Quantum Computation For Minimal Cut Set Identification In Fault Trees

Gabriel San Martín Silva[a,b,c], Enrique López Droguett[a,b,c]
[a]Department of Civil & Environmental Engineering, University of California, Los Angeles, USA
[b]Center for Reliability Science and Engineering, University of California, Los Angeles, USA
[c]The B. John Garrick Institute for the Risk Sciences, University of California, Los Angeles, USA

Abstract: Fault Tree Analysis is a widely used framework in risk, reliability, and safety engineering for identifying the root causes of system-level failures. By translating the relationship between component failures and asset behavior into a Boolean function, Fault Tree models enable the qualitative and quantitative assessment of complex engineering systems. However, as the number of components and their interconnectivity within the system increases, so do the computational resources required for their assessment. While significant progress has been achieved in adapting traditional techniques to keep up with system growth, it is still important to explore the potential advantages that new computational paradigms could bring to the field. In this regard, quantum computation stands out as an ideal candidate, due to its potential to tackle problems that are intractable by traditional techniques. To explore this research path, we tackle the problem of minimal cut set identification in coherent Fault Trees, a fundamental task within Fault Tree analysis. To this end, this study proposes a novel quantum circuit capable of detecting minimal cut set configurations, enabling its use as an oracle function in the Grover algorithm. Using the proposed quantum-based approach, we demonstrate theoretical and numerical evidence of a quadratic reduction in the number of samples required to identify all minimal cut set configurations when compared against random sampling.

Keywords: Fault Tree Analysis, Quantum Computation, Minimal Cut Set Identification, Grover Algorithm.

1. INTRODUCTION

Fault Tree Analysis is a fundamental tool in the reliability and safety assessment of engineering systems. Originally developed for assessing defense systems, Fault Tree Analysis is now used across several industries, including nuclear, power, transportation, and health sectors. Its widespread applicability can be attributed to several key reasons, such as the simplicity in its formulation, its flexibility in representing a variety of systems, and most importantly, its capability to perform both qualitative and quantitative assessments.

Among the tasks performed by FTA, one of the most important is the identification of irreducible configurations that can cause the system to fail. These configurations are known as Minimal Cut Sets (MCS). Formally, they are defined as a minimal set of basic components, whose simultaneous failure triggers the failure of the overall system. Identifying minimal cut set configurations in an engineering system is crucial for two main reasons. First, it allows practitioners to harden the system by either building redundancies around weaker components or improving their reliability. Second, identifying minimal cut sets that cannot be eliminated through system hardening enables the preparation of strategies to improve the system's resilience. However, as the demand for more efficient, safe, and sustainable engineering increases, so does their size and complexity. This makes the identification of minimal cut sets a resource-intensive task, due to the exponentially large space of possible component configurations and the increased complexity of the resulting Boolean function, which restricts the possible advantages that can be achieved by using traditional algorithms.

The FTA community has achieved important advances in addressing this relevant challenge by updating existing ideas and developing new ones to match the ever-increasing scale of engineering systems. Examples of these efforts include recent developments in Boolean Decision Diagrams [1], and approaches based on Boolean Satisfiability algorithms (SAT) [2]. However, the expected scale of future systems and the potential interconnectivity between them demands the exploration of novel techniques outside of traditional computation.

In this regard, Quantum Computation stands out as a promising field to tackle these types of challenges. The reasons for this are multiple. First, quantum computing algorithms present attractive capabilities in the parallel evaluation of Boolean functions, which can be useful in the determination of the state of the system. Second, through its formulation, quantum computation is naturally predisposed to work with binary random variables and categorical distributions, such as the ones that commonly arise in standard Fault Trees. Third, quantum computing enables the use of a powerful set of techniques for unstructured search algorithms, which can be readily applied in the context of minimal cut set identification. Indeed, approaches that use quantum computation for minimal cut set identification have recently started to be explored by the community. However, these early approaches are either too resource-intensive [3], or do not target minimal cut set configurations efficiently [4], targeting cut sets instead and therefore requiring an additional reduction step that adds complexity to the overall algorithm.

To overcome these existing disadvantages, this study proposes a quantum-based approach to directly identify all minimal cut set configurations in a standard, coherent Fault Tree. This is done through the combination of a novel quantum circuit capable of identifying whether a configuration is a minimal cut set and the well-known Grover algorithm [5] to increase the likelihood of sampling these configurations. We present compelling numerical and theoretical evidence of a decrease in the number of queries required to identify all minimal cut set configurations in a Fault Tree.

The study is organized as follows. Section 2 presents a math-based overview of quantum computation and the Grover algorithm. Then, Section 3 reviews the concept of a *Quantum Fault Tree*, presenting an algorithm to translate any standard Fault Tree model into a quantum operation. Section 4 builds on top of the Quantum Fault Tree technique to present a quantum operation that can detect, in parallel, all configurations of a Fault Tree that are minimal cut sets. This quantum operation, denoted as U_{mcs}, is then used alongside the Grover algorithm to develop a novel quantum algorithm for minimal cut set identification. Section 5 presents theoretical and numerical evidence of a quadratic reduction in the number of samples required to identify all minimal cut set configurations when compared against random sampling. Finally, Section 6 presents the concluding remarks of the study, highlighting avenues for future research in the field.

2. BACKGROUND: QUANTUM COMPUTATION

This section first reviews the general principles of quantum computation from a mathematical point of view. Then, it relates those principles with the hardware implementation known as “gate-based” quantum computation. Finally, it describes the Grover algorithm [5], a fundamental result that serves as one of the main components in the proposed approach for minimal cut set identification.

