Detecting crosstalk errors in quantum information processors

Mohan Sarovar, Timothy Proctor, Kenneth Rudinger, Kevin Young, Erik Nielsen, and Robin Blume-Kohout

Quantum Performance Laboratory, Sandia National Laboratories, Albuquerque, NM 87185 and Livermore, CA 94550

Crosstalk occurs in most quantum computing systems with more than one qubit. It can cause a variety of correlated and nonlocal *crosstalk errors* that can be especially harmful to fault-tolerant quantum error correction, which generally relies on errors being local and relatively predictable. Mitigating crosstalk errors requires understanding, modeling, and detecting them. In this paper, we introduce a comprehensive framework for crosstalk errors and a protocol for detecting and localizing them. We give a rigorous definition of crosstalk errors that captures a wide range of disparate physical phenomena that have been called "crosstalk", and a concrete model for crosstalk-free quantum processors. Errors that violate this model are crosstalk errors. Next, we give an equivalent but purely operational (model-independent) definition of crosstalk errors. Using this definition, we construct a protocol for detecting a large class of crosstalk errors in a multi-qubit processor by finding conditional dependencies between observed experimental probabilities. It is highly efficient, in the sense that the number of unique experiments required scales at most cubically, and very often quadratically, with the number of qubits. We demonstrate the protocol using simulations of 2-qubit and 6-qubit processors.

1 Introduction

Quantum computing has grown from a theoretical concept into a nascent technology. Cloud-accessible quantum information processors (QIPs) with 20+ qubits exist today, and ones with around 100 qubits may appear in the next few years [45]. Fundamen-

Mohan Sarovar: mnsarov@sandia.gov

tal operations – gates, state preparation and measurements (SPAM) – are approaching the demanding error rates required by the theory of fault-tolerance on a number of physical platforms, including superconducting qubits and trapped ions [29]. However, as experimentalists and engineers have begun to build systems of 10-20 qubits, it is becoming clear that emergent failure modes may be an even bigger problem than errors in elementary operations. The most obvious failure mode that emerges at scale is *crosstalk*.

"Crosstalk" describes a wide range of physical phenomena that vary significantly across physical platforms used for quantum computing. We will focus, instead, on the visible *effects* of crosstalk on the quantum logical behavior of a physical system that is used and treated like a quantum computer. We refer to these hardware-agnostic effects as *crosstalk errors* – deviations from the ideal behavior of quantum gates and circuits, which can be formalized and captured in an architecture-independent way. Crosstalk errors violate either of two key assumptions that go into any well-behaved model of QIP dynamics: spatial locality, and independence of operations. Gates and other operations are supposed to act non-trivially only in a specific "target" region of the QIP, and their action on that region is supposed to be independent of the context in which they are applied. These assumptions enable tractable models for quantum computing, and crosstalk errors violate them. Here, we give a rigorous definition of crosstalk errors that captures the effects of crosstalk, while avoiding the need to engage deeply with the physical phenomena themselves.

We begin in Sec. 2 and Sec. 3 by defining what it means for a quantum processor to be "crosstalk-free" at the quantum logic level. In Sec. 4, we construct an explicit error model for Markovian crosstalk-free behavior. Markovian dynamics that are not consistent with that model constitute crosstalk errors. Then in Sec. 5 we discuss the difficulty of detecting arbitrary unknown crosstalk errors and define a class of low-weight crosstalk errors that can be efficiently detected. In Sec. 6 and Sec. 7, we take an *operational* approach and show how to detect low-weight

crosstalk errors using only correlations between experimental variables – the settings and the outcomes of experiments. The protocol we develop specifies a set of at most $\tilde{\mathcal{O}}(n^3)$, and often $\tilde{\mathcal{O}}(n^2)$, experiments for detecting crosstalk on an n-qubit QIP. The analysis of the data from these experiments uses techniques adapted from causal inference on probabilistic graphical models [27, 52].

Much recent work has been published on detecting, quantifying, and modeling crosstalk and crosstalk errors in quantum computing devices [6, 42, 18, 44, 51, 39, 16, 23, 49, 50, 48, 17, 2, 19, 21, 11, 35]. Variants of Ramsey sequences have been used to detect and quantify coherent coupling between qubits [2]. This technique is very hardware-specific and typically limited to detecting crosstalk in the form of unwanted Hamiltonian couplings of known form. Several groups have also demonstrated mitigation of crosstalk in readout lines by detailed characterization and compensation [23, 49, 11, 35] (see also Supplementary Information in Refs. [6, 42, 16, 19, 21]). A very different approach, which is platform-independent and model-free like the work we present here, is the *simultaneous randomized benchmarking* (SRB) technique for detecting and quantifying crosstalk between pairs of qubits [18, 51]. The crosstalk detection protocol we present here is similar in motivation to SRB, and is meant to be used as a light-weight diagnostic for the presence of crosstalk. It is specifically designed to be run efficiently on many-qubit QIPs and identify the crosstalk structure (*i.e.*, which qubits have crosstalk errors between them), whereas we are not aware of an application of SRB that reveals crosstalk structure in a many-qubit QIP as efficiently. Moreover, our protocol is designed to detect a wide range of crosstalk errors and is more flexible in terms of the experiments that are performed, allowing it to be tailored towards detection of certain types of crosstalk errors. However, SRB has at least one clear advantage over our protocol; it measures the quantitative *rate* of certain crosstalk errors, whereas our protocol is just designed to detect and localize them, and has limited quantitative ability. Finally, we note that in a previous paper [50] we gave a protocol for detecting context dependence, including crosstalk, that can be seen as a precursor to the protocol given here.

2 Crosstalk and crosstalk errors

Before we embark on defining things precisely, a brief discussion of exactly *what* we are defining is apropos. In particular, the distinction between "crosstalk" and "crosstalk errors" needs further explanation.

Crosstalk is an imprecise but widely used term that appears primarily in electrical engineering and communication theory, and generally refers to "unwanted coupling between signal paths" [38]. In experimental quantum computing, the word has been adapted to describe a range of physical phenomena in which some subsystem of an experimental device – a qubit, field, control line, resonator, photodetector, etc. – unintentionally affects another subsystem.

A specific quantum computing device will generally display more than one such effect. For example, a transmon-based quantum processor might experience

- Residual coherent couplings between transmons that should be uncoupled,

- Traditional electromagnetic (EM) crosstalk between microwave lines,

- Stray on-chip EM fields due to imperfect microwave hygiene,

- Coupling between readout resonators attached to distinct qubits,

- 60Hz line noise that influences all the qubits.

Any and all of these phenomena could legitimately be termed crosstalk. All of them are architecture-specific; a trapped-ion processor would have its own endemic crosstalk effects, some analogous to these and some not.

Our goal is to understand and address crosstalk in a platform-independent way that facilitates comparisons between quantum processors without reference to the underlying physics. This is clearly inconsistent with the established use of the term crosstalk to describe specific physics phenomena. There is no reasonable direct comparison between an unwanted 2-transmon coupling (measured in MHz) and the intensity of a control laser spillover in a trapped-ion setup (measured in W/m^2). But we can legitimately compare their *effects* at the quantum logic level of abstraction, where each device is required to behave like a quantum computer, performing quantum logic gates and quantum circuits.

We introduce the new term "crosstalk errors" for this purpose. It means any observable effect at the quantum logic level (qubits, gates, quantum circuits, and their associated probabilities) that stems uniquely from some form of physical crosstalk. Some forms of physical crosstalk may result in purely local errors – *e.g.*, independent bit flips – at the quantum logic level; these are not crosstalk errors (despite their source) because they *could* have been produced by local noise. Similarly, if physical crosstalk exists but has no effect at the quantum logic level (perhaps because of intentional mitigation) then we say that the system is "crosstalk-free".

3 Definition of crosstalk errors

Crosstalk errors are undesired dynamics that violate either (or both) of two principles: locality and independence. In an ideal QIP, each qubit is completely isolated from the rest of the universe, and evolves independently of it, *except* when an operation is applied. Operations, including gates and measurements, couple qubits to other systems, such as external control fields and/or other qubits. This coupling is supposed to be precise and limited in scope.

Unfortunately, real QIPs are not ideal. They experience all manner of noise and errors. Of course, not all errors constitute crosstalk errors. Errors can cause deviations from ideal behavior yet still respect locality and independence. Unwanted dynamics that *do* violate locality or independence constitute crosstalk errors. We now make this precise by defining locality and independence.

> **Locality of operations**: A QIP has local operations if and only if the physical implementation of any quantum circuit does not create correlation between any qubits, or disjoint subsets of qubits, unless that circuit contains multiqubit operations that intentionally couple them.

If a processor obeys locality, then it makes sense to talk about the action of operations on their target qubits, and we can go further and define independence. If locality is violated, then operations do not necessarily have well-defined actions on their targets, and independence may not be well-defined.

> **Independence of local operations**: When an operation (gate, measurement, etc) appears in a quantum circuit acting on target qubits q at time t, the dynamical evolution of q at time t does not depend on what other operations (acting on disjoint qubits) appear in the circuit at the same time t.

> **Defintion 1:** A QIP's behavior is *crosstalk-free* if its behavior, when implementing arbitrary circuits, satisfies locality and independence.

4 An explicit error model for crosstalk-free processors

The definitions in the previous section are abstract. They neither rely upon nor define a concrete model for crosstalk errors or for crosstalk-free processors. In this section, we specialize to *Markovian* processors and construct an explicit model for crosstalk-free Markovian processors. By assuming Markovianity we are able to rule out many conceivable failures and define a model in which only finitely many things can go wrong. By defining crosstalk-free within this framework, we get a division of Markovian errors into crosstalk-free or *local, independent* errors, and everything else (*i.e.*, crosstalk errors).

4.1 Defining *crosstalk-free* for Markovian QIPs

We place Markovianity in context within a hierarchy of models for quantum hardware, based on increasing levels of *modularity* (see Fig. 1): *stable* quantum circuit, *Markovian* quantum circuit, and *Markovian, crosstalk-free* quantum circuit. We define each layer in this hierarchy in the following.

4.1.1 Stable QIPs

We call a QIP *stable* if every circuit's outcome probability distribution (over n-bit strings) is independent of external contexts [58]. Contexts on which these probabilities might depend include the time at which the circuit is run, the identity of the circuit that was run before it, or even the phase of the moon. Stability is the weakest notion of modularity: a stable QIP is modular only in the sense that its output distribution is independent of any external contexts, so that each circuit run on the QIP forms a "module". If a QIP is not stable, then modeling or probing its behavior becomes much more difficult. Importantly for this work, protocols for detecting crosstalk will likely be corrupted by this instability, and any results will be unreliable or inconclusive. Fortunately, explicit stability tests for QIPs can often be applied directly to data from other characterization protocols with only minimal modifications to the experiment design. For instance, by repeating a characterization
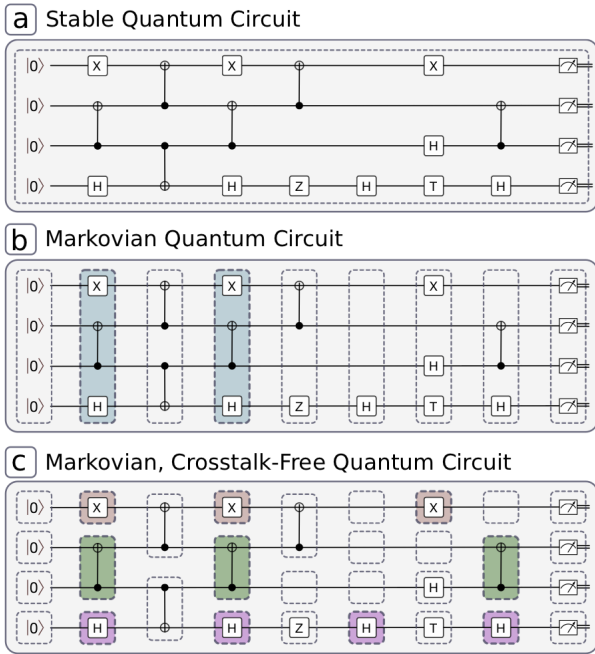
Figure 1: A hierarchy of modularity for QIPs. The dotted lines indicate the modular components in each layer of the hierarchy. (a) A quantum circuit is specified by a schedule of quantum gates on target qubits. It is *stable* if the associated measurement outcome distribution does not depend on any external context. (b) The dynamics of a stable QIP are *Markovian* if each layer in the circuit, including state preparation and measurement, can be represented by a CPTP map that depends only on the operations that comprise that layer, and not on any external context. For example, the two shaded circuit layers are identical and therefore must be represented by the same CPTP map. (c) The dynamics are Markovian and *crosstalk-free* if the gate operations are modular: the CPTP map describing a given circuit layer can be written as a tensor product of CPTP maps describing each of the component gates (locality), and these component maps do not depend on the other gates in the layer (independence). For example, each appearance of the shaded X, CNOT, or H gates must be represented by the same CPTP maps.

protocol in at least two different contexts, the techniques from Ref. [50] can verify whether the associated data sets are statistically consistent with one another. To check for the presence of drift, Ref. [46] applies Fourier analysis methods to the timestamped output data for each quantum circuit and checks if the resulting spectra are consistent with a stable measurement outcome distribution. Taken together, these approaches can help establish trust that a device is stable, or provide evidence that it is not.

4.1.2 Markovian QIPs

We need a stronger notion of modularity to predict how a QIP will perform on new quantum circuits that have not been run before. Circuits have a well-defined notion of time, which usually defines a natural division into consecutive *layers* [1] of parallel operations (gates, state preparations or measurements). See Fig. 1 for an example circuit with 9 layers that we notate $L_0, ..., L_8$. Operations within a single layer are effectively simultaneous. A layer is uniquely defined by the list of operations applied to each qubit during that layer, where "operations" can include idles, measurements, and initialiation/reset operations as well as elementary gates. Figure 1(b) shows a circuit partitioned into layers.

We call the QIP *Markovian* if we can describe and model each unique layer by a CPTP map acting on all n qubits in the system. We use a broad definition of CPTP map here, in which the input and output spaces need not be the same, and can include classical systems. Typically, an initialization operation is represented by a density matrix, which is a CPTP map from a trivial (1-dimensional) state space to a d^2-dimensional vector in quantum state space (Hilbert-Schmidt space); here $d = 2^n$. Elementary gates are represented by "square" CPTP maps from a quantum state space to itself. Terminating measurements are represented by POVMs, which are CPTP maps that map a d^2-dimensional quantum state to a d-outcome classical distribution. Intermediate measurements are represented by quantum instruments [15], which are CPTP maps that map a d^2-dimensional quantum state to a d^2-dimensional quantum state *and* a d-dimensional classical distribution. Layers involving multiple kinds of operations are represented by CPTP maps whose input and output spaces correspond to the tensor product of the input/output spaces of all the component operations.

There are many ways to violate the Markovian condition. For example, a layer might appear multiple times in the circuit, and act differently each time. But if a QIP is Markovian, then the CPTP map representing each layer depends *only* on the identity of the layer, not on any external context (*e.g.*, the time, or which layers occurred previously). Hereafter, we will only consider Markovian QIPs. The abstract definitions of locality and independence can presumably

---

[1]The word "cycle" (*i.e.*, clock cycle) has been used for the same concept. We use "layer" to describe a slice of a *circuit*, and reserve "cycle" to describe what a specific quantum processor actually *does* in a unit of time.

be instantiated in a non-Markovian model, but since no general model for non-Markovian processors is known we leave this for future work.

We use L to denote a CPTP map describing a given circuit layer. To specify *which* layer L describes, we will either index its position in the circuit or specify its component operations explicitly. For example, in Fig. 1(b), layers 1 and 3 (highlighted) are identical; each involves an X gate on qubit 0, a CNOT gate from qubit 2 to qubit 1, and a Hadamard gate on qubit 3. So we denote the CPTP map for this layer by:

$$\mathsf{L}_1 = \mathsf{L}_3 = \mathsf{L}(\mathrm{X}_0, \mathrm{CNOT}_{1 \leftarrow 2}, \mathrm{H}_3). \qquad (1)$$

The probabilities of the measurement outcomes for a quantum circuit are determined entirely by the CPTP maps describing the circuit's layers. For a depth-N circuit that begins with an initialization layer ρ, ends with a POVM measurement layer $\{\mathsf{M_i}\}$, and includes $N-2$ gate layers in the middle, the probability of the $\mathbf{i}$th possible result is

$$\Pr(\mathbf{i}) = \mathrm{tr}\left[\mathsf{M_i}\,\mathsf{L}_{N-2} \circ \cdots \circ \mathsf{L}_2 \circ \mathsf{L}_1(\rho)\right]. \qquad (2)$$

Here $\mathbf{i}$ is an n-bit string denoting the measurement result.

Markovianity ensures that the QIP's behavior is modular *in time*. It is the layers that are modular; each layer's effect on the QIP's state must be well-defined, only dependent on identity, and independent of temporal or other contexts. This is a powerful assumption. It makes modeling possible – we can now predict the results of *new* circuits as long as they are composed of layers that we have characterized already.

But *efficient* modeling of n-qubit circuits requires a stronger modularity condition. Representing every possible layer by an n-qubit CPTP map is neither compact nor tractable. Exponentially many layers need to be described, and each one requires $\mathcal{O}(16^n)$ real numbers. Even storing that model is impractical for large n, and learning it from data becomes infeasible for as few as three qubits. Stronger modularity assumptions, like the absence of crosstalk errors, enable efficient models like the one we present below.

Although general n-qubit Markovian models are intractable to reconstruct, Markovianity (like stability) can be tested. Published protocols include those in Refs. [12, 8, 59, 3, 32, 60]. Violations of the Markovian model – generally termed *non-Markovianity* – may result from a number of underlying causes, including time-dependence, persistent bath memories, or even serial context dependence,

where the performance of a layer operation is influenced by the layers that immediately precede (or even follow) it due to the finite bandwidth of control pulses.

We expect that *all* QIPs are at least a little bit non-Markovian, but we also expect that our Markovian model for crosstalk errors will (like the Markovian CPTP map model itself) fail gracefully, and continue to work well for slightly non-Markovian QIPs. However, our experience is that crosstalk detection protocols (including the one we develop in the second half of the paper) can confuse violations of Markovianity for crosstalk. So, in practice, it is important to test for non-Markovian effects before (or simultaneously with) testing for crosstalk.

4.1.3 Crosstalk-free Markovian QIPs

Whereas Markovianity allows modularity in time, a processor without crosstalk is modular in *space* – *i.e.*, across qubits and regions. Layer operations can reliably be composed by combining even smaller operations, that act locally and independently. Earlier, we said that a processor is crosstalk-free if it obeys locality and independence. If it is also Markovian, then the conditions for locality and independence can be stated as explicit conditions on the CPTP maps describing circuit layers.

