible circuits are circuits whose action can be undone by running the circuit backwards. More formally, n-bit reversible circuits implement permutations on the set of all 2^n bit strings. This, in particular, implies that the number of input and output bits must be the same ($n = m$). Despite these restrictions, reversible circuits are universal, i.e., *any* logical function on n bits can be implemented by a reversible circuit [13] and efficient mapping techniques are readily available [14]–[16] (this implementation may require strictly more than n bits, though). Negation (NOT), exclusive or (CNOT) and the Toffoli gate (CCNOT) are examples of simple reversible functionalities. Viewed as a logic gate, CCNOT is also universal. Every reversible circuit can be constructed from Toffoli gates alone [13].

To summarize, reversible circuits bear strong similarities with classical (irreversible) circuits, but there are some notable additional characteristics. Chief among them is reversibility itself which implies that information cannot easily escape. Here, we show that this has profound implications for error detection with random inputs. More precisely,

(i) reversible circuits can never mask single reversible errors (rigorous result, see Proposition 1)
(ii) the probability of detecting a single reversible error only depends on its size, i.e., on the number of bits it affects, not the total number of bits (unsurprising rigorous result, see Lemma 2)
(iii) multiple reversible errors are typically even easier to detect (empirical studies, see Fig. 3 and discussions in Section IV)

The first two insights are mathematical statements that address single errors only. They readily follow from reversibility and fundamental properties of uniformly random input strings. We refer to Section III for details and Fig. 2 for illustrative caricatures. When combined, they imply the following confidence bound for detecting single errors with random inputs.

Theorem 1. *Suppose that a general reversible circuit is affected by a single reversible error of size k and fix $\delta \in (0, 1)$ (confidence). Then, at most $\lceil \log(1/\delta)2^{k-1} \rceil$ randomly selected inputs suffice to witness this error with probability (at least) $1 - \delta$.*

For $k = 1$—a single bit-flip error (NOT) *anywhere* within the circuit—this statement can be further improved (see Theorem 2) and actually becomes deterministic: already a single (random) input is guaranteed to detect this error with certainty. We emphasize that this statement is true irrespective of the number of lines and the circuit's size. It is simply impossible to hide a single bit-flip inside a reversible circuit. Such a behavior is strikingly different from irreversible circuit architectures. There it can routinely happen that order 2^n random inputs are necessary to detect even a single bit-flip error, see e.g. Fig. 1.
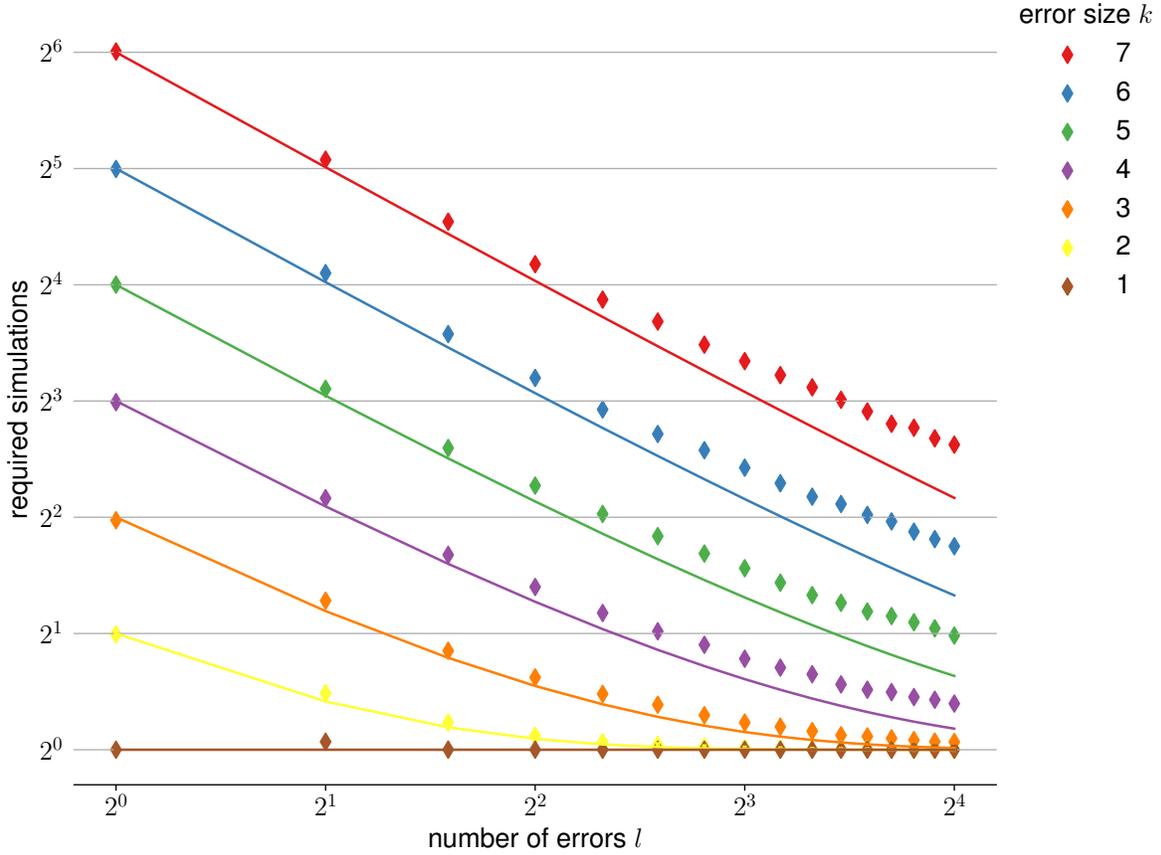
Fig. 3. *Typical accumulation effects for multiple errors (log-log plot):* number l of randomly injected reversible errors (x-axis) vs. average number of random inputs required to detect erroneous behavior (y-axis) in a generic $n = 20$-bit reversible circuit with 4000 gates. Different colors denote worst-case errors of increasing size k. Solid lines track the theoretical best-case behavior (independent errors, see Eq. (5) below). For small l, the plot highlights an excellent agreement between typical (diamonds) and best-case (solid lines) behavior.

