

Causally consistent dynamic slicing

Roly Perera^{1,2}, Deepak Garg³, and James Cheney¹

1 Laboratory for Foundations of Computer Science, University of Edinburgh, Edinburgh, UK
rperera@inf.ed.ac.uk, jcheney@inf.ed.ac.uk

2 School of Computing Science, University of Glasgow, Glasgow, UK
roly.perera@glasgow.ac.uk

3 Max Planck Institute for Software Systems, Saarbrücken, Germany
dg@mpi-sws.org

Abstract

We offer a lattice-theoretic account of the problem of dynamic slicing for π -calculus, building on prior work in the sequential setting. For any particular run of a concurrent program, we exhibit a Galois connection relating forward and backward slices of the initial and terminal configurations. We prove that, up to lattice isomorphism, the same Galois connection arises for any causally equivalent execution, allowing an efficient concurrent implementation of slicing via a standard interleaving semantics. Our approach has been formalised in the dependently-typed programming language Agda.

1998 ACM Subject Classification D.1.3 Concurrent Programming; D.2.5 Testing and debugging

Keywords and phrases π -calculus; dynamic slicing; causal equivalence; Galois connection

Digital Object Identifier 10.4230/LIPIcs...

1 Introduction

Dynamic slicing, due originally to Weiser [20], is a dynamic analysis technique with applications in debugging, security and provenance tracking. The basic goal is to identify a sub-program, or *program slice*, that may affect an outcome of interest called the *slicing criterion*, such as the value of a variable. Dynamic slicing in concurrent settings is often represented as a graph reachability problem, thanks to influential work by Cheng [4]. However, most prior work on dynamic slicing for concurrency does not yield minimum slices, nor allow particularly flexible slicing criteria, such as arbitrary parts of configurations. Systems work on concurrent slicing [10, 15, 19] tends to be largely informal.

Perera et al [16] developed an approach where backward dynamic slicing is treated as a kind of (abstract) reverse execution or “rewind” and forward slicing as a kind of (abstract) re-execution or “replay”. Forward and backward slices are related by a Galois connection, ensuring the existence of minimal slices. This idea is straightforward in the sequential setting of the earlier work. However, generalising it to concurrent programs is non-trivial. Suppose we run a concurrent computation, discover a bug, and then wish to compute a dynamic slice. It would likely be prohibitively slow, or impractical for other reasons, to insist that the slice be computed using the exact interleaving of the original run. On the other hand, computing the slice using a brand-new concurrent execution may make different non-deterministic choices, producing a slice of a different computation than the one intended.

Intuitively, any execution which exhibits the same causal structure should be adequate for computing the slice, and any practical approach to concurrent slicing should take advantage of this. Cristescu et al [5] make a similar observation about reversible concurrency, arguing



© Roly Perera, Deepak Garg and James Cheney;

licensed under Creative Commons License CC-BY

Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

<i>Scheduler node 1</i>	<i>Scheduler node 2</i>	A_1	A_2
$\mathbf{a_1.c_1.(b_1.\overline{c_2.r_1} + \overline{c_2}.b_1.\overline{r_1})}$	$\overline{c_1}.a_2.c_2.(b_2.\overline{c_1.r_2} + \overline{c_1}.b_2.\overline{r_2})$	$\overline{\mathbf{a_1}}.\overline{b_1}.\overline{p_1}$	$\overline{a_2}.\overline{b_1}.\overline{p_2}$
→ $\mathbf{c_1.(b_1.\overline{c_2.r_1} + \overline{c_2}.b_1.\overline{r_1})}$	$\overline{c_1}.a_2.c_2.(b_2.\overline{c_1.r_2} + \overline{c_1}.b_2.\overline{r_2})$	$\overline{b_1}.\overline{p_1}$	$\overline{a_2}.\overline{b_1}.\overline{p_2}$
→ $b_1.\overline{c_2.r_1} + \overline{c_2}.b_1.\overline{r_1}$	$\mathbf{a_2.c_2.(b_2.\overline{c_1.r_2} + \overline{c_1}.b_2.\overline{r_2})}$	$\overline{b_1}.\overline{p_1}$	$\overline{\mathbf{a_2}}.\overline{b_1}.\overline{p_2}$
→ $\mathbf{b_1.\overline{c_2.r_1} + \overline{c_2}.b_1.\overline{r_1}}$	$c_2.(b_2.\overline{c_1.r_2} + \overline{c_1}.b_2.\overline{r_2})$	$\overline{b_1}.\overline{p_1}$	$\overline{\mathbf{b_1}}.\overline{p_2}$
→ $\overline{c_2}.\overline{r_1}$	$\mathbf{c_2.(b_2.\overline{c_1.r_2} + \overline{c_1}.b_2.\overline{r_2})}$	$\overline{b_1}.\overline{p_1}$	$\overline{a_2}.\overline{b_1}.\overline{p_2}$
→ $a_1.c_1.(b_1.\overline{c_2.r_1} + \overline{c_2}.b_1.\overline{r_1})$	$b_2.\overline{c_1.r_2} + \overline{c_1}.b_2.\overline{r_2}$	$\overline{b_1}.\overline{p_1}$	$\overline{a_2}.\overline{b_1}.\overline{p_2}$