2.1. Mathematical Foundations

Let us start by considering ψ as a quantum system. The state of ψ is represented using the ket notation as $|\psi\rangle = \sum_{i=0}^{N-1} c_i |e_i\rangle$, where $c_i \in \mathbb{C}$ and $|e_i\rangle$ is the i-th standard basis vectors of $\mathbb{C}^N$. The collection $\{|e_i\rangle\}_{i=0}^{N-1}$ is used to represent all possible states of the system $|\psi\rangle$, enumerated from $i = 0$ to $N - 1$. Here, we assume that these states follow an arbitrary, but consistent order.

In quantum computation, defining the state of a quantum system as a linear combination of its possible states is an approach used to convey uncertainty about the system: the probability that ψ is in state e_i is given by $p(e_i) = |c_i|^2$, where $|\cdot|$ is the Euclidean norm operation. For this reason, the collection of complex coefficients $\{c_i\}_{i=0}^{N-1}$ is referred to as *probability amplitudes*, and they represent one of the distinctive features of quantum computation when compared to the approach used in traditional systems to model stochasticity. An important corollary from this definition is that all valid quantum states $|\psi\rangle$ must be unit-length vectors since the sum of their squared coefficients is required to add up to unity.

The coefficients of a quantum state can be modified through the application of *quantum operations.* These operations are represented as the multiplication of a unitary matrix, U, with the quantum state $|\psi\rangle$. Quantum operations modify the underlying probability distribution encoded in the quantum state. Algorithmic design in quantum computation is equivalent to selecting the set of unitary matrices that are going to be multiplied by the quantum state. By carefully modifying the coefficients of a quantum state, its stochasticity can be controlled. This is a core interpretation of quantum computation: controlling the behavior of a stochastic system to guide it towards areas of the probability space that represent a solution to a certain computational task.

Three important considerations should be noted about unitary matrices. First, they preserve the norm of the vectors on which they are applied, maintaining the validity of quantum states. Second, they are easily invertible, $UU^\dagger = U^\dagger U = I$. Third, unitary matrices are closed under matrix multiplication. This enables the composition of intricate unitary operations by multiplying simpler unitary matrices together.

These are the mathematical principles of quantum computation required for this study. The next section is devoted to explaining its current implementation into a physical device known as a "gate-based quantum computer".

2.2. Gate-based Quantum Computation

Quantum computers are physical machines composed of multiple two-dimensional quantum systems called *qubits*. Let us represent a qubit with the ket vector $|q\rangle = [c_0 \;\; c_1]^T = c_0|e_0\rangle + c_1|e_1\rangle \in \mathbb{C}^2$. In this case, the ket $|e_0\rangle$ and $|e_1\rangle$, also written as $|0\rangle$ and $|1\rangle$, are defined as $[1 \;\; 0]^T$and $[0 \;\; 1]^T$, respectively.

Quantum states are formed by combining multiple qubits. A group of qubits is commonly referred to as a qubit *registry*. The mathematical mechanism for this composition is the Kronecker product, $|\psi\rangle = \otimes_{i=0}^{N-1}|q_i\rangle \in \mathbb{C}^{2^N}$. The resulting vector from the Kronecker operation is also a unit-length vector and therefore will constitute a valid quantum state. Its dimension, however, scales exponentially with the size of the qubit registry used in its generation. As a result, while quantum states can in theory have an arbitrary dimension, in the context of gate-based quantum computation they are restricted to dimensionalities equal to 2^N, $N \in \mathbb{N}_{>0}$.

Similarly, quantum operations are also subjected to physical restrictions that limit their generality. While in theory any unitary matrix can be used as a quantum operation, in practice quantum computers only implement a reduced number of quantum operations, applied over one or two qubits at most. These operations are known as *quantum gates*. In what follows, only those quantum gates that are required for the algorithms proposed in this study are presented. For a complete review, the reader is referred to [6].

Table 1. Quantum Gates used in this paper.

Quantum Gate	Symbol	Matrix Expression	Effect
Pauli-X (Bit Flip)	σ_X	$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$	Changes the state of a qubit from $|0\rangle$ to $|1\rangle$, and vice versa.
Pauli-Z	σ_Z	$\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$	When applied over qubit $|\psi\rangle = c_0|0\rangle + c_1|1\rangle$, applies a negative sign over the second coefficient, c_1.
Hadamard	H	$\frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$	When applied over a qubit initialized as $|0\rangle$, it makes its measurement as either $|0\rangle$ or $|1\rangle$ equally likely.
Control-NOT	$CNOT$	$\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}$	Entangles a *control* and a *target* qubit together. The entanglement is as follows: if the control qubit is measured as $|1\rangle$, then a Bit Blip gate is applied to the target qubit.

The behavior of the CNOT gate can be extended from two to N qubits, with the first $N-1$ qubits acting as *control* qubits, and the last one acting as the *target* qubit. The target qubit is only flipped through the action of a Bit Flip gate if all the controls are measured as $|1\rangle$. Note that this "multi CNOT" gate is the quantum equivalent of an AND gate over $N-1$ inputs. We will use it extensively for the quantum encoding process of fault trees, described in Section 3.