The multiple-error case is much more intricate, because error locations and circuit structure start to matter. This leads to drastically different behaviors of best case (independent errors) and worst case (severe masking) behavior. To better understand the typical behavior of multiple errors, we resort to numerical simulations. These indicate a (close-to) best-case behavior: the probability of failing to detect a total of l reversible errors is exponentially suppressed in l, see Fig. 3. Additional simulation results and details are provided in Section IV.

Note that a similar line of thought has recently been presented for the domain of quantum computing (which bears many similarities to reversible circuits). More precisely, a verification scheme heavily based on simulation has been proposed in [17] and refined in [18]. A similar theoretical result has been presented in [19].

III. RIGOROUS THEORY FOR SINGLE ERRORS

A. Reversible circuits and error model

We will work in the reversible circuit model for n input bits (and n output bits). A high-level of mathematical abstraction already suffices to deduce powerful consequences. An n-bit reversible circuit implements a permutation $R : \{0,1\}^n \to \{0,1\}^n$ of all 2^n bit strings. Reversing the circuit, that is running it backwards, produces the unique permutation $R^T : \{0,1\}^n \to \{0,1\}^n$ that undoes the original circuit: $R^T \circ R = R \circ R^T = \mathrm{id}$, where $\mathrm{id}(\vec{x}) = \vec{x}$ for all $\vec{x} \in \{0,1\}^n$ is the identity permutation ("do nothing"). This defining feature suffices to deduce three elementary properties that will form the basis of our proof strategy.

Lemma 1 (Characteristics of reversible circuits). *Consider reversible circuits $R_1, R_2, R_3 : \{0,1\}^n \to \{0,1\}^n$ and an n-bit string $\vec{x} \in \{0,1\}^n$. Then,*

(i) output equivalence is unaffected by composition:
$R_1(\vec{x}) = R_2(\vec{x}) \Leftrightarrow (R_3 \circ R_1)(\vec{x}) = (R_3 \circ R_2)(\vec{x})$

(ii) invariance of the uniform distribution:
$\vec{x} \sim \mathrm{UNIF}(\{0,1\}^n)$ *implies* $R_1(\vec{x}) \sim \mathrm{UNIF}(\{0,1\}^n)$

(iii) non-trivial action: suppose $R_1 \neq \mathrm{id}$. Then, there are at least two bit strings such that $R_1(\vec{x}) \neq \vec{x}$.
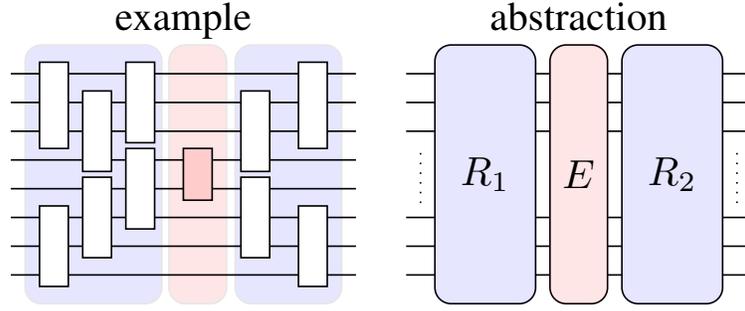
Fig. 4. *(Single) error model and compatible circuit decomposition:* An ideal reversible circuit (blue) is corrupted by a single reversible error (red). The error location begets a decomposition of ideal and corrupted circuit into matching constituents: $R = R_2 \circ R_1$ (ideal) and $\tilde{R} = R_2 \circ E \circ R_1$ (corrupted).

Proof. All proofs utilize the fact that reversible circuits act like permutations on the set of all 2^n bit strings.

(i) Permutations are *invertible* transformations. As such, they preserve equivalence: $y = y'$ if and only if $R(y) = R(y')$ for any reversible circuit R. The claim follows from setting $y = R_1(\vec{x})$, $y' = R_2(\vec{x})$ and $R = R_3$.

(ii) The uniform distribution over n-bit strings assigns the same weight to each of the 2^n bit strings. Permuting the bit strings cannot affect the weights and, by extension, the uniform distribution itself.

(iii) The number of invariant bit strings $(\vec{x} \in \{0,1\}^n : R_1(\vec{x}) = \vec{x})$ is equal to the number of fix points of the underlying permutation. A non-trivial permutation of 2^n elements can have most $2^n - 2$ fix points (transposition). □

Different reversible circuits of compatible bit-size n can be combined to yield another (larger) circuit: $(R_2 \circ R_1)(\vec{x}) = R_2(R_1(\vec{x}))$ for input $\vec{x} \in \{0,1\}^n$ ("composition"). The reverse direction is also possible ("decomposition") and, arguably, more interesting. Circuit diagrams provide a well-established tool that does precisely that. They decompose a possibly complicated circuit into a structured sequence of simpler building blocks. We use circuit decomposition on a rather high level to reason about *single* reversible errors in reversible circuits. Suppose that an n-bit reversible circuit R is affected by a reversible error E that produces a functionally different circuit $\tilde{R}$. Then, the location of this error within the circuit suggests a compatible decomposition into three parts:

(i) $R_1 : \{0,1\}^n \to \{0,1\}^n$ describes the original functionality up to the location where the error occurs ("past"),

(ii) $E : \{0,1\}^n \to \{0,1\}^n$ captures the error as an additional circuit layer on all n bits ("present"),

(iii) $R_2 : \{0,1\}^n \to \{0,1\}^n$ describes the original functionality from the error location onwards ("future").

