Polynomial Size, Short-Circuit Resilient Circuits for NC

Yael Tauman Kalai
Microsoft Research, Cambridge, MA, USA
MIT, Cambridge, MA, USA

Raghuvansh R. Saxena
Tata Institute of Fundamental Research, Mumbai, India

Abstract

We show how to convert any circuit of poly-logarithmic depth and polynomial size into a functionally equivalent circuit of polynomial size (and polynomial depth) that is resilient to adversarial short-circuit errors. Specifically, the resulting circuit computes the same function even if up to ϵd gates on every root-to-leaf path are short-circuited, *i.e.*, their output is replaced with the value of one of its inputs, where d is the depth of the circuit and $\epsilon > 0$ is a fixed constant.

Previously, such a result was known for formulas (Kalai-Lewko-Rao, FOCS 2012). It was also known how to convert general circuits to error resilient ones whose size is quasi-polynomial in the size of the original circuit (Efremenko et al. STOC 2022). The reason both these works do not extend to our setting is that there may be many paths from the root to a given gate, and the resilient circuits needs to "remember" a lot of information about these paths, which causes it to be large. Our main idea is to reduce the amount of this information at the cost of increasing the depth of the resilient circuit.

2012 ACM Subject Classification Theory of computation → Communication complexity; Theory of computation → Circuit complexity

Keywords and phrases Error-resilient computation, short-circuit errors

Digital Object Identifier 10.4230/LIPIcs.ITCS.2025.90

Funding *Raghuvansh R. Saxena*: Supported by the Department of Atomic Energy, Government of India, under project no. RTI4001.

1 Introduction

Constructing error-resilient circuits is one of the most fundamental problems in theoretical computer science, dating back to von Neumann [30]. In the study of error resilient circuits, the goal is to convert any circuit C into a circuit C' that computes the same function as C even if some of the gates of C' are faulty. Moreover, the goal is to do so with a small overhead in size, *i.e.*, constructing a circuit C' whose size is bounded by a polynomial in the size of C.

In this paper, we show how to convert any polynomial size circuit of poly-logarithmic depth into a polynomial size circuit that is functionally equivalent and is resilient to short-circuit errors, an error model that was introduced by Kleitman, Leighton, and Ma [20], and was soon adopted as a central model in the study of error-resilient circuits [18, 4, 7]. In this error model, the adversary can replace any gate in the circuit with an arbitrary gate $g : \{0,1\}^2 \to \{0,1\}$ as long as $g(0,0) = 0$ and $g(1,1) = 1$ (and in particular can replace an AND gate with an OR gate, and vice versa)[1]. Equivalently, we can say that the value of the gate g is replaced by the value of one of its children (the wire to the other child is "cut

[1] Note that the restriction that $g(b,b) = b$ for all $b \in \{0,1\}$ is necessary since otherwise an adversary can simply flip the result of the output gate, and thus no circuit can be resilient to even a single error.

out"). This model is motivated by applications; indeed, as noted in [20], "stuck-at" and "power-ground" failures resulting from short-circuits or broken connections are more common than other types of errors.

Similar to the works of [18, 4, 7], we consider omniscient adversaries that have full information about the entire circuit, and can corrupt ϵd of the gates on every root-to-leaf path, where d is the depth of the circuit and $\epsilon > 0$ is some fixed constant. These prior works were either restricted to formulas, *i.e.*, logarithmic depth circuits [18, 4], or incurred a quasi-polynomial blowup to the circuit size [7]. Moreover, it was conjectured in [7] that, unlike the case of formulas [18, 4], a quasi-polynomial blowup in the circuit size is necessary when converting general circuits into noise resilient ones.

1.1 Our Results

We show that how to convert any circuit of polynomial size and poly-logarithmic depth (i.e., any NC circuit) into a functionally equivalent circuit of polynomial size (and polynomial depth) that is error resilient.

▶ **Theorem 1** (Informal, see formal version in Theorem 8)**.** *Let C be a Boolean circuit in the class NC that[2] computes a function f. There exists a Boolean circuit C' whose size and depth are polynomial in the size of C that computes f even when a constant fraction of gates in any input to output path in C' are short-circuited.*

In terms of comparison, Theorem 1 improves on the main result of [18] in the sense that it works for all circuits with polynomial size and poly-logarithmic depth while [18] only works for formulas. It also improves upon [7] in the sense that it gets a polynomial blowup in the size instead of the quasi-polynomial blowup obtained in [7]. However, it does not subsume either of these results. For example, [18] also preserve the depth of the original formula up to constant factors but we do not provide such a guarantee, and [7] works for general circuits, including those not in the class NC.

We also mention that [7] believed that the quasi-polynomial blowup they obtain is inherent and conjectured that a polynomial blowup will not suffice like it does in the case of formulas [18, 4]. Theorem 1 disproves this conjecture for the restricted case of circuits in NC. Proving/disproving the conjecture for general circuits is an amazing problem that we leave open. However, proving this conjecture, even in an existential manner, would imply $\mathsf{P/poly} \not\subseteq \mathsf{NC}$, and thus may be extremely hard. Indeed, if $\mathsf{P/poly} \subseteq \mathsf{NC}$, then a general polynomial size circuit has a functionally equivalent circuit in NC and we can use Theorem 1 to make it error resilient with only a polynomial blowup.

Error-resilient dag-protocols

Computation and communication have a deep connection. Just like prior work [18, 4, 7], the current work also exploits this connection via the Karchmer-Wigderson transformation. Specifically, Karchmer and Wigderson [19] show that any circuit computing a function f can be converted into a communication protocol for a related search problem KW_f such that the depth of the circuit matches the length of the communication protocol, and vice versa.

[2] Thus, if f takes n variables as input, the number of gates in C is bounded by a polynomial in n and the depth of C is bounded by a polynomial in $\log n$. The other constants in the theorem statement depend on the degree of these polynomials; see the formal version of the theorem (Theorem 8) for the precise dependence.

An analogous transformation can be shown in the reverse direction, *i.e.* a transformation that takes error-resilient communication protocols to error resilient circuits, and was used in [18, 4] to build error resilient formulas.

The Karchmer-Wigderson transformation does not preserve the size of the circuit and in fact, blows it up to be exponential in the depth. To circumvent this blowup in size, [25, 29] showed how to convert any circuit into a "**dag**-protocol" that preserves both its size and depth. **dag**-protocols are a generalization of communication protocols, and can be represented by the following pebble game: There is a rooted directed acyclic graph, and each non-sink node in the graph belongs to one of the two parties, while each sink node is associated with an output. The game starts with a "pebble" at the root of graph, that is moved along the edges of the graph. At every step, if the pebble is not already at a sink node, the party who owns that node will move the pebble along one of its out-edges. This will end when the pebble reaches a sink node, and the output will be the output of the sink node.

[25, 29] showed that every circuit computing a function f can be converted to a **dag**-protocol with the same size and depth that solves KW_f with a strong correctness guarantee called rectangular correctness (see Definition 5 for a precise definition). The reverse direction also holds, and in fact even holds in the error-resilient setting. This direction was used to get error-resilient circuits from error-resilient **dag**-protocols in [7]. We adopt a similar approach and use the [25, 29] transformation to go from circuits to **dag**-protocols and back. Indeed, our main technical result is the following theorem about error-resilient **dag**-protocols.

▶ **Theorem 2** (Informal, see formal version in Theorem 9)**.** *Let S be a search problem and Π be a* **dag***-protocol that solves S with rectangular correctness. Assume that the depth of Π is poly-logarithmic in the size. There exists a* **dag***-protocol Π' with a polynomially larger size that solves S with rectangular correctness even if a constant fraction of nodes on any root to leaf path in the protocol Π' are adversarially corrupted.*

1.2 Related Work

The works most closely related to ours are the works [18, 4, 7] cited above. We discuss additional related work below.

Other noise models for circuits

Even though the short-circuit error model we consider is the most studied one, other models have also been considered. For example, von Neumann [30] initiated the study of the stochastic noise model, where the noise flips the value of each gate in the circuit independently with some small fixed probability. Von Neumann's model was studied by a long sequence of works, including [6, 22, 23, 11, 8, 9, 17, 13, 12, 10]. In this model, it is known that a circuit of size s can converted to a noise resilient circuit of size $O(s \log s)$, and that a function with sensitivity s' requires a resilient circuits of size $\Omega(s' \log s')$ [30, 6, 22, 13, 12].