Quantum gates are also combined using the Kronecker product to form unitary operations that match the dimension of the quantum state. The generation of these quantum operations from a subset of quantum gates applied over a subset of qubits in the system is represented in a diagram known as *quantum circuit*. The quantum circuit fully describes a quantum algorithm, indicating the type of quantum gate, the location, and order in which they should be applied to the system.

A final implementation limitation of quantum computation is the observability of quantum states. Due to restrictions derived from quantum mechanics, the set of complex coefficients $\{c_i\}_{i=0}^{N-1}$ cannot be experimentally observed. However, the underlying probability distribution represented by this set can be characterized by repeatedly preparing and measuring the quantum state. In this measurement process, the quantum state $|\psi\rangle$ is *collapsed* towards one of its feasible states e_i, following the probability distribution $p(e_i) = |c_i|^2$.

Measurement is a destructive operation in quantum computation, meaning that a quantum state cannot be measured more than once without reconstructing it. By rebuilding the quantum state and repeating this process enough times, the squared amplitudes of the complex coefficients can be estimated with arbitrary accuracy. This experimental limitation enables the understanding of quantum states as probability distributions over bitstrings, a fundamental interpretation for its application to Fault Tree Analysis. To see this, notice that each qubit, when measured, will collapse towards one of its basis states given by the set $\{|0\rangle, |1\rangle\}$. Consequently, the measurement of a quantum state composed of N qubits can be read as an ordered sequence of $0s$ and $1s$, i.e., a bitstring of length N.

2.3. Grover Algorithm

The Grover algorithm [5] is one of the few quantum computing algorithms with a proven computational advantage against traditional approaches. Its purpose is to selectively increase the likelihood of measuring bitstrings that satisfy a Boolean function f known as the *oracle function*. This is done through the repeated application of a quantum operation known as the Grover operation, U_G. To explain its formulation, let us first assume that a quantum state $|\psi\rangle$ is prepared by applying the unitary operation U_A to an N-qubit registry, each one of them originally prepared in the $|0\rangle$ basis state. In mathematical terms, $|\psi\rangle$ is written as $U_A|0\rangle^{\otimes N}$.

The Grover operation is given by $U_G = U_A^\dagger S_0 U_A S_f$, where the dagger symbol indicates conjugate transpose, S_f is a unitary operation implementing an oracle function that applies a negative sign to states that satisfy the oracle function f (see Section 4 for details), and S_0 is a diagonal matrix given by $S_0 = diag([1, -1, -1, \dots, -1]) \in \mathbb{C}^{2^N \times 2^N}$, which can be easily generated using the set of quantum gates described in Section 2 [7]. The Grover operator has the effect of updating the probability of sampling a marked state from p_a to $\sin^2\big((2k+1)\theta_a\big)$, where $\theta_a = \text{asin}\big(\sqrt{p_a}\big)$ and k is the number of times U_G is applied to the quantum state $U_A|0\rangle^{\otimes N}$. As the reader can infer, the probability of sampling marked states can be, in theory, equal to 1. This occurs when the value of k is equal to $k^* = \pi/(4\theta_a)$. However, in practice U_G can only be applied a discrete number of times, and therefore k^* is selected as $\lfloor \pi/(4\theta_a) \rfloor$.

This study uses the Grover algorithm to increase the likelihood of measuring a state that represents a minimal cut set configuration. A key aspect of the proposed approach is rooted in the fact that fault tree diagrams are just graphical representations of Boolean functions, and as such can be used to form oracles that guide the measurement process towards failure states, avoiding unnecessary computational expense. The implementation of fault trees as quantum circuits is described in the following section.

3. FAULT TREES AS QUANTUM CIRCUITS

Fault tree models represent Boolean functions connecting the operational status of basic components to the potential failure of the overall system. In this section, an approach to translating a traditional Fault Tree model into a quantum operation is described. This exposition is based on a previous PSAM16 conference article presented by the authors [8].

The procedure for translating a fault tree model into a quantum operation consists of three key steps. First, a registry composed of $N = N_{BE} + N_{IE} + 1$ qubits is prepared in the $|0\rangle^{\otimes N}$ state. Here, N_{BE} and N_{IE} are the number of basic and intermediary events in the fault tree, while the extra qubit is used to store the status of the TOP event. Note how independent qubits are used to store the final state of the system. For this, we shall use the convention $\{|0\rangle, |1\rangle\} \rightarrow \{operational, failed\}$ when interpreting the bitstring that results from measuring the quantum circuit.

The second step in the proposed methodology is the application of a quantum operation to encode the initial stochastic state of the basic events. In this step, the methodology makes full use of the superposition property to induce a probability distribution over the $2^{N_{BE}}$ possible basic event configurations. The quantum operation is defined as $U_{BE} = U_p \otimes I$, where $U_p \in \mathbb{C}^{2^{N_{BE}} \times 2^{N_{BE}}}$ is applied to the basic events qubit registry, while the identity matrix $I \in \mathbb{C}^{2^{N_{IE}+1} \times N_{IE}+1}$ is used as a "padding" matrix to make the operation match the dimension of the overall quantum state. While in principle any unitary operation U_p can be used to encode a probability distribution into the basic events' configuration, this study uses $U_p = \bigotimes_{i=0}^{N_{BE}-1} H = H^{\otimes N_{BE}}$. This operation has the practical effect of assigning a probability of failure equal to 0.5 to all basic components, resulting in all possible configurations of basic events being equally likely and allowing us to explore all of them uniformly.