In summary,

$$\tilde{R} = R_2 \circ E \circ R_1, \quad \text{while} \quad R = R_2 \circ R_1, \tag{2}$$

and we refer to Fig. 4 for a visual illustration.

B. No masking for random inputs

We now have all building blocks in place to present and derive the main conceptual result of this work. It addresses the probability of detecting single reversible errors in arbitrary reversible circuits (2) based on a single random input $\vec{x} \sim \mathrm{Unif}(\{0,1\}^n)$.

Proposition 1 (no masking). *Fix $R = R_2 \circ R_1$ (ideal circuit) and $\tilde{R} = R_2 \circ E \circ R_1$ (single, reversible error). Then, the probability of detecting this discrepancy with a random input $\vec{x} \sim \mathrm{Unif}(\{0,1\}^n)$ only depends on the error E, not the actual circuit. More precisely,*

$$\Pr\left[\tilde{R}(\vec{x}) \neq R(\vec{x})\right] = \Pr\left[E(\vec{x}) \neq \vec{x}\right],$$

where the probability is taken with respect to the uniform distribution over all 2^n possible input strings.

Proof. This statement is an immediate consequence of two elementary characteristics of reversible circuit architectures. Apply Lemma 1 (i) to remove the effect of R_2,

$$\Pr\left[\tilde{R}(\vec{x}) = R(\vec{x})\right] = [R_2 \circ E \circ R_1(\vec{x}) = R_2 \circ R_1(\vec{x})]$$
$$= \Pr\left[E(R_1(\vec{x})) = R_1(\vec{x})\right],$$

and note that, according to Lemma 1 (ii), $\vec{x} \sim \mathrm{Unif}(\{0,1\}^n)$ implies $R_1(\vec{x}) \sim \mathrm{Unif}(\{0,1\}^n)$. □

Although simple to prove, Proposition 1 pinpoints remarkable differences between reversible and irreversible circuits. As illustrated in Fig. 2, the former cannot hide errors from randomly sampled inputs ("no masking").

We emphasize that a uniformly random selection of input strings is crucial to arrive at such a powerful conclusion. Reversibility alone is enough to ignore the final portion of the circuit R_2 (after the error has occurred). Reversible circuits always map (non-)equal bit strings to (non-)equal bit strings. In contrast, the first portion of the circuit R_1 (before the error has occurred) can affect concrete inputs $\vec{x} \in \{0,1\}^n$. But if $\vec{x}$ is sampled randomly, then $R_1(\vec{x})$ will be a different, but still random, bit string. The uniform distribution is special in the sense that it is invariant under reversible transformations. The circuit R_1 may affect every concrete input, but it does not affect the underlying distribution.

C. Only error size matters

We have seen that uniformly random inputs can uncover single reversible errors in a general reversible circuit. According to Proposition 1, the probability of witnessing a discrepancy only depends on the error, not the underlying circuit structure.

We say that an error $E : \{0,1\}^n \to \{0,1\}^n$ has *size* k if it only affects k bits in a nontrivial fashion. The remaining $n - k$ bits are not touched at all. We refer to Fig. 2 for a visual illustration of this summary parameter. Intuitively, we would expect that "large" errors are easier to detect than "small" ones and that the number of lines n plays an active role. However, the following simple statement shows that the probability of detecting an error in the worst case is exponentially suppressed with respect to the error size k, but is independent of the actual number of bits n.

Lemma 2 (only error size matters). *Suppose that $E : \{0,1\}^n \to \{0,1\}^n$ is a reversible error that only affects k bits in a non-trivial fashion and $\vec{x} \sim \mathrm{Unif}(\{0,1\}^n)$ is sampled from the uniform distribution. Then,*

$$\Pr\left[E(\vec{x}) \neq \vec{x}\right] \geq 2^{-(k-1)}.$$

Proof. Suppose, without loss of generality, that the error E only affects the least-significant k bits, i.e., $E(\vec{x}) = E(x_n, \ldots, x_1) = (x_n, \ldots, x_{k+1}, y_k, \ldots, y_1)$, where $(y_k, \ldots, y_1) = \tilde{E}(x_k, \ldots, x_1)$. Since E is reversible, its restriction $\tilde{E} : \{0,1\}^k \to \{0,1\}^k$ to the k relevant bits must also be reversible. Moreover, $\tilde{E} \neq \mathrm{id}$, because E is non-trivial. Lemma 1 (iii) then implies that there must be at least 2 bit strings of size k that are affected by $\tilde{E}$. Finally, we use the fact that $\vec{x} = (x_n, \ldots, x_1) \sim \mathrm{Unif}(\{0,1\}^n)$ implies that the least-significant k bits are also distributed uniformly: $(x_k, \ldots, x_1) \sim \mathrm{Unif}(\{0,1\}^k)$. Therefore,

$$\Pr\left[E(\vec{x}) \neq \vec{x}\right] = \Pr\left[\tilde{E}(x_k, \ldots, x_1) \neq (x_k, \ldots, x_1)\right] \geq \frac{2}{2^k}.$$

$\square$

This probability bound is actually sharp. Worst-case reversible errors of size k permute exactly 2 out of the 2^k possible k-bit inputs on which they act. Concrete examples of such a behavior are NOT ($k = 1$), CNOT ($k = 2$), CCNOT ($k = 3$) and, more generally, a $(k-1)$-fold controlled NOT gate on k bits (general k). The numerical simulations shown in Fig. 3 are based on injecting such worst-case errors at random circuit locations.