In this paper, we study the short-circuit error model of [20] but there is also a different adversarial model studied by [14], where the adversary may corrupt the output of a small constant fraction of the gates at each layer of the circuit in an arbitrary way. They show how to construct error resilient circuits for symmetric functions in this model, by exploiting interesting connections between their model and probabilistically checkable proofs. However, the obtained circuit is only guaranteed to compute, what they call, a "loose version" of

the function, and may err on many inputs. Lastly, we mention the orthogonal direction of constructing testable circuits which are circuits on which errors can detected (but not necessarily corrected) efficiently [2, 1].

Codes for interactive communication

As our work develops error-resilient circuits via error-resilient **dag**-protocols, it is also connected with the vast literature on codes for interactive communication, or simply "interactive codes", that are used to make communication protocols resilient to noise. This field, initiated by the seminal works [26, 27, 28] has received a lot of attention over the last three decades, and various aspects of interactive codes have been extensively studied. A work loosely related to circuits is [5] but see [16] for an excellent survey. We note that all constructions of error resilient **dag**-protocols are heavily inspired by an underlying interactive coding scheme (recall that **dag**-protocols are a generalization of communication protocols).

dag-protocols

Razborov [25] introduced a model of *PLS communication protocols*, and used it generalize the Karchmer-Wigderson transformation [19] in a way that preserves both the size and the depth of the circuit. This connection was used by Krajícek [21], who introduced the technique of *monotone feasible interpolation*, which became a popular method for proving lower bounds on the refutation size in propositional proof systems such as Resolution, and Cutting Planes [3], by reducing to monotone circuit lower bounds. The notion of PLS communication protocols was simplified by Pudlak [24] and Sokolov [29] to the notion of **dag**-*like communication protocols*, which we call **dag**-protocols. Subsequently, a "converse" to monotone feasible interpolation was established in [15] to prove new lower bounds on monotone circuits by lifting lower bounds on Resolution refutations. To the best of our knowledge, our work is the only work, other than [7], to use the notion of **dag**-protocols in a "positive" manner, namely to construct error-resilient circuits, in contrast to prior work which mainly used this notion to prove lower bounds in proof complexity and circuit complexity.

1.3 Additional Discussion

An obvious problem left open by our work is whether general circuits can be made error resilient with only a polynomial overhead. As mentioned above, the only work in this regard is [7], that showed that general circuits can be made error-resilient by incurring a quasi-polynomial overhead. Specifically, if the original circuit has size s and depth d, they construct a resilient version of this circuit that has size $s^{\mathcal{O}(\log d)}$ and depth $\mathcal{O}(d)$. [7] further say that the power of $\mathcal{O}(\log d)$ seems inherent.

The current work manages to get around the $\mathcal{O}(\log d)$ in the exponent, at the cost of increasing depth to be a polynomial (note that $d = (\log s)^{\mathcal{O}(1)}$ in our setting). It is interesting to see if our techniques can be combined with [7] to shave off a small factor off of this exponent for general circuits, at the expense of a larger depth. We leave this investigation to future work. Another great question for future research is whether the polynomial blowup in depth that we suffer is inherent, or is there a way to make circuits in NC error resilient while preserving their depth up to a constant factor.

Another interesting direction for future work is whether the resilient circuit can be computed efficiently, say in time polynomial in the size of the circuit. The initial work of [18] guaranteed efficiency when the initial circuit is a formula although the later work [7]

was not able to generalize it to circuits[3]. The current work for circuits in NC also leaves this direction open. Finally, even though we get resilience to some constant fraction of adversarial errors, we make no efforts to optimize this constant. Finding the right constant is a great question for future work (see [4]).

2 Technical Overview

We now overview the ideas behind Theorem 1 and contrast it with prior work [18, 4, 7]. Roughly speaking, Theorem 1 says that there exist polynomial sized resilient circuits for every circuit C in NC. We let s and d be the size and the depth of C, respectively, and use f to denote the Boolean function computed by C. We will also assume without loss of generality that the circuit C is layered and alternating. Just like prior work, our construction is in three-steps via **dag**-protocols[4]:

1. **Circuits → dag-protocols.** This step is the same as prior work, and uses the Karchmer-Wigderson transformation [19] to transform the circuit that computes f to a **dag**-protocol that solves the Karchmer-Wigderson game KW_f associated with f with rectangular correctness. This transformation satisfies the property that short-circuit errors in the original circuit correspond to running the protocol over a channel with corruption noise and perfect feedback. We will call such channels feedback channels for simplicity and use this property in the Step 3 below.
2. **dag-protocols → error-resilient dag-protocols.** This is the main step in the proof where the input is a **dag**-protocol and the goal is to output a functionally equivalent **dag**-protocol that works on feedback channels. This is the main contribution of this work and is detailed below.
3. **Error-resilient dag-protocols → error-resilient circuits.** Having constructed a protocol that solves KW_f with rectangular correctness even on feedback channels, we use the "inverse" Karchmer-Wigderson transformation to transform it to a Boolean circuit computing f that is resilient to short-circuit errors. This step in our proof can also be found in prior work.

The main contribution of this work is Step 2 above, which transforms **dag**-protocols to error-resilient **dag**-protocols. Analogous transformations in prior work were usually done using interactive coding schemes, an area of research initiated by the seminal works [26, 27, 28] that is dedicated to making communication protocols resilient to errors. At an extremely high level, this is done by designing mechanisms to detect errors fast and then using rewind mechanisms to rewind to a point in the communication history that was before these errors occurred.

However, to rewind to a point in the communication history before the errors occurred, the resilient protocol needs to remember the state at that point. After the Karchmer-Wigderson transformation, this extra memory shows up as a blowup in the size of the resilient circuit, which we want to minimize. To be more specific, prior work on resilient circuits [7], remembers $\mathcal{O}(\log d)$ different states at every point in the resilient protocol which means that each layer in the resilient circuit is quasi-polynomially larger (*i.e.*, has size $s^{\mathcal{O}(\log d)}$) than each layer in the original circuit. Thus, even though [7] increase the number of layers by only a constant factor, the blowup in the size of each layer makes the overall blowup quasi-polynomial.

[3] Note that the error-resilient circuit of [7] is of quasi-polynomial size, and thus polynomial in this size of this circuit means that it is quasi-polynomial in the size of the original circuit.
[4] Readers not familiar with **dag**-protocols may find it useful to first go over Section 2.1 of [7].

Our approach is to greatly reduce the size of every layer, at the expense of the increasing the total number of layers. This means that we cannot even afford to remember $\mathcal{O}(\log d)$ many states at any point. As any interactive coding scheme we are aware of remembers at least logarithmically many states, they would not suit us, and new ideas are needed.

2.1 Our Approach For NC^1

We start by considering a restricted case where the original circuit C lies in NC^1. Even though circuits in NC^1 can be transformed into formulas with only a polynomial blowup, we will stick with the circuit view as it will allow us to generalize later on. Note that prior work already gives a resilient circuit for this case that has size polynomial in s and has depth $\mathcal{O}(d)$. We construct a different resilient circuit that has the same size guarantee but relaxes the depth guarantee[5]. This alternate construction allows us to extend to all of NC.

Our approach

We devise a way to make **dag**-protocols corresponding to NC^1 circuits resilient while remembering only a constant number (in fact, 2) states at each point, at the cost of blowing up the number of layers exponentially. As the depth of NC^1 circuits is logarithmic in n, this exponential blowup is still polynomial in n. Moreover, as we only remember a constant number of states at each point, the blowup in the size of each layer is also polynomial, and combining this with the blowup in the number of layers still gives a polynomial blowup.