The final step in the translation of a fault tree model to a quantum circuit is the encoding of the logic gates that generate the intermediary and TOP events. For this, we shall provide quantum circuits that are equivalent to the NOT and AND logical operators. Due to the functional completeness of these two operations, all other logical relationships can be created by properly combining them. In particular, we shall use this approach to generate a quantum version of the OR logical gate.

First, the quantum counterpart of applying a NOT gate over a binary variable i is simply applying the Pauli-X quantum gate over the corresponding qubit $|q_i\rangle$, as described in Section 2.2. For the logical AND gate, its behavior can be emulated by applying an MCNOT quantum gate to the quantum circuit, using as control qubits those representing the set of basic and intermediary events that act as inputs of the original AND gate in the fault tree, and storing the result in the qubit representing the output of the gate.

These two quantum equivalent gates can be combined to form the logical OR gate. For this, let us first note the Boolean relationship $OR(\{x_i\}) = \overline{AND}(\{\bar{x}_i\})$, where $\{x_i\}$ is a set of binary variables and a bar indicates negation. This relationship can be readily implemented in a quantum circuit by using the quantum versions of the NOT and AND gates to generate a quantum equivalent of the OR gate.

The prior steps enable the implementation of an iterative strategy that encodes the intermediary events of the fault tree in the N_{IE} qubits initially allotted for their statuses. A similar approach can be followed to encode the TOP event. Let us denote the resulting unitary operations as U_{IE} and U_{TOP}, respectively. This allows the composition of a unitary operation $U_{ft} = U_{TOP}U_{IE}U_{BE}$ representing the original fault tree. The number of qubits used by the proposed methodology scales linearly with the number of basic and intermediary events in the fault tree. This is an important improvement when compared to other approaches that propose the encoding of fault trees by first transforming them to a conjunctive normal form.

The resulting quantum state, $|\psi\rangle = U_{ft}|0\rangle^{\otimes N}$ holds a superposition of all possible system outcomes. In generating this quantum state, the quantum computer has calculated the result of the fault tree Boolean function for all the configurations of basic events using only one application of U_{ft}. This is a direct consequence of the linearity of quantum operations, and particularly of the distributive property of matrix multiplication. Furthermore, it is evidence of the parallelism capabilities of quantum computation. Nevertheless, while the computation can be executed in parallel, the restrictions in the observability of quantum states prevent us from retrieving and accessing all these results at once. Instead, the measurement process only enables the retrieval of one of these outcomes per execution of the circuit before the quantum state is collapsed. Consequently, searching for minimal cut sets by repeatedly executing and measuring the fault tree's quantum circuit has an equivalent query complexity to that of traditional Monte Carlo sampling. In other words, no advantage can be achieved by exclusively implementing a fault tree as a quantum model. However, the encoding of the fault tree as a quantum circuit enables the utilization of additional quantum operations in the data transformation pipeline, such as the Grover algorithm.

By applying the Grover algorithm using U_{ft} as the oracle operation, we can artificially enhance the likelihood of measuring states that satisfy the underlying Boolean function, i.e., those configurations that correspond to the cut sets of the original fault tree. While this provides an advantage over traditional Monte Carlo approaches, it is unlikely to have a significant impact on practical fault trees. The reason is that this approach does not increase the likelihood of minimal cut set configurations, but rather that of cut sets instead. In a practical situation, the number of cut sets in a fault tree is likely to be orders of magnitude higher than the number of minimal cut sets. This makes necessary the use of an additional reduction step to obtain minimal cut sets from cut sets, increasing the complexity of the algorithm, and wasting significant amounts of computation, since one cut set can span multiple minimal cut sets. This is a current shortcoming of many existing approaches. In the next section, we propose an algorithm to generate an alternative oracle operation that targets configurations identified as minimal cut sets. This novel approach allows us to selectively increase the likelihood of measuring only those configurations, resulting in a significant decrease in query complexity.

4. PROPOSED ALGORITHM FOR MINIMAL CUT SET IDENTIFICATION.

Section 3 reviewed the approach to generating the oracle operation U_{ft}, which can be used in combination with the Grover algorithm to increase the likelihood of obtaining cut-set configurations from the measurement process. In this section, we improve on this formulation by selectively increasing the likelihood of only those configurations that correspond to minimal cut sets. To this end, we first propose a novel quantum operation,

U_{mcs}, capable of verifying whether a given configuration of basic events is a minimal cut set. Then, by using this operation as U_A in the Grover algorithm, we propose a quantum-based algorithm for minimal cut set identification. Our formulation is inspired by the following definition of a minimal cut set in a standard, coherent fault tree:

Definition 1: Minimal cut set configuration

A given configuration of basic events, represented as a bitstring b, is a minimal cut set if and only if:

1. *Configuration b causes the occurrence of the TOP event, i.e., b is a cut set.*
2. *Preventing the failure of any failed basic event in b also prevents the failure of the system, i.e., configuration b is irreducible as a cut set.*

Mathematically, Definition 1 can be written as the Boolean function shown in Eq. (1).

$$f_{mcs}(b) = f_{ft}(b) \wedge \bigwedge_{i \in F(b)} \overline{f_{ft}}\big(s(b,i)\big) \tag{1}$$

where f_{ft} is the underlying Boolean function representing the fault tree, $F(b)$ is a collection of the indices of failed basic events in configuration b, $F(b) = \{\, j | b_j = 1\}$, and $s(b,i)$ is a *switching* function that flips the state of basic event i in configuration b.