Each *ideal* circuit layer defines a locality (tensor product) structure that partitions the qubits into disjoint and uncoupled target subsets. A Markovian QIP satisfies locality if and only if the CPTP map describing each layer obeys that structure. (The proof is trivial: for any bipartite system AB and operation $\mathcal{G}_{AB}$, there exists an initial product state $\rho_A \otimes \rho_B$ such that $\mathcal{G}_{AB}[\rho_A \otimes \rho_B]$ is correlated – *i.e.*, *not* a product state – if and only if $\mathcal{G}_{AB}$ is not a tensor product of operations).

Therefore, a Markovian model obeys locality if and only if each layer can be represented by a tensor product of local operations. For such a model, independence is well-defined. A local, Markovian QIP satisfies independence of operations if and only if each local operation (gate, initialization, measurement) is represented by the *same* local CPTP map in every layer where it appears. (No proof is needed – this is just a restatement of the definition of independence above in terms of CPTP maps).

If a Markovian model satisfies both of these conditions, then we say it is crosstalk-free. Its behavior is consistent with the absence of physical crosstalk, and its dynamics contain no crosstalk errors. Con-

versely, any violation of these conditions constitutes a crosstalk error.

If a QIP satisfies Condition 1 (locality), then each layer's CPTP map is a tensor product over the target subsets implied by that layer. The CPTP map for the layer described above, in the example of Markovianity, would be

$$\mathsf{L}(\mathrm{X}_0, \mathrm{CNOT}_{1\leftarrow 2}, \mathrm{H}_3) =$$
$$\mathcal{G}(\mathrm{X}_0) \otimes \mathcal{G}(\mathrm{CNOT}_{1\leftarrow 2}) \otimes \mathcal{G}(\mathrm{H}_3), \quad (3)$$

where $\mathcal{G}$, indexed by the gate operation and qubit target, represents a component CPTP map for that gate.

To satisfy Condition 2 (independence), a gate that appears in multiple layers must act identically in each of them. For example, in Fig. 1, layers 1, 3, 5, and 7 all contain a Hadamard gate acting on the fourth qubit. So the CPTP map describing layer 5 then must take the form

$$\mathsf{L}_5 = \mathsf{L}(\mathrm{H}_3) = \mathcal{G}(\mathrm{I}_0) \otimes \mathcal{G}(\mathrm{I}_1) \otimes \mathcal{G}(\mathrm{I}_2) \otimes \mathcal{G}(\mathrm{H}_3), \quad (4)$$

where $\mathcal{G}(\mathrm{H}_3)$ is the *same* local map that appeared in the other layers (although this map does not have to be the same as the same gate on another qubit, *e.g.*, $\mathcal{G}(\mathrm{H}_4)$).

Initialization and measurement operations must obey the same structure. For the four-qubit QIP shown in Fig. 1, this means that:

$$\rho = \rho_0 \otimes \rho_1 \otimes \rho_2 \otimes \rho_3 \quad (5)$$
$$M_{\mathbf{i}} = M_{0,i_0} \otimes M_{1,i_1} \otimes M_{2,i_2} \otimes M_{3,i_3}, \quad (6)$$

where M_{j,i_j} is the POVM effect operator for outcome i_j on qubit j. If the initial state is correlated, or the output bit on one qubit depends on another qubit's state, then the QIP is not crosstalk-free.

4.2 Discussion of the crosstalk-free QIP model

In classical systems, crosstalk usually refers to a signal in one channel influencing the signal on another channel. For example, inductive coupling between adjacent copper telephone wires may cause a conversation on one line to be heard on another. Analogous effects occur in QIPs – laser beams have finite width and may illuminate neighboring ions, superconducting transmission line resonators may capacitively couple to each other, or qubits themselves may interact directly. These interactions can be modeled by coupling Hamiltonians. So it is tempting to say that "crosstalk" is nothing more than a coupling Hamiltonian, and the complex abstraction that we have introduced is unnecessary.

But this misses three key points. First, those Hamiltonians appear in low-level device modeling, and are specific to particular physical implementations. Second, like all low-level device Hamiltonians, they fluctuate in time and with the state of the environment. Third, the systems that they couple are often ancillary ones – control wires, ambient fields, etc – that would not normally appear in an effective description of the processor and its qubits. Defining, detecting, and modeling crosstalk at this low level is possible – and even desirable for device physicists – but not portable across many devices.

We have presented a high-level, hardware-agnostic *effective* model. This approach is common. It is present when qubits are described as 2-dimensional Hilbert spaces, when elementary gates are described by CPTP maps, and when errors are modeled as depolarization or T_1 processes. Our model, like all of those techniques, trades the conceptual simplicity of Hamiltonian dynamics on very large system-specific Hilbert spaces for the practical tractability of an effective model on n qubits. The CPTP map formalism strikes a good balance between rigorous, low-level device models and cross-platform, high-level abstraction – but as a picture of the underlying physics, it is coarse-grained and can sometimes be counterintuitive.

For example, consider two qubits in a magnetic field along the Z axis whose strength varies slowly in time, see Fig. 2. The field causes both qubit states to rotate around the Z axis. Clearly, there is neither coupling nor communication between the qubits. So, if we include the magnetic field in our model, then it seems that there should be no crosstalk between the qubits. But if we only model the two qubits, and integrate out the field, then the CPTP map describing the *effective* dynamics of the two qubits violates the crosstalk free model – they experience correlated Z errors, which violate locality. This may appear counterintuitive, since the qubits are not coupled, and neither has any *causal* effect on the other. But it reflects the fact that there *is* crosstalk in the system, between each qubit and the magnetic field. Even when the field is eliminated from the model, it still mediates an effect that creates unexpected correlations between the qubits. Crosstalk errors can occur at the coarse-grained level even between two qubits that are not directly coupled by the underlying physics.

The stable/Markovian/crosstalk-free hierarchy of models given above is based on strict criteria that, as stated, are either true or false. One might object that
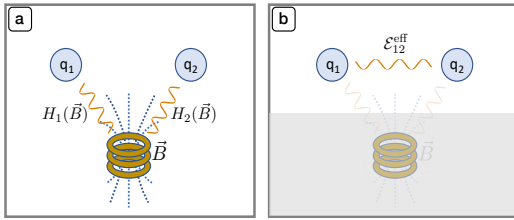
Figure 2: Two qubits influenced by the same fluctuating magnetic field $(\vec{B})$. **(a)** If the field's state is modeled and tracked, then there is no crosstalk between the two qubits the correlations between their states and errors are fully explained by the field and its coupling to them. **(b)** But if we do not track the field, focusing on the two qubits only then there *is* crosstalk between the qubits, in the form of correlated stochastic errors mediated by the (untracked) magnetic field.

these conditions are practically useless – no processor is *perfectly* Markovian or crosstalk-free, and could not be proven so even if it were. While this objection is strictly speaking true, it dismisses the utility of idealized models. No operation is perfectly unitary, yet unitary dynamics is both well-defined and highly useful *as an ideal*. In the same way, what matters is not whether a QIP is perfectly crosstalk-free, but how *close* it is to the ideal. The definitions given above lay the groundwork for metrics that quantify that closeness, and thus for measuring how much crosstalk is present.

Similarly, perfect Markovianity is not required. In a real and slightly non-Markovian QIP, we can confidently detect crosstalk as long as the violations of Markovianity (or stability) are small compared to the violations of the crosstalk-free conditions. An experiment to detect crosstalk has a certain duration and a certain statistical power. If it detects crosstalk, that conclusion is reliable as long as the QIP's instability and non-Markovianity do not rise above the experiment's level of sensitivity over its duration.

Finally, note that the CPTP maps describing experimental operations are only unique up to a gauge freedom [7, 47, 37]. In multiqubit QIPs, this gauge freedom is non-local. Gauge transformations – which simply change the *description* of the QIP, and have no observable consequences – can change the tensor product structure of operations, transforming a CPTP map that respects a tensor product structure to one that does not, and vice versa. This raises the question of whether the "crosstalk-freeness" of a model is real and experimentally testable, since it appears to be not gauge-invariant.

Fortunately, there is a simple resolution: a stable,

Markovian QIP is crosstalk-free if there exists *some* gauge in which Conditions 1-2 hold. This is directly analogous to the definition of a perfectly error-free gate set. An ideal target set of operations can be written down in many gauges. In all but one of them, the CPTP maps *appear* to be different from the original "ideal" ones. But this is the nature of gauge theories. What matters are the observable probabilities predicted from the theory. Those are identical in all gauges. So if there exists *any* gauge in which a gate set coincides with its ideal target, then no experiment (with this gate set) will ever detect any error. Similarly, if there exists *any* gauge in which a set of n-qubit operations is crosstalk-free, then no experiment (with this gate set) will detect evidence of crosstalk. A processor is crosstalk-free if and only if it admits *some* crosstalk-free model.

4.3 Examples

We now consider some examples of crosstalk phenomena, and the crosstalk errors they induce. All the examples in this section involve a QIP with just two qubits, which we label A and B. The examples can be generalized easily to more qubits.

1. Pulse spillover: Quantum gates should act only on their target qubits, but control pulses may spill over onto neighboring qubits and affect them. This is the most widely discussed form of crosstalk, *e.g.*, [57, 44, 43, 10, 34]. For example, consider two qubits that experience no errors when both are idle. But whenever an X_π gate is applied to qubit A, the control field spills over onto qubit B and induces a small X rotation. Each layer still respects the tensor product structure of the two qubits, so locality is not violated. However, the effect of the idle operation on qubit B depends on whether an idle or an X_π gate was applied to qubit A at the same time, so this scenario violates independence.

2. Always-on Hamiltonian: Suppose that when both qubits are idle, they experience an unwanted XX Hamiltonian. Thus, if A is in the $|+\rangle$ (respectively, $|-\rangle$) state, B undergoes a slow rotation around the $+X$ (respectively, $-X$) axis. Each qubit is influenced by the state of the other. This example violates locality, because the map describing the global idle is an entangling unitary operation, which is not a tensor product of two single-qubit CPTP maps.

3. Correlated stochastic errors from common causes: Correlated dynamics caused by a common influence can violate locality. For example

(see Fig. 2), suppose both qubits interact with a common magnetic field along the quantization axis, and that field undergoes white-noise fluctuations. This produces correlated (weight-2) dephasing or ZZ errors while the qubits are idle. This is not a tensor product map, and violates locality. Note that a constant field would only cause local unitary rotations, which respect the tensor product structure and does note result in crosstalk errors.

4. Detection crosstalk: Measurements of a qubit's state may be influenced by the state of neighboring qubits. As an example, consider measuring trapped-ion qubits A and B simultaneously using resonance fluorescence. If light scattered from qubit B spills over onto the detector for qubit A, then the result of measuring qubit A will depend on the state of qubit B. We refer to this type of crosstalk error as detection (or readout) crosstalk, because it specifically affects measurement results. This example violates locality – the POVM describing the measurement does not respect the QIP's tensor product structure, because the marginal effects corresponding to "0" and "1" on qubit A act nontrivially on qubit B.

5. Correlated state preparation: Correlated errors in the controls used to prepare the qubits can create correlated, or even entangled, initial states. This violates locality. For example, consider initializing qubits A and B to the $|0\rangle$ state using a common control field. Occasionally, some noise in the common control field may increase the state preparation error for both qubits. For any single trial, the resulting state would be a product state, but when averaged over many initializations the density matrix describing the initial state can no longer be factorized, so locality is violated.

This list of examples is not exhaustive, but we hope it helps to connect common notions of crosstalk to the conditions that define the crosstalk-free model.

4.4 Useful terminology for crosstalk errors

Any violation of the crosstalk-free model results in crosstalk errors, but there are many ways to violate the model. Some of them are quite distinct from others, both in the physical phenomena that typically produce them, and in their consequences and behavior. It is useful to identify the most common categories and give them names, if only to facilitate answering the question "What kind of crosstalk do you see?" We suggest some useful categories here, based on our experience examining data and modeling noise.

First, we observe a fundamental difference between errors that violate locality, and those that only violate independence. Any violation of locality can be traced to at least one specific layer operation that creates unexpected correlations. We call these crosstalk errors *absolute*. In contrast, violations of independence *cannot* be isolated to a specific layer operation. Some local operation just behaves differently in different layers, and no one layer defines the correct behavior of that operation. We call these crosstalk errors *relative*.

In addition to these terms, which are relatively rigorous, we have found the following less-precise categories to be useful. These categories are not intended to be exhaustive, and may not prove over time to be the most useful classification. For example, the "correlated state preparation" example given in the previous section does not fall into any of these categories (it could define another category, but it is not clear that it is sufficiently common or important). Other violations of the crosstalk-free model can be invented that fall into none of these three categories, or bridge them. Furthermore, we do not yet have specific protocols for rigorously distinguishing these categories. Nonetheless, we have found them useful, and so we propose them to the research community.

Idle crosstalk is any violation of locality when all qubits are idle. The unique layer in which no nontrivial operations are performed corresponds to a CPTP map that we call the global idle, and if the global idle is *not* a tensor product of 1-qubit CPTP maps, then we say there is idle crosstalk. Any error occurring during the global idle that produces correlation between qubits (an error of weight 2 or higher) is an idle crosstalk error. Examples 2 and 3 in the previous section are examples of idle crosstalk errors. The same physical phenomena (always-on Hamiltonians, correlated decoherence, *etc.*) can also cause high-weight errors during nontrivial gates, but their effects are usually strongest and easiest to detect during the global idle.

Operation crosstalk refers to violations of independence caused by particular elementary operations. A QIP displays operation crosstalk if the act of performing an operation on qubits in region A changes the dynamics of qubits in a disjoint region B. It is not always possible to unambiguously ascribe a crosstalk error to an operation (*i.e.*, to define operation crosstalk orthogonally to idle crosstalk), but we have found it useful to have terminology for crosstalk errors that change as (non-idle) operations are ap-

plied to a QIP. Operation crosstalk is a special case of relative crosstalk, corresponding to cases where the change in region B's dynamics can confidently be blamed on a particular operation.

Detection crosstalk refers to violations of locality in the outcomes or results of measurement operations. If the result of a measurement on one qubit depends on the pre-measurement state of another qubit, that is detection crosstalk. We avoid the term "measurement crosstalk" because it is ambiguous; it could also refer to errors on spectator qubits that are caused by measuring a target qubit in the middle of a circuit, which would be an instance of operation crosstalk instead of detection crosstalk. Example 4 in the previous section is an instance of detection crosstalk.

5 Crosstalk errors are too diverse to detect without assumptions

Having given a definition of crosstalk errors, we would like to be able to test a QIP to detect their presence, and for further characterization purposes, reveal the structure of crosstalk in the QIP (*i.e.*, map out which qubits are most impacted by the crosstalk errors so as to focus the next level of detailed characterization on this subset).

5.1 Detecting arbitrary crosstalk errors is hard

Comprehensive characterization of crosstalk errors is extraordinarily demanding. Even just detecting any possible crosstalk error is hard (it requires resources that scale super-polynomially with the number of qubits). Let us demonstrate this. To begin, we need to define and exclude "weak" errors that can be arbitrarily hard to detect. We say that an experiment E *detects crosstalk* in a stable, Markovian model M if performing E on a QIP described by M has a high probability of producing data that rules out every crosstalk-free model with high confidence. Crosstalk in a model M is "strong" if it can be detected by an experiment using a small number of layers, and "weak" if all the experiments that detect it require a large number of layers. (These concepts are easy to state quantitatively, but it is tedious and not necessary here).

Even detecting strong crosstalk errors is hard because crosstalk models M are combinatorially diverse. Each given one can be detected easily by a tailored experiment, but no small set of experiments can detect them all efficiently. To illustrate this, we consider two examples, one for relative crosstalk and one for absolute crosstalk.

To see that arbitrary relative crosstalk is hard to detect, consider a QIP that allows any parallel combination of either an X_π rotation or I (the identity) on each qubit. Index these possible layers by n-bit strings, where "0" and "1" on the kth bit indicate (respectively) that I or X_π should be performed on the kth qubit. Let s be a randomly selected n-bit string, and suppose that every layer *except* the one indexed by s acts perfectly, while applying the one indexed by s depolarizes all qubits. While each layer respects locality, this model has strong violation of independence because on each qubit there is a gate that causes an error if (and only if) the other $n-1$ qubits are controlled in a particular way.

This crosstalk error is hard to detect because it only occurs if a particular layer (out of exponentially many possible layers) is performed – but it constitutes strong crosstalk because it is easy to demonstrate by using that layer.

A second example illustrates an analogous problem for absolute crosstalk. Consider the "idle layer", where no gates are performed on the qubits. It should act as the n-qubit identity channel. Again, let s be a randomly selected n-bit string, and let the idle layer act as the unitary that applies a phase -1 to $|s\rangle$ and acts trivially on its complement. This unitary can easily correlate qubits, so it violates locality. It is strong, because if s is known, then the correlation can be detected using just a few very short circuits. But it is also, of course, a Grover oracle for the unknown s. Detecting that it is not the identity is known to be as hard as finding s, which requires $O(\sqrt{2^n})$ uses of the layer [5].

This sort of crosstalk is hard to detect because it is very weak on almost all input states. It only manifests as a significant effect if the input state has high overlap with $|s\rangle$. So there is a bit of a catch-22: this crosstalk is strong because it *could* have a dramatic impact on a particular input state, but hard to detect because it has almost no effect on most input states.

Detecting arbitrary strong crosstalk errors is impossible to do efficiently, because it requires testing an exponential number of configurations. Going further, and characterizing those errors (even partially) is strictly harder. Designing a protocol to detect crosstalk errors and learn something about them requires specifying something more about the *kind* of errors to be detected, and accepting that other kinds of errors may not be detected.

5.2 Low-weight crosstalk errors

We expect characterizing crosstalk in QIPs to require device-specific protocols, informed both by theoretical models of a specific QIP's behavior and the specific tasks or applications that it will run. But generic protocols are also important. They provide cross-platform benchmarks, and may detect unexpected errors that tailored protocols miss because of their design. In the next section, we present a candidate protocol of this type, whose purpose is to (1) detect a significant (but not universal) class of crosstalk errors, and (2) localize those errors, by characterizing which qubits they affect (but not how they act on those qubits). Since no efficient protocol can be completely generic, some sort of assumptions are necessary to limit the diversity of crosstalk errors.

Our protocol targets *low-weight* crosstalk errors – ones that result from interactions of just a few subsystems that are supposed to be independent. In a processor that is crosstalk-free, distinct subsystems *never* interact or develop correlations (note that by "distinct" we mean "not intentionally" coupled – two qubits undergoing a 2-qubit gate form a single subsystem for this purpose). So if the weight of a crosstalk error is (informally) defined as the number of distinct subsystems that it couples together, then all the errors in a crosstalk-free QIP have weight 1.

In contrast, the two examples in the previous subsection illustrated high-weight crosstalk errors. Each example constructs an input/output function that depends on *all* of its inputs. In the first example, that function was the map from layer specifications (represented as n-bit strings) to CPTP maps. In the second example, that function was the CPTP map for a single layer, which applied a phase that depended inextricably on every qubit of the input state. Functions or maps that depend on all their inputs in arbitrary ways are (demonstrably) too diverse, allowing even strong crosstalk to be hidden from efficient detection.

