Embedding cyclic causal structures in acyclic spacetimes: no-go results for process matrices

Working Paper

Author(s):
Vilasini, Venkatesh; Renner, Renato

Publication date:
2022-03-21

Permanent link:
https://doi.org/10.3929/ethz-b-000593322

Originally published in:
arXiv, https://doi.org/10.48550/ARXIV.2203.11245

Embedding cyclic causal structures in acyclic spacetimes: no-go results for process matrices

V. Vilasini and Renato Renner

Institute for Theoretical Physics, ETH Zurich, 8093 Zürich, Switzerland

Causality can be defined in terms of a space-time structure or based on information-theoretic structures, which correspond to very different notions of causation. When analysing physical experiments, these notions must be put together in a compatible manner. The process matrix framework describes quantum indefinite causal structures in the information-theoretic sense, but the physicality of such processes remains an open question. At the same time, there are several experiments in Minkowski spacetime (which implies a definite spacetime notion of causality) that claim to have implemented indefinite information-theoretic causal structures, suggesting an apparent tension between these notions. To address this, we develop a general framework that disentangles these two notions and characterises their compatibility in scenarios where quantum systems may be delocalised over a spacetime. The framework first describes a composition of quantum maps through feedback loops, and then the embedding of the resulting (possibly cyclic) signalling structure in an acyclic spacetime. Relativistic causality then corresponds to the compatibility of the two notions of causation. We reformulate the process matrix framework here, establishing a number of connecting results as well as no-go results for physical implementations of process matrices in a spacetime. These reveal that it is impossible to physically implement indefinite causal order processes with spacetime localised systems, and also characterise the degree to which they must be delocalised. Further, we show that any physical implementation of an indefinite order process can ultimately be fine-grained to one that admits a fixed acyclic information-theoretic causal order that is compatible with the spacetime causal order, thus resolving the apparent paradox. Our work sheds light on the operational meaning of indefinite causal structures which we discuss in detail.

V. Vilasini: vilasini@phys.ethz.ch

Renato Renner: renner@phys.ethz.ch

Contents

1 Introduction

The notion of *causality* in physics has a striking resemblance to that of *entropy*: Everyone uses the term, but no one knows what it really is.[1] And like entropy, causality doesn't have a single definition. Rather, depending on the branch of physics, the definitions that are commonly used highlight different aspects. In the context of relativity theory, causality is a property of the geometry of spacetime and thus ultimately determined by Einstein's field equations. To distinguish it from other notions, we will call it *spacetime causality*. In the context of quantum theory, for instance, the notion of causality is more commonly used to describe how interventions, such the choices of parameters by an experimenter, influence observations. Since this is related to how information flows between these values, we refer to it as *information-theoretic causality*. Note that this information-theoretic notion does not refer to the structure of spacetime. Conversely, spacetime causality can be defined without referring to information theory. Nonetheless, like in the case of entropy, whose use in thermodynamics is deeply connected to the — a priori rather different — information-theoretic definition, the two notions of causality are related to each other. This distinction and connection has also been highlighted and studied in a recent work involving one of us [2].

Both spacetime causality and information-theoretic causality define a relation between certain *events*. That is, they assert whether or not an event B lies in the *causal future* of an event A. In the case of spacetime causality, the events are points (or regions) in spacetime, and the relation is specified by the metric tensor together with a time direction.[2] Conversely, for information-theoretic causality, the events are associated to random variables (or, more generally, quantum systems) and the order relation is specified by information-thoretic channels connecting them, which can generate correlations between them.

In any real physical experiment, these two notions of causality play together. For a concrete example, consider an experiment where a source emits particles whose properties can be measured by a detector, the property to be measured being specified by the position of a knob A_s that is turned at a particular time. The detector measures the property of the incoming particles according to the knob setting A_s and displays the result as a position of a pointer A_o at a certain time. A_s and A_o can be regarded as random variables and thus give raise to an information-theoretic notion of causality. But A_s and A_o are also associated to locations in spacetime, and they thus inherit the causality relation that is specified by the spacetime metric. Naturally, we would require that these two notions of causal order are compatible with each other, in the sense that any information-theoretic causality relation implies a corresponding spacetime causality relation. Concretely, A_o can only lie in the information-theoretic future of A_s if A_o lies in the spacetime future of A_s.

These considerations hint at a more general principle, which one may regard as an instance of Landauer's famous slogan that "information is physical" [3]. Landauer was referring to thermodynamics and what he meant is that any realistic information-processing system is also a physical system and hence has to obey the laws of thermodynamics. This principle plays a key role in the modern understanding of the notion of entropy that we alluded to above. It emphasises that information-theoretic and thermodynamic entropy are unavoidably related. Repurposing Landauer's slogan, we may say that any processing of information takes place in a spacetime.[3] In particular, there must be an association between information-theoretic events and spacetime events. Consequently, the information-theoretic causal structure is *embedded* in the spacetime causality structure. This leads us back to the compatibility requirement mentioned above.

The interplay between spacetime and information-theoretic causality can however be quite subtle, especially when one considers quantum experiments. The probably most famous example that illustrates this point is Einstein-Podolski-Rosen (EPR) experiment [4], where two agents, Alice and Bob, carry out

[1] According to a widely circulated story, Claude Shannon, who was looking for a name for his measure of information, received the following advise from his close friend John von Neumann: "Why don't you call it entropy? [...] no one understands entropy very well, so in any discussion you will be in a position of advantage" [1].

[2] Considering the fact that the metric tensor can be recovered from the causal structure up to a scaling factor, one can basically identify the causal relation with the spacetime geometry.

[3] By this we do not mean that there must exist a background spacetime on which physics is embedded, the notion of spacetime may very well be given solely by physical reference systems such as rods and clocks possessed by agents participating in the protocol. The point is that we must instantiate our information-processing protocols with some notion of spacetime and characterise the compatibility of the two causal notions, in order to analyse physical experiments.

measurements on two separate quantum systems, which are mutually entangled. The analysis of the experiment relies on both notions of causality. Let A_s and B_s be the choices of measurements made by Alice and Bob, and let A_o and B_o be the corresponding outcomes they observe. For the spacetime causality, this means that A_o must lie in the causal future of A_s and, similarly, B_o must lie in the causal future of B_s. In addition, EPR demand that the two measurements are spacelike separated, i.e., that B_o does *not* lie in the spacetime future of A_s, and that A_o does *not* lie in the spacetime future of B_s. The compatibility requirement described above now implies that B_o cannot lie in the information-theoretic future of A_s, nor can A_o lie in the information-theoretic future of B_s. The compatibility requirement thus immediately leads to the important question whether quantum theory satisfies such an information-theoretic criterion. And it was precisely this question that Einstein answered very colorfully when he wrote to Max Born that there should be no "spooky action" at a distance [5].

Recall that information-theoretic causality relates to the flow of information. That is B_o lies in the causal future of A_s if information could be transmitted from A_s to B_o. Whether this is the case for the EPR experiment ultimately depends on the theory that one uses to describe it. According to standard quantum theory, the information-theoretic and spacetime causal futures coincide, and B_o does not causally depend on the choice of A_s. Consequently, we may say that within the EPR experiment quantum theory satisfies the compatibility criterion between spacetime and information-theoretic causality. Conversely, as shown by Bell [6], any classical, deterministic theory that correctly[4] describes the EPR experiment with A_s and B_s chosen independently of the parameters specifying the states of particles, cannot possibly satisfy this information-theoretic criterion, and consequently B_o must causally depend on the choice of A_s. An example of such a theory is *Bohmian mechanics* [7, 8]. This theory thus violates the compatibility criterion between spacetime and information-theoretic causality.[5]

We have thus seen that some of the key claims surrounding the EPR experiment and Bell's theorem can be understood as instances of the compatibility requirement between spacetime and information-theoretic causality. This also shows, as highlighted in previous works [2, 9], that disentangling these notions is key to analysing the different causal explanations of quantum correlations, in light of Bell's theorem, which tells us that in order to preserve the free choice of settings and the compatibility with spacetime causality and still explain the results of quantum experiments, the information theoretic notion of causality cannot be entirely classical. In the case of standard quantum theory, free choice and compatibility are preserved by generalising the information-theoretic notion of causation from a classical to a quantum information theoretic formulation. The relatively recent but seminal advancements in developing genuinely quantum frameworks for causal modelling, provide a compatible description of causation in quantum experiments (even beyond the EPR and Bell experiments) where information-theoretic events are localised in spacetime, and quantum operations occur in a fixed acyclic causal order [10–23]. Within this modern formulation, Bell's theorem can be simply seen as showing that within a given information-theoretic causal structure (that accounts for free choice and compatibility), classical causal models [24] cannot possibly explain quantum correlations. This non-classicality of the causal model can be certified by observing that the correlations between the settings and outcomes in the experiment violates a Bell inequality.

More generally, in quantum relativistic experiments, quantum systems may be delocalised over a spacetime or travel in a superposition of different trajectories through spacetime. Further, in quantum gravitational settings, the assumption of a fixed background spacetime structure might be too restrictive, and one can envisage thought-experiments where the spacetime causal structure is in a quantum superposition, leading to the concept of *indefinite causal structures* [25–27]. Motivated by this, the process matrix framework [28] was proposed for defining indefinite causal structures in a purely information-theoretic manner, by considering information processing protocols where parties act within local quantum laboratories in the absence of a global definite acyclic causal order connecting their local operations. For instance, like in the case of the EPR experiment, consider two parties, Alice and Bob, who receive quantum systems A_I and B_I, apply a local operation (such as a measurement) on these systems depending on a choice of classical setting A_s and B_s, and obtain the corresponding classical outcome A_o and B_o, along with a final quantum system A_O and B_O after the operation is applied. Unlike the EPR case, they are allowed to communicate and can send the final quantum system A_O and B_O out of the labs.

[4] Correctness means here that the description is in agreement with the observations of actual quantum experiments.

[5] This problem can be rectified by changing the notion of spacetime causality: If one replaces special relativistic spacetime by the Newtonian notion of space and time, compatibility is reestablished.

The outside environment of these labs (modelled by a process matrix) may contain quantum channels that route the outputs of one party to the input of another. This means that there is a well defined information-theoretic order of local events, A_o and A_O lie in the causal future of A_s and A_I, while B_o and B_O lie in the causal future of B_s and B_I. For such a process to be compatible with a *fixed order* $A \prec B$ or $B \prec A$ over Alice and Bob, would mean that it is possible to simulate the resulting correlations by one-way communication, either from Alice to Bob ($A \prec B$) or from Bob to Alice ($B \prec A$). This would exclude causal relations where A_o lies in the future of B_s *and* B_o lies in the future of A_s. A typical example of a process *without* fixed order between Alice and Bob's operations is the *quantum switch* [29]. Here a quantum system is either first routed to Alice and then to Bob or vice versa, and the decision of the route is also controlled by a quantum system.

In the process matrix framework, analogous to Bell inequalities, so-called causal inequalities have been proposed to certify the non-classicality of the causal order itself and there exist theoretical indefinite causal order processes that violate these causal inequalities [28]. The causal separability of a process matrix (analogous to separability of a quantum state) [30, 31] is considered another indication of the indefiniteness of its causal structure and can be certified through so-called causal witnesses. The quantum switch (QS) motivated above is an example of a causally non-separable processes that does not violate causal inequalities [29, 30]. Further theoretical examples of causally non-separable processes that do violate causal inequalities have also been obtained in the process matrix framework [28, 32]. Causally non-separable processes have also been shown to provide significant advantages over processes admitting a definite causal order in several information processing tasks in the fields of quantum communication complexity [33, 34], quantum channels [35], quantum metrology [36], quantum computing [37] as well as quantum thermodynamics [38, 39].

Again in the spirit of Landauer's slogan, a natural question in the study of processes with indefinite causal order is whether they can be implemented physically. In the language developed here, the question is whether there exists an embedding of the given information-theoretic causal structure into a spacetime causal structure such that the compatibility requirement described above is satisfied. Indeed, while indefinite causal order processes are intriguing in theory, there are several longstanding open questions and debates regarding the their physicality. Here we aim to address such questions with our approach. We describe these open questions, challenges and previous works in more detail below.

A first open question is regarding the set of process matrices that can be implemented in a lab using quantum systems in Minkowski spacetime, whether causal inequality violating processes can be physically implemented and what this would imply for our understanding of causality in quantum theory. A general class of quantum circuits modelling quantum controlled suerpositions of orders such as the quantum switch have been proposed and shown to not violate causal inequalities [67, 68], however, it remains unclear whether these are the largest set of physically realisable process matrices as these are constructed through a bottom-up approach of starting with simple fixed order processes and building up generality. A top down approach that starts with general processes and imposes conditions for compatibly embedding the process in a spacetime is needed for fully addressing this question.

On the other hand, there are numerous experiments and experimental proposals that claim to physically implement indefinite causal structures such as the quantum switch in Minkowski spacetime [37, 40–49]. This leads to an apparent paradox– while Minkowski spacetime implies a definite acyclic spacetime notion of causality, indefinite causal order processes imply an indefinite information-theoretic notion of causality, but we expect these notions to be compatible in any physical experiment. How can we consistently describe both notions of causality as well as the results of such experiments, as quantum causal models do for Bell experiments? What does this consistent description imply for the physical meaning of indefinite causal order processes? What does this tell us about the resource responsible for the information-theoretic advantages offered by causally non-separable processes?

There has been a longstanding and continuing debate on these matters (see for instance [50]) but a clear answer to the above questions would require a general framework for quantum causality that has the following features.

1. The framework must clearly disentangle the information-theoretic and spacetime notions of causality, and characterise both under minimal but operational assumptions. This means that the information-theoretic causal structure can in general be cyclic.

2. While the disentangling offers a vast generality, it does not tell us about physical experiments, for which we must relate the two notions by embedding one in the other and imposing a compatibility

condition.

3. The embedding must be general enough to allow for a description of experiments with spacetime delocalised quantum systems.

4. The framework should offer the ability of analysing the causal structure at different levels of detail, as the cyclicity or acyclicity of a causal structure (under both notions of causation) can depend on the information captured by the nodes of the causal structure.

Here we develop a framework that meets all these criteria, and apply it to characterise physical implementations of process matrices. We briefly describe previous works in this direction.

In recent work [51], it has been shown that a large class of process matrices (unitary processes) can be described using cyclic quantum causal models. This provides significant insights into the information-theoretic causal structure of such processes, this along with previous works relating indefinite and cyclic causation [29, 52] provide useful insights for achieving the first criterion above. However, these works focus on the details of the information-theoretic notion and do not consider spacetime causality or compatibility between these notions as required by the second criterion. For compatibility, we would not need a full framework for cyclic information-theoretic causal models and will formulate it under rather minimal operational assumptions.

In a recent work involving one of us [2], a framework that meets both the first and second criteria has been developed. This allows for modelling a general class of cyclic and non-classical (quantum or post-quantum) causal models and conditions for characterising whether or not a causal model leads to signalling outside the future with respect to an embedding in a spacetime structure have been proposed. However, this framework mainly focusses on a more device independent notion of signalling between classical settings and outcomes (even if these may be generated by measuring quantum or post-quantum systems), and the criteria 3 and 4 were not considered there, which are necessary for analysing process matrix implementations.

Cyclic causal models are often used in classical data sciences for modelling physical scenarios with feedback [53], for instance the demand for a commodity may causally influence the price of the commodity which may in turn causally influence the demand. The cyclic causal structures modelling physical feedback and those modelling exotic closed-time like curves comes from analysing the spacetime embedding. In the former case, we know that we ultimately have an acyclic causal structure where the demand D^{t_1} at time t_1 influences the price P^{t_2} at time $t_2 > t_1$ which in turn influences the demand D^{t_3} at time $t_3 > t_2$ and so on. If we coarse grain over the time information in this acyclic causal structure, we recover the original cyclic causal structure. This coarse-graining corresponds to combining multiple nodes D^{t_1}, D^{t_3} of the acyclic causal structure into a single node D, and this example illustrates point number 4 above, that the cyclicity or acyclicity of the causal structure depends on the level of detailed information encoded in its nodes. Notice that this scenario can also be described in terms of a more general spacetime embedding of the variables D and P of the information-theoretic structure, one where each information-theoretic event is assigned a set of spacetime events and not a single spacetime event.

Here we adopt a similar approach to [2], in that we characterise the operational causal structure through minimal assumptions, by analysing the effect of interventions on physical systems to infer the possibilities for signalling offered by the underlying causal structure and impose conditions on the spacetime embedding based these signalling possibilities. Here we go a step further and allow for signalling at the level of quantum systems in addition to classical settings and outcomes, and also allow for more general spacetime embeddings where systems may be delocalised over the spacetime. Further, apart from the spacetime embedding, we also define the general notion of fine-graining which allow for an analysis of causal structures, quantum channels and systems and their properties at different levels of detail. As illustrated by the above example with demand and price (and further intuitive Examples 1, 2 and 3 of the main text), both the embedding and fine-graining are crucial to gaining physical insights about the causal structure.

Applying this framework, we provide an answer to the question regarding the physicality of indefinite causal order processes in terms of no-go theorems. They assert that certain assumptions, which one would naturally want to make about such an embedding, are mutually contradictory. Concretely, a first assumption concerns the information-theory side and demands that the process has no fixed order. This assumption restricts our attention to genuinely indefinite-order processes. Another assumption concerns the spacetime side and demands that there are no closed timelike curves, i.e., the spacetime causality

relation has no cycles. This restricts our attention to considering compatibility with spacetimes of physical interest (Minkowski spacetime being a particular example). Finally, a third assumption concerns the embedding of the information-theoretic causal structure into the spacetime causal structure. It limits the degree to which information-theoretic events can be delocalised in the spacetime. For instance, a particular embedding that satisfies this would be one where information-theoretic events also correspond to well localised spacetime events. Our main no-go theorem now asserts that these three assumptions are contradictory. As corollaries, it implies that it is impossible to physically implement indefinite causal order processes in an acyclic spacetime, solely using spacetime localised systems, or using systems that are time-localised in a global reference frame if we wish to preserve relativistic causality in the spacetime.

This no-go theorem sheds light on recent experiments which claim to have implemented processes with indefinite causal order, such as the quantum switch. Here, the first assumption holds by construction, the second is also satisfied for any presently feasible lab experiment (it is safe to assume that the spacetime within the lab is well-behaved). The no-go theorem thus implies that, in any physical implementation, the information-theoretic events are spread out in spacetime.

A second result that we are going to present deals with exactly this situation. It asserts that any process satisfying the first two assumptions above can be fine-grained into a process over a larger number of parties such that the new process has a fixed order. Further, in the new process, the information-theoretic events are also well localised spacetime events, even though this was not the case in the original coarse-grained process. This means that any experiment that implements an indefinite-order process can, according to a sufficiently fine-grained description, again be regarded as a fixed-order process. We note that such a result was first suggested in [50] where they showed this for the special case of the quantum switch processes, and for certain types of spacetime implementations of this processes. Our result generalises this to arbitrary processes and arbitrary spacetime implementations. This is analogous to our simple example with demand and price where we fine-grained the cyclic causal structure into an acyclic one by adding time information and demanding compatibility of the two notions. Our results show that physical implementations of indefinite causal order processes follow the same intuition, and we can resolve the apparent tension between the definite spacetime causal structure and indefinite information-theoretic causal structure once we look at a sufficiently fine-grained description. Our results also reveal a tight connection between non fixed order processes and cyclic signalling structures. Consequently, our results imply that none of these experiments truly implement an indefinite causal structure, the fine-grained description is acyclic while the coarse-grained description is cyclic.

We now summarise the main results of the paper that were alluded to in the above, along with the cross references to the relevant sections and theorems, before diving into the technical part of the paper. This long introduction and summary of contributions can be treated as a coarse-grained version of the more fine-grained results of the main paper.

Summary of contributions We summarise the main contributions and results of this paper below.

- We first develop, in Section 2, a general and purely information-theoretic framework for describing cyclic quantum networks formed by composition of quantum CPTP maps. We characterise causation in such networks by focussing on the operationally verifiable property of signalling. We define two new concepts, an *embedding* and a *fine-graining*, which allows us to embed quantum networks and their signalling relations in an abstract causal structure (modelled as a directed graph) and to analyse the compatibility between these at different levels of detail.

- In Section 3 we apply these concepts to particular case of a relativistic causal structure corresponding to a spacetime, and define what it means for a cyclic quantum network to be implemented in a spacetime, and for it to satisfy relativistic causality therein. This also includes implementations where quantum systems are classically or quantumly delocalised over multiple spacetime locations.

- In Theorem 3.6 we show that any signalling structure (possibly cyclic) can be embedded in an acyclic spacetime without violating relativistic causality, if we allow for the quantum systems of the network to be associated with sufficiently large spacetime regions. We then show in Lemma 3.9 that such set of spacetime embedded signalling relations can be ultimately fine-grained to a set of acyclic signalling relations, whose edges flow from past to future in the spacetime.

- We review the process matrix framework in Section 4 and in Section 5, we reformulate the process matrix framework in terms of the more general framework developed here and derive a number

of connecting results (Lemmas 5.2 and 5.3) that enables the expression of processes in terms of composition through feedback loops.

- A key ingredient in our reformulation is the notion of an *extended local map*, which is a single quantum map that captures all possible interventions an agent can potentially perform, the different choices being encoded in the values of a classical setting. We use this to establish an equivalence between device dependent and device-independent notions of signalling in Theorem 5.4, where the former is at the level of quantum states and the latter at the level of observed probabilities over classical settings and outcomes of the parties.

- In Section 6.1 we derive a number of no-go results for physical implementations of process matrices in a fixed spacetime. We show in Theorem 6.2 that any implementation of a non-fixed order process that does not violate relativistic causality in a spacetime, will necessarily violate a certain condition on the spacetime embedding that limits the degree to which the systems of the information-theoretic causal structure are spread out in the spacetime. Corollary 6.3 establishes that the signalling structure of a process is cyclic if and only if it is not a fixed order process. Corollaries 6.4 and 6.6 of this theorem imply that it is impossible to physically implement non-fixed order processes in a fixed spacetime using spacetime localised systems or using systems that are time-localised in a global reference frame.

- In Section 6.2, we show that physical implementations of all process in a fixed spacetime can be ultimately fine-grained to a fixed order process over a larger number of parties. This result captures the fact that even in scenarios where quantum systems take a superposition of different trajectories through spacetime, an agent has the potential to intervene at any of the spacetime locations to verify the probability of detecting the particle there (even if they may choose to not do so in a particular experiment where they wish to main the coherence). Further, if agents choose to perform such interventions, this will not enable them to signal outside the future of the spacetime.

- Our results have implications for several table-top experiments in Minkowski spacetime that claim to implement an indefinite causal order process, the quantum switch. In Section 7 we analyse this process in detail and discuss these implications. We also analyse theoretical proposals for quantum gravitational implementations of the quantum switch, and outline properties of this gravitational implementation that might differ from physics in a fixed spacetime. In this regard, we find that the property of events being time-localised for each agent is not special to the quantum gravitational implementation, we show this by constructing an explicit quantum switch protocol in Minkowski spacetime with this property in Section 7.4.

- Our framework and results shed light on several open questions and debates surrounding the meaning of indefinite causal structures, the notion of events in these settings and their relation to cyclic causal structures. We discuss these points in detail in Section 8, while analysing implicit assumptions of the process framework more explicitly in the language of our framework. We conclude with the main take home messages in Section 9.

2 Composition of quantum maps and signalling structure

As motivated in the introduction, it is desirable to describe causation, spacetime and their relationships in an operational manner, and using minimal assumptions. The operational formulation of quantum theory modelling quantum states and operations in terms of density matrices and completely positive maps suggests a way to describe causation operationally, without reference to a spacetime. When we combine such quantum maps together to form quantum circuits/networks, we are forming a causal structure that enables the in/output of one map to causally influence the in/output of another. However, in standard operational quantum theory, only combinations of operations resulting in an acyclic causal structure are considered. This means that every circuit singles out a direction of "time" even though no notion of spacetime was alluded to in the construction and we can always describe such a circuit as being immersed in a background spacetime (see for instance [50]). More generally, from a purely operational perspective, there is no reason to restrict to acyclic causal structures. We often have physical scenarios with feedback

where the output of a physical device is looped back and fed in to its input which are modelled by cyclic causal structures even in classical settings [53].

Here, we describe how quantum maps can be composed together through feedback loops to form a network of maps that can in general be associated with a cyclic causal structure. Noting that it is not possible to fully characterise causation under minimal assumptions, in the interest of generality, we consider the more operational notion of signalling which enables certain causal influences to be operationally detected. This allows us to talk about the signalling structure of any network of quantum maps, which too can in general be cyclic.

2.1 Composition of maps through feedback loops

To model composition of completely positive and trace preserving (CPTP) quantum maps, we adapt a definition proposed in [54]. While [54] describes a different framework (the causal box framework) that alludes to spacetime in its very construction and explicitly ensures that feedback loops only connect outputs at earlier spacetime locations to inputs in their future, their definition of composition for causal boxes can be formulated independently of spacetime, referring only to the in and output Hilbert spaces associated with quantum maps. Here we extract this purely operational part of this definition, applying it to the case of finite-dimensional Hilbert spaces. In later sections we will separately define a spacetime structure and then relate the operational and spacetime notions by characterising what it means for such cyclic compositions of maps to be compatible with relativistic causality in a spacetime.

Given two CPTP maps $\hat{\Phi}$ and $\hat{\Psi}$, we can consider three types of composition operations: *parallel composition*, *sequential composition* and *loop composition*. The second can be defined entirely in terms of the other two. The parallel composition of $\hat{\Phi}$ and $\hat{\Psi}$ is given in the obvious way, by their tensor product $\hat{\Phi} \otimes \hat{\Psi}$, then the set of input/output systems of the parallel composition is simply the union of the input/output systems of the individual maps. Loop composition is an operation on a single CPTP map $\hat{\Phi}$ where an output system O of the map is looped back and connected to an input system I of the same map that has matching dimensions, and is denoted as $\hat{\Phi}^{O \hookrightarrow I}$. In a slight, but harmless abuse of notation, we will, in the rest of the paper use the system label, e.g., I to also denote the state space of the system i.e., the set of all linear operators on the Hilbert space $\mathcal{H}_I$ associated with the system I. Then sequential composition $\hat{\Phi}_2 \circ \hat{\Phi}_1$ of two maps $\hat{\Phi}_1 : I_1 \mapsto O_1$ and $\hat{\Phi}_2 : I_2 \mapsto O_2$ with the former applied before the latter corresponds to first composing the maps in parallel to obtain the map $\hat{\Phi}_1 \otimes \hat{\Phi}_2 : I_1 \otimes I_2 \mapsto O_1 \otimes O_2$ followed by a loop composition connecting the output system O_1 to the input system I_2 to obtain $\hat{\Phi}_2 \circ \hat{\Phi}_1 := (\hat{\Phi}_1 \otimes \hat{\Phi}_2)^{O_1 \hookrightarrow I_2}$, as shown in Figure 1. More formally, we have the following definition. In the rest of the paper, whenever we refer to a CPTP map, it should be understood that this is a linear CPTP map, which is the case with all valid, normalised quantum operations. Furthermore we work with finite dimensional Hilbert spaces and assume that each Hilbert space $\mathcal{H}$ of dimension d has a well defined *computational basis* $\{|i\rangle\}_{i \in \{0,\ldots,d-1\}}$ consisting of orthonormal vectors $|i\rangle$. In the following, when we say an orthonormal basis, we typically mean such a computational basis.

Definition 2.1 (Loop composition of CPTP maps [54]) *Consider a CPTP map $\hat{\Phi} : \mathcal{L}(\mathcal{H}_{AB}) \mapsto \mathcal{L}(\mathcal{H}_{CD})$ with input systems A and B and output systems C and D, with $\mathcal{H}_B \cong \mathcal{H}_D$, and where $\mathcal{L}(\mathcal{H})$ denotes the set of linear operators on the Hilbert space $\mathcal{H}$. Let $\{|k\rangle_D\}_k$ and $\{|l\rangle_D\}_l$ be orthonormal bases of $\mathcal{H}_D$, and $\{|k\rangle_B\}_k$ and $\{|l\rangle_B\}_l$ denote the corresponding bases of $\mathcal{H}_B$ i.e., for all k and l, $|k\rangle_D \cong |k\rangle_B$ and $|l\rangle_D \cong |l\rangle_B$. The action of the new map $\hat{\Psi} = \hat{\Phi}^{D \hookrightarrow B}$ resulting from looping the output system D to the input system B, on basis elements $|i\rangle\langle j|_A$ of $\mathcal{L}(\mathcal{H}_A)$ is given as*

$$\hat{\Psi}(|i\rangle\langle j|_A) = \sum_{k,l} \langle k|_D \Big(\hat{\Phi}(|i\rangle\langle j|_A \otimes |k\rangle\langle l|_B) \Big) |l\rangle_D. \tag{1}$$

Note that the final map $\hat{\Psi}$ obtained after loop composition need not be CPTP, see Remark 2.2. The original definition of loop composition proposed in [54] is in terms of a Choi representation of CPTP maps on infinite dimensional systems, which is modelled as a sesquilinear positive semidefinite form. Here we restrict to the finite dimensional case and have therefore extracted a simpler but equivalent version of this definition. The reduction from the original definition to the above one is given in Appendix A.

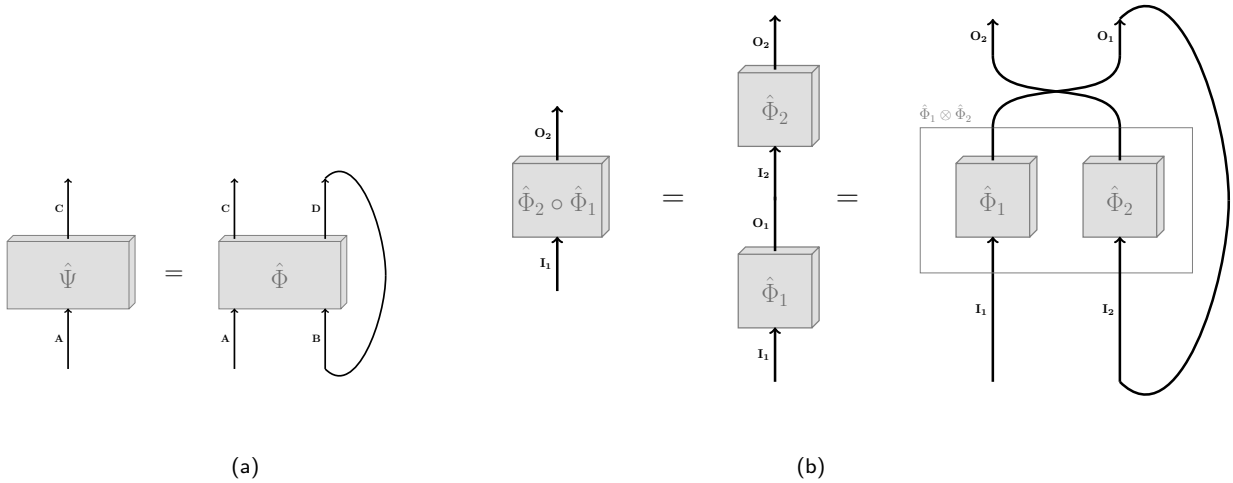

Figure 1: **Loop and sequential compositions** (a) The output system D is loop composed with the input system B of a map $\hat{\Phi}$ from A and B to C and D, to yield a new map $\hat{\Psi}$ with input A and output C. (b) Sequential composition of two maps, $\hat{\Phi}_1$ followed by $\hat{\Phi}_2$ can be obtained by first parallel composing them to obtain $\hat{\Phi}_1 \otimes \hat{\Phi}_2$ and then using loop composition, as explained in the main text and depicted here.

Networks of CPTP maps and causal structure Composing multiple CPTP maps together results in what we will refer to as a *network* of CPTP maps, which in itself could be described by a single map. The way in which we connect the in and output systems of a set of maps describes a causal structure that one would associate with such a network. We will use squiggly arrows $\rightsquigarrow$ to denote the causal arrows of such an underlying causal structure associated with a network of maps. We can also decompose each map in the network in terms of maps on smaller subsystems, and thus uncover the internal causal structure of each map. There are two ways to describe such a causal structure. For instance, sequentially composing $\hat{\Phi}_1 : I_1 \mapsto O_1$ and $\hat{\Phi}_2 : I_2 \mapsto O_2$ by connecting O_1 to I_2 can be viewed as a network of two maps, or a single map $\hat{\Phi} = (\hat{\Phi}_1 \otimes \hat{\Phi}_2)^{O_1 \hookrightarrow I_2}$. This allows inputs on I_1 to causally influence the outputs on O_2 and we can view this scenario as a causal structure in two ways, either by taking the maps $\hat{\Phi}_1$ and $\hat{\Phi}_2$ to be the nodes with a directed edge $\hat{\Phi}_1 \rightsquigarrow \hat{\Phi}_2$ (denoting causal influence) or by taking the in/output systems I_1, O_1, I_2 and O_2 to be the nodes, in which case we have $O_1 \rightsquigarrow I_2$ and additionally $I_1 \rightsquigarrow O_1$ and/or $I_2 \rightsquigarrow O_2$ depending on whether or not each map allows its input system to causally influence the output system (see the following paragraph for an example). Note that the latter view is in fact a more detailed description of the former, since it splits each map further into all its in and output systems and looks for causal dependences between these, while still containing all the causal dependences between different maps in the former view. In this paper, we will therefore adopt the latter view where the in/output systems form the nodes of a causal structure as this will allow us to formalise the existence of a causal influence $\rightsquigarrow$ from input to output of a map operationally without making assumptions about the internal structure of the maps.

Causation vs signalling The "causal structure" implied by a network of composed maps as described above does not fully capture what is meant by causation $\rightsquigarrow$. This is because the existence of a connecting map between two systems I and O does not imply that I is a cause of O, the connecting map may be a trivial one that discards the input on I and independently reprepares a state on O. Thus it is impossible to define the meaning of causation (or the edge $\rightsquigarrow$) without getting into the internal structure of the maps. However, we can consider a more operational way of detecting whether there is a causal influence– if inputting different input states on I (while keeping all other inputs fixed), results in different output states on O, then we can use I to *signal* to O, which we will denote as $I \to O$ and consequently we know that I causally influences O, which would justify $I \rightsquigarrow O$. Note crucially that while signalling implies causation, the converse is not true. Imagine a classical channel from I to O (both carrying classical bits), that takes a bit on I, internally generates a uniformly random bit K and outputs $O = I \oplus K$ (where $\oplus$ denotes modulo 2 addition). This is operationally equivalent to the trivial channel above that

discards I and reprepares a uniformly distributed O. However, here we have $I \rightsquigarrow O$ since I is indeed used to produce O (along with K), while we did not have this for the implementation that discards and reprepares. Further, I does not signal to O in both cases, since we have $P(O|I) = P(O)$. More generally, we may have signalling $I \to \{O_1, O_2\}$ from an input I jointly to a set of outputs O_1 and O_2 of a map $\hat{\Phi}$ without having signalling from I to O_1 or O_2 individually [2, 51]. Such a signalling relation suggests that we must have at least one of $I \rightsquigarrow O_1$ or $I \rightsquigarrow O_2$ even though we have neither $I \to O_1$ nor $I \to O_2$. Therefore, signalling between two systems is sufficient but not necessary for causation between them and while the causal structure may be a directed graph (with edges $\rightsquigarrow$) over the nodes $I_1, O_2, I_2, O_2, ...$, the signalling structure is in general a directed graph (with edges $\to$) over the set of all subsets of these nodes. To rigorously define causation, one would need to go into a full causal modelling framework such as [22]. In this paper, in the interest of generality, we aim to characterise causation under minimal assumptions and without imposing any unnecessary constraints on the internal structure of the maps. We will therefore focus our attention on signalling, instead of causation and this will be sufficient for all the general results we wish to establish. We define this more formally in the following section.

Remark 2.2 (Closedness under composition and linearity) *We note that linear CPTP maps are not closed under arbitrary compositions in general, since arbitrary loop composition may not result in systems producing valid normalised probabilities [54] and normalising them can introduce non-linearities.*[6] *Surprisingly, demanding that a composition of a set of linear CPTP maps results in a linear CPTP map does not rule out cyclic dependences. In Section 5 we show that the general framework we develop here can in particular be used to describe process matrices [28] which are regarded as high-order quantum maps that act on the space of standard quantum maps. There, we will see that the action of a process matrix on quantum maps can be described through a loop composition which in general results in a cyclic causal/signalling structure. However, by definition of process matrices, they map valid linear CPTP maps to valid linear CPTP maps [37] (or valid normalised probability distributions).*[7] *Some interesting (and not so easy) questions that we leave for future work are: given a set of linear CPTP maps, what is the largest set of allowed compositions under which the result of composition is also a linear CPTP map? What is the largest set of linear CPTP maps that is closed under arbitrary composition? We note that linear CPTP maps can be viewed as a subset of more general multi-time operators which can act non-linearly on quantum states (and other multi-time objects) and are closed under arbitrary composition [56]. In case of compositions of linear CPTP maps that result in a non-linear or trace decreasing map, the resulting object can be interpretted in the multi-time formalism (that involves measurements on pre and post-selected quantum states). The formal connection between our framework and the multi-time formalism is a subject of future work.*

2.2 Signalling structure of maps

We now define what it means to have signalling from a set of input systems to a set of output systems of a CPTP map. Consider a CPTP map $\hat{\Phi}$ with n input systems $I = \{I_1, ..., I_n\}$ and m output systems $O = \{O_1, ..., O_m\}$, where we use the system label $S \in I, O$ to also denote the state space associated with the system i.e., the set of all linear operators on the Hilbert space $\mathcal{H}_S$ of the system S. Then, we have the following, where $\mathrm{Tr}^{\rho}_{S_I} : S_I \to S_I$ denotes the operation that traces out the input on the systems in S_I and replaces it with some fixed state ρ_{S_I},

Definition 2.3 (Signalling structure of a CPTP map) *We say that there is a* signalling relation *from a set $S_I \subseteq I$ of input systems to a set $S_O \subseteq O$ of output systems of a CPTP map $\hat{\Phi}$ and denote it as $S_I \to S_O$ if there exists a state $\rho_{S_I} \in \bigotimes_{I_i \in S_I} I_i$ such that,*

$$\mathrm{tr}_{O \setminus S_O} \circ \hat{\Phi} \neq \mathrm{tr}_{O \setminus S_O} \circ \hat{\Phi} \circ \mathrm{Tr}^{\rho}_{S_I} \,. \tag{2}$$

[6] For instance, the projector $|0\rangle\langle 0|$ is a linear and completely positive but trace decreasing map. When applied to some state $\alpha|0\rangle + \beta|1\rangle$, it produces an un-normalised state $\alpha|0\rangle$. If we normalise this output state by dividing by its norm, the resulting map (apply projector and then renormalise) is completely positive and trace preserving but it is non-linear.

[7] A related result is that process matrices are operationally equivalent to a special linear subset of closed time-like curves [52] as well as a linear subset of pre and post-selected quantum states [55].

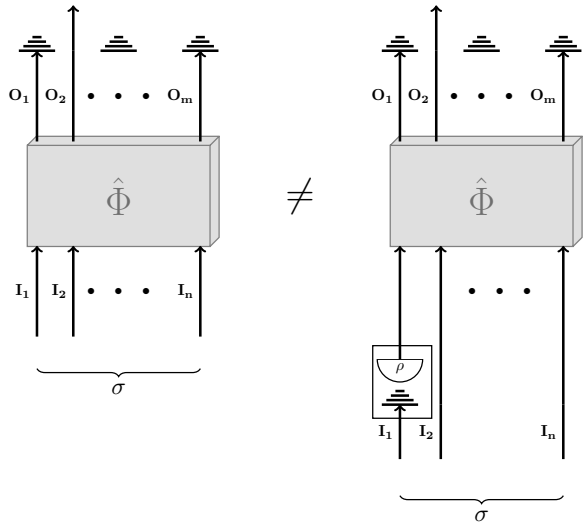

Figure 2: **Diagrammatic representation of Definition 2.3 of signalling** This figure represents the condition $I_1 \to O_2$ in the map $\hat{\Phi}$, whereby we require a global state σ on all inputs and a local state ρ on the input I_1 such that the two sides fail to be equal.

The set of all signalling relations of a CPTP map $\hat{\Phi}$ forms a directed graph $\mathcal{G}^{sig}_{\hat{\Phi}}$ the nodes of which correspond to arbitrary subsets of $I \cup O$, with an edge $\to$ between two nodes whenever there is signalling between those sets. We refer to this graph as the signalling structure of $\hat{\Phi}$.

The above definition is illustrated diagrammatically in Figure 2. Conversely, if $\mathrm{tr}_{O\setminus S_O} \circ \hat{\Phi} = \mathrm{tr}_{O\setminus S_O} \circ \hat{\Phi} \circ \mathrm{Tr}^{\rho}_{S_I}$ for all choices of ρ_{S_I} then we say that S_I does not signal to S_O and denote it as $S_I \not\to S_O$. Wherever the map $\hat{\Phi}$ being referred to is evident from context, or more generally when we refer to the signalling structure of a network of maps (see below), we will simply use $\mathcal{G}^{sig}$ instead of $\mathcal{G}^{sig}_{\hat{\Phi}}$.