■ **Figure 1** Stuck configuration, overlaid with backward slice with respect to final state of node 1

that the most liberal notion of reversibility is one that just respects causality: an action can be undone precisely after all the actions that causally depend on it have also been undone.

In this paper we give a formal account of dynamic slicing for π -calculus, and show that any causally equivalent execution generates precisely the same slicing information. We do this by formalising dynamic slicing with respect to a particular execution \tilde{t} , and then proving that slicing with respect to any causally equivalent computation \tilde{u} yields the same slice, after a unique “rewiring” which interprets the path witnessing $\tilde{t} \simeq \tilde{u}$ as a lattice isomorphism relating the two slices. The isomorphism is constructive, rewriting one slice into the other: this allows non-deterministic metadata (e.g. memory addresses or transaction ids) in the slicing execution to be aligned with the corresponding metadata in the original run. To achieve this we build on an earlier “proof-relevant” formalisation of causal equivalence for π -calculus in Agda [17]. As long as causality is respected, an implementation of our system can safely use any technique (e.g. redex trails, proved transitions, or thread-local memories) to implement rewind and replay.

Example: scheduler with non-compliant task. While dynamic slicing cannot automatically isolate bugs, it can hide irrelevant detail and yield compact provenance-like explanations of troublesome parts of configurations. As an example we consider Milner’s scheduler implementation [14, p. 65]. The scheduler controls the repeated performance of a set of n tasks, executed by agents A_1, \dots, A_n . Agent A_i sends the message a_i (*announce*) to the scheduler to start its task, and message b_i (*break*) to end its task. The scheduler ensures that the actions a_i occur cyclically starting with a_1 , and that for each i the actions a_i and b_i alternate, starting with a_i . Although started sequentially, once started the tasks are free to execute in parallel.

Figure 1 shows five transitions of a two-node scheduler, with the redex selected at each step highlighted in bold. The parts of the configuration which contribute to the final state of node 1 are in black; the grey parts are discarded by our backward-slicing algorithm. Assume prefixing binds more tightly than either \cdot or $+$. To save space, we omit the ν -binders defining the various names, and write $x.0$ simply as x . The names r_1 , r_2 , p_1 and p_2 are used to make recursive calls [14, p. 94]: a recursive procedure is implemented as a server which waits for an invocation request, spawns a new copy of the procedure body, and then returns to the wait state. Here we omit the server definitions, and simply replace a successful invocation by the spawned body; thus in the final step of Figure 1, after the synchronisation on c_2 the invocation $\overline{r_1}$ is replaced by a fresh copy of the initial state of scheduler node 1.

The final state of Figure 1 has no redexes, and so is stuck. The slice helps highlight the fact that by the time we come to start the second loop of scheduler 1, the task was terminated by message $\overline{b_1}$ from A_2 , before any such message could be sent by A_1 . We can understand the slice of the initial configuration (computed by “rewinding”, or backward-slicing) as *sufficient* to explain the slice of the stuck configuration by noting that the former is able to

compute the latter by “replay”, or forward-slicing. In other words, writing a sliced part of the configuration as \square , and pretending the holes \square are sub-computations which diverge, we can derive

$$\mathbf{a}_1.c_1.(b_1.\square + \square) \mid \overline{c_1}.a_2.\square \mid \overline{a_1}.\square \mid \overline{a_2}.\overline{b_1}.\overline{p_2} \longrightarrow^* \overline{a_2}.\square$$

without diverging. The slice on the left may of course choose to take the right-hand branch of the choice instead. But if we constrain the replay of the sliced program to follow the causal structure of the original unsliced run – to take the same branches of internal choices, and have the same synchronisation structure – then it will indeed evolve to the slice on the right. This illustrates the correctness property for backward slicing, which is that forward-slicing its result must recompute (at least) the slicing criterion.