D. General confidence bound for detecting single reversible errors

We now have all necessary ingredients to establish a rigorous performance guarantee for reversible error detection with (uniformly) random inputs. The following statement bounds the number of uniformly random inputs that may be required to detect a single reversible error of size k.

Theorem 2. *Fix $R = R_2 \circ R_1$ (ideal circuit), $\tilde{R} = R_2 \circ E \circ R_1$ (single, reversible error) and E has size k. Suppose that $\vec{x}_1, \ldots, \vec{x}_N$ are N (independent) uniformly random inputs. Then,*

$$\Pr\left[\bigwedge_{1 \leq i \leq N} \left\{\tilde{R}(\vec{x}_i) = R(\vec{x}_i)\right\}\right] \leq \exp\left(-N/2^{k-1}\right)$$

In words, the probability of failing to detect a single error is exponentially suppressed in the number N of random test inputs.

Theorem 1 above is a streamlined consequence of this observation: setting $N = \lceil \log(1/\delta) 2^{k-1} \rceil$ provides a concrete number of repetitions that ensures that we detect the discrepancy with probability (at least) $1 - \delta$.

Proof of Theorem 2. For $N = 1$ (one random input), the claim readily follows from combining Proposition 1 and Lemma 2 (more precisely, their contrapositions):

$$\Pr\left[\tilde{R}(\vec{x}_1) = R(\vec{x}_1)\right] = \Pr\left[E(\vec{x}_1) = \vec{x}_1\right] \leq 1 - 2^{-(k-1)}.$$

This bound readily extends to the general N-case by using the assumption that the individual input strings $\vec{x}_1, \ldots, \vec{x}_N$ are all sampled independently. Joint probabilities of independent events factorize and we conclude

$$\Pr\left[\bigwedge_{1 \leq i \leq N} \left\{\tilde{R}(\vec{x}_i) = R(\vec{x}_i)\right\}\right] = \prod_{i=1}^{N} \Pr\left[\tilde{R}(\vec{x}_i) = R(\vec{x}_i)\right]$$

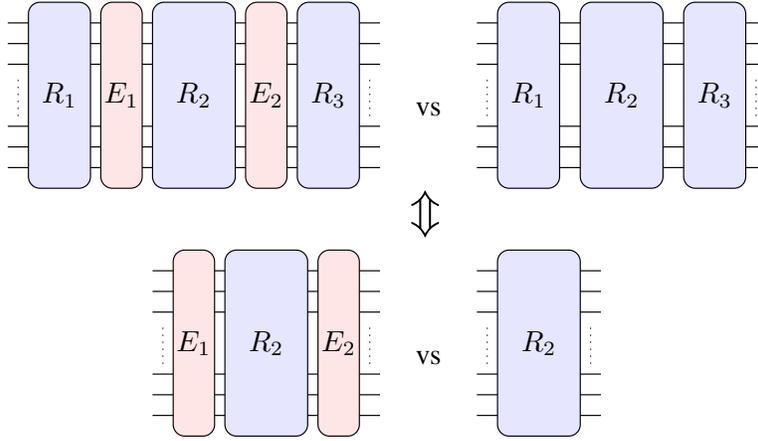$$\leq \left(1 - 2^{-(k-1)}\right)^N. \tag{3}$$

Fig. 5. *Partial simplification for multiple errors:* Simulation with uniformly random inputs exposes multiple errors only partially. Everything before the first error (R_1) and after the last error (R_3) can be safely ignored, but the part in between (R_2) does matter. Different circuit structures can lead to strikingly different error detection probabilities.
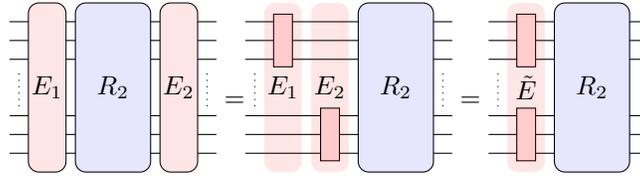

Fig. 6. *Best-case scenario for two errors:* One of the errors, say E_2, commutes with the relevant circuit part R_2. Reordering allows us to treat the two errors as a single effective error $\tilde{E} = E_1 \circ E_2$. In addition, E_1 and E_2 affect disjoint bit collections (independence) and $\tilde{E}$ factorizes nicely into two disjoint components: $\Pr\big[\tilde{R}(\vec{x}) \neq R(\vec{x})\big] = \Pr\big[\tilde{E}(\vec{x}) \neq \vec{x}\big] \geq 1 - (1 - 2^{-(k-1)})^2$ (quadratic improvement).