Defining "weight" precisely for an absolute crosstalk error, which appears in a specific layer x, is straightforward. That layer's action is represented by a CPTP map $\mathsf{L}(x)$. Its ideal error-free action can be represented by a CPTP map $\mathsf{L}_0(x)$, and so the error in that layer is represented by $\mathcal{E}(x) = \mathsf{L}_0(x)^{-1}\mathsf{L}(x)$. The layer's ideal behavior defines a decomposition into distinct subsystems $r_1 \otimes r_2 \otimes ...$ that should not interact. An error map has weight k if it can be written as a convex combination of maps that act trivially (as the identity) on all but k of those subsystems, and cannot be written this way for any smaller k.

Typical error maps are not exactly weight-k for any finite k – e.g., a tensor product of local weak error channels has terms of every weight – but can be approximated very well by low-weight channels, because the magnitude of the weight-k terms declines exponentially with k above some value. Henceforth we will take this for granted, and by "low-weight error map", we will mean "error map that can be approximated to high precision by a sum of low-weight terms."

Quantifying the weight of relative crosstalk errors is slightly more technical. To do so, we consider a larger state space describing a register of n qubits $\mathcal{Q} = \bigotimes Q_i$ *and* a register of n classical digits $\mathcal{C} = \bigotimes C_i$. Each C_i specifies what operation is to be performed on the corresponding Q_i. Every possible layer is represented by a distinct state of $\mathcal{C}$, and an entire stable Markovian model can be represented by a *single* operation $\mathcal{M}$ acting on $\mathcal{C} \otimes Q$, of the form

$$\mathcal{M} = \sum_{\text{possible layer specs } x} |x\rangle\langle x|_{\mathcal{C}} \otimes \mathsf{L}(x)_{\mathcal{Q}}. \quad (7)$$

This is simply a conditional operation, which applies CPTP map $\mathsf{L}(x)$ to the qubits, conditional on the classical control register being in state x.

Now, as above, we can write $\mathcal{M} = \mathcal{M}_0\mathcal{E}_{\mathcal{M}}$ so that

$$\mathcal{E}_{\mathcal{M}} = \sum_x |x\rangle\langle x|_{\mathcal{C}} \otimes \mathcal{E}(x)_{\mathcal{Q}} \quad (8)$$

is the *entire* model's error operation, and perform the same decomposition into weight-k terms. Now the ith subsystem is not just Q_i but $C_i \otimes Q_i$, and a relative crosstalk error that causes Q_i to evolve differently conditional on another qubit's control line C_j is represented by a weight-2 term in $\mathcal{E}_{\mathcal{M}}$ [2].

Many natural and expected forms of crosstalk have low weight. Note that low weight does not mean that a single qubit is not perturbed by many other qubits or control lines – it just means that it is perturbed *independently* by them. So low-weight crosstalk encompasses many simultaneous few-body interactions. Moreover, low-weight crosstalk errors are not very diverse. Simple counting shows that there are

---

[2]We note that it is arguably more elegant to represent error maps including $\mathcal{E}_{\mathcal{M}}$ by their logarithms or generators, and apply the same weight decomposition to them. The logarithm of a tensor product $\mathcal{E}_1 \otimes \mathcal{E}_2$ is a sum of weight-1 terms, $\log(\mathcal{E}_1)\otimes\mathrm{I}+\mathrm{I}\otimes\log(\mathcal{E}_2)$. So the error generator of a crosstalk-free model is *exactly* weight-1, whereas the error map itself is only approximately weight-1. However, this representation is less common and requires more machinery that seems unjustified for our purposes here.

only $\mathcal{O}(n)^k$ errors of weight at most k on n qubits, so we can hope to detect any low-weight crosstalk error without devoting exponential resources to the task.

We do not expect that *all* crosstalk errors will have low weight, but we expect that high-weight errors will stem directly from specific features of the QIP (especially its control architecture, where classical correlations can flourish and induce highly complex dependencies), and are best addressed by tailored, device-specific protocols. Because low-weight errors are plausible in almost any architecture, a generic protocol to detect and localize them is desirable.

6 An operational protocol for detecting crosstalk errors

We now return to the abstract definitions of locality and independence presented in Sec. 3 to build a protocol for detecting crosstalk errors, based on the fact that violations of these conditions can be observed directly from operational phenomena.

In Sec. 6.1 we present the model-free and operational definition of crosstalk-free QIPs that forms the basis of the protocol. In Sec. 6.2-6.4 we develop the ingredients of the protocol, including an efficient set of experiments to be performed and tractable data analysis based on statistical tools originally developed for inference on probabilistic graphical models. In Sec. 6.5 we discuss in detail the assumptions behind our protocol and its limitations, especially the crosstalk errors it can and cannot detect. Finally, in Sec. 6.6 we present some guidance on how to choose the parameters that define our crosstalk detection protocol based on the physics of the QIP under test.

6.1 Model-free framework and definitions

Consider a QIP comprising n qubits. Let r be a partition of the n qubits into $M < n$ disjoint subsets, $r_i \subset \{0, ..., n-1\}$, which we call *regions*, and let $\mathsf{n}(r_i)$ be the number of qubits in region r_i. We assume no model, only that for each region r_i we (1) apply operations that ideally should only affect qubits in r_i and should not affect qubits in any other region, and (2) make measurements whose results should only depend on the state of qubits in r_i. We will define crosstalk errors in terms of the *settings* that denote the operations applied to a region, and the *results* of measurements on qubits in a region. An experiment is defined by a tuple $\Omega \equiv (\mathsf{S}_{r_0}, \mathsf{S}_{r_1}, ..., \mathsf{S}_{r_{M-1}}, \mathsf{R}_{r_0}, \mathsf{R}_{r_1}, ..., \mathsf{R}_{r_{M-1}})$, where S_{r_i}
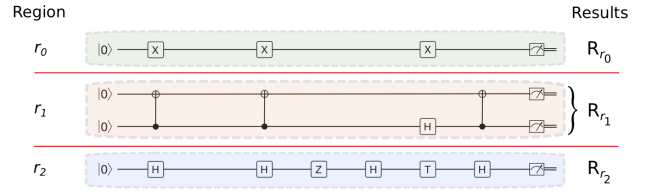

Figure 3: Illustration of the type of circuits used in our protocol. A 4-qubit QIP is partitioned into three regions, labeled r_0, r_1, r_2, and the goal is to detect crosstalk errors between these regions. To do so, we perform circuits that *only* apply coupling operations between qubits within a region, never between regions (across the red lines in the figure). The random variable outcome from measuring the qubits in region r_i is denoted R_{r_i}. In this example R_{r_0} and R_{r_2} are 1-bit-valued while R_{r_1} is 2-bit-valued.

are the settings assigned to the qubits in region r_i and R_{r_i} are the measurement results from the qubits in region r_i. We treat each member of this tuple as a random variable drawn from some sample space, $\mathsf{S}_{r_i} \in \mathbb{S}_{r_i}, \mathsf{R}_{r_i} \in \mathbb{R}_{r_i}$. It is clear that the results are random variables; they are the results of measurements on quantum systems, which are always random variables. We also treat the settings as random variables, but for a different reason. In a large QIP, it is not feasible to perform an exhaustive set of experiments that enumerates all the possible experimental settings. So, in practice, observed data constitute a sample over all the possible settings. As we shall see, a random sampling over settings often yields good results. The random variable R_{r_i} takes values that are bit strings of length $\mathsf{n}(r_i)$, obtained by measuring all qubits in region r_i in some basis. More complicated scenarios, *e.g.*, involving detection of leakage, are possible but we restrict ourselves to the simplest case here. Fig. 3 illustrates these definitions.

The settings S_{r_i} are random variables that describe (i) what state is prepared natively on the qubits in r_i, (ii) what gates are applied to the qubits in r_i, and (iii) what basis the qubits in r_i are measured in. So S_{r_i} labels a *quantum circuit* for that region (defined here as the state preparation, applied gates and measurement basis choice for a region). We note that most quantum computing architectures have only one qubit state that is natively prepared (*e.g.*, the ground state) and only one measurement basis (*e.g.*, the Z basis). Therefore the only setting that can be varied is the gates applied to the qubits in between state preparation and measurement. Hence in most quantum computing architectures, the settings will be synonymous with "gates applied to qubits in r_i".

Definition 2: We say that a region r_i is free of crosstalk errors to/from other regions if conditional distributions over the measurement results on this region satisfy:

$$P(\mathsf{R}_{r_i}|\mathsf{S}_{r_i}, \mathsf{T}) = P(\mathsf{R}_{r_i}|\mathsf{S}_{r_i}), \quad \text{with}$$
$$\mathsf{T} \subseteq \Omega \setminus \{\mathsf{R}_{r_i}, \mathsf{S}_{r_i}\} \quad (9)$$

This means that the distribution of measurement results on region r_i depends *only* on the settings for r_i; conditioned on those settings, it is independent of all the other random variables in Ω. Any violation of these conditions is a witness to some kind of crosstalk error.

It is preferable to define the crosstalk-free condition in terms of conditional independence as opposed to *marginal* independence – *i.e.*, $P(\mathsf{R}_{r_i}, \mathsf{S}_{r_j}) = P(\mathsf{R}_{r_i})P(\mathsf{S}_{r_j})$ – because it is more robust to confounding by hidden (or intentional) correlations in settings, which can become an issue when detecting crosstalk errors in large QIPs. Appendix A discusses this further.

This model-free definition of crosstalk errors is equivalent to our model-based definition of crosstalk errors (Definition 1) stated in Sec. 3; see Appendix B for proof. The two definitions capture the same notions of locality and independence of quantum operations – the model-based definition does so in terms of conditions on models of quantum operations (*i.e.*, CPTP maps), while the model-free definition does so in terms of conditions on operational random variables that arise naturally in a QIP.

Example. Here is an example to illustrate the notation introduced above. We wish to detect crosstalk errors induced by single qubit operations on a QIP with 3 qubits, partitioned into two regions $r_0 = \{0\}$ and $r_1 = \bar{r}_0 = \{1, 2\}$. The following elementary single-qubit operations can be performed: initialization in $|0\rangle$; initialization in $|+\rangle$; idle gate (*i.e.*, do nothing for one clock cycle); $X_{\pi/2}$ gate; $Z_{\pi/2}$ gate; and measurement in the computational basis. Circuits can be performed that comprise (1) parallel initialization of all 3 qubits, (2) a sequence of k layers built from arbitrary single-qubit gates on each qubit in parallel, and (3) measurement of all qubits in the computational basis. Then the sample space of settings on region r_0 – which includes only qubit 0 – is

$$\mathbb{S}_{r_0} = \mathbb{S}_p \times \mathbb{S}_g$$
$$= \{\text{Prep}_{|0\rangle}, \text{Prep}_{|+\rangle}\} \times \{G_I, G_X, G_Z\}^k$$

where we have distinguished prep settings ($\mathbb{S}_p$) and gate settings ($\mathbb{S}_g$). Only one measurement layer is allowed and the only measurement basis accessible is the computational basis, so there are no measurement settings. The space of settings for region r_1 is isomorphic to two copies of the settings for r_0: $\mathbb{S}_{r_1} = \mathbb{S}_{r_0} \times \mathbb{S}_{r_0}$. The spaces of possible results for each of the two regions are simply $\mathbb{R}_{r_0} = \{0, 1\}$ and $\mathbb{R}_{r_1} = \{0, 1\}^2$. In this example, each experiment is labeled by the following tuple of nine random variables,

$$(\mathsf{S}_{r_0}, \mathsf{S}_{r_1}, \mathsf{R}_{r_0}, \mathsf{R}_{r_1})$$
$$= ((P_0, G_0), ((P_1, G_1), (P_2, G_2)), R_0, (R_1, R_2)), \quad (10)$$

where $P_i \in \mathbb{S}_p, G_i \in \mathbb{S}_g$ and $R_i \in \{0, 1\}$ label (respectively) the preparation, sequence of gates, and measurements results for qubit i.

The model-free definition given by Eq. (9) leads directly to practical tests for crosstalk, because if we draw a circuit at random from the distribution defined by $P(\mathsf{S}_{r_0}, \mathsf{S}_{r_1}, ..., \mathsf{S}_{r_{M-1}})$ and perform it on the QIP, the result is a sample from the joint probability distribution $P(\Omega) \equiv P(\mathsf{S}_{r_0}, \mathsf{S}_{r_1}, ..., \mathsf{S}_{r_{M-1}}, \mathsf{R}_{r_0}, \mathsf{R}_{r_1}, ..., \mathsf{R}_{r_{M-1}})$. These samples can be used to statistically test the conditions implied by Eq. (9). This is, in fact, a general procedure for detecting crosstalk errors. There is always some partitioning of the QIP into regions, some circuit family that can be executed, and some data sample size that will detect *any* crosstalk error using this method. However, as discussed in Sec. 5, detecting any possible crosstalk error requires exponential resources, and thus is not a scalable goal. Therefore, our aim is to use this model-independent definition to formulate an efficient protocol that targets low-weight crosstalk errors.

Developing this protocol requires three ingredients: (i) defining a set of region partitions for a QIP, (ii) defining a set of experiments to perform on the QIP, and (iii) defining an analysis technique on the data produced by these experiments to detect crosstalk using Definition 2. The following subsections tackle each of these ingredients.

6.2 Defining regions

Our crosstalk detection protocol looks for correlations between regions of a QIP that should be uncoupled. This requires partitioning the QIP into disjoint regions. No single partition into regions will suffice – for example, we might need to test whether the 2-qubit region $\{1, 2\}$ has crosstalk with the 2-qubit

region $\{3,4\}$, but also whether $\{2,3\}$ has crosstalk with $\{4,5\}$. So we need multiple partitions, and for each one, we will define a set of circuits that respect it.

We cannot test every possible partition – the total number of ways to partition n qubits is B_n, the nth Bell number, which scales super-exponentially in n. However, testing all possible partitions is unnecessary. Crosstalk errors are associated with individual layers of elementary operations. In almost every QIP architecture, each elementary operation targets only 1 or 2 qubits. So, since we focus on low-weight crosstalk errors, it is sufficient to consider partitions into disjoint one- and two-qubit regions. These allow us to ask (and detect) whether correlations emerge between any two such regions, in circuits that never couple them intentionally.

Let us first set some terminology. We refer to a region containing exactly k qubits as a *k-region*, and a partition of the entire QIP into regions that each contain k or fewer qubits as a *k-partition*. We say that a region is *allowed* if it is possible to define circuits that couple all the qubits within that region, without involving any other qubits. So a 2-region is allowed only if the QIP has a 2-qubit gate directly between those qubits. We say that a region is *in* a given partition if it is one of the regions making up the partition – *e.g.*, region $\{1,2\}$ is in the 6-qubit partition $\{\{1,2\},\{3,4\},\{5,6\}\}$. Further, a tuple of regions is in a given partition if each region in the tuple is in the partition – *e.g.*, the pair $\{1,2\},\{5,6\}$ is also in the above partition.

There is exactly one unique 1-partition of an n-qubit QIP. So if we wish to detect crosstalk errors associated with single-qubit gates, this is the only partition we need to use.

However, we must also detect crosstalk errors associated with 2-qubit gates, which requires 2-partitions. The number of possible 2-partitions scales exponentially with n; the number of ways to partition n elements into distinct sets of size k (assuming k divides n) is $\#(n,k) = \frac{n!}{(k!)^{n/2}(n/k)!}$, and hence $\#(n,2) \approx (\sqrt{n/e})^n$, via the Stirling approximation. This assumes that two-qubit operations are possible between any two qubits in the QIP, however even in the more realistic case of limited connectivity, the number of 2-regions grows exponentially in n. So it is impractical to even test all 2-partitions exhaustively. Fortunately, since we are focused on detecting low-weight crosstalk errors, it suffices to detect pairwise crosstalk between 2-regions (see Sec. 6.5 for a

discussion of the resulting limitations), and doing this only requires that we guarantee that every pair of 2-regions is in at least one 2-partition of the QIP that is tested.

Since there are at most $n(n-1)/2$ allowed 2-regions, there are $\mathcal{O}(n^4)$ pairs of 2-regions. Therefore, it is easy to define a set of $\mathcal{O}(n^4)$ 2-partitions that contain every such pair (*e.g.*, for each of the $\mathcal{O}(n^4)$ possible pairs of 2-regions, define a partition by starting with those two 2-regions and then arbitrarily partitioning the remaining qubits into 2-regions). This requires only poly(n) resources, but is clearly wasteful. If we are satisfied with a high probability guarantee, a randomized partitioning strategy is more efficient.

Theorem 1 *Given an n-qubit QIP, let r_i, $0 < i < R = \frac{n(n-1)}{2}$, be a labeling of all 2-regions in the QIP, and let $\mathcal{P}$ be a set of independent and uniformly sampled random 2-partitions of the QIP. For any $0 < \epsilon < 1$, the probability that any pair of distinct 2-regions is in at least one 2-partition is bounded below by $(1-\epsilon)$ if $|\mathcal{P}| \geq n^2(2\log(R) - \log(\epsilon))$.*

Proof: We will follow the logic of the proof of Theorem 3.1 in Ref. [36]. Let p be a lower bound on the probability that any 2-partition of the QIP contains a pair of 2-regions (we will see below this bound is the same for any pair). Then, the probability that this pair of 2-regions is not in $|\mathcal{P}|$ random 2-partitions is at most $(1-p)^{|\mathcal{P}|}$. By applying the union bound, we see that the probability that any one of the $R(R-1)/2$ possible pairs of 2-regions in the QIP is not in any partition in $\mathcal{P}$ is at most $\frac{R(R-1)}{2}(1-p)^{|\mathcal{P}|}$. We would like this probability to be at most ϵ,

$$\frac{R(R-1)}{2}(1-p)^{|\mathcal{P}|} \leq \epsilon$$

$$\Rightarrow \quad |\mathcal{P}| \geq \frac{\log\left(\frac{\epsilon}{R^2}\right)}{\log(1-p)}$$

$$\Rightarrow \quad |\mathcal{P}| \geq p^{-1}(2\log(R) - \log(\epsilon)),$$

where we have used the fact that $-p > \log(1-p)$. It remains to compute the lower bound p in terms of the system parameters. The probability that any pair of 2-regions is in a 2-partition is given by the ratio $\frac{\#(n-4,2)}{\#(n,2)}$ since there are $\#(n-4,2)$ possible partitions of the remaining $n-4$ qubits once the four qubits in the pair of 2-regions have been removed. Computing this ratio, we get $\frac{1}{(n-1)(n-3)}$, and hence a lower bound on this probability is $p > \frac{1}{n^2}$. Substituting this into the above bound on $|\mathcal{P}|$ gives the desired result. $\square$

Either of these approaches – the brute-force array of $\mathcal{O}(n^4)$ partitions, or the $\mathcal{O}(n^2 \log(n))$ randomized strategy – defines a list of poly(n) partitions of the QIP that will detect pairwise crosstalk between any pair of 2-regions with high probability.