For this example, the tasks are entirely atomic and so fixing the outcome of $+$ has the effect of making the computation completely sequential. Less trivial systems usually have multiple ways they can evolve, even once the causal structure is fixed. A confluence lemma typically formalises the observational equivalence of two causally equivalent runs. However, a key observation made in [17] is that requiring causally equivalent runs to reach exactly the same state is too restrictive for π -calculus, in particular because of name extrusion. As we discuss in §3, two causally unrelated extrusion-rendezvous lead to states which differ in the relative position of two ν -binders, reflecting the two possible orderings of the rendezvous. Although technically unobservable to the program, interleaving-sensitive metadata, such as memory locations in a debugger or transaction ids in a financial application, may be important for domain-specific reasons. In these situations being able to robustly translate between the target states of the two executions may be useful.

Summary of contributions. §2 defines the core forward and backward dynamic slicing operations for π -calculus transitions and sequences of transitions (traces). We prove that they are related by a Galois connection, showing that backward and forward slicing, as defined, are minimal and maximal with respect to each other. §3 extends this framework to show that the Galois connections for causally equivalent traces compute the same slices up to lattice isomorphism. §4 discusses related work and §5 offers closing thoughts and prospects for follow-up work. Appendix A summarises the Agda module structure and required libraries; the source code can be found at <https://github.com/rolyp/concurrent-slicing>, release 0.1.

2 Galois connections for slicing π -calculus programs

To summarise informally, our approach is to interpret, functorially, every transition diagram in the π -calculus into the category of lattices and Galois connections. For example the interpretation of the transition diagram on the left is the commutative diagram on the right:



where $\downarrow P$ means the lattice of slices of P , and $\overline{\text{step}}_t : \downarrow P \longrightarrow \downarrow Q$ is a *Galois connection*, a kind of generalised order isomorphism. An order isomorphism between posets A and B is a pair of monotone functions $f : A \longrightarrow B$ and $g : B \longrightarrow A$ such that $f \circ g = \text{id}_B$ and $g \circ f = \text{id}_A$. Galois connections require only $f \circ g \geq \text{id}_B$ and $g \circ f \leq \text{id}_A$ where \leq means the pointwise order. Galois connections are closed under composition.

The relationship to slicing is that these properties can be unpacked into statements of sufficiency and minimality: for example $f \circ g \geq \text{id}_B$ means g (backward-slicing) is “sufficient” in that f (forward-slicing) is able to use the result of g to restore the slicing criterion, and $g \circ f \leq \text{id}_A$ means g is “minimal” in that it computes the smallest slice with that property. One can dualise these statements to make similar observations about f .

We omit a treatment of structural congruence from our approach, but note that it slots easily into the framework, generating lattice isomorphisms in a manner similar to the “bound braid” relation \times discussed in §3, Definition 12.

2.1 Lattices of slices

The syntax of names, processes and actions is given in Figure 2. Slices are represented syntactically, via the \square notation introduced informally in §1. Our formalisation employs de Bruijn indices [7], an approach with well-known strengths and weaknesses compared to other approaches to names such as higher-order abstract syntax or nominal calculi.

Name	$x, y, z ::= \square$	erased	Process	$P, Q, R, S ::= \square$	erased
	$0 \mid 1 \mid \dots$			$\mathbf{0}$	inactive
Action	$a ::= \square$	erased		$\underline{x}.P$	input
	\underline{x}	input		$\bar{x}(y).P$	output
	$\bar{x}(y)$	output		$P + Q$	choice
	\bar{x}	bound output		$P \mid Q$	parallel
	τ	silent		νP	restriction
				$!P$	replication

■ **Figure 2** Syntax of names, processes and actions

Names. The erased name \square gives rise to a (trivial) partial order \leq over names, namely the least partial order containing $\square \leq x$ for any x . The set of *slices* of x is written $\downarrow x$ and defined to be $\{y \mid y \leq x\}$; because names are atomic $\downarrow x$ is simply the two-element set $\{\square, x\}$. The set $\downarrow x$ is a finite lattice with meet and join operations \sqcap and \sqcup , and top and bottom elements x and \square respectively. The meet and join are related to the partial order by $x \leq y \iff x \sqcup y = y \iff x \sqcap y = x$. Lattices are closed under component-wise products, justifying the notation $\downarrow(x, y)$ for $\downarrow x \times \downarrow y$.

Processes. The \leq relation and \downarrow operation extend to processes, via names which may be \square , and a special undefined process also written \square . A slice of P is simply P with some sub-processes replaced by \square . The relation \leq is the least compatible partial order which has \square as least element; all process constructors both preserve and reflect \leq , so we assume an equivalent inductive definition of \leq when convenient. A process has a closing context Γ enumerating its free variables; in the untyped de Bruijn setting Γ is just a natural number. Often it is convenient to conflate Γ with a set of that cardinality.