Signalling structure of a network of maps Consider a set $\{\hat{\Phi}_i : I_i \mapsto O_i\}_{i=1}^n$ of CPTP maps, where $I_i = \{I_i^1, I_i^2, ..., I_i^{n_i}\}$ and $O_i = \{O_i^1, O_i^2, ..., O_i^{m_i}\}$ denote the set of all input and output systems respectively of the map $\hat{\Phi}_i$ which has n_i input systems and m_i output systems. When we form a network through arbitrary composition of such a set of maps, we allow the signalling structure of the network to be general enough to include all signalling relations coming from each individual map, as well as signalling relations indicating the connections specified by the composition (since the compositions connect systems through identity channels). This is because, in the most general case, a set of agents could potentially isolate a map from a network by discarding all the inputs to the map coming from other parts of the network, and instead freely chose input states to send in to the map and verify the signalling relations of the individual map. Thus, the signalling structure of a network that can be accessed in any physical setting depends on what assumptions are made about the power of the agents in that setting, whether they can make such "global interventions" on all inputs of a map or on multiple maps, or are restricted to sending receiving states associated with only certain in/output systems of a map/set of maps. The signalling structure of the network will in general be a directed graph over the nodes $\mathrm{Powerset}[(\bigcup_{i=1}^n I_i) \cup (\bigcup_{i=1}^n O_i)]$ and can contain directed cycles. For instance, we may compose two identity maps $\hat{\mathcal{I}}_1 : I_1 \mapsto O_1$ and $\hat{\mathcal{I}}_2 : I_2 \mapsto O_2$ (where all systems have the same dimension) by connecting O_1 to I_2 and O_2 to I_1. Then we have $I_1 \to O_1$ and $I_2 \to O_2$ coming from the union of individual signalling structures, and $O_1 \to I_2$ and $O_2 \to I_1$ coming from the composition, which gives the directed cycle $I_1 \to O_1 \to I_2 \to O_2 \to I_1$. The signalling structure can also include further direct signalling relations. For instance, using the network, we can also directly signal from I_1 to O_2 and we can also choose to separately include $I_1 \to O_2$ in the network's signalling structure, noting that signalling $\to$ need not always be a transitive relation, even in classical causal networks [2]. In the interest of generality, we allow all these possibilities in the signalling structure of the network. Note that the same network may admit multiple decompositions in terms of individual maps, and it is important to specify the set of all in and output systems over which the signalling structure is defined.

2.3 Compatibility of a signalling structure with a causal structure

Imagine we are told that there exists a CPTP map where $I \cup O$ denotes the set of all in and output systems of the map and $\mathcal{G}^{sig}$ is the signalling structure associated with that map. What can we then infer about the causal influences $\rightsquigarrow$ associated with the map, given no further information about it? While we have not fully and formally defined what $\rightsquigarrow$ means, we have motivated it with examples in the previous section and we would expect any meaningful notion of causation associated with this map to have the following necessary features. For a map $\hat{\Phi}$ with the set of input systems I and output systems O, we would like this notion of causation to specify, given any pair of systems $I_i \in I$ and $O_j \in O$, whether or not I_i causally influences O_j through the map $\hat{\Phi}$ i.e., the causal structure $\mathcal{G}^{caus}_{\hat{\Phi}}$ of the map is a directed graph over the nodes $I \cup O$ with the edges $\rightsquigarrow$. Secondly, whenever we find that a subset $S_I \subseteq I$ signals to a subset $S_O \subseteq O$ in $\hat{\Phi}$, then there must be at least one in and output pair, $I_i \in S_I$ and $O_j \in S_O$ such that $I_i \rightsquigarrow O_j$ in this causal structure. More generally, given a set of signalling relations over a set of systems $\mathcal{S}$, we can consider its compatibility with any directed graph, which may not necessarily be a directed graph over the nodes $\mathcal{S}$. For instance, we may wish to consider a directed graph corresponding to a spacetime, where nodes correspond to spacetime locations and the light cone structure of the spacetime specify the directed edges. We can then consider an assignment of systems in $I \cup O$ to the nodes of this causal structure in a one-to-one manner, which induces a causal structure over $I \cup O$, and we can then ask whether this is compatible with a given signalling structure over these systems.

With these minimal expectations, we propose the following definition of what it means for a signalling structure (possibly arising from an unknown network of maps) to be compatible with an arbitrary causal structure which we model as a directed graph. This allows us to consider whether the operational notion of signalling is compatible with different notions of causality. In the next section, we will apply these concepts to define relativistic causality in a spacetime.

Definition 2.4 (Causal structure) *A causal structure is any directed graph $\mathcal{G}^{caus}$, where* $\mathrm{Nodes}(\mathcal{G}^{caus})$ *denotes the set of all nodes of this graph and* $\mathrm{Edges}(\mathcal{G}^{caus})$ *denote the set of all edges. Unless specified otherwise, we will denote the edges of a directed graph $\mathcal{G}^{caus}$ as $\xrightarrow{C}$.*

Definition 2.5 (Embedding systems in a causal structure) *An embedding $\mathcal{E}$ of a set of systems $\mathcal{S}$ in a causal structure $\mathcal{G}^{caus}$ is an injective map $\mathcal{E} : \mathcal{S} \mapsto \mathrm{Nodes}(\mathcal{G}^{caus})$. For each system $S \in \mathcal{S}$, we will use S^{N_S} to denote the system embedded on the node $\mathcal{E}(S) = N_S$ of $\mathcal{G}^{caus}$, and refer to S^{N_S} as the $\mathcal{G}^{caus}$-embedded system. The subgraph of $\mathcal{G}^{caus}$ restricted to the nodes in the range of $\mathcal{E}$ is denoted as $\mathcal{G}^{caus}_{\mathcal{S}}$ and can be equivalently viewed as a graph over the node set $\{S^{N_S}\}_{S\in\mathcal{S}}$, we will then refer to such a causal structure as a causal structure over systems.*

Definition 2.6 (Implementing a CPTP map in a causal structure) *An implementation of a CPTP map $\hat{\Phi}$ (over a set $\mathcal{S}$ of in and output systems) in a causal structure $\mathcal{G}^{caus}$ with respect to an embedding $\mathcal{E}$ is a CPTP map $\hat{\Phi}^{\mathcal{G},\mathcal{E}}$ that is equivalent to $\hat{\Phi}$ up to a relabelling of input/output systems $S \Leftrightarrow S^{N_S}$ through an embedding $\mathcal{E}$ of $\mathcal{S}$ in $\mathcal{G}^{caus}$, where $\mathcal{E}(S) = N_S \in \mathrm{Nodes}(\mathcal{G}^{caus})$.*

It is important to note that a CPTP map $\hat{\Phi}$ may come with its own causal structure $\mathcal{G}^{caus}_{\hat{\Phi}}$ that specifes the internal connections entailed in $\hat{\Phi}$, but the above definition defines what it means to implement the map in any arbitrary causal structure $\mathcal{G}^{caus}$ which is achieved by embedding the in/output systems of $\hat{\Phi}$ in this new causal structure. Under such an implementation, two order relations come into play, $\rightsquigarrow$ associated with the causal structure $\mathcal{G}^{caus}_{\hat{\Phi}}$ of the map (some of which can be inferred by the signalling structure of the map) and $\xrightarrow{C}$ associated with the new causal structure $\mathcal{G}^{caus}$, which is a priory independent of $\hat{\Phi}$ until further constraints relating them are imposed. For instance, we may require that the signalling relations produced by $\hat{\Phi}$ are compatible with the causal structure $\mathcal{G}^{caus}$ in which it is implemented, which would in turn connect its internal causal structure $\mathcal{G}^{caus}_{\hat{\Phi}}$ with the new causal structure $\mathcal{G}^{caus}$. Such a compatibility condition is defined below.

Definition 2.7 (Compatibility of a signalling structure with a causal structure) *Let $\mathcal{G}^{sig}$ be a signalling structure associated with a network of maps where I and O denote the set of all input and output systems appearing in the network, represented as a directed graph over* $\mathrm{Powerset}(I \cup O)$ *with the directed edges*

$\rightarrow$. *Suppose that* $\mathcal{G}^{caus}$ *is some causal structure and* $\mathcal{E}$ *is an embedding of the systems* $I \cup O$ *in* $\mathcal{G}^{caus}$. *Then we say that the signalling structure* $\mathcal{G}^{sig}$ *is compatible with the causal structure* $\mathcal{G}^{caus}$ *with respect to the embedding* $\mathcal{E}$ *if the following holds.*

$$\exists \text{ directed path from } \mathcal{S}_1 \text{ to } \mathcal{S}_2 \text{ in } \mathcal{G}^{sig}$$

$$\Downarrow$$

$$\exists S_1 \in \mathcal{S}_1, S_2 \in \mathcal{S}_2 \text{ such that there is a directed path from } \mathcal{E}(S_1) \text{ to } \mathcal{E}(S_2) \text{ in } \mathcal{G}^{caus}.$$

Remark 2.8 *We note that a recent work [2] involving one of the authors also proposes a condition for compatibility between a set of (possibly cyclic) signalling relations and a causal structure (taken to be acyclic). As a clarification, we briefly point out the distinctions between the definitions and conditions of this and the present work. Firstly [2] adopts a causal modelling approach formulated under minimal conditions that need not necessarily arise from a composition of valid quantum maps, consequently the definition of signalling differs in these two works, even though they can be related (details left for future work). Further, the compatibility condition between a set of signalling relations (in the framework of [2], these are called* affects relations*) and a causal structure proposed in [2] are necessary and sufficient conditions for avoiding signalling outside the future of the causal structure through the embedded signalling relations. This is a weaker condition than our compatibility condition above which is not strictly necessary for "no signalling outside the future" in* $\mathcal{G}^{caus}$. *This is because we may have* A *signals jointly to* $\{B, C\}$ *without signalling individually to* B *or* C. *Then the compatibility condition of [2] only requires that the joint future of the locations of* B *and* C *in the causal structure is contained in the future of the location of* A *as this is where the joint signalling can be verified. This is possible even when* A*'s location in* $\mathcal{G}^{caus}$ *has no directed paths to either of* B *or* C*'s locations i.e.,* A *does not lie in the past of* B *or* C *with respect to the causal structure* $\mathcal{G}^{caus}$. *The above condition requires that* A*'s location in the causal structure must be in the past of at least one of* B *or* C*'s location if* A *signals to* $\{B, C\}$, *which is a stronger condition. This is because here, we require a more physically motivated condition for compatibiity with a causal structure while in [2] the goal was to identify the minimal conditions for avoiding signalling outside the future of a causal structure. And in fact, [2] uses this to establish the surprising mathematically possibility of causal loops embedded in Minkowski spacetime that do not lead to superluminal signalling, even though the existence of the causal loop can be operationally verified through suitable interventions. Such loops however violate the present definition of compatibility where* $\mathcal{G}^{caus}$ *is taken to represent the locations and lightcone structure of Minkowski spacetime.*[8]

2.4 Fine-graining causal structures, systems and maps

We now introduce an important concept that will feature in many of our results, namely that of fine-graining which can be applied to a causal structure, a set of systems, or a CPTP map. This captures the idea that the same physical protocol can be analysed in different levels of detail depending on the information that one wishes to capture (or has access to), which can in turn alter what we consider to be the "nodes" of a causal structure associated with such a scenario. Thus one may describe the same physical protocol through different causal structures, depending on the level of detail, and we will see that this affects the structural properties of the causal structure such as is cyclicity or acyclicity. We illustrate the concept with a few intuitive examples before proceeding to define it formally in the general case.

One example of a causal structure is an operational one that arises in a causal modelling framework [24] and provides a purely operational way of defining causation through the existence of connecting functions/maps between variables/systems. In a classical, deterministic causal model, the nodes of the causal structure are random variables and each variable X with a non-empty set of parents is obtained by applying a deterministic function f_X to the set of variables that are the parents of X in the causal structure. For parentless variables X, a probability distribution $P(X)$ is specified in the model. The

[8] These loops allow the causal influences $\rightsquigarrow$ of the operational causal structure to flow outside the future lightcone of the spacetime even though signalling stays within the future. The present compatibility condition is necessary for ensuring the edges $\rightsquigarrow$ of the underlying operational causal structure leading to the signalling relations also align with the edges of the new causal structure $\mathcal{G}^{caus}$ (which for instance could represent Minkowski spacetime).

following two examples illustrate two ways of fine-graining a model associated with a cyclic causal structure into an acyclic one. In the non-deterministic classical case, we have stochastic maps (from the set of parents of each variable to that variable) instead of deterministic functions.

2.4.1 Fine-graining by splitting into smaller subsystems

Consider a classical channel $P(Y|X)$ with input system X and output system Y which correspond to random variables. Suppose the channel is used multiple times and X_n, Y_n denote the input and output of the nth use of the channel, and let $X^n = \{X_1, ..., X_n\}$ (and $Y^n = \{Y_1, ..., Y_n\}$) denote the set of all inputs (and outputs) until the nth round. If the channel is used with feedback, the input X_n of nth round can in general depend on previous outputs Y^{n-1}. When there is feedback, we can in general view the cumulative in and outputs X^n and Y^n as the nodes of a cyclic causal structure $X^n \rightsquigarrow Y^n$ and $Y^n \rightsquigarrow X^n$ with Y^n computed from X^n through a stochastic map (given by the channel) and X^n computed from Y^n through another map (given by the feedback mechanism). However, if we describe the same situation through a causal model over the individual rounds' in and outputs, $\{X_i\}_i$ and $\{Y_i\}_i$, we know that we would have an acyclic causal structure $X_1 \rightsquigarrow Y_1 \rightsquigarrow X_2 \rightsquigarrow Y_2 ... Y_{n-1} \rightsquigarrow X_n \rightsquigarrow Y_n$. This shows that we can describe the same protocol through two different causal structures, one cyclic and another acyclic depending on the information captured by the nodes of the causal structure. The following example illustrates this with a more concrete construction.

Example 1 *Consider the cyclic causal structure over three random variables A, B and D shown in Figure 3a with $A \in \{0, 1, 2, 3\}$ and $B, D \in \{0, 1\}$. A classical causal model over this causal structure would be specified as $A = f_A(B, D)$, $B = f_B(A)$ and D distributed according to some distribution $P(D)$. Suppose f_A sets $A = 0, 1, 2, 3$ depending on whether $(D, B) = (0, 0), (0, 1), (1, 0), (1, 1)$ and f_B sets $B = 0$ whenever $A \in \{0, 1\}$ and $B = 1$ whenever $A \in \{2, 3\}$. If we now look carefully at the model, we can see that we can map A into two bits A_1, A_2 by identifying $A = 0, 1, 2, 3$ with $(A_1, A_2) = (0, 0), (0, 1), (1, 0), (1, 1)$. Then D specifies the first bit A_1 while B specifies the second bit A_2, and in addition B is itself the first bit A_1. Thus we can obtain an equivalent acyclic model where the new causal structure would be that of Figure 3b and the model would have $A_1 = D$, $B = A_1$ and $A_2 = B$, with the same distribution $P(D)$ over the parentless D. Combining the two binary A_1 and A_2 into a single 4-valued variable A using the above mentioned identification, we get back the original model.*

2.4.2 Fine-graining through uncertainty in location

Suppose that Alice and Bob share a classical channel, local random number generators (RNGs) along with a common source of randomness $\Lambda \in \{0, 1\}$ distributed according to a distribution $P(\Lambda)$ and they execute the following protocol. Whenever $\Lambda = 0$, Alice gets a random bit R_A from her RNG, sets her output $A = R_A$, and also forwards this value to Bob through the channel. Bob sets his output B to the value received from Alice. Whenever $\Lambda = 1$, Bob obtains R_B from his RNG, sets $B = R_B$, forwards the same to Alice who sets A to the value received from Bob. This can be modelled within a cyclic causal structure over R_A, R_B, Λ, A and B with a directed cycle $A \rightsquigarrow B$, $B \rightsquigarrow A$, as shown in Example 2 below.

This is a physically plausible protocol, even though it corresponds to a cyclic causal model. Whenever $\Lambda = 0$, Alice acts before Bob and whenever $\Lambda = 1$, Alice acts after Bob and whenever Λ is unknown, Alice's "location" with respect to Bob is uncertain. We can model the same protocol within an acyclic causal structure $A_1 \rightsquigarrow B \rightsquigarrow A_2$ with Alice's output A split into two nodes A_1 and A_2 corresponding to the cases where she acts before or after Bob.

When $\Lambda = 0$, $A_1 = R_A$, $B = A_1$ and A_2 is trivial, we can denote this by a "vacuum state", Ω corresponding to the absence of a physical message. When $\Lambda = 1$, $B = R_B$ and $A_2 = B$ while $A_1 = \Omega$ is trivial. Thus the bit which was earlier denoted by A and related to B through a cyclic causal model is now a bit that is "delocalised" in an acyclic structure over a greater number of nodes– there is a non-trivial bit valued message only at A_1 or at A_2 (and a vacuum state Ω at the other), depending on the value of Λ. Example 2 below describes this protocol first as a cyclic causal model over R_A, A, Λ, R_B and B where A and B are causes of each other and then as an acyclic causal model where the bit $A \in \{0, 1\}$ from the original model is associated with an uncertain location (depending on Λ) in the new acyclic causal structure.

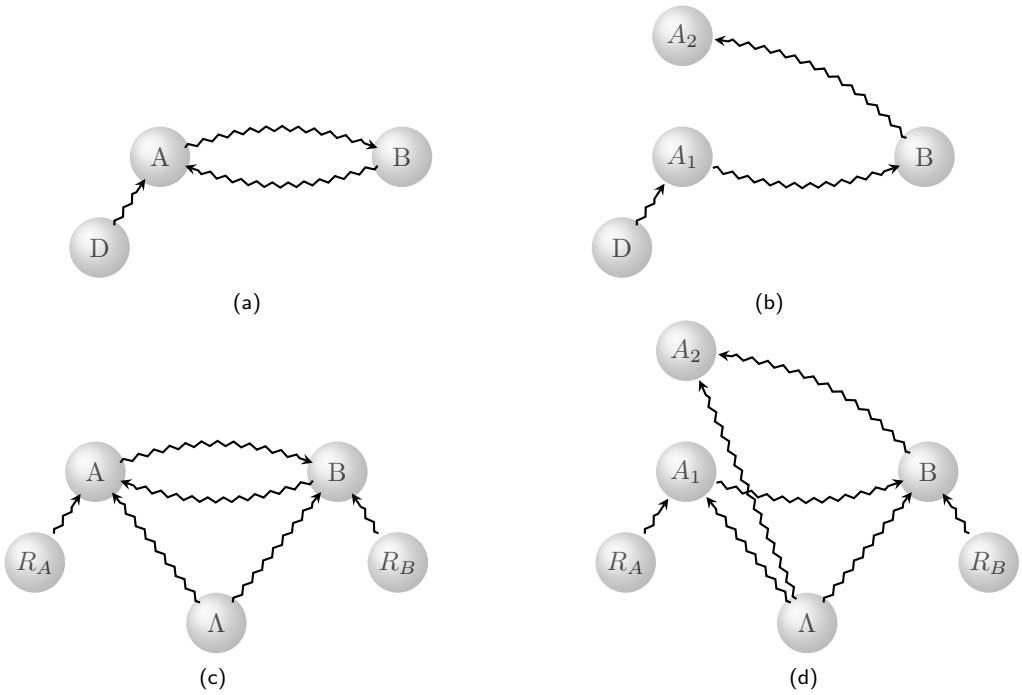

Figure 3: **Causal structures of Examples 1 and 2** (a) and (b) illustrate the original and fine-grained causal structures for the former while (c) and (d) illustrate the original and fine-grained causal structure for the latter example respectively

Example 2 *Consider the cyclic causal structure of Figure 3c and the following causal model over this causal structure. The parentless nodes R_A, R_B and Λ are distributed according arbitrary, non-deterministic distributions, $A = (\Lambda \oplus 1).R_A \oplus \Lambda.B$ and $B = (\Lambda \oplus 1).A \oplus \Lambda.R_B$, where $\oplus$ denotes modulo-2 addition. It is easy to check that this model implements the protocol described in the above paragraphs where the order in which Alice and Bob act (to generate A and B respectively) is decided by Λ. We can view the same situation as a scenario where the location of Alice's bit A in an acyclic causal structure (Figure 3d) is uncertain whenever Λ is unknown. A causal model describing this would be one where R_A, R_B and Λ have the same distributions as before, $B \in \{0, 1\}$ is still binary and the fine-grained nodes A_1 and A_2 take values in $\{\Omega, 0, 1\}$ which represent the absence ($A_1, A_2 = \Omega$) or presence ($A_1, A_2 \in \{0, 1\}$) of a bit-valued message at these locations. The values of each node can calculated given the values of its parents as $A_1 = (\Lambda \oplus 1).R_A \oplus \Lambda.\Omega$, $B = (\Lambda \oplus 1).A_1 \oplus \Lambda.R_B$, $A_2 = (\Lambda \oplus 1).\Omega \oplus \Lambda.B$, where $0.\Omega = 0$ and $1.\Omega = \Omega$. We can see that the acyclic causal structure coarse-grains to the original cyclic causal structure when we combine A_1 and A_2 into a single node A while preserving the in and outgoing causal arrows of this set. The original cyclic causal model is obtained from the fine-grained acyclic model by setting $A = i \in \{0, 1\}$ whenever $(A_1, A_2) \in \{(\Omega, i), (i, \Omega)\}$ i.e., when we are not interested in the "location" information but only in the "value" information of A. More explicitly, we can separately model the causal structure representing the locations as a directed graph $\mathcal{G}^{caus}$ containing a subgraph $N^1 \xrightarrow{C} N^2 \xrightarrow{C} N^3$ and we can view the process of fine-graining above as a more general type of "embedding" of the system A in the causal structure $\mathcal{G}^{caus}$ whereby it is assigned an uncertain location in the causal structure, N^1 or N^2, depending on Λ), while assigning B a fixed location N^2 in the new causal structure. Physically, $\mathcal{G}^{caus}$ might correspond to a spacetime structure (such as Minkowski spacetime) for instance, such that N^1 and N^2 specify the spacetime location of the message A. We will formalise this in the following sections.*

Messages vs systems: Here as well, as in the previous example, we have fine-grained a cyclic causal model for a protocol into an acyclic one describing the same protocol (or a particular implementation of it), by splitting a node A into multiple nodes A_1 and A_2. The difference however lies in the cardinality or size of the nodes. In Example 1, a system A of cardinality 4 was fine-grained into two systems A_1 and A_2

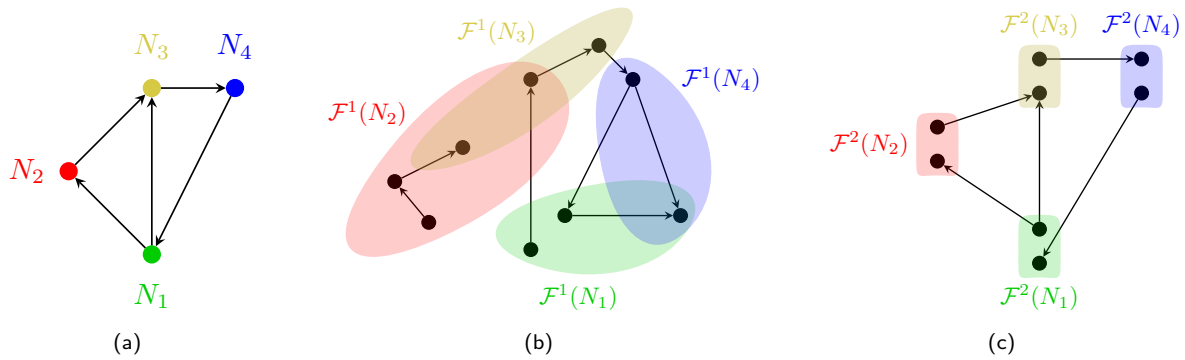

Figure 4: **Fine-grainings of a directed graph (color online)** A directed graph $\mathcal{G}$ (a) and two possible fine-grainings $\mathcal{G}^{\mathcal{F}^1}$ (b) and $\mathcal{G}^{\mathcal{F}^2}$ (c) of the same graph, both of which satisfy Definition 2.9. Each node N_i in $\mathcal{G}$ (with each $i \in \{1,2,3,4\}$ shown in a distinct color) maps to a set of nodes $\mathcal{F}^j(N_i)$ ($j \in \{1,2\}$) in the fine-grained graph $\mathcal{G}^{\mathcal{F}^j}$ lying within a blob of the same color.

each of cardinality 2. In Example 2, a system A of cardinality 2 is fine-grained into two systems A_1 and A_2 that each carry either zero or one 2-dimensional message, with the case of zero messages being represented by the vacuum state Ω. In such scenarios where there is uncertainty in location of a non-vacuum state with respect to a causal structure, we will refer to the vacuum state as representing zero messages and a d-dimensional non-vacuum state as a single d-dimensional message. In this case the systems A_1 and A_2 are associated with zero or one message of dimension 2. Notice that we can also take the view that each of A_1 and A_2 are three dimensional systems if the vacuum state is viewed as just another possible value of the variables. To avoid this ambiguity, we will interpret the vacuum state as "zero messages" in the rest of the paper and model values of variables and basis elements of non-trivial quantum spaces using natural numbers. In the above example, we had a classical bit A in a probabilistic mixture of different locations N^1 and N^3 with respect to a causal structure $\mathcal{G}^{caus}$. More generally, we can consider a state $|\psi\rangle$ of a d_S-dimensional quantum system S being at a superposition of different locations N^1 and N^3 in $\mathcal{G}^{caus}$. We can model this as a fine-graining of S into two systems S^{N_1} and S^{N_2}, and associate the state $\alpha|\psi\rangle^{S^{N_1}}|\Omega\rangle^{S^{N_2}} + \beta|\Omega\rangle^{S^{N_1}}|\psi\rangle^{S^{N_2}}$ with the fine-grained description. Under the fine-graining, each basis element $|v\rangle$ of $\mathcal{H}_S$, gets associated with the fine-grained state space $\text{Span}\{|v\rangle^{S^{N_1}}|\Omega\rangle^{S^{N_2}}, |\Omega\rangle^{S^{N_1}}, |v\rangle^{S^{N_2}}\}$ that corresponds to the message $|v\rangle$ being in an arbitrary superposition of the locations N^1 and N^2. This is similar to how temporal superposition states are modelled in previous frameworks such as [54].

State spaces under fine-graining: In both Examples 1 and 2, an original system S taking values in $\{0, 1, ..., d_S - 1\}$ splits into multiple systems S_1, S_2,... where the allowed states on each S_k is a subset of $\{\Omega, 0, 1, ..., d_S - 1\}$. Further, in both cases, we partition the state space of the fine-grained systems S_1, S_2,.. into disjoint subspaces, each of which map back to a value in the state space of the coarse-grained system S. More generally, nodes of a causal structure may be associated with quantum systems in which case a system S associated with a d_S-dimensional Hilbert space $\mathcal{H}_S$ can be fine-grained to a set of systems $\mathcal{S} = \{S_k\}_k$ where the state-space $\mathcal{H}_{S_k}$ of each S_k is isomorphic to $|\Omega\rangle \oplus \mathcal{H}_S$ and the overall state space $\mathcal{H}_{\mathcal{S}}$ of all the fine-grained systems is some subspace $\mathcal{H}_{\mathcal{S}} \subseteq \bigotimes_{S_k \in \mathcal{S}} \mathcal{H}_{S_k}$. We are now ready to define fine-graining of a general causal structure, and induced notions of fine-graining on systems embedded in this graph and on signalling relations over these embedded systems.

Definition 2.9 (Fine-graining of a directed graph) *A directed graph $\mathcal{G}_{\mathcal{F}}$ is called a fine-graining of another directed graph $\mathcal{G}$ if there exists a map $\mathcal{F} : \text{Nodes}(\mathcal{G}) \mapsto \text{Powerset}\big(\text{Nodes}(\mathcal{G}_{\mathcal{F}})\big)$ that maps each node $N \in \text{Nodes}(\mathcal{G})$ to a set of nodes $\mathcal{F}(N) \subseteq \text{Nodes}(\mathcal{G}_{\mathcal{F}})$ such that the following property holds. For any pair of distinct nodes $N_i, N_j \in \text{Nodes}(\mathcal{G})$, if there exists a directed path from N_i to N_j in $\mathcal{G}$ then there exists at least one pair of nodes $n_i \in \mathcal{F}(N_i)$ and $n_j \in \mathcal{F}(N_j)$ with a directed path from n_i to n_j in $\mathcal{G}_{\mathcal{F}}$.*

If the above definition is satisfied, we will refer to $\mathcal{F}$ as the fine-graining map and $\mathcal{G}$ as a coarse-graining of $\mathcal{G}_{\mathcal{F}}$. The concept is illustrated in Figure 4. Based on the motivation set out by the above

examples, we now define the fine-graining of quantum systems, with $\subseteq$ below to be read as "is a subspace of" and $\cong$ as "is isomorphic to".

Definition 2.10 (Fine-graining of quantum systems) *A set of quantum systems $\mathcal{S}_{\mathcal{F}}$ is called a fine-graining of another set $\mathcal{S}$ of quantum systems if there exists a map $\mathcal{F}^{sys} : \mathcal{S} \mapsto \text{Powerset}(\mathcal{S}_{\mathcal{F}})$ that maps each $S \in \mathcal{S}$ to a set of systems $\mathcal{F}^{sys}(S) \subseteq \mathcal{S}_{\mathcal{F}}$ where the joint state-space $\mathcal{H}_{\mathcal{F}^{sys}(S)}$ associated with the systems $\mathcal{F}^{sys}(S)$ is given as follows.*

$$\begin{aligned} \mathcal{H}_{\mathcal{F}^{sys}(S)} &\subseteq \bigotimes_{f_S \in \mathcal{F}^{sys}(S)} \mathcal{H}_{f_S}, \\ \mathcal{H}_{f_S} &\cong |\Omega\rangle \oplus \mathcal{H}_S. \end{aligned} \tag{3}$$

Furthermore, for every orthonormal basis $\{|v\rangle\}_{v\in\{0,1,\ldots,d_S-1\}}$ of $\mathcal{H}_S$ there exists a partition of $\mathcal{H}_{\mathcal{F}^{sys}(S)} = \bigoplus_{v\in\{0,1,\ldots,d_S-1\}} \mathcal{H}^v_{\mathcal{F}^{sys}(S)}$ into corresponding orthogonal subspaces, where the subspace $\mathcal{H}^v_{\mathcal{F}^{sys}(S)}$ of $\mathcal{H}_{\mathcal{F}^{sys}(S)}$ is the fine-graining of the basis element $|v\rangle$ of $\mathcal{H}_S$. For any $\mathcal{R} \subseteq \mathcal{S}$, $\mathcal{F}^{sys}(\mathcal{R})$ will be used as short form for $\{\mathcal{F}^{sys}(R)\}_{R\in\mathcal{R}}$

In particular, let $\mathcal{E} : S \in \mathcal{S} \mapsto N_S \in \text{Nodes}(\mathcal{G})$ be an an embedding of $\mathcal{S}$ in a causal structure $\mathcal{G}^{caus}$ (Definition 2.5). Any fine-graining $\mathcal{G}^{caus}_{\mathcal{F}}$ of $\mathcal{G}^{caus}$ associated with a fine-graining map $\mathcal{F}$ induces the following fine-graining $\mathcal{S}_{\mathcal{F}}$ of $\mathcal{S}$ by associating a fine-grained system S^{n_S} with each fine-grained node $n_S \in \mathcal{F}(N_S) \subseteq \text{Nodes}(\mathcal{G}^{caus}_{\mathcal{F}})$ through $\mathcal{F}^{sys}(S) := \{S^{n_S}\}_{n_S\in\mathcal{F}(N_S)}$.

$$\mathcal{S}_{\mathcal{F}} := \{S^{n_S}\}_{n_S\in\mathcal{F}(N_S), S\in\mathcal{S}} = \mathcal{F}^{sys}(\mathcal{S}).$$

We now define the fine-graining of a map, after setting out some notation. Let $\hat{\Phi}$ be a CPTP map with a set $I = \{I_1, ...I_n\}$ of input and a set $O = \{O_1, ..., O_m\}$ of output systems with $\mathcal{S} = I \cup O$, and $\{|v_i\rangle_{I_i}\}_{v_i}$ and $\{|u_j\rangle_{O_j}\}_{u_j}$ be orthonormal bases of $\mathcal{H}_{I_i}$ and $\mathcal{H}_{O_j}$ for $i \in \{1, ..., n\}$ and $j \in \{1, ..., m\}$. Then $\{|\mathbf{v}\rangle_I = \bigotimes_{i=1}^n |v_i\rangle_{I_i}\}_{v_1,\ldots,v_n}$ and $\{|\mathbf{u}\rangle_O = \bigotimes_{j=1}^m |u_j\rangle_{O_j}\}_{u_1,\ldots,u_m}$ are orthonormal bases of $\mathcal{H}_I$ and $\mathcal{H}_O$. For a fine-graining of $\mathcal{S}$ associated with the fine-graining map $\mathcal{F}^{sys}$, we will use $\mathcal{H}^{\mathbf{v}}_{\mathcal{F}^{sys}(I)} = \bigotimes_{i=1}^n \mathcal{H}^{v_i}_{\mathcal{F}^{sys}(I_i)}$ and $\mathcal{H}^{\mathbf{u}}_{\mathcal{F}^{sys}(O)} = \bigotimes_{j=1}^m \mathcal{H}^{u_j}_{\mathcal{F}^{sys}(O_j)}$ to denote the cumulative fine-grained spaces of I and O and $|\psi^{\mathbf{v}}\rangle \in \mathcal{H}_{\mathcal{F}^{sys}(I)}$ should be read as $|\psi^{\mathbf{v}}\rangle = \bigotimes_{i=1}^n |\psi^{v_i}\rangle$, and similarly for the output space.

Definition 2.11 (Fine-graining of quantum maps) *We say that the CPTP map $\hat{\Phi}_{\mathcal{F}}$ is a fine-graining of $\hat{\Phi}$ if there exists a fine-graining $\mathcal{F}^{sys}$ of the in/output systems $\mathcal{S}$ of $\hat{\Phi}$ such that the set of all in/output systems of $\hat{\Phi}_{\mathcal{F}}$ is given by $\mathcal{F}^{sys}(\mathcal{S})$, and $\hat{\Phi}_{\mathcal{F}}$ has a set $\mathcal{F}^{sys}(S)$ of in/output systems corresponding to each in/output system S of $\hat{\Phi}$. Further, for each basis state $|\mathbf{v}\rangle_I$ of $\mathcal{H}_I$*

$$\hat{\Phi} : |\mathbf{v}\rangle\langle\mathbf{v}| \mapsto \rho_O$$

for some $\rho_O = \sum_{\mathbf{u},\mathbf{u}'} p_{\mathbf{u},\mathbf{u}'} |\mathbf{u}\rangle\langle\mathbf{u}'|_O$ expressed in an orthonormal basis $\{|\mathbf{u}\rangle_O\}_{\mathbf{u}}$ of $\mathcal{H}_O$ if and only if for all states $|\psi^{\mathbf{v}}\rangle \in \mathcal{H}^{\mathbf{v}}_{\mathcal{F}^{sys}(I)}$ in the fine-grained subspace of $|\mathbf{v}\rangle_I$

$$\hat{\Phi}_{\mathcal{F}} : |\psi^{\mathbf{v}}\rangle\langle\psi^{\mathbf{v}}|_{\mathcal{F}^{sys}(I)} \mapsto \rho_{\mathcal{F}^{sys}(O)}$$

where $\rho_{\mathcal{F}^{sys}(O)} = \sum_{\mathbf{u},\mathbf{u}'} p_{\mathbf{u},\mathbf{u}'} |\psi^{\mathbf{u}}\rangle\langle\psi^{\mathbf{u}'}|_{\mathcal{F}^{sys}(O)}$ for some $|\psi^{\mathbf{u}}\rangle \in \mathcal{H}^{\mathbf{u}}_{\mathcal{F}^{sys}(O)}$ and $|\psi^{\mathbf{u}'}\rangle \in \mathcal{H}^{\mathbf{u}'}_{\mathcal{F}^{sys}(O)}$ belonging to the corresponding fine-grained subspaces of $|\mathbf{u}\rangle_O$ and $|\mathbf{u}'\rangle_O$.

Note that the signalling structure associated with a CPTP map or a network of such maps is also a directed graph, over the set of all subsets of all the in and output quantum systems associated with the map/network. The notion of fine-graining for signalling structures hence follows from Definition 2.9. In particular, if $\mathcal{G}^{sig}$ is a signalling structure over the set of systems $\mathcal{S}$, then $\text{Nodes}(\mathcal{G}^{sig}) = \text{Powerset}(\mathcal{S})$ as we can consider signalling (cf. Definition 2.3) between arbitrary subsets of quantum systems. Under a fine-graining, $\mathcal{S}$ transforms into a larger set $\mathcal{S}_{\mathcal{F}} := \{\mathcal{F}^{sys}(S)\}_{S\in\mathcal{S}}$ of systems and we would then be interested in a fine-grained signalling structure $\mathcal{G}^{sig}_{\mathcal{F}}$ where $\text{Nodes}(\mathcal{G}^{sig}_{\mathcal{F}}) = \text{Powerset}(\mathcal{S}_{\mathcal{F}})$. We then have the following lemma that relates the signalling structure of a CPTP map with that of its fine-graining, showing that signalling relations in a map are preserved under fine-graining of the map.

Lemma 2.12 *[Fine-graining a map preserves its signalling relations] Given a map $\hat{\Phi}$ and a fine-graining $\mathcal{F}^{sys}$ of its in/output systems $\mathcal{S}$, for every signalling relation $\mathcal{S}_I \to \mathcal{S}_O$ in $\hat{\Phi}$ between some subsets $\mathcal{S}_I \subset \mathcal{S}$ and $\mathcal{S}_O \subset \mathcal{S}$ of its input and output systems, there exists a corresponding signalling relation $\mathcal{F}^{sys}(\mathcal{S}_I) \to \mathcal{F}^{sys}(\mathcal{S}_O)$ in the fine-grained map $\hat{\Phi}_{\mathcal{F}}$. Consequently, the signalling structure $\mathcal{G}^{sig}_{\mathcal{F}}$ associated with $\hat{\Phi}_{\mathcal{F}}$ is a fine-graining of the signalling structure $\mathcal{G}^{sig}$ associated with $\hat{\Phi}$.*

The converse of the above lemma does not hold in general, as illustrated by the following simple, classical example. In general, there can be additional signalling relations at a fine-grained level that may get washed out under sufficient coarse-graining.

Example 3 *Let $\hat{\Phi}$ be a classical channel from an input bit $I \in \{0,1\}$ to an output bit $O \in \{0,1\}$ that discards the input I and deterministically prepares $O = 0$ as the output, then we immediately have $I \not\to O$ in $\hat{\Phi}$. We now construct a fine-graining of $\hat{\Phi}$ where $\mathcal{F}^{sys}(I) \to \mathcal{F}^{sys}(O)$ holds. Consider the channel $\hat{\Phi}_{\mathcal{F}}$ with the input systems $I_1, I_2 \in \{0,1\}$ and output systems $O_1, O_2 \in \{0,1\}$ acting as $\hat{\Phi}_{\mathcal{F}} : (I_1, I_2) \mapsto (O_1 = I_1 \oplus I_2, O_2 = I_1 \oplus I_2)$ (where $\oplus$ denotes modulo-2 addition). Let $\mathcal{F}^{sys}(I) = \{I_1, I_2\}$ and $\mathcal{F}^{sys}(O) = \{O_1, O_2\}$ be a fine-graining of the systems I and O, where the input value $I = 0$ in $\hat{\Phi}$ is identified with the subspace $(I_1, I_2) \in \{(0,0),(1,1)\}$ of possible input values in $\hat{\Phi}_{\mathcal{F}}$ and $I = 1$ with the orthogonal subspace $(I_1, I_2) \in \{(0,1),(1,0)\}$, and the output values $O = 0$ and $O = 1$ are similarly identified with the orthogonal subspaces $(O_1, O_2) \in \{(0,0),(1,1)\}$ and $(O_1, O_2) \in \{(0,1),(1,0)\}$ in the fine-grained map. In other words the coarse grained variables encode the parity of the corresponding two fine-grained variables. We can see that $\hat{\Phi}_{\mathcal{F}}$ is indeed a fine-graining of $\hat{\Phi}$, since $\hat{\Phi}$ maps every $I \in \{0,1\}$ to $O = 0$ while $\hat{\Phi}_{\mathcal{F}}$ maps every input to outputs in the fine-grained subspace $(O_1, O_2) \in \{(0,0),(1,1)\}$ associated with $O = 0$. Further, we can see that $\{I_1, I_2\} \to \{O_1, O_2\}$ in $\hat{\Phi}_{\mathcal{F}}$ since $\hat{\Phi}_{\mathcal{F}}(I_1 = 0, I_2 = 0) \neq \hat{\Phi}_{\mathcal{F}}(I_1 = 0, I_2 = 1)$.*

In the next section, we will apply these general concepts to define the implementations of CPTP maps and networks of maps in a fixed spacetime where Equation (3) will be used to explicitly define the state-spaces of quantum systems embedded in a spacetime.