We note that in QIPs with local connectivity restrictions, *e.g.*, a planar array of qubits where intentional coupling operations are only possible between nearest neighbors qubits, $p^{-1} = \mathcal{O}(n)$, and therefore the scaling of the randomized strategy is improved to $\mathcal{O}(n \log(n))$. Similarly, the scaling of the brute-force partitioning improves to $\mathcal{O}(n^2)$ under local connectivity.

6.3 Lightweight experiment design

Given a partition of a QIP into regions, we must define a set of circuits to run on the QIP that constitute the crosstalk detection experiment. We only consider circuits that do not (intentionally) couple regions, which means that for each region there is a well-defined *subcircuit* comprising all operations applied to it. We also assume, for the sake of simplicity, that the QIP has unique initialization and measurement operations (in $|0\rangle^{\otimes n}$ and the $\{\langle 0|, \langle 1|\}^{\otimes n}$ basis). Thus, the settings for a region correspond precisely to the gates in the subcircuit on that region.

Each possible circuit on the QIP is composed of the parallel application of multiple subcircuits, one on each region. The simplest approach is to choose a collection of N_{circ} subcircuits for each region, and then perform *all* combinations of those subcircuits. We refer to this collection of subcircuits as the "bag" of circuits applied to a region. (We postpone the question of what subcircuits to place in the bag to the end of this subsection). We call this the *exhaustive experiment*, in which each subcircuit on region r_i gets performed in an exhaustive variety of different *contexts – i.e.*, in parallel with all N_{circ} circuits on other regions – and so violations of independence are easy to detect in the data. Unfortunately, this experiment defines a hypercube containing N_{circ}^M distinct circuits, which grows too rapidly with M (the number of regions) to be feasible.

However, we observe that in the exhaustive experiment, each subcircuit on every region r_i is performed in exponentially many distinct contexts (defined by the settings on the other regions $r_j \neq r_i$). This is arduous and overkill; since crosstalk errors are not likely to only be present in one or few of this exponential number of contexts (this is discussed further in Sec. 6.5), we can subsample from this exhaustive

experiment. So we will choose a sparse subset of the experiments in the hypercube defining the exhaustive experiment, with the goal of defining a small set of experiments that allow low-weight crosstalk errors to be detected.

6.3.1 An explicit construction

The sparse sampling of the hypercube should maintain two important properties of the exhaustive experiment. First, each subcircuit in the bag for each region r_i must appear in multiple contexts (but not *exponentially* many). Second, that set of contexts in which each subcircuit gets performed must vary on *each* of the other regions. These properties ensure that – whatever subcircuits we select for each region's bag – the data will reveal whether the local results of those subcircuits are significantly influenced by the settings (choice of subcircuit) on any other region.

The construction we outline now ensures that these properties are preserved, even with much fewer experiments. It is defined by three adjustable integer parameters:

- L is the length or depth of all the subcircuits. Subcircuits on different regions are applied in parallel, so they must all be the same length, so L must be chosen and fixed.

- N_{circ} is the number of circuits in the bag for each region. This number can be chosen to be a constant (independent of n), as we argue below.

- N_{con} is the number of random contexts in which each subcircuit will be tested.

First, we choose a bag of N_{circ} depth-L subcircuits for each of the M regions (see below for their construction). Now, for each region $m \in [0 \ldots M-1]$ and each of the subcircuits ν^m in that region's bag, we define N_{con} *different* circuits that perform ν^m in different contexts, by choosing a subcircuit for each of the other regions at random from the corresponding bag, and performing all those subcircuits (including ν^m) in parallel. This circuit selection procedure is illustrated in Fig. 4.

This design ensures that (1) for each region, approximately N_{circ} different subcircuits are studied in detail, (2) each of these subcircuits is performed in N_{con} different contexts, and (3) those contexts vary independently across all the other regions.

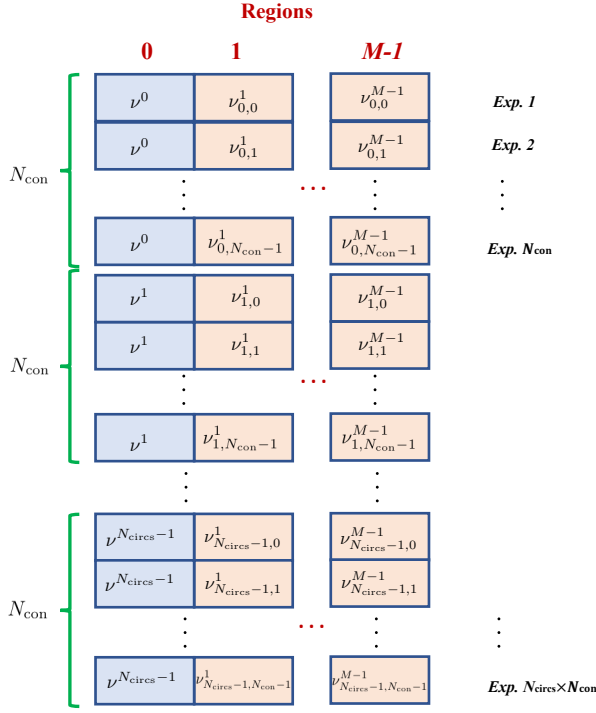We have found that a small refinement improves the protocol in practice. Often, the idle gate is less

Figure 4: Illustration of the circuits performed in the lightweight crosstalk error detection experiment design. The QIP is partitioned into M non-overlapping regions, and subcircuits of fixed length L are applied to each region. This diagram represents one "epoch", during which the N_{circs} subcircuits in the bag for region 0 (denoted ν^i) are iterated over. Each ν^i is repeated N_{con} times on region 0, and in each instance (each line in the diagram) the subcircuits on the other $M-1$ regions are randomly sampled (with replacement) from the subcircuits in the bag for that region. In the diagram $\nu_{i,j}^m$ denotes the subcircuit applied to region m in the jth context when ν^i is applied to region 0. The experiment design prescribes M such epochs; in epoch m the N_{circs} subcircuits in the bag for region m are iterated over while the subcircuits for all the other regions are randomly sampled. It is assumed that all qubits are initialized in their ground states, and the measurements after each of the prescribed circuits are performed simultaneously on all qubits in the computational basis. Finally, each of the experiments is repeated N_{rep} times in order to collect statistics.

noisy than others, and the depth-L idle circuit is the least noisy depth-L circuit *and* most sensitive to crosstalk. We have found it useful to artificially boost the probability of sampling all-idle circuits when the random contexts are defined (*not* when ν^m is drawn). To do this, we sample context subcircuits normally, but replace each sample by the depth-L idle circuit with probability p_{idle}. This experiment design is also described by pseudocode in Appendix C.

Two additional variations are useful in some circumstances. First, in certain regimes, we find that crosstalk can be comprehensively detected without testing all N_{circ} subcircuits in each region's bag. Iterating ν^m over a randomly chosen subset is sufficient. Second, it is sometimes easier to *sample* the subcircuits ν^m at random (with replacement) – just as the contexts are sampled randomly – than to iterate over them.

The experiment defined above can be seen as a sparse filling of the hypercube defined by the exhaustive experiment, as long as N_{con} does not grow exponentially with M. Ideally, it should not grow with M at all. In practice, we find a constant N_{con} (with respect to M) to be sufficient to detect crosstalk errors. The protocol also requires specifying both N_{circ} and how to construct the subcircuits, which we discuss at the end of this section. The total number of experimental configurations to be performed for any fixed length L is $N_{\text{exp}} \approx M \times N_{\text{circ}} \times N_{\text{con}}$ [3], which scales linearly in the number of regions, $M = \mathcal{O}(n)$.

Each of the circuits prescribed for this protocol should be repeated N_{rep} times to collect statistics. Each repetition yields a single datum, comprising a label for the circuit applied to each region (S_{r_i}) and a bit string describing the measurement results from each region (R_{r_i}). This is a single sample from the distribution over settings and results $P(\Omega)$ that we seek to test for correlations that signal crosstalk errors.

As much as possible, the repetitions of the various circuits should be distributed uniformly over the entire time of the experiment – *not* performed all at once in a single chunk. They may be *rasterized* (each circuit is performed once, in succession, and this is repeated N_{rep} times), or *randomized* (all the circuit repetitions are shuffled and performed in completely random order). This minimizes the probability of systematic false positives caused by drift. If behavior of the device (*e.g.*, error rates) is correlated with time – *i.e.*, it drifts – then if the settings are also correlated with time, this will produce spurious evidence of correlation between settings and results. Time is an unobserved, or latent, variable; *e.g.*, in the simplest case an unobserved classical degree of freedom (*e.g.*, a two-state fluctuator) may cause drift by providing a fluctuating local potential. When a variable that is a common cause for multiple other variables

---

[3]This is an approximately equal to and not a strict equality because any accidentally duplicate circuits generated during the sampling of contexts are removed.

Accepted in ⟨ ⟩uantum 2020-09-07, click title to verify. Published under CC-BY 4.0.

15

is not observed, it can create a fictitious conditional dependence between these variables [24]. Randomization and rasterization reduce or destroy correlations between the settings and time, reducing the risk of conflation; see also related discussion about conflation in Sec. 6.4.2. Finally, we note that rasterizing also facilitates concurrent drift detection with the same data [46].

We saw above that the number of distinct experiments required to detect crosstalk errors for one partition of the QIP is $\mathcal{O}(n)$. This is multiplied by the number of partitions to get the overall cost of detecting crosstalk errors for the QIP. As discussed above, the number of required partitions for a QIP scales as $\mathcal{O}(n^2 \log(n))$ in the worst case. Therefore the number of distinct experiments required for crosstalk detection with our protocol scales as $\mathcal{O}(n^3 \log(n))$ (this reduces to $\mathcal{O}(n^2 \log(n))$ if qubits are only coupled locally).

6.3.2 Choosing the subcircuits for each region's bag

Exactly what subcircuits to choose or define for each region is a critical component. We have left it open for now for a simple reason: there are many reasonable, yet very different, possible choices. For the sake of concreteness, we specify particular circuits here. But we also expect that new, creative, and perhaps objectively better choices can be usefully explored.

The subcircuits run on each region have two purposes: to *manifest* crosstalk, and to *detect* crosstalk. It may be that only certain circuits, when run on r_i, cause or amplify errors on r_j. And it may also be that certain circuits on r_j are more sensitive to these effects. Our goal is to detect whatever crosstalk errors exist. Therefore, in principle, the subcircuits chosen for each region's bag should be those that (1) cause the greatest effects on other regions, and (2) are the most sensitive to effects caused by other regions.

Given a specific physical model of crosstalk, it is possible to design subcircuits with these properties. Or, given a *parameterized* model of the sorts of crosstalk that might occur, it is possible to design a rather larger set of subcircuits that *collectively* amplify all the effects appearing in that model. (This is how the circuits for gate set tomography (GST) are chosen).

An entirely different approach is to switch from trying to detect *all* low-weight crosstalk errors to focusing on the crosstalk errors that impact a specific application. This motivates choosing subcircuits that are representative of the subroutines that appear in specific algorithms, and would emphasize detection of crosstalk errors that impact execution of those particular algorithms. There is no obvious general way to doing this since most algorithms will couple many qubits and not respect the region boundaries defined for crosstalk detection. However, if there are certain subcircuits, or *circuit motifs*, that occur repeatedly in an algorithm or application, these can be used to define regions and subcircuits for crosstalk detection. It should be noted though that this approach will not detect crosstalk errors that might occur *only* when these circuit motifs are all put together into an application circuit. Detecting such errors requires an application- or architecture-specific test.

But we have intentionally assumed no specific model and no specific application. In the absence of any other guidance, *random* circuits – like those used in randomized benchmarking – are a sensible choice. These have certain known drawbacks; they are less sensitive than periodic circuits (*e.g.*, those used in GST or robust phase estimation) to some forms of noise because they twirl it [14, 20, 7]. But random circuits are both common and hard to fool – their sensitivity to noise is not always high, but it is reliable. Therefore, we propose that the bag for each region be constructed by choosing N_{circ} subcircuits uniformly at random from an ensemble of random sequences of the processor's elementary gates. The simulations presented in Sec. 7 use this construction.

6.4 Analyzing data to detect crosstalk errors

The experiment described in Sec. 6.3 generates data, which can be analyzed to detect and quantify crosstalk errors in a QIP. In this section we explain how to do this.

Running the circuits described above generates samples from a joint probability distribution of settings and results over the M regions, $P(\mathsf{S}_{r_0}, \mathsf{S}_{r_1}, ..., \mathsf{S}_{r_{M-1}}, \mathsf{R}_{r_0}, \mathsf{R}_{r_1}, ..., \mathsf{R}_{r_{M-1}})$. Testing this joint distribution for violations of the conditions in Eq. (9) enables detecting whether where are crosstalk errors between any of the regions in the QIP. But we can go further, by determining the *structure* of the crosstalk errors – *i.e.*, which pairs of regions experience crosstalk errors. This can be achieved using techniques from causal inference that discover conditional dependence relationships between the $2M$ variables in this distribution. Specifically, we show how to adapt techniques developed to learn causal structure in Bayesian networks [41, 27], to detect the

structure of crosstalk errors.

A Bayesian network is a directed graph where each node represents a random variable and the edges represent joint probabilistic relationships between the variables. It is a concise representation of the joint distribution over the variables, with an edge indicating a conditional dependence between the variables – an edge from node i to node j indicates that variable j is dependent on variable i, when conditioned on the other variables in the graph. This is notated $(i \not\perp j) \mid A$, where A is a set representing all other nodes/variables in the graph.

Identifying causal network structure from data is an active and rapidly evolving area of research in the field of causal inference, and there are many algorithms available to do such causal network discovery [55]. These algorithms fall into two broad classes. The first, termed *search-and-score methods*, enumerate or search through graph structures (each of which corresponds to a particular form of the joint distribution over the variables) and evaluates the how well each fits the data according to a score (which is often its likelihood [4]). The second class of algorithms, referred to as *constraint-based methods*, operate by reconstructing a graph that is consistent with the conditional dependencies seen in the data by performing a series of hypothesis tests. Search-and-score methods are typically very computationally expensive, especially for datasets with a large number of variables, so we focus on constraint-based algorithms in this work.

Any constraint-based algorithm for causal network structure discovery can be split into two parts: (i) a statistical test that tests for conditional independence (CI) of some sets of random variables, given samples and at a level of statistical significance, and (ii) a network discovery algorithm that repeatedly applies this CI test to determine the edges in the network. The keys to developing a "good" network learning algorithm are to formulate a CI test that is efficient and powerful, and to formulate a network discovery algorithm that is efficient, in the sense of needing to applying as few CI tests as possible. Given the mature body of research in this field, we seek to apply a previously developed network discovery algo-

rithm to reveal the conditional dependence structure between the random variables we have in the context of crosstalk error detection. In the following subsections we present specific choices for the CI test and network discovery algorithm.

We emphasize that although we are using tools traditionally used in causal inference, we are not making claims about causality. Specifically, an edge between nodes S_{r_i} and R_{r_j} (or R_{r_i} and R_{r_j}) for $i \neq j$ does not imply a direct causal relationship between the regions r_i and r_j, just that there is some crosstalk error between these regions. This is an important caveat. Even in the context of classical physics, it is well known that statistical causal discovery algorithms are only heuristics for revealing causal relationships (especially in the presence of latent, or unobserved, variables) [54, 27, 55]. In quantum theory, even defining causality and a definite causal order between random variables is thorny [9, 61]. So we emphasize that we are simply using causal inference tools to efficiently assess conditional independence relationships that form the basis of our model-free definition of crosstalk errors.

6.4.1 Statistical tests for conditional independence

There are many statistical tests for conditional independence. In the protocol described in Sec. 6.3 the random variables of interest represent experimental settings and measurement outcomes. Both are drawn from a finite set. Therefore, all random variables in a data set resulting from such experiments will be *categorical*. For such variables, a well-motivated test for conditional independence is the log-likelihood ratio test, or G^2 test [4].

To describe the test statistic associated with this test, let us first describe the data. The dataset consists of samples from K random variables $\mathbb{X} = \{\mathsf{X}_k\}_{k=0}^{K-1}$ some of which represent experimental settings (S_{r_i}) and some of which represent measurement outcomes (the R_{r_i}). We assume that each X_k takes values from a finite set $\mathbb{X}_k$ of size $|\mathbb{X}_k|$. Then the G^2 test statistic that tests for the conditional dependence between variable X_i and X_j, conditioned on the variables in the set $A \subset \mathbb{X}$ is defined as [4]

---

[4]In practice one uses an information criterion as opposed to just the likelihood in order to avoid overfitting the data.

$$G^2(i,j \mid A) = 2 \sum_{x_i, x_j, x_A} n_{ijA}(x_i, x_j, x_A) \log \frac{n_{ijA}(x_i, x_j, x_A) n_A(x_A)}{n_{iA}(x_i, x_A) n_{jA}(x_j, x_A)}, \tag{11}$$

where $n_{ijA}(x_i, x_j, x_A)$ is the frequency of the random variables $(\mathsf{X}_i, \mathsf{X}_j, \mathsf{X}_A)$ taking on the values (x_i, x_j, x_A) in the dataset, and similarly for the other quantities. Note that X_A is a composite random variable since one may want to condition on several variables, $i.e.$, $|A| > 1$. Under the null hypothesis, where $(\mathsf{X}_i \perp\!\!\!\perp \mathsf{X}_j) \mid A$, this test statistic is asymptotically distributed as chi-squared with degrees of freedom $df = (|\mathbb{X}_i| - 1)(|\mathbb{X}_j| - 1)(|\mathbb{X}_A|)$. This test statistic is a scaled version of the empirical estimate of the conditional mutual information between variables X_i and X_j, given A. Thus this quantity also has a convenient information theoretic interpretation [4].

Finally, we note that in the simplest case where the conditioning set is null, $A = \emptyset$, this statistical test is often referred to as a *homogeneity* or independence test (with $df = (|\mathbb{X}_i| - 1)(|\mathbb{X}_j| - 1)$) since it tests whether the distribution of variable i is the same (homogeneous) regardless of the value of the variable j.

6.4.2 Network discovery algorithms

The second part of a constraint-based causal network structure learning algorithm applies a CI test on data to reconstruct a network consistent with the data. The PC algorithm by Spirtes and Glymour [53] is a popular network discovery algorithm that has been widely implemented and tested. Appendix D has a detailed description of the algorithm, but here we outline its basic steps. The PC algorithm starts with a complete undirected graph with edges between all nodes (each of which represents a variable in the dataset). Then each edge is tested for conditional independence, given some conditioning set A comprising neighbors of the nodes connected by the edge, for conditioning sets of increasing size (starting from an empty set). The resulting undirected graph is called the *skeleton*, and the last step applies certain edge orientation rules in order to estimate a directed acyclic graph (DAG) representing the causal relations in the data.

For crosstalk error detection, we will omit the last, edge orientation, step of the PC algorithm and will focus on the graph skeleton. We do this because we are not interested in identifying causal relationships (for reasons mentioned at the beginning of the section) and simply wish to detect conditional dependence relationships that signal violation of the crosstalk-free model.