Apply $1 + x < \exp(x)$ for all $x \in \mathbb{R}$ (convexity of the exponential function) with $x = -2^{-(k-1)}$ to complete the argument. $\qquad\square$

The bound provided in Theorem 2 is simple, but not sharp (the inequality $1 + x \leq \exp(x)$ is never tight). As such, it always under-estimates the actual confidence level. This discrepancy is most pronounced for small error sizes k. The extreme case is a single NOT error ($k = 1$). For $k = 1$, the bound in Eq. (3) becomes (exactly) zero. By contraposition, *every possible input bit string is guaranteed to detect a single bit-flip error that is hidden anywhere within the circuit.*

IV. EMPIRICAL ANALYSIS FOR MULTIPLE ERRORS

In the previous section, we have established strong theoretical support for detecting single reversible errors. At its heart has been the decomposition $\tilde{R} = R_2 \circ E \circ R_1$ illustrated in Fig. 4. Reversibility and uniformly random inputs have subsequently allowed us to discuss away the circuit portions R_2 and R_1 completely. In turn, we were able to focus exclusively on the error itself.

For more than one error, this is in general not an option anymore. While we can safely ignore circuit contributions before the first and after the last error, the circuit in between cannot be ignored, see Fig. 5. The relation between errors and intermediate circuit parts governs how likely it is to witness the overall error.

In this section, we analyze error accumulation effects in generic reversible circuits. To obtain guiding intuition, we will first isolate and discuss the two extreme cases. Independent errors (best case, see Section IV-A) and maximal masking (worst case, see Section IV-B) turn out to behave in a radically different fashion. Subsequent numerical studies demonstrate that typical error accumulation effects closely follow the best-case trajectory: Multiple errors are typically *much* easier to detect than a single error.

A. Best-case behavior: Commuting and independent errors

Let us first discuss $l = 2$ reversible errors of size k. An extension to multiple errors ($l \geq 3$) and different sizes will be straightforward. Fig. 5 provides valuable guidance for potential best-case behavior. Suppose that one of the errors, say E_2, can be pulled through the central circuit part R_2 without affecting it: $E_2 \circ R_2 = R_2 \circ E_2$. If circuit and error commute in such a fashion, we can group both errors into a single layer and have effectively reduced the problem to the single-error case which we already understand:

$$\tilde{R} = R_3 \circ E_2 \circ R_2 \circ E_1 \circ R_1 = (R_3 \circ R_2) \circ (E_2 \circ E_1) \circ R_1.$$
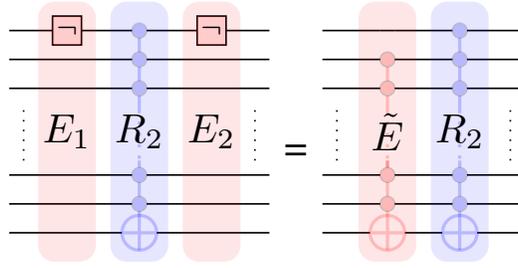
Fig. 7. *Worst-case scenario for two errors:* Two bit-flip errors ($k = 1$) affect one control line of a $(n-1)$-fold controlled NOT-gate. These errors do not commute with the relevant circuit part R_2. Quite the opposite: two errors with size $k = 1$ produce an effective error $\tilde{E}$ of size $k = (n-1)$. To make matters even worse, such a $(n-2)$-fold controlled NOT error is extremely difficult to detect: $\Pr\big[\tilde{R}(\vec{x}) \neq R(\vec{x})\big] = \Pr\big[\tilde{E}(\vec{x}) \neq \vec{x}\big] = 4/2^n$ (masking).

The only remaining question is: what is the probability of failing to detect the cumulative error $E_2 \circ E_1$ with a single random input? This failure probability is smallest if the two errors are *independent* in the sense that they act on disjoint sets of k bits each. A uniformly random input $\vec{x} \in \mathrm{Unif}(\{0,1\}^n)$ then ensures that the failure probability factorizes:

$$\Pr\left[(E_2 \circ E_1)(\vec{x}) = \vec{x}\right] = \prod_{i=1}^{2} \Pr\left[E_i(\vec{x}) = \vec{x}\right] \leq \left(1 - 2^{-(k-1)}\right)^2.$$

This argument readily extends to multiple errors ($l \geq 3$). Taking the complement ensures

$$\Pr\left[\tilde{R}(\vec{x}) \neq R(\vec{x})\right] = 1 - \Pr\left[E_l \circ \cdots \circ E_1(\vec{x}) = \vec{x}\right]$$
$$\geq 1 - \left(1 - 2^{-(k-1)}\right)^l, \tag{4}$$

provided that all l errors commute with the circuit (first equality) and act on different subsets of k bits each (second inequality). Rel. (4) highlights that the probability of (best case) error detection increases substantially with the number of errors l. Intuitively, this makes sense: more errors should be easier to detect. This insight has implications for the number N of random inputs that are required to detect l best-case errors of size k each. To pinpoint them, it is instructive to view a single simulation run as a biased coin toss: we detect a discrepancy with probability $p = \Pr\big[\tilde{R}_2(\vec{x}) \neq R_2(\vec{x})\big]$ ("heads") and fail to detect it with probability $1 - p = \Pr\big[\tilde{R}_2(\vec{x}) = R_2(\vec{x})\big]$ ("tails"). When attempting to detect a discrepancy, we input new randomly generated inputs until we find a mismatch. This is equivalent to tossing the biased coin until "heads" appears. The expected number of required coin tosses to achieve this goal is $1/p$ (geometric distribution). Together with Rel. (4), this analogy allows us to conclude that we expect to require

$$N_{\mathrm{expect}}^{(\downarrow)} \leq \frac{1}{1 - \left(1 - 2^{-(k-1)}\right)^l} \qquad \text{(best case)} \tag{5}$$

random inputs to detect l commuting and independent errors of size k each. This bound is sharp. It holds with equality if each of the l errors is a worst-case error of size k, e.g. a $(k-1)$-fold controlled NOT gate.