3 Spacetime structure and relativistic causality

Definition 3.1 (Fixed acyclic spacetime) *We model a fixed acyclic spacetime by a partially ordered set $\mathcal{T}$ associated with the order relation $\preceq$, without assuming any further structure/symmetries. $P \preceq Q$ is denoted as $P \prec Q$ whenever $P, Q \in \mathcal{T}$ are distinct elements. Then $P \prec Q$, $P \succ Q$ and $P \not\prec\not\succ Q$ represent P being in the past of, future of and neither in the past nor future of Q respectively, with respect to this order relation.*

Spacetime as an abstract causal structure Modelling spacetime as an abstract partially ordered set means that we can regard $\mathcal{T}$ as a directed graph with the order relation $\prec$ playing the role of the directed edges and this would define a causal structure (as in Definition 2.7). By virtue of being a partial order relation, this would correspond to a directed acyclic graph. For example, in the particular case of Minkowski spacetime, the partial order would correspond to the lightcone structure. More generally, one may also consider spacetime structures that are not partially ordered such as those arising from exotic closed timelike curve solutions to Einstein's equations, here $\prec$ would be a pre-order relation. Our framework and definitions would easily generalise to this case. However, we focus on the case of partially ordered spacetimes in this paper, as our goal is to apply this framework to consider the properties of networks of maps that can be experimentally implemented in a physical spacetime, such as Minkowski spacetime. Once the operational and spacetime notions of causation are disentangled as done here, the question of how cyclic signalling structures can be embedded compatibly in an acyclic spacetime becomes a much more interesting question than the case of a cyclic spacetime. Understanding what kind of tasks are fundamentally impossible to implement in definite acyclic spacetimes (even when allowing for quantum systems to be delocalised in space and time), would also shed light on how physics can differ in more exotic spacetimes, be it cyclic spacetimes or quantum indefinite spacetimes. This will be the aim of this paper, and we therefore focus on fixed acyclic spacetimes. While the causal structure with spacetime locations as nodes is a directed acyclic graph in this case, we can also consider a causal structure with spacetime regions as the nodes and naturally define an order relation on these, which will in general not be a partial order.

Ordering and fine-graining spacetime regions: We formally define the order relation on spacetime regions below, this will help in understanding how a set of cyclic signalling relations can be implemented in an acyclic spacetime.

Definition 3.2 (Order relation on spacetime regions) *Let $\mathcal{P}^1, \mathcal{P}^2 \subseteq \mathcal{T}$ be two distinct subsets of locations (or "regions") in a spacetime $\mathcal{T}$. We say that $\mathcal{P}^1 \xrightarrow{R} \mathcal{P}^2$ if there exists at least one pair of locations $P^1 \in \mathcal{P}^1$ and $P^2 \in \mathcal{P}^2$ such that $P^1 \prec P^2$. More generally, we will refer to a directed graph $\mathcal{G}^R_{\mathcal{T}}$ as a region causal structure of $\mathcal{T}$ if $\mathrm{Nodes}(\mathcal{G}^R_{\mathcal{T}}) \subseteq \mathrm{Powerset}(\mathcal{T})$ and its edges are given by $\xrightarrow{R}$.*

Remark 3.3 *Note that $\xrightarrow{R}$ is neither a pre-order nor a partial order relation as it is non-transitive and we can in general have $\mathcal{P}_1 \xrightarrow{R} \mathcal{P}_2$ as well as $\mathcal{P}_2 \xrightarrow{R} \mathcal{P}_1$ for any two spacetime regions. Further, any partition of each spacetime region $\mathcal{P} \subseteq \mathcal{T}$ into mutually disjoint sub-regions $\mathcal{P} = \cup_i \mathcal{P}_i$ defines a fine-graining of $\mathcal{G}^R_{\mathcal{T}}$ (Definition 2.9) since it follows from Definition 3.2 that $\mathcal{P} = \cup_i \mathcal{P}_i \xrightarrow{R} \mathcal{Q} = \cup_j \mathcal{P}_j$ implies that there exist $\mathcal{P}_i$ and $\mathcal{Q}_j$ such that $\mathcal{P}_i \xrightarrow{R} \mathcal{Q}_j$. The spacetime $\mathcal{T}$ itself when represented as a directed acyclic graph corresponds to a fine-graining of any such $\mathcal{G}^R_{\mathcal{T}}$ which we will refer to as the* maximal fine-graining *of $\mathcal{G}^R_{\mathcal{T}}$. This simply corresponds to writing each region $\mathcal{P} \in \mathrm{Nodes}(\mathcal{G}^R_{\mathcal{T}})$ in terms of the individual spacetime locations in $\mathcal{T}$ that comprise it $\mathcal{P} = \cup_{P \in \mathcal{P}} P$. While $\mathcal{G}^R_{\mathcal{T}}$ can in general be cyclic, its maximal fine-graining, with nodes corresponding to elements of $\mathcal{T}$, would always be acyclic since $\mathcal{T}$ is a partially ordered set.*

3.1 Implementing quantum maps in a spacetime

The spacetime structure itself is devoid of any operational meaning until we embed physical systems in it, until this point, it is simply an abstract causal structure by virtue of being a directed graph. An agent may assign physical meaning to a spacetime point P only when it can be associated with some operational event (one that can in principle be operationally verified) such as "I received message from Bob at the spacetime location P". In our case, the physical systems are the in and output Hilbert spaces of the quantum maps. When we implement a CPTP map in a spacetime structure, we are associating spacetime regions with the in and output systems of the map i.e., we are embedding the systems in a causal structure $\mathcal{G}^R_{\mathcal{T}}$ whose nodes are subsets of $\mathcal{T}$ and edges $\xrightarrow{R}$ are given by Definition 3.2, as formalised below.

Definition 3.4 (Fixed spacetime implementation of a CPTP map) *A fixed spacetime implementation of $\hat{\Phi}$ in a spacetime $\mathcal{T}$ with respect to an embedding $\mathcal{E}$ is an implementation $\hat{\Phi}^{\mathcal{G}^R_{\mathcal{T}},\mathcal{E}}$ of $\hat{\Phi}$ in the causal structure $\mathcal{G}^R_{\mathcal{T}}$ over spacetime regions in $\mathcal{T}$ (cf. Definition 2.6), which we will denote as $\hat{\Phi}^{\mathcal{T},\mathcal{E}}$ for short.*

The same CPTP map $\hat{\Phi}$ can have different implementations $\hat{\Phi}^{\mathcal{T},\mathcal{E}}$, $\hat{\Phi}^{\mathcal{T},\mathcal{E}'}$ in the same spacetime $\mathcal{T}$, corresponding to different choices of spacetime embeddings $\mathcal{E}$, $\mathcal{E}'$.

Fine-graining fixed spacetime implementations We can now use the concepts previously defined to consider the fine-graining of a spacetime implemented map. Each partition of the regions forming the nodes of $\mathcal{G}^R_{\mathcal{T}}$ into subregions, defines a fine-graining $\mathcal{G}^R_{\mathcal{T},\mathcal{F}}$ of $\mathcal{G}^R_{\mathcal{T}}$ (cf. Remark 3.3), which defines a corresponding fine-graining $\hat{\Phi}^{\mathcal{T},\mathcal{E}}_{\mathcal{F}}$ of the spacetime implemented map $\hat{\Phi}^{\mathcal{T},\mathcal{E}}$. Explicitly, expressing a region $\mathcal{P} \subseteq \mathcal{T}$ in terms of disjoint subregions $\mathcal{P}^i \subseteq \mathcal{T}$ as $\mathcal{P} = \cup_i \mathcal{P}^i$ induces a fine-graining $\mathcal{F}^{sys}(S) = \{S^{\mathcal{P}^i}\}_i$ of the systems $S \in \mathcal{S}$ which in turn defines a fine-graining $\hat{\Phi}^{\mathcal{T},\mathcal{E}}_{\mathcal{F}}$ of the map $\hat{\Phi}^{\mathcal{T},\mathcal{E}}$ (or equivalently of the original map $\hat{\Phi}$ since this acts equivalently upto relabelling of the in/output systems as $S \Leftrightarrow S^{\mathcal{P}^S}$), as per Definition 2.10. The state-spaces of the fine-grained systems are given according to Equation (3). In particular, when we maximally fine-grain each region in terms of the individual spacetime points, $\mathcal{P}^S = \{P^S \in \mathcal{T}\}_{P^S \in \mathcal{P}^S}$, each system S is correspondingly fine-grained to a set of systems $\mathcal{F}^{sys}_{max}(S) = \{S^{P^S}\}_{P^S \in \mathcal{P}^S}$, such that each spacetime location $P^S \in \mathcal{P}^S$ becomes associated with a Hilbert space $\mathcal{H}_{S^{P^S}} \cong |\Omega\rangle \oplus \mathcal{H}_S$. This means that whenever a system S is embedded in a spacetime region $\mathcal{P}_S \subseteq \mathcal{T}$, when we sufficiently fine-grain, we can associate a copy of the state space of the original system (augmented with a vacuum state) at each spacetime location in $\mathcal{P}_S$. We will refer to each such system S^{P^S} associated with a single spacetime location $P^S \in \mathcal{T}$ as an *elemental subsystem* of the implementation $\hat{\Phi}^{\mathcal{T},\mathcal{E}}$. The set of all elemental subsystems is then given as $\mathcal{F}^{sys}_{max}(\mathcal{S}) = \{\mathcal{F}^{sys}_{max}(S)\}_{S \in \mathcal{S}}$.

3.2 Relativistic causality

Once we implement a map in a spacetime as described above, we can then use the signalling relations associated with the map to signal between different spacetime regions. Now two order relations come into play, one is the order relation between in/output systems of the map implied by the signalling relations $\rightarrow$, and the other is the order relation $\preceq$ of the spacetime $\mathcal{T}$ and the induced order $\xrightarrow{R}$ on spacetime regions. A priory, these order relations are completely independent of each other, and it is the notion of relativistic causality that connects the two by demanding compatibility of the signalling relations of a map with the abstract causal structure of a spacetime. This can be seen as a special case of the more general Definition 2.7 which defines compatibility of a signalling structure with a causal structure. But we state it below explicitly for completeness, and for the ease of cross referencing.

Definition 3.5 (Relativistic causality (special case of Definition 2.7)) *Let $\mathcal{G}^{sig}$ be a set of signalling relations over a set $\mathcal{S}$ of systems and $\mathcal{E} : S \mapsto S_{\mathcal{P}_S}$ be an embedding of the systems $S \in \mathcal{S}$ in the region causal structure $\mathcal{G}^R_{\mathcal{T}}$ associated with a spacetime $\mathcal{T}$. Then, we say that the signalling structure $\mathcal{G}^{sig}$ does not violate relativistic causality with respect to the embedding $\mathcal{E}$ in the spacetime $\mathcal{T}$ if the signalling relations in $\mathcal{G}^{sig}$ are compatible with the causal structure $\mathcal{G}^R_{\mathcal{T}}$ under the embedding $\mathcal{E}$, according to Definition 2.7.*

The above relativistic causality condition is necessary for insuring that the causal dependencies given by the information-theoretic causal structure of the CPTP map(s) flow from past to future in the spacetime in which it is implemented. It need not be sufficient, since we are working under minimal assumptions for characterising causation. The following theorem shows that it is possible to embed any set of (possibly cyclic) signalling relations in a partially ordered spacetime, such as Minkowski spacetime without violating the relativistic causality condition.

Theorem 3.6 *[Embedding arbitrary, cyclic signalling relations in spacetime] For every signalling structure $\mathcal{G}^{sig}$, there exists a fixed acyclic spacetime $\mathcal{T}$ and an embedding $\mathcal{E}$ of $\mathcal{G}^{sig}$ in a region causal structure $\mathcal{G}^R_{\mathcal{T}}$ of $\mathcal{T}$ that respects relativistic causality.*

Remark 3.7 *We note that Definition 3.1 is a very minimal definition of spacetime that does not assume any of the symmetries or differentiable manifold structure that is usually associated with spacetime in relativistic physics. This means that according to this definition, we would regard two distinct partially ordered sets $\mathcal{T}$ and $\mathcal{T}'$ are two different spacetimes. However, if we model spacetime instead as globally hyperbolic manifold, then we could sample a finite number of points from the same spacetime to form different partially ordered sets $\mathcal{T}$ and $\mathcal{T}'$ from the same manifold. And indeed, if we model the spacetime as a globally hyperbolic manifold, the statement of the above theorem becomes a stronger one, we will instead have that for every signalling structure and every spacetime manifold, there exists an embedding with the said properties, and the region causal structure in this case will have nodes that correspond to a finite collection of points in the manifold. The proof of this statement is also included in the proof of the above theorem.*

For the spacetime implementation $\hat{\Phi}^{\mathcal{T},\mathcal{E}}$ of a map $\hat{\Phi}$ to respect relativistic causality, we would naturally require the signalling relations at each level of fine-graining to respect relativistic causality and this naturally extends to a network of maps, as formalised below.

Definition 3.8 (Relativistic causality for a network of spacetime embedded maps) *A fixed spacetime implementation of a network of maps is said to satisfy relativistic causality only if every map $\hat{\Phi}_i^{\mathcal{T},\mathcal{E}}$ in the spacetime implemented network is such that under every fine-graining $\hat{\Phi}_i^{\mathcal{T},\mathcal{E},\mathcal{F}}$ of the map, the corresponding signalling relations $\mathcal{G}^{\mathcal{T},\mathcal{E},\mathcal{F}}$ satisfy the relativistic causality condition of Definition 3.5.*

For example, suppose we have a map $\hat{\Phi} : A \mapsto B$ with the signalling relation $A \rightarrow B$ and the systems A and B are embedded in spacetime such that $\mathcal{P}^A = \{P_1, P_2\}$ and $\mathcal{P}^B = \{P\}$ for some spacetime locations P_1, P_2 and P. This defines a fixed spacetime implementation $\hat{\Phi}^{\mathcal{T},\mathcal{E}}$ of the map and its maximal fine-graining $\Phi^{\mathcal{T},\mathcal{E},\mathcal{F}_{max}}$ which has the elemental subsystems $\mathcal{F}^{sys}_{max}(A) = \{A^{P_1}, A^{P_2}\}$ as the input systems and the elemental subsyste, $\mathcal{F}^{sys}_{max}(B) = \{B^P\}$ as the output system. Then, since we have $A \rightarrow B$, relativistic causality would require that $\mathcal{P}^A \xrightarrow{R} \mathcal{P}^B$ which is equivalent to saying that either $P_1 \prec P$ or $P_2 \prec P$ must hold. In the fine-grained map, the corresponding signalling structure must now specify whether or not

there is signalling between the elemental subsystems. We could have $A^{P_1} \to B^P$, $A^{P_2} \to B^P$ or both in this signalling structure. All three cases would coarse-grain to give back the original signalling relation $A \to B$ (when we combine A^{P_1} and A^{P_2} to a single node A and relabel B^P to B while preserving the edge structure), but notice that we would in general have more relativistic causality constraints on the embedding, from each level of the fine-graining. The following lemma then immediately follows from the above definitions.

Lemma 3.9 *[Fine-graining to an acyclic signalling structure] Every network of CPTP maps that admits an implementation in a fixed spacetime $\mathcal{T}$ that does not violate relativistic causality in that spacetime admits a fine-graining that has a definite acyclic signalling structure, whose edges $\to$ align with the partial order relation $\prec$ of the spacetime.*

Composition of fixed spacetime implementations As fixed spacetime implementations of CPTP maps are themselves CPTP maps, we can consider compositions of these. Let $\hat{\Phi}_1 : I_1 \mapsto O_1$ and $\hat{\Phi}_2 : I_2 \mapsto O_2$ be two CPTP maps and let $\hat{\Phi} = \hat{\Phi}_2 \circ \hat{\Phi}_1$ be their sequential composition obtained by connecting O_1 to I_2. Then we can describe the fixed spacetime implementation $\hat{\Phi}^{\mathcal{T},\mathcal{E}}$ of $\hat{\Phi}$ in terms of the fixed spacetime implementations $\hat{\Phi}_1^{\mathcal{T},\mathcal{E}_1}$ and $\hat{\Phi}_2^{\mathcal{T},\mathcal{E}_2}$ of its components $\hat{\Phi}_1$ and $\hat{\Phi}_2$ by requiring that the embeddings $\mathcal{E}_1 : I_1 \mapsto I_1^{\mathcal{P}^{I_1}}, O_1 \mapsto O_1^{\mathcal{P}^{O_1}}$ and $\mathcal{E}_2 : I_2 \mapsto I_2^{\mathcal{P}^{I_2}}, O_2 \mapsto O_2^{\mathcal{P}^{O_2}}$ assign the same spacetime region to the systems being connected i.e., $\mathcal{P}^{O_1} = \mathcal{P}^{I_2} := \mathcal{P}$ and setting $\hat{\Phi}^{\mathcal{T},\mathcal{E}} := \hat{\Phi}_2^{\mathcal{T},\mathcal{E}_2} \circ \hat{\Phi}_2^{\mathcal{T},\mathcal{E}_2}$ where the composition now connects each elemental system O_1^P, $P \in \mathcal{P}$ associated with O_1 to the corresponding elemental system I_2^P. In this manner, we can define the spacetime implementation of any network formed by arbitrary composition of another set of maps, in terms of the fixed spacetime implementations of its constituent maps.

Remark 3.10 (Single vs multiple uses of a map in spacetime) *When we implement a CPTP map $\hat{\Phi}$ in a spacetime by assigning regions to its in/output systems, do we allow the map the be used multiple times in the assigned spacetime regions or is the map used only once on a non-vacuum state, that may be in a superposition of different spacetime locations? A priory, Definition 3.4 does not forbid multiple uses. For example, suppose that $\hat{\Phi} : I \mapsto O$ is an identity channel with $\mathcal{H}_I = \mathcal{H}_O = |\Omega\rangle \oplus \mathbb{C}^2$. Consider a fixed spacetime implementation of this map where $\mathcal{P}_I := \{(\mathbf{r}, t_i)\}_{i\in\{1,3,5,\ldots\}}$ and $\mathcal{P}_O := \{(\mathbf{r}, t_j)\}_{j\in\{2,4,6,\ldots\}}$ are the spacetime regions assigned to I and O in Minkowski spacetime, expressed in spatial temporal co-ordinates with respect to some inertial frame. Now the spacetime implementation $\hat{\Phi}^{\mathcal{T},\mathcal{E}}$ could apply the identity map independently at each spacetime location, such that a state arriving in I at $(\mathbf{r}, t_i)$ is mapped to the same state on O at $(\mathbf{r}, t_{i+1})$. When we input a non-vacuum state at each of the possible input locations, this would appear as though the map is used multiple times, once between each pair $(\mathbf{r}, t_i)$ and $(\mathbf{r}, t_{i+1})$ of in and output locations. On the other hand, if we wish to restrict to a single use of the map on a non-vacuum state, we can simply restrict state space of each in/output system of the map to be a suitable subspace of the space defined in Equation (3). In particular, we would have the following subspace that models a d_S dimensional quantum message in a superposition of spacetime locations in $\mathcal{P}^S \subseteq \mathcal{T}$.*

$$\mathrm{Span}\{|v\rangle^{S^P} \bigotimes_{R\in\mathcal{P}_S\setminus\{P\}} |\Omega\rangle^{S^R}\}_{v\in\{0,1,\ldots d_S-1\}, P\in\mathcal{P}_S}.$$

Restricting to this subspace would mean that whenever a system S of the original map $\hat{\Phi}$ carries a single d^S dimensional state, then the corresponding system $\mathcal{E}(S)$ of the spacetime implementation $\hat{\Phi}^{\mathcal{T},\mathcal{E}}$ would also carry exactly a single d^S dimensional state, but the state may be delocalised (classically or quantumly) over spacetime locations in the region $\mathcal{P}^S$ specified by the embedding $\mathcal{E}$.

4 Review of the process matrix framework

The process matrix framework [28] describes multi-partite scenarios where the parties act within local quantum laboratories compatible with a local ordering of events within each laboratory, in the absence of a global order or spacetime structure conneting different laboratories. The global behaviour is characterised by a *process matrix*, which models the "outside environment" of these local labs (which the parties cannot access) and encodes information about how the local labs interact. Each party is associated with a corresponding local quantum laboratory and there are certain setup assumptions made about

the operations implemented in these local laboratories, which justify the mathematical definitions of the framework. Here we briefly review the process matrix framework, starting with these assumptions and then moving onto the formal definition of process matrices, and different classes of process matrices.

4.1 Assumptions and framework preliminaries

Assumptions: There are three relevant events associated with each party (or corresponding local laboratory), A_I which corresponds to the local input event at which the party receives a quantum system from the outside environment, A_U which is the event at which they apply their local operation to this system and A_O which is the local output event at which they send a quantum system to the outside environment. It is assumed that there is an ordering between these local events (even in the absence of a global order between events corresponding to different parties), each party receives an input system to their lab, then applies their local operation and then sends an output to the environment i.e, the input event A_I precedes the operation event A_U which precedes the output event A_O for each party A. We will refer to this as the local order (LO) assumption. Further, the labs are assumed to be closed to external in/outputs otherwise, that is A_I and A_O are the only events through which the lab interacts with the outside environment. This is called the closed lab (CL) assumption. Note that no spacetime information is explicitly considered but even in the absence of information about the absolute time of occurrence of these operational events, it is in principle possible to ensure that they are ordered in a certain way. Finally, it is assumed that the parties can freely chose the local operation performed in their lab, this choice can be encoded in a classical setting as we will see later. This corresponds to the free choice (FC) assumption.

Local behaviour: local quantum experiments. Each party A acts within their respective local laboratory, performing a *local quantum experiment* associated with the input Hilbert space $\mathcal{H}_{A_I}$ of dimension d_{A_I} and output Hilbert space $\mathcal{H}_{A_O}$ of dimension d_{A_O}.[9] The operations performed by agents during the course of their local experiments are described by *quantum instruments* $\mathcal{J}^A = \{\mathcal{M}^A_x\}_{x=1}^m$ with $\mathcal{M}^A_x : A_I \to A_O$ [28, 30]. Here, x parametrizes the possible local measurement outcomes, and A_I and A_O represent the set of all linear operators over $\mathcal{H}_{A_I}$ and $\mathcal{H}_{A_O}$ respectively. A classical setting a can be used to specify the choice of operation implemented on the input Hilbert space, for instance, this can act as a measurement setting. Then the corresponding instrument is denoted as $\mathcal{J}^A_a = \{\mathcal{M}^A_{x|a}\}_{x=1}^m$ (for some set $\{1, ..., m\}$ of possible outcome values). Quantum instruments being a set of CP maps have a corresponding Choi-Jamiołkowski representation [30], and a quantum instrument $\mathcal{J}^A_a = \{\mathcal{M}^A_{x|a}\}_{x=1}^m$ can be equivalently represented by the set of Choi-Jamiołkowski states $\{M^{A_I A_O}_{x|a} = \left[\mathcal{I} \otimes \mathcal{M}^A_{x|a}\left(|\mathbb{1}\rangle\langle\langle\mathbb{1}|\right)\right]^T\}_{x=1}^m$, where $|\mathbb{1}\rangle\rangle := \sum_j |j\rangle^{A_I}|j\rangle^{A_I}$ and T denotes matrix transposition with respect to the chosen orthonormal basis $|j\rangle^{A_I}$ of $\mathcal{H}_{A_I}$.

Global behaviour: process matrices. The probability $P(x_1, ..., x_N|a_1, ..., a_N)$ that the N agents $\{A^i\}_i$ observe the outcomes $(x_1, ..., x_N)$ for a choice of measurement settings $(a_1, ..., a_N)$ is a function of the corresponding local maps $\mathcal{M}^{A^1}_{x_1|a_1}, ..., \mathcal{M}^{A^N}_{x_N|a_N}$ and a global behaviour, called the process matrix. This can be expressed using the Choi-Jamiołkowski representation of the maps as follows [28, 30],

$$P(x_1, ..., x_N|a_1, ..., a_N) = P\left(\mathcal{M}^{A^1}_{x_1|a_1}, ..., \mathcal{M}^{A^N}_{x_N|a_N}\right) = \mathrm{tr}\left[\left(M^{A^1_I A^1_O}_{x_1|a_1} \otimes ... \otimes M^{A^N_I A^N_O}_{x_N|a_N}\right) W\right], \tag{4}$$

for a Hermitian operator $W \in A^1_I \otimes A^1_O \otimes ... \otimes A^N_I \otimes A^N_O$, known as the *process matrix*. The above equation plays the role of the Born rule in these general scenarios, with W playing the role of the quantum state (or a density matrix). The set of valid process matrices is characterised by the set of all such Hermitian operators W that yield positive normalised probabilities for all possible CP maps $\{\mathcal{M}^{A^k}_{x_k|a_k}\}_{k=1}^N$. This is required to hold also for CP maps that act on ancillary quantum systems (in addition to the in/output systems A^k_I and A^k_O associated with the process), where the ancillas between multiple labs may be entangled [28]. This imposes certain conditions on W, such as non-negativity [30] and implies that an N-partite process matrix $W = \mathcal{I} \otimes \hat{W}\left(|\mathbb{1}\rangle\rangle\langle\langle\mathbb{1}|\right)$ can be viewed as the Choi representation of a completely

[9] These Hilbert spaces are assumed to be finite dimensional.

positive and trace preserving map $\hat{W}$ from the input systems A_O^1,...,A_O^N (corresponding the outputs of the parties) to the output systems A_I^1,...,A_I^N (corresponding to the inputs of the parties), where $|\mathbb{1}\rangle\rangle$ corresponds to the unnormalised maximally entangled state over two copies of the input Hilbert space of $\hat{W}$. Note that the Choi representation of $\hat{W}$ and that of the local operations defined in the previous paragraphs differs by a transpose, this is a choice of convention made in the process matrix framework, that makes the notation and calculations more convenient.

4.2 Different classes of processes

Fixed order processes Process matrices in general need not be compatible with a global (acyclic) ordering between the operations of the parties. It is therefore useful to identify the subset of processes that are compatible with a definite acyclic causal order. Here we review a formal definition of such processes, that is adapted from [30, 31].

Definition 4.1 (Fixed order processes) *An N-partite process matrix W is said to be a fixed order process, if there exists a partial order $\mathcal{K}(\mathcal{S}_{I/O})$ on the set $\mathcal{S}_{I/O} = \{A_I^1, A_O^1, ..., A_I^N, A_O^N\}$ of the input and output systems of the N parties and associated with the binary relations $\prec_\mathcal{K}$ (first element precedes the second), $\succ_\mathcal{K}$ (first element succeeds the second) and $\not\prec\not\succ_\mathcal{K}$ (the elements are unordered) such that the following conditions are satisfied*

1. *For any $i \in \{1, ..., N\}$, $A_I^i \prec_\mathcal{K} A_O^i$.*
2. *For any party A^i and a subset $A^\mathcal{S}$ of the remaining parties, such that $A_O^i \not\prec_\mathcal{K} A_I^S$, $\forall A^S \in A^\mathcal{S}$ (which is denoted in short as $A_O^i \not\prec_\mathcal{K} A_I^\mathcal{S}$) with respect to the partial order $\mathcal{K}(\mathcal{S}_{I/O})$, the joint probability distribution $P(x_1, ..., x_N|a_1, ..., a_N)$ (cf. Equation (4)) obtained from W for any choice of local measurements of the N parties does not allow the outcome of any of the parties in $A^\mathcal{S}$ to depend on the setting of party A^i. That is, taking $x_\mathcal{S}$ to denote the set of outcomes $\{x_S\}_{S\in\mathcal{S}}$ of the parties in $A^\mathcal{S}$, we have the following whenever $A_O^i \not\prec_\mathcal{K} A_I^\mathcal{S}$ with respect to $\mathcal{K}(\mathcal{S}_{I/O})$*

$$\begin{aligned} P(x_\mathcal{S}|a_1, ...a_{i-1}, a_i, a_{i+1}, ..., a_N) =& \sum_{x_1,...,x_{j-1},x_{j+1},...,x_N} P(x_1, ..., x_N|a_1, ..., a_N) \\ =& P(x_\mathcal{S}|a_1, ..., a_{i-1}, a_{i+1}, ..., a_N). \end{aligned} \tag{5}$$

Not all processes are compatible with a fixed partial order (in the above sense), or a probabilistic mixture thereof. This incompatibility can be witnessed at a device-dependent manner (at the level of the process matrix) or a device-independent manner (at the level of the probabilities produced by a process matrix) leading to two distinct notions of causality in the process framework— causal separability and causal inequalities. Below, we review these concepts for the bipartite case with the parties A and B and refer the reader to [30, 31] for more general definitions.

Causal non-separability and causal inequalities: Consider a bipartite fixed causal order process compatible with the order $A \prec B$ (cf. Definition 4.1) and denote it as $W^{A\prec B}$. Then [31] shows that $W^{A\prec B}$ must be such that tracing out the output system of B leaves the process invariant i.e.,

$$W^{A\prec B} = W^{A_I A_O B_I} \otimes \mathbb{1}_{B_O},$$

where $W^{A_I A_O B_I} \geq 0$ (with $\operatorname{Tr} W^{A_I A_O B_I} = d_{A_O}$) is a valid process matrix for the case where Bob has a trivial output $d_{B_O} = 1$. Similarly, we can define $W^{B\prec A}$ to be a fixed order process that is compatible with the order $B \prec A$.

Definition 4.2 (Bipartite causally separable process [30]) *A bipartite process matrix W is said to be* causally separable *iff it decomposes as*

$$W = qW^{A\prec B} + (1-q)W^{B\prec A}, \tag{6}$$

for some $q \in [0, 1]$, where $W^{A\prec B}$ and $W^{B\prec A}$ are process matrices compatible with the fixed ordering between parties indicated in the respective superscripts. W is said to be causally non-separable *otherwise.*

Definition 4.3 (Bipartite causal process/distribution [28]) *A bipartite process matrix W is said to be* causal *iff for all choices of local operations, the joint probability $P(xy|ab)$ generated by W (for outcomes x and y and settings a and b of parties A and B respectively) decomposes as follows for some $q \in [0,1]$*

$$P(xy|ab) = qP^{A\prec B}(xy|ab) + (1-q)P^{B\prec A}(xy|ab), \tag{7}$$

where $P^{A\prec B}$ is a probability distribution compatible with the causal order $A \prec B$ by disallowing signalling from B to A i.e., $P^{A\prec B}(x|ab) = P^{A\prec B}(x|a)$, and similarly for $P^{B\prec A}$. The process W is called non-causal *otherwise. Similarly, distributions $P(xy|ab)$ are said to be causal/non-causal depending on whether they can be decomposed as above.*

[28] derives a linear inequality constraint on the joint probabilities $P(xy|ab)$ under four assumptions, this is referred to as a *causal inequality* and is shown to be a necessary constraint on causal distributions (analogous to Bell inequalities which are necessarily satisfied by local-causal distributions). The first three assumptions are the set-up assumptions LO, CL and FC of the process matrix framework. The fourth is an additional assumption referred to as *causal structure* (CS) which states that the input and output events of the parties are localised in a fixed partial order such as spacetime that prohibits signalling outside the future. Imposing these assumptions implies that the underlying process is compatible with a fixed order between the parties' operations, in the sense of Definition 4.1. This is because localisation of events in a partial order allows us to view the events as elements of the partial order, the LO assumption then implies that $A^i_I \prec A^i_O$ for each party A^i and the rest ensure that A^i can signal to A^j only if $A^i_O \prec A^j_I$ in the partial order. Hence, correlations produced by bipartite fixed order processes necessarily satisfy the causal inequality of [28]. Since this inequality is linear in the probabilities, probabilistic mixtures of such correlations also satisfy them, such mixtures correspond to causal distributions as we have seen in Definition 4.3.

Definition 4.3 formalises what it means for in/output events of parties to be localised in a partially ordered causal structure (Minkowski spacetime being an example). However, physical implementations of causally non-separable processes such as the quantum switch involve spacetime delocalised systems in Minkowski spacetime. We therefore require a way to formalise what it means for in/output events to be delocalised in a fixed, partially ordered spacetime structure. Such a formalisation is lacking in the previous literature, and in the following, we develop a framework that enables us to describe such scenarios by disentangling the operational aspects of the process framework (such as in/output events in the local labs) from the spacetime structure. We will revisit these assumptions in Section 9 where we discuss in detail the operational meaning of causal inequality violations, in light of our results.

Remark 4.4 *The set of causal processes is a strict superset of the set of causally separable processes, which is in turn a strict superset of the set of fixed order processes. For example, the quantum switch process (reviewed in Section 7) is causally non-separable but causal [30] and the classical switch, $W^{CS} := \frac{1}{2}W^{A\prec B} + \frac{1}{2}W^{B\prec A}$ is causally separable by definition but is not a fixed order process.*

5 Reformulating the process framework in terms of composition

5.1 The process map, extended local maps and their composition

The process matrix is typically viewed as a higher-order map that acts on the space of CPTP maps i.e., it maps the local operations of the parties into another channel (from some states in the global past to those in the global future of all parties [57]) or to the joint probabilities over the settings and outcomes of all the parties (See Figure 5). In this section, we show that this action of the process matrix on the local operations can be described through yet another view— as a composition of CPTP maps that involves feedback loops. We then derive the generalised Born rule (4) of the process matrix framework in this picture, derive reduced processes from partial composition of the process with local operations of a subset of parties, and prove the equivalence between two notions of signalling (one at the level of probabilities and the other at the level of the underlying map and quantum states).

As mentioned before, any process matrix W over N parties $\{A^1, ..., A^N\}$ can be seen as the Choi representation of a CPTP map $\hat{W}$ from the set of all output systems of the parties $\{A^1_O, ..., A^N_O\}$ to the set of all their input systems $\{A^1_I, ..., A^N_I\}$ [28, 30]. We will refer to $\hat{W}$ as the process map, in the rest

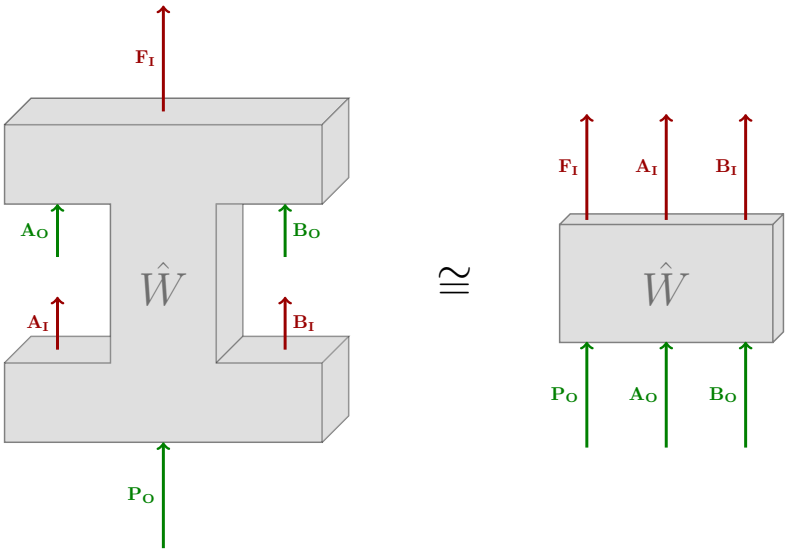

Figure 5: **A process as a higher order transformation and a CPTP map:** A process matrix W can be seen as the Choi representation of a CPTP map $\hat{W}$ from the output systems (green) of all parties to the input systems (red) of all parties (right side). The process shown here involves four parties A, B, P and F with P having a trivial input and F having a trivial output (hence not pictured) i.e., none of the parties can signal to P and this party can be seen as preparing states on the system P in the global past of the rest and similarly F can be seen as acting in the global future of the rest, possibly measuring the final states on F. $\hat{W}$ can also be seen as a higher order transformation maps the local transformations of the parties A and B (acting between the systems A_I and A_O, and B_I and B_O respectively) to a channel between the systems P and F. Given the preparation of P and measurement of F, it can also be seen as higher order transformation acting on the operations of all four parties, and mapping them to a joint probability distribution.

of the paper. As before, for each in/output system S of a CPTP map we will use the same label S to denote the state space of S i.e. the set of linear operators on the Hilbert space $\mathcal{H}_S$.

The local operation of each party A in the process framework is modelled as a map from their input system A_I to their output system A_O, possibly labelled by a classical setting a and outcome x. Here we extend the set of in and output systems of the local map to explicitly include the setting and outcome. We model the local operation $\mathcal{M}^{\mathcal{A}}$ of a party A as a CPTP map from the input systems A_I and A_s to the output systems A_O and A_o, where A_I and A_O are the quantum in/output systems we have seen before, A_s and A_o model the local in/output systems carrying the classical setting a (that specifies the choice of operation to be applied on the input system A_I) and a possible measurement outcome x of the party. In other words, $\mathcal{M}_a^A(\cdot) := \mathcal{M}^A(|a\rangle\langle a|_{A_s} \otimes \cdot)$ acts as a CPTP map from A_I to A_O and A_o for each choice of classical setting a.[10] If the map does not implement a measurement but implements a transformation from A_I to A_O depending on a setting choice a on A_s, the output state on A_o will correspond to the deterministic outcome represented by a fixed state $|\perp\rangle$, and the output system A_o can simply be ignored in this case. We will say that an outcome is non-trivial if it does not equal $\perp$. Then the action of each of the CP maps $\mathcal{M}_{x|a}^A$ (defining the quantum instrument $\mathcal{J}_a^A$ in the process framework) on an input state ρ_{A_I} can be described as

$$\mathcal{M}_{x|a}^A(\rho_{A_I}) := \mathrm{tr}_{A_o}\left[\left(|x\rangle\langle x|_{A_o} \otimes \mathcal{I}_{A_O}\right)\left[\mathcal{M}^{\mathcal{A}}(|a\rangle\langle a|_{A_s} \otimes \rho_{A_I})\right]\right]. \tag{8}$$

We will then refer to the map $\mathcal{M}^{A^k} : A_I^k \otimes A_s^k \mapsto A_O^k \otimes A_o^k$ for each party A^k as the *extended local map* or *extended local operation* of that party.

Remark 5.1 *Note that the set of all possible local operations $\{\mathcal{M}_a^A\}_a : A_I \mapsto A_O \otimes A_o$ for each party A is captured by a single extended map $\mathcal{M}^A : A_I \otimes A_s \mapsto A_O \otimes A_o$ by considering arbitrary classical initial states $|a\rangle\langle a|$ on A_s, through the correspondence $\mathcal{M}_a^A(\cdot) = \mathcal{M}^A(|a\rangle\langle a|_{A_s} \otimes \cdot)$.*

[10] The classical setting and outcome, a and x are encoded as quantum states $|a\rangle$ and $|x\rangle$ in the computational basis.

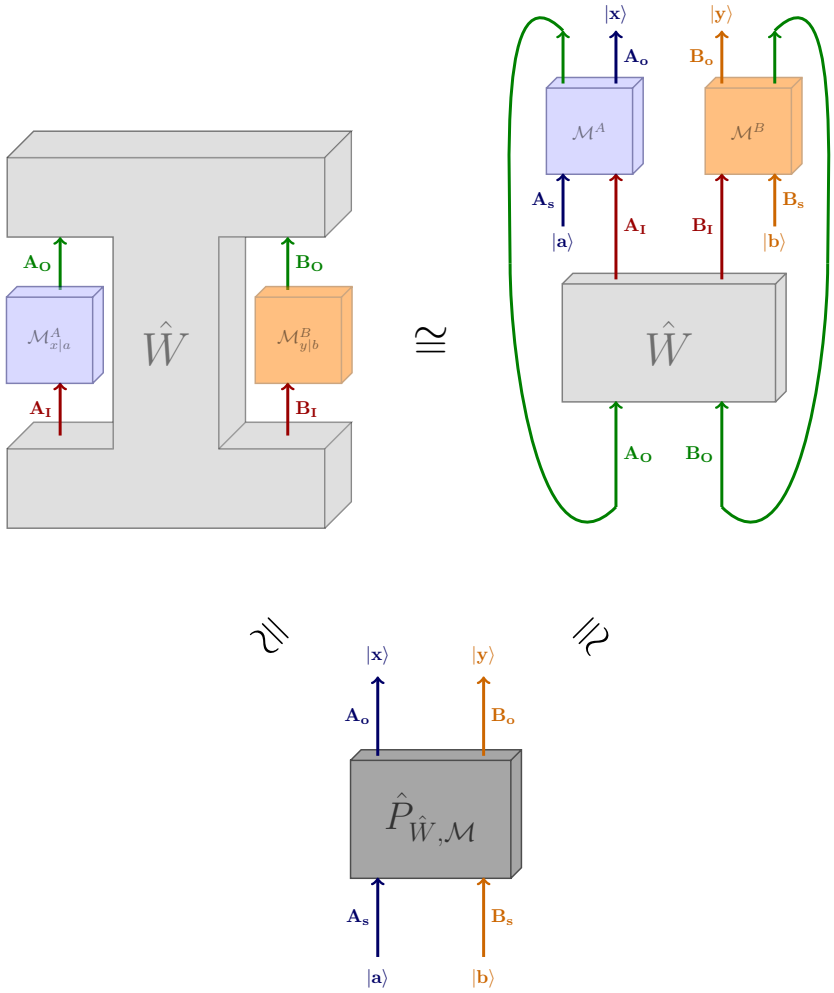

Figure 6: **Probabilities through composition of process map and local operations** In the process matrix framework, the joint probability of Alice and Bob obtaining measurement outcomes x and y upon measuring the settings a and b on their respective input systems A_I and B_I is associated with the CP (but trace decreasing) maps $\mathcal{M}^A_{x|a}$ and $\mathcal{M}^B_{y|b}$, and is given by Equation (4). This is illustrated in the top left. The probability rule of Equation (4) can also be obtained by viewing the action of the process map $\hat{W}$ on the local operations are a composition operation, as shown on the top right: where the three maps are composed by connecting the output systems A_I and B_I of $\hat{W}$ to the corresponding input systems of the local maps, and the output systems A_O and B_O of the local maps to the corresponding input systems of $\hat{W}$, through loop composition (Definition 2.1). The local maps of Alice and Bob are given additional input and output systems associated with their setting and outcome choices. This composition then leads to a new map $\hat{P}_{\hat{W},\mathcal{M}}$ (bottom middle) that encodes the joint probability distribution as shown in Lemma 5.2.