In the worst case, the runtime of the PC algorithm grows exponentially with the number of variables. However, graph sparsity greatly reduces computational cost, and the algorithm has been demonstrated on data with hundreds and thousands of variables [30]. Furthermore, detecting crosstalk errors is simpler than general causal network learning, because we can enforce some sparsity by encoding physically motivated information into the graph from the start. For example, the edge between any two experimental settings can be removed if they are randomized according to the experiment design outlined in Sec. 6.3.

The PC algorithm performs multiple hypothesis tests to determine conditional independence relationships between random variables. In such multiple hypothesis testing scenarios one typically applies a significance adjustment, such as the Bonferroni correction, to control the number of false positives (type-I errors). These corrections are not done in the standard PC algorithm, because controlling the family-wise error rate is complicated by the structure of the PC algorithm: one does not know how many hypothesis tests will be performed *a priori*. However, we note that there have been recent attempts to incorporate statistical methods for controlling the false discovery rate by modifying the PC algorithm [31, 56]. Implementing this more complex algorithm may increase the reliability and statistical rigor of the crosstalk error detection protocol. Alternatively, α-significance weak control of the family-wise error rate [5] can be maintained by setting the input significance of the standard algorithm to α/K where K is the number of edges in the initial graph.

---

[5]Weak control of the family-wise error rate with a significance level of α means that the probability of rejecting one or more null hypothesis is at most α when all the null hypotheses are true. It is more common to demand strong control of the family-wise error rate at significance α, meaning that the probability of rejecting one or more true null hypothesis is at most α regardless of which of the hypotheses are true.

6.4.3 Quantifying crosstalk errors

Applying the PC algorithm to a dataset reporting the experimental settings and measurement outcomes for regions will reconstruct a graph whose edges can be used to detect crosstalk errors at a specified significance level. However, we can also use this analysis to statistically quantify the amount of crosstalk error across any edge that represents crosstalk.

Let the edge that represents crosstalk in a reconstructed graph be between variables X and Y, *i.e.*, X → Y. Recall that X takes values in the set $\{x_0, ... x_{|\mathbb{X}|-1}\}$ and Y takes values in $\{y_0, ... y_{|\mathbb{Y}|-1}\}$. We compute the following total variation distance (TVD) estimates

$$d_{ij}^{\mathsf{X} \to \mathsf{Y}}$$
$$= \sum_{z=y_0}^{y_{|\mathbb{Y}|-1}} \left| \frac{n_{xy}(x_i, z)}{\sum_{z=y_0}^{y_{|\mathbb{Y}|-1}} n_{xy}(x_i, z)} - \frac{n_{xy}(x_j, z)}{\sum_{z=y_0}^{y_{|\mathbb{Y}|-1}} n_{xy}(x_j, z)} \right|,$$

for $0 \leq i, j \leq |\mathbb{X}| - 1$. This quantity is a measure of the difference between the distribution of Y when $X = x_i$ and when $X = x_j$. We quantify the amount of crosstalk error across the edge X → Y as the maximum over all i, j, since this represents the maximum deviation in the distribution of Y when X is varied:

$$\mathcal{C}_{\mathsf{X} \to \mathsf{Y}} = \max_{i,j} d_{ij}^{\mathsf{X} \to \mathsf{Y}}. \tag{12}$$

Often, we also calculate the median over these TVDs to understand how much of an outlier the maximum is.

One has to be a little careful with this definition when Y is a result random variable and X is a setting random variable. To see this, suppose $X = \mathsf{S}_1$ and $Y = \mathsf{R}_2$. Then, the most sensible thing is to compare the distributions of R_2 generated by the same setting S_2, as S_1 is varied. This requires that S_2 take on a value s when $\mathsf{S}_1 = i$ and $\mathsf{S}_2 = j$ in the above definition. In this case, we calculate the above TVD for every such common setting S_2 for a pair $\mathsf{S}_1 = i$ and $\mathsf{S}_2 = j$, and maximize over these. If no such common settings exist (which can happen for example, in the experiment specified in Sec. 6.3 since it is a randomized design), then we fail to compute a TVD for that edge.

6.5 Discussion and limitations

The crosstalk error detection protocol developed in the previous subsections is efficient in terms of experiment number and has tractable post-processing complexity for QIPs with hundreds of qubits [30]. However, we have made several assumptions and restrictions in order to obtain this efficiency. Our assumptions stem from the fact that we are targeting low-weight crosstalk errors, see Sec. 5. This greatly restricts the set of realistic crosstalk errors, and we concentrate on detecting these. Here we discuss the implications of our assumptions, and the associated limitations of the protocol.

Given a circuit layer in an n-qubit Markovian QIP that has crosstalk errors (or a family of layers for relative crosstalk), the question of whether our protocol will detect it or not is dictated by the following factors:

1. The partition of the QIP into regions, since the protocol detects crosstalk errors across regions.

2. Whether the particular layer(s) that exhibit the crosstalk error is (are) sampled in the lightweight experiment design.

3. Whether the detection procedure, using a network discovery algorithm and pairwise conditional independence tests is sufficient to detect the error.

4. Statistical power; do we collect enough samples to determine the signal from noise?

In the following, we will discuss each of these in turn.

Factor 1: We partition a QIP into regions based on the elementary operations in the device and this implies a poly(n) number of necessary partitions. Such a partitioning ignores regions of larger size that are composed of $k > 2$ qubits. Operations on such regions will be composed out of one- and two-qubit operations, and therefore any crosstalk error would still be generated by the elementary operations. The assumption behind ignoring partitions with such "composite" regions is that any crosstalk error present between such regions will also be present between some set of regions composed of $k \leq 2$ qubits.

More fundamentally, a crosstalk error that exists between a region with $k - 1$ qubits and another with one qubit, and does not exist when the $k - 1$-qubit region is sub-partitioned in any way, must be a weight-k crosstalk error. Given this, by choosing regions of size at most 2 as suggested above, we are accepting that we have no guarantee of detecting crosstalk errors of weight 4 or higher.

Factor 2: Since the lightweight experiment design is based on random sampling, L, N_{circ} and

N_{con} dictate the probability that any particular layer will be present in the crosstalk detection experiment. There are an exponential number of possible layers – if there are g elementary gates that can be applied to each region in an M-region QIP, this results in g^M possible layers that can be executed in this QIP. Therefore, in order to guarantee that any possible layer is included in the experiment with high probability, L or N_{circ} are required to scale exponentially in M. Our lightweight experiment does not provide this guarantee. However, such a guarantee of including every possible layer should be unnecessary for realistic quantum computing architectures, where if crosstalk is present in one layer it is also present in many others since the source of the crosstalk is one or several of the operations within a layer, and these are present in exponentially many layers. It is easy to imagine adversarial crosstalk error models that do not fit this bill – *e.g.*, there are crosstalk errors on qubit 1 if and only if there is an X_π gate is applied on *all* the other qubits. Our protocol would almost certainly not detect this crosstalk error because of the low probability of sampling this particular layer. However, this is an adversarial error model that is high-weight (since the operation on qubit 1 is conditioned on the classical register recording the operation applied to all the other qubits). We sacrifice detecting such crosstalk errors in order to derive an efficient protocol.

Factor 3: Using the PC algorithm to identify crosstalk structure in a QIP implies some subtle assumptions about the crosstalk errors. To clarify these, we first note that the PC algorithm is known to fail to detect causal network structure when the probability distribution being sampled from is not *faithful* to the underlying causal graph [54, 26, 55]. In our context, faithfulness means that if there exists crosstalk between regions r_i and r_j, then there exist at least some random variables in r_i that exhibit dependence to some random variables r_j, vice versa, or both. The classic example [28, 26] where the faithfulness assumption is violated and the PC algorithm fails is with three random variables $\mathsf{X}_1, \mathsf{X}_2, \mathsf{X}_3$, that are pairwise independent; *e.g.*, if $\mathsf{X}_1, \mathsf{X}_2, \mathsf{X}_3$ are binary, and $\mathsf{X}_3 = \mathsf{X}_1 \oplus \mathsf{X}_2$. This means that $\mathsf{X}_i \perp\!\!\!\perp \mathsf{X}_j$, for any i, j, but $(\mathsf{X}_i \not\!\perp\!\!\!\perp \mathsf{X}_j) \mid \mathsf{X}_k$ (for $i \neq j \neq k$). So each pair X_i and X_j are *conditionally dependent*, but *marginally independent*. The PC algorithm's first step tests each pair of variables for marginal independence [54]. This step would indicate that all pairs are marginally independent, and therefore all edges

would be removed and the algorithm would terminate. Therefore the PC algorithm evaluated on samples from this distribution (even in the infinite sample size limit) would yield a graph with three nodes and no edges, despite the fact that these variables are clearly dependent. An analogue of this example in the context of crosstalk detection in QIPs is the following: suppose one is trying to detect crosstalk caused by single qubit gates in an n-qubit QIP. The regions are composed of single qubits, and suppose that the crosstalk errors are such that with the circuits that are tested, one ends up preparing an entangled state of the n qubits with any two-local marginal density matrix that is completely mixed (*e.g.*, multiparty data hiding states [22]). Then the results of measuring any qubit will be uncorrelated with the results from any other qubit (if all other measurement results are ignored) and the PC algorithm would not indicate any crosstalk between regions. The basic problem is that the marginal/local states do not produce distributions over measurement outcomes that are faithful to the underlying dependence (and correlation) between local subsystems. Testing the dependence (or correlation) between a large number of subsystems would reveal strong dependence. But the PC algorithm orders its tests by increasing number of variables (increasing size of conditioning set) for efficiency, and declares two variables to be independent as soon as it fails to detect a dependence. Therefore, it would never perform the necessary tests to reveal the dependence, which is also the root cause of the failure in the pairwise independent, three variable example given above.

Fortunately, producing unfaithful distributions over the random variables in the crosstalk error detection setting appears to be extremely artificial. Every case where we have been able to manufacture such distributions requires either (i) high-weight crosstalk errors acting non-trivially on several regions, (ii) extremely large crosstalk errors (*e.g.*, errors causing $\frac{\pi}{2}$-rotations), or (iii) fine-tuned crosstalk that cancels or adds up in precise ways. Moreover, we have not encountered this issue in any of the physically-relevant crosstalk error models that we have simulated. Therefore, we note it as an issue to be aware of when using the PC algorithm, but something that does not seem to practically affect the performance of the crosstalk detection protocol developed here. Moreover, we note that the PC algorithm is not the only option for the graph discovery portion of the protocol; it is possible to apply other approaches to detect the condi-

Accepted in 〈 〉uantum 2020-09-07, click title to verify. Published under CC-BY 4.0.

20

tional dependency relationships between the operational variables [55], although we have not explored these.

Factor 4: If the experiment design and data analysis technique are sufficient to detect the faulty layer, the statistical power of each of the hypothesis tests underlying the network discovery algorithm increases with N_{rep} (*i.e.*, the distribution of the test statistic narrow with sample size). However, as mentioned above, since the overall analysis involves an *a priori* unknown number of hypothesis tests, it is difficult to estimate the detection accuracy of the whole procedure as a function of sample size. Hence, we simply advice as large an N_{rep} as possible to minimize statistical error.

6.6 Guidance for selecting protocol parameters

This protocol has several user-adjustable parameters. Their values can be chosen, but not arbitrarily – they control the reliability and power of the experiment. Here, we provide some heuristic guidance on how to choose them.

1. N_{rep} is the number of repetitions of each experiment. Increasing N_{rep} reduces statistical noise, at the cost of requiring more time to take data. We suggest that this should be as large as possible, and no less than 1000.

2. L is the the length or depth of the circuits, and two useful rules of thumb suggest what L should be. The first is that longer circuits (all else being equal) can exhibit increased crosstalk effects and therefore permit more sensitive detection. However, once L becomes greater than $1/\epsilon$, where ϵ is the rate of stochastic errors or decoherence, generic noise tends to swamp the effect sought. Therefore, L should be as large as feasible, but no greater than $O(1/\epsilon)$.

3. N_{circ} is the number of circuits in each region's bag, which in turn are randomly selected from the population of all depth-L subcircuits on each region. This parameter can be chosen to be a constant, independent of M, of order $10 - 30$. At a minimum, it needs to be large enough to guarantee that all possible elementary gates that can be performed in a region appear in at least one of the subcircuits chosen for that region.

4. N_{con} is the number of random contexts in which each subcircuit is intentionally performed. Empirically, we find that it should be $\mathcal{O}(N_{circs})$ –

but the best value for this parameter depends on the relative strength of crosstalk errors and local errors, which we refer to as *signal-to-noise ratio*. When crosstalk errors are comparable to local errors (*low* signal-to-noise), we require $N_{con} \sim N_{circs}/2$. But if crosstalk errors dominate (*high* signal-to-noise), we find that $N_{con} \sim N_{circs}/4$ is sufficient.

5. p_{idle} is the probability of sampling the length L idle circuit on any of the $M - 1$ regions when constructing a context. The recommended value of p_{idle} depends on whether the idle operation has a significantly lower local error rate than other operations. If it does, then we recommend choosing $p_{idle} \sim 1/M$, so that there is probability $\sim (M - 1)/M \approx 1$ for large M that the idle circuit is among the contexts provided. Otherwise, p_{idle} should be smaller – but even in this case, we find that $p_{idle} > 0$ is often advantageous, although we do not have a good rule of thumb for how the optimal value varies with the idle error rate.

As mentioned, the guidance for these parameters is based on empirical studies (except for the guidance for L, which is fairly standard in quantum benchmarking). It might be possible to develop more rigorous estimates for these parameters based on an analysis of the statistical convergence of the PC algorithm. However, the complexity of the PC algorithm makes analysis of its convergence difficult [25, 56], and hence we leave this as an avenue for potential future work.

is fairly standard, the guidance for N_{con} and p_{dile} is

7 Simulations

In this section we illustrate the crosstalk error detection and quantification procedure developed above by simulation. The analysis of simulated data is performed using crosstalk error detection routines in the pyGSTi package [1] that implement the PC algorithm as described above.

7.1 Two-qubit simulations

We first simulate several kinds of crosstalk error on a two-qubit system, with qubits (which form the regions in this case) labeled 0 and 1. The settings for each qubit enumerate the subcircuits applied (which

are just gate sequences in this case since each region is composed of a single qubit), and the experiments simulated correspond to the experiment design outlined in Sec. 6.3. In addition to the crosstalk error models, we also simulate local errors through a local depolarization channel (after every gate, including the idle) with depolarization rate p_{local}. To illustrate the efficacy of the technique, in all the simulations below, we operate in the low signal-to-noise regime where the local error rates are comparable or larger than the crosstalk error rates. This is where we expect that it is most challenging to detect the crosstalk errors.

The elementary one-qubit gates are assumed to be $X_{\pi/2}, Y_{\pi/2}, I$, where I is an idle or identity gate that takes the same time as the other gates. The native state preparation is always ideally the $|0\rangle$ state for both qubits, and the measurements are in the computational basis. The gate sequences are determined according to the experiment design outlined in Sec. 6.3. In all of the simulated experiments, we follow the guidance in Sec. 6.6 and use the suggested value $N_{\text{con}} = N_{\text{circs}}/2$ (since the parameters chosen are in the low signal-to-noise regime). The values of $N_{\text{circs}}, N_{\text{rep}}, L$ and p_{idle} vary and are specified below for each case.

Finally, since the regions in this case are composed of single qubits, in this section we simplify notation and dispense with the additional r subscript when denoting results and settings; *i.e.*, $\mathsf{S}_{r_i} \rightarrow \mathsf{S}_i$ and $\mathsf{R}_{r_i} \rightarrow \mathsf{R}_i$.

7.1.1 Operation crosstalk error 1

The first error model we simulate is what we refer to as operation crosstalk error, and is also sometimes termed classical, or control line, crosstalk in literature. An $X_{\pi/2}$ gate on qubit 0 induces a depolarization channel on qubit 1 with depolarization rate p, *i.e.*,

$$\mathcal{X}_{\pi/2} \otimes \mathcal{I} \rightarrow \mathcal{X}_{\pi/2} \otimes \mathcal{D}_p, \qquad (13)$$

where $\mathcal{X}_{\pi/2}(\rho) = e^{-i\frac{\pi}{4}\sigma_x}\rho e^{i\frac{\pi}{4}\sigma_x}$ denotes a superoperator representation of a $X_{\pi/2}$ unitary rotation, $\mathcal{I}$ denotes an identity superoperator, and $\mathcal{D}_p(\rho) = (1-p)\rho + pI$ denotes a depolarization channel with depolarization probability p.

Fig. 5(a) shows the reconstructed crosstalk graph for this error model. The error model parameters used are: $p = 10^{-2}, p_{\text{local}} = 10^{-2}$. The parameters defining the simulated experiment are $L = 30, N_{\text{circs}} = 10, p_{\text{idle}} = 0.1, N_{\text{rep}} = 10^4$. The maximum number

of unique circuits is $N_{\text{exp}} = M \times N_{\text{circs}} \times \frac{N_{\text{circs}}}{2} = 100$. Finally, the significance level of the hypothesis tests used to test for conditional independence was set to $\alpha = 0.01$. The red edge between settings in region 0 and results in region 1 in the graph signals the crosstalk between the qubits.

7.1.2 Operation crosstalk error 2

The next error model we simulate is an example of what is sometimes called coherent, or Hamiltonian, crosstalk. We model an $X_{\pi/2}$ gate on qubit 0 as inducing the desired rotation on qubit 0, but with an additional small two-qubit $Z \otimes Z$ Hamiltonian rotation as well, *i.e.*,

$$X_{\pi/2} \otimes I \rightarrow \exp\left(-\frac{i}{2}\left[\frac{\pi}{2}X \otimes I + \frac{\epsilon}{2}Z \otimes Z\right]\right). \qquad (14)$$

Fig. 5(b) shows the reconstructed crosstalk graph for this error model. The error model parameters used are: $\epsilon = 2 \cdot 10^{-2}, p_{\text{local}} = 10^{-2}$. The parameters defining the simulated experiment are $L = 30, N_{\text{circs}} = 10, p_{\text{idle}} = 0, N_{\text{rep}} = 10^5$ (therefore, $N_{\text{exp}} = 100$), and the significance level of the hypothesis tests was set to $\alpha = 0.01$. Note that the coherent crosstalk error shows up at $\sim \epsilon^2$ in the measurement probabilities since we are using random gate sequences, and this is why more samples are required to detect this crosstalk error.

The red edges in the crosstalk graph indicate crosstalk errors between the qubits. In this case there are conditional dependencies between settings and results in different regions, and *also* between the results in different regions. There is no clear causal direction for this type of crosstalk error (and one can show using a model of this kind of crosstalk error and calculations such as in Appendix B that conditional dependencies between results are expected for this kind of crosstalk error).