Alternatively, Eq. (1) can be interpreted as the serial evaluation of multiple copies of the original fault tree. The first of these copies is evaluated using the original configuration b, while the rest are evaluated using the modified configurations $s(b,i),\ i \in F$. During the evaluation of $f_{ft}(b)$, the Boolean formula verifyies condition (1) of Definition 1, to test whether b is a cut set of the fault tree. During the posterior evaluation of $\overline{f_{ft}}\big(s(b,i)\big), i \in F$, the algorithm verifies condition (2) of Definition 1, to test whether preventing the failure of basic event i prevents the failure of the overall fault tree. Finally, a series of conjunctive operations are used to verify the overall minimal cut set condition.

The approach to translate Eq. (1) into a quantum operation U_{MCS} is presented in Figure 1, divided into 3 distinct stages. As shown, the proposed quantum circuit requires six qubit registries of lengths: $\big||q_{aux}\rangle\big| = 1$, $\big||Q_{BE}\rangle\big| = N_{BE}$, $\big||Q_{IE}\rangle\big| = N_{IE}$, $\big||q_{TOP}\rangle\big| = 1$, $\big||Q'_{TOP}\rangle\big| = N_{BE}$, and $\big||q_{mcs}\rangle\big| = 1$, where N_{BE} and N_{IE} are the number of basic and intermediary events in the fault tree, respectively. As such, the total number of qubits required to generate U_{mcs} is $2N_{BE} + N_{IE} + 3$, maintaining a linear scaling in the number of basic and intermediary events.

Stage #1 corresponds to the implementation of the first clause of Eq. (1). For this, the operation U_{ft} is applied to the quantum circuit, followed by the application of the adjoint $U_{IE}^{\dagger}$, effectively reversing the encoding of the intermediary events and freeing the qubit registry $|Q_{IE}\rangle$ for future computation. The second stage consists of the repeated application of a set of quantum operations that encode the set of Boolean functions $f_{ft}\big(s(b,i)\big)$. For each repetition $i \in \{1, \ldots, N_{BE}\}$, a modified version of U_{ft} is applied to the quantum state, where two changes are made. The first change consists of replacing basic event i by $|q_{aux}\rangle$, which is initialized and always maintained in the $|0\rangle$ state, effectively simulating the configuration where the basic event i does not fail. Second, the result of the TOP event under those conditions is stored in the i-th qubit of registry $|Q'_{TOP}\rangle$. To finalize each repetition, the adjoint $U_{IE_i}^{\dagger}$ operation is applied, once again reverting the state of the qubit registry $|Q_{IE}\rangle$ and allowing its reutilization. The final step, stage 3, applies a series of CNOT gates and a global AND gate to combine the states of the combined registry $\{|q_{TOP}\rangle, |Q'_{TOP}\}$ into a signal that identifies minimal cut set configurations, storing the result in qubit $|q_{mcs}\rangle$.

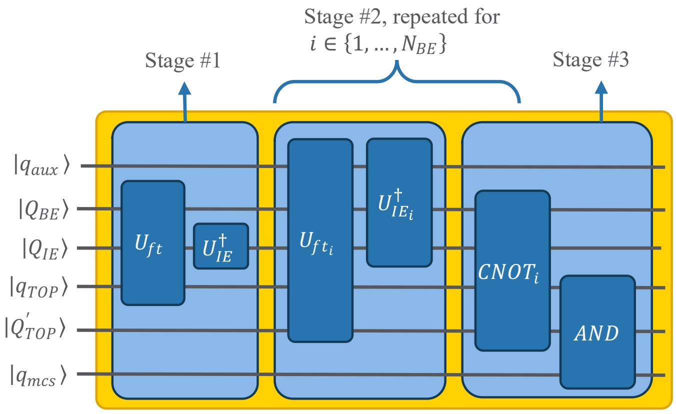

Figure 1. Quantum circuit representation of U_{mcs}.

With this formulation, we can define $U_A = U_{mcs}$ in the definition of the Grover operation, resulting in $U_G = U_{mcs}^{\dagger} S_0 U_{mcs} S_f$. In this case, the oracle operation S_f is defined as a single Pauli-Z gate applied to the qubit registry $|q_{mcs}\rangle$. The proposed procedure to find minimal cut set configurations involves preparing the quantum state $|\psi\rangle = U_G^k U_{mcs}$, with k defined following Section 2.3. We hypothesize that measuring this quantum state should result in a stochastic process that requires fewer samples to produce all minimal cut set configurations in the fault tree. The next section tests this hypothesis and shows both numerical and theoretical proof of its validity.

5. THEORETICAL AND EXPERIMENTAL VALIDATION

In Section 2, we established that preparing and measuring a quantum state is equivalent to sampling from a probability distribution with a support composed of all possible bitstrings of length N. Then, in Section 4 we proposed a quantum-based methodology that utilizes the Grover algorithm to increase the probability mass associated with bitstrings representing minimal cut sets configurations. This section shows our validation approach. For this, we divide this section into two stages. For the first stage, we show numerical evidence that our proposed approach is capable of selectively increasing the likelihood of minimal cut set configurations. For the second stage, we show numerical and theoretical evidence to support that our proposed approach can provide a quadratic reduction in the number of samples required to find all minimal cut sets when compared against a random sampling approach.