$\hat{W}$ acts on the local maps $\{\mathcal{M}^{A^k}\}_{k=1}^N$ through composition, as shown in Figure 6: first the local maps and the process map $\hat{W}$ are sequentially composed by connecting the output systems A_I^k of $\hat{W}$ to the input system of the corresponding local map $\mathcal{M}^{A^k}$, then the output system A_O^k of each local map is connected back to the corresponding input system of $\hat{W}$ through loop composition. Note that the systems being connected through loops will always have isomorphic state spaces of the same dimension, by construction. Explicitly, distinguishing the isomorphic state spaces associated with $\hat{W}$ and the local operations by adding a bar on top of the external in/output spaces of the local operations, the action of the process map $\hat{W} : \bigotimes_{k=1}^N A_O^k \mapsto \bigotimes_{k=1}^N A_I^k$ on a set of local operations $\{\mathcal{M}^{A^k}\}_{k=1}^N$, $\mathcal{M}^{A^k} : \bar{A}_I^k \otimes A_s^k \mapsto \bar{A}_O^k \otimes A_o^k$ corresponds to the following composition

$$\mathfrak{C}(\hat{W}, \{\mathcal{M}^{A^k}\}_{k=1}^N) := (\hat{W} \otimes \mathcal{M}^{A^1} \otimes ... \otimes \mathcal{M}^{A^N})^{(A_I^1 \hookrightarrow \bar{A}_I^1, ..., A_I^N \hookrightarrow \bar{A}_I^N, \bar{A}_O^1 \hookrightarrow A_O^1, ..., \bar{A}_O^N \hookrightarrow A_O^N)}.$$

We will refer to $\mathfrak{C}(\hat{W}, \{\mathcal{M}^{A^k}\}_{k=1}^N)$ as the *complete composition* of the process with the local operations. This corresponds to a new map, with the classical input systems $\{A_s^k\}_k$ carrying the measurement setting choices of all parties and classical output systems $\{A_o^k\}_k$ carrying the measurement outcomes of the parties. It is also useful to define the *partial composition* of $\hat{W}$ with the local maps $\{\mathcal{M}^{A^k}\}_{k=1}^l$ of the first $l < N$ parties, which is given as

$$\mathfrak{C}(\hat{W}, \{\mathcal{M}^{A^k}\}_{k=1}^l) := (\hat{W} \otimes \mathcal{M}^{A^1} \otimes ... \otimes \mathcal{M}^{A^N})^{(A_I^1 \hookrightarrow \bar{A}_I^1, ..., A_I^N \hookrightarrow \bar{A}_I^N, \bar{A}_O^1 \hookrightarrow A_O^1, ..., \bar{A}_O^l \hookrightarrow A_O^l)}.$$

Note that the partial composition $\mathfrak{C}(\hat{W}, \{\mathcal{M}^{A^k}\}_{k=1}^l)$ is a map with the input systems $\{A_O^k\}_{k=l+1}^N$ along with $\{A_s^k\}_{k=1}^N$ and output systems $\{A_I^k\}_{k=l+1}^N$ and $\{A_o^k\}_{k=1}^N$. In particular, we denote the complete composition $\mathfrak{C}(\hat{W}, \{\mathcal{M}^{A^k}\}_{k=1}^N)$ in short as $\hat{P}_{\hat{W},\mathcal{M}}$ since this map encodes the probabilities of local measurements– if we input a choice $\{a_k\}_k$ for settings to $\hat{P}_{\hat{W},\mathcal{M}}$ and post-select on a set $\{x_k\}_k$ of outcomes on the output of $\hat{P}_{\hat{W},\mathcal{M}}$, we get the joint probability of obtaining those outcomes given those setting choices as the success probability of that post-selection. Explicitly, Given an N-partite process map $\hat{W}$ and a set $\{\mathcal{M}_{A_k}\}_{k=1}^N$ of extended local maps, the joint probability of obtaining a set $\{x_k\}_k$ of outcomes given a choice $\{a_k\}_k$ of the settings is obtained from the complete composition $\hat{P}_{\hat{W},\mathcal{M}}$ as follows

$$P(x_1, ..., x_N | a_1, ..., x_N) = \frac{\text{tr}\left[\Pi_{x_1} \otimes ... \otimes \Pi_{x_N} \left(\hat{P}_{\hat{W},\mathcal{M}}(|a_1\rangle\langle a_1| \otimes ... \otimes |a_N\rangle\langle a_N|)\right)\right]}{\sum_{x_1,...,x_N} \text{tr}\left[\Pi_{x_1} \otimes ... \otimes \Pi_{x_N} \left(\hat{P}_{\hat{W},\mathcal{M}}(|a_1\rangle\langle a_1| \otimes ... \otimes |a_N\rangle\langle a_N|)\right)\right]}, \quad (9)$$

where the projector $\Pi_{x_k} = |x_k\rangle\langle x_k|$ projects the state on the system A_o^k to $|x_k\rangle$. Notice that in the general case, the denominator is needed to ensure that we get normalised probabilities, since a map such as $\hat{P}_{\hat{W},\mathcal{M}}$ formed by loop composition of CPTP maps could in general be trace decreasing (cf. Remark 2.2) and considering the numerator alone may not result in a valid normalised distribution. In the special case where the denominator of the above expression is unity, we would have the following, which we explicitly state below, as we show in the next section that this expression is equivalent to the process matrix probability rule (4).

$$P(x_1, ..., x_N | a_1, ..., x_N) = \text{tr}\left[\Pi_{x_1} \otimes ... \otimes \Pi_{x_N} \left(\hat{P}_{\hat{W},\mathcal{M}}(|a_1\rangle\langle a_1| \otimes ... \otimes |a_N\rangle\langle a_N|)\right)\right]. \quad (10)$$

5.2 Probabilities and reduced processes

The following lemma illustrates that this formulation of process matrices in terms of composition recovers the probability rule of the process matrix framework, and by construction, process matrices lead to valid probabilities under this probability rule.

Lemma 5.2 *[Probabilities from composition] For every process map $\hat{W}$, the joint probabilities obtained through the complete composition $\hat{P}_{\hat{W},\mathcal{M}}$ as in Equation (10) are equivalent to those obtained in the process matrix framework through Equation (4).*

In the following lemma, we relate the partial composition with the notion of a reduced process [30]. Given an N-party process matrix W and a CPTP map $\mathcal{M}_{a_j}^{A^j}$ for the jth party with Choi representation $M_{a_j}^{A_I^j A_O^j}$, the reduced process matrix [30] for the remaining $N-1$ parties is given as

$$\bar{W}(M_{a_j}^{A_I^j A_O^j}) := \text{Tr}_{A_I^j A_O^j}\left[\left(\mathbb{1}^{A_I^1 A_O^1} \otimes ... \otimes M_{a_j}^{A_I^j A_O^j} \otimes ... \otimes \mathbb{1}^{A_I^N A_O^N}\right).W\right], \quad (11)$$

Lemma 5.3 *[Partial composition and reduced process] Consider an N-partite process map $\hat{W}$ and the local operations $\{\mathcal{M}_{a_k}^{A^k}\}_{k=1}^l$ for the first $l < N$ parties for a fixed set of settings $\{a_k\}_{k=1}^l$. Then the partial composition $\mathfrak{C}(\hat{W}, \{\mathcal{M}_{a_k}^{A^k}\}_{k=1}^l)$ corresponds to a CPTP map whose Choi representation is the reduced process matrix $\bar{W}(M_{a_1}^{A_I^1 A_O^1}, ..., M_{a_l}^{A_I^l A_O^l})$.*

5.3 Equivalence of device dependent and independent notions of signalling

In the present paper, we focus on CPTP maps that correspond to process maps $\hat{W}$ and local operations and on the network $\hat{P}_{\hat{W},\mathcal{M}}$ corresponding to the complete composition of a process map $\hat{W}$ with a set of local operations. The partial compositions of $\hat{W}$ with a subset of local operations will be useful for defining signalling relations between different parties in such a network. For example, in a tripartite fixed order process with the parties A, B and C and causal order $A \prec B \prec C$, there would be no signalling from A_O to C_I in $\hat{W}$ but there would be signalling from A_O to C_I in the partial composition $\mathfrak{C}(\hat{W}, \mathcal{M}^B)$ of $\hat{W}$ with some local operation $\mathcal{M}^B$ of B, indicating that the party C indeed acts after the party A in the network. More generally, when we want to check for signalling from a party A^i to another party A^j in the network $\hat{P}_{\hat{W},\mathcal{M}}$, we can check whether $A^i_O \to A^j_I$ in the partial composition of $\hat{W}$ with the local maps of the remaining $N-2$ parties.

The following theorem shows that signalling relations between in/output quantum systems in partial compositions are in fact equivalent to signalling between classical settings and outcomes at the level of the joint probabilities that are obtained from the complete composition (cf. Lemma 5.2), which shows that the signalling relations in the different partial compositions indeed fully capture all the ways in which the parties can signal in such a network. This establishes a tight connection between a device-dependent (at the level of the underlying quantum maps and states) and device-independent (at the level of the observed probabilities) notions of signalling. The key to this equivalence lies in encoding all possible choices of the local operations for a party within the setting of that party, by extending the local operations (see also Remark 5.1).

Consider an N-partite process map $\hat{W}$ involving the parties $\mathcal{A} := \{A^1, ..., A^N\}$ and the extended local maps $\{\mathcal{M}^{A^k}\}_{k=1}^N$. Let $A^i \in \mathcal{A}$ be a party, $A^{\mathcal{S}} \subseteq \mathcal{A}\backslash A^i$ be a subset of the remaining parties, and $\mathfrak{C}(\hat{W}, \{\mathcal{M}^{A^k}\}_{k=1,k\notin i\cup\mathcal{S}}^{k=N})$ denote the partial composition of $\hat{W}$ with the local maps of parties in $\mathcal{A}\backslash(A^i \cup A^{\mathcal{S}})$, which, in the case where $A^i \cup A^{\mathcal{S}} = \mathcal{A}$ is defined to be $\mathfrak{C}(\hat{W}, \{\mathcal{M}^{A^k}\}_{k=1,k\notin i\cup\mathcal{S}}^{k=N}) = \hat{W}$. Then, taking $A^{\mathcal{S}}_I$ to denote the set of all (quantum) input systems of parties in $A^{\mathcal{S}}$, we have the following theorem.

Theorem 5.4 *[Equivalence of two notions of signalling] A^i_O does not signal to $A^{\mathcal{S}}_I$ in $\mathfrak{C}(\hat{W}, \{\mathcal{M}^{A^k}\}_{k=1,k\notin i\cup\mathcal{S}}^{k=N})$ if and only if the set of outcomes $x_{\mathcal{S}} := \{x_S\}_{S\in\mathcal{S}}$ of the parties in A^S do not depend on the setting a_i of the party A^i i.e., the corresponding joint probability distribution satisfies Equation (5).*

Remark 5.5 (Non-unitary processes) *One might wonder why we need to consider signalling relations in the partial composition instead of considering a directed graph of signalling relations in $\hat{W}$. For instance, if we say that party A^i signals to party A^j whenever $A^i_O \to A^j_I$ in $\hat{W}$, then doesn't the absence of a directed path of signalling relations from A^i to A^k imply that A^i's setting cannot be correlated with A^k's outcome in any network that the agents implement using $\hat{W}$? While this might seem intuitive, this need not be true in non-unitary processes where it is possible to have signalling relations from one party A to a set of parties $\{B, C\}$ without any signalling relation from A to B or from A to C. Then A could signal to D through the set $\{B, C\}$. The partial composition already takes into account these possibilities, and therefore our results apply to all processes, not just unitary ones.*

6 Characterising physical implementations of process matrices

6.1 No-go results

Using the general framework developed here along with the above results connecting this framework to process matrices, we establish a number of related no-go results for fixed spacetime implementations of process matrices. In Theorem 3.6, we have already seen that any signalling structure can be embedded in a spacetime consistently with relativistic causality. This means that in order to make any non-trivial statements, we must impose some constraints on the spacetime embeddings. Our no-go results below reveal the constraints on the embeddings that are necessary for implementing non-fixed order processes consistently with relativistic causality in a spacetime. These results also show the necessity of spacetime or time localisation of the in/output events of parties in physical implementations of non-fixed order

processes. Before stating the theorem and corollaries, we explicitly define the condition on the spacetime embedding that will appear there.

Definition 6.1 (Cycle-free spacetime regions) *We say that a region causal structure $\mathcal{G}^R_{\mathcal{T}}$ of a spacetime $\mathcal{T}$ is* cycle-free *if there is no sequence $\mathcal{P}^{S_1}, \mathcal{P}^{S_2}, ..., \mathcal{P}^{S_n}$ of nodes in $\mathcal{G}^R_{\mathcal{T}}$ such that $\mathcal{P}^{S_1} \xrightarrow{R} \mathcal{P}^{S_2} \xrightarrow{R} ... \xrightarrow{R} \mathcal{P}^{S_n} \xrightarrow{R} \mathcal{P}^{S_1}$, where $\xrightarrow{R}$ is the order relation of Definition 3.2.*

Let $\hat{P}_{\hat{W},\mathcal{M}}$ be a network formed by the composition of an N-partite process map $\hat{W}$ with the extended local maps $\{\mathcal{M}^{A^k}\}_{k=1}^N$, we then have the following no-go theorem for such a process network.

Theorem 6.2 *[No-go theorem for physical implementations of processes] No fixed spacetime implementation (Definition 3.4) $\hat{P}^{\mathcal{T},\mathcal{E}}_{\hat{W},\mathcal{M}}$ of the process network $\hat{P}_{\hat{W},\mathcal{M}}$ within a fixed spacetime structure $\mathcal{T}$ (Definition 3.1) can simultaneously satisfy the following three assumptions.*

1. *W is not a fixed order process (Definition 4.1).*
2. *$\hat{P}^{\mathcal{T},\mathcal{E}}_{\hat{W},\mathcal{M}}$ satisfies the relativistic causality condition of Definition 3.5.*
3. *The region causal structure given by the embedding $\mathcal{E}$ with $\mathrm{Nodes}(\mathcal{G}^R_{\mathcal{T}}) := \{\mathcal{E}(S)\}_{S\in\mathcal{S}}$ is cycle-free.*

Relativistic causality and free choice of local operations Note that the network $\hat{P}_{\hat{W},\mathcal{M}}$ in the theorem allows for arbitrary choices of local operations to be implemented, as it allows for arbitrary choices of settings on the input systems $\{A^k_s\}_{k=1}^N$ (cf. Remark 5.1) of the local maps $\{\mathcal{M}^{A^k}\}_k$ being composed with $\hat{W}$. Hence assumption 2 captures the condition that relativistic causality is not violated in the spacetime implementation of the process $\hat{W}$ with respect to the given spacetime embedding $\mathcal{E}$, irrespective of the choice of local operations that it is composed with. If we required no violation of relativistic causality to hold only for a particular set of local operations, then one can trivially implement any process $\hat{W}$ by restricting to local operations of the form $\mathcal{M}^{A^i} = \rho_{A^i_O} \circ \mathrm{Tr}_{A^i_I}$ i.e., those that discard the input system and independently reprepare an output. This allows parties to send out a system from their lab before they receive a system into their lab, without violating relativistic causality. A spacetime embedding that would enable such an implementation would be one where all the outputs $A^1_O, ..., A^N_O$ are associated with a spacetime location P and all the inputs $A^1_I, ..., A^N_I$ with a spacetime location $Q \succ P$, then $\hat{W}$ is quantum CPTP map that evolves spacetime localised quantum states at P to spacetime localised quantum states at $Q \succ P$ and can always be implemented.

Non-fixed order processes must have a cyclic causal structure A useful corollary that follows from Theorem 5.4 and the proof of the above theorem is given below. It characterises the causal structure of the process network $\hat{P}_{\hat{W},\mathcal{M}}$ associated with any process $\hat{W}$, and is irrespective of any spacetime embedding. It follows because the proof of the above theorem only uses the fact that $\{\mathcal{E}(S)\}_{S\in\mathcal{S}}$ together with the order relation $\xrightarrow{R}$ defines a directed graph (i.e., a causal structure according to Definition 2.7) and the third assumption implies that this is a directed acyclic graph.

Corollary 6.3 *Let $\hat{P}_{\hat{W},\mathcal{M}}$ be a network as defined in Theorem 6.2. The signalling relations of this network are compatible with an acyclic causal structure (Definition 2.7) if and only if $\hat{W}$ is a fixed order process. Moreover, if $\hat{W}$ is not a fixed order process, there exists a cyclic causal structure that the signalling relations of $\hat{P}_{\hat{W},\mathcal{M}}$ are compatible with.*

The above corollary establishes a tight connection between non-fixed order processes and cyclicity of signalling relations, and can be seen as a generalisation of a result from [51] which shows that for all unitary processes, causal non-separability and cyclicity of the causal structure are equivalent notions. This is established within a full causal modelling framework, which describes the causal structure of processes in much more detail but is difficult to formalise for non-unitary processes (see their paper for discussions on this point). In contrast, we have here characterised causal structures under very minimal assumptions which has allowed us to show that more generally it is the non-fixed orderedness of a process that is equivalent to cyclicity of the causal structure, causal non-separability is sufficient but

not necessary in this case. For instance, process matrices that can be expressed as a non-deterministic probabilistic mixture of fixed order processes (such as the classical switch) are causally separable by definition but also have a cyclic causal structure.

Fine-graining of spacetime regions It follows from Lemma 3.9 that any set of spacetime regions can be fine-grained such that the order relation $\xrightarrow{R}$ acts as a partial order over the fine-grained set of regions. The most fine-grained description is in terms of individual spacetime locations, in which case $\xrightarrow{R}$ reduces to the spacetime partial order $\prec$, and we immediately have Corollary 6.4 for this extreme case. But often, we do not need to fine-grain all the way to single spacetime locations to reduce $\xrightarrow{R}$ to a partial order and the above theorem is therefore more general, and still allows for systems to be delocalised in spacetime.

Corollary 6.4 (Spacetime localisation) *For every process $\hat{W}$ that is not a fixed order process, it is impossible to implement the corresponding network $\hat{P}_{\hat{W},\mathcal{M}}$ in a fixed spacetime without violating relativistic causality through an embedding that localises all the in/output systems in the spacetime.*

Note that all the above statements hold irrespective of the choice of reference frame used to describe the spacetime, since they only depend on the order relation between spacetime points which is an agent/frame independent notion according to Definition 3.1. We can also obtain a frame-dependent statement by considering the spatial and temporal coordinates of all the spacetime locations from the perspective of a single agent. For this, we must first add some more structure to our definition of spacetime to include details about spacetime co-ordinates.

Definition 6.5 (spacetime co-ordinates and time localisation) *Let $\mathcal{T}$ be a spacetime structure according to Definition 3.1. Each agent A can express every point $P \in \mathcal{T}$ in terms of a spatial co-ordinate $\mathbf{r}_P^A \in \mathbb{R}^n$ and temporal co-ordinate $t_P^A \in \mathbb{R}$ as $(\mathbf{r}_P^A, t_P^A)$. We require that agents always agree on the order relations even if they disagree on the co-ordinate assignment. Explicitly, the order relation $\preceq$ must be such that whenever $P \prec Q$ for some $P, Q \in \mathcal{T}$, $t_P^A < t_Q^A$ for all agents A. We say that a set of spacetime points $\mathcal{P} \subseteq \mathcal{T}$ are time localised with respect to an agent A if $t_P^A = t_{P'}^A$ for any $P, P' \in \mathcal{P}$.*

We then have the following corollary.

Corollary 6.6 *[Time localisation in a global frame] Under the same notation as Theorem 6.2, no fixed spacetime implementation $\hat{P}_{\hat{W},\mathcal{M}}^{\mathcal{T},\mathcal{E}}$ of the network $\hat{P}_{\hat{W},\mathcal{M}}$ within a fixed spacetime structure $\mathcal{T}$ can simultaneously satisfy the following three assumptions,*

1. *W is not a fixed order process (Definition 4.1).*
2. *$\hat{P}_{\hat{W},\mathcal{M}}^{\mathcal{T},\mathcal{E}}$ satisfies the relativistic causality condition of Definition 3.5.*
3. *The spacetime embedding $\mathcal{E}$ has the property that each of the spacetime regions $\mathcal{P}^S \subseteq \mathcal{T}$ are time-localised from the perspective of some agent A.*

We note that the above theorem no longer holds if we replace the requirement that all in/output systems are time-localised in a single global frame with the weaker requirement that the in/output systems A_I^k, A_O^k of each agent A^k are time-localised in their own frame. This is because spacetime localisation is an absolute (i.e., agent independent) notion in a fixed spacetime, but time localisation is an agent-dependent notion even in a fixed spacetime. For instance, consider two space-tike separated locations $P_1 \not\prec\not\succ P_2$ in Minkowski spacetime which are simultaneous (i.e, have the same time co-ordinate) in one frame (say that if agent A) but not in another (that of agent B). Then a quantum state arriving at a superposition of P_1 ad P_2 would be time-localised (but spatially delocalised) from A's perspective but time delocalised (and also spatially delocalised) from B's perspective. Using this observation, in Section 7.4 we propose an explicit protocol for realising the quantum switch (a causally non-separable process) using two agents in relative motion in a fixed spacetime where the agents perceive their respective in/output systems to be time-localised in their own frame.

Remark 6.7 (Cyclic signalling relations, correlations and composability) *The present work focusses on the signalling relations in a CPTP map, which indicate the presence of absence of signalling but we have not considered the nature or strength of the correlations associated with those signalling relations here. It can be the case that once we also take into account the correlations, then certain scenarios can no longer be implemented in accordance with relativistic causality even though their signalling relations alone may not indicate this. Indeed, we can have cyclic signalling relations arising both from causally separable processes (such as the classical switch), causally non-separable but causal processes (such as the quantum switch) as well as non-causal processes (such as the Lugano process [32]), but these processes lead to very different correlations. In the present work, we have characterised how the signalling relations of the process behave when the process is implemented in a spacetime, to characterise how the correlations should behave under spacetime embeddings, further work is required in formalising the assumptions required to rule out trivial causal inequality violations in situations with spacetime delocalised systems and addressing issues regarding the composability of physical systems. It has been shown [58] that process matrices are not closed under composition, for instance the parallel composition of two bipartite, fixed order processes is no longer a valid bipartite process as the composition can lead to a paradoxical causal loop. On the other hand, composability is fundamental to how we understand physics, we would expect to be able to compose physically implementable processes in an arbitrary manner and use them as sub-routines in other physical protocols. How is this familiar notion of composability restored in physical implementations of process matrices? These questions are a subject of upcoming works which are based on the current work and results presented in the extended abstract/talk [59], which suggests that once these assumptions are formalised, non-trivial causal inequality violations might be ruled out by relativistic causality in a fixed spacetime. Composability is also recovered by noting that paradoxical causal loops that assign contradicting values to a single outcome simply translate into agents acting at multiple times to produce two different outcomes, in physical implementations. Modelling this requires an extension of process matrices and quantum circuits to allow for multiple messages or a superposition of different number of messages to be exchanged between agents such that the in and output spaces are Fock spaces and not just Hilbert spaces. The details will be formalised follow-up works.*

6.2 Unravelling indefinite order processes into fixed order processes

Theorem 6.2 characterises the necessary condition on the spacetime regions required for implementing non-fixed order processes compatibly with relativistic causality in a fixed spacetime, which corresponds to the region causal structure being cycle-free. In the following, when we say *physical implementation*, we mean a fixed spacetime implementation that satisfies our relativistic causality condition of Definition 3.8. Now suppose that we have physically implemented a process protocol corresponding to a non-fixed order process in a region causal structure that is not cycle-free, i.e., we resolve Theorem 6.2 by giving up the third assumption. What can we say about such an implementation? In this section, we show that such an implementation of a non-fixed order process can always be unravelled into a physical implementation of a fixed order process (but over a larger number of parties) under fine-graining of the implementation. For this, the following property of spacetime regions will be useful.

Definition 6.8 (Pairwise correspondence of regions) *Given two spacetime regions $\mathcal{P}^1, \mathcal{P}^2 \subseteq \mathcal{T}$, we say that there exists a* pairwise correspondence *from $\mathcal{P}^1$ to $\mathcal{P}^2$ if there exists an invertible map $\mathcal{O} : \mathcal{P}^1 \mapsto \mathcal{P}^2$ such that for every $P^1 \in \mathcal{P}^1$, $\mathcal{O}(P^1) = P^2$ is such that $P^1 \prec P^2$.*

Now consider an N-partite process $\hat{W}$ and a fixed spacetime implementation $\hat{W}^{\mathcal{T},\mathcal{E}}$ of this process in a spacetime $\mathcal{T}$ with respect to an embedding $\mathcal{E}$. Suppose that each party A involved in this process is assigned the spacetime regions $\mathcal{P}^{A_I}$ and $\mathcal{P}^{A_O}$ under the embedding such that there is a pairwise correspondence $\mathcal{O}^A : \mathcal{P}^{A_I} \mapsto \mathcal{P}^{A_O}$ from the input to the output region. Consider the maximal fine-graining of the regions which induces the maximal fine-graining $\hat{W}^{\mathcal{T},\mathcal{E}}_{max}$ of $\hat{W}^{\mathcal{T},\mathcal{E}}$. Under the maximal fine-graining, each in/output system S of $\hat{W}$ transforms into a set of $|\mathcal{P}^S|$ systems, with one system S^{P^S} associated with each spacetime location $P^S \in \mathcal{P}^S$ (cf. Definition 3 and Section 3.1). That is, $\mathcal{F}^{sys} : S \mapsto \{S^{P^S}\}_{P^S \in \mathcal{P}^S}$ is a fine-graining of the in and output systems $S \in \{A^1_I, A^1_O, ..., A^N_I, A^N_O\}$ of $\hat{W}$. Denoting $\mathcal{F}^{sys}(S)$ explicitly as $S^{\mathcal{P}^S}$, this means that each $S^{P^S} \in S^{\mathcal{P}^S}$ is associated with a Hilbert space $\mathcal{H}_{S^{P^S}} \cong |\Omega\rangle \oplus \mathcal{H}_S$ where $\mathcal{H}_S$ is the Hilbert space of S in $\hat{W}$ and $|\Omega\rangle$ is the vacuum state. The joint state space of the fine-grained systems $S^{\mathcal{P}^S}$ is some subspace

$$\mathcal{H}_{S^{\mathcal{P}^S}} \subseteq \bigotimes_{S^{P^S} \in S^{\mathcal{P}^S}} \mathcal{H}_{S^{P^S}}.$$

We can now consider a local map $\mathcal{M}^A : A_I \otimes A_s \mapsto A_O \otimes A_o$ of an agent A, that can be composed with $\hat{W}$. This map can be physically implemented and fine-grained accordingly to $\hat{W}^{\mathcal{T},\mathcal{E}}_{max}$, by assigning the region $\mathcal{P}^{A_I}$ to the input systems A_I and A_s and the region $\mathcal{P}^{A_O}$ to the output systems A_O and A_o of this map. We can similarly define the maximal fine-graining $\mathcal{M}^{A,\mathcal{T},\mathcal{E}}_{max}$ of $\mathcal{M}^A$ under the same spacetime embedding and this can compose with $\hat{W}^{\mathcal{T},\mathcal{E}}_{max}$ by connecting each input elemental system $A_I^{P^{A_I}}$ of the two maps $\mathcal{M}^{A,\mathcal{T},\mathcal{E}}_{max}$ and $\hat{W}^{\mathcal{T},\mathcal{E}}_{max}$, and each elemental output system $A_O^{P^{A_O}}$ of the two maps. Notice that the pairwise correspondence from $\mathcal{P}^{A_I}$ to $\mathcal{P}^{A_O}$, ensures that the implementations of the local maps will also satisfy relativistic causality.

Then a natural way to describe the action of $\mathcal{M}^{A,\mathcal{T},\mathcal{E}}_{max}$ is to say that it independently applies the original map $\mathcal{M}^A$ between each pair of in and output locations $P^{A_I} \prec P^{A_O}$ where $P^{A_O} = \mathcal{O}^A(P_I)$ is given by the pairwise correspondence between regions. This is a special case of a more general map that acts independently between the pairs of locations, but applies a possibly different map at each location. That is a map of the form

$$\mathcal{M}^{A,\mathcal{T},\mathcal{E}}_{max} = \bigotimes_{P^{A_I} \in \mathcal{P}^{A_I}} \mathcal{M}^A_{P^{A_I}}, \tag{12}$$

where each $\mathcal{M}^A_{P^{A_I}} : A_I^{P^{A_I}} \otimes A_s^{P^{A_I}} \mapsto A_O^{P^{A_O}} \otimes A_o^{P^{A_O}}$ can in principle be any valid quantum CPTP map. In specific implementations, we may need to impose further constraints on these maps (e.g., for them to apply the original map $\hat{M}^A$ at each location, independent of the location at which it is applied), and we may also need to specify how these maps act on vacuum states as this is not specified in the process matrix framework. However, we already obtain the following general theorem under these minimal requirements. We explain the further constraints after the theorem, in relation to previous works.

Theorem 6.9 *[Unravelling physical process implementations into fixed order processes] Let $\hat{W}^{\mathcal{T},\mathcal{E}}_{max}$ be a maximally fine-grained physical implementation of an N-partite process $\hat{W}$ in a spacetime, where each pair of in and output regions $\mathcal{P}^{A_I^k}$ and $\mathcal{P}^{A_O^k}$ have a pairwise correspondence $\mathcal{O}^{A^k} : \mathcal{P}^{A_I^k} \mapsto \mathcal{P}^{A_O^k}$. Then $\hat{W}^{\mathcal{T},\mathcal{E}}_{max}$ acts as an $\tilde{N}$-partite fixed order process with $\tilde{N} = \sum_{k=1}^N |\mathcal{P}^{A_I^k}|$, upon composition with corresponding maximally fine-grained local maps $\{\mathcal{M}^{A^k,\mathcal{T},\mathcal{E}}_{max}\}_{k=1}^N$, where each $\mathcal{M}^{A^k,\mathcal{T},\mathcal{E}}_{max}$ acts independently between the pairs of points $P^{A_I^k} \in \mathcal{P}^{A_I^k}$ and $\mathcal{O}^{A^k}(P^{A_I^k}) \in \mathcal{P}^{A_O^k}$, as described by Equation (12).*

The above theorem shows that if a process network can be physically implemented in a fixed spacetime, then it can always be fined grained to a fixed order process with a larger number of parties. In other words, any process implementation satisfying assumptions 1 and 2 but not 3 of our no-go theorem 6.2 can be fine-grained into an process implementation satisfying assumptions 2 and 3 but not 1, but with a larger number of parties.

This generalises a result of [50] where they show that for the particular case of the quantum switch process which is a causally non-separable process (which we discuss in the next section), certain types of physical implementation of the quantum switch process can be described by a fixed order process matrix over a larger number of parties, with a further assumption that parties act trivially on vacuum states. In our framework, this assumption corresponds to requiring that each map $\mathcal{M}^A_{P^{A_I}}$ maps a vacuum state $|\Omega\rangle$ at the input $A_I^{P^{A_I}}$, for any setting a on $A_s^{P^{A_I}}$ to a vacuum state on the output $A_O^{P^{A_O}}$, while leaving the outcome variable in $A_o^{P^{A_O}}$ to denote the rest state $|\perp\rangle$ of the measurement device (denoting the absence of a non-trivial measurement outcome),

$$\mathcal{M}^A_{P^{A_I}}(|\Omega\rangle^{A_I^{P^{A_I}}} \otimes |a\rangle^{A_s^{P^{A_I}}}) = |\Omega\rangle^{A_O^{P^{A_O}}} \otimes |\perp\rangle^{A_o^{P^{A_O}}}, \quad \forall a \tag{13}$$

This condition ensures that any non-vacuum output sent by a party A at P^{A_O} must necessarily be preceeded by a non-vacuum input at the corresponding $P^{A_I} \prec P^{A_O}$ i.e., a party must receive a non-vacuum input before sending out a non-vacuum output. This is the local order condition that is implicit in the process matrix framework. In the absence of this condition, we can easily construct classical protocols with a definite time order, that cannot be regarded as fixed order processes and which

trivially generate correlations that violate causal inequalities (as pointed out in [59]). Note that further assumptions are also required, for instance that each local map is only used once. Multiple rounds of communication can also be used to trivially violate a causal inequality. This assumption can be imposed in our framework as described in Remark 3.10.

Our framework is not restricted to process matrices and can therefore model more general physical scenarios where the implicit assumptions of the process framework are not satisfied, therefore our main theorem above does not make these additional assumptions. In a follow up work based on [59], we further characterise the correlations realisable in physical process implementations, for which it is necessary to impose these constraints to make any interesting statements about non-causal correlations.

In the case of general networks of maps, not necessarily associated with process matrices, we have already shown in Lemma 3.9 that any signalling structure arising in such a network, once physically implemented in a spacetime can be fine-grained to a set of acyclic signalling relations. These results have consequences for several table-top experiments in Minkowski spacetime that claim to physically implement an indefinite causal order process, the quantum switch. We discuss our results for the quantum switch in the following section.

Physical meaning of the fine-grained maps Our results show that ultimately, any physical implementation of non-fixed order processes in a definite spacetime must necessarily involve quantum systems taking a superposition of different trajectories through a fixed spacetime. Physically, we always have the potential to intervene at any location in spacetime to check for the presence of a quantum system there. This potential is captured by the maximally fine-grained process and local maps, since we now have an "agent" associated with each possible spacetime location in the implementation, who can choose to perform any measurement or operation at that location. Physical processes are such that even under arbitrary interventions that could potentially be performed at any of the spacetime locations over which our quantum systems are delocalised, we can still not signal outside the future in the spacetime. This means that we ultimately have a fixed order process over all these "agents". The maximally fine-grained process and local map implementations in our framework capture this idea. Note that in a physical experiment, a single experimentalist may play the role of the multiple parties/agents associated with the process map. This is indeed the case with physical implementations of processes through table-top experiments which are performed in a single lab, which we will discuss in the next section. The maximal fine-graining also establishes a one-to-one connection between operational and spacetime events even in experiments involving highly spacetime delocalised quantum system.

7 Causality in the quantum switch

We now apply our framework to the particular example of the quantum switch (QS), which is a particularly popular example of an indefinite causal order and has been repeatedly claimed to be experimentally implemented [40, 42, 60] in Minkowski spacetime. What does it physically mean to implement an indefinite causal order process (in particular, one that is causally non-separable) in a fixed spacetime, which itself implies a notion of a fixed acyclic causal structure? Here, we address and clarify this question by analysing information-theoretic and relativistic notions of causality in various fixed spacetime implementations of the quantum switch and show that all physical implementations of the quantum switch can indeed be explained within a definite acyclic operational causal structure that is compatible with the light-cone structure of Minkowski spacetime (as one would expect). The quantum switch corresponds to a unitary, causally non-separable process, and we can therefore make stronger statements for this case than Theorem 6.2 which applies to general processes.

7.1 The quantum switch as a process matrix

The quantum switch is originally defined as a supermap QS that maps a pair of quantum channels U^A and V^B to a new channel $W(U^A, V^B)$ that implements a coherently controlled superposition of the orders of U^A and V^B on a target system (details of the original definition can be found in Appendix B). In the process matrix framework, the quantum switch can be represented as a four party process matrix [57] between the labs A, B, C and D. Here, C prepares the control and target subsystems in her lab and outputs it to the process matrix W^{QS}, which acts as follows: if the control is in state $|0\rangle$, it sends

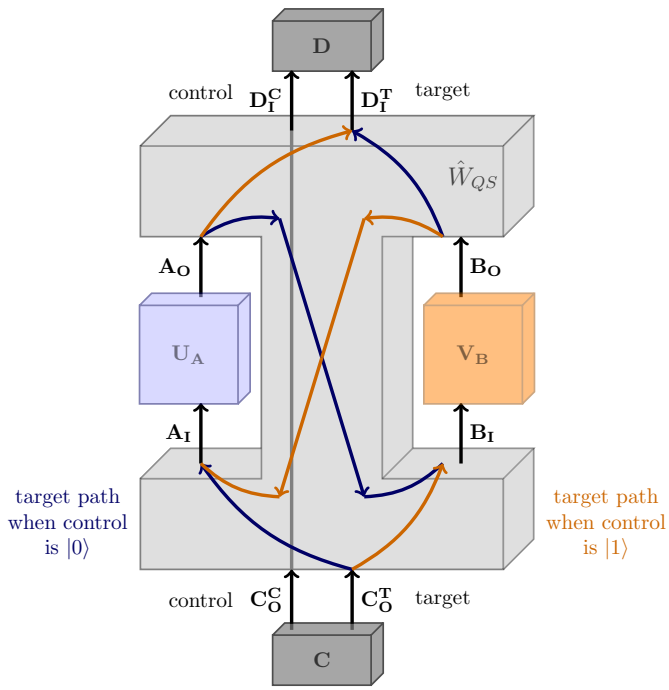

Figure 7: **4 party process matrix for the quantum switch:** A lab C in the past of all others (with trivial input space) prepares the control and target states and sends the target to A if the control is in state $|0\rangle$ and to B if the state is $|1\rangle$. After A and B have operated on the target in an order depending on the control state, a lab D in the future of all others (with trivial output space) receives the target from A or B and control directly from C (gray path). D therefore holds the final state of the control and target where the order of A's and B's operation on the target is entangled with the control state. The process matrix, W for the quantum switch in this case represents a controlled superposition of the orders $C \prec A \prec B \prec D$ (blue path) and $C \prec B \prec A \prec D$ (orange path) of operations on the target system.

only the target subsystem (prepared in a state $|\psi\rangle$) to A and then to B (after A's operation) and finally to D (after B's operation), and if the control is in state $|1\rangle$, it sends the target (prepared in a state $|\psi\rangle$) to B first, then to A (after B's operation) and finally to D. Further, W^{QS} also sends the control subsystem unchanged, directly from C to D. Note that C lies in the global past of all parties and cannot be signalled to by any of them while D lies in the global future of all parties and cannot signal to any of them; thus C has a trivial input space and D has a trivial output space. Further, C and D send or receive both the control and target systems, while A and B only receive, operate on and send out the target system. So it is convenient to decompose the output space of C and input space of D as $C_O = C_O^C \otimes C_O^T$ and $D_I = D_I^C \otimes D_I^T$ corresponding to the control and target systems. When the local operations of A and B are qubit channels, the dimensions of input and output systems of the local laboratories are $d_{A_I} = d_{A_O} = d_{B_I} = d_{B_O} = d_{C_O^T} = d_{D_I^T} = d_{C_O^C} = d_{D_I^C} = 2, d_{C_I} = d_{D_O} = 1$. The corresponding process matrix is pure i.e., is rank one and given as $W^{QS} = |W^{QS}\rangle\langle W^{QS}|$ where

$$|W^{QS}\rangle = |\mathbb{1}\rangle\!\rangle^{C_O^T A_I} |\mathbb{1}\rangle\!\rangle^{A_O B_I} |\mathbb{1}\rangle\!\rangle^{B_O D_I^T} |00\rangle^{C_O^C D_I^C} + |\mathbb{1}\rangle\!\rangle^{C_O^T B_I} |\mathbb{1}\rangle\!\rangle^{B_O A_I} |\mathbb{1}\rangle\!\rangle^{A_O D_I^T} |11\rangle^{C_O^C D_I^C} \tag{14}$$

The situation is illustrated in Figure 7. If the lab C prepares the suitable input state, labs A and B perform the respective operations U^A and V^B, the final state arriving at lab D is given as follows.

$$\begin{aligned} &\left((\alpha\langle 0| + \beta\langle 1|)^{C_O^C} \otimes \langle\psi^*|^{C_O^T} \otimes \langle\!\langle U^{A^*}|^{A_I A_O} \otimes \langle\!\langle V^{B^*}|^{B_I B_O}\right) \cdot |W^{QS}\rangle \\ =&\alpha|0\rangle^{D_I^C} \otimes (V^B U^A|\psi\rangle)^{D_I^T} + \beta|1\rangle^{D_I^C} \otimes (U^A V^B|\psi\rangle)^{D_I^T}, \end{aligned} \tag{15}$$

where $\langle\psi^*|$ denotes the complex conjugate of $\langle\psi| = |\psi\rangle^\dagger$ in the computational basis $\{|0\rangle, |1\rangle\}$, such that $\langle\psi^*|i\rangle = \langle i|\psi\rangle, i \in \{0,1\}$. $|U^{A^*}\rangle\!\rangle^{A_I A_O} = (\mathbb{I} \otimes U^{A^*})|\mathbb{1}\rangle\!\rangle^{A_I A_I}$ and similarly for V^B, where $*$ denotes the complex conjugate in the chosen orthonormal basis.