Note that any protocol must at least remember the current state, and thus a protocol that remembers only 2 states actually remembers only 1 extra state! Interestingly, for our protocol, this extra state is not a state in the past but actually the next state in the future that the protocol is considering going to. Specifically, if the protocol is currently in state in layer $i-1$ (for some $i > 0$), and wants to advance to a state in layer i, then instead of jumping to that state directly, it is remembered as a "candidate" state until the protocol has enough "confidence" in the candidate to advance. This confidence is implemented via a variable `cnf` that is incremented every time the parties want to advance and decremented every time they suspect that the states in memory are incorrect. Only after the confidence `cnf` hits a certain threshold $\mathfrak{C}_i$ does the protocol actually advance to the state in layer i.

We will view the thresholds as being cumulative, *i.e.*, $\mathfrak{C}_i$ is the total confidence required to go from layer 0 to layer i. This means that $\mathfrak{C}_i - \mathfrak{C}_{i-1}$ is the additional confidence needed to advance from layer $i-1$ to layer i. With this said, the obvious question is how large does $\mathfrak{C}_i - \mathfrak{C}_{i-1}$ need to be? In other words, how much confidence should the protocol have in the candidate before they advance to a state in layer i? To answer this question, one needs to consider the total "loss" of the protocol in case it advances wrongly and reaches an incorrect state. As no state in the past is remembered, the only way the protocol can recover from this incorrect state is if it starts all over again and reaches the state in layer $i-1$, which requires a total increase in confidence of $\mathfrak{C}_{i-1}$. Thus, the increase (which equals $\mathfrak{C}_i - \mathfrak{C}_{i-1}$) of the confidence required to jump from layer $i-1$ to layer i should be the same order of magnitude as the confidence $\mathfrak{C}_{i-1}$ needed to reach layer $i-1$ from the beginning of the protocol. Solving, we get that $\mathfrak{C}_i$ should be (at least) exponential in i. This is the choice that we make, setting $\mathfrak{C}_i = c^i$ for some suitable constant $c > 1$.

[5] The depth is at most the size, so it must still be polynomial in s.

Protocol details

We now provide more details of our protocol and argue its correctness. The protocol starts at the initial (root) node in layer 0 with confidence $\texttt{cnf} = \mathfrak{C}_0 = 0$. At any given time, the protocol is in some layer, say layer $i-1$ with some confidence $\texttt{cnf}$, and trying to continue the simulation from some state in this layer. For this, the protocol first checks[6] if the current state has any errors. If so, the protocol decreases the confidence $\texttt{cnf}$ by 1. This may cause $\texttt{cnf}$ to hit 0, which is taken as a signal to forget all progress and restart the simulation. If no errors are found, then the current state is assumed to be correct and the action to be taken is determined by the value of the confidence $\texttt{cnf}$.

If we have $\texttt{cnf} < \mathfrak{C}_{i-1}$, then it must be the case that the confidence was decreased after reaching layer $i-1$. The protocol simply increases it by 1 trying to get it back to $\texttt{cnf}_{i-1}$ so that further progress can be made. If the confidence $\texttt{cnf} = \mathfrak{C}_{i-1}$, the confidence in layer $i-1$ is high enough for the protocol to compute a candidate in layer i. A candidate $\texttt{nxt}$ for the state to advance to is computed and the confidence is further increased by 1.

Finally, if the confidence $\texttt{cnf} > \mathfrak{C}_{i-1}$, then the protocol must already have a candidate $\texttt{nxt}$ that it is considering going to. The protocol then checks if this candidate was caused due to errors, and if so, decreases the confidence by 1. If the confidence is lowered all the way back to $\mathfrak{C}_{i-1}$, the protocol "forgets" about the candidate $\texttt{nxt}$ (and will compute a new one in the future if needed). On the other hand, if the candidate $\texttt{nxt}$ was not due to errors, the protocol increases the confidence $\texttt{cnf}$ by 1. If this increased confidence reaches $\texttt{cnf}_i$, it is high enough to advance. The protocol forgets about the current state in layer $i-1$ and continue the rest of the simulation from the state $\texttt{nxt}$ in layer i.

Arguing correctness

We now show that the above protocol indeed simulates the original protocol correctly. For this, we first show that if the protocol has high confidence at the end of the protocol, specifically, if $\texttt{cnf} \geq \mathfrak{C}_d$, where d is the total number of layers (depth) of the original protocol, then the output of the protocol must be correct. Indeed, if the confidence is that high, the protocol is in a state in layer d or higher[7], and the only reason this state would not be correct is if the protocol advanced to it incorrectly. However, this means that there must have been at least $\mathfrak{C}_d - \mathfrak{C}_{d-1}$ errors while this state was the next candidate. As the confidence thresholds are exponential, we have that the total number of errors is at least $\mathfrak{C}_d - \mathfrak{C}_{d-1} = \Omega(\mathfrak{C}_d)$, which is more than a constant fraction and therefore, unaffordable.

Now, all we need to show is that the confidence is indeed high at the end of the protocol. For this, we fix any execution of the protocol and partition the rounds into rounds where the protocol wants to "go back" and decrease the confidence $\texttt{cnf}$ (this happens when either the current state or the next state $\texttt{nxt}$ has any errors) and rounds where it does not. We collect the former type of rounds in the set Back. By definition, the confidence will increase in any round not in Back unless there are errors, and will decrease in any round in Back

[6] As in prior work, assuming that the original **dag**-protocol is rectangular correct is crucial for the protocol to be able to check whether or not there are errors. As this is not a novel idea for the current work, we omit a precise description in this sketch. (Very) roughly speaking (see formal definition in Definition 5), a state is said to be rectangular correct for inputs x and y for Alice and Bob respectively, if there exists inputs x' for Alice and y' for Bob such that the protocol reaches that state when executed with inputs x and y' and also reaches it with inputs x' and y. This means that the parties can go over all possible inputs of the other party in order to check if the state is (rectangular) correct or not.

[7] We append the original protocol with dummy states to get states beyond layer d.

unless there are errors. As the total number of errors is only a small constant fraction, this means that all we need to do to show that the confidence is high at the end of the protocol is to show that the size of Back is small.

To show this, we argue that the size of Back is at most a constant factor more than the number of rounds in Back where the protocol did not go back. To see why, note that the first round in Back must be a round where the protocol added an incorrect next state as their candidate `nxt`. As mentioned above, if the next state is in layer i, this only happens when the confidence is exactly $\mathfrak{C}_{i-1}$.

As long as the protocol does not advance to this next state (in other words, it remains a candidate), the confidence must stay above $\mathfrak{C}_{i-1}$, as otherwise the protocol will forget the incorrect candidate and look for a new one. This means that the total number of rounds where the protocol does not go back and increase the confidence exceeds the total number of rounds where the protocol did go back, and we are done. On the other hand, if the protocol does advance to the candidate, there must have been at least $\mathfrak{C}_i - \mathfrak{C}_{i-1}$ errors, and by definition of $\mathfrak{C}$, these are high enough (up to constant factors) for $\mathfrak{C}_i$ rounds of progress. Thus, the protocol can wait until the confidence hits 0 and restart the simulation from there.

2.2 From NC^1 to NC

We now show how we can extend the ideas above to get a simulation for all of NC. Note that the Karchmer-Wigderson transformation transforms general NC circuits to **dag**-protocols of poly-logarithmic depth. As our arguments above increase the depth exponentially, using them on general NC circuits would make the depth super-polynomial.

To reduce the depth, we extend our protocol above to also remember some states in the past. Specifically, if the current state is layer $i-1$, the protocol also remembers the state it encountered in the most recent layer that was a multiple of $(\log s)^j$, for all $j \in \mathbb{N}$. As the total number of layers we have is poly-logarithmic, we get that when j is larger than some fixed constant, the most recent layer that is a multiple of $(\log s)^j$ is just layer 0. This means that the parties only remember a constant number of distinct states, implying as above that the blowup in the size of each layer is still polynomial.

We now argue why remembering these extra states ("meeting points") in the past will also reduce the total number of layers back to polynomial. The reason is that if the protocol remembers meeting point in the near future, then, when the confidence is too low, the protocol does not have to forget all progress and restart the simulation from scratch. Instead, the protocol only has to forget the progress since the last meeting point and restart the simulation from said meeting point. Because of the loss in progress is now smaller, we can allow the confidence increase (which equals $\mathfrak{C}_i - \mathfrak{C}_{i-1}$) required to advance to also be smaller. As this happens for almost all i, the value of $\mathfrak{C}_d$, the highest confidence threshold is significantly smaller, and we can bound it by a polynomial.