We conclude this section with a simplified interpretation of Rel (5). For small l (in comparison to $2^{(k-1)}$), the claim is comparable to $N_{\mathrm{expect}}^{(\downarrow)} \approx 2^{k-1}/l$, which can also be observed in Fig. 3: the slopes of the solid lines match this estimate rather well whenever the number of errors l is small compared to $2^{(k-1)}$. Under best-case assumptions, detecting l size k-errors is l-times easier than detecting a single error of the same size.

B. Worst-case: anti-commuting errors and masking

We expect that worst case error accumulation should occur when errors and relevant circuit portion do not commute at all ("anti-commutation"). If this is the case, the probability of detecting errors can become exponentially small in the total number of bits. We illustrate this by means of an example that is illustrated in Fig. 7: E_1 and E_2 are bit-flip errors ($k = 1$) that affect the first bit while $R_2 : \{0,1\}^n \to \{0,1\}^n$ is a $(n-1)$-fold controlled NOT-gate. It is easy to check that

$$E_2 \circ R_2 \circ E_1 = \tilde{E} \circ R_2,$$

where $\tilde{E}$ is a $(n-2)$-fold controlled NOT gate that acts on all bits, except the very first one ($k = n-1$). This is a single worst-case error of almost maximal size. Proposition 1 and Lemma 2 assert

$$\Pr\big[\tilde{R}(\vec{x}) \neq R(\vec{x})\big] = \Pr\left[\tilde{E}(\vec{x}) \neq \vec{x}\right] = \frac{4}{2^n}.$$

This success probability is exponentially small in the total number of bits and we expect to require a total of

$$N_{\mathrm{expect}}^{(\uparrow)} \geq 2^{(n-2)} \qquad \text{(worst case)} \tag{6}$$

random inputs in order to detect the discrepancy. Even worse error accumulation effects can occur for more errors ($l \geq 3$) and/or larger error sizes ($k \geq 2$). But already Rel. (6) is almost as bad as it can be. It is only a factor of two away from 2^{n-1}—the absolute worst case for distinguishing *any* pair of reversible circuits, see Lemma 1 (iii).

C. Empirical studies

The multiple-error case is intricate by comparison, because the interplay between error (locations) and underlying circuit geometry starts to matter. We have seen that this leads to strikingly different best- (commuting errors, Sub. IV-A) and worst-case (anticommuting errors, Sub. IV-B) behavior. Concrete problem instances fall into the wide range between these extreme cases. In this section, we employ numerics to delineate *typical* behavior.

We study the effect of size-k reversible errors in reversible circuits with n lines. For a given number of lines n, we construct random reversible circuits with $g \approx \mathcal{O}(n^2)$ arbitrary multi-controlled NOT gates. When injecting errors of size k, we always consider $(k-1)$-fold controlled NOT gates which represent the worst case behavior, as discussed in Section III-C. Without loss, we assume that these errors are geometrically local, i.e., they only affect neighbouring lines. All experiments were repeated $10\,000$ times with different random seeds in order to ensure adequate statistical uniformity.
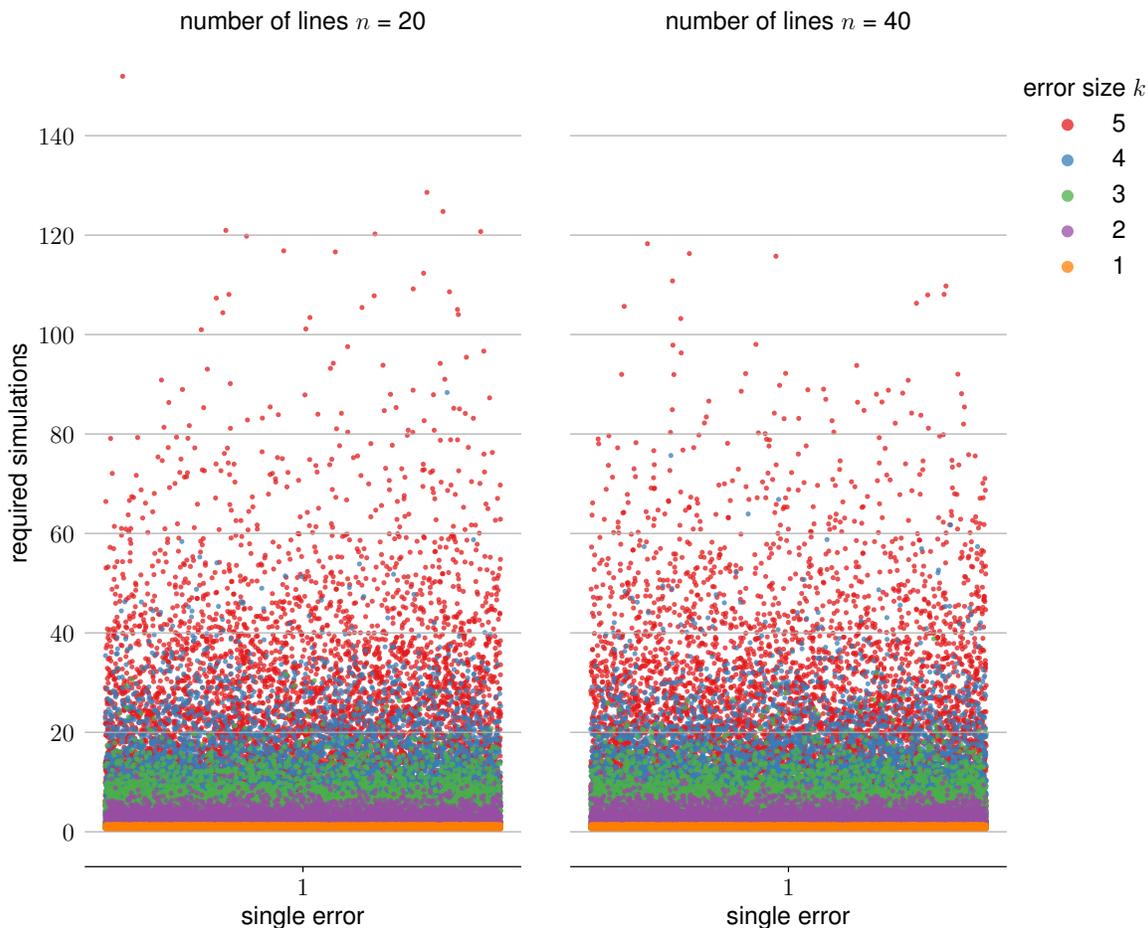