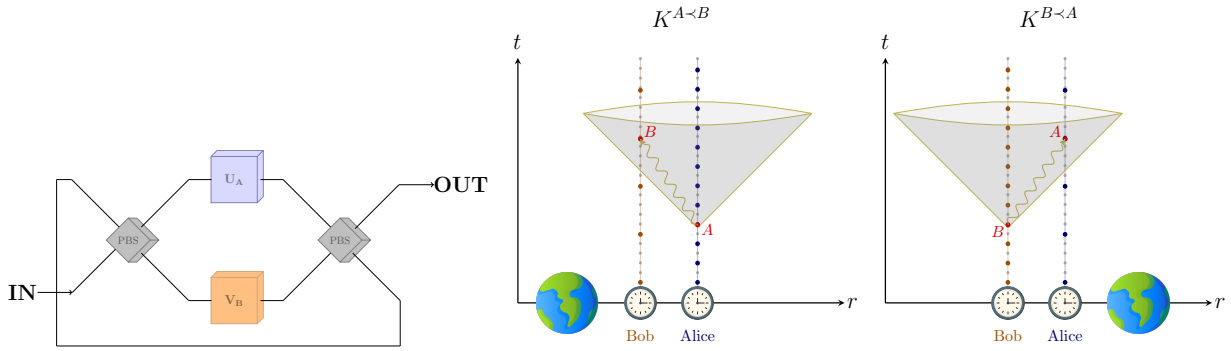

(a) Optical quantum switch (b) Gravitational quantum switch

Figure 8: **Definite and indefinite spacetime implementations of the quantum switch.** **(a)** This is a schematic of a linear optical experimental implementation of QS within a fixed spacetime structure. Here, the control qubit is encoded in the polarisation of a photon and the target qubit is encoded in a different degree of freedom of the same photon e.g., angular momentum modes. Then a horizontally polarized photon is transmitted by the all the polarizing beam splitters (PBSs) and takes the path where the unitary U^A is applied first and then U_B while a vertically polarized photon is reflected by all PBSs and follows the path where U_B is applied before U^A. The unitaries U^A and U_B act on the target degrees of freedom. This was proposed in [37] and implemented in [40, 60]. **(b)** This is a theoretical proposal [27] for implementing the QS transformation using a quantum superposition of gravitating masses, which results in a superposition of spacetime structures. Here, the two parties Alice and Bob are each in possession of their own clock C_A and C_B which are initially synchronised. A gravitating mass is prepared in a quantum superposition of macroscopically distinguishable (relative to the agents) spatial configurations depending on the state of a control qubit. If the control is in the state $|0\rangle$, the mass is placed closer to Bob such that the clock C_B ticks slower that C_A due to gravitational time dilation enabling Alice to send a physical system (the target) to Bob at a proper time $t_A = 3$, such that it is received by Bob at his proper time $t_B = 3$. This mass configuration is labelled as $\kappa_{A\prec B}$. If the control is in the state $|1\rangle$, the mass is placed closer to Alice, enabling Bob to send a physical system to Alice at $t_B = 3$ in his local reference frame, with the system being received by Alice at $t_A = 3$ in her local reference frame. This mass configuration is labelled as $\kappa_{B\prec A}$. That is, irrespective of the control, both parties receive the target system in their lab at the same local time, but due to the mass superposition, their clocks are experiencing a superposition of different time dilations and consequently ticking at a superposition of different rates. The co-ordinate axes are in the frame of a distant observer Charlie for whom the effect of the gravitational field is negligible. Notice that Charlie's time intervals (small gray dots) are unaffected by the mass configuration, and this observer would perceive a fixed spacetime.

The process matrix W^{QS} is known to be causally non-separable (i.e., cannot be decomposed as in Equation (6)) but nevertheless causal (i.e., always produces probabilities that decompose as per Equation (7)) [30].

Physical implementations of QS We discuss two implementations of the quantum switch transformation that indeed query each operation not more than once (on a non-vacuum state), the optical and the gravitational implementations. The former corresponds to experiments that claim to implement the supermap QS and these are performed through table-top optical interferometric setups within a fixed spacetime structure (which can be safely approximated to be a Minkowski spacetime). The latter corresponds to a theoretical proposal for implementing QS through a quantum superposition of gravitating masses [27], which would in turn result in a superposition of spacetime geometries i.e., it involves an indefinite spacetime structure. The main features and intuition behind these implementations are illustrated and explained in Figure 8.

7.2 No-go result for the quantum switch

We now derive a slightly stronger version of our general no-go result, theorem 6.2 for the particular case of the quantum switch, which corresponds to a causally non-separable process (and hence not a fixed order process) that nevertheless does not violate causal inequalities [37]. In the following, we will call a local operation $\mathcal{M}_{A^k}$ of some party A^k non-trivial if $A^k_I \to A^k_O$ in $\mathcal{M}_{A^k}$.

Lemma 7.1 *[No-go result for the quantum switch] Consider the process map $\hat{W}^{QS}$ whose Choi representation is the process matrix W^{QS} of the quantum switch. Let $\hat{P}_{QS,U,V}$ be the quantum switch network where W^{QS} acts on two non-trivial local operations $U^A : A_I \mapsto A_O$ and $V^B : B_I \mapsto B_O$ of Alice and Bob. Then any fixed spacetime implementation $\hat{P}^{\mathcal{T},\mathcal{E}}_{QS,U,V}$ of this network cannot simulataneously satisfy both of the following assumptions*

1. *$\hat{P}^{\mathcal{T},\mathcal{E}}_{QS,U,V}$ satisfies relativistic causality*
2. *The subgraph of the region causal structure given by the embedding $\mathcal{E}$ with $\mathrm{Nodes}(\mathcal{G}^R_{\mathcal{T}}) := \{\mathcal{E}(S)\}_{S\in\mathcal{S}}$, restricted to $S \in \{A_I, A_O, B_I, B_O\}$ is cycle-free.*

The above is a stronger statement than that of Theorem 6.2 applied to this process because the theorem deals with the extended local maps that include all possible choices of local operations for each party. On the other hand, the above statement follows for any fixed (but non-trivial) choice of local operations for A and B.[11]

7.3 Consequences for experimental implementations

Due to Lemma 7.1, we know that any physical implementation of the quantum switch protocol satisfying relativistic causality in Minkowski spacetime must be such that Alice and Bob's in/output systems are delocalised within large enough spacetime regions to enable bidirectional signalling between these regions. Several experiments [37, 40–43, 45–49] claim to physically implement the quantum switch process and therefore an indefinite causal structure in Minkowski spacetime. Our general results of Theorem 6.9 and Lemma 3.9 can be applied to provide further insights into the causal structure of such experiments. They tell us that these experiments, if we believe that they do not violate relativistic causality in the spacetime, can be ultimately described in terms of a fixed order process over a larger number of agents, and they necessarily admit an acyclic causal structure, even though they aim to implement a causally non-separable process matrix. This means that the causal structures of the process map of QS and that of a fixed spacetime implementation of this map compatible with relativistic causality are distinct, as shown in Figure 9. The following corollary formalises this.

While the rest of the paper focuses on the more operational notion of signalling, rather than causation (as motivated in Section 2), in the case of the quantum switch which corresponds to a unitary process, these two notions coincide [51]. We can therefore make statements about the causal structure of the quantum switch based on our framework and results. The equivalence between causation and signalling in this case is further explained in Appendix C. Applying Lemma 3.9, it follows that, by virtue of being a network of CPTP maps implemented in a fixed spacetime (which happens to be Minkowski spacetime in all these experiments) that does not violate relativistic causality, all these implementations must necessarily admit a signalling structure that is ultimately acyclic and compatible with the spacetime partial order. Since signalling and causation are equivalent in the case of the quantum switch, this implies that all these implementations admit an explanation in terms of a definite acyclic quantum causal structure as stated in the corollary below.

Corollary 7.2 *[Experimental implementations] Any implementation of the quantum switch process map $\hat{W}^{QS}$ in a fixed spacetime that does not violate relativistic causality in the spacetime admits an explanation in terms of a fixed order process matrix $\hat{W}^{\mathcal{T},\mathcal{E}}_{QS,max}$ that is associated with a definite acyclic causal structure whose edges flow from past to future in the spacetime in which it is implemented.*

A possible acyclic causal structure capable of explaining a fixed spacetime implementation of QS is illustrated in Figure 9, Section 7.4 and Appendix B describe an explicit protocol for for QS that yields this acyclic causal structure but nevertheless implements the QS transformation with only one query (on a non-vacuum state) to each local operation.

[11] This is because, for general processes, whether or not agent A^i can signal to an agent A^j can depend on the choice of local operations of A^j as well as those of the remaining agents. In the quantum switch, signalling from A to B can be verified by A by suitable local choices of operations alone, for any choice of operation for B (and symmetrically for B), whenever the control qubit is in a non-trivial superposition state $\alpha|0\rangle_C + \beta|1\rangle_C$ with $\alpha, \beta \neq 0$.

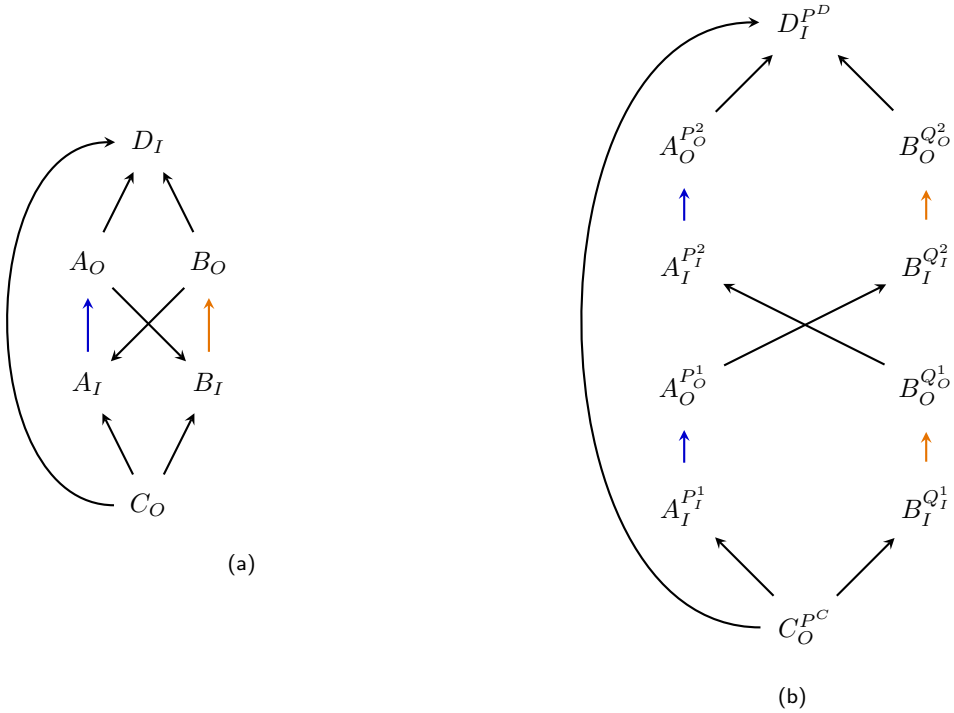

Figure 9: **Causal structure of the quantum switch protocol and that of its fixed spacetime implementation** Black arrows denote signalling relations coming from the process matrix W^{QS}, while blue and orange arrows are signalling relations coming from Alice and Bob's (non-trivial) local operations. (a) The signalling relations of the process W^{QS} acting on non-trivial local operations yields a cyclic causal structure shown here. Here, the input C_O and D_I denote the combined output $C_O^C \otimes C_O^T$ of C and input $D_I^C \otimes D_I^T$ of D respectively. Being a unitary process, causation and signalling are equivalent notions in this case and this gives the causal structure of the process [51]. (b) The causal structure of a fixed spacetime implementation of QS satisfying relativistic causality must necessarily involve spacetime delocalised systems (cf. Corollary 6.4). This causal structure would therefore involve more nodes than that of (a), where the nodes in this case correspond to elemental subsystems, each elemental subsystem S^P is a quantum system S (with a Hilbert space $\mathcal{H}_S = |\Omega\rangle \oplus \mathbb{C}^{d_S}$) associated with a fixed spacetime location $P \in \mathcal{T}$. This causal structure must be acyclic if the implementation does not violate relativistic causality in the spacetime, and we have that whenever there is an edge $\rightarrow$ from one node to other in this graph, the spacetime location of the first precedes the spacetime location of the second with respect to the order relation $\prec$ i.e., the edges of the information theoretic causal structure flow from past to future with respect to the spacetime causal structure. For an explicit description of an implementation of QS corresponding to this causal structure see Appendix C. The causal structure of (a) is essentially identical to the causal structure of W^{QS} obtained in the framework of [51] with the distinction that the in/output systems of each party are associated with a single node such that only the black arrows are relevant.

On the notion of agents and interventions Let us briefly comment upon the notion of an "agent" or "party" in the theoretical process description is comparison to the physical observer or experimentalist. All these experiments are performed on a table-top within a single laboratory, even though they aim to implement the quantum switch with is modelled as a 4 partite process matrix.[12] In particular, the main agents Alice and Bob who are part of the superposition in the theoretical description simply correspond to circuit elements that implement the operations U^A and V^B in the experiment. Furthermore, in the theoretical model, the process matrix is supposed to describe an outside environment that is inaccessible to the parties in the "local laboratories", while in the physical implementation, the whole process takes place on a table-top and the experimentalist can in principle intervene upon and control any part of the experiment (in fact they must do so in order to set up the very experiment). Therefore, we argue that the more appropriate theoretical model for such experiments is the fine-grained fixed order process over the larger number of parties as this captures the ability of the physical experimenter to control or intervene at any location within the experiment (even if they may choose not to do so in certain runs of

[12] Or a tripartite process if the initial preparation is encoded within the process.

experiment that they wish to report). However in this theoretical model of the experiment, the process being implemented is no longer causally non-separable, it is a fixed order process as we have shown in Theorem 6.9.

Remark 7.3 (Local distinguishability of the order) *The fixed spacetime implementation of QS described in Figure 9 as well that of [54] (explained in Appendix B) are such that the spacetime location at which Alice or Bob's operation is applied is perfectly correlated with the control: if the control is $|0\rangle$, the Alice's operation is applied (to a non-vacuum state) arriving at the location P_I^1 and Bob's at is applied at $Q_I^2 \succ P_I^1$, while if the control is $|1\rangle$, Bob's operation is applied on the (non-vacuum) state arriving at the location Q_I^1 and Alice's at $P_I^2 \succ Q_I^1$. In the experimental implementation of [42, 46], this is not the case and local measurements by Alice and Bob of the (space)time location at which they receive a non-vacuum state would not reveal significant information about the control since there is a large uncertainty in this spacetime location even when the operations are applied in a fixed order (i.e., where the control is in one of the computational basis states). Our results are general and apply to both these type of implementations. In the former case, the in/output systems of Alice and Bob each split into two elemental systems in the spacetime implementation while in the latter case, there can be many more elemental systems since the in/output systems can be delocalised over many more spacetime points. This means that we would need to consider a larger number of parties in the latter case in order to describe the experiment using a fixed order processes (cf. Theorem 6.9), which changes the number of nodes one would consider in the fine-grained causal structure, but does not alter the fact that the causal structure would still be a definite acyclic one once we finegrain and look in sufficient level of detail. Further, we also note that while the order of operations may not be locally indistinguishable in [42, 46], they are globally distinguishable. If Alice and Bob, in addition to applying their local operations choose to send a photon to a friend as soon as they apply these operations, then it is possible for the friend to distinguish the orders in which the operations were applied, by measuring the arrival times of the photons [50]. This protocol and its use in distinguishing between such physical optical implementations and theoretical, quantum gravitational implementations in discussed in Section 7.5.*

7.4 Minkowski quantum switch with time localised systems

Here, we propose a new quantum switch protocol in Minkowski spacetime with the property that the in and output events are time localised for both Alice and Bob in their respective frames. The protocol requires Alice and Bob to be in relative motion with respect to each other, at a constant velocity. Previous quantum switch protocols typically involve spatial localisation and time delocalisation, in contrast, ours will involve spatial delocalisation and time localisation (with respect to the local reference frames). It also demonstrates that Corollary 6.6 no longer holds when only requiring time localisation with respect to each local frame, instead of time localisation with respect to a single global reference frame. Consider the following protocol where the spacetime $\mathcal{T}$ is taken to be Minkowski spacetime and the partial order corresponds to the light cone structure. In this spacetime implementation of the process map $\hat{W}^{QS}$, we have the following embedding.

1. Charlie prepares the initial state of the control and target system $(\alpha|0\rangle_C + \beta|1\rangle_C) \otimes |\psi\rangle_T$ at a spacetime location P^C and sends the target to Alice at spacetime location $P_I^1 \succ P^C$ (and a vacuum state $|\Omega\rangle$ to Bob at spacetime location $Q_I^1 \succ P^C$) or the target state to Bob at spacetime location Q_I^1 (and the vacuum to Alice at P_I^1) depending coherently on the control being $|0\rangle$ or $|1\rangle$.

2. Alice and Bob apply their local operations U and V on the state arriving at P_I^1 and Q_I^1 respectively, mapping it to a state on their output wires at the spacetime location $P_O^1 \succ P_I^1$ and $Q_O^1 \succ Q_I^1$ respectively. We assume that U and V act trivially on the vacuum i.e., $U|\Omega, P_I^1\rangle = |\Omega, P_O^1\rangle$ and similarly for V.

3. The state on Alice's output at P_O^1 is forwarded to Bob's input at another spacetime location $Q_I^2 \succ P_O^1$ and the state on Bob's output is forwarded to Alice's input at some spacetime location $P_I^2 \succ Q_O^1$.

4. Alice and Bob again apply their local oerations U and V to the states incoming at P_I^2 and Q_I^2 mapping it to a corresponding state on their output systems at spacetime locations $P_O^2 \succ P_I^2$ and $Q_O^2 \succ Q_I^2$ respectively, again while acting trivially on the vacuum state.

5. Finally, depending on the control, either the state from Bob's output at Q_O^2 or the state from Alice's output at P_O^2 is forwarded to Danny's input at $P_F \succ P_O^2, Q_O^2$, along with the control.

That is, the above protocol corresponds to a Minkow spacetime implementation of the process map $\hat{W}^{QS}$ associated with the following spacetime embedding.

$$\begin{aligned}
\mathcal{E}(C_O^C) = \mathcal{E}(C_O^T) = \mathcal{P}^C &:= \{P^C\} \\
\mathcal{E}(A_I) = \mathcal{P}^{A_I} &:= \{P_I^1, P_I^2\}, \\
\mathcal{E}(A_O) = \mathcal{P}^{A_O} &:= \{P_O^1, P_O^2\}, \\
\mathcal{E}(B_I) = \mathcal{P}^{B_I} &:= \{Q_I^1, Q_I^2\}, \\
\mathcal{E}(B_O) = \mathcal{P}^{B_O} &:= \{Q_O^1, Q_O^2\}, \\
\mathcal{E}(D_I^C) = \mathcal{E}(D_I^T) = \mathcal{P}^D &:= \{P^D\}.
\end{aligned} \tag{16}$$

The explicit CPTP maps for this protocol are given in Appendix C. There, it is also shown that in this protocol each party acts exactly once on a non-vacuum state and that it admits the definite acyclic causal structure of Figure 9 (b) that respects relativistic causality in Minkowski spacetime.

Figure 10 shows that we can arrange the situation such that $P_I^1 = (\mathbf{r}_1^{A_I}, t^{A_I})$ and $P_I^2 = (\mathbf{r}_2^{A_I}, t^{A_I})$ have the same time co-ordinate t^{A_I} in Alice's frame, similarly P_O^1 and P_O^2 have the same time co-ordinate $t^{A_O} > t^{A_I}$ in Alice's frame, Q_I^1 and Q_I^2 have the time co-ordinate t^{B_I} while Q_O^1 and Q_O^2 have the time co-ordinate $t^{B_O} > t^{B_I}$ in Bob's frame, where Alice and Bob's frames are related by a Lorentz transformation. This establishes the claim.

7.5 Gravitational quantum switch vs fixed spacetime implementations

In this paper, we have focused our attention on process implementations in a fixed spacetime. As we have seen in Figure 8, the gravitational switch is a theoretical proposal for a quantum switch implementation in an "indefinite spacetime structure" achieved through a quantum superposition of gravitating masses. It is then natural to ask, what, if at all, are the operational differences between these two implementations of the same process? Are there some features of the gravitational switch that are impossible to achieve in a fixed spacetime?

Some properties that are impossible to achieve in fixed spacetime implementations At first sight, it might seem that the property of Alice and Bob receiving the target system at the same time (in their local reference frames) irrespective of the order in which they act is unique to the gravitational implementation. However, our protocol of Section 7.4 (in particular, Figure 10) illustrates that this property can also be achieved in fixed spacetime implementations with classical agents who are in relative motion with a constant relative velocity. Our results also reveal that a property that is impossible to achieve in fixed spacetime implementations is where Alice and Bob receive the target system at the same spacetime location irrespective of the order in which they act (cf. Corollary 6.4). While it is possible to have two types of QS implementations in Minkowski spacetime a) where the in/output events are spatially localised but time delocalised b) where they are temporally localised but spatially delocalised, a further observation that we can make is that it is impossible for a single implementation to have both these properties (depending on the choice of reference frame used to describe the spatio-temporal co-ordinates), if spacetime co-ordinates in different frames are related as described in Definition 6.5. For this, notice that whenever we have temporal localisation of A's input event A_I in some frame, it means that the spacetime region $\mathcal{P}^{A_I}$ associated with that event consists of spacetime locations that are pairwise space-like separated (as they have the same time co-ordinate but different spatial co-ordinates in some frame). Then there can exist no other frame in which all locations in $\mathcal{P}^{A_I}$ have the spatial co-ordinate and different time co-ordinates as this would allow us to order certain points in this region in the future of others, which violates the condition that the order relation between spacetime points must be frame independent in a fixed spacetime (cf. Definition 6.5). On the other hand, consider the gravitational implementation in indefinite spacetime [27]. Here, the in/output events of Alice are time localised in her own frame, as are Bob's events in his frame. However, when described in the reference

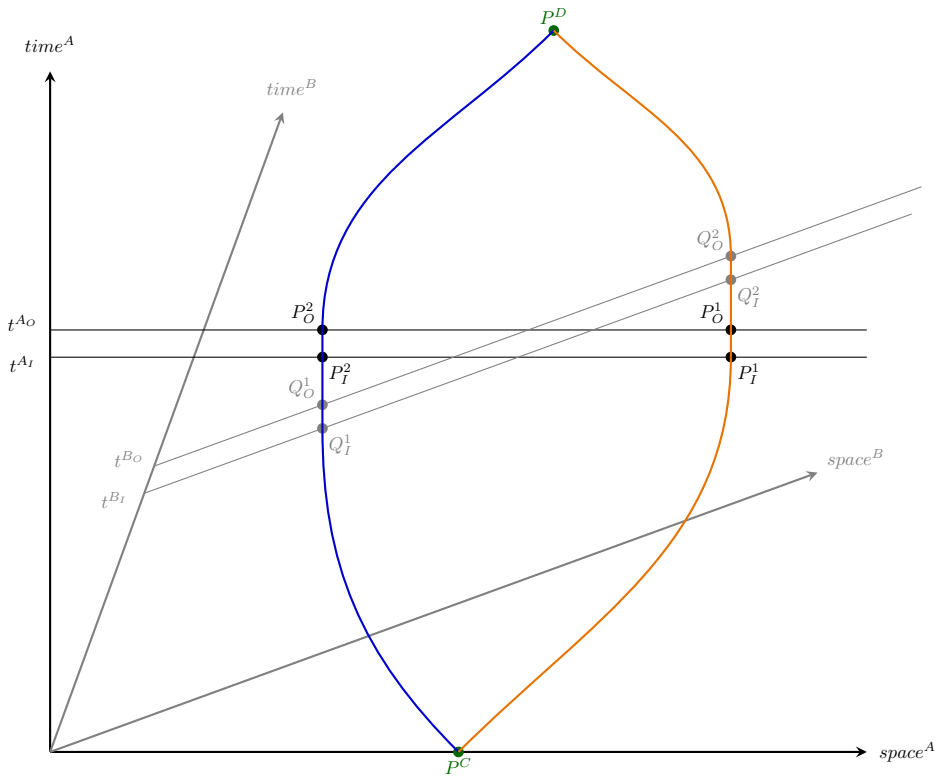

Figure 10: **Minkowski spacetime diagram for new quantum switch protocol with time localised systems** Alice labels her spacetime locations with respect to the co-ordinate system associated with the black axes and Bob with respect to the gray axes. Alice and Bob's reference frames are related by a Lorentz transformation. Here, the 45 degree line (with respect to Alice's axes) would correspond to a light-like surface. The black spacetime points are associated with Alice's in and output events and gray ones are associated with Bob's in and output events. The green points P^C and P^D are associated with the parties in the global past and future respectively, P^C is in the past light cone of P_I^1 and Q_I^1 while P^D in the future light cone of P_O^2 and Q_O^2. When the control is in the state zero, the target follows the blue path through the spacetime, going to Alice first and then to Bob and when the control is in the state $|1\rangle$, the target follows the orange path through spacetime going to Bob first and then to Alice. Irrespective of the order, Alice's in and output events are localised at times t^{A_I} and t^{A_O} in her frame and Bob's in and output events are localised at time t^{B_I} and t^{B_O} in his frame. However, depending on the order, Alice and Bob apply their local operation to a non-vacuum target state are two different space-like separated locations on the same time-slice.

frame of a distant party, Charlie, Alice's events are localised in space but delocalised in time, and the same for Bob's events [61].

Another relevant aspect in above comparison is the number of "elemental events" (i.e., quantum in/output events associated with a single spacetime locations) in each perspective. Note that this is preserved in fixed spacetime implementations as defined in our framework, even when spacetime delocalisation is involved. This is because spacetime localisation/delocalisation is an agent-independent concept in our framework. A physical assumption that can be associated with this property is that the reference frames of different agents share a common origin (e.g., the agents agree on a common event that is taken to be localised in all their frames and describe spacetime distances relative to this event) with respect to which all agents can measure spacetime distances using their respective spatio-temporal reference frames. If there was no common origin, agents would naturally make the choice of localising their own events at their "origin" and describing other agents' events relative to this, such that each event would be localised in spacetime with respect to the corresponding agent, but may be delocalised with respect to another agent.

In the gravitational implementation, the in/output events of Alice and Bob are typically regarded as each being single "spacetime" events from the local perspective, even though these appear to be different spacetime events from the perspective of a distant observer, Charlie who does not see the gravitational field. Whether these must be regarded as single events from the local perspective also depends on whether

or not the events appear to be spatially localised in these perspectives (as they are temporally localised by construction) which in turn depends on how spatial co-ordinates are measured. For instance, if the agents measure spatial distances with respect to the gravitating mass, then by construction, their spatial co-ordinate would depend on the branch of the superposition they are in (cf. Figure 8) and they would not be spatially localised. As Alice and Bob's clocks are initially synchronised before they enter the superposed gravitational fields, their temporal origin i.e, where $t^A = t^B = 0$ coincides with that of the distance observer. The spatial origin is however unclear. A natural question would therefore be whether there exists a common spatial and temporal origin for Alice and Bob in the gravitational scenario and a physical way for them to measure spacetime distances from this origin event using their respective reference frames such that they perceive their in/output events as being spacetime localised? This will depend on how spatial distances can be measured in such quantum gravitational settings. If the answer is yes in some model of measuring spatial distances in this gravitational scenario, we know from our results, that this achieves a property that is not possible in fixed spacetime implementations (cf. Corollary 6.4). If the answer is no, then, we have a set of events that are time localised and spatially delocalised in Alice's frame while the same events are time delocalised and spatially localised in Charlie's frame, which is also impossible in fixed spacetime implementations as we have argued above.

We note that even in classical general relativity, while the proper time is a well defined operational concept, this is not always the case for spatial distance. In order to measure one's spatial distance with respect to an event, one would typically calculate the light crossing time in both directions between oneself and that event, but this property is not locally defined at the observer's location. We leave a further investigation of the above questions for future work.

Protocol for distinguishing definite and indefinite spacetime implementations of QS In [50], a protocol has been proposed for distinguishing between the physically realised optical implementations and theoretically proposed gravitational implementations of QS in a manner that does not disturb the coherence between the different branches of the superposition. For the definite spacetime case, [50] focuses on a QS implementation in Minkowski spacetime described with respect to a single global frame where the in/output events of Alice and Bob are spatially localised but temporally delocalised. Here we recast their protocol in our framework and show that the core argument can be generalised to arbitrary fixed spacetime implementations.

The main idea behind the protocol of [50] is to introduce an additional agent F, hereby known as "Friend" to whom Alice and Bob send out photons in addition to performing their usual operations in the QS scenario. F is assumed to be spatially localised and can measure information regarding the times of arrival of the photons arriving from Alice and Bob to decide whether or not the local operation in each lab was a single spacetime event. In the optical implementation of QS, Alice and Bob act (on a non-vacuum state) either at an earlier or later time depending coherently on the control qubit. In this case, each of their photons are in a coherent superposition of arriving to F at different times. In gravitational implementations, the superposition of spacetime metrics can in principle be used to ensure that any photon from Alice always arrives to F at the same time t^A and any photon from Bob always arrives to F at the same time t^B. A non-demolition measurement is then performed by F to distinguish these two scenarios without collapsing the superposition of orders (see [50] for details).

We can model this protocol by considering a new 5-partite process map $\hat{W}_F^{QS}$ obtained from the 4-party quantum switch map $\hat{W}^{QS}$ by including the party F with input systems F_I^A and F_I^B and a trivial output space. We can give additional outputs A_O^F and B_O^F to Alice and Bob for the photons being sent to F and the process vector $|W_F^{QS}\rangle$ is then identical to $|W^{QS}\rangle$ of Equation (14), but with the additional factor $|\mathbb{1}\rangle\!\rangle^{A_O^F F_I^A}|\mathbb{1}\rangle\!\rangle^{B_O^F F_I^B}$ in both terms of the superposition, representing the identity channels from A_O^F to F_I^A and B_O^F to F_I^B. We can then see that any Minkowski spacetime embedding of this protocol where F_I^A is assigned a single spacetime location $(\mathbf{r}_F, t_F^A)$ and F_I^B is assigned a single spacetime location $(\mathbf{r}_F, t_F^B)$ would not satisfy relativistic causality if we require Alice's outputs A_O^F and A_O to be embedded into the spacetime region $\mathcal{P}^{A_O}$ and Bob's outputs B_O^F and B_O to be embedded in the region $\mathcal{P}^{B_O}$, such that $\mathcal{P}^{A_O}$ and $\mathcal{P}^{B_O}$ belong to the past light-like surface of $(\mathbf{r}_F, t_F^A)$ and $(\mathbf{r}_F, t_F^B)$ respectively. Without loss of generality, let $t_F^A > t_F^B$. Then the past light cone of $(\mathbf{r}_F, t_F^B)$ is fully contained in the past light cone of $(\mathbf{r}_F, t_F^A)$, which implies that $\mathcal{P}^{A_O} \overset{R}{\not\to} \mathcal{P}^{B_O}$. Now, in any fixed spacetime implementation of the network consisting of the action of $\hat{W}_{QS}^F$ on a set of local maps of the parties, we will have signalling

from Alice's output A_O to Bob's output B_O (through Bob's input B_I) and this would therefore violate relativistic causality in this embedding.

More generally, we can consider an arbitrary spacetime $\mathcal{T}$, and assume that F_I^A is embedded at a spacetime location $P^{F_I^A}$ and F_I^B at the location $P^{F_I^B}$, taking $P^{F_I^B} \prec P^{F_I^A}$ (the argument for $P^{F_I^A} \prec P^{F_I^B}$ is analogous). Denoting $\text{Past}(P) := \{Q \in \mathcal{T} | Q \prec P\}$ to be the past of a spacetime point P, we have $\text{Past}(P^{F_I^B}) \subseteq \text{Past}(P^{F_I^A})$. Then, if we impose that the output spacetime locations $\mathcal{P}^{A_O}$ of Alice must lie in $\text{Past}(P^{F_I^A}) \backslash \text{Past}(P^{F_I^B})$ and $\mathcal{P}^{B_O} \subseteq \text{Past}(P^{F_I^B})$, the implementation of the QS protocol would not satisfy relativistic causality since have $\mathcal{P}^{A_O} \overset{R}{\not\to} \mathcal{P}^{B_O}$ even though $A_O \to B_O$. However, in a gravitational implementation where the spacetime locations are described with respect to local quantum reference frames [61], it might nevertheless be possible to satisfy all these conditions in an implementation of QS, the gravitational implementations proposed in [50] being particular examples.

8 Demystifying indefinite causation

The framework and results presented here can be used to provide a clearer operational interpretation of the notion of indefinite causal structures. They allow for the analysis of several open questions relating indefinite causation to cyclic causation, and quantum information processing in spacetime through new tools and insights. Many of the possible future steps stemming from this work, relating to the composability of physical processes, physicality of causal inequality violations, and causal models for processes have been outlined in the main text (cf. Remarks 2.2, 2.8, 6.7). Here we discuss the broader outlook provided by our work, and also clarify the operational meaning of indefinite causation suggested by our work, by explicitly analysing the assumptions behind causal inequalities and the process framework.

Disentangling the information-theoretic and spacetime causal structures Our results bring clarity to the apparently paradoxical situation– while a fixed spacetime structure defines a fixed causal structure in relativistic physics, there continue to be numerous claims [37, 40–43, 45–49] that indefinite causal structures have been physically implemented in tabletop experiments within Minkowski spacetime. In the absence of a clear resolution to this apparent paradox, the operational meaning of an indefinite or a cyclic causal structure remains obscured. By disentangling the information-theoretic and spacetime notions of causation, and then characterising how they fit together we have shown that relativistic causality is indeed upheld in such experiments. On the other hand, it is the claim regarding the indefiniteness of the implemented causal structure that is to be questioned (as also noted in previous works such as [50]). Our results indicate that no experiment satisfying relativistic causality in a fixed background spacetime can potentially implement an indefinite causal structure, any such implementation could ultimately be described by a fixed acyclic information-theoretic causal structure involving (space-)time delocalised quantum systems that is compatible with the relativistic causal structure of the spacetime. Therefore the notion of an indefinite causal structure could potentially only make sense in quantum gravitational settings. No-go results characterising what is impossible to achieve in a fixed spacetime are important also for understanding how physics in these more exotic scenarios may differ. To this effect, we have applied our no-go results to discuss the similarities and differences between the fixed spacetime and quantum gravitational implementations of the quantum switch in Section 7.5. This disentangling of the two notions of causation also sheds light on the notion of events as discussed below.

Notion of events The operational notion of causality characterised through the possibility of signalling between physical systems yields a definite (cyclic or acyclic) causal structure, even when it involves quantum systems. On the other hand, with respect to a spacetime structure, causality is understood as the condition that causes must be in the past of the corresponding effect with respect to this spacetime structure (for any suitable definition of cause and effect). A key feature distinguishing the "causal structure" associated with these two notions is what the nodes or the "events" in the causal structure represent. For instance, in the operational causal structure, the nodes correspond to operational events such as "Alice received a quantum system in her lab", "Alice sent a quantum system out of her lab", "Alice set the knob of her measuring device to a certain position". In the spacetime notion of causation, the events may correspond to spacetime locations, which may apriori have no operational meaning until we associate physical systems with these points (in the language of our framework, "embed systems in

the spacetime"). Another spacetime related causal structure could be one where the nodes correspond to spacetime regions, rather than individual spacetime locations. But we may not wish to regard this as a separate notion of causation as we know that if we look at the latter in a sufficient level of detail, we would get back the former description in terms of individual spacetime points. Therefore, what constitutes an operational or spacetime event also depends on the level of detail at which one analyses the situation. This is captured by the notion of a fine-graining introduced here, which can be applied to operational, spacetime or any other abstract notion of causation that can be represented by a directed graph, and also to systems and quantum maps that may give rise to such a causal structure. In physical implementation of an indefinite causal order process, we associate spacetime regions with the quantum in and output systems of the parties, and these systems may be delocalised over the associated region. The fine-graining of spacetime regions into individual spacetime points induces a fine-graining of an information theoretic causal structure embedded in the spacetime. Formalising this, allowed us to show that any such implementation of an indefinite causal order process can be ultimately fine-grained to a fixed order process over a larger number of information-theoretic events (in the process language, this corresponds to a larger number of parties), where each information theoretic event is associated with a single spacetime event. While we have established that such a fine-graining is always possible, and thereby resolved the apparent tension between the two notions of causation at play in such experiments, further work is needed in characterising the exact fine-graining map. It would be interesting to characterise how correlations behave under fine-graining, we have only focussed on the possibility of signalling here and not the strength of signalling or of correlations.

Indefinite vs cyclic causal structures In the process matrix framework, the lack of a definite acyclic (information-theoretic) causal order is interpretted as indicating the indefiniteness of the causal structure. However, the lack of a definite acyclic causal structure (according to any notion of causation) need not imply that the causal structure is indefinite, but can also mean that the causal structure is definite but cyclic. Our work (building upon previous insights on the relation between indefinite and cyclic causal structures [29, 51, 52]) shows a tight correspondence between (classically or quantumly) indefinite causal structures in the sense of the process framework and cyclic causal structures as characterised by signalling relations. Interestingly, cyclic causal structures can capture both physical scenarios with feedback as well model physics in the presence of exotic closed timelike curves (CTCs). Indeed, we have derived the process matrix probability rule and construction under the former view while this has been done using CTCs in previous works such as [52]. The physical distinction between these two situations come from how the cyclic causal structure is instantiated with spacetime information. We have shown that in the case that we view the cyclic causal structure as being implemented, compatibly with relativistic causality in a fixed acyclic spacetime, then we can ultimately unravel the causal structure into an acyclic one through a fine-graining, even if the spacetime implementation allows the quantum systems in the causal structure to be delocalised in the spacetime. On the other hand, if we assign each node of the cyclic (information-theoretic) causal structure a single location in a spacetime, then we would have bidirectional causation between two spacetime locations which would correspond to a CTC.

Finally we note that there are several previous results linking indefinite and cyclic causal structures [29, 51, 52]. While these results provide significant insights on the simulability of the former in terms of the latter, they do not fully solve the problem regarding the operational meaning of a given process matrix, its physical realisability in a spacetime or relativistic causality therein. In particular, they leave open the questions of what it means to implement a process matrix in a spacetime, and which physical implementations of a process matrix in a spacetime would violate relativistic causality in the spacetime, which we have addressed here.

Quantum spacetime and quantum reference frames In Section 7.5, we have compared and contrasted the physical Minkowski spacetime realisations of the quantum switch with the hypothetical gravitational implementation in a quantum indefinite spacetime, in light of our framework and results. We note that the very interpretation of a spacetime as being definite or indefinite can depend on the choice of reference frames used to characterise spacetime information. For instance, spacetime as a fixed partial order as we have done here only makes sense when the identity of locations in this partial order and the corresponding order relations are preserved under reference frame transformations. This is often the case with classical reference frame transformations, such as the Lorentz transformations

in Minkowski spacetime. When describing spacetime co-ordinates using classical reference frames, a spacetime point P that is assigned the co-ordinates $(\mathbf{r}_A, t_A)$ in agent A's frame transforms to some coordinates $(\mathbf{r}_B, t_B)$ in agent B's frame, but A and B agree on the identity of P as a single spacetime location and on the order relation $P \prec Q$ between any P and Q. This property need not be preserved in scenarios (possibly classical) where for instance A and B do not share a common origin, in which case agents would tend to regard different events as being localised at the origin of their co-ordinates and can disagree on which events are localised or delocalised (see discussion in Section 7.5). This can also be the case where agents use quantum reference frames, an event that is spacetime localised at some location P with respect to one agent may appear to be a highly delocalised event with respect to another agent [61–64]. Quantum versions of the equivalence principle [65] suggest a correspondence between definite spacetime with quantum agents/reference frames (e.g., in a superposition of accelerations) and indefinite spacetimes with possibly classical agents, and [66] explicitly demonstrates such a correspondence for the quantum switch.

We believe that the framework introduced here can also be relevant to these more general scenarios beyond fixed acyclic spacetimes and classical reference frames. The localisation/delocalisation properties perceived by different agents could be modelled by considering different fine-grainings or spacetime embeddings of the same operational causal structure. This way, one can distinguish between the physical information (the coarse-grained operational causal structure) that all agents agree on, and the frame dependent information (the fine-graining, and embedding) that may differ from agent to agent. For instance, is it possible to formalise the condition that the order relations between "spacetime events" is preserved for in all reference frames even though their localisation may be agent or reference frame dependent? Such questions provide an interesting avenue for future research.

Operational meaning of causal inequality violations As discussed in Section 4, causal inequalities are shown to be necessary conditions on correlations generated by protocols satisfying the assumptions: free choice (FC), local order (LO), closed labs (CL) and causal structure (CS) i.e., a violation of causal inequalities under FC, LO and CL implies a violation of CS. This is often interpreted as implying the indefiniteness of the causal structure.

However, a more careful look reveals that CS is essentially two assumptions: (CS1) there exists a global partial order such that signalling is only possible from past to future with respect to this partial order (CS2) input/output events of every party in a multi-partite process are localised with respect to this partial order. In our framework, (CS1) applied to a protocol implies the existence fixed spacetime implementation of the protocol (Definitions 3.1 and 3.4) that satisfies the necessary condition for relativistic causality (Definition 3.5). (CS2) then corresponds to a constraint on the spacetime embedding, requiring the spacetime region $\mathcal{P}_S$ assigned to each system S to be a single spacetime point. FC and CL are implicit in the construction, as they are in the process framework: A_I and A_O are the only external in/outputs of each party A and their input systems carrying settings are never composed with other systems. Demanding relativistic causality for local maps ensures that the input spacetime location precedes the output spacetime location and is a necessary (but not sufficient) condition for LO. In the presence of spacetime delocalisation, we must also demand that the local maps act trivially on vacuum states in order to preserve the LO condition (see Section 6.2 for an explanation). Our no-go result of Corollary 6.4 therefore implies the following: under FC, LO, CL and CS (=CS1+CS2), the only physical processes that can be implemented are fixed order processes. Theorem 6.2 makes a much more general statement that establishes the above under a weaker assumption on the embedding that CS2 which requires perfect spacetime localisation, this weaker assumption allows in/output events to be delocalised in spacetime. These results rule out not only causal inequality violating processes, but also causally non-separable ones and classical mixtures of causal orders and is thus a stronger statement than the derivation of causal inequalities in [28] which rule out bipartite non-causal processes under the assumptions FC, LO, CL and CS (=CS1+CS2).