The new confidence thresholds

We are yet to specify the precise increase in confidence (which equals $\mathfrak{C}_i - \mathfrak{C}_{i-1}$) required to advance to a state in layer i. For this, we look at the representation of i as an integer[8] in base $\log s$. If $a_1, a_2, \ldots$ are the digits in this representation, we set $\mathfrak{C}_i - \mathfrak{C}_{i-1}$ to be proportional to $c^{\sum_i a_i}$. As there is a meeting point in memory for every power of $\log s$, it can be seen

[8] The actual proof uses the representation of $i-1$ for technical reasons.

that $c^{\sum_i a_i}$ takes into account the distance from all the meeting points stored in memory, while maintaining the exponential growth required between two meeting points. We defer the remaining details to the proof.

Getting rectangular correctness

Just like prior work, for Step 2 above to work, it is crucial that our simulation protocol is rectangular correct, which is a stronger notion of correctness than standard correctness, and is necessary for Step 3 to work. Without going into the precise definition, this means that we need to account for the errors in the rounds where the parties (Alice and Bob) are talking separately. That is, the errors in rounds where Alice is talking cannot be used to offset the rounds where Bob is going back and decreasing the confidence, and vice versa.

Other than doubling the notation we use (*e.g.*, we now need to remember a pair of meeting points, one for Alice and one for Bob, *etc.*), this also introduces a more subtle challenge regarding the fraction of errors that we can be resilient to. To understand this, recall that the layers in the original circuit, and therefore the layers in the original **dag**-protocol, are alternating. This means that the party controlling the layer changes every time the protocol advances a layer. Now, imagine what happens when the adversary spends $\mathfrak{C}_i - \mathfrak{C}_{i-1}$ errors in a layer controlled by Alice to advance it incorrectly to a layer controlled by Bob.

As the errors for the both the parties are accounted for separately, Bob cannot detect and correct these errors, and the protocol will run as if there were no errors till the next time Alice gets to speak. This happens only after $\mathfrak{C}_{i+1} - \mathfrak{C}_i$ rounds, which due to our definition of $\mathfrak{C}$, may be a factor of c larger than $\mathfrak{C}_i - \mathfrak{C}_{i-1}$, the total number of errors inserted by the adversary. Thus, it is possible for the adversary to waste c rounds of simulation per error, implying that there has to be a multiplicative $1/c$ loss in the overall error resilience. Moreover, c cannot be too small, as then the parties will not be able to correctly maintain their meeting points. We are able to show that this loss is a constant proportional to the total number of meeting points; see the formal version of our main theorem (Theorem 8).

3 Model and Preliminaries

We now describe the model and our result more formally. Note that our description closely follows [7].

3.1 Circuits and dag-protocols

Circuits

We will consider Boolean circuits with negations at the inputs. That is, a Boolean circuit will be a directed acyclic graph C where each non-source node (or gate) is labelled as an $\wedge$ gate or an $\vee$ gate, while each source node is labelled with an input variable or the negation of the input variable. The source nodes are also called input gates and one designated node in C is called the output gate. We use $V_\wedge$ and $V_\vee$ to denote the set of all the nodes in C that are labelled $\wedge$ and $\vee$ respectively. We define how computation over a circuit C takes place by inductively defining the functions $f_{C,v}$ computed at each gate v. We have:

$$f_{C,v} = \begin{cases} z, & \text{if } v \text{ is an input gate with label } z \\ \bigvee_{u:(u,v) \text{ is an edge in } C} f_{C,u}, & \text{if } v \in V_\vee \\ \bigwedge_{u:(u,v) \text{ is an edge in } C} f_{C,u}, & \text{if } v \in V_\wedge \end{cases} . \tag{1}$$

When v is the output gate, we omit it from the notation and simply write f_C, which we call the function computed by C. We will assume that every gate in C that is not an input gate has at least one edge coming to it. We define the size of C to be the number of nodes in C and the depth of C to be the length of the longest path in C starting from an input gate and ending at the output gate. We use $\|C\|$ to denote the size of C.

Note that edges in C are going "up" from the inputs to the output gate. This contrasts with the edges in **dag**-protocols that are going "down" from the root to the leaves as defined below. These two directions, although different, follow the conventions for both circuits and **dag**-protocols.

Search problems and KW-games

Let $\mathcal{X}$ and $\mathcal{Y}$ be input sets for Alice and Bob respectively and $\mathcal{O}$ be a set of outputs. A search problem on these input and output sets is defined by a relation $S \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{O}$, and the goal of Alice and Bob is to determine, given inputs $x \in \mathcal{X}$ and $y \in \mathcal{Y}$ respectively, an element $o \in \mathcal{O}$ such that $(x, y, o) \in S$.

KW-games are special type of search problems that are defined using a Boolean function f on n variables (for some n). Here, the set $\mathcal{X}$ is the set $f^{-1}(1)$ of all inputs for which f evaluates to 1, $\mathcal{Y}$ is the set $f^{-1}(0)$ of all inputs for which f evaluates to 0, and $\mathcal{O}$ is the set $[n]$. As $\mathcal{X}$ and $\mathcal{Y}$ are disjoint by definition, for any $x \in \mathcal{X}$, $y \in \mathcal{Y}$, there exists $o \in [n]$ such that $x_o \neq y_o$. The search problem KW_f is the problem of finding such an o, namely, it is the subset:

$$\mathsf{KW}_f = \{(x, y, o) \in \mathcal{X} \times \mathcal{Y} \times \mathcal{O} \mid x_o \neq y_o\}.$$

dag-protocols

We now define the notion of **dag**-protocols. Let $\mathcal{X}$, $\mathcal{Y}$, and $\mathcal{O}$ be input and output sets as above. A **dag**-protocol is defined by a tuple:

$$\Pi = \left(G = (V_A \cup V_B \cup V_O, E), \mathsf{rt}, \{h_v\}_{v \in V_A \cup V_B}, \{o_v\}_{v \in V_O}\right).$$

Here, G is a directed acyclic graph with vertices partitioned into V_A, V_B, and V_O respectively, and edges E. We will use $V = V_A \cup V_B \cup V_O$ to be the set of all vertices in G and $\mathsf{rt} \in V$ is a special vertex called the root. For all $v \in V_A$, the function h_v is a "message function" that maps Alice's input set $\mathcal{X}$ to a vertex in V that v has an edge to. Similarly, for all $v \in V_B$, the function h_v maps Bob's input set $\mathcal{Y}$ to a vertex in V that v has an edge to. We often conflate the two and write $h_v : \mathcal{X} \times \mathcal{Y} \to V$ with the understanding that the second argument is redundant if $v \in V_A$ and the first argument is redundant if $v \in V_B$. Finally, for all $v \in V_O$, the value $o_v \in \mathcal{O}$ is the output value for the vertex v. We will assume that vertices in V_O do not have any out-edges, all other vertices have at least one outgoing edge (as otherwise the message functions will not be defined) and rt does not have any incoming edges. We define the size of Π to be $|V|$, the number of nodes in Π, and the depth of Π to be the length of the longest path in G starting from rt. We use $\|\Pi\|$ to denote the size of Π.

As mentioned above, edges in a **dag**-protocol go "down" from the root to the leaves, and are thus opposite to our convention for circuits. We are explicit about the direction in our exposition whenever there is room for ambiguity.

Execution of a dag-protocol

A **dag**-protocol Π as above is executed by start with inputs $x \in \mathcal{X}$, $y \in \mathcal{Y}$ and a vertex $v = \mathsf{rt}$, and continuing as follows: If $v \in V_A \cup V_B$ then the execution simply uses the message function and updates the current vertex to $h_v(x, y)$[9]. Otherwise, we have $v \in V_O$, and the execution terminates with the output o_v. As this output is determined by x and y, we will denote it using $\Pi(x, y)$. For a search problem $S \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{O}$, we say that Π solves S if for all $x \in \mathcal{X}$, $y \in \mathcal{Y}$, we have $(x, y, \Pi(x, y)) \in S$.