To start our validation, Figure 3 explores how the use of the Grover algorithm modifies the probability distribution encoded in the quantum state. For this, we utilize as a case study the fault tree shown in Figure 2, with $N_{BE} = 8$, $N_{IE} = 4$, and $U_p = H^{\otimes N}$. That is, we assign a failure probability equal to 0.5 to all basic events. This resulting fault tree consists of $2^{N_{BE}} = 256$ total possible configurations, all equally likely, with 16 of them corresponding to minimal cut sets. We prepare and measure three versions of the quantum state $|\psi\rangle = \left(U_A^{\dagger} S_0 U_A S_f\right)^k U_A |0\rangle^N$. The first one, shown in Figure 3a, uses $k = 0$ (no Grover operations) and $U_A = U_{ft}$. The second version, shown in Figure 3b, uses $k = 1$ and $U_A = U_{ft}$, i.e., increasing the likelihood of cut sets in general. The final version, shown in Figure 3c, represents our proposed approach. Here, we use $k = 3$ and $U_A = U_{mcs}$, selectively increasing the probability mass of those configurations that are recognized as minimal cut sets. Note that Figure 3 is in logarithmic scale.

Figure 3a shows that the lack of a Grover operator in conjunction with setting $U_p = H^{\otimes N}$ results in a quantum state that assigns equal probability to all possible system outcomes. On the other hand, Figure 3b shows that the application of the Grover operator with operator with $U_A = U_{ft}$ results in a selective increase in the sampling probability of failure outcomes. While this is expected to provide an advantage against random sampling approaches, it is still an important inefficiency since most of the computation will be wasted in producing configurations that are not minimal cut sets. Finally, Figure 3c shows the effect of applying a Grover operator using the proposed circuit U_{mcs} as the state preparation operation U_A. From the results, it is clear that the proposed approach is capable of selectively increasing the likelihood of only those states that are recognized as minimal cut sets, treating all the rest (operational outcomes and cut sets) in the same manner.

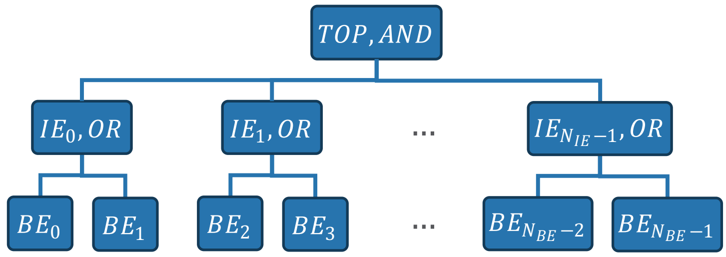

Figure 2. Fault tree structure used in this paper. As mentioned in Section 3, we use an equal failure probability for all basic events, $p_i = 0.5.\ \forall i \in \{1, \ldots, N_{BE} - 1\}$.