An open question that is not fully answered by our results is whether it is possible to violate causal inequalities in a fixed spacetime where events are delocalised over the spacetime i.e., where (CS2) is not satisfied. Our results nevertheless provide insights into this question. The fixed spacetime corresponds to a partial order (Definition 3.1) and relativistic causality (Definition 3.5) forbids signalling outside the spacetime future, this formalises CS1 independently of CS2 in our framework. Then the above question amounts to asking whether is it possible to violate causal inequalities whenever FC, LO, CL and CS1

are satisfied. But such a violation, even if possible would not certify the indefiniteness of the causal structure, but only the violation of CS2 i.e., that the in/output events are not localised in the relativistic causal structure (i.e., the spacetime $\mathcal{T}$). But we have already shown in Theorem 6.2 that in order to implement any process that is not a fixed order process, without violating relativistic causality in a fixed spacetime, we must necessarily violate CS2, and any such implementation will ultimately have a definite acyclic signalling structure by Theorem 6.2. This suggests that such a violation of causal inequalities, even if possible, would not tell us much more than what we have already shown here. It is nevertheless interesting to consider whether this is possible at all, as it would then correspond to asking whether the fact that the information-theoretic events are not well localised spacetime events can be certified in a device independent manner. Our Theorem 6.9 regarding the ability to fine-grain any physical process into a fixed order process strongly suggests that the answer to this question is negative, it implies that the initial non-fixed order process that we did physically implement in the spacetime is simply a coarse-graining of a fixed order process (which by construction does not violate causal inequalities) and we would not expect to be able to violate causal inequalities under coarse-graining. However proving this requires a few more steps, such as the assumptions for ruling our trivial causal inequality violations in the presence of spacetime delocalised systems, and connecting causality conditions for process matrices over different numbers of parties. Formally showing that relativistic causality in a fixed spacetime rules out non-trivial causal inequality violations, as indicated by our results here, is a subject of ongoing work. Recent results [59, 67, 68] suggest that causal inequality violations are not possible for a general class of protocols implementable in Minkowski spacetime with quantum systems, which corroborate with these observations.

Finally, we note that there is a proposal for violating causal inequalities in Minkowski spacetime (i.e., under CS1) using quantum fields [44]. Since these are infinite dimensional systems, our proofs do not directly cover this case, but the overall intuition discussed above nevertheless does. A particular problem in this case is that it is unclear how the closed lab assumption can be formulated since there is no clear notion of subsystems. Therefore it is ambiguous if this causal inequality violation corresponds to a violation of only of CS2 or also of CL (or LO), and whether it conflicts with the above intuition that causal inequalities cannot be violated under FC, CL, LO and CS1. This remains to be further analysed. In either case, this would not certify the indefiniteness of the causal structure since this would require a violation of CS1 which is satisfied by construction as the proposed implementation is in Minkowski spacetime with classical reference frames.

Non-causal processes, time delocalised systems and CTCs A recent work [69] describes a method realise several process matrices using time delocalised systems and in particular shows that there exist non-causal processes that admit realisations through quantum circuits on time-delocalised systems. This might seem to be in apparent contradiction with our results and what was said in the above discussions, and we clarify this here. While [69] considers time delocalised systems, they do not consider a background causal structure with respect to which these systems are delocalised or relativistic causality with respect to this causal structure. We have on the other hand modelled implementations of processes in a fixed spacetime while allowing systems to be be delocalised over this fixed spacetime structure as long as relativistic causality in this spacetime is not violated. In particular, the spacetime is a partially ordered one and does not admit "closed timelike curves". Then our related work [59] suggests that it is impossible to physically implement non-causal processes compatibly with relativistic causality in such a spacetime even if we allow for systems to be arbitrarily delocalised in space and in time. Allowing the spacetime to be a pre-order rather than a partial order would evade such a result and enable constructions such as those of [69] to be physically implemented in such a spacetime. However, such a pre-ordered spacetime can only arise from exotic solutions in GR that include closed timelike curves which may itself be considered unphysical. This distinction also depends on how the assumptions for ruling out trivial causal inequality violations are formulated in the presence of spacetime delocalised systems (such as: what does it mean to apply a local operation only "once"). For instance, [59] provides such a set of assumptions. We leave a more detailed investigation of this point, along with a comparison of the underlying assumptions to future work.

9 Conclusions

We have developed a general framework for clearly distinguishing between and formulating different notions of causality, which allows the analysis of their relationships at different levels of detail, and meets the criteria outlined in the introduction. We have also shown that this approach sheds light on the physicality of so-called indefinite causal order processes. This flexibility allows us to model physical scenarios with feedback as cyclic causal structures, when we only care about the information-theoretic properties, or fine-grain them using spacetime embeddings into acyclic causal structures compatible with the spacetime, in scenarios where we are also interested in the spacetime information. The advantage is that we can be very general, and model possibly post-quantum scenarios and but can also make a connection to (in principle) physically realisable quantum experiments in a definite spacetime, as per our need and interest. Our main results have been summarised in the introduction, so we will not repeat that here. We conclude here with the main take home messages of this work.

Our work highlights the need to clearly distinguish between information-theoretic and spacetime related notions of causation and events, which was also stressed in a recent paper involving one of us [2]. Even within each of these broad notions, there are further distinctions to be made, for instance even at a purely information theoretic level, causation and signalling and not equivalent, and we can analyse a causal structure at different levels of details and therefore draw different conclusions. Such a disentangling of concepts provides a vast level of generality and flexibility as it involves minimal assumptions. However, this disentangling alone is not enough, for it provides no insights into physical implementations. We must therefore also establish formal connections between these different notions in order to be able to talk about physical experiments.

In the natural process of science, we may often need to update our preconceived notions in light of new experiments. For instance, Bell experiments pose a serious challenge to a classical description of cause and effect at the operational level, while they are compatible with relativistic notions of causation such as no superluminal signalling [16]. Consequently in light of Bell experiments, we are forced to update our prior understanding of the interface between information-theoretic and spacetime notions of causality and events. A theory of quantum gravity may further challenge our understanding of these notions and their interface, but we cannot fully anticipate how. The approach of disentangling different notions and carefully reconnecting them, as proposed in the present paper and in [2] would help us better prepare for such challenges, as it would enable us to identify the notions of causation that are retained and those which are challenged in light of new experiments. As we have seen, this also helped us reinterpret the results and claims of existing experiments.

Bell's theorem has set an unprecedented example in highlighting the power of no-go theorems both for foundational and practical purposes— establishing what is impossible to achieve within certain physical regimes, tells us how physics in new regimes deviates from our prior intuitions and how we can exploit these new physical phenomena for useful practical tasks. Our work outlines a number of no-go results for the characterisation of quantum causal structures and their interplay with a definite spacetime structure, and we have discussed how hypothetical quantum gravitational scenarios without a definite spacetime structure might deviate from some of these results. Apart from these foundational considerations, this framework could have potential applications for the study of quantum information processing tasks in spacetime that involve spacetime delocalised quantum systems. The methods developed here also have applications for the study of cyclic quantum causal models, which can be used to describe both physical quantum scenarios with feedback as well as more exotic closed timelike curves. There are several fascinating questions that still remain open, as we have outlined throughout this text, many of which could be addressed by building on the tools and ideas presented in this paper, and we leave this for future work.

Acknowledgements VV thanks Esteban Castro-Ruiz, Ämin Baumeler and Lorenzo Maccone for enlightening discussions. VV's research is supported by the ETH Postdoctoral Fellowship. RR acknowledges support from the Swiss National Science Foundation via project No. 200021_188541 and via the National Centre of Competence in Research "QSIT" (grant No. 51NF40-185902).

Appendix

A Loop composition

Here we show how we obtain Definition 2.1 from the original definition of loop composition proposed in [54] for CPTP maps on infinite dimensional systems. Given a CPTP map $\hat{\Phi} : \mathcal{L}(\mathcal{H}_A) \mapsto \mathcal{L}(\mathcal{H}_B)$, the Choi representation of $\hat{\Phi}$ is an operator on $\mathcal{L}(\mathcal{H}_A) \otimes \mathcal{L}(\mathcal{H}_B)$ are is given as $\sum_{i,j} |i\rangle\langle j|_A \otimes \hat{\Phi}(|i\rangle\langle j|_A)$, which is a positive semi-definite operator and where $\{|i\rangle\langle j|_A\}_{i,j}$ is a basis of $\mathcal{L}(\mathcal{H}_A)$. If $\mathcal{H}_A$ and $\mathcal{H}_B$ are infinite dimensional, the Choi operator can be unbounded, and the Choi representation is instead given as a sesquilinear positive semi-definite form $R_{\hat{\Phi}}$ on

$$\mathcal{H}_B \times \mathcal{H}_A = \mathrm{Span}\{\psi_B \otimes \psi_A : \psi_B \in \mathcal{H}_B, \psi_A \in \mathcal{H}_A\},$$

satisfying

$$R_{\hat{\Phi}}(\psi_B \otimes \psi_A; \phi_B \otimes \phi_A) := \langle \overline{\psi}_B | \hat{\Phi}(|\overline{\psi}_A\rangle\langle\overline{\phi}_A|) | \overline{\phi}_B\rangle, \tag{17}$$

where $|\overline{\psi}\rangle = \sum_i |i\rangle\overline{\langle i|\psi\rangle}$ for some fixed basis $\{|i\rangle\}_i$ of $\mathcal{H}_A$.

Then the original definition of loop composition given in [54] is as follows. In this definition, it will be more convenient to put system labels inside the bras and kets rather than using them as subscripts outside the bras and kets as we have been doing so far, e.g., $|k_B\rangle$ instead of $|k\rangle_B$. The meaning is however the same.

Definition A.1 (Loop composition of infinite dimensional CPTP maps [54]) *Consider a CPTP map $\hat{\Phi} : \mathfrak{T}(\mathcal{H}_{AB}) \mapsto \mathfrak{T}(\mathcal{H}_{CD})$ with input systems A and B and output systems C and D, of dimensions d_A, d_B, d_C and d_D with $d_B = d_D$, and where $\mathfrak{T}(\mathcal{H})$ denotes the set of trace class operators on the (possibly infinite dimensional) Hilbert space $\mathcal{H}$. Let $\{|k_D\rangle\}_k$ and $\{|l_D\rangle\}_l$ be any orthonormal bases of $\mathcal{H}_D$, and $\{|k_B\rangle\}_k$ and $\{|l_B\rangle\}_l$ denote the corresponding bases of $\mathcal{H}_B$ i.e., for all k and l, $|k_D\rangle \cong |k_B\rangle$ and $|l_D\rangle \cong |l_B\rangle$. Then the Choi representation of the new map $\hat{\Psi} = \hat{\Phi}^{D\hookrightarrow B} : \mathfrak{T}(\mathcal{H}_A) \mapsto \mathfrak{T}(\mathcal{H}_C)$, resulting from looping the output system D to the input system B in the map $\hat{\Phi}$ is given as*

$$R_{\hat{\Psi}}(\psi_C \otimes \psi_A; \phi_C \otimes \phi_A) = \sum_{k,l} R_{\hat{\Phi}}(\psi_C \otimes k_D \otimes \psi_A \otimes \overline{k}_B; \phi_C \otimes l_D \otimes \phi_A \otimes \overline{l}_B), \tag{18}$$

where $|\overline{k}_B\rangle = \sum_i |i\rangle\overline{\langle i|k\rangle}$ for the basis $\{|i_B\rangle\}_i$ of $\mathcal{H}_B$ used in the Choi representation of $\hat{\Phi}$.

Using Equation (17), the right hand side of Equation (18) becomes

$$R_{\hat{\Psi}}(\psi_C \otimes \psi_A; \phi_C \otimes \phi_A) = \langle \overline{\psi}_C | \hat{\Psi}(|\overline{\psi}_A\rangle\langle\overline{\phi}_A|) | \overline{\phi}_C\rangle.$$

The left hand side of Equation (18) is

$$\sum_{k,l} R_{\hat{\Phi}}(\psi_C \otimes k_D \otimes \psi_A \otimes \overline{k}_B; \phi_C \otimes l_D \otimes \phi_A \otimes \overline{l}_B) = \sum_{k,l} \langle \psi_C | \langle k_D | \hat{\Phi}\Big(|\overline{\psi}_A\rangle\langle\overline{\phi}_A| \otimes |k_B\rangle\langle l_B|\Big) |\phi_C\rangle |l_D\rangle.$$

Noting that both the side of the equation are of the form $\langle\psi_C|(...)|\phi_C\rangle$, the expression corresponding to the dots in the parenthesis must be the same, and we have

$$\hat{\Psi}(|\overline{\psi}_A\rangle\langle\overline{\phi}_A|) = \sum_{k,l} \langle k_D | \hat{\Phi}\Big(|\overline{\psi}_A\rangle\langle\overline{\phi}_A| \otimes |k_B\rangle\langle l_B|\Big) |l_D\rangle,$$

Now, if $|\psi_A\rangle$ and $|\phi_A\rangle$ happen to be two basis states $|i^1\rangle$ and $|i^2\rangle$ of the same basis $\{|i\rangle\}_i$ of $\mathcal{H}_A$, then we have $|\overline{\psi_A}\rangle = \sum_i |i\rangle\overline{\langle i|\psi_A\rangle} = \sum_i |i\rangle\overline{\langle i|i^1\rangle} = \sum_i |i\rangle\delta_{i,i^1} = |i^1\rangle$ and similarly $|\overline{\phi_A}\rangle = |i^2\rangle$. Then the action on basis states is given as follows, which is the same as Equation (1) of the main text, where we now revert to our original convention of putting system labels as subscripts outside the bras and kets and shorten expressions such as $|k_B\rangle\langle l_B| \cong |k\rangle_B\langle l|_B$ to $|k\rangle\langle l|_B$.

$$\hat{\Psi}(|i\rangle\langle j|_A) = \sum_{k,l} \langle k|_D \hat{\Phi}\Big(|i\rangle\langle j|_A \otimes |k\rangle\langle l|_B)\Big) |l\rangle_D,$$

where $|i\rangle$ and $|j\rangle$ are elements of the same orthonormal basis of $\mathcal{H}_A$. The above is a convenient form for calculating the Choi operator $\sum_{i,j} |i\rangle\langle j|_A \otimes \hat{\Psi}(|i\rangle\langle j|_A)$ of the final map, in the finite dimensional case (which is what we focus on in this paper).

B QS as a higher-order transformation

Here we briefly review the quantum switch transformation as originally proposed in [29]. The quantum switch is originally defined as a supermap QS, or higher-order transformation, that acts on the space of quantum channels (which are themselves linear maps) mapping a pair of quantum channels U^A and V^B to a new channel $W(U^A, V^B)$. The channel $W(U^A, V^B)$ thus obtained implements a coherently controlled superposition of the orders of U^A and V^B on a target system. In particular, given two unitary channels U^A and V^B that act on a target system T, the quantum switch maps them to the channel $W(U^A, V^B)$ acting on $\mathcal{H}_C \otimes \mathcal{H}_T$ (the joint Hilbert space of a control qubit and the target system) and given as

$$W(U^A, V^B) = |0\rangle\langle 0|_C \otimes V^B U^A + |1\rangle\langle 1|_C \otimes U^A V^B \tag{19}$$

For example, $W(U^A, V^B)$ acts on the initial state $(\alpha|0\rangle + \beta|1\rangle)_C \otimes |\Psi\rangle_T$ (where $|\Psi\rangle_T \in \mathbb{C}^d$ is an arbitrary pure state of the target qudit) as

$$W(U^A V^B): \quad (\alpha|0\rangle + \beta|1\rangle)_C \otimes |\Psi\rangle_T \quad \longrightarrow \quad \alpha|0\rangle_C \otimes \left(V^B\ U^A|\Psi\rangle\right)_T \ + \ \beta|1\rangle_C \otimes \left(U^A\ V^B|\Psi\rangle\right)_T . \tag{20}$$

More generally, one can consider the quantum switch operation on non-unitary local channels U^A and V^B. In this case, the action of $W(U^A, V^B)$ can be defined by constructing a set of Kraus operators for the channel $W(U^A, V^B)$ (which can be obtained given a set of Kraus operators for U^A and V^B) and specifying the action of each of the Kraus operators of $W(U^A, V^B)$ analogously to Equation (19). Equation (20) remains the same in the non-unitary case. An interested reader may refer to [70] for further details on the general definition. Importantly, it is required that each of the operations U^A and V^B are queried only once. With this requirement, it is known that it is not possible to implement the transformation QS using a standard quantum circuit acting on non-vacuum systems (this would require at least one of the operations to be queried twice) [29].

C An implementation of QS with a definite acyclic causal structure

Here we describe the fixed spacetime QS protocol of Section 7.4 more explicitly in terms of the underlying CPTP maps and their composition, and show that it implements a definite acyclic causal structure compatible with relativistic causality in Minkowski spacetime, even though each agent perceives events to be time localised in their own reference frame. For this, we employ the causal box framework [54] which models information processing protocols involving systems and operations that may be delocalised in a fixed spacetime that are compatible with relativistic causality in that spacetime.

We will not review the CB framework here as this is quite involved. For our purposes, it will suffice to say that a CB can be viewed as a fixed spacetime implementation of a CPTP map (Definition 3.4) that satisfies relativistic causality (Definition 3.5), it also takes into account vacuum and spacetime information. In general, causal boxes can have a larger state space (allowing for superpositions of different numbers of physical systems of a given dimension), and are closed under composition with arbitrary protocols involving multiple rounds of communication between parties, and not just protocols that can be viewed as the action of a process on local operations where each party acts once on a physical system. A more precise comparison and mapping between these frameworks is a subject of a follow-up work. Here, we present the causal box representation of the QS protocol of Section 7.4, as an example of a definite spacetime implementation of the process W^{QS} that admits a fixed acyclic causal structure.

Our quantum switch protocol of Section 7.4 is described by a maximally fine-grained fixed spacetime implementation $\hat{W}^{QS,\mathcal{T},\mathcal{E}}_{max}$ of the quantum switch process matrix $\hat{W}^{QS}$, where the embedding $\mathcal{E}$ is as given in Equation (16), which means that the in and output systems of $\hat{W}^{QS,\mathcal{T},\mathcal{E}}_{max}$ are $\{C_O^{C,P^C}, C_O^{C,P^C}, A_O^{P_O^1}, A_O^{P_O^2}, B_O^{Q_O^1}, B_O^{Q_O^2}\}$ and $\{D_I^{C,P^D}, D_I^{C,P^D}, A_I^{P_I^1}, A_I^{P_I^2}, B_I^{Q_I^1}, B_I^{Q_I^2}\}$ respectively. $\hat{W}^{QS,\mathcal{T},\mathcal{E}}_{max}$ corresponds to a causal box as it is a CPTP map respecting relativistic causality in the spacetime (we also explicitly show this below by describing its *sequence representation*, which is a defining feature of a causal box). We will there refer to it as the causal box $\hat{\mathbf{QS}}$ for short.

The local operations U^A and V^B are similarly implemented in spacetime and fine-grained to give $U^{A,\mathcal{T},\mathcal{E}}_{max}$ and $V^{B,\mathcal{T},\mathcal{E}}_{max}$ which also correspond to causal boxes, which we denote as $\hat{\mathbf{U}}^A$ and $\hat{\mathbf{V}}^B$ for short. $\hat{\mathbf{U}}^A$ has the inputs $A_I^{P_I^1}$ and $A_I^{P_I^1}$, and outputs $A_O^{P_O^1}$ and $A_O^{P_O^1}$ and applies U^A independently between the

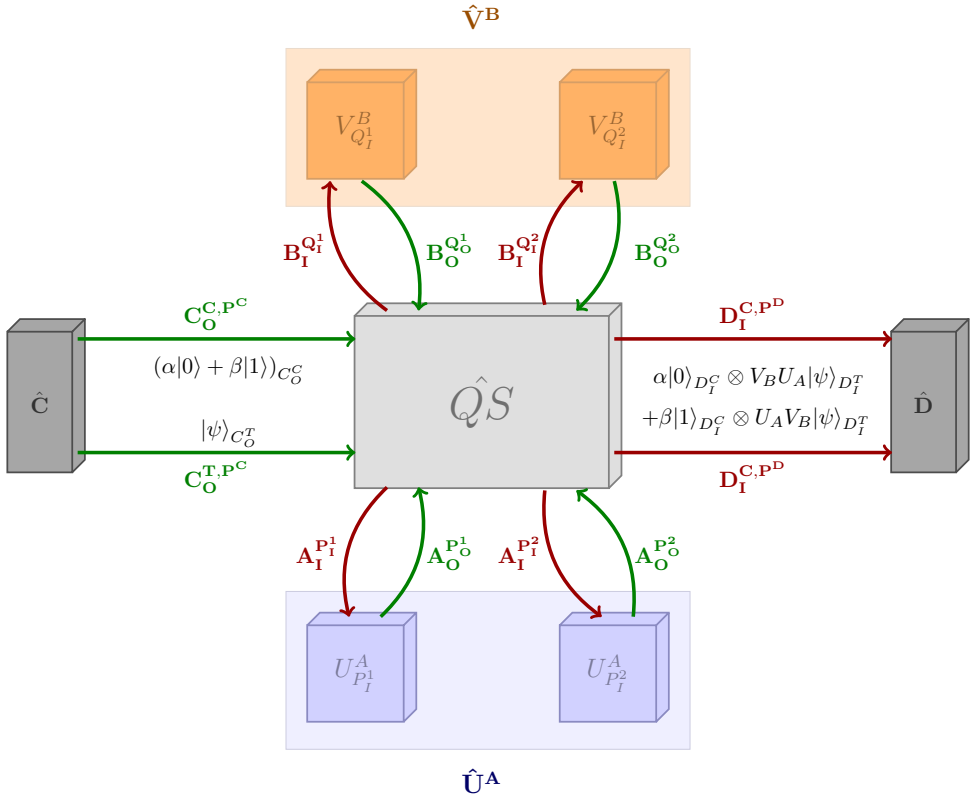

Figure 11: **The quantum switch protocol of Figure 10 as a composition of maximally fine-grained maps: $\hat{\mathbf{QS}}$** models the maximal fine-graining of the spacetime implemented process map of the protocol, the input and output systems of $\hat{\mathbf{QS}}$ are depicted in green and red respectively. $\hat{\mathbf{U}}^A$ and $\hat{\mathbf{V}}^B$ correspond to the maximally fine-grained spacetime implementations of the local operations U^A and V^B. These act independently between the respective pairs of in and output systems, and trivially on vacuum states (cf. Equation (21). A party C in the global past can prepare input states of the control and target systems, while a party D in the global future can receive the final state of the control and target from $\hat{\mathbf{QS}}$ and perform measurements on it. The composition of $\hat{\mathbf{QS}}$ with $\hat{\mathbf{U}}^A$ and $\hat{\mathbf{V}}^B$ yields the desired transformation (Equations (19) and (20)) from the global past to the global future. Furthermore, $\hat{\mathbf{QS}}$ acts as a fixed order process over 6 parties, in contrast to the original process $\hat{W}^{QS}$ (of which $\hat{\mathbf{QS}}$ is a physical implementation) which was an indefinite causal order process over 4 parties. This is also witnessed by the fact that causal structure of this physical protocol with $\hat{\mathbf{QS}}$ corresponds to the directed acyclic graph given in part (b) of Figure 9 while that of a protocol involving the process map $\hat{W}^{QS}$ would be the cyclic graph given in part (a) of Figure 9.

input-output pairs $A_I^{P_I^i}$ and $A_O^{P_O^i}$ for $i \in \{1,2\}$ whenever a non-vacuum state is received and acts trivially on the vacuum state i.e., $\hat{\mathbf{U}}^A = U^A_{P_I^1} \otimes U^A_{P_I^2}$, where the following holds for $i \in \{1,2\}$.

$$\begin{aligned} U^A_{P_I^i}|\Omega\rangle_{A_I^{P_I^i}} &= |\Omega\rangle_{A_O^{P_O^i}} \\ U^A_{P_I^i}|\psi\rangle_{A_I^{P_I^i}} &= |U_A(\psi)\rangle_{A_O^{P_O^i}} \end{aligned} \tag{21}$$

The description of $\hat{\mathbf{V}}^B$ is analogous. The composition of the causal boxes $\hat{\mathbf{QS}}$, $\hat{\mathbf{U}}^A$ and $\hat{\mathbf{V}}^B$ is a causal box that implements the channel of Equation (19) from an initial state of control and target on C_O^{C,P^C} and C_O^{T,P^C} to a corresponding final state on D_I^{C,P^D} and D_I^{T,P^D}, which may be prepared and measured by parties C and D in the global past and global future (associated with the local maps $\hat{\mathbf{C}}$ and $\hat{\mathbf{D}}$). This composition is depicted in Figure 11. Note that the quantum switch QS itself cannot be described as a standard quantum circuit, while it is a causal box [54].

The action of each of the causal boxes, $\hat{\mathbf{QS}}$, $\hat{\mathbf{U}}^A$, $\hat{\mathbf{V}}^B$ as well as their composition is can then be specified by a sequence of operations that the box implements at each time-step (i.e., through a

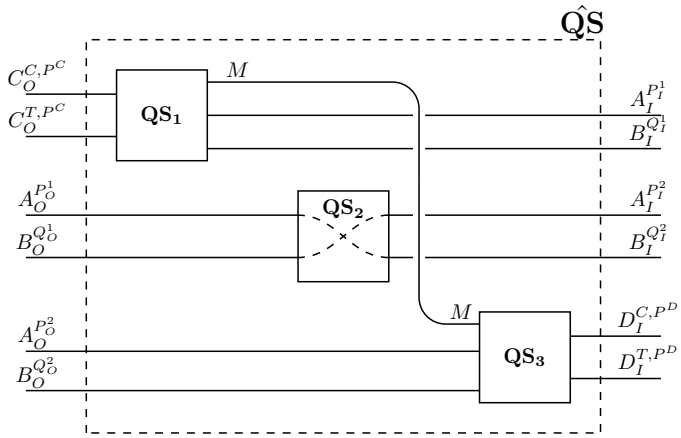

Figure 12: **Sequence representation of the causal box $\hat{\mathbf{QS}}$ of Figure 11:** If valid control and target states (non-vacuum qubit states) are sent in systems C_O^{C,P^C} and C_O^{T,P^C} respectively, QS_1 sends the control on the quantum memory system M and the target on $A_I^{P_I^1}$ (and $|\Omega\rangle$ on $B_I^{Q_I^1}$) if the control was $|0\rangle$, and on $B_I^{Q_I^1}$ (and $|\Omega\rangle$ on $A_I^{P_I^1}$) if the control was $|1\rangle$. QS_2 merely connects $A_O^{P_O^1}$ to $B_I^{Q_I^2}$ and $B_O^{Q_O^1}$ to $A_I^{P_I^2}$ as shown. QS_3 takes in the quantum memory M connects $B_O^{Q_O^2}$ to D_I^{T,P^D} (ignoring $A_O^{P_O^2}$) if the memory was $|0\rangle$ and $A_O^{P_O^2}$ to D_I^{T,P^D} (ignoring $B_O^{Q_O^2}$) if it was $|1\rangle$ and forwards M to D_I^{C,P^D}. If we consider composition of $\hat{\mathbf{QS}}$ with causal boxes $\hat{\mathbf{U}}_A$ and $\hat{\mathbf{V}}_B$, where the output systems $A_I^{P_I^1}$ and $A_I^{P_I^2}$ will connect to the input of $\hat{\mathbf{U}}_A$ and the input systems $A_O^{P_O^1}$ and $A_O^{P_O^2}$ connect to its output (and analogously for $\hat{\mathbf{V}}_B$), and we assume that these boxes satisfy Equation (21), we get the desired transformation on the joint state of the control and target that corresponds to a quantum controlled superposition of orders of U^A and V^B.

"sequence representation" of the causal box). The causal structure of this network corresponds to the directed acyclic graph given in part (b) of Figure 9. The decomposition of the causal box $\hat{\mathbf{QS}}$ (Figure 11) in terms of its action on these elemental systems is illustrated in Figure 12, this corresponds to the sequence representation [54] of the causal box $\hat{\mathbf{QS}}$.

One can then easily verify that the causal box described in the figure indeed implements the desired transformation (Equations (19) and (20)) (see also [54]). Further, to check that the operations U^A and V^B are indeed only queried once each in this implementation of the quantum switch, the corresponding boxes $\hat{\mathbf{U}}_A$ and $\hat{\mathbf{V}}_B$ can be provided with (internal) quantum counters each of which increment their value by one every time the corresponding operation U^A or V^B is applied to a non-vacuum state as explicitly shown in [54].[13]

Figure 12 also illustrates the causal structure of the causal box $\hat{\mathbf{QS}}$. We can see that the output $C_O^{P^C} = (C_O^{C,P^C}, C_O^{T,P^C})$ of Charlie is a cause of the first inputs $A_I^{P_I^1}$ and $B_I^{Q_I^1}$ of Alice and Bob through the map $\mathbf{QS}_1$. Further, $A_O^{P_O^1}$ causally influences $B_I^{Q_I^2}$, and $B_O^{Q_O^1}$ causally influences $A_I^{P_I^2}$ through the map $\mathbf{QS}_2$, and both $A_O^{P_O^2}$ and $B_O^{Q_O^2}$ causally influence Danny's input $D_I^{P^D} = (D_I^{C,P^D}, D_I^{T,P^D})$ through the map $\mathbf{QS}_3$. One can easily check that the signalling relations corresponding to each of these causal influences also holds.

D Proofs of all results

Lemma 2.12 *[Fine-graining a map preserves its signalling relations] Given a map $\hat{\Phi}$ and a fine-graining $\mathcal{F}^{sys}$ of its in/output systems $\mathcal{S}$, for every signalling relation $\mathcal{S}_I \to \mathcal{S}_O$ in $\hat{\Phi}$ between some subsets $\mathcal{S}_I \subset \mathcal{S}$ and $\mathcal{S}_O \subset \mathcal{S}$ of its input and output systems, there exists a corresponding signalling relation $\mathcal{F}^{sys}(\mathcal{S}_I) \to \mathcal{F}^{sys}(\mathcal{S}_O)$ in the fine-grained map $\hat{\Phi}_{\mathcal{F}}$. Consequently, the signalling structure $\mathcal{G}^{sig}_{\mathcal{F}}$ associated with $\hat{\Phi}_{\mathcal{F}}$ is a fine-graining of the signalling structure $\mathcal{G}^{sig}$ associated with $\hat{\Phi}$.*

[13]The action of the unitaries on the vacuum state $|\Omega\rangle$ is not counted since it represents "nothing" being given as input to the black-boxes and remains invariant under all operations.

Proof: Let us denote the set of all inputs I as $\{I_1, ..., I_n\}$, and without loss of generality, take $\mathcal{S}_I \subseteq I$ to be the set of the first $|\mathcal{S}_I|$ subsystems.

Firstly, we show that $\mathcal{S}_I$ signals to $\mathcal{S}_O$ in $\hat{\Phi}$ implies that there exists a basis state $|\mathbf{v}\rangle_I = \bigotimes_{i=1}^n |v_i\rangle_{I_i}$ belonging to an orthonormal basis $\{|\mathbf{v}\rangle_I\}_{\mathbf{v}}$ of $\mathcal{H}_I$ and another element $|\mathbf{v}'\rangle_{\mathcal{S}_I} = \bigotimes_{I_i \in \mathcal{S}_I} |v_i'\rangle_{I_i}$ of the same orthonormal basis restricted to the subset of inputs $\mathcal{S}_I$ such that $\mathrm{Tr}_{O\setminus\mathcal{S}_O} \circ \hat{\Phi}(|\mathbf{v}\rangle\langle\mathbf{v}|_I) \neq \mathrm{Tr}_{O\setminus\mathcal{S}_O} \circ \hat{\Phi}(|\mathbf{v}'\rangle\langle\mathbf{v}'|_{\mathcal{S}_I} \otimes \mathrm{Tr}_{\mathcal{S}_I}(|\mathbf{v}\rangle\langle\mathbf{v}|_I))$. This follows readily from linearity of the map $\hat{\Phi}$. $\mathcal{S}_I$ signals to $\mathcal{S}_O$ in $\hat{\Phi}$ implies that there exist states σ_I and $\rho_{\mathcal{S}_I}$ such that $\mathrm{Tr}_{O\setminus\mathcal{S}_O} \circ \hat{\Phi}(\sigma_I) \neq \mathrm{Tr}_{O\setminus\mathcal{S}_O} \circ \hat{\Phi}(\rho_{\mathcal{S}_I} \otimes \mathrm{Tr}_{\mathcal{S}_I}(\sigma_I))$. But if we cannot detect signalling in a linear map at the level of a complete orthonormal basis of the input space of the map, then we cannot do so using any pure states which are a linear combination of these, and consequently, we cannot do so using any mixed states which are convex mixtures of pure states, which contradicts the above equation for the existence of signalling.

Noting that $|\mathbf{v}'\rangle\langle\mathbf{v}'|_{\mathcal{S}_I} \otimes \mathrm{Tr}_{\mathcal{S}_I}(|\mathbf{v}\rangle\langle\mathbf{v}|_I)$ can simply be expressed as a new basis element

$$|\tilde{\mathbf{v}}\rangle := \bigotimes_{I_i \in \mathcal{S}_I} |v_i'\rangle_{I_i} \bigotimes_{I_j \in I\setminus\mathcal{S}_I} |v_j\rangle_{I_j} \in \mathcal{H}_I,$$

we have that

$$\mathrm{Tr}_{O\setminus\mathcal{S}_O} \circ \hat{\Phi}(|\mathbf{v}\rangle_I\langle\mathbf{v}_I|) \neq \mathrm{Tr}_{O\setminus\mathcal{S}_O} \circ \hat{\Phi}(|\tilde{\mathbf{v}}\rangle_I\langle\tilde{\mathbf{v}}_I|). \tag{22}$$

Using Definitions 2.10 and 2.11 (and the paragraph in between these definitions, setting out the notations therein), we know that for each basis state $\mathbf{v}_I = \bigotimes_{i=1}^n |v_i\rangle_{I_i}$, we have a corresponding fine-grained subspace $\mathcal{H}^{\mathbf{v}}_{\mathcal{F}^{sys}(I)}$ consisting of states of the form $|\psi^{\mathbf{v}}\rangle = \bigotimes_i |\psi^{v_i}\rangle$. Then it follows from applying Definition 2.11 to Equation (22) that in the fine-grained map $\hat{\Phi}_{\mathcal{F}}$, there exists $|\psi^{\mathbf{v}}\rangle \in \mathcal{H}^{\mathbf{v}}_{\mathcal{F}^{sys}(I)}$ and $|\psi^{\tilde{\mathbf{v}}}\rangle_{\mathcal{F}^{sys}(I)}$ such that

$$\mathrm{Tr}_{\mathcal{F}^{sys}(O\setminus\mathcal{S}_O)} \circ \hat{\Phi}_{\mathcal{F}}(|\psi^{\mathbf{v}}\rangle\langle\psi^{\mathbf{v}}|_{\mathcal{F}^{sys}(I)}) \neq \mathrm{Tr}_{\mathcal{F}^{sys}(O\setminus\mathcal{S}_O)} \circ \hat{\Phi}_{\mathcal{F}}(|\psi^{\tilde{\mathbf{v}}}\rangle\langle\psi^{\tilde{\mathbf{v}}}|_{\mathcal{F}^{sys}(I)}), \tag{23}$$

where $|\psi^{\mathbf{v}}\rangle = \bigotimes_i |\psi^{v_i}\rangle \in \mathcal{H}^{\mathbf{v}}_{\mathcal{F}^{sys}(I)}$ and

$$|\psi^{\tilde{\mathbf{v}}}\rangle_{\mathcal{F}^{sys}(I)} = \bigotimes_{I_i \in \mathcal{S}_I} |\psi^{v_i'}\rangle_{\mathcal{F}^{sys}(I_i)} \bigotimes_{I_j \in I\setminus\mathcal{S}_I} |\psi^{v_j}\rangle_{\mathcal{F}^{sys}(I_j)} \in \mathcal{H}^{\tilde{\mathbf{v}}}_{\mathcal{F}^{sys}(I)}$$

Using the above in Equation (23) shows that we have $\mathcal{F}^{sys}(\mathcal{S}_I)$ signals to $\mathcal{F}^{sys}(\mathcal{S}_O)$ in $\hat{\Phi}_{\mathcal{F}}$.

□

Theorem 3.6 *[Embedding arbitrary, cyclic signalling relations in spacetime] For every signalling structure $\mathcal{G}^{sig}$, there exists a fixed acyclic spacetime $\mathcal{T}$ and an embedding $\mathcal{E}$ of $\mathcal{G}^{sig}$ in a region causal structure $\mathcal{G}^R_{\mathcal{T}}$ of $\mathcal{T}$ that respects relativistic causality.*

Proof: In our framework, a signalling structure $\mathcal{G}^{sig}$ is in general a directed graph where the nodes belong to $\mathrm{Nodes}(\mathcal{G}^{sig}) := \mathrm{Powerset}[I \cup O]$, where I and O denote the set of all input and output systems in some network of CPTP maps. An embedding $\mathcal{E}$ of $\mathcal{G}^{sig}$ in a spacetime $\mathcal{T}$ corresponds to an assignment of spacetime regions to each system in $I \cup O$. This immediately implies an embedding for all systems in $\mathrm{Powerset}[I \cup O]$, for any subset $\mathcal{S}_{I/O}$ of $I \cup O$ the corresponding spacetime region is simply the union $\mathcal{P}^{\mathcal{S}_{I/O}} = \bigcup_{S \in \mathcal{S}} \mathcal{P}^S$ of all the spacetime regions assigned to the individual elements $S \in \mathcal{S}_{I/O}$ under the embedding $\mathcal{E}$. Thus, in order to establish the theorem statement, we first find an embedding for systems in $I \cup O$, such the signalling relations in $\mathcal{G}^{sig}$ over these subset of nodes respects relativistic causality. We will later see that this immediately implies an embedding of all the nodes of $\mathcal{G}^{sig}$ such that relativistic causality is still preserved. So for the purpose of the next few paragraphs, we will treat $\mathcal{G}^{sig}$ as a directed graph over the nodes $I \cup O$ and generalise the result at the end, and we will refer to it as $\mathcal{G}$ for short.

We first need set out some nomenclature. For every node N of a directed graph $\mathcal{G}$, the set $\mathrm{Par}(N) := \{N' \in \mathrm{Nodes}(\mathcal{G}) : N' \to N \in \mathcal{G}\}$ denotes the set of all parents of the node N, the set $\mathrm{Ch}(N) := \{N' \in \mathrm{Nodes}(\mathcal{G}) : N \to N' \in \mathcal{G}\}$ denotes the set of all children of the node N, the set $\mathrm{Anc}(N) := \{N' \in \mathrm{Nodes}(\mathcal{G}) : \exists \text{ directed path } N' \to ... \to N \in \mathcal{G}\}$ denotes the set of all ancestors of N in $\mathcal{G}$ and the set $\mathrm{Desc}(N) := \{N' \in \mathrm{Nodes}(\mathcal{G}) : \exists \text{ directed path } N \to ... \to N' \in \mathcal{G}\}$ denotes the set of all descendants of N in $\mathcal{G}$. Then the set of all nodes in $\mathcal{G}$ that are involved in at least one "loop" is defined as

$$\mathrm{Loop}(\mathcal{G}) = \{N \in \mathrm{Nodes}(\mathcal{G}) : N \in \mathrm{Anc}(N) \cap \mathrm{Desc}(N)\}.$$

That is, every node N that is both its own ancestor and its own descendant belongs to the set $\text{Loop}(\mathcal{G})$). Note that whenever $\mathcal{G}$ is a directed acyclic graph $\text{Loop}(\mathcal{G}) = \emptyset$.

We now show that we can fine-grain any signalling structure $\mathcal{G}$ into a directed acyclic graph $\mathcal{G}'$ by "splitting" nodes in $\text{Loop}(\mathcal{G})$, such that when the split nodes in $\mathcal{G}'$ are recombined, we get back the original graph $\mathcal{G}$ i.e., $\mathcal{G}$ is a coarse-graining of $\mathcal{G}'$. Once we have a directed acyclic graph $\mathcal{G}'$, we can always embed it in a partially ordered spacetime $\mathcal{T}$ though an embedding $\mathcal{E}'$ that respects relativistic causality. We can then coarse-grain the embedding $\mathcal{E}'$ of $\mathcal{G}'$ to an embedding $\mathcal{E}$ of the original structure $\mathcal{G}$ that also respects relativistic causality. We explicitly carry out these steps below.