3.2 Rectangular Correctness & Equivalence of Circuits and dag-protocols

The models of Boolean circuits and **dag**-protocols described above have a deep connection. To understand this connection, we first define the notion of rectangular correctness for **dag**-protocols.

▶ **Definition 3** (Rectangles for **dag**-protocols). *Let Π be a* **dag***-protocol with input sets $\mathcal{X}$, $\mathcal{Y}$, and output set $\mathcal{O}$. We inductively define the (combinatorial) rectangle associated with each vertex. For the root rt, define the rectangle to $R_{\mathsf{rt}} = \mathcal{X} \times \mathcal{Y}$. For a node $v \neq \mathsf{rt}$, define the (combinatorial) rectangle $R_v \subseteq \mathcal{X} \times \mathcal{Y}$ to be the smallest rectangle containing $\bigcup_{u:(u,v)\in E}\{(x, y) \in R_u \mid h_u(x, y) = v\}$.*

For a combinatorial rectangle $R \subseteq \mathcal{X} \times \mathcal{Y}$, we will use R_A to denote the projection of the rectangle on the set $\mathcal{X}$ and R_B to denote the projection on the set $\mathcal{Y}$. As we defined R_v to be the smallest rectangle in Definition 3, we have the following observation.

▶ **Observation 4.** *Let Π be a* **dag***-protocol with input sets $\mathcal{X}$, $\mathcal{Y}$, and output set $\mathcal{O}$ and $\{R_v\}_{v\in V}$ be the associated rectangles. For all $v \neq \mathsf{rt}$ and all $x \in (R_v)_A$, there exists a $u \in V$ such that $(u, v) \in E$ and $y \in \mathcal{Y}$ such that $(x, y) \in R_u$ and $h_u(x, y) = v$. An analogous claim holds with the roles of x and y reversed.*

Additionally, observe that for all vertices v, the rectangle R_v contains all the pairs (x, y) such that the execution of Π goes to v on inputs x and y. Thus, the following notion of rectangular correctness is stronger than the "standard" notion of correctness.

▶ **Definition 5** (Rectangular Correctness). *Let Π be a* **dag***-protocol with inputs sets $\mathcal{X}$, $\mathcal{Y}$, and output set $\mathcal{O}$, and $\{R_v\}_{v\in V}$ be the associated rectangles. For a search problem $S \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{O}$, we say that Π solves S with rectangular correctness if for all $v \in V_O$ and all $(x, y) \in R_v$, we have $(x, y, o_v) \in S$.*

We are now ready to state the equivalence between circuits and **dag**-protocols formally. This theorem can also be found in [7]. We include a proof for completeness.

▶ **Theorem 6** ([25, 29]). *Let $f : \{0, 1\}^n \to \{0, 1\}$ be a Boolean function.*

1. *For any circuit C that computes f, there is a* **dag***-protocol Π with the same size and depth as C that solves KW_f with rectangular correctness.*
2. *For any* **dag***-protocol Π that solves KW_f with rectangular correctness, there is a circuit C with $\|C\| \leq \|\Pi\|$ that computes f.*

[9] Recall that $h_v(x, y)$ is independent of y if $v \in V_A$ and is independent of x if $v \in V_B$.

Proof. We only prove Item 1 as Item 2 is implied by Theorem 7 proved later. Fix a circuit C and define a **dag**-protocol Π for KW_f by setting the graph G to be the same graph as C except that the directions of all the edges are reversed. Let V_A be the set of all $\vee$ gates in C, V_B be the set of all the $\wedge$ gates in C and V_O be the set of all the inputs gates. Let rt be the output gate of C. For all $v \in V_O$, we define o_v to be the variable that (possibly after negation) is input at that gate. It remains to define the message transmission functions $\{h_v\}_{v \in V_A \cup V_B}$.

Fix $v \in V_A \cup V_B$ and recall that the input set $\mathcal{X} = f^{-1}(1)$ and $\mathcal{Y} = f^{-1}(0)$. For $x \in \mathcal{X}$ and $y \in \mathcal{Y}$, we define $h_v(x, y)$ as follows: If $v \in V_A$, then $h_v(x, y)$ is just a function of x. If $f_{C,v}(x) = 1$, then, as v is a $\vee$ gate[10], by Equation (1), there exists a neighbor u (in-neighbor in C and out-neighbor in G) of v such that $f_{C,u}(x) = 1$, and we set $h_v(x, y)$ to be the lexicographically smallest such neighbor u. Otherwise, we set $h_v(x, y)$ to be the lexicographically smallest neighbor u of v. Observe that such a neighbor always exists as v is not an input gate. Similarly, if $v \in V_B$, then $h_v(x, y)$ is just a function of y. If $f_{C,v}(y) = 0$, then, as v is a $\wedge$ gate, by Equation (1), there exists a neighbor u (in-neighbor in C and out-neighbor in G) of v such that $f_{C,u}(y) = 0$, and we set $h_v(x, y)$ to be the lexicographically smallest such neighbor u. Otherwise, we set $h_v(x, y)$ to be the lexicographically smallest neighbor u of v. As before, such a neighbor always exists as v is not an input gate.

We clearly have $\|\Pi\| = \|C\|$ and we just have to show that Π solves KW_f with rectangular correctness. By Definition 5, we have to show that, if $\{R_v\}_{v \in V}$ are the rectangles associated with Π as in Definition 3, then, for all $v \in V_O$ and all $(x, y) \in R_v$, we have $(x, y, o_v) \in \mathsf{KW}_f$. In fact, we will show the stronger statement that for all $v \in V$ and all $(x, y) \in R_v$, we have:

$$f_{C,v}(x) = 1 \qquad \text{and} \qquad f_{C,v}(y) = 0.$$

This is indeed stronger, as if $v \in V_O$, then Equation (1) says that $f_{C,v}$ is just the (possibly negated) input variable o_v implying that $(x, y, o_v) \in \mathsf{KW}_f$. We show this by induction starting from the root rt downwards to the nodes reachable from it[11]. For the root rt, we use the fact that it corresponds to the output gate of C implying that $f_{C,\mathsf{rt}} = f$. The result follows as $\mathcal{X} = f^{-1}(1)$ and $\mathcal{Y} = f^{-1}(0)$.

We now show it for $v \neq \mathsf{rt}$ assuming it holds for all u that have edges to v in Π (equivalently, all u that v has an edge to in the circuit C). Let $(x, y) \in R_v$ be arbitrary. We only show that $f_{C,v}(x) = 1$ as the proof that $f_{C,v}(y) = 0$ is analogous. As $(x, y) \in R_v$, we have by Observation 4 that there exists u that has an edge to v in Π and a $y' \in \mathcal{Y}$ such that $(x, y') \in R_u$ and $h_u(x, y') = v$. Applying the induction hypothesis on u, it follows that $f_{C,u}(x) = 1$. Now, if $u \in V_\wedge$, then the fact that v has an edge to u in C together with Equation (1) implies that $f_{C,v}(x) = 1$, as desired. On the other hand, if $u \in V_\vee$, then, use the fact that $h_u(x, y') = v$ to get that $f_{C,v}(x) = 1$, finishing the proof. ◀

3.3 Error Models

Circuits

We consider short-circuit errors. Let C be a Boolean circuit as above. An error pattern for C is defined by a function $e : V \to V \cup \{*\}$ satisfying the property that $e(v)$ maps v to an in-neighbor of v, *i.e.* to a neighbor of v that is "below" it, or to $*$, for all $v \in V$. In particular,

[10] We abuse notation and use v to denote both the vertices in G and the corresponding vertices in C.
[11] Observe that if v is not reachable from rt, then Definition 3 says that $R_v = \emptyset$ and there is nothing to show.