Fig. 8. *Confirmation of theoretical results*: Scatter-plot of required simulations (y-axis) for detecting a single reversible error of size k in a circuit with $n = 20$ (left plot) and $n = 40$ (right plot) lines. Different colors denote varying values of $k \in \{1, 2, 3, 4, 5\}$. This experimentally confirms that the distribution of simulations does not depend on the number of lines and that the number of required simulations grows exponentially with the size of the error.

First and foremost, we confirm interesting aspects of the theory developed in Sec. III. To this end, we considered the injection of a single size-k-error and count the required number of simulations for detecting this error. The results are depicted in Fig. 8. In contrast to classical intuition, the probability of detecting a single reversible error of size k is (1) completely independent of the circuit under consideration, and (2) diminishes exponentially in the error size k, i.e., the smaller the error, the greater its impact. This is in excellent agreement with Theorem 2. On average, the required simulations exactly follow the predicted 2^{k-1} trajectory with no apparent variation. Additionally, the distributions of results is the same when simulating the circuits $R = R_2 \circ R_1$ and $\tilde{R} = R_2 \circ E \circ R_1$ as compared to only simulating the error E itself.

The next set of numerical experiments pilots us in more interesting territory. Namely, the multiple-error case. We have already teased the results in the introduction and summarized them in Fig. 3. The averaged number of inputs highlights an

excellent agreement between the observed behavior and the best-case scenario discussed in Section IV-A. The deviation from this optimum for higher numbers of errors can be explained by accumulation affects of errors not acting independently (see Section IV-B).
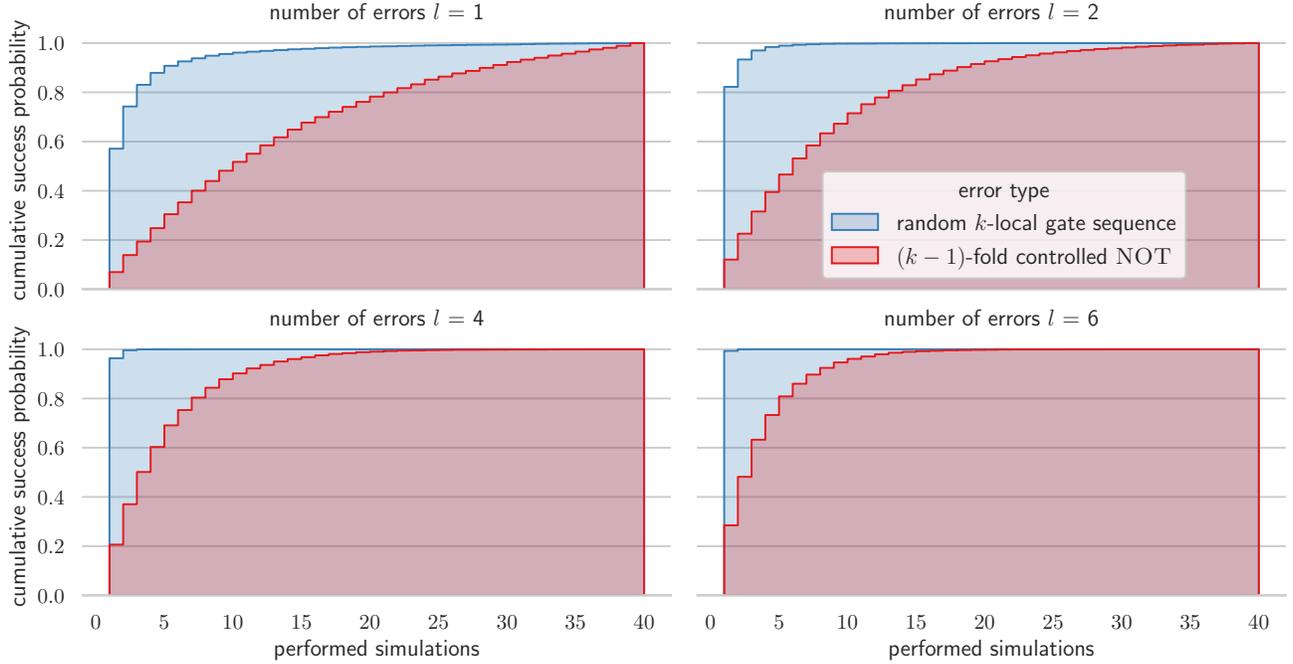

Fig. 9. *Comparison of worst-case and average-case errors* : performed simulations (x-axis) vs. cumulative distribution function (cdf) for detecting $l = 1, 2, 4, 6$ reversible errors of size $k = 5$ (y-axis). The red curve corresponds to injecting worst-case errors, while the blue curve delineates the cdf for detecting *randomly generated* errors of the same size. This goes to show, that average-case errors require far less simulations than worst-case ones.