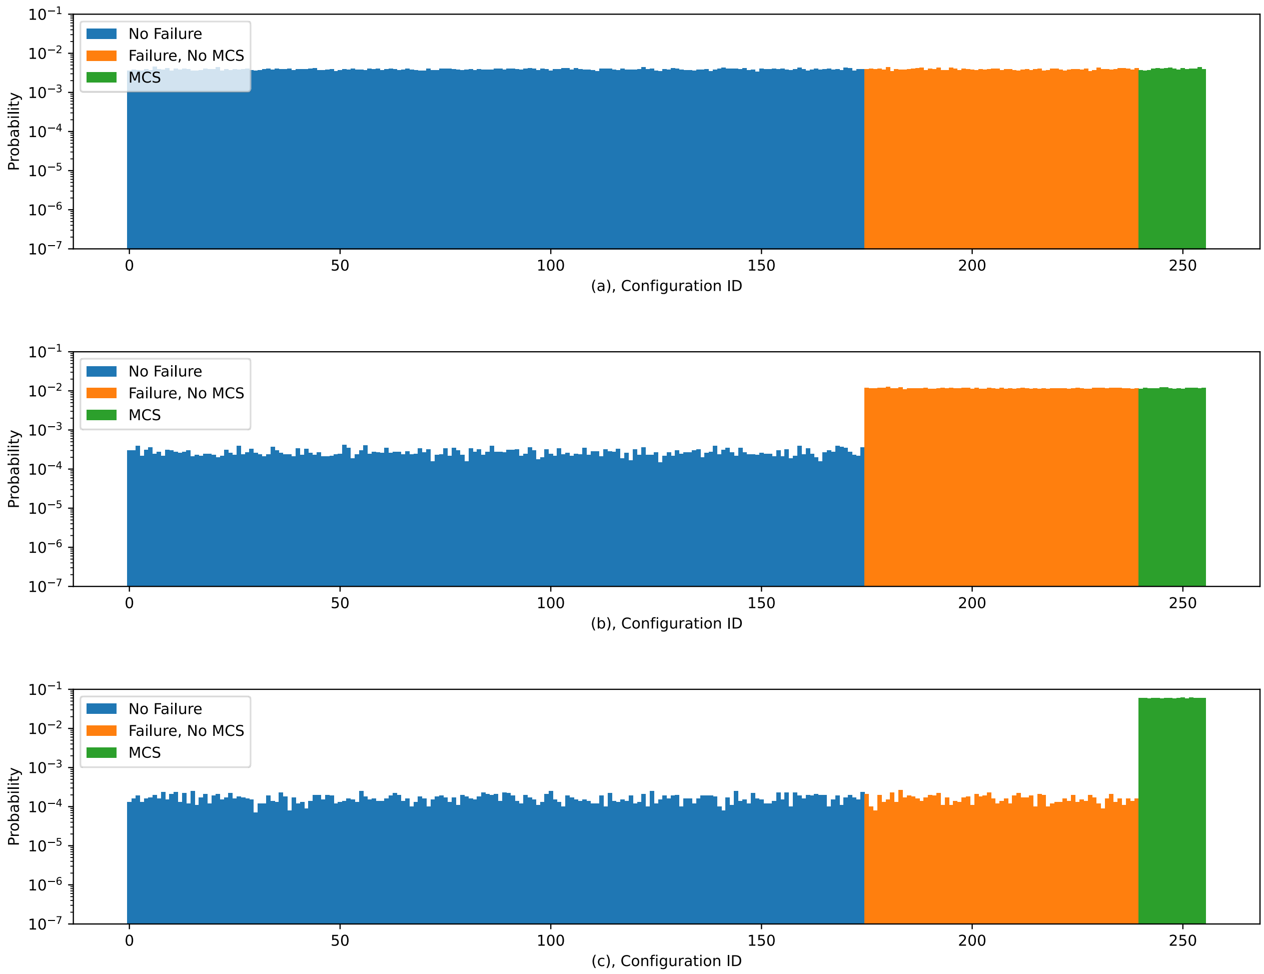

Figure 3. Histograms depicting the underlying probability distribution of three distinct quantum states: (a) $|\psi\rangle = U_{ft}|0\rangle^N$, (b) $|\psi\rangle = \left(U_{ft}^{\dagger} S_0 U_{ft} S_f\right)^1 U_{ft}|0\rangle^N$, and (c) $|\psi\rangle = \left(U_{mcs}^{\dagger} S_0 U_{mcs} S_f\right)^3 U_{mcs}|0\rangle^N$. Each quantum state is prepared and measured 1E5 times to create these histograms.

For the second stage of validation, we implement a series of fault trees according to the structure in Figure 2, each with an increasing number of basic events. We formally compare three methods for minimal cut set identification. As a baseline, we first consider a traditional random sampling approach, where all configurations are sampled with equal probability and evaluated to identify whether they are a minimal cut set or not. While inefficient, this approach gives us a point of comparison to see if the quantum-based approach can provide an advantage versus a naïve, conventional solution methodology. We henceforth refer to this approach as *FT*. The second approach is the one described in Section 3, where the cut sets of the fault tree are targeted by the Grover algorithm. We shall henceforth identify this approach as *QFT*. Finally, the last approach, henceforth identified as *MCS,* uses the Grover algorithm to target minimal cut sets directly using the proposed unitary operation U_{mcs}. All experiments were executed in a traditional computer using quantum simulation software (Python + Pennylane) equipped with 128 GB of RAM. The use of a quantum simulator instead of a quantum computer has the advantage of enable theoretical research without the consideration of hardware errors that affect negatively the accuracy of the results. Currently, hardware errors are a significant issue in contemporary quantum computers, and large research efforts are devoted to devise methods to correct and tackle these issues. While hardware related topics are not the focus of this paper, the interested reader can review the strategies to perform quantum error correction and the advances achieved in the field in the following reference: [9].

Figure 4a shows the average number of queries required to find all minimal cut set configurations by the three approaches tested in this study. These were calculated by repeatedly preparing and measuring the corresponding quantum states or executing the traditional random sampling approach until all minimal cut sets were found. We can see that for all the trees tested in the quantum simulator, the performance of the proposed approach surpasses both the existing quantum and traditional approaches, QFT and FT, respectively.

Due to limitations related to the exponential scaling of quantum states and their impact on the memory usage for their simulation in traditional hardware, only systems up to 32 qubits can be simulated. This severely limits the fault tree sizes that can be tested numerically to $N_{BE} \leq 10$. However, as seen in Section 2.3, the probability increase given by the Grover algorithm only depends on the initial probability of sampling a target configuration, denoted as p_a. Due to the relatively simple structure of the fault tree shown in Figure 2, these initial probabilities can be easily computed for both the QFT and MCS approaches, and therefore their performances can be projected for an arbitrarily large fault tree. For other systems, such as a more complex fault tree, where this probability may not be known before hand, an iterative approach can be used to find the ideal number of Grover operator applications. This approach is fully described in [5], Theorem 3. However, the estimation of p_a is an active area of research in the field and further improvements will be required to make the technique fully applicable in complex systems.

The results of this projected performance are shown in Figure 4b, where we have also included an exponential fitting for each curve.

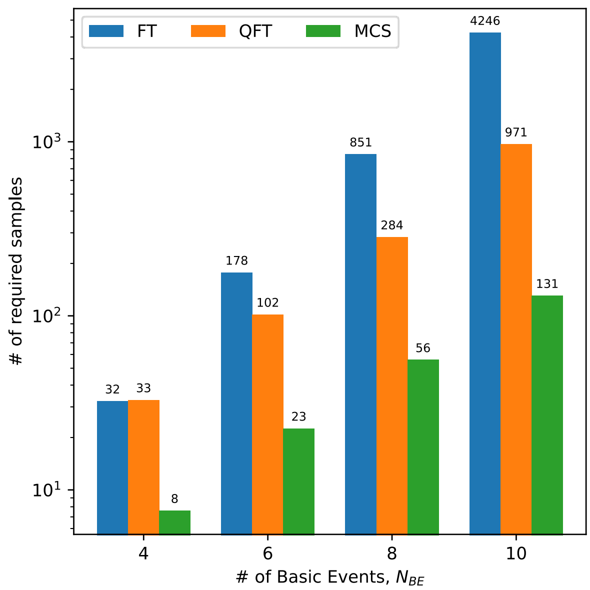

(a) Results obtained through numerical simulation of quantum states.