if $v \notin V_\wedge \cup V_\vee$, we have $e(v) = *$. It will often be convenient to separate e into two functions and write $e = (e_\vee, e_\wedge)$, where $e_\vee$ is the function e restricted to the vertices in $V_\vee$ and $e_\wedge$ is the function e restricted to the vertices in $V_\wedge$. Let C be a circuit and e be an error pattern for C. Let $v \in C$ be a gate. We define how computation takes place in C in the presence of errors e by inductively defining the functions $f_{C,e,v}$ computed at each gate v. We have[12]:

$$f_{C,e,v} = \begin{cases} z, & \text{if } v \text{ is an input gate with label } z \\ \bigvee_{u:(u,v) \text{ is an edge in } C} f_{C,e,u}, & \text{if } v \in V_\vee \text{ and } e(v) = * \\ \bigwedge_{u:(u,v) \text{ is an edge in } C} f_{C,e,u}, & \text{if } v \in V_\wedge \text{ and } e(v) = * \\ f_{C,e,e(v)}, & \text{if } e(v) \neq * \end{cases}. \tag{2}$$

We omit v from the notation if v is the output gate of C. Let C be a circuit and $\mathcal{E}$ be a set of error patterns for C. For a function f, we say that C computes f despite $\mathcal{E}$, if for all $e \in \mathcal{E}$, we have $f_{C,e} = f$. We say that $\mathcal{E}$ is rectangular if viewing every $e \in \mathcal{E}$ as a pair $e = (e_\vee, e_\wedge)$ gives a combinatorial rectangle. When this happens, we use $\mathcal{E}_\vee$ to denote the projection of this rectangle on the first coordinate and $\mathcal{E}_\wedge$ to denote the projection on the second coordinate (thus, $\mathcal{E} = \mathcal{E}_\vee \times \mathcal{E}_\wedge$)

We now define the set of error patterns that we work with. Let $\Theta \geq 0$ be an integer parameter. For a circuit C, define the set $\mathcal{E}_\Theta(C)$ to be the set of all error patterns e such that for any path in C that starts at an input gate and ends the output gate, at most Θ gates $v \in V_\wedge$ on the path satisfy $e(v) \neq *$ and at most Θ gates $v \in V_\vee$ on the path satisfy $e(v) \neq *$. Observe that the set $\mathcal{E}_\Theta(C)$ is rectangular for all $\Theta \geq 0$.

dag-protocols

Let Π be a **dag**-protocol as above. An error pattern ξ for Π is defined by a function $\xi : V \to V \cup \{*\}$ satisfying the property that $\xi(v)$ maps v to an out-neighbor of v, *i.e.* to a neighbor of v that is "below" it[13], or to $*$. In particular, if $v \in V_O$, we have $\xi(v) = *$. It will often be convenient to separate ξ into two functions $\xi = (\xi_A, \xi_B)$, where ξ_A is the function ξ restricted to the vertices in V_A, and ξ_B is the function ξ restricted to the vertices in V_B. The error pattern ξ affects the execution of Π as follows: If the current vertex is v and $\xi(v) = *$, the execution proceeds as before. On the other hand, if $\xi(v) \neq *$ (observe that this can only happen if $v \in V_A \cup V_B$), the execution updates the current vertex to $\xi(v)$ instead of updating using the function h_v as before.

Let Π be a **dag**-protocol and Ξ be a set of error patterns for Π. Similarly to before, we say that Ξ is rectangular if viewing every $\xi \in \Xi$ as a pair $\xi = (\xi_A, \xi_B)$ gives a combinatorial rectangle. When this happens, we use Ξ_A to denote the projection of this rectangle on the first coordinate and Ξ_B to denote the projection on the second coordinate (thus, $\Xi = \Xi_A \times \Xi_B$). For a rectangular Ξ, define the **dag**-protocol Π_Ξ to be the same as Π except that the input sets are now $\mathcal{X} \times \Xi_A$ and $\mathcal{Y} \times \Xi_B$ and the message functions $h_{\Xi,v}$ are defined as (recall that $\xi = (\xi_A, \xi_B)$):

$$h_{\Xi,v}((x, \xi_A), (y, \xi_B)) = \begin{cases} h_v(x, y), & \text{if } \xi(v) = * \\ \xi(v), & \text{if } \xi(v) \neq * \end{cases}. \tag{3}$$

[12] Recall that the label of an input gate is either an input variable or the negation of an input variable.
[13] Recall that our conventions for the directions of edges in circuits is opposite that in **dag**-protocols.

For a search problem $S \subseteq \mathcal{X} \times \mathcal{Y} \times \mathcal{O}$ and a rectangular Ξ, we say that Π solves S with rectangular correctness despite Ξ if Π_Ξ solves S_Ξ with rectangular correctness, where $S_\Xi \subseteq (\mathcal{X} \times \Xi_A) \times (\mathcal{Y} \times \Xi_B) \times \mathcal{O}$ is the search problem satisfying $((x, \xi_A), (y, \xi_B), o) \in S_\Xi \iff (x, y, o) \in S$ for all values of x, y, o, ξ_A, ξ_B.

We now define the set of error patterns that we work with. Let $\Theta \geq 0$ be an integer parameter. For a **dag**-protocol Π, define the set $\Xi_\Theta(\Pi)$ to be the set of all error patterns ξ such that for any path in G that starts at rt and ends at a node in V_O, at most Θ nodes $v \in V_A$ on the path satisfy $\xi(v) \neq *$ and at most Θ nodes $v \in V_B$ on the path satisfy $\xi(v) \neq *$. Observe that the set $\Xi_\Theta(\Pi)$ is rectangular for all $\Theta \geq 0$. For convenience, we will often abbreviate $\Pi_{\Xi_\Theta(\Pi)}$ to Π_Θ.

3.4 Connecting the Error Models

We now show that error resilient **dag**-protocols are stronger than error-resilient circuits. Observe that the theorem below implies Item 2 of Theorem 6 as can be seen by setting $\Theta = 0$. Additionally, note that the theorem below differs from the analogous theorem in [7] in that the number of allowed errors Θ is the same for every path, while in [7], it was a constant fraction of the length of the path. Having it as a constant fraction creates some minor complications while trimming the "empty edges" (in the first paragraph of the proof), that were overlooked in [7]. These complications do not affect the rest of the proof of [7], as [7] actually works even when the allowed number of errors is the same for every path (and equals a constant fraction of the length of the *longest* path). Nonetheless, we flesh out the details of trimming for our model formally in the full version. We also mention that having the same number of errors on all the paths, irrespective of their length, is a stronger error model than having the number of errors on each path be proportional to its length.

▶ **Theorem 7.** *Let $\Theta \geq 0$ be given and $f : \{0,1\}^n \to \{0,1\}$ be a Boolean function. For any* **dag***-protocol Π that solves KW_f with rectangular correctness despite $\Xi_\Theta(\Pi)$, there is a circuit C with size and depth at most that of Π that computes f despite $\mathcal{E}_\Theta(C)$.*

Proof. Fix a protocol Π that solves KW_f with rectangular correctness despite $\Xi_\Theta(\Pi)$. We can assume without loss of generality that there are no empty edges in Π, as defined in the full version. Equivalently, if $\{R_{\Theta,v}\}_{v \in V}$ are the rectangles associated with Π_Θ according to Definition 3 and $\{h_{\Theta,v}\}_{v \in V_A \cup V_B}$ are the message functions, for all edges $e = (u, v) \in E$ we have:

$$R_{\Theta,u} = \emptyset \quad \text{or} \quad \exists ((x, \xi_A), (y, \xi_B)) \in R_{\Theta,u} : \qquad h_{\Theta,u}((x, \xi_A), (y, \xi_B)) = v. \tag{4}$$

As Π that solves KW_f with rectangular correctness despite $\Xi_\Theta(\Pi)$, we have that Π_Θ solves $\mathsf{KW}_{f,\Xi_\Theta(\Pi)}$ with rectangular correctness. From Definition 5, we get that for all $v \in V_O$ and all $((x, \xi_A), (y, \xi_B)) \in R_{\Theta,v}$, we have $((x, \xi_A), (y, \xi_B), o_v) \in \mathsf{KW}_{f,\Xi_\Theta(\Pi)}$. Simplifying using the definition of $\mathsf{KW}_{f,\Xi_\Theta(\Pi)}$, we get:

$$\forall v \in V_O, ((x, \xi_A), (y, \xi_B)) \in R_{\Theta,v} : \quad x_{o_v} \neq y_{o_v}. \tag{5}$$