We obtain the directed acyclic graph $\mathcal{G}'$ from the directed graph $\mathcal{G}$ as follows. If $\text{Loop}(\mathcal{G}) = \emptyset$, set $\mathcal{G}' = \mathcal{G}$. If $\text{Loop}(\mathcal{G}) \neq \emptyset$, then we split every node $N \in \text{Loop}(\mathcal{G})$ into two nodes N_1 and N_2 such that $\text{Par}(N_1) = \text{Par}(N)\backslash\text{Loop}(\mathcal{G})$ and $\text{Ch}(N_1) = \text{Ch}(N)$ and $\text{Par}(N_2) = \text{Par}(N)$ and $\text{Ch}(N_2) = \text{Ch}(N)\backslash\text{Loop}(\mathcal{G})$ i.e., N_1 contains no incoming arrows from nodes in $\text{Loop}(\mathcal{G})$ but all the same outgoing arrows as N does in $\mathcal{G}$ while N_2 contains no outgoing arrows to nodes in $\text{Loop}(\mathcal{G})$ but all the same incoming arrows as N does in $\mathcal{G}$. Nodes that do not belong to $\text{Loop}(\mathcal{G})$, and other edges not featuring in the above construction remain unaffected. This fully defines $\mathcal{G}'$. Then two things are clear. Firstly that when we recombine N_1 and N_2 back into a single node N, for each pair of split nodes and without altering the edge structure, we recover the original graph $\mathcal{G}$ since $\text{Par}(N_1 \cup N_2) = \text{Par}(N)$ and $\text{Ch}(N_1 \cup N_2) = \text{Ch}(N)$ for every loop node, and the non-loop nodes were not split or altered in going from $\mathcal{G}$ to $\mathcal{G}'$. Secondly, $\mathcal{G}'$ is an acyclic graph since every node that was part of a loop in $\mathcal{G}$ is now split such that no single node contains both incoming and outgoing arrows from another node in a loop.

In our case, each node N corresponds to a quantum Hilbert spaces $\mathcal{H}_N$, and splitting a node corresponds to creating two copies of the Hilbert space $\mathcal{H}_{N_1} \cong \mathcal{H}_N$ and $\mathcal{H}_{N_2} \cong \mathcal{H}_N$. Note that this can always be done also at the level of the network of CPTP maps that gives rise to the signalling relations. Going back to our simple example with $\hat{\Phi}_2 : I_2 \mapsto O_2$ sequentially composed after $\hat{\Phi}_1 : I_1 \mapsto O_1$, we have the signalling structure $\{I_1 \to O_1, O_1 \to I_2, I_2 \to O_2\}$. We can for instance split the node O_1 into two nodes O_1^1 and O_1^2 giving the signalling structure $\{I_1 \to O_1^1, O_1^2 \to I_2, I_2 \to O_2\}$ where we have the map $\hat{\Phi}_1$ acting between I_1 and O_1^1, while O_1^2 is directly connected to I_2 (through an identity channel) and I_2 to O_2 through the map $\hat{\Phi}_2$ as before.

Getting back to the main proof, since $\mathcal{G}'$ is a directed acyclic graph, there exists an embedding $\mathcal{E}' : \text{Nodes}(\mathcal{G}') \mapsto \mathcal{T}$ of $\mathcal{G}'$in a partially ordered set $\mathcal{T}$ (associated with the order relation $\prec$) such that $N_i \to N_j$ in $\mathcal{G} \Leftrightarrow \mathcal{E}'(N_i) \prec \mathcal{E}'(N_j)$. By virtue of being a partial order, $\mathcal{T}$ satisfies our minimal definition of spacetime structure, according to Definition 3.1. Then the required embedding $\mathcal{E}$ of $\mathcal{G}$ in the spacetime $\mathcal{T}$ simply associates two spacetime locations with each node $N \in \text{Loop}(\mathcal{G})$, the two locations being precisely those assigned by $\mathcal{E}'$ to each of the split nodes i.e., $\mathcal{E}(N) := \{\mathcal{E}'(N_1), \mathcal{E}'(N_2)\}$. For all nodes $N \notin \text{Loop}(\mathcal{G})$, $\mathcal{E}(N) = \mathcal{E}'(N)$ noting that these nodes never got split. Then it is clear that the embedding $\mathcal{E}$ of $\mathcal{G}$ respects relativistic causality whenever the embedding $\mathcal{E}'$ of $\mathcal{G}'$ respects relativistic causality, which it does by construction.

We now explain how the proof generalises to case where $\text{Nodes}(\mathcal{G}) = \text{Powerset}[I \cup O]$. For this, note that our above proof covers all cases where $\mathcal{G}$ has the property that $\mathcal{S}_1 \to \mathcal{S}_2$ for two subsets $\mathcal{S}_1, \mathcal{S}_2$ of $I \cup O$, then there exists $S_1 \in \mathcal{S}_1$ and $S_2 \in \mathcal{S}_2$ such that $S_1 \to S_2$. However, suppose that we have a signalling relation $\mathcal{S}_1 \to \mathcal{S}_2$ in $\mathcal{G}$ such that there is no signalling relation between individual elements of these two sets. The relativistic causality condition implied by this signalling relation on the corresponding spacetime embedding is that $\mathcal{P}^{\mathcal{S}_1} \xrightarrow{R} \mathcal{P}^{\mathcal{S}_2}$ (cf. Definition 3.5). Since $\mathcal{P}^{\mathcal{S}_1} = \bigcup_{S_1 \in \mathcal{S}_1} \mathcal{P}^{S_1}$ (and similarly for $\mathcal{S}_2$), this is equivalent to saying that there exists $S_1 \in \mathcal{S}_1$ and $S_2 \in \mathcal{S}_2$ such that the corresponding spacetime regions satisfy $\mathcal{P}^{S_1} \xrightarrow{R} \mathcal{P}^{S_2}$. In other words, the relativistic causality constraints on the spacetime embedding of $\mathcal{G}$ are the same irrespective of whether or not $\mathcal{G}$ satisfies the aforementioned property. Thus the above proof also applies to establish the theorem statement for signalling structures $\mathcal{G}$ not satisfying this property i.e., hence it applies to all signalling structures.

Finally, we note that according to Definition 3.1 any partially ordered set corresponds to a spacetime. This rather minimal definition allows us to derive general results that only depend on the order relation between spacetime points and does not require the spacetime to have any further symmetries, or a smooth differentiable structure. However, under this minimal definition, one might regard two different partially ordered sets $\mathcal{T}$ and $\mathcal{T}'$ as two different "spacetimes". On the other hand, if we consider the more standard method of modelling spacetime as a differentiable manifold $\mathcal{M}$, as done in relativistic physics, we could

sample different sets of points on the same manifold to generate different partially ordered sets[14] $\mathcal{T}$ and $\mathcal{T}'$ from the same spacetime. If we model spacetime as a globally hyperbolic manifold that ensures the absence of closed timelike curves, then the statement of the present theorem would instead become "For every signalling struc- ture $\mathcal{G}^{sig}$ and every globally hyperbolic manifold $\mathcal{M}$, there exists an embedding $\mathcal{E}$ of $\mathcal{G}^{sig}$ in a region causal structure $\mathcal{G}^R_{\mathcal{M}}$ of $\mathcal{M}$ that respects relativistic causality, where each node of $\mathcal{G}^R_{\mathcal{M}}$ is a finite set of points in $\mathcal{M}$." This can be shown as follows. If $\mathcal{G}$ is a directed acyclic graph, then it can be embedded in any globally hyperbolic manifold $\mathcal{M}$ through an embedding $\mathcal{E}$: Nodes($\mathcal{G}$) $\mapsto \mathcal{M}$ that assigns a point in $\mathcal{M}$ to each node of $\mathcal{G}$ (see [50] for an explicit construction of such an embedding for the acyclic case). This is because the graph has a finite number of nodes and we can always sample a suitable set of points in the manifold having the required order relations. One can apply this embedding to the acyclic graph $\mathcal{G}'$ constructed in the proof above, this would define the embedding $\mathcal{E}'$: Nodes($\mathcal{G}$) $\mapsto \mathcal{M}$. The rest of the proof will be the same as the above case for partially ordered sets $\mathcal{T}$. □

Lemma 3.9 *[Fine-graining to an acyclic signalling structure] Every network of CPTP maps that admits an implementation in a fixed spacetime $\mathcal{T}$ that does not violate relativistic causality in that spacetime admits a fine-graining that has a definite acyclic signalling structure, whose edges $\rightarrow$ align with the partial order relation $\prec$ of the spacetime.*

Proof: A spacetime implementation of any map can be maximally fine-grained in terms of its elemental subsystems which are quantum systems S associated with a spacetime location $P^S \in \mathcal{T}$. Relativistic causality (cf. Definition 3.5) then requires that whenever $S_1 \rightarrow S_2$ for two elemental subsystems, then $P^{S_1} \prec P^{S_2}$ must hold. Then the signalling structure of any network of such maps that satisfy relativistic causality cannot contain a directed cycle of signalling relations between elemental systems as $\prec$ is a partial order and therefore the signalling structure over the elemental subsystems of any network of spacetime implemented maps that satisfy relativistic causality must be acyclic. □

Lemma 5.2 *[Probabilities from composition] For every process map $\hat{W}$, the joint probabilities obtained through the complete composition $\hat{P}_{\hat{W},\mathcal{M}}$ as in Equation (10) are equivalent to those obtained in the process matrix framework through Equation (4).*

Proof: We first construct the complete composition $\hat{P}_{\hat{W},\mathcal{M}}$ step by step to make explicit how it can be obtained from the process and the local operations. Denoting the quantum input and output spaces of the local operations $\mathcal{M}^{A^k}$ with a bar on top ($\bar{A}^k_I$ and $\bar{A}^k_O$) to distinguish them from the corresponding input and output spaces of $\hat{W}$ before composition, we first compose these $N+1$ CPTP maps in parallel to obtain the CPTP map

$$\mathcal{M}^{A^1} \otimes ... \otimes \mathcal{M}^{A^N} \otimes \hat{W}.$$

This map has the $3N$ input systems $\{\bar{A}^1_I, A^1_s, ..., \bar{A}^N_I, A^N_s, A^1_O, ..., A^N_O\}$ and the $3N$ output systems $\{\bar{A}^1_O, A^1_o, ..., \bar{A}^N_O, A^N_o, A^1_I, ..., A^N_I\}$. We now loop each of the output systems $\bar{A}^k_O$ to the corresponding inputs A^k_O, and similarly the output systems A^k_I get looped back to the inputs $\bar{A}^k_I$ and this is possible since by construction, the "barred" systems are copies of their "unbarred" versions with the same state-spaces. Performing this loop composition yields a map $\hat{P}$ with the uncontracted input and output systems, namely the inputs $\{A^1_s, ..., A^N_s\}$ and outputs $\{A^1_o, ..., A^N_o\}$, which as we will now show, encodes the joint probabilities of possible measurements implemented by the local maps. In the following, for brevity, we detail the proof for the bipartite case. However, the proof readily generalises to the N party case.In the bipartite case, taking the parties to be A and B with local settings associated with input systems A_s, B_s and outcomes associated with output systems A_o, B_o, the parallel composition yields the map $\mathcal{M}^A \otimes \mathcal{M}^B \otimes \hat{W}$ with input systems $\{\bar{A}_I, A_s, \bar{B}_I, B_s, A_O, B_O\}$ and output systems $\{\bar{A}_O, A_o, \bar{B}_O, B_o, A_I, B_I\}$. Applying the loop formula of Equation (1) to describe the final map $\hat{P}_{\hat{W},\mathcal{M}}$ (with classical inputs A_s and B_s and classical outputs A_o ad B_o), we obtain

[14] That is, if the manifold is globally hyperbolic, in more exotic spacetimes with closed timelike curves, we can also obtain pre-ordered sets from sampling suitable points.

$$
\begin{aligned}
&\hat{P}_{\hat{W},\mathcal{M}}(|a\rangle\langle a|_{A_s}\otimes|b\rangle\langle b|_{B_s})\\
&=\sum_{k\ldots r}\langle k|_{A_I}\langle m|_{\bar{A}_O}\langle o|_{B_I}\langle q|_{\bar{B}_O}\left(\mathcal{M}_a^A\big(|k\rangle\langle l|_{\bar{A}_I}\big)\otimes\mathcal{M}_b^B\big(|o\rangle\langle p|_{\bar{B}_I}\big)\otimes\hat{W}\big(|mq\rangle\langle nr|_{A_OB_O}\big)\right)|l\rangle_{A_I}|n\rangle_{\bar{A}_O}|p\rangle_{B_I}|r\rangle_{\bar{B}_O}\\
&=\sum_{k\ldots r}\langle k|_{A_I}\langle n|_{\bar{A}_O}\langle o|_{B_I}\langle r|_{\bar{B}_O}\left(\big[\mathcal{M}_a^A\big(|k\rangle\langle l|_{\bar{A}_I}\big)\big]^T\otimes\big[\mathcal{M}_b^B\big(|o\rangle\langle p|_{\bar{B}_I}\big)\big]^T\otimes\hat{W}\big(|mq\rangle\langle nr|\big)\right)|l\rangle_{A_I}|m\rangle_{\bar{A}_O}|p\rangle_{B_I}|q\rangle_{\bar{B}_O},
\end{aligned}
$$

where we have used the notation $\mathcal{M}_a^A\big(|k\rangle\langle l|_{\bar{A}_I}\big):=\mathcal{M}^A\big(|a\rangle\langle a|_{A_s}\otimes|k\rangle\langle l|_{\bar{A}_I}\big)$ (and similarly for B's operation). Denoting the factor $\left(\big[\mathcal{M}_a^A\big(|k\rangle\langle l|_{\bar{A}_I}\big)\big]^T\otimes\big[\mathcal{M}_b^B\big(|o\rangle\langle p|_{\bar{B}_I}\big)\big]^T\otimes\hat{W}\big(|m\rangle\langle n|_{A_O}\otimes|q\rangle\langle r|_{B_O}\big)\right)$ by (...), introducing factors of the identity $\mathcal{I}=\sum_j|j\rangle\langle j|$, and then rearranging the resultant inner products we have

$$
\begin{aligned}
&\hat{P}_{\hat{W},\mathcal{M}}(|a\rangle\langle a|_{A_s}\otimes|b\rangle\langle b|_{B_s})\\
&=\sum_{ijst}\sum_{k\ldots r}\langle k|_{A_I}\langle n|j\rangle\langle j|_{\bar{A}_O}\langle o|_{B_I}\langle r|s\rangle\langle s|_{\bar{B}_O}\left(\ldots\right)|i\rangle_{A_I}\langle i|l\rangle_{A_I}|m\rangle_{\bar{A}_O}|t\rangle_{B_I}\langle t|p\rangle_{B_I}|q\rangle_{\bar{B}_O}\\
&=\sum_{ijst}\sum_{k\ldots r}\langle i|_{A_I}|l\rangle\langle k|_{A_I}\langle j|_{\bar{A}_O}\langle t|_{B_I}|p\rangle\langle o|_{B_I}\langle s|_{\bar{B}_O}\left(\ldots\right)|i\rangle_{A_I}|m\rangle\langle n|_{\bar{A}_O}|j\rangle_{\bar{A}_O}|t\rangle_{B_I}|q\rangle\langle r|_{\bar{B}_O}|s\rangle_{\bar{B}_O}\\
&=\mathrm{tr}_{A_I\bar{A}_OB_I\bar{B}_O}\left[\sum_{k\ldots r}|l\rangle\langle k|_{A_I}\otimes|p\rangle\langle o|_{B_I}\otimes\left(\ldots\right)\otimes|m\rangle\langle n|_{\bar{A}_O}\otimes|q\rangle\langle r|_{\bar{B}_O}\right]\\
&=\mathrm{tr}_{A_I\bar{A}_OB_I\bar{B}_O}\left[\left(\sum_{kl}|l\rangle\langle k|_{A_I}\otimes\big[\mathcal{M}_a^A\big(|k\rangle\langle l|_{\bar{A}_I}\big)\big]^T\right)\otimes\left(\sum_{op}|p\rangle\langle o|_{B_I}\otimes\big[\mathcal{M}_b^B\big(|o\rangle\langle p|_{\bar{B}_I}\big)\big]^T\right)\right.\\
&\qquad\left.\otimes\left(\sum_{mnqr}|mq\rangle\langle nr|_{\bar{A}_O\bar{B}_O}\otimes\hat{W}\big(|mq\rangle\langle nr|_{A_OB_O}\big)\right)\right]
\end{aligned}
$$

Now, we wish to calculate the probability that the output of $\hat{P}_{\hat{W},\mathcal{M}}(|a\rangle\langle a|_{A_s}\otimes|b\rangle\langle b|_{B_s})$ is $|x\rangle_{A_o}\otimes|y\rangle_{B_o}$ i.e., the outcome x obtained by Alice upon measuring the setting a and outcome y for Bob upon measuring the setting b. This is simply $P(xy|ab)=\mathrm{tr}\left[\Big(|x\rangle\langle x|_{A_o}\otimes|y\rangle\langle y|_{B_o}\Big)\Big(\hat{P}_{\hat{W},\mathcal{M}}(|a\rangle\langle a|_{A_s}\otimes|b\rangle\langle b|_{B_s})\Big)\right]$, which is the probability that the projection of the output space into $|x\rangle\langle x|_{A_o}\otimes|y\rangle\langle y|_{B_o}$ succeeds. Combining this with the above equation for $\hat{P}(|a\rangle\langle a|_{A_s}\otimes|b\rangle\langle b|_{B_s})$ and using Equation (8) to absorb the outcome projectors into the definition of the maps $\mathcal{M}_{x|a}^A$ and $\mathcal{M}_{y|b}^B$, we immediately obtain the required result.

$$
P(xy|ab)=\mathrm{tr}\left[\Big(M_{x|a}^{A_IA_O}\otimes M_{y|b}^{B_IB_O}\Big)W\right],\tag{24}
$$

where $W=\mathcal{I}\otimes\hat{W}|\mathbb{1}\rangle\rangle\langle\langle\mathbb{1}|$ is the process matrix and $M_{x|a}^{A_IA_O}=\Big[\mathcal{I}\otimes\mathcal{M}_{x|a}^A\Big(|\mathbb{1}\rangle\rangle\langle\langle\mathbb{1}|\Big)\Big]^T$ is the Choi representation of the local map $\mathcal{M}_{x|a}^A$ (and similarly for B's operation) as defined in Section 4.

□

Lemma 5.3 *[Partial composition and reduced process] Consider an N-partite process map $\hat{W}$ and the local operations $\{\mathcal{M}_{a_k}^{A^k}\}_{k=1}^l$ for the first $l<N$ parties for a fixed set of settings $\{a_k\}_{k=1}^l$. Then the partial composition $\mathfrak{C}(\hat{W},\{\mathcal{M}_{a_k}^{A^k}\}_{k=1}^l)$ corresponds to a CPTP map whose Choi representation is the reduced process matrix $\bar{W}(M_{a_1}^{A_I^1A_O^1},\ldots,M_{a_l}^{A_I^lA_O^l})$.*

Proof: The proof method is very similar to that of Lemma 5.2 but we provide it here for completeness and follow the same notation as the previous proof. Again, we restrict to the bipartite case for simplicity

and brevity but the proof easily generalises to the multipartite case. Consider a bipartite process map $\hat{W}$ with input systems A_O, B_O and output systems A_I, B_I. Let $\mathcal{M}_b^B : \bar{B}_I \mapsto \bar{B}_O$ be a local operation of the party B labelled by the setting b, where the barred systems are isomorphic to the corresponding unbarred ones. We now show that for every choice of setting b, the Choi representation of the corresponding partial composition $\mathfrak{C}(\hat{W}, \mathcal{M}_b^B)$ is the reduced process matrix $\bar{W}(M_b^{B_I \bar{B}_O})$ (cf. Equation (11)). Denoting $\mathfrak{C}(W, \mathcal{M}_b^B)$ as $\hat{C}_{W,b}$ for short and noting that $\hat{C}_{W,b}$ is a map from A_O to A_I, the Choi representation of $\hat{C}_{W,b}$ is given as

$$C_{W,b} = \sum_{m,n} |m\rangle\langle n|_{A_O} \otimes \hat{C}_{W,b}(|m\rangle\langle n|_{A_O}).$$

Using the composition operation, we can write $\hat{C}_{W,b}(|m\rangle\langle n|_{A_O})$ as

$$\begin{aligned}\hat{C}_{W,b}(|m\rangle\langle n|_{A_O}) &= \sum_{o,p,q,r} \langle o|_{B_I}\langle q|_{\bar{B}_O}\left(\mathcal{M}_b^B \otimes \hat{W}\right)\left(|o\rangle\langle p|_{\bar{B}_I} \otimes |mq\rangle\langle nr|_{A_O B_O}\right)|p\rangle_{B_I}|r\rangle_{\bar{B}_O} \\ &= \sum_{o,p,q,r} \langle q|_{\bar{B}_O}\mathcal{M}_b^B\left(|o\rangle\langle p|_{\bar{B}_I}\right)|r\rangle_{\bar{B}_O} \otimes \langle o|_{B_I}\hat{W}\left(|mq\rangle\langle nr|_{A_O B_O}\right)|p\rangle_{B_I}\end{aligned}$$

Plugging this back into the Choi representation, inserting factors of the identity and rearranging, we have

$$\begin{aligned}C_{W,b} &= \sum_{o,p,q,r,m,n} |m\rangle\langle n|_{A_O} \otimes \langle q|_{\bar{B}_O}\mathcal{M}_b^B\left(|o\rangle\langle p|_{\bar{B}_I}\right)|r\rangle_{\bar{B}_O} \otimes \langle o|_{B_I}\hat{W}\left(|mq\rangle\langle nr|_{A_O B_O}\right)|p\rangle_{B_I} \\ &= \sum_{o,p,q,r,m,n} |m\rangle\langle n|_{A_O} \otimes \langle r|_{\bar{B}_O}\left[\mathcal{M}_b^B\left(|o\rangle\langle p|_{\bar{B}_I}\right)\right]^T|q\rangle_{\bar{B}_O} \otimes \langle o|_{B_I}\hat{W}\left(|mq\rangle\langle nr|_{A_O B_O}\right)|p\rangle_{B_I} \\ &= \sum_{o,p,q,r,m,n,j,k} |m\rangle\langle n|_{A_O} \otimes \langle r|_{\bar{B}_O}|j\rangle\langle j|_{\bar{B}_O}\left[\mathcal{M}_b^B\left(|o\rangle\langle p|_{\bar{B}_I}\right)\right]^T|q\rangle_{\bar{B}_O} \otimes \langle o|_{B_I}|k\rangle\langle k|_{B_I}\hat{W}\left(|mq\rangle\langle nr|_{A_O B_O}\right)|p\rangle_{B_I} \\ &= \sum_{o,p,q,r,m,n,j,k} |m\rangle\langle n|_{A_O} \otimes \langle j|_{\bar{B}_O}\left[\mathcal{M}_b^B\left(|o\rangle\langle p|_{\bar{B}_I}\right)\right]^T|q\rangle\langle r|_{\bar{B}_O}|j\rangle_{\bar{B}_O} \otimes \langle k|_{B_I}\hat{W}\left(|mq\rangle\langle nr|_{A_O B_O}\right)|p\rangle\langle o|_{B_I}|k\rangle \\ &= \sum_{j,k}\langle jk|_{B_I\bar{B}_O}\left(\mathbb{1}^{A_I A_O} \otimes \sum_{o,p}|p\rangle\langle o| \otimes \left[\mathcal{M}_b^B\left(|o\rangle\langle p|\right)\right]^T\right)\left(\sum_{m,q,n,r}|mq\rangle\langle nr| \otimes \hat{W}\left(|mq\rangle\langle nr|\right)\right)|jk\rangle_{B_I\bar{B}_O} \\ &= \mathrm{Tr}_{B_I B_O}\left(\left(\mathbb{1}^{A_I A_O} \otimes M_b^{B_I B_O}\right).W\right) = \bar{W}(M_b^{B_I B_O}),\end{aligned}$$

where we have used the definition of the Choi representations of the local operations and the process map (see Section 4) in the last line, along with that of the reduced process matrix (Equation (11)). This completes the proof. □

Theorem 5.4 *[Equivalence of two notions of signalling] A_O^i does not signal to $A_I^{\mathcal{S}}$ in $\mathfrak{C}(\hat{W}, \{\mathcal{M}^{A^k}\}_{k=1,k\notin i\cup\mathcal{S}}^{k=N})$ if and only if the set of outcomes $x_{\mathcal{S}} := \{x_S\}_{S\in\mathcal{S}}$ of the parties in A^S do not depend on the setting a_i of the party A^i i.e., the corresponding joint probability distribution satisfies Equation (5).*

Proof: Sufficiency: Here, we establish that no signalling in the partial composition implies no signalling in joint probabilities. We first prove the result for the bipartite case where $A^{\mathcal{S}}$ reduces to a single party and the corresponding partial composition $\mathfrak{C}(\hat{W}, \{\mathcal{M}^{A^k}\}_{k=1,k\notin i\cup\mathcal{S}}^{k=N})$ reduces to the bipartite process map $\hat{W}$ itself. The proof immediately generalises to the multipartite case with arbitrary $A^{\mathcal{S}}$, with $\mathfrak{C}(\hat{W}, \{\mathcal{M}^{A^k}\}_{k=1,k\notin i\cup\mathcal{S}}^{k=N})$ taking the place of $\hat{W}$.

For the bipartite case with two parties A and B, the statement we need to prove is that A_O does not signal to B_I in $\hat{W}$ implies that for all choices of settings a and b on the input systems A_s and B_s of the local maps $\mathcal{M}^A$ and $\mathcal{M}^B$ and corresponding outcomes x and y obtained on the output systems A_o and B_o of these maps, the joint probability distribution (cf. Lemma 5.2) satisfies $P(y|ab) = P(y|b)$. We establish

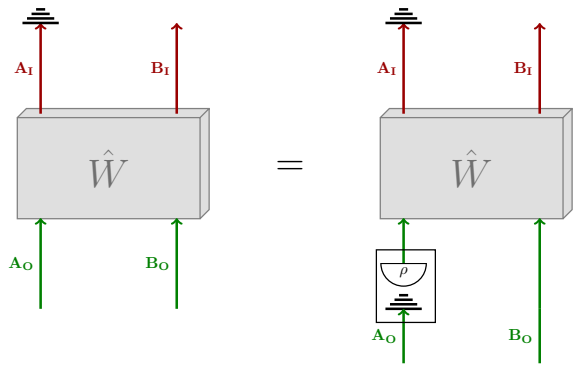

Figure 13: Diagrammatic representation of Equation (25). The equality holds for all states ρ on A_O.

below and show, and then explain how the argument readily generalises to the multipartite case where the signalling relation under consideration is with respect to the partial composition $\mathfrak{C}(\hat{W}, \{\mathcal{M}^{A^k}\}_{k=1,k\neq i,j}^{k=N})$.

Writing out the condition that A_O does not signal to B_I in $\hat{W}$ explicitly, we have the following, which is illustrated diagrammatically in Figure 13

$$\mathrm{Tr}_{A_I} \circ \hat{W}(\sigma_{A_O B_O}) = \mathrm{Tr}_{A_I} \circ \hat{W}(\rho_{A_O} \otimes \mathrm{Tr}_{A_O}(\sigma_{A_O B_O})), \qquad \forall \sigma_{A_O B_O}, \rho_{A_O}. \tag{25}$$

This in turn implies that $\forall \sigma_{A_O B_O}, \rho_{A_O}, \mathcal{M}^B, b$,

$$\mathcal{M}^B \circ \mathrm{Tr}_{A_I} \circ \hat{W}(\sigma_{A_O B_O} \otimes |b\rangle\langle b|_{B_s}) = \mathcal{M}^B \circ \mathrm{Tr}_{A_I} \circ \hat{W}(\rho_{A_O} \otimes \mathrm{Tr}_{A_O}(\sigma_{A_O B_O} \otimes |b\rangle\langle b|_{B_s})),$$

where $\mathcal{M}^B$ is a local map of B with inputs B_I and B_s and outputs B_O and B_o. Since, the above statement holds for all $\sigma_{A_O B_O}$, we could choose $\sigma_{A_O B_O}$ to be $\sigma_{A_O B_O} = (\mathcal{M}^A \otimes \mathbb{1}^{B_O})(\tilde{\sigma}_{A_I B_O} \otimes |a\rangle\langle a|_{A_s})$ for some map $\mathcal{M}^A$ (with inputs A_I and A_s and outputs A_O and A_o) and state $\tilde{\sigma}_{A_I B_O} \otimes |a\rangle\langle a|_{A_s}$. Then noting that the order of $\mathcal{M}_B$ and Tr_{A_I} does not matter as they act on different subsystems, we have $\forall \tilde{\sigma}_{A_I B_O}, \rho_{A_O}, \mathcal{M}^B, \mathcal{M}^A, b, a$

$$\begin{aligned} &\mathrm{Tr}_{A_o A_I B_O} \circ \mathcal{M}^B \circ \hat{W}\left((\mathcal{M}^A \otimes \mathbb{1}^{B_O})(\tilde{\sigma}_{A_I B_O} \otimes |a\rangle\langle a|_{A_s} \otimes |b\rangle\langle b|_{B_s}) \right) \\ =&\mathrm{Tr}_{A_o A_I B_O} \circ \mathcal{M}^B \circ \hat{W}\left(\rho_{A_O} \otimes \mathrm{Tr}_{A_O}\left((\mathcal{M}^A \otimes \mathbb{1}^{B_O})(\tilde{\sigma}_{A_I B_O} \otimes |a\rangle\langle a|_{A_s} \otimes |b\rangle\langle b|_{B_s}) \right) \right). \end{aligned} \tag{26}$$

This condition is illustrated in Figure 14. Note that the right hand side of the above equation is equal to the expression $\mathrm{Tr}_{B_O A_I} \circ \mathcal{M}^B \circ \hat{W}\left(\rho_{A_O} \otimes \mathrm{Tr}_{A_o A_O}\left((\mathcal{M}^A \otimes \mathbb{1}^{B_O})(\tilde{\sigma}_{A_I B_O} \otimes |a\rangle\langle a|_{A_s} \otimes |b\rangle\langle b|_{B_s}) \right) \right)$, this equality is much more apparent from the figure. Then writing out Equation (26) with $|a'\rangle\langle a'|_{A_s}$ instead of $|a\rangle\langle a|_{A_s}$ and using the fact that for all a, a',

$$\mathrm{Tr}_{A_O A_o}\left((\mathcal{M}^A \otimes \mathbb{1}^{B_O})(\tilde{\sigma}_{A_I B_O} \otimes |a\rangle\langle a|_{A_s} \otimes |b\rangle\langle b|_{B_s}) \right) = \mathrm{Tr}_{A_O A_o}\left((\mathcal{M}^A \otimes \mathbb{1}^{B_O})(\tilde{\sigma}_{A_I B_O} \otimes |a'\rangle\langle a'|_{A_s} \otimes |b\rangle\langle b|_{B_s}) \right),$$

we have $\forall \tilde{\sigma}_{A_I B_O}, \rho_{A_O}, \mathcal{M}_b^B, \mathcal{M}^A$, a, a', b

$$\begin{aligned} &\mathrm{Tr}_{A_o A_I B_O} \circ \mathcal{M}^B \circ \hat{W}\left((\mathcal{M}^A \otimes \mathbb{1}^{B_O})(\tilde{\sigma}_{A_I B_O} \otimes |a\rangle\langle a|_{A_s} \otimes |b\rangle\langle b|_{B_s}) \right) \\ =&\mathrm{Tr}_{A_o A_I B_O} \circ \mathcal{M}^B \circ \hat{W}\left((\mathcal{M}^A \otimes \mathbb{1}^{B_O})(\tilde{\sigma}_{A_I B_O} \otimes |a'\rangle\langle a'|_{A_s} \otimes |b\rangle\langle b|_{B_s}) \right). \end{aligned} \tag{27}$$

The above is illustrated in Figure 15 and is simply the condition that A_s does not signal to B_o in the map $\mathcal{M}^B \circ \hat{W} \circ \mathcal{M}^A$. The sequential composition $\mathcal{M}^B \circ \hat{W} \circ \mathcal{M}^A$ has the input systems A_I, A_s,

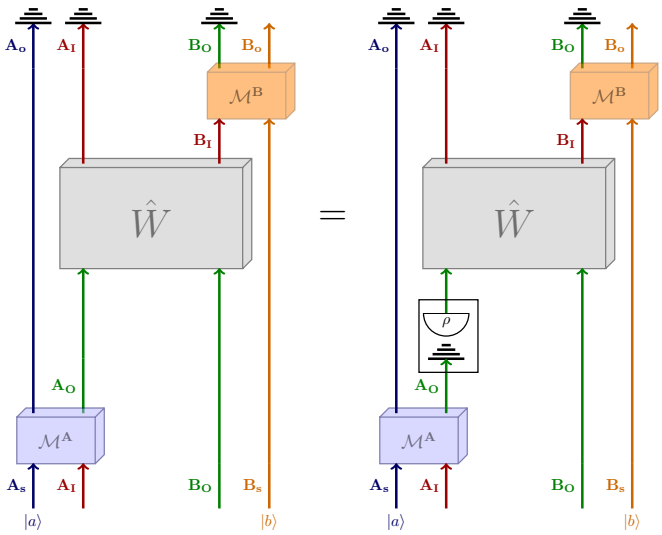

Figure 14: Diagrammatic representation of Equation (26). The equality holds for all states ρ on A_O, $|a\rangle$ on A_s, $|b\rangle$ on B_s and for all maps $\mathcal{M}_A$ and $\mathcal{M}_B$.

B_O and B_s and the output systems A_I, A_o B_O and B_o. In order to obtain the complete composition $\mathfrak{C}(\hat{W}, \mathcal{M}^A, \mathcal{M}^B)$ from the sequential composition $\mathcal{M}^B \circ \hat{W} \circ \mathcal{M}^A$, we must connect the systems with the same label through loop composition, this involves two loop compositions, one connecting the input-output pair labelled A_I and one for the pair labelled B_O. We now show that the non-signalling relation $A_s \not\to B_o$ is preserved under these two compositions. The first composition connecting the input-output pair A_I in $\mathcal{M}^B \circ \hat{W} \circ \mathcal{M}^A$ yields to map $\mathcal{M}^B \circ \mathfrak{C}(\hat{W}, \mathcal{M}^A)$, then connecting the input-output pair B_O in $\mathcal{M}^B \circ \mathfrak{C}(\hat{W}, \mathcal{M}^A)$ yields the complete composition $\hat{P}_{\hat{W},\mathcal{M}} := \mathfrak{C}(\hat{W}, \mathcal{M}^A, \mathcal{M}^B)$.

The fact that the non-signalling relation $A_s \not\to B_o$ is preserved under the first composition is apparent from the condition established in Figure 14. From the right hand side of the figure, we can see that irrespective of whether the output A_I is connected back to the input A_I or simply discarded, the output on B_o is obtained by first applying $\hat{W}$ to $\rho_{A_O} \otimes \sigma_{B_O}$, then applying $\mathcal{M}^B$ with the setting $|b\rangle\langle b|_{B_s}$ and tracing out B_O. This means that the output on B_o is independent of the choice of classical input on A_s in both $\mathcal{M}^B \circ \hat{W} \circ \mathcal{M}^A$ and $\mathcal{M}^B \circ \mathfrak{C}(\hat{W}, \mathcal{M}^A)$, $\forall \mathcal{M}^A$, $\mathcal{M}^B$.

The fact that it is also preserved under the second composition, i.e., in going from $\mathcal{M}^B \circ \mathfrak{C}(\hat{W}, \mathcal{M}^A)$ to $\hat{P}_{\hat{W},\mathcal{M}} := \mathfrak{C}(\hat{W}, \mathcal{M}^A, \mathcal{M}^B)$, can be shown as follows. From $A_O \not\to B_I$ in $\hat{W}$, we can establish that $A_s \not\to B_O$ in $\mathcal{M}^B \circ \hat{W} \circ \mathcal{M}^A$, $\forall \mathcal{M}^A$ and $\mathcal{M}^B$, in exactly the same manner as we established that $A_O \not\to B_I$ in $\hat{W}$ implies $A_s \not\to B_o$ in $\mathcal{M}^B \circ \hat{W} \circ \mathcal{M}^A$, $\forall \mathcal{M}^A$ and $\mathcal{M}^B$ (cf. Figure 16). We can then apply the same argument of the previous paragraph to conclude that $A_s \not\to B_O$ also in $\mathcal{M}^B \circ \mathfrak{C}(\hat{W}, \mathcal{M}^A)$, $\forall \mathcal{M}^A$, $\mathcal{M}^B$. Since the final loop composition of interest connects the output B_O (which we have established to be independent of inputs on A_s) to the corresponding input of the same name, we know that after the composition, the set of allowed states that can flow on the output system B_O of $\mathcal{M}^B$ is independent of the choice of state $|a\rangle$ on A_s. These set of allowed states can be no larger than the set of possible states on B_O that we considered before the composition for establishing the non-signalling relation $A_s \not\to B_o$. Therefore $A_s \not\to B_o$ before the composition implies $A_s \not\to B_o$ after this loop composition.

In the above, we have shown that $A_O \not\to B_I$ in $\hat{W}$ implies that $A_s \not\to B_o$ in the complete composition $\hat{P}_{\hat{W},\mathcal{M}} := \mathfrak{C}(\hat{W}, \mathcal{M}^A, \mathcal{M}^B)$ for all choices of local operations $\mathcal{M}^A$ and $\mathcal{M}^B$, this is illustrated in Figure 16. Since $\mathfrak{C}(\hat{W}, \mathcal{M}^A, \mathcal{M}^B)$ only has classical input system A_s and B_s and classical output systems, A_o and B_o, it is a classical channel between these systems and the fact that A_s does not signal to B_o immediately implies that $\sum_x P(xy|ab) = \sum_x P(xy|a)$ or $P(y|ab) = P(y|a)$ as required.

The proof for the multi-partite case with arbitrary $A^{\mathcal{S}}$ proceeds in the same manner as the bipartite case shown above. To see this, first note that the proof for the multipartite case with $A^i \cup A^{\mathcal{S}} = \mathcal{A}$ is identical to the bipartite case, where A^i plays the role of the party A and the set of parties $A^{\mathcal{S}}$ jointly play the role of the party B. In such cases, the partial composition $\mathfrak{C}(\hat{W}, \{\mathcal{M}^{A^k}\}_{k=1,k\notin i\cup\mathcal{S}}^{k=N})$ reduces to $\hat{W}$ whose input systems are A_O^i, $A_O^{\mathcal{S}}$ (set of quantum outputs of parties in $A^{\mathcal{S}}$) and output systems are

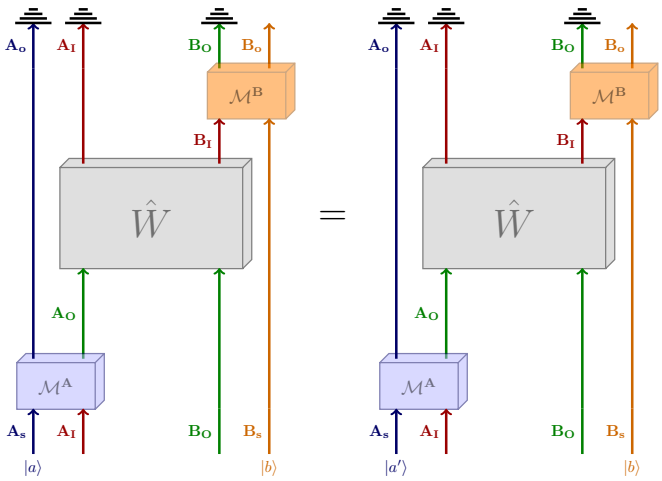

Figure 15: Diagrammatic representation of Equation (27). The equality holds for all states $|a\rangle$, $|a'\rangle$ on A_s, $|b\rangle$ on B_s and for all maps $\mathcal{M}_A$ and $\mathcal{M}_B$.

A_I^i and $A_I^{\mathcal{S}}$ (set of quantum inputs of parties in $A^{\mathcal{S}}$).

For the case of a general $A^{\mathcal{S}}$ such that $A^i \cup A^{\mathcal{S}} \subset \mathcal{A}$, we must use the partial composition $\mathfrak{C}(\hat{W}, \{\mathcal{M}^{A^k}\}_{k=1,k\notin i\cup\mathcal{S}}^{k=N})$ instead of the process map $\hat{W}$ in the proof. The partial composition $\mathfrak{C}(\hat{W}, \{\mathcal{M}^{A^k}\}_{k=1,k\notin i\cup\mathcal{S}}^{k=N})$ has additional in and output systems $\{A_s^k\}_{k=1,k\notin i\cup\mathcal{S}}^N$ and $\{A_o^k\}_{k=1,k\notin i\cup\mathcal{S}}^N$ carrying the settings and outcomes of all the parties other than those in $A^i \cup A^{\mathcal{S}}$. When applying the above proof method to this general case, all the statement then apply for all choices of settings on the additional inputs $\{A_s^k\}_{k=1,k\notin i\cup\mathcal{S}}^N$ i.e., for all choices of local operations of the remaining $N-1-|A^{\mathcal{S}}|$ parties, and we must trace out the additional outputs $\{A_o^k\}_{k=1,k\notin i\cup\mathcal{S}}^N$ when considering signalling between A_O^i and $A_I^{\mathcal{S}}$ in the partial composition. Since these outputs carry classical measurement outcomes, the trace can be performed in the basis in which they are encoded (the computational basis) and then simply corresponds to marginalising over them in the resulting probability distribution. Noting these points, the above proof readily generalises to the general multi-partite case where the bipartite no signalling condition $\sum_x P(xy|ab) = \sum_x P(xy|a)$ or $P(y|ab) = P(y|a)$ established above generalises to the condition of Equation 5.

Necessity: To establish that no signalling at the level of joint probabilities implies no signalling at the level of the partial composition, we prove the contrapositive of the statement which is that $A_O^i \to A_I^{\mathcal{S}}$ in the partial composition $\mathfrak{C}(\hat{W}, \{\mathcal{M}^{A^k}\}_{k=1,k\notin i\cup\mathcal{S}}^{k=N})$ implies that the set of outcomes of parties in $A^{\mathcal{S}}$ depends on the setting of A^i i.e., the associated joint probability distribution does not satisfy the independence of Equation (5). We again, for the sake of simplicity, demonstrate the proof for the bipartite case where $A^i = A$, $A^{\mathcal{S}} = B$ and $\mathfrak{C}(\hat{W}, \{\mathcal{M}^{A^k}\}_{k=1,k\notin i\cup\mathcal{S}}^{k=N}) = \hat{W}$. The generalisation to multipartite case follows from this in straight forward manner.