7.1.3 Detection crosstalk error

The final error model we simulate is a model of crosstalk during the qubit measurement process. The measurement effects, indexed by the outcome values, are:

$$\begin{aligned}
E_{00} &= |00\rangle\langle 00| \\
E_{01} &= |01\rangle\langle 01| \\
E_{10} &= (1-p_{\text{m}})|10\rangle\langle 10| + p_{\text{m}}|11\rangle\langle 11| \\
E_{11} &= (1-p_{\text{m}})|11\rangle\langle 11| + p_{\text{m}}|10\rangle\langle 10|
\end{aligned}$$
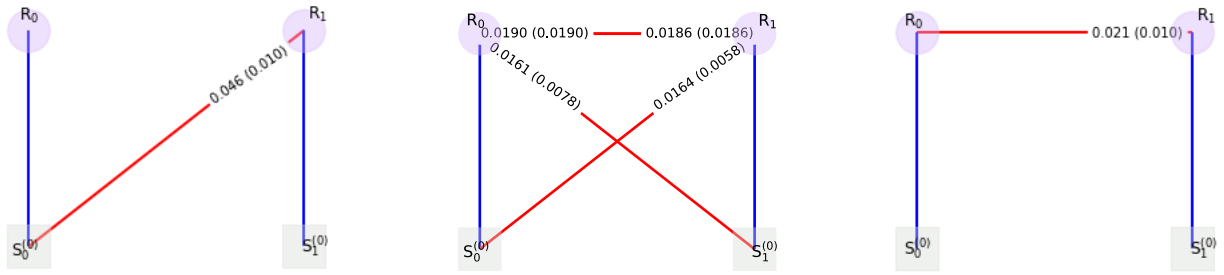
(a) Operation crosstalk error 1   (b) Operation crosstalk error 2   (c) Detection crosstalk error

Figure 5: Reconstructed graphs for various crosstalk error models in a systems of two qubits; see Sec. 7.1 for details of error models. The regions in this case are composed on one qubit each. R_i represents the measurement result on qubit i and $S_i^{(0)}$ represents the setting on qubit i (the superscript (0) indexes the settings for a region – in all our examples there is only one setting per region since only the applied gate sequence is varied). The blue edges indicate conditional dependencies between variables that are expected (*i.e.*, both variables belong to the same region). The red edges indicate conditional dependencies between variables in different regions, and these represent crosstalk. The red edges are labeled with the maximum TVD (and median TVD in parentheses) for that conditional dependence (see main text for definitions of these quantities).

In other words, if the measured value for the first qubit is 1, there is a p_{m} probability that the measured value of the second qubit is flipped. This could, for example, model detection crosstalk due to scattered photons that flip the neighboring qubit state.

Fig. 5(c) shows the reconstructed crosstalk graph for this error model. The error model parameters used are: $p_{\mathrm{m}} = 10^{-2}, p_{\mathrm{local}} = 10^{-2}$. The parameters defining the simulated experiment are $L = 10$, $N_{\mathrm{circs}} = 20, p_{\mathrm{idle}} = 0, N_{\mathrm{rep}} = 10^5$ (therefore, $N_{\mathrm{exp}} = 400$), and the significance level of the hypothesis tests was set to $\alpha = 0.01$. Unlike the previous crosstalk error models, the effects of this error do not potentially build up over a gate sequence, and thus only impact the outcome probabilities weakly. Moreover, its effect is reduced in the experiments where the first qubit's outcome is 0 with high probability. For these reasons, we found that a larger N_{rep} and N_{circs} are required to detect this crosstalk error.

The reconstructed graph in this case shows a red edge between the results on the two qubits indicating a conditional dependence that should not exist without some form of crosstalk error.

7.1.4 Crosstalk error quantification

It is important to keep in mind that the weights on the edges of a crosstalk graph are estimated maximum TVDs of outcome distributions, and not necessarily physical quantities like error rates. To illustrate this, in Fig. 6 we return to the second operation crosstalk error model in Sec. 7.1.2 and plot the weight of the edge from R_0 to R_1 and S_1 to R_0 as the amount of

crosstalk, *i.e.*, the magnitude of the coherent $Z \otimes Z$ coupling term, is varied. The experiment sampling and physical parameters are the same as in Sec. 7.1.2, except that we use $N_{\mathrm{circs}} = 5$. We see that while the max TVD increases (up to statistical variation) with increasing ϵ for one edge, it does not for the other. Even if the max TVD should vary monotonically with some crosstalk parameter when computed over all random circuits, in practice it is a function of the experiments that are sampled and therefore sensitive to finite sampling variations in these experiments.

Therefore, the maximum TVD quantification should not be thought of as a direct measure of the *physical degree* of crosstalk. It should instead be used as a way to identify the regions of a multiqubit device that require the most attention in terms of needing crosstalk mitigation. In addition, examining the experimental configuration that led to this maximum TVD often lends insight into the source of crosstalk errors.

7.2 Six-qubit simulations

In this section we illustrate the scalability of the crosstalk error detection protocol by simulations on a 6-qubit device. The hypothetical device has a ladder layout, as shown in Fig. 7(a) and we are interested in detecting the crosstalk errors caused by single qubit gates. So we partition the QIP into six regions with a single qubit in each. The settings for each qubit enumerate the gate sequences applied, and the experiments simulated correspond to the experiment design outlined in Sec. 6.3.
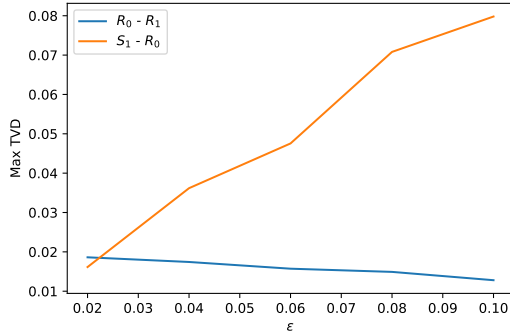
Figure 6: The weight of the edge from R_0 to R_1 and S_1 to R_0 versus the the physical crosstalk error magnitude, ϵ, for the second operation crosstalk error model detailed in Sec. 7.1.2.

The crosstalk error model is similar to the first operation crosstalk error model detailed in Sec. 7.1.1; all single qubit gates on any of the qubits in the bottom line (qubits 3,4,5) result in a depolarizing channel with depolarization rate p on the vertical neighboring qubit. In addition to these crosstalk errors we also simulate local errors through a depolarization channel with depolarization rate p_{local} after every gate and rate p_{idle} after every idle clock cycle. All other details (elementary gate set, state preparation and measurement, and form of experimental gate sequences used) are the same as in the two-qubit simulations.

Fig. 7(b) shows the reconstructed crosstalk graph for this error model. The parameters used were: $p = 10^{-2}, p_{\text{local}} = 10^{-2}, p_{\text{idle}} = 5 \times 10^{-3}, N_{\text{rep}} = 10^4$. The simulated experiment used $l = 20$, $N_{\text{circs}} = 10$ (resulting in $N_{\text{exp}} = 300$), $p_{\text{idle}} = 0.1$, and the significance level of the hypothesis tests was set to $\alpha = 0.01$.

The red edges in the crosstalk graph indicate crosstalk between the qubits that are vertical neighbors, as expected. We emphasize that this accurate crosstalk detection is achieved with just 300 distinct experiments, which highlights the benefits of using a technique with experimental burden that scales essentially linearly with the number of qubits.

8   Conclusions

We make two contributions in this paper. First, we provide a universal and hardware-agnostic definition of crosstalk errors in terms of a model for QIP dynamics based on representations of gates, state preparations and measurements on the device. Second, we provide a model-free definition of crosstalk in terms of operational variables (QIP settings and measurement results), and develop a protocol for detecting crosstalk errors based on it.
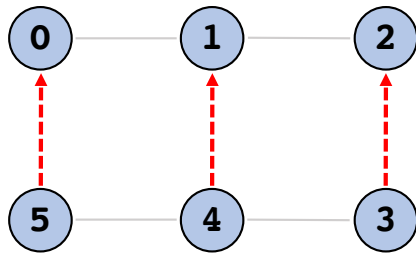
The protocol is based on testing conditional independence relations between a potentially large number of random variables, and targets detection of low-weight crosstalk errors, which are a major concern for existing QIPs. We have tested the protocol and associated data processing on simulated experiments on QIPs with up to six qubits. The technique shows promise for crosstalk error detection on medium-scale QIPs since it requires a number of experiments that scales as $\tilde{\mathcal{O}}(n^3)$ in the worst-case, and scales as $\tilde{\mathcal{O}}(n^2)$ in realistic scenarios where qubit connectivity is limited.

An avenue for future research is to explore the utility of alternatives to the PC algorithm for discovering the crosstalk structure in a QIP. The PC algorithm is arguably the most established constraint-based algorithm for causal network structure discovery, but there is an active field of study developing new approaches to causal network structure discovery, *e.g.*, the new kernel-based learning methods in Refs. [33, 40], and it would be interesting to study whether any of these present any advantages when post-processing lightweight experimental data for crosstalk error detection.
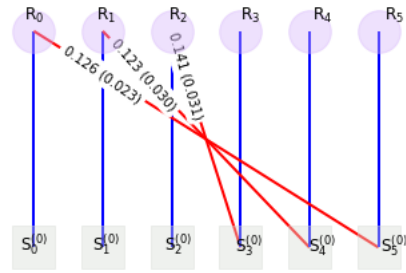
Of course, detecting crosstalk is just the first step. One would ideally like to also characterize crosstalk errors once detected in order to learn their form and possibly also their origin. In future work we will utilize the model-based definition of crosstalk developed here to construct efficient protocols for characterizing crosstalk errors.

Acknowledgements

Accepted in ⟨ Quantum 2020-09-07, click title to verify. Published under CC-BY 4.0.

24

(a) 6 qubit QIP layout

(b) Reconstructed crosstalk graph

Figure 7: Reconstructed graphs for 6 qubit QIP with operation crosstalk errors along vertical lines; see Sec. 7.2 for details of error models. The regions in this case are composed on one qubit each. R_i represents the measurement results on qubit i and $S_i^{(0)}$ represents the setting on qubit i (the superscript (0) indexes the settings for a region – in all our examples there is only one setting per region since only the gate sequence applied is varied). The blue edges indicate conditional dependencies between variables that are expected (*i.e.*, both variables belong to the same region). The red edges indicate conditional dependencies between variables in different regions, and this represent crosstalk. The red edges are labeled with the maximum TVD (and median TVD in parentheses) for that conditional dependence (see main text for definitions of these quantities).

References

[1] PyGSTi. A python implementation of Gate Set Tomography. URL: http://www.pygsti.info/.

[2] IBM Qiskit/ibmq-device-information/tenerife, 2019. URL: https://github.com/Qiskit/ibmq-device-information/tree/master/backends/tenerife/V1.

[3] Carole Addis, Francesco Ciccarello, Michele Cascio, G Massimo Palma, and Sabrina Maniscalco. Dynamical decoupling efficiency versus quantum non-markovianity. *New Journal of Physics*, 17(12):123004, 2015. doi:10.1088/1367-2630/17/12/123004.

[4] D Bacciu, T A Etchells, P J G Lisboa, and J Whittaker. Efficient identification of independence networks using mutual information. *Computational Statistics*, 28:621, 2013. doi:10.1007/s00180-012-0320-6.

[5] Robert Beals, Harry Buhrman, Richard Cleve, Michele Mosca, and Ronald de Wolf. Quan-

tum lower bounds by polynomials. *Journal of the ACM*, 48(4):778–797, July 2001. doi:10.1145/502090.502097.

[6] R. C. Bialczak, M. Ansmann, M. Hofheinz, E. Lucero, M. Neeley, A. D. O'Connell, D. Sank, H. Wang, J. Wenner, M. Steffen, A. N. Cleland, and J. M. Martinis. Quantum process tomography of a universal entangling gate implemented with Josephson phase qubits. *Nature Physics*, 6(6):409–413, 2010. doi:10.1038/nphys1639.

[7] Robin Blume-Kohout, John King Gamble, Erik Nielsen, Kenneth Rudinger, Jonathan Mizrahi, Kevin Fortier, and Peter Maunz. Demonstration of qubit operations below a rigorous fault tolerance threshold with gate set tomography. *Nature Communications*, 8:1, 2017. doi:10.1038/ncomms14485.

[8] Heinz-Peter Breuer, Elsi-Mari Laine, and Jyrki Piilo. Measure for the degree of non-markovian behavior of quantum processes in open systems. *Phys. Rev. Lett.*, 103:210401, 2009. doi:10.1103/PhysRevLett.103.210401.

[9] Caslav Brukner. Quantum causality. *Nature Physics*, 10(4):259–263, 2014. doi:10.1038/nphys2930.

[10] Donovan Buterakos, Robert E. Throckmorton, and S. Das Sarma. Crosstalk error correction through dynamical decoupling of single-qubit gates in capacitively coupled singlet-triplet semiconductor spin qubits. *Physical Re-*

view B, 97(4):045431, January 2018. doi:10.1103/PhysRevB.97.045431.

[11] Yanzhu Chen, Maziar Farahzad, Shinjae Yoo, and Tzu-Chieh Wei. Detector tomography on IBM quantum computers and mitigation of an imperfect measurement. *Physical Review A*, 100(5):052315, November 2019. Publisher: American Physical Society. doi:10.1103/PhysRevA.100.052315.

[12] W. A. Coish and Daniel Loss. Hyperfine interaction in a quantum dot: Non-markovian electron spin dynamics. *Phys. Rev. B*, 70:195340, 2004. doi:10.1103/PhysRevB.70.195340.

[13] Diego Colombo and Marloes H Maathuis. Order-Independent Constraint-Based Causal Structure Learning. *J. Machine Learning Research*, 15:3921, 2014. URL: http://jmlr.org/papers/volume15/colombo14a/colombo14a.pdf.

[14] Christoph Dankert, Richard Cleve, Joseph Emerson, and Etera Livine. Exact and approximate unitary 2-designs and their application to fidelity estimation. *Physical Review A*, 80(1):012304, 2009. doi:10.1103/PhysRevA.80.012304.

[15] E. B. Davies and J. T. Lewis. An operational approach to quantum probability. *Communications in Mathematical Physics*, 17(3):239–260, 1970. doi:10.1007/BF01647093.

[16] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe. Demonstration of a small programmable quantum computer with atomic qubits. *Nature*, 536(7614):63–66, August 2016. doi:10.1038/nature18648.

[17] Alexander Erhard, Joel J. Wallman, Lukas Postler, Michael Meth, Roman Stricker, Esteban A. Martinez, Philipp Schindler, Thomas Monz, Joseph Emerson, and Rainer Blatt. Characterizing large-scale quantum computers via cycle benchmarking. *Nature Communications*, 10(1):5347, 2019. doi:10.1038/s41467-019-13068-7.

[18] Jay M. Gambetta, A D Córcoles, S T Merkel, B R Johnson, John A Smolin, Jerry M Chow, Colm A Ryan, Chad Rigetti, S Poletto, Thomas A Ohki, Mark B Ketchen, and M Steffen. Characterization of Addressability by Simultaneous Randomized Benchmarking. *Phys-*

ical Review Letters, 109(24):240504–5, 2012. doi:10.1103/PhysRevLett.109.240504.

[19] Ming Gong, Ming-Cheng Chen, Yarui Zheng, Shiyu Wang, Chen Zha, Hui Deng, Zhiguang Yan, Hao Rong, Yulin Wu, Shaowei Li, Fusheng Chen, Youwei Zhao, Futian Liang, Jin Lin, Yu Xu, Cheng Guo, Lihua Sun, Anthony D. Castellano, Haohua Wang, Chengzhi Peng, Chao-Yang Lu, Xiaobo Zhu, and Jian-Wei Pan. Genuine 12-Qubit Entanglement on a Superconducting Quantum Processor. *Physical Review Letters*, 122(11):110501, 2019. doi:10.1103/PhysRevLett.122.110501.

[20] D. Gross, K. Audenaert, and J. Eisert. Evenly distributed unitaries: On the structure of unitary designs. *Journal of Mathematical Physics*, 48(5):052104, 2007. doi:10.1063/1.2716992.

[21] Vojtech Havlicek, Antonio D. Corcoles, Kristan Temme, Aram W. Harrow, Abhinav Kandala, Jerry M. Chow, and Jay M. Gambetta. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747):209–212, 2019. doi:10.1038/s41586-019-0980-2.

[22] Patrick Hayden, Debbie Leung, and Graeme Smith. Multiparty data hiding of quantum information. *Physical Review A*, 71(6):062339, 2005. doi:10.1103/PhysRevA.71.062339.

[23] Johannes Heinsoo, Christian Kraglund Andersen, Ants Remm, Sebastian Krinner, Theodore Walter, Yves Salathe, Simone Gasparinetti, Jean-Claude Besse, Anton Potocnik, Andreas Wallraff, and Christopher Eichler. Rapid High-fidelity Multiplexed Readout of Superconducting Qubits. *Physical Review Applied*, 10(3), 2018. doi:10.1103/PhysRevApplied.10.034040.

[24] Miguel A Hernan, David Clayton, and Niels Keiding. The Simpson's paradox unraveled. *International Journal of Epidemiology*, 40(3):780–785, 2011. doi:10.1093/ije/dyr041.

[25] Markus Kalisch and Peter Bhlmann. Estimating High-Dimensional Directed Acyclic Graphs with the PC-Algorithm. *Journal of Machine Learning Research*, 8(Mar):613–636, January 2007. URL: http://www.jmlr.org/papers/v8/kalisch07a.html.

[26] Anna Klimova, Caroline Uhler, and Tams Rudas. Faithfulness and learning hypergraphs from discrete distributions. *Computational Statistics & Data Analysis*, 87:57–72, 2015. doi:10.1016/j.csda.2015.01.017.

[27] D Koller and N Friedman. *Probabilistic Graphical Models*. MIT Press, 2009.

[28] Timo J.T. Koski and John Noble. A Review of Bayesian Networks and Structure Learning. *Mathematica Applicanda*, 40(1):51–103, 2012. doi:10.14708/ma.v40i1.278.

[29] T D Ladd, F Jelezko, R Laflamme, Y Nakamura, C Monroe, and J L O'Brien. Quantum computers. *Nature*, 464:45, 2010. doi:10.1038/nature08812.

[30] Thuc Duy Le, Tao Hoang, Jiuyong Li, Lin Liu, Huawen Liu, and Shu Hu. A Fast PC Algorithm for High Dimensional Causal Discovery with Multi-Core PCs. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 16(5):1483–1495, 2019. doi:10.1109/TCBB.2016.2591526.

[31] Junning Li and Z Jane Wang. Controlling the False Discovery Rate of the Association/Causality Structure Learned with the PC Algorithm. *J. Machine Learning Research*, 10:475, 2009. URL: http://www.jmlr.org/papers/v10/li09a.html.

[32] Li Li, Michael J.W. Hall, and Howard M. Wiseman. Concepts of quantum non-markovianity: A hierarchy. *Physics Reports*, 759:1 – 51, 2018. doi:10.1016/j.physrep.2018.07.001.