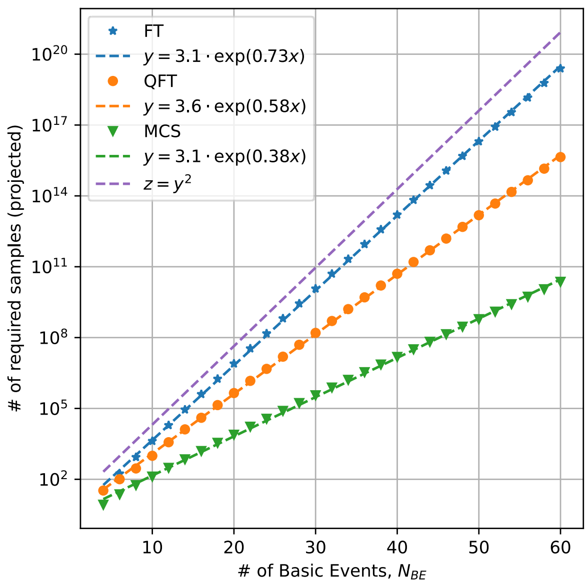

(b) Results obtained through the theoretical projection of performance.

Figure 4: Number of required samples by the tested approaches as a function of the number of basic events.

Figure 4b clearly shows that the proposed approach uses a significantly smaller number of samples to identify all minimal cut set configurations when compared against the FT and QFT approaches. This advantage scales non-linearly with the size of the fault tree, represented in this example by the number of basic events. Moreover, the fitted curves clearly show an approximately quadratic reduction in the number of samples required by our approach, MCS, when compared against random sampling.

6. CONCLUSION

In this paper, we introduced a quantum operation that can selectively mark configurations that are recognized as minimal cut sets in a standard, coherent fault tree. Using this quantum operation in combination with the Grover algorithm, we proposed a novel methodology that can achieve a quadratic reduction in the number of samples required to identify all minimal cut sets when compared against a random sampling approach. Numerical and theoretical evidence is provided to validate this speedup.

While the proposed approach is shown to provide an advantage against the classical approach used as a baseline, it is not without its drawbacks. Indeed, we mention three aspects that can be improved in successive investigations and that can lead to exciting novel research in this area. First, the initial probability p_a was computed using the structure of a pre-defined fault tree. However, in practical fault trees, this probability will not be known a priori. A suggestion is to generate alternative approaches to the estimation of the number of Grover operation applications, k. As a second avenue for future research, we acknowledge that the traditional baseline tested in this paper is not the state-of-the-art approach to identifying minimal cut sets in fault trees. More advanced methods exist that make better use of existing structures within the fault tree itself. In this aspect, a thorough comparison of the proposed approach against such approaches is left for future work. Finally, this study defines performance as the number of samples required to fulfill a certain task. This definition plays nicely with a quantum-based approach due to the interpretation of quantum states as probability distributions. Nonetheless, it is still an open question how to compare quantum-based approaches against algorithms that are, by their nature, not stochastic, or that do not employ a sampling strategy. Within the context of fault trees, this is particularly important given that the current state of the art includes approaches that perform Boolean manipulation instead of sampling to identify minimal cut sets.

References

[1] E. Ruijters and M. Stoelinga, “Fault tree analysis: A survey of the state-of-the-art in modeling, analysis and tools,” *Comput. Sci. Rev.*, vol. 15–16, pp. 29–62, Feb. 2015, doi: 10.1016/j.cosrev.2015.03.001.

[2] W. Luo, O. Wei, and H. Wan, “SATMCS: An Efficient SAT-Based Algorithm and Its Improvements for Computing Minimal Cut Sets,” *IEEE Trans. Reliab.*, vol. 70, no. 2, pp. 575–589, Jun. 2021, doi: 10.1109/TR.2020.3014012.

[3] S. Brand and A. Laarman, “D6.6: Divide and quantum open source software,” NEASQC, 951821, Feb. 2022. [Online]. Available: https://www.neasqc.eu/about-the-project/

[4] M. Hibti, A. Zaiou, Y. Bennani, and B. Matei, “D6.8: State of the art of SAT and PSA solvers in the light of quantum computing,” 30/09/2022, 951821. [Online]. Available: https://www.neasqc.eu/about-the-project/

[5] G. Brassard, P. Hoyer, M. Mosca, and A. Tapp, “Quantum Amplitude Amplification and Estimation,” vol. 305, 2002, pp. 53–74. doi: 10.1090/conm/305/05215.

[6] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information*, 10th anniversary ed. Cambridge ; New York: Cambridge University Press, 2010.

[7] G. San Martin and E. Lopez, “Quantum Fault Trees and Minimal Cut Sets Identification,” 2024, doi: 10.48550/ARXIV.2404.05853.

[8] G. S. M. San Martin, T. Parhizkar, and E. L. Droguett, “Quantum Fault Trees,” 2022, doi: 10.48550/ARXIV.2204.10877.

[9] A. Jayashankar and P. Mandayam, “Quantum Error Correction: Noise-Adapted Techniques and Applications,” *J. Indian Inst. Sci.*, vol. 103, no. 2, pp. 497–512, Apr. 2023, doi: 10.1007/s41745-022-00332-x.