We first show that whenever there is a signalling relation $A_O \to B_I$ in the map $\hat{W}$, this signalling relation can be verified only using product states over the input wires A_O and B_O of the map. This follows from linearity. Suppose that the signalling is undetectable using product states i.e., for all product states $\rho_{A_O} \otimes \sigma_{B_O}$ on A_O and B_O, and for all states $\tilde{\rho}_{A_O}$ on A_O alone,

$$\mathrm{Tr}_{A_I} \circ \hat{W}(\rho_{A_O} \otimes \sigma_{B_O}) = \mathrm{Tr}_{A_I} \circ \hat{W}(\tilde{\rho}_{A_O} \otimes \sigma_{B_O}).$$

In particular, we can consider a set of product states $\{|i\rangle_{A_O} \otimes |j\rangle_{B_O}\}_{i,j}$ that form a complete basis for $\mathcal{H}_{A_O} \otimes \mathcal{H}_{B_O}$. This means that the map $\mathrm{Tr}_{A_I} \circ \hat{W}$ acts identically on the states $|i\rangle_{A_O} \otimes |j\rangle_{B_O}$ and $|i'\rangle_{A_O} \otimes |j\rangle_{B_O}$ for all possible i, i' and j, and in particular we can set $i' = 0$. Then applying this to each term in an arbitrary pure state $|\psi\rangle = \sum_{i,j} \alpha_{ij} |i\rangle_{A_O} \otimes |j\rangle_{B_O}$ and invoking linearity, we know that the action of the map $\mathrm{Tr}_{A_I} \circ \hat{W}$ on this state would be identical to its action on the product state $|0\rangle_{A_O} \otimes \sum_j (\sum_i \alpha_{ij}) |j\rangle_{B_O}$. The argument immediately extends also to arbitrary mixed states on $A_O \otimes B_O$, as these are convex mixtures of pure states.

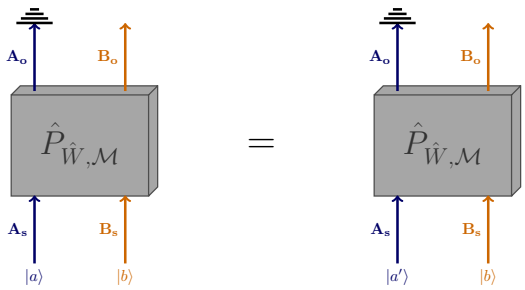

Figure 16: A consequence of the condition of Equation (27) (or equivalently Figure 15). The equality holds for all states $|a\rangle$, $|a'\rangle$ on A_s. $\hat{P}_{\hat{W},\mathcal{M}}$ corresponds to the complete composition of the process with the two local operations, as illustrated in Figure 6.

Using the above, we show that whenever $A_O \to B_I$ in the map $\hat{W}$, it can be used to implement a corresponding signalling relation from the classical input system A_s of A to the classical output system B_o of B, in the complete composition of $\hat{W}$ with the parties' local maps, which in turn implies signalling at the level of the probabilities as these are encoded in the complete composition (cf. Lemma 5.2). We have established that $A_O \to B_I$ in the map $\hat{W}$ implies that there exists a product state $\rho_{A_O} \otimes \sigma_{B_O}$ on A_O and B_O, and a state $\tilde{\rho}_{A_O}$ on A_O alone such that $\mathrm{Tr}_{A_I} \circ \hat{W}(\rho_{A_O} \otimes \sigma_{B_O}) \neq \mathrm{Tr}_{A_I} \circ \hat{W}(\tilde{\rho}_{A_O} \otimes \sigma_{B_O})$.

In other words, the state on B_I when $\hat{W}$ acts on $\rho_{A_O} \otimes \sigma_{B_O}$ is distinct from the state on B_I when it acts on $\tilde{\rho}_{A_O} \otimes \sigma_{B_O}$ and hence there exists a measurement $\{\mathcal{M}^B_{y|b}\}_y$ that B can perform to distinguish these states with a non-zero probability. Similarly, we can define A's local operation $\mathcal{M}^A$ to be such that whenever the setting $a = 0$ is input on A_s, it discards the input on A_I and prepares the state ρ_{A_O} to send to $\hat{W}$ and whenever the setting $a = 1$ is input on A_s, it discards the input on A_I and prepares the state $\tilde{\rho}_{A_O}$ to send to $\hat{W}$. This in turn implies that $P(y|a, b) \neq P(y|b)$ i.e., that the outcome y of B depends on the setting a of A, which establishes the claim.

□

Theorem 6.2 *[No-go theorem for physical implementations of processes] No fixed spacetime implementation (Definition 3.4) $\hat{P}^{\mathcal{T},\mathcal{E}}_{\hat{W},\mathcal{M}}$ of the process network $\hat{P}_{\hat{W},\mathcal{M}}$ within a fixed spacetime structure $\mathcal{T}$ (Definition 3.1) can simultaneously satisfy the following three assumptions.*

1. *W is not a fixed order process (Definition 4.1).*
2. *$\hat{P}^{\mathcal{T},\mathcal{E}}_{\hat{W},\mathcal{M}}$ satisfies the relativistic causality condition of Definition 3.5.*
3. *The region causal structure given by the embedding $\mathcal{E}$ with $\mathrm{Nodes}(\mathcal{G}^R_{\mathcal{T}}) := \{\mathcal{E}(S)\}_{S\in\mathcal{S}}$ is cycle-free.*

Proof: Consider the following property that signalling relations associated with the network $\hat{P}_{\hat{W},\mathcal{M}}$ may satisfy– every signalling relation $\mathcal{S}_1 \to \mathcal{S}_2$ between some subsets $\mathcal{S}_1, \mathcal{S}_2$ of $\{A^1_I, A^1_O, ..., A^N_I, A^N_O\}$ is such that there exists $S_1 \in \mathcal{S}_1$, and $S_2 \in \mathcal{S}_2$ with a corresponding signalling relation $S_1 \to S_2$. Notice that for every network $\hat{P}_{\hat{W},\mathcal{M}}$ that does not have this property, there exists network that does satisfy the property such that the signalling relations of both networks impose the same relativistic causality constraints on any spacetime embedding $\mathcal{E}$. A proof of this statement can be found in the proof of Theorem 3.6, and we will not repeat this here. This means that without loss of generality, we can proceed with assuming that our network always leads to signalling relations satisfying this property. This also means that, without loss of generality, we only need to consider signalling relations between individual systems (and not subsets of systems) in $\{A^1_I, A^1_O, ..., A^N_I, A^N_O\}$.

Now, notice that the network $\hat{P}_{\hat{W},\mathcal{M}}$ can either be viewed as a complete composition of $\hat{W}$ with the local maps, or equivalently, as a composition of the partial composition $\mathfrak{C}(\hat{W}, \{\mathcal{M}^{A^k}\}_{k=1,A^k\notin A\cup\mathcal{S}})$ (which itself is a composition of $\hat{W}$ with the local operations of the agents in $A \cup \mathcal{S}$), where A is some agent and $\mathcal{S}$ is a subset of the $N-1$ agents excluding A, with the local maps of the remaining $N-1-|\mathcal{S}|$ agents $\{A^1, ..., A^N\}\backslash\{A \cup \mathcal{S}\}$. This means that all the signalling relations coming from these maps can be exploited in an implementation of the network. Furthermore, since the "extended" local map $\mathcal{M}^{A^k}$

is also part of the network, we must also consider the signalling relations that it generates. Noting that each extended map encodes all possible choices of local operations that could be implemented from A_I^k to A_O^k (and not just maps that discard states on A_I^k and reprepare states on A_O^k), we will have a signalling relation $A_I^k \to A_O^k$ for each k.

Keeping the above in mind, we construct a signalling structure $\mathcal{G}^{sig}$ over the systems $\{A_I^1, A_O^1, ..., A_I^N, A_O^N\}$ as follows, which captures the signalling relations possible in the network $\hat{P}_{\hat{W},\mathcal{M}}$, as we have argued above. The nodes of $\mathcal{G}^{sig}$ correspond to elements in $\{A_I^1, A_O^1, ..., A_I^N, A_O^N\}$. Whenever $A_O \to A_I^{\mathcal{S}}$ in the partial composition $\mathfrak{C}(\hat{W}, \{\mathcal{M}^{A^k}\}_{k=1, A^k \notin A \cup A^{\mathcal{S}}})$, where $A_I^{\mathcal{S}}$ denotes the set of all input systems of agents in $A^{\mathcal{S}}$, we pick any individual system $S_I \in \mathcal{S}_I$ and include an edge $\to$ from A_O to A_I^S in $\mathcal{G}^{sig}$. Further, we include $A_I^i \to A_O^i$ in $\mathcal{G}^{sig}$ for each $i \in \{1, ..., N\}$. The reason for not considering $\mathcal{G}^{sig}$ to a directed graph over Powerset$[\{A_I^1, A_O^1, ..., A_I^N, A_O^N\}]$ and including an edge from A_O to $A_I^{\mathcal{S}}$ is that we only care about the relativistic causality constraints imposed by these signalling relations on a spacetime embedding. And in both these cases, we have the same constraints (see proof of Theorem 3.6 for a more detailed explanation of this statement). By Theorem 5.4, we know that this is equivalent to the procedure where we check whether the set of outcomes of agents in $A^{\mathcal{S}}$ depends on the setting of agent A through the joint probability distribution (i.e., whether Equation 5 fails when setting $A := A^i$) and then picking $S \in \mathcal{S}$ and including an edge $A \to A^S$ in $\mathcal{G}^{sig}$.

It is then easy to see that if W is not a fixed order process (as required by assumption 1), then $\mathcal{G}^{sig}$ must contain a directed cycle i.e. a set of systems $S_1, S_2, ..., S_n \in \{A_I^1, A_O^1, ..., A_I^N, A_O^N\}$ such that $S_1 \to S_2 \to ... \to S_n \to S_1$. Suppose that $\mathcal{G}^{sig}$ does not contain a directed cycle, then it would be a directed acyclic graph and we can define a partial order relation $\prec_{\mathcal{K}}$ on the systems $\{A_I^1, A_O^1, ..., A_I^N, A_O^N\}$ such that

$$S_1 \to S_2 \in \mathcal{G} \quad \Leftrightarrow \quad S_1 \prec_{\mathcal{K}} S_2.$$

It follows that W satisfies Definition 4.1 and is therefore a fixed order process.

We now impose relativistic causality (assumption 2) and show that assumption 3 must be violated for any spacetime implementation $\hat{P}^{\mathcal{T},\mathcal{E}}_{\hat{W},\mathcal{M}}$ of the network $\hat{P}_{\hat{W},\mathcal{M}}$. This also follows readily. Relativistic causality requires that whenever $S_1 \to S_2$ we must have $\mathcal{P}^{S_1} \xrightarrow{R} \mathcal{P}^{S_2}$, where $\mathcal{P}^{S_1}$ and $\mathcal{P}^{S_2}$ are the spacetime regions assigned to the systems S_1 and S_2 by the spacetime embedding $\mathcal{E}$. Applied this to the directed cycle $S_1 \to S_2 \to ... \to S_n \to S_1$ in $\mathcal{G}^{sig}$ (which we have proven to exist whenever assumption 1 is satisfied), we require $\mathcal{P}^{S_1} \xrightarrow{R} \mathcal{P}^{S_2} \xrightarrow{R} ... \xrightarrow{R} \mathcal{P}^{S_n} \xrightarrow{R} \mathcal{P}^{S_1}$ in order to satisfy assumption 2. However, this is a direct violation of assumption 3, which establishes the claim.

□

Corollary 6.6 *[Time localisation in a global frame] Under the same notation as Theorem 6.2, no fixed spacetime implementation $\hat{P}^{\mathcal{T},\mathcal{E}}_{\hat{W},\mathcal{M}}$ of the network $\hat{P}_{\hat{W},\mathcal{M}}$ within a fixed spacetime structure $\mathcal{T}$ can simultaneously satisfy the following three assumptions,*

1. *W is not a fixed order process (Definition 4.1).*
2. *$\hat{P}^{\mathcal{T},\mathcal{E}}_{\hat{W},\mathcal{M}}$ satisfies the relativistic causality condition of Definition 3.5.*
3. *The spacetime embedding $\mathcal{E}$ has the property that each of the spacetime regions $\mathcal{P}^S \subseteq \mathcal{T}$ are time-localised from the perspective of some agent A.*

Proof: This corollary follows from noting that the first two assumptions are identical to the first two assumptions of Theorem 6.2, while the third assumption here implies the third assumption of the theorem. More explicitly, assumption 3 here requires that for each system S, the corresponding spacetime region $\mathcal{P}^S$ assigned to S by the embedding $\mathcal{E}$ is such that all spacetime points in $\mathcal{P}^S$ have the same time coordinate, say t^S in a global reference frame. Then, Definitions 3.1 and 3.2 tell us that $\mathcal{P}^{S_1} \xrightarrow{R} \mathcal{P}^{S_2}$ implies $t^{S_1} < t^{S_2}$. Since we can never have a sequence of times in some global reference frame such that $t^{S_1} < t^{S_2} < ... < t^{S_n} < t^{S_1}$, it follows that the set of spacetime regions satisfying assumption 3 of this corollary can never contain a sequence of regions such that $\mathcal{P}^{S_1} \xrightarrow{R} \mathcal{P}^{S_2} \xrightarrow{R} ... \xrightarrow{R} \mathcal{P}^{S_n} \xrightarrow{R} \mathcal{P}^{S_1}$, i.e., the regions satisfy assumption 3 of Theorem 6.2. □

Theorem 6.9 *[Unravelling physical process implementations into fixed order processes] Let $\hat{W}^{\mathcal{T},\mathcal{E}}_{max}$ be a maximally fine-grained physical implementation of an N-partite process $\hat{W}$ in a spacetime, where each pair of in and output regions $\mathcal{P}^{A^k_I}$ and $\mathcal{P}^{A^k_O}$ have a pairwise correspondence $\mathcal{O}^{A^k} : \mathcal{P}^{A^k_I} \mapsto \mathcal{P}^{A^k_O}$. Then $\hat{W}^{\mathcal{T},\mathcal{E}}_{max}$ acts as an $\tilde{N}$-partite fixed order process with $\tilde{N} = \sum_{k=1}^{N} |\mathcal{P}^{A^k_I}|$, upon composition with corresponding maximally fine-grained local maps $\{\mathcal{M}^{A^k,\mathcal{T},\mathcal{E}}_{max}\}_{k=1}^{N}$, where each $\mathcal{M}^{A^k,\mathcal{T},\mathcal{E}}_{max}$ acts independently between the pairs of points $P^{A^k_I} \in \mathcal{P}^{A^k_I}$ and $\mathcal{O}^{A^k}(P^{A^k_I}) \in \mathcal{P}^{A^k_O}$, as described by Equation (12).*

Proof: We prove the bipartite case here for simplicity, but the proof method readily generalises to the N-partite case. Let $\hat{W}$ be a bipartite process matrix over the parties A and B, and $\hat{W}^{\mathcal{T},\mathcal{E}}$ be a spacetime implementation of $\hat{W}$ where the spacetime regions assigned by the embedding are $\mathcal{P}^{A_I} = \{P^1_I, ..., P^n_I\}$, $\mathcal{P}^{A_O} = \{P^1_O, ..., P^n_O\}$, $\mathcal{P}^{B_I} = \{Q^1_I, ..., Q^m_I\}$ and $\mathcal{P}^{B_O} = \{Q^1_O, ..., Q^m_O\}$ with the pairwise correspondence $\mathcal{O}^A(P^i_I) = P^O_i \succ P^i_I$ and $\mathcal{O}^B(Q^j_I) = Q^O_i \succ Q^j_I$. While $\hat{W}$ has the inputs $\{A_O, B_O\}$ and outputs $\{A_I, B_I\}$, the maximal fine-graining $\hat{W}^{\mathcal{T},\mathcal{E}}_{max}$ now has $n+m$ input systems, $\{A^1_O, ..., A^n_O, B^1_O, ..., B^m_O\}$ and $n+m$ output systems $\{A^1_I, ..., A^n_I, B^1_I, ..., B^m_I\}$, where $A^i_{O/I}$ is short for $A^{P^i_O}_{O/I}$, and similarly for B.

Let $\mathcal{M}^A$ and $\mathcal{M}^B$ be local maps and $\mathcal{M}^{A,\mathcal{T},\mathcal{E}}_{max}$ and $\mathcal{M}^{B,\mathcal{T},\mathcal{E}}_{max}$ be the maximal fine-graining of a spacetime implementation of these maps with the same embedding $\mathcal{E}$ considered above i.e., the input systems of $\mathcal{M}^{A,\mathcal{T},\mathcal{E}}_{max}$ consist of n quantum inputs $\{A^1_I, ..., A^n_I\}$ and n classical inputs $\{A^1_s, ..., A^n_s\}$, and outputs consist of the n quantum outputs $\{A^1_O, ..., A^n_O\}$ and n classical outputs $\{A^1_o, ..., A^n_o\}$, and similarly for $\mathcal{M}^{B,\mathcal{T},\mathcal{E}}_{max}$, which will have m quantum and m classical in and outputs. Given that these fine-grained local maps to act according to Equations (12) and (13), $\mathcal{M}^{A,\mathcal{T},\mathcal{E}}_{max}$ corresponds to a tensor product of n local maps $\{\mathcal{M}^A_{P^i_I}\}_{i=1}^n$ and $\mathcal{M}^{B,\mathcal{T},\mathcal{E}}_{max}$ corresponds to a tensor product of m local maps $\{\mathcal{M}^B_{Q^j_I}\}_{j=1}^m$.

Then composing $\mathcal{M}^{A,\mathcal{T},\mathcal{E}}_{max}$ and $\mathcal{M}^{B,\mathcal{T},\mathcal{E}}_{max}$ with $\hat{W}^{\mathcal{T},\mathcal{E}}_{max}$ through loop composition of wires with identical labels is equivalent to composing the $n+m$ local maps $\{\mathcal{M}^A_{P^i_I}\}_{i=1}^n \bigcup \{\mathcal{M}^B_{Q^j_I}\}_{j=1}^m$ with $\hat{W}^{\mathcal{T},\mathcal{E}}_{max}$ through loop composition. We can therefore regard this as an $n+m$-partite process over the parties $\{A^1, ..., A^n, B^1, ..., B^m\}$. Now note that the spacetime implemented process network formed by the composition of $\hat{W}^{\mathcal{T},\mathcal{E}}_{max}$ with these $n+m$ maps is one where the in and output systems of all $n+m$ parties is localised in spacetime. Then noting that we want this to be a physical implementation satisfying relativistic causality, it immediately follows by Corollary 6.4 that $\hat{W}^{\mathcal{T},\mathcal{E}}_{max}$ is a fixed order process over these $n+m$ parties.

□

Lemma 7.1 *[No-go result for the quantum switch] Consider the process map $\hat{W}^{QS}$ whose Choi representation is the process matrix W^{QS} of the quantum switch. Let $\hat{P}_{QS,U,V}$ be the quantum switch network where W^{QS} acts on two non-trivial local operations $U^A : A_I \mapsto A_O$ and $V^B : B_I \mapsto B_O$ of Alice and Bob. Then any fixed spacetime implementation $\hat{P}^{\mathcal{T},\mathcal{E}}_{QS,U,V}$ of this network cannot simulataneously satisfy both of the following assumptions*

1. *$\hat{P}^{\mathcal{T},\mathcal{E}}_{QS,U,V}$ satisfies relativistic causality*
2. *The subgraph of the region causal structure given by the embedding $\mathcal{E}$ with $\mathrm{Nodes}(\mathcal{G}^R_{\mathcal{T}}) := \{\mathcal{E}(S)\}_{S\in\mathcal{S}}$, restricted to $S \in \{A_I, A_O, B_I, B_O\}$ is cycle-free.*

Proof: We first show that in $\hat{W}^{QS}$, when the initial state of the control and target systems is $|\psi_C\rangle := \alpha|0\rangle + \beta|1\rangle$ and $|\psi_T\rangle$ respectively where α and β are both non-zero amplitudes and $|\psi_T\rangle$ is an arbitrary qubit state, then we have $A_O \to B_I$ irrespective of the state input on B_O and $B_O \to A_I$ irrespective of the state input on A_O, in $\hat{W}$. This implies that both signalling relations can be realised in a single instance of $\hat{P}_{QS,U,V}$, irrespective of the choice of local operations U^A and V^B, in contrast to the general case of Theorem 6.2 where the signalling relation being realised may depend on the choice of local operations, such that not all signalling relations allowed by $\hat{W}$ are realised when it is composed with a given set of local operations (which is why the Theorem evokes the extended local maps which encode all these choices).

We show that the network $\hat{P}_{QS,U,V}$ gives rise to a directed cycle of signalling relations $A_O \to B_I \to B_O \to A_I \to A_O$ for any non-trivial local operations U^A and V^B. Then this directed cycle of signalling relations implies (by relativistic causality) that the spacetime regions must satisfy $\mathcal{P}^{A_O} \xrightarrow{R} \mathcal{P}^{B_I} \xrightarrow{R}$

$\mathcal{P}^{B_O} \xrightarrow{R} \mathcal{P}^{A_I} \xrightarrow{R} \mathcal{P}^{A_O}$, which violates assumption 3. Therefore to complete the proof, we only need to establish the statement about the directed cycle of signalling relations in $\hat{W}^{QS}$ outlined at the beginning of this paragraph, which we do below.

Consider the action of the process map $\hat{W}^{QS} : P^C \otimes P^T \otimes A_O \otimes B_O \mapsto F^C \otimes F^T \otimes A_I \otimes B_I$ on the input state $|\Psi\rangle_{P^C P^T A_O B_O} := (\alpha|0\rangle + \beta|1\rangle)_{P^C} \otimes |\psi_T\rangle_{P^T} \otimes |\psi_A\rangle_{A_O} \otimes |\psi_B\rangle_{B_O}$, where $|\psi_A\rangle$ and ψ_B are arbitrary qubits states. We have

$$\hat{W}^{QS}.|\Psi\rangle_{P^C P^T A_O B_O} = \alpha|0\rangle_{F^C}|\psi_B\rangle_{F^T}|\psi_T\rangle_{A_I}|\psi_A\rangle_{B_I} + \beta|1\rangle_{F^C}|\psi_A\rangle_{F^T}|\psi_B\rangle_{A_I}|\psi_T\rangle_{B_I} := |\Phi\rangle_{F^C F^T A_I B_I}.$$

Then

$$\begin{aligned} \mathrm{Tr}_{F^C F^T A_I}\left[\hat{W}^{QS}(|\Psi\rangle\langle\Psi|_{P^C P^T A_O B_O})\right] &= |\alpha|^2|\psi_A\rangle\langle\psi_A|_{B_I} + |\beta|^2|\psi_T\rangle\langle\psi_T|_{B_I} \\ \mathrm{Tr}_{F^C F^T B_I}\left[\hat{W}^{QS}(|\Psi\rangle\langle\Psi|_{P^C P^T A_O B_O})\right] &= |\alpha|^2|\psi_T\rangle\langle\psi_T|_{A_I} + |\beta|^2|\psi_B\rangle\langle\psi_B|_{A_I} \end{aligned} \tag{28}$$

Notice that $\mathrm{Tr}_{F^C F^T A_I}\left[\hat{W}^{QS}(|\Psi\rangle\langle\Psi|)\right]$ which is the output on B_I depends only on ψ_A i.e., the input on A_O and not on ψ_B, the input on B_O and similarly $\mathrm{Tr}_{F^C F^T B_I}\left[\hat{W}^{QS}(|\Psi\rangle\langle\Psi|)\right]$ only depends on ψ_B. This implies that given the knowledge of the initial control and target states, $\alpha|0\rangle + \beta|1\rangle$ and $|\psi_T\rangle$ Alice and Bob can signal to each other by suitable choices of ψ_A and ψ_B on their respective output systems A_O and B_O, irrespective of the local operation of the other party. The above proof easily generalises to arbitrary input states $\rho_{P^C P^T A_O B_O} := |\psi_C\rangle\langle\psi_C|_{P^C} \otimes |\psi_T\rangle\langle\psi_T|_{P^T} \otimes \rho_{A_O B_O}$, where the input state $\rho_{A_O B_O}$ on A_O and B_O may be an entangled state, $\mathrm{Tr}_{F^C F^T A_I}\left[\hat{W}^{QS}(|\Psi\rangle\langle\Psi|)\right]$ depends only on the marginal of the initial state over A_O which is unaffected by local operations on B_O.

□

References

[1] John Avery. Information Theory And Evolution, chapter 4, pages 73–94. *World Scientific*, 2003.

[2] V. Vilasini and Roger Colbeck. Possibility of causal loops without superluminal signalling – a general framework, 2021. https://arxiv.org/abs/2109.12128.

[3] Rolf Landauer. *Physics Today*, 44.

[4] A. Einstein, B. Podolsky, and N. Rosen. Can quantum-mechanical description of physical reality be considered complete? *Phys. Rev.*, 47:777–780, May 1935.

[5] Albert B. Einstein, Max Born, and Hedwig Born. The Born-Einstein letters: Correspondence between Albert Einstein and Max and Hedwig Born from 1916-1955, with commentaries by Max Born;. 1971.

[6] J.S. Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1:195–200, 1964.

[7] Louis de Broglie. La mécanique ondulatoire et la structure atomique de la matière et du rayonnement. *J. Phys. Radium*, 8:225–241, 1927.

[8] David Bohm. A suggested interpretation of the quantum theory in terms of "hidden" variables. i. *Phys. Rev.*, 85:166–179, Jan 1952.

[9] Howard Wiseman and Eric G. Cavalcanti. Causarum Investigatio and the Two Bell's Theorems of John Bell, 2014. https://arxiv.org/abs/1503.06413.

[10] Robert R. Tucci. Quantum bayesian nets. *International Journal of Modern Physics B*, 09(03):295–337, Jan 1995. http://dx.doi.org/10.1142/S0217979295000148.

[11] M. S. Leifer. Quantum dynamics as an analog of conditional probability. *Physical Review A*, 74(4), Oct 2006. http://dx.doi.org/10.1103/PhysRevA.74.042310.

[12] Kathryn B. Laskey. Quantum causal networks, 2007. https://arxiv.org/abs/0710.1200.

[13] M.S. Leifer and D. Poulin. Quantum graphical models and belief propagation. *Annals of Physics*, 323(8):1899–1946, Aug 2008. http://dx.doi.org/10.1016/j.aop.2007.10.001.

[14] M. S. Leifer and Robert W. Spekkens. Towards a formulation of quantum theory as a causally neutral theory of Bayesian inference. *Physical Review A*, 88(5):052130, nov 2013. https://link.aps.org/doi/10.1103/PhysRevA.88.052130.

[15] Joe Henson, Raymond Lal, and Matthew Pusey. Theory-independent limits on correlations from generalized bayesian networks. *New Journal of Physics*, 16(11):113043, nov 2014.

[16] 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):33002, 2015. https://iopscience.iop.org/article/10.1088/1367-2630/17/3/033002.

[17] Jacques Pienaar and Časlav Brukner. A graph-separation theorem for quantum causal models. *New Journal of Physics*, 17(7):73020, 2015. https://iopscience.iop.org/article/10.1088/1367-2630/17/7/073020.

[18] Katja Ried, Megan Agnew, Lydia Vermeyden, Dominik Janzing, Robert W. Spekkens, and Kevin J. Resch. A quantum advantage for inferring causal structure. *Nature Physics*, 11(5):414–420, Mar 2015. http://dx.doi.org/10.1038/nphys3266.

[19] Fabio Costa and Sally Shrapnel. Quantum causal modelling. *New Journal of Physics*, 18(6):63032, 2016. https://iopscience.iop.org/article/10.1088/1367-2630/18/6/063032.

[20] Tobias Fritz. Beyond bell's theorem ii: Scenarios with arbitrary causal structure. *Communications in Mathematical Physics*, 341(2):391–434, Nov 2015. http://dx.doi.org/10.1007/s00220-015-2495-5.

[21] John-Mark A. Allen, Jonathan Barrett, Dominic C. Horsman, Ciarán M. Lee, and Robert W. Spekkens. Quantum Common Causes and Quantum Causal Models. *Physical Review X*, 7(3):031021, jul 2017. http://link.aps.org/doi/10.1103/PhysRevX.7.031021.

[22] Jonathan Barrett, Robin Lorenz, and Ognyan Oreshkov. Quantum causal models, 2020. https://arxiv.org/abs/1906.10726.

[23] Jacques Pienaar. Quantum causal models via quantum bayesianism. *Physical Review A*, 101(1), Jan 2020. http://dx.doi.org/10.1103/PhysRevA.101.012104.

[24] Judea Pearl. Causality: Models, reasoning, and inference. *Second edition, Cambridge University Press*, 2009.

[25] Lucien Hardy. Probability Theories with Dynamic Causal Structure: A New Framework for Quantum Gravity, sep 2005. http://arxiv.org/abs/gr-qc/0509120.

[26] Lucien Hardy. Towards quantum gravity: a framework for probabilistic theories with non-fixed causal structure. *Journal of Physics A: Mathematical and Theoretical*, 40(12):3081–3099, Mar 2007. http://dx.doi.org/10.1088/1751-8113/40/12/S12.

[27] Magdalena Zych, Fabio Costa, Igor Pikovski, and Časlav Brukner. Bell's theorem for temporal order. *Nature Communications*, 10(1):3772, 2019. https://doi.org/10.1038/s41467-019-11579-x.

[28] Ognyan Oreshkov, Fabio Costa, and Časlav Brukner. Quantum correlations with no causal order. *Nature Communications*, 3:1092, oct 2012. https://www.nature.com/articles/ncomms2076.

[29] Giulio Chiribella, Giacomo Mauro D'Ariano, Paolo Perinotti, and Benoit Valiron. Quantum computations without definite causal structure. *Physical Review A*, 88(2):022318, aug 2013. https://link.aps.org/doi/10.1103/PhysRevA.88.022318.

[30] Mateus Araújo, Cyril Branciard, Fabio Costa, Adrien Feix, Christina Giarmatzi, and Časlav Brukner. Witnessing causal nonseparability. *New Journal of Physics*, 17(10):102001, 2015. https://iopscience.iop.org/article/10.1088/1367-2630/17/10/102001.

[31] Ognyan Oreshkov and Christina Giarmatzi. Causal and causally separable processes. *New Journal of Physics*, 18(9):93020, 2016. https://iopscience.iop.org/article/10.1088/1367-2630/18/9/093020.

[32] Ämin Baumeler and Stefan Wolf. The space of logically consistent classical processes without causal order. *New Journal of Physics*, 18(1):013036, jan 2016.

[33] Philippe Allard Guérin, Adrien Feix, Mateus Araújo, and Časlav Brukner. Exponential communication complexity advantage from quantum superposition of the direction of communication. *Phys. Rev. Lett.*, 117:100502, Sep 2016. https://link.aps.org/doi/10.1103/PhysRevLett.117.100502.

[34] Giulio Chiribella, Manik Banik, Some Sankar Bhattacharya, Tamal Guha, Mir Alimuddin, Arup Roy, Sutapa Saha, Sristy Agrawal, and Guruprasad Kar. Indefinite causal order enables perfect quantum communication with zero capacity channels. *New Journal of Physics*, 23(3):033039, Mar 2021. http://dx.doi.org/10.1088/1367-2630/abe7a0.

[35] Giulio Chiribella. Perfect discrimination of no-signalling channels via quantum superposition of causal structures. *Physical Review A*, 86(4), Oct 2012. http://dx.doi.org/10.1103/PhysRevA.86.040301.

[36] Xiaobin Zhao, Yuxiang Yang, and Giulio Chiribella. Quantum metrology with indefinite causal order. *Physical Review Letters*, 124(19), May 2020. http://dx.doi.org/10.1103/PhysRevLett.124.190503.

[37] Mateus Araújo, Adrien Feix, Fabio Costa, and Časlav Brukner. Quantum circuits cannot control unknown operations. *New Journal of Physics*, 16(9):93026, 2014. https://iopscience.iop.org/article/10.1088/1367-2630/16/9/093026.

[38] Tamal Guha, Mir Alimuddin, and Preeti Parashar. Thermodynamic advancement in the causally inseparable occurrence of thermal maps. *Physical Review A*, 102(3), Sep 2020. http://dx.doi.org/10.1103/PhysRevA.102.032215.

[39] David Felce and Vlatko Vedral. Quantum refrigeration with indefinite causal order. *Physical Review Letters*, 125(7), Aug 2020. http://dx.doi.org/10.1103/PhysRevLett.125.070603.

[40] Lorenzo M Procopio, Amir Moqanaki, Mateus Araújo, Fabio Costa, Irati Alonso Calafell, Emma G Dowd, Deny R Hamel, Lee A Rozema, Časlav Brukner, and Philip Walther. Experimental superposition of orders of quantum gates. *Nature Communications*, 6:7913, August 2015. https://www.nature.com/articles/ncomms8913.

[41] Giulia Rubino, Lee A. Rozema, Adrien Feix, Mateus Araújo, Jonas M. Zeuner, Lorenzo M. Procopio, Časlav Brukner, and Philip Walther. Experimental verification of an indefinite causal order. *Science Advances*, 3(3), 2017. http://advances.sciencemag.org/content/3/3/e1602589.

[42] K. Goswami, C. Giarmatzi, M. Kewming, F. Costa, C. Branciard, J. Romero, and A.G. White. Indefinite causal order in a quantum switch. *Physical Review Letters*, 121(9), Aug 2018. http://dx.doi.org/10.1103/PhysRevLett.121.090503.

[43] Kejin Wei, Nora Tischler, Si-Ran Zhao, Yu-Huai Li, Juan Miguel Arrazola, Yang Liu, Weijun Zhang, Hao Li, Lixing You, Zhen Wang, and et al. Experimental quantum switching for exponentially superior quantum communication complexity. *Physical Review Letters*, 122(12), Mar 2019. http://dx.doi.org/10.1103/PhysRevLett.122.120504.

[44] C. T. Marco Ho, Fabio Costa, Christina Giarmatzi, and Timothy C. Ralph. Violation of a causal inequality in a spacetime with definite causal order, 2019. https://arxiv.org/abs/1804.05498.

[45] Yu Guo, Xiao-Min Hu, Zhi-Bo Hou, Huan Cao, Jin-Ming Cui, Bi-Heng Liu, Yun-Feng Huang, Chuan-Feng Li, Guang-Can Guo, and Giulio Chiribella. Experimental transmission of quantum information using a superposition of causal orders. *Physical Review Letters*, 124(3), Jan 2020. http://dx.doi.org/10.1103/PhysRevLett.124.030502.

[46] K. Goswami and J. Romero. Experiments on quantum causality. *AVS Quantum Science*, 2(3):037101, Oct 2020. http://dx.doi.org/10.1116/5.0010747.

[47] Márcio M. Taddei, Jaime Cariñe, Daniel Martínez, Tania García, Nayda Guerrero, Alastair A. Abbott, Mateus Araújo, Cyril Branciard, Esteban S. Gómez, Stephen P. Walborn, and et al. Computational advantage from the quantum superposition of multiple temporal orders of photonic gates. *PRX Quantum*, 2(1), Feb 2021. http://dx.doi.org/10.1103/PRXQuantum.2.010320.

[48] Giulia Rubino, Lee A. Rozema, Daniel Ebler, Hlér Kristjánsson, Sina Salek, Philippe Allard Guérin, Alastair A. Abbott, Cyril Branciard, Časlav Brukner, Giulio Chiribella, and et al. Experimental quantum communication enhancement by superposing trajectories. *Physical Review Research*, 3(1), Jan 2021. http://dx.doi.org/10.1103/PhysRevResearch.3.013093.

[49] David Felce, Nicetu Tibau Vidal, Vlatko Vedral, and Eduardo O. Dias. Indefinite causal orders from superpositions in time, 2021. https://arxiv.org/abs/2107.08076.

[50] Nikola Paunković and Marko Vojinović. Causal orders, quantum circuits and spacetime: distinguishing between definite and superposed causal orders. *Quantum*, 4:275, May 2020. http://dx.doi.org/10.22331/q-2020-05-28-275.

[51] Jonathan Barrett, Robin Lorenz, and Ognyan Oreshkov. Cyclic quantum causal models. *Nature Communications*, 12(1), Feb 2021. http://dx.doi.org/10.1038/s41467-020-20456-x.

[52] Mateus Araújo, Philippe Allard Guérin, and Ämin Baumeler. Quantum computation with indefinite causal structures. *Physical Review A*, 96(5), Nov 2017. http://dx.doi.org/10.1103/PhysRevA.96.052315.

[53] Stephan Bongers, Patrick Forré, Jonas Peters, and Joris M. Mooij. Foundations of structural causal models with cycles and latent variables. *The Annals of Statistics*, 49(5):2885 – 2915, 2021. https://doi.org/10.1214/21-AOS2064.

[54] Christopher Portmann, Christian Matt, Ueli Maurer, Renato Renner, and Bjorn Tackmann. Causal Boxes: Quantum Information-Processing Systems Closed under Composition. *IEEE Transactions on Information Theory*, 63(5):3277–3305, 2017. http://ieeexplore.ieee.org/document/7867830/.

[55] Ralph Silva, Yelena Guryanova, Anthony J Short, Paul Skrzypczyk, Nicolas Brunner, and Sandu Popescu. Connecting processes with indefinite causal order and multi-time quantum states. *New Journal of Physics*, 19(10):103022, Oct 2017. http://dx.doi.org/10.1088/1367-2630/aa84fe.

[56] Yakir Aharonov, Sandu Popescu, Jeff Tollaksen, and Lev Vaidman. Multiple-time states and multiple-time measurements in quantum mechanics. *Physical Review A*, 79(5), May 2009. http://dx.doi.org/10.1103/PhysRevA.79.052110.

[57] Mateus Araújo, Adrien Feix, Miguel Navascués, and Časlav Brukner. A purification postulate for quantum mechanics with indefinite causal order. *Quantum*, 1:10, nov 2016.

[58] Philippe Allard Guérin, Marius Krumm, Costantino Budroni, and Časlav Brukner. Composition rules for quantum processes: a no-go theorem. *New Journal of Physics*, 21(1):012001, Jan 2019. http://dx.doi.org/10.1088/1367-2630/aafef7.

[59] V. Vilasini, Lídia del Rio, and Renato Renner. Causality in definite and indefinite spacetimes, 2020. Extended abstract from QPL 2020, https://wdi.centralesupelec.fr/users/valiron/qplmfps/papers/qs01t3.pdf, talk version at https://www.youtube.com/watch?v=iBCntE0zxg8.

[60] Giulia Rubino, Lee A. Rozema, Francesco Massa, Mateus Araújo, Magdalena Zych, Časlav Brukner, and Philip Walther. Experimental entanglement of temporal orders. In *Quantum Information and Measurement (QIM) V: Quantum Technologies*, page S3B.3. Optical Society of America, 2019. http://www.osapublishing.org/abstract.cfm?URI=QIM-2019-S3B.3.

[61] Esteban Castro-Ruiz, Flaminia Giacomini, Alessio Belenchia, and Časlav Brukner. Quantum clocks and the temporal localisability of events in the presence of gravitating quantum systems. *Nature Communications*, 11(1), May 2020. http://dx.doi.org/10.1038/s41467-020-16013-1.

[62] Flaminia Giacomini, Esteban Castro-Ruiz, and Časlav Brukner. Quantum mechanics and the covariance of physical laws in quantum reference frames. *Nature Communications*, 10(1):494, 2019. https://doi.org/10.1038/s41467-018-08155-0.

[63] Augustin Vanrietvelde, Philipp A. Hoehn, Flaminia Giacomini, and Esteban Castro-Ruiz. A change of perspective: switching quantum reference frames via a perspective-neutral framework. *Quantum*, 4:225, January 2020. https://quantum-journal.org/papers/q-2020-01-27-225/.

[64] Philipp A Höhn and Augustin Vanrietvelde. How to switch between relational quantum clocks. *New Journal of Physics*, 22(12):123048, dec 2020. https://doi.org/10.1088/1367-2630/abd1ac.

[65] Flaminia Giacomini and Časlav Brukner. Quantum superposition of spacetimes obeys einstein's equivalence principle. *AVS Quantum Science*, 4(1):015601, 2022. https://doi.org/10.1116/5.0070018.

[66] Aleksandra Dimić, Marko Milivojević, Dragoljub Gočanin, Natália S. Móller, and Časlav Brukner. Simulating indefinite causal order with rindler observers. *Frontiers in Physics*, 8:470, 2020. https://www.frontiersin.org/article/10.3389/fphy.2020.525333.

[67] Julian Wechs, Hippolyte Dourdent, Alastair A. Abbott, and Cyril Branciard. Quantum circuits with classical versus quantum control of causal order, 2021. https://arxiv.org/abs/2101.08796.

[68] Tom Purves and Anthony J. Short. Quantum theory cannot violate a causal inequality, 2021. https://arxiv.org/abs/2101.09107.

[69] Julian Wechs, Cyril Branciard, and Ognyan Oreshkov. Existence of processes violating causal inequalities on time-delocalised subsystems, 2022. https://arxiv.org/abs/2201.11832.

[70] Giulio Chiribella, Giacomo Mauro D'Ariano, and Paolo Perinotti. Theoretical framework for quantum networks. *Physical Review A*, 80(2), Aug 2009. http://dx.doi.org/10.1103/PhysRevA.80.022339.