[33] David Lopez-Paz, Krikamol Muandet, Bernhard Schölkopf, and Ilya Tolstikhin. Towards a learning theory of cause-effect inference. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37*, ICML'15, pages 1452–1461, 2015. URL: http://dl.acm.org/citation.cfm?id=3045118.3045273.

[34] Ruichao Ma, Brendan Saxberg, Clai Owens, Nelson Leung, Yao Lu, Jonathan Simon, and David I. Schuster. A dissipatively stabilized Mott insulator of photons. *Nature*, 566(7742):51, February 2019. doi:10.1038/s41586-019-0897-9.

[35] Filip B. Maciejewski, Zoltán Zimborás, and Michał Oszmaniec. Mitigation of readout noise in near-term quantum devices by classical post-processing based on detector tomography. *Quantum*, 4:257, April 2020. doi:10.22331/q-2020-04-24-257.

[36] Rupak Majumdar and Filip Niksic. Why is random testing effective for partition tolerance bugs? *Proceedings of the ACM on Programming Languages*, 2(POPL):1–24, January 2018. doi:10.1145/3158134.

[37] S. Mavadia, C. L. Edmunds, C. Hempel, H. Ball, F. Roy, T. M. Stace, and M. J. Biercuk. Experimental quantum verification in the presence of temporally correlated noise. *npj Quantum Information*, 4(1):7, 2018. doi:10.1038/s41534-017-0052-0.

[38] F Mazda. *Telecommunications Engineer's Reference Book*. Elsevier, 1993. doi:10.1016/C2013-0-06529-2.

[39] David C. McKay, Sarah Sheldon, John A. Smolin, Jerry M. Chow, and Jay M. Gambetta. Three-Qubit Randomized Benchmarking. *Physical Review Letters*, 122(20):200502, 2019. doi:10.1103/PhysRevLett.122.200502.

[40] Jovana Mitrovic, Dino Sejdinovic, and Yee Whye Teh. Causal inference via kernel deviance measures. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18, pages 6986–6994, 2018. URL: https://dl.acm.org/doi/10.5555/3327757.3327802.

[41] Richard E Neapolitan. *Learning Bayesian Networks*. Prentice Hall, 2004.

[42] Matthew Neeley, Radoslaw C. Bialczak, M. Lenander, E. Lucero, Matteo Mariantoni, A. D. O'Connell, D. Sank, H. Wang, M. Weides, J. Wenner, Y. Yin, T. Yamamoto, A. N. Cleland, and John M. Martinis. Generation of three-qubit entangled states using superconducting phase qubits. *Nature*, 467(7315):570–573, 2010. doi:10.1038/nature09418.

[43] C. Neill, P. Roushan, K. Kechedzhi, S. Boixo, S. V. Isakov, V. Smelyanskiy, A. Megrant, B. Chiaro, A. Dunsworth, K. Arya, R. Barends, B. Burkett, Y. Chen, Z. Chen, A. Fowler, B. Foxen, M. Giustina, R. Graff, E. Jeffrey, T. Huang, J. Kelly, P. Klimov, E. Lucero, J. Mutus, M. Neeley, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White,

H. Neven, and J. M. Martinis. A blueprint for demonstrating quantum supremacy with superconducting qubits. *Science*, 360(6385):195–199, April 2018. `doi:10.1126/science.aao4309`.

[44] C. Piltz, T. Sriarunothai, A.F. Varon, and C. Wunderlich. A trapped-ion-based quantum byte with 10^{-5} next-neighbour cross-talk. *Nature Communications*, 5(1):4679, December 2014. `doi:10.1038/ncomms5679`.

[45] John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, 2018. `doi:10.22331/q-2018-08-06-79`.

[46] Timothy Proctor, Melissa Revelle, Erik Nielsen, Kenneth Rudinger, Daniel Lobser, Peter Maunz, Robin Blume-Kohout, and Kevin Young. Detecting, tracking, and eliminating drift in quantum information processors. July 2019. arXiv:1907.13608. URL: `http://arxiv.org/abs/1907.13608`.

[47] Timothy Proctor, Kenneth Rudinger, Kevin Young, Mohan Sarovar, and Robin Blume-Kohout. What Randomized Benchmarking Actually Measures. *Physical Review Letters*, 119(13), 2017. `doi:10.1103/PhysRevLett.119.130502`.

[48] Timothy J. Proctor, Arnaud Carignan-Dugas, Kenneth Rudinger, Erik Nielsen, Robin Blume-Kohout, and Kevin Young. Direct Randomized Benchmarking for Multiqubit Devices. *Physical Review Letters*, 123(3):030503, 2019. `doi:10.1103/PhysRevLett.123.030503`.

[49] Matthew Reagor, Christopher B. Osborn, Nikolas Tezak, Alexa Staley, Guenevere Prawiroatmodjo, Michael Scheer, Nasser Alidoust, Eyob A. Sete, Nicolas Didier, Marcus P. da Silva, Ezer Acala, Joel Angeles, Andrew Bestwick, Maxwell Block, Benjamin Bloom, Adam Bradley, Catvu Bui, Shane Caldwell, Lauren Capelluto, Rick Chilcott, Jeff Cordova, Genya Crossman, Michael Curtis, Saniya Deshpande, Tristan El Bouayadi, Daniel Girshovich, Sabrina Hong, Alex Hudson, Peter Karalekas, Kat Kuang, Michael Lenihan, Riccardo Manenti, Thomas Manning, Jayss Marshall, Yuvraj Mohan, William O'Brien, Johannes Otterbach, Alexander Papageorge, Jean-Philip Paquette, Michael Pelstring, Anthony Polloreno, Vijay Rawat, Colm A. Ryan, Russ Renzas, Nick Rubin, Damon Russel, Michael Rust, Diego Scara-belli, Michael Selvanayagam, Rodney Sinclair, Robert Smith, Mark Suska, Ting-Wai To, Mehrnoosh Vahidpour, Nagesh Vodrahalli, Tyler Whyland, Kamal Yadav, William Zeng, and Chad T. Rigetti. Demonstration of universal parametric entangling gates on a multi-qubit lattice. *Science Advances*, 4(2):eaao3603, 2018. `doi:10.1126/sciadv.aao3603`.

[50] Kenneth Rudinger, Timothy Proctor, Dylan Langharst, Mohan Sarovar, Kevin Young, and Robin Blume-Kohout. Probing Context-Dependent Errors in Quantum Processors. *Physical Review X*, 9(2):021045, 2019. `doi:10.1103/PhysRevX.9.021045`.

[51] Sarah Sheldon, Easwar Magesan, Jerry M Chow, and Jay M. Gambetta. Procedure for systematically tuning up cross-talk in the cross-resonance gate. *Phys. Rev. A*, 93(6):060302, 2016. `doi:10.1103/PhysRevA.93.060302`.

[52] Peter Spirtes. Introduction to Causal Inference. *Journal of Machine Learning Research*, 11:1643, 2010. URL: `http://www.jmlr.org/papers/v11/spirtes10a.html`.

[53] Peter Spirtes and Clark Glymour. An Algorithm for Fast Recovery of Sparse Causal Graphs. *Social Science Computer Review*, 9(1):67, 1991. `doi:10.1177/089443939100900106`.

[54] Peter Spirtes, Clark N. Glymour, and Richard Scheines. *Causation, prediction, and search*. MIT Press, Cambridge, Mass, 2nd ed edition, 2000.

[55] Peter Spirtes and Kun Zhang. Causal discovery and inference: concepts and recent methodological advances. *Applied Informatics*, 3:3, 2016. `doi:10.1186/s40535-016-0018-x`.

[56] Eric V. Strobl, Peter L. Spirtes, and Shyam Visweswaran. Estimating and Controlling the False Discovery Rate of the PC Algorithm Using Edge-specific P-Values. *ACM Transactions on Intelligent Systems and Technology*, 10(5):1–37, 2019. `doi:10.1145/3351342`.

[57] L. M. K. Vandersypen and I. L. Chuang. NMR techniques for quantum control and computation. *Reviews of Modern Physics*, 76(4):1037–1069, January 2005. `doi:10.1103/RevModPhys.76.1037`.

[58] Andrzej Veitia, Marcus P. da Silva, Robin Blume-Kohout, and Steven J. van Enk. Macroscopic instructions vs microscopic operations.

2017. arXiv: 1708.08173. URL: http://arxiv.org/abs/1708.08173.

[59] Joel J Wallman and Steven T Flammia. Randomized benchmarking with confidence. *New Journal of Physics*, 16(10):103032, 2014. doi:10.1088/1367-2630/16/10/103032.

[60] Adam Winick, Joel J. Wallman, and Joseph Emerson. Phenomenological measure of quantum non-markovianity. 2019. arXiv:1901.00267. URL: http://arxiv.org/abs/1901.00267.

[61] Christopher J Wood and Robert W Spekkens. The lesson of causal discovery algorithms for quantum correlations: causal explanations of Bell-inequality violations require fine-tuning. *New Journal of Physics*, 17(3):033002, 2015. doi:10.1088/1367-2630/17/3/033002.

Accepted in 〈 〉uantum 2020-09-07, click title to verify. Published under CC-BY 4.0.

29

A Conditional versus marginal independence

Our definition of crosstalk errors is based conditional independence of random variables. As mentioned in the main text, this is motivated by the central role played by conditional independence in defining causality in graphical models. In this Appendix we explain why we prefer to use the notion of conditional independence over marginal independence, which might seem simpler to work with. To do this, we refer to Fig. 8, which represents the random variables involved in a two qubit example: S_1, S_2, R_1, R_2. Fig. 8(a) represents the true dependence relationships between the variables. The arrow between S_1 and S_2 could be due to poor experiment design, whereby the settings on qubit 2 are not selected independently of those on qubit 1.

Excepting statistical issues (*i.e.*, given the underlying probability distribution over these random variables), the dependency relationships given by the graph in Fig. 8(a) would be reconstructed by examining conditional dependence relations. In contrast, if we only assess marginal independence between R_1 and S_2, the fact that both random variables have a common cause (S_1) would create a fictitious dependence between these variables (see Fig. 8(b)). This is of course well known in statistics as the confounding of statistical association by an unobserved common cause [24].

Now, a marginal independence test would be sufficient if the experiment design was suitably randomized such that S_1 and S_2 are independent. However, for quantum computing platforms with many qubits a suitably randomized experiment design may be difficult to guarantee, and evaluation of association between random variables in distinct regions based on conditional independence testing is more reliable.
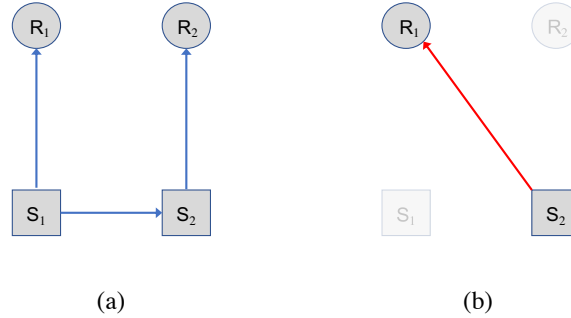

(a) (b)

Figure 8: An example illustrating the difference between conditional and marginal independence between random variables. (a) shows the true causal relationships between the settings (S_1, S_2) and results (R_1, R_2) in a two qubit experiment. (b) shows the causal relationship inferred if a marginal independence test is performed on just variables R_1 and S_2.

B Equivalence of two definitions of crosstalk

In the main text we presented two definitions of crosstalk-free QIPs, at two different layers of abstraction. The first, **Definition 1**, was stated in terms of properties of quantum operations (locality and independence) assuming a Markovian QIP, and the second, **Definition 2**, was stated operationally in terms of conditional independence between of classical random variables associated with experiments. In this Appendix we prove that these two definitions are in fact, equivalent; *i.e.*, A QIP is crosstalk-free according to **Definition 1** $\iff$ A QIP is crosstalk-free according to **Definition 2**.

B.1 Conditional independence in terms of quantum operations

Before we prove this equivalence we need to define the probabilities that arise in assessing conditional independence in terms of the quantum operations that form a model for the QIP.

We restrict ourselves to a QIP composed of two qubits for simplicity, and because it suffices to demonstrate the points we wish to make. In this case there are four random variables in the problem: $R_i \in \{r_i^k\}_{k=1}^K, i = 1, 2$, the measurement results on the two qubits, and $S_i \in \{r_i^l\}_{l=1}^L, i = 1, 2$, the settings on the two qubits. The settings will enumerate over some set of single-qubit gate sequences on each qubit.

Consider the conditional independence statement $P(\mathsf{R}_i|\mathsf{S}_i, \mathsf{S}_j, \mathsf{R}_j) = P(\mathsf{R}_i|\mathsf{S}_i)$ for $i, j \in \{1, 2\}, i \neq j$, which captures the crosstalk-free condition between regions 1 and 2. What is this condition equivalent to in terms of physical states, operations and measurements? To address this question, we note that for any experiment, the Born rule dictates

$$P(r_1^i, r_2^j|s_1^k, s_2^l) = (E(r_1^i, r_2^j)|\mathcal{M}(s_1^k, s_2^l)|\rho_0), \tag{15}$$

where $|E(r_1^i, r_2^j))$ is a POVM indexed by the measurement results, $\mathcal{M}(s_1^k, s_2^l)$ is a general two qubit CPTP map indexed by the settings, and $|\rho_0)$ is the initial state of the two qubits. Note that we are using Hilbert-Schmidt representations of all of these quantities for notational simplicity. In fact, this is the only relation we can write down without making further assumptions about crosstalk, factorizability of operations, *etc.* To obtain other probabilities we need to apply the usual conditioning and marginalization rules, *e.g.*,

$$P(r_1^i|s_1^k, s_2^l) = \sum_n (E(r_1^i, r_2^n)|\mathcal{M}(s_1^k, s_2^l)|\rho_0)$$

$$P(r_1^i|s_1^k) = \sum_n P(r_1^i|s_1^k, s_2^n)P(s_2^n|s_1^k)$$

$$= \sum_n \sum_m (E(r_1^i, r_2^m)|\mathcal{M}(s_1^k, s_2^n)|\rho_0) \ P(s_2^n|s_1^k). \tag{16}$$

Using such relations we can write the crosstalk-free condition as

$$P(\mathsf{R}_1|\mathsf{S}_1, \mathsf{S}_2, \mathsf{R}_2) = P(\mathsf{R}_1|\mathsf{S}_1)$$

$$\Rightarrow P(r_1^i|s_1^k, s_2^l, r_2^j) = P(r_1^i|s_1^k) \qquad \forall i, j, k, l$$

$$\Rightarrow \frac{P(r_1^i, r_2^j|s_1^k, s_2^l)}{P(r_2^j|s_1^k, s_2^l)} = \sum_{n,m} P(r_1^i, r_2^n|s_1^k, s_2^m)P(s_2^m|s_1^k) \qquad \forall i, j, k, l$$

$$\boxed{\Rightarrow P(r_1^i, r_2^j|s_1^k, s_2^l) = \left(\sum_z P(r_1^z, r_2^j|s_1^k, s_2^l)\right)\left(\sum_{n,m} P(r_1^i, r_2^n|s_1^k, s_2^m)P(s_2^m|s_1^k)\right) \qquad \forall i, j, k, l}$$

$$\tag{17}$$

$$\Rightarrow (E(r_1^i, r_2^j)|\mathcal{M}(s_1^k, s_2^l)|\rho_0) = \left(\sum_z (E(r_1^z, r_2^j)|\mathcal{M}(s_1^k, s_2^l)|\rho_0)\right)\left(\sum_{n,m} (E(r_1^i, r_2^n)|\mathcal{M}(s_1^k, s_2^m)|\rho_0)P(s_2^m|s_1^k)\right)$$

$$\forall i, j, k, l$$

We can write a similar explicit equation for the condition $P(\mathsf{R}_2|\mathsf{S}_1, \mathsf{S}_2, \mathsf{R}_1) = P(\mathsf{R}_2|\mathsf{S}_2)$.

B.2 Definition 1 $\Rightarrow$ Definition 2

A crosstalk-free Markovian QIP according to model-based definition, see Sec. 4, has state preparations, gate operations, and measurements that satisfy:

$$|\rho_0) = |\rho_0^1) \otimes |\rho_0^2)$$

$$\mathcal{M}(\mathsf{S}_1, \mathsf{S}_2) = \mathcal{M}_1(\mathsf{S}_1) \otimes \mathcal{M}_2(\mathsf{S}_2)$$

$$|E(\mathsf{R}_1, \mathsf{R}_2)) = |E(\mathsf{R}_1)) \otimes |E(\mathsf{R}_2)) \tag{18}$$

We proceed to show that given such a model for operations, the QIP is also crosstalk-free under **Definition 2**; *i.e.*, $P(\mathsf{R}_i|\mathsf{S}_i, \mathsf{S}_j, \mathsf{R}_j) = P(\mathsf{R}_i|\mathsf{S}_i)$, for $i, j \in \{1, 2\}, i \neq j$. To do so, we substitute the factorized forms in

Eq. (18) into the explicit form of the condition $P(\mathsf{R}_1|\mathsf{S}_1,\mathsf{S}_2,\mathsf{R}_2) = P(\mathsf{R}_1|\mathsf{S}_1)$ given in Eq. (17):

$$P(r_1^i, r_2^j | s_1^k, s_2^l) \overset{?}{=} \left(\sum_z P(r_1^z, r_2^j | s_1^k, s_2^l) \right) \left(\sum_{n,m} P(r_1^i, r_2^n | s_1^k, s_2^m) P(s_2^m | s_1^k) \right)$$

$$\forall i, j, k, l$$

$$\Rightarrow (E(r_1^i)|\mathcal{M}_1(s_1^k)|\rho_0^1)(E(r_2^j)|\mathcal{M}_2(s_2^l)|\rho_0^2) \overset{?}{=} \left(\sum_z (E(r_1^z)|\mathcal{M}_1(s_1^k)|\rho_0^1)(E(r_2^j)|\mathcal{M}_2(s_2^l)|\rho_0^2) \right) \cdot$$

$$\left(\sum_{n,m} (E(r_1^i)|\mathcal{M}_1(s_1^k)|\rho_0^1)(E(r_2^n)|\mathcal{M}_2(s_2^m)|\rho_0^2) P(s_2^m | s_1^k) \right)$$

$$\forall i, j, k, l$$

$$\Rightarrow (E(r_1^i)|\mathcal{M}_1(s_1^k)|\rho_0^1)(E(r_2^j)|\mathcal{M}_2(s_2^l)|\rho_0^2) \overset{?}{=} (E(r_2^j)|\mathcal{M}_2(s_2^l)|\rho_0^2)(E(r_1^i)|\mathcal{M}_1(s_1^k)|\rho_0^1), \tag{19}$$

which is obviously true. To arrive at the last line we have used the properties $\sum_z (E(r^z)|\mathcal{M}(s_k)|\rho_0) = 1$ for any k, ρ_0, and $\sum_m P(s_2^m | s_1^k) = 1$. We can verify that the equality $P(\mathsf{R}_2|\mathsf{S}_2,\mathsf{S}_1,\mathsf{R}_1) = P(\mathsf{R}_2|\mathsf{S}_2)$ also holds for quantum operations that satisfy Eq. (18). This concludes the proof that the model-based definition of a crosstalk-free QIP implies our model-free definition of a crosstalk-free QIP.