Last but not least, we emphasize that—up to this point—theoretical and empirical results have been contingent on a worst-case assumption: each injected size-k error is a $(k-1)$-fold controlled NOT-gate. In a final series of evaluations, we analyzed the success probability after conducting a certain number of simulations when choosing errors *at random*. More precisely, each size-k error is a randomly selected gate sequence with the additional constraint that none of the k relevant lines remain unaffected (such a scenario would produce an error of size (at most) $(k-1)$). We expect that this error model captures typical behavior in a more accurate fashion. The results are shown in Fig. 9 and highlight a considerable discrepancy between random (blue) and worst-case (red) errors. This is not at all surprising. Random errors of size k tend to factorize into several independent contributions and the probabilities of detecting them with random inputs factorizes accordingly, see Sub. IV-A. Such factorizations lead to an increased error detection probability within (very) few simulation runs.

V. CONCLUSION

In this work, we have shown the impact of the reversible circuit paradigm on the probability of detecting errors in circuits. Our rigorous analysis shows, that, as opposed to classical/irreversible circuits, reversible circuits can never mask single errors and, that the probability of detecting a single reversible error only depends on the error's size and not at all on the surrounding circuit. Empirical evaluations have shown that, in case of multiple errors, the detection probability is very close to the theoretical best case. Finally, we have observed that, in case the assumption of worst-case errors is dropped, the probability of detecting these errors is increased even more.

ACKNOWLEDGMENTS

REFERENCES

[1] S. Disch and C. Scholl, "Combinational equivalence checking using incremental SAT solving, output ordering, and resets," in *Asia and South Pacific Design Automation Conf.*, 2007, pp. 938–943.

[2] J. Marques-Silva and T. Glass, "Combinational equivalence checking using satisfiability and recursive learning," in *Design, Automation and Test in Europe*, 1999.

[3] P. Molitor and J. Mohnke, *Equivalence checking of digital circuits: Fundamentals, principles, methods*. Springer, 2010.

[4] S. Jha, Y. Lu, M. Minea, and E. M. Clarke, "Equivalence checking using abstract BDDs," in *Int'l Conf. on Comp. Design*, 1997.

[5] E. M. Clarke, O. Grumberg, D. Kroening, D. A. Peled, and H. Veith, *Model Checking*. MIT Press, 2018.

[6] A. Biere, A. Cimatti, E. M. Clarke, and Y. Zhu, "Symbolic model checking without BDDs," in *Tools and Algorithms for the Construction and Analysis of Systems*, 1999, pp. 193–207.

[7] J. Yuan, C. Pixley, and A. Aziz, *Constraint-based verification*. Springer, 2006.

[8] A. Biere and W. Kunz, "SAT and ATPG: Boolean engines for formal hardware verification," in *Int'l Conf. on CAD*, 2002, pp. 782–785.

[9] R. Wille, D. Große, F. Haedicke, and R. Drechsler, "SMT-based stimuli generation in the SystemC Verification library," in *Forum on Specification and Design Languages*, 2009.

[10] N. Kitchen and A. Kuehlmann, "Stimulus generation for constrained random simulation," in *Int'l Conf. on CAD*, 2007, pp. 258–265.

[11] K. Gent and M. S. Hsiao, "Fast multi-level test generation at the RTL," in *IEEE Annual Symp. on VLSI*, 2016, pp. 553–558.

[12] K. Laeufer, J. Koenig, D. Kim, J. Bachrach, and K. Sen, "RFUZZ: Coverage-directed fuzz testing of RTL on FPGAs," in *Int'l Conf. on CAD*, 2018.

[13] T. Toffoli, "Reversible computing," in *Automata, Languages and Programming*, vol. 85, Springer, 1980, pp. 632–644.

[14] A. Zulehner and R. Wille, "Make it reversible: Efficient embedding of non-reversible functions," in *Design, Automation and Test in Europe*, 2017.

[15] D. Maslov and G. W. Dueck, "Reversible cascades with minimal garbage," *IEEE Trans. on CAD of Integrated Circuits and Systems*, vol. 23, no. 11, pp. 1497–1509, 2004.

[16] Z. Zilic, K. Radecka, and A. Kazamiphur, "Reversible circuit technology mapping from non-reversible specifications," in *Design, Automation and Test in Europe*, 2007.

[17] L. Burgholzer and R. Wille, "The power of simulation for equivalence checking in quantum computing," in *Design Automation Conf.*, 2020.

[18] L. Burgholzer, R. Kueng, and R. Wille, "Random stimuli generation for the verification of quantum circuits," in *Asia and South Pacific Design Automation Conf.*, 2021.

[19] N. Linden and R. de Wolf. "Lightweight detection of a small number of large errors in a quantum circuit." arXiv: 2009.08840 [quant-ph]. (2020).