We now define the circuit C. The gates in the circuit C are exactly the nodes in V and we abuse notation slightly by using u, v *etc.* to refer to both. For any edge $(u, v) \in E$, we add the flipped edge (v, u) to the circuit C. We define all gates in V_A to be $\vee$ gates and all gates in V_B to be $\wedge$ gates while rt is chosen as the output gate. The remaining gates are in V_O are input gates whose labels are defined next. As $R_{\Theta,v}$ in Equation (5) is a combinatorial rectangle, we get that there exists $b \in \{0,1\}$ such that

$$\forall ((x, \xi_A), (y, \xi_B)) \in R_{\Theta,v} : \quad x_{o_v} = b \quad \text{and} \quad y_{o_v} = \overline{b}.$$

If $b = 1$, we label the input gate v by the variable o_v unnegated, and we label it by the negated variable $\overline{o_v}$ otherwise. This finishes the description of C. As the claim about the size $\|C\|$ is straightforward, we focus on showing that C computes f despite $\mathcal{E}_\Theta(C)$. As both C and Π have the same graph upto the directions of the edges, observe that any error pattern for C can also be seen as a pattern for Π and that $\mathcal{E}_\Theta(C) = \Xi_\Theta(\Pi)$. Because of this, we interpret any error pattern $\xi = (\xi_A, \xi_B)$ as an error pattern $e = (e_\vee, e_\wedge)$ for C and we ξ and subscripts A and B to refer to both. We also define the error patterns $*_A$ and $*_B$ to be those that map all gates in V_A (equivalently, $V_\vee$) and V_B (equivalently, $V_\wedge$) respectively to $*$. To show that C computes f despite $\mathcal{E}_\Theta(C)$, we show inductively that for all $u \in V$, we have:

$$\forall((x,\xi_A),(y,\xi_B)) \in R_{\Theta,u}: \quad f_{C,(\xi_A,*_B),u}(x) = 1 \quad \text{and} \quad f_{C,(*_A,\xi_B),u}(y) = 0. \tag{6}$$

This suffices as plugging $v = \mathsf{rt}$ and using the definition of $R_{\Theta,\mathsf{rt}} = \left(f^{-1}(1) \times (\Xi_\Theta(\Pi))_A\right) \times \left(f^{-1}(0) \times (\Xi_\Theta(\Pi))_B\right)$ from Definition 3, we get that:

$$\forall((x,\xi_A),(y,\xi_B)) \in R_{\Theta,\mathsf{rt}}: \quad f_{C,(\xi_A,*_B)}(x) = f(x) \quad \text{and} \quad f_{C,(*_A,\xi_B)}(y) = f(y).$$

As short-circuiting a gate in $V_\wedge$ can never change the output from 1 to 0 and short-circuiting a gate in $V_\vee$ can never change the output from 0 to 1, we get that:

$$\forall((x,\xi_A),(y,\xi_B)) \in R_{\Theta,\mathsf{rt}}: \quad f_{C,\xi}(x) = f(x) \quad \text{and} \quad f_{C,\xi}(y) = f(y).$$

Thus, we get that $f_{C,\xi} = f$ for all $\xi \in \Xi_\Theta(\Pi)$. As $\mathcal{E}_\Theta(C) = \Xi_\Theta(\Pi)$, we get that C computes f despite $\mathcal{E}_\Theta(C)$, as desired. Having shown that Equation (6) is sufficient, we now prove that Equation (6) holds by induction. For the base case, when $u \in V_O$, this follows by our labels and Equations (2) and (5). We show the results for $u \notin V_O$ assuming it holds for all v that have edges to u in C. If $R_{\Theta,u} = \emptyset$, there is nothing to show, so we assume otherwise. As the proof in the other case is similar, assume without loss of generality that $u \in V_A$. From Equation (4), we get that for all v that have edges to u in C, we have that there exists $((x,\xi_A),(y,\xi_B)) \in R_{\Theta,u}$ such that $h_{\Theta,u}((x,\xi_A),(y,\xi_B)) = v$. As $u \in V_A$, the message function is determined by the first argument and we get from Definition 3 that:

$$(R_{\Theta,u})_A \subseteq \bigcup_{v:(v,u) \text{ is an edge in } C} (R_{\Theta,v})_A \quad \text{and} \quad (R_{\Theta,u})_B \subseteq \bigcap_{v:(v,u) \text{ is an edge in } C} (R_{\Theta,v})_B. \tag{7}$$

We now show Equation (6). Fix $((x,\xi_A),(y,\xi_B)) \in R_{\Theta,u}$.

- **Showing that $f_{C,(\xi_A,*_B),u}(x) = 1$ when $\xi(u) \neq *$:** By Equation (3), we get that $h_{\Theta,u}((x,\xi_A),(y,\xi_B)) = \xi(u)$ implying from Definition 3 that $((x,\xi_A),(y,\xi_B)) \in R_{\Theta,\xi(u)}$. By the induction hypothesis on $\xi(u)$, we get that $f_{C,(\xi_A,*_B),\xi(u)}(x) = 1$. From Equation (2), we get:

$$f_{C,(\xi_A,*_B),u}(x) = f_{C,(\xi_A,*_B),\xi(u)}(x) = 1.$$

- **Showing that $f_{C,(\xi_A,*_B),u}(x) = 1$ when $\xi(u) = *$:** We have from $((x,\xi_A),(y,\xi_B)) \in R_{\Theta,u}$ that $(x,\xi_A) \in (R_{\Theta,u})_A$. From Equation (7), there exists an in-neighbor v' of u in C and $(y_{v'},\xi_{B,v'})$ such that $((x,\xi_A),(y_{v'},\xi_{B,v'})) \in R_{\Theta,v'}$. By the induction hypothesis on v', we have $f_{C,(\xi_A,*_B),v'}(x) = 1$. From Equation (2), we get:

$$f_{C,(\xi_A,*_B),u}(x) = \bigvee_{v:(v,u) \text{ is an edge in } C} f_{C,(\xi_A,*_B),v}(x) = 1.$$

- **Showing that $f_{C,(*_A,\xi_B),u}(y) = 0$:** We have from $((x,\xi_A),(y,\xi_B)) \in R_{\Theta,u}$ that $(y,\xi_B) \in (R_{\Theta,u})_B$. It follows from Equation (7) that for all in-neighbors v of u in C, we have that there exists $(x_v,\xi_{A,v})$ such that $((x_v,\xi_{A,v}),(y,\xi_B)) \in R_{\Theta,v}$. From the induction hypothesis on v, we get that $f_{C,(*_A,\xi_B),v}(y) = 0$ for all such v. From Equation (2), we get:

$$f_{C,(*_A,\xi_B),u}(y) = \bigvee_{v:(v,u) \text{ is an edge in } C} f_{C,(*_A,\xi_B),v}(y) = 0. \qquad \blacktriangleleft$$

3.5 Our Results

We are now ready to state our main result formally. We show that:

▶ **Theorem 8** (Formal version of Theorem 1). *Let $c > 1$, $n > 0$ be integers and $f : \{0,1\}^n \to \{0,1\}$ be a Boolean function. Let C be a Boolean circuit of size $\|C\| = s$ and depth*[14] *$d < (\log n)^{c-1}$ that computes f. There exists $\Theta \leq s^{\mathcal{O}(c \log c)}$ and a Boolean circuit C' with size at most $s^{\mathcal{O}(c \log c)}$ and depth $\mathcal{O}(c) \cdot \Theta$ that computes f despite $\mathcal{E}_\Theta(C')$.*

To prove Theorem 8, we need the following result about **dag**-protocols, which forms the technical core of Theorem 8.

▶ **Theorem 9** (Formal version of Theorem 2). *Let S be a search problem and Π be a* **dag***-protocol of size s and depth d that solves S with rectangular correctness. Let $c > 1$ be an integer satisfying*[15] *$d < (\log s)^{c-1}$. There exists $\Theta \leq s^{\mathcal{O}(c \log c)}$ (defined in the full version) and a* **dag***-protocol Π' (as defined in the full version) with size $s' = s^{\mathcal{O}(c \log c)}$ and depth $\mathcal{O}(c) \cdot \Theta$ that solves S with rectangular correctness despite $\Xi_\Theta(\Pi')$.*

We prove Theorem 9 in the full version, but use it here to prove Theorem 8.