B.3 Definition 2 $\Rightarrow$ Definition 1

To prove this direction, we will actually prove its contrapositive, namely,

$$(\neg\text{locality}) \text{ or } (\neg\text{independence}) \Rightarrow \text{violation of } \textbf{Definition 2}. \tag{20}$$

We proceed by showing that quantum operations that do not satisfy locality or independence lead to violations of the explicit form of $P(\mathsf{R}_1|\mathsf{S}_1,\mathsf{S}_2,\mathsf{R}_2) = P(\mathsf{R}_1|\mathsf{S}_1)$ stated in Eq. (17). The proofs straightforwardly generalize to the violation of $P(\mathsf{R}_2|\mathsf{S}_2,\mathsf{S}_1,\mathsf{R}_1) = P(\mathsf{R}_2|\mathsf{S}_2)$.

B.3.1 Locality

A Markovian QIP that does not satisfy the locality principle has state preparations, gate operations, or measurements (or all of these) that do not factorize as in Eq. (18).

First consider the case where the initial state does not factorize, but all other operations do. Then, expanding Eq. (17), we get:

$$P(r_1^i, r_2^j | s_1^k, s_2^l) \overset{?}{=} \left(\sum_z P(r_1^z, r_2^j | s_1^k, s_2^l) \right) \left(\sum_{n,m} P(r_1^i, r_2^n | s_1^k, s_2^m) P(s_2^m | s_1^k) \right)$$

$$\forall i, j, k, l$$

$$\Rightarrow (E(r_1^i)| \otimes (E(r_2^j)|\mathcal{M}_1(s_1^k) \otimes \mathcal{M}_2(s_2^l)|\rho_0) \overset{?}{=} \left(\sum_z (E(r_1^z)| \otimes (E(r_2^j)|\mathcal{M}_1(s_1^k) \otimes \mathcal{M}_2(s_2^l)|\rho_0) \right) \cdot$$

$$\left(\sum_{n,m} (E(r_1^i)| \otimes (E(r_2^n)|\mathcal{M}_1(s_1^k) \otimes \mathcal{M}_2(s_2^m)|\rho_0) P(s_2^m | s_1^k) \right)$$

$$\forall i, j, k, l$$

$$\Rightarrow (E(r_1^i)| \otimes (E(r_2^j)|\mathcal{M}_1(s_1^k) \otimes \mathcal{M}_2(s_2^l)|\rho_0) \overset{?}{=} \left((E(r_2^j)||\bar{\rho}_2(k,l)) \right) \left(E(r_1^i)||\bar{\rho}_1(k)) \right) \qquad \forall i, j, k, l, \tag{21}$$

where we have defined $|\bar{\rho}_1(k)) \equiv \text{tr}_2 \left(\mathcal{M}_1(s_1^k) \otimes \sum_m P(s_2^m | s_1^k) \mathcal{M}_2(s_2^m) |\rho_0) \right)$ and $|\bar{\rho}_2(k,l)) \equiv \text{tr}_1 \left(\mathcal{M}_1(s_1^k) \otimes \mathcal{M}_2(s_2^l)|\rho_0) \right)$. This last equality cannot be true in general since it is expressing a joint distribution over r_1^i, r_2^j on the left hand side with a product of marginal distributions over these two variables

on the right hand side. To see this more clearly, note that the equality must hold for all i, j, k, l, so suppose $\mathcal{M}_1(s_1^k) = \mathcal{M}_2(s_2^l) = \mathbf{1}, \ \forall k, l$. In this case, the condition simplifies to:

$$(E(r_1^i)| \otimes (E(r_2^j)||\rho_0) \overset{?}{=} \left((E(r_2^j)||\bar{\rho}_2) \right) (E(r_1^i)||\bar{\rho}_1) \qquad \forall i, j, \tag{22}$$

where $|\bar{\rho}_{1(2)}) = \text{tr}_{2(1)} (|\rho_0))$. This equality cannot be true for any state ρ_0 that is not separable and therefore we conclude that conditional independence condition is violated in the case where the initial state does not factorize.

Next, consider the case where the gate operations do not factorize:

$$\exists a : \mathcal{M}(s_1^a, s_2^l) \neq \mathcal{M}_1(s_1^a)\mathcal{M}_2(s_2^l), \ \ \forall l, \tag{23}$$

i.e., that the operation induced when the setting on the first qubit is a is an entangling (non-factorizable) operation between the two qubits (regardless of what the setting is on qubit 1). We assume that the initial state and all measurement POVM elements factorize.

Then, returning to the condition in Eq. (17), and considering the case $k = a$,

$$P(r_1^i, r_2^j|s_1^a, s_2^l) \overset{?}{=} \left(\sum_z P(r_1^z, r_2^j|s_1^a, s_2^l) \right) \left(\sum_{n,m} P(r_1^i, r_2^n|s_1^k, s_2^m)P(s_2^m|s_1^k) \right)$$
$$\forall i, j, l$$

$$\Rightarrow (E(r_1^i)| \otimes (E(r_2^j)|\mathcal{M}(s_1^a, s_2^l)|\rho_0^1) \otimes |\rho_0^2) \overset{?}{=} \left(\sum_z (E(r_1^z)| \otimes (E(r_2^j)|\mathcal{M}(s_1^a, s_2^l)|\rho_0^1) \otimes |\rho_0^2) \right) \cdot$$
$$\left(\sum_{n,m} (E(r_1^i)| \otimes (E(r_2^n)|\mathcal{M}(s_1^a, s_2^m)|\rho_0^1) \otimes |\rho_0^2)P(s_2^m|s_1^k) \right)$$
$$\forall i, j, l \tag{24}$$

Now, suppose this equality holds. Then we can perform a weighted sum over L on both sides to get

$$(E(r_1^i)| \otimes (E(r_2^j)|\bar{\mathcal{M}}(s_1^a)|\rho_0^1) \otimes |\rho_0^2) = (E(r_1^i)||\bar{\rho}_1(a))(E(r_2^j)||\bar{\rho}_2(a)), \qquad \forall i, j,$$

where

$$\bar{\mathcal{M}}(s_1^a) \equiv \sum_l P(s_2^l|s_1^a)\mathcal{M}(s_1^a, s_2^l), \qquad \text{and}$$
$$|\bar{\rho}_{1(2)}(a)) \equiv \text{tr}_{2(1)} \left(\bar{\mathcal{M}}(s_1^a)|\rho_0^1) \otimes |\rho_0^2) \right)$$

Again, we have an expression with a joint distribution over r_1^i, r_2^j on the left hand side and a product of marginals over the same variables on the right hand side. This cannot be true if $\mathcal{M}(s_1^a, s_2^l)$ does not factorize for all L, and therefore we conclude the original assumption of Eq. (24) holding is false. Therefore, we have shown that the model-free crosstalk-free condition is violated when gate operations violate locality and do not factorize.

Finally, the proof that violation of locality in measurements – *i.e.*, $|E(\mathsf{R}_1, \mathsf{R}_2) \neq |E(\mathsf{R}_1)) \otimes |E(\mathsf{R}_2))$ – results in a violation of the model-free crosstalk-free conditions follows straightforwardly from the corresponding proof for non-factorizable initial states since state preparation and measurement are dual operations.

B.3.2 Independence

Now we proceed to show that even if locality is respected by a QIP, violations of independence in the model-based framework result in violations of the model-free definition of a crosstalk-free QIP. For simplicity we have assumed that the only settings correspond to gate operations, and we only have one choice for state preparation and measurement basis. Therefore violation of independence within a model that respects locality can only manifest in one way:

$$\exists a, b \quad : \quad \mathcal{M}(s_1^a, s_2^b) = \mathcal{E}_1(s_2^b)\mathcal{M}_1(s_1^a) \otimes \mathcal{M}_2(s_2^b). \tag{25}$$

In other words, for some combination of settings, the operation done on the first qubit depends on the setting of the second qubit. Here, $\mathcal{E}_1()$ is some CPTP map on qubit 1.

Let us determine if this violation of independence results in a violation of the model-free condition in Eq. (17), by assuming all other operations respect locality and independence, and considering the case $k = a, l = b$:

$$P(r_1^i, r_2^j|s_1^a, s_2^b) \overset{?}{=} \left(\sum_z P(r_1^z, r_2^j|s_1^a, s_2^b)\right)\left(\sum_{n,m} P(r_1^i, r_2^n|s_1^a, s_2^m)P(s_2^m|s_1^a)\right)$$
$$\forall i, j$$

$$\Rightarrow (E(r_1^i)|\mathcal{E}_1(s_2^b)\mathcal{M}_1(s_1^a)|\rho_0^1)(E(r_2^j)|\mathcal{M}_2(s_2^b)|\rho_0^2) \overset{?}{=} \left(\sum_z (E(r_1^z)|\mathcal{E}_1(s_2^b)\mathcal{M}_1(s_1^a)|\rho_0^1)(E(r_2^j)|\mathcal{M}_2(s_2^b)|\rho_0^2)\right) \cdot$$
$$\left(\sum_n \left[(E(r_1^i)|\mathcal{E}_1(s_2^b)\mathcal{M}_1(s_1^a)|\rho_0^1)(E(r_2^n)|\mathcal{M}_2(s_2^b)|\rho_0^2)P(s_2^b|s_1^a)\right.\right.$$
$$\left.\left. + \sum_{m \neq b} (E(r_1^i)|\mathcal{M}_1(s_1^a)|\rho_0^1)(E(r_2^n)|\mathcal{M}_2(s_2^m)|\rho_0^2)P(s_2^m|s_1^a)\right]\right)$$
$$\forall i, j$$

$$\Rightarrow (E(r_1^i)|\mathcal{E}_1(s_2^b)\mathcal{M}_1(s_1^a)|\rho_0^1)(E(r_2^j)|\mathcal{M}_2(s_2^b)|\rho_0^2) \overset{?}{=} (E(r_2^j)|\mathcal{M}_2(s_2^b)|\rho_0^2) \cdot$$
$$\left[(E(r_1^i)|\mathcal{E}_1(s_2^b)\mathcal{M}_1(s_1^a)|\rho_0^1)P(s_2^b|s_1^a) + (E(r_1^i)|\mathcal{M}_1(s_1^a)|\rho_0^1)\left(1 - P(s_2^b|s_1^a)\right)\right] \qquad \forall i, j,$$

where we have used the completeness property of the POVM elements to perform the sums over z and n in the last line. The equality on the last line holds if $(E(r_2^j)|\mathcal{M}_2(s_2^b)|\rho_0^2) = 0$, so consider the cases where this quantity is non-zero (it cannot be zero for all j), and divide through both sides by this non-zero value. So for j such that $(E(r_2^j)|\mathcal{M}_2(s_2^b)|\rho_0^2) \neq 0$, the condition we are evaluating becomes:

$$(E(r_1^i)|\mathcal{E}_1(s_2^b)\mathcal{M}_1(s_1^a)|\rho_0^1)\left(1 - P(s_2^b|s_1^a)\right) \overset{?}{=} (E(r_1^i)|\mathcal{M}_1(s_1^a)|\rho_0^1)\left(1 - P(s_2^b|s_1^a)\right) \qquad \forall i$$
$$\Rightarrow (E(r_1^i)|\mathcal{E}_1(s_2^b)\mathcal{M}_1(s_1^a)|\rho_0^1) \overset{?}{=} (E(r_1^i)|\mathcal{M}_1(s_1^a)|\rho_0^1) \qquad \forall i, \tag{26}$$

where we have assumed $1 - P(s_2^b|s_1^a) \neq 0$. Under what conditions is this last equality true when $\mathcal{E}_1(s_2^b) \neq \mathbf{1}$? The only other way for this equality to hold is if $(E(r_1^i)|\mathcal{E}_1(s_2^b) = (E(r_1^i)|$, $\forall i$; *i.e.*, all the error maps act trivially on the measurement effects. Note that we could have written the violation of independence as a premultiplication error map (*i.e.*, $\mathcal{M}(s_1^a, s_2^b) = \mathcal{M}_1(s_1^a)\mathcal{E}_1(s_2^b) \otimes \mathcal{M}_2(s_2^b)$), in which case the equality would hold if the initial state is invariant under the error map. However, note that these CPTP maps represent the action of gate sequences, and the error map $\mathcal{E}_1(s_2^b)$ is the *effective* error on qubit 1 when some sequence $\mathcal{M}_2(s_2^b)$ is performed on qubit 2, after the desired sequence on qubit 1, $\mathcal{M}_1(s_1^a)$ has been factored out. These sequences are composed of elementary gates, some subset of which violate the independence condition, which leads the whole map to violate the independence condition. However, for a sufficiently rich set of sequences if one sequence violates independence, then it is likely that others do as well (in other words, there are a set of (a, b) satisfying Eq. (25). And the probability that the measurement effects are invariant under all the effective error maps $\mathcal{E}_1(s_2^b)$ is extremely unlikely. Therefore we conclude that for a sufficiently rich set of settings, violation of independence results in violation of model-free definition of a crosstalk-free QIP.

B.4   Definition 1 $\Longleftrightarrow$ Definition 2

The above subsections prove the two directions of implication required to establish equivalence between the model-based definition (**Definition 1**) and the model-free definition (**Definition 2**) of crosstalk-free QIPs.

C Pseudocode for lightweight experiment design

Algorithm 1 Lightweight crosstalk detection experiment generation. The output is a set of roughly $M \times N_{\text{circs}} \times N_{\text{con}}$ experiments on an M region QIP, with each experiment consisting of length L circuits on each region.

1: **procedure** CROSSTALKEXPERIMENTS($M, l, N_{\text{circs}}, N_{\text{con}}, p_{\text{idle}}$)
2:     **for** $0 < m < M - 1$ **do**   ▷ Sample N_{circs} circuits for each region
3:         $\texttt{bag}_m \leftarrow$ Sample of N_{circs} circuits length L, composed of elementary gates on region m
4:     $\texttt{Expts} \leftarrow \{\}$ ▷ Initialize with empty list of experiments
5:     **for** $0 < m < M - 1$ **do**
6:         **for** $0 < n < N_{\text{circs}} - 1$ **do**   ▷ For each region, iterate over the N_{circs} circuits sampled for that region
7:             $s_m \leftarrow$ Circuit number n from $\texttt{bag}_m$
8:             **for** $0 < c < N_{\text{con}}$ **do**   ▷ Generate N_{con} experiment with s_m circuit on region m
9:                 **for** $0 < k < M - 1$ **do**
10:                     **if** $k \neq m$ **then**
11:                         **if** Unif(0,1) $< p_{\text{idle}}$ **then**   ▷ With probability p_{idle} region k gets idle circuit
12:                           $s_k \leftarrow$ the length L idle circuit on region k
13:                       **else**
14:                           $s_k \leftarrow$ Sample a circuit from $\texttt{bag}_k$
15:                 Append to $\texttt{Expts}$ the experiment defined by parallel application of s_n (for $0 < n < M - 1$) to the M regions
16:     $\texttt{Expts} \leftarrow$ **RemoveDuplicates**($\texttt{Expts}$)   ▷ Remove duplicate experiments
17:     **return** $\texttt{Expts}$

D Summary of the PC algorihtm

The PC algorithm is described in detail in references [54, 13], but we describe its main steps here, and comment on its implementation in the crosstalk detection context.

Exhaustively checking for conditional dependence relations between N data variables is exponential cost in terms of computational difficulty and required dataset size. To get around this, the PC algorithm uses insights from graph theory to perform a hierarchy of tests that can result in reduced costs, particularly in sparsely connected graphs (*i.e.*, datasets with sparse conditional dependence relations). The fundamental property exploited by the PC algorithm to reduce the number of conditional independency tests is this: two nodes (X, Y) are conditionally independent given some subset of remaining nodes **S**, if and only if they are conditionally independent given pa(X) or pa(Y), where pa(X) are the parent nodes of X.

Algorithm 2 presents pseudocode for the PC algorithm, adapted from Ref. [54]. Each variable is represented by a node in a graph G, and **Adj**(G, X) is the set of nodes adjacent to node X. The algorithm initializes by constructing the complete undirected graph with N nodes. Then it prunes edges on this graph in a hierarchical manner: for every edge in the graph connecting nodes (X, Y), it examines subsets of neighbors of one of the nodes of increasing size, n, (starting from the null set, $n = 0$) and tests whether the two nodes are conditionally independent given the nodes in the subset. If so, it removes the edge. This procedure is repeated for every pair of connected nodes for increasing n, until no nodes have adjacency sets of size equal to or greater than n. At the end of this procedure, we have a pruned undirected graph with edges between nodes that are not conditionally independent under any conditioning set of adjacent nodes; this is often referred to as the *skeleton* graph. Under the PC algorithm this undirected graph is passed to a subroutine OrientEdges that orients each edge in the graph using several orienting rules, resulting in a directed acyclic graph. We do not execute this portion of the algorithm for crosstalk detection and therefore do not provide details on that step. Interested readers are referred to [54, 53].

Algorithm 2 Pseudocode for the PC algorithm executing on a graph with N nodes.

1: **procedure** $\mathrm{PC}(N)$
2:      $G \leftarrow$ the complete undirected graph on a vertex set of size N
3:      $n \leftarrow 0$
4:      **repeat**
5:        **repeat**
6:          $(X, Y) \leftarrow$ an ordered pair of variables that are adjacent in G such that $\mathbf{Adj}(G, X) \backslash \{Y\}$ has cardinality $\geq n$
7:            $\mathbf{S} \leftarrow$ a subset of $\mathbf{Adj}(G, X) \backslash \{Y\}$ of cardinality n
8:            **if** (X, Y) are conditional independent given $\mathbf{S}$ **then**
9:              delete edge $X - Y$ from G
10:             add $\mathbf{S}$ to $\mathbf{SepSet}(X)$ and $\mathbf{SepSet}(Y)$
11:        **until** all ordered pairs of adjacent variables (X, Y) such that $|\mathbf{Adj}(G, X) \backslash \{Y\}| \geq n$ and all $\mathbf{S} \subset \mathbf{Adj}(G, X) \backslash \{Y\}$ such that $|\mathbf{S}| = n$ have been tested for conditional independence.
12:        $n \leftarrow n + 1$
13:      **until** for each order pair of adjacent vertices (X, Y), $|\mathbf{Adj}(G, X) \backslash \{Y\}| < n$
14:      $G_{\mathrm{orient}} \leftarrow \mathrm{OrientEdges}(G, \mathbf{SepSet}(G))$
15:      **return** G_{orient}

The PC algorithm is stated in algorithm 2 in terms of abstract conditional independence tests. These are implemented statistically in most implementations. Also, one might be concerned that the order in which pairs of variables are considered will effect the resulting undirected graph, and indeed this is true. However, one can formulate an order-independent version of the PC algorithm [13] that removes this issue, and this is the version we use for crosstalk detection.