Proof of Theorem 8 assuming Theorem 9. Fix a circuit C as in the theorem statement. Applying Theorem 6, we get that there exists a **dag**-protocol Π with size s and depth d that solves KW_f with rectangular correctness. Applying Theorem 9, we get that there exists $\Theta \leq s^{\mathcal{O}(c \log c)}$ and a **dag**-protocol Π' with size $s' = s^{\mathcal{O}(c \log c)}$ and depth $\mathcal{O}(c) \cdot \Theta$ that solves KW_f with rectangular correctness despite $\Xi_\Theta(\Pi')$. Now, applying Theorem 7 on Π', we get that there exists a circuit C' with size at most $s^{\mathcal{O}(c \log c)}$ and depth at most $\mathcal{O}(c) \cdot \Theta$ that computes f despite $\mathcal{E}_\Theta(C')$, as desired. ◀

References

1. Mirza Ahad Baig, Suvradip Chakraborty, Stefan Dziembowski, Małgorzata Gałązka, Tomasz Lizurej, and Krzysztof Pietrzak. Efficiently testable circuits. In *14th Innovations in Theoretical Computer Science Conference (ITCS 2023)*, 2023. doi:h10.4230/LIPIcs.ITCS.2023.10.
2. Mirza Ahad Baig, Suvradip Chakraborty, Stefan Dziembowski, Małgorzata Gałązka, Tomasz Lizurej, and Krzysztof Pietrzak. Efficiently testable circuits without conductivity. In *Theory of Cryptography Conference*, 2023. doi:10.1007/978-3-031-48621-0_5.
3. Maria Luisa Bonet and Samuel R. Buss. Size-depth tradeoffs for boolean fomulae. *Information Processing Letters*, 49(3):151–155, 1994. doi:10.1016/0020-0190(94)90093-0.
4. Mark Braverman, Klim Efremenko, Ran Gelles, and Michael A. Yitayew. Optimal short-circuit resilient formulas. In *Computational Complexity Conference (CCC)*, volume 137, pages 10:1–10:22, 2019. doi:10.4230/LIPICS.CCC.2019.10.

[14] The "−1" in this expression is for technical convenience.
[15] The "−1" in this expression is for technical convenience.

5 T.-H. Hubert Chan, Zhibin Liang, Antigoni Polychroniadou, and Elaine Shi. Small memory robust simulation of client-server interactive protocols over oblivious noisy channels. In *Symposium on Discrete Algorithms (SODA)*, pages 2349–2365, 2020. doi:10.1137/1.9781611975994.144.

6 Roland L'vovich Dobrushin and SI Ortyukov. Upper bound on the redundancy of self-correcting arrangements of unreliable functional elements. *Problemy Peredachi Informatsii*, 13(3):56–76, 1977.

7 Klim Efremenko, Bernhard Haeupler, Yael Tauman Kalai, Pritish Kamath, Gillat Kol, Nicolas Resch, and Raghuvansh R. Saxena. Circuits resilient to short-circuit errors. In *54th Annual ACM SIGACT Symposium on Theory of Computing, Rome, Italy, June 20 - 24, 2022*, 2022.

8 William S. Evans and Nicholas Pippenger. On the maximum tolerable noise for reliable computation by formulas. *IEEE Transactions on Information Theory*, 44(3):1299–1305, 1998. doi:10.1109/18.669417.

9 William S. Evans and Leonard J. Schulman. Signal propagation and noisy circuits. *IEEE Transactions on Information Theory*, 45(7):2367–2373, 1999. doi:10.1109/18.796377.

10 William S. Evans and Leonard J. Schulman. On the maximum tolerable noise of k-input gates for reliable computation by formulas. *IEEE Transactions on Information Theory*, 49:3094–3098, 2003. doi:10.1109/TIT.2003.818405.

11 Tomás Feder. Reliable computation by networks in the presence of noise. *IEEE Transactions on Information Theory*, 35(3):569–571, 1989. doi:10.1109/18.30978.

12 Péter Gács and Anna Gál. Lower bounds for the complexity of reliable boolean circuits with noisy gates. *IEEE Transactions on Information Theory*, 40(2):579–583, 1994. doi:10.1109/18.312190.

13 Anna Gál. Lower bounds for the complexity of reliable boolean circuits with noisy gates. In *Foundations of Computer Science (FOCS)*, pages 594–601, 1991. doi:10.1109/SFCS.1991.185424.

14 Anna Gál and Mario Szegedy. Fault tolerant circuits and probabilistically checkable proofs. In *Proceedings of Structure in Complexity Theory. Tenth Annual IEEE Conference*, pages 65–73, 1995. doi:10.1109/SCT.1995.514728.

15 Ankit Garg, Mika Göös, Pritish Kamath, and Dmitry Sokolov. Monotone circuit lower bounds from resolution. In *Symposium on Theory of Computing (STOC)*, pages 902–911, 2018. doi:10.1145/3188745.3188838.

16 Ran Gelles. Coding for interactive communication: A survey. *Foundations and Trends in Theoretical Computer Science*, 13(1–2):1–157, 2017. doi:10.1561/0400000079.

17 Bruce E. Hajek and Timothy Weller. On the maximum tolerable noise for reliable computation by formulas. *IEEE Transactions on Information Theory*, 37(2):388–391, 1991. doi:10.1109/18.75261.

18 Yael Tauman Kalai, Allison B. Lewko, and Anup Rao. Formulas resilient to short-circuit errors. In *Foundations of Computer Science (FOCS)*, pages 490–499, 2012. doi:10.1109/FOCS.2012.69.

19 Mauricio Karchmer and Avi Wigderson. Monotone circuits for connectivity require super-logarithmic depth. In *Symposium on Theory of Computing (STOC)*, pages 539–550, 1988. doi:10.1145/62212.62265.

20 Daniel J. Kleitman, Frank Thomson Leighton, and Yuan Ma. On the design of reliable boolean circuits that contain partially unreliable gates. *Journal of Computer and System Sciences*, 55(3):385–401, 1997. doi:10.1006/JCSS.1997.1531.

21 Jan Krajícek. Interpolation theorems, lower bounds for proof systems, and independence results for bounded arithmetic. *Journal of Symbolic Logic*, 62(2):457–486, 1997. doi:10.2307/2275541.

22 Nicholas Pippenger. On networks of noisy gates. In *Foundations of Computer Science (FOCS)*, pages 30–38, 1985. doi:10.1109/SFCS.1985.41.

23 Nicholas Pippenger. Reliable computation by formulas in the presence of noise. *IEEE Transactions on Information Theory*, 34(2):194–197, 1988. doi:10.1109/18.2628.

24 Pavel Pudlák. On extracting computations from propositional proofs (a survey). In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS 2010)*, volume 8, pages 30–41, 2010. doi:10.4230/LIPICS.FSTTCS.2010.30.

25 Alexander Razborov. Unprovability of lower bounds on circuit size in certain fragments of bounded arithmetic. *Izvestiya of the RAN*, pages 201–224, 1995.

26 Leonard J Schulman. Communication on noisy channels: A coding theorem for computation. In *Foundations of Computer Science (FOCS)*, pages 724–733. IEEE, 1992. doi:10.1109/SFCS.1992.267778.

27 Leonard J. Schulman. Deterministic coding for interactive communication. In *Symposium on Theory of Computing (STOC)*, pages 747–756, 1993. doi:10.1145/167088.167279.

28 Leonard J Schulman. Coding for interactive communication. *IEEE Transactions on Information Theory*, 42(6):1745–1756, 1996. doi:10.1109/18.556671.

29 Dmitry Sokolov. Dag-like communication and its applications. In *Proceedings of the 12th Computer Science Symposium in Russia (CSR)*, pages 294–307. Springer, 2017. doi:10.1007/978-3-319-58747-9_26.

30 John von Neumann. Probabilistic logics and the synthesis of reliable organisms from unreliable components. *Automata Studies*, pages 43–98, 1956.