Actions. An action a labels a transition (Figure 3 below), and is either a *bound* or *non-bound*. A bound action b is of the form \underline{x} or \bar{x} and opens a process with respect to x , taking it from Γ to $\Gamma + 1$. A non-bound action c of the form $\bar{x}(y)$ or τ and preserves the free variables of the process. The \leq relation and \downarrow operation extend to actions via \square names, plus a special undefined action also written \square .

Renamings. In the lattice setting, a renaming $\rho : \Gamma \rightarrow \Gamma'$ is any function from Γ to $\Gamma' \uplus \{\square\}$; we also allow σ to range over renamings. Renaming application ρ^*P is extended with the equation $\rho^*\square = \square$. The \leq relation and \downarrow operation apply pointwise.

Labelled transition semantics. The late-style labelled transition semantics is given in

$$\begin{array}{c}
\frac{}{\underline{x}.P \xrightarrow{x} P} \quad \frac{}{\bar{x}(y).P \xrightarrow{\bar{x}(y)} P} \quad \frac{P \xrightarrow{a} R}{P + Q \xrightarrow{a} R} \quad \frac{P \xrightarrow{c} R}{P | Q \xrightarrow{c} R | Q} \\
(*) \frac{P \xrightarrow{b} R}{P | Q \xrightarrow{b} R | \text{push}^* Q} \quad (\S) \frac{P \xrightarrow{x} R \quad Q \xrightarrow{\bar{x}(y)} S}{P | Q \xrightarrow{\tau} (\text{pop } y)^* R | S} \quad \frac{P \xrightarrow{(x+1)(0)} R}{\nu P \xrightarrow{\bar{x}} R} \\
\frac{P \xrightarrow{x} R \quad Q \xrightarrow{\bar{x}} S}{P | Q \xrightarrow{\tau} \nu(R | S)} \quad (\dagger) \frac{P \xrightarrow{\text{push}^* c} R}{\nu P \xrightarrow{c} \nu R} \quad (\#) \frac{P \xrightarrow{\text{push}^* b} R}{\nu P \xrightarrow{b} \nu(\text{swap}^* R)} \quad \frac{P | !P \xrightarrow{a} R}{!P \xrightarrow{a} R}
\end{array}$$

■ **Figure 3** Labelled transition relation $P \xrightarrow{a} R$ (symmetric variants omitted)

Figure 3, and is distinguished only by its adaptation to the de Bruijn setting. The primary reference for a de Bruijn formulation of π -calculus is [11]; the consequences such an approach are explored in some depth in [17]. One pleasing consequence of a de Bruijn approach is that the usual side-conditions associated with transition rules can be operationalised via renamings. We briefly explain this, along with other uses of renamings in the transition rules, and refer the interested reader to these earlier works for more details. Definition 1 defines the renamings used in Figure 3 and Definition 2 the application $\rho^* a$ of ρ to an action a .

► **Definition 1** (push, pop, and swap).

$$\begin{array}{lll}
\text{push}_\Gamma & : & \Gamma \longrightarrow \Gamma + 1 \\
\text{pop}_\Gamma \mathbf{y} & : & \Gamma + 1 \longrightarrow \Gamma \\
\text{pop } x & = & x + 1 \\
\text{pop } y \ 0 & = & y \\
\text{pop } y \ (x + 1) & = & x \\
\text{swap}_\Gamma & : & \Gamma + 2 \longrightarrow \Gamma + 2 \\
\text{swap } 0 & = & 1 \\
\text{swap } 1 & = & 0 \\
\text{swap } (x + 2) & = & x + 2
\end{array}$$

► **Definition 2** (Action renaming). Define the following lifting of a renaming to actions.

$$\begin{array}{ll}
\rho^* & : (\Gamma \longrightarrow \Gamma') \longrightarrow \text{Action } \Gamma \longrightarrow \text{Action } \Gamma' \\
\rho^* \square & = \square \\
\rho^* x & = \rho x \\
\rho^* \bar{x} & = \overline{\rho x} \\
\rho^* \tau & = \tau \\
\rho^* \bar{x}(y) & = \overline{\rho x}(\rho y)
\end{array}$$

- push occurs in the transition rule which propagates a bound action through a parallel composition $P | Q$ (rule $(*)$ in Figure 3), and rewires Q so that the name 0 is reserved. The effect is to ensure that the binder being propagated by P is not free in Q .
- push also occurs in the rules which propagate an action through a ν -binder (rules (\dagger) and $(\#)$), where it is applied to the action being propagated using the function defined in Definition 2. This ensures the action does not mention the binder it is propagating through. The use of $\cdot + 1$ in the name extrusion rule can be interpreted similarly.
- pop y is used in the event of a successful synchronisation (rule (\S)), and undoes the effect of push, substituting the communicated name y for index 0.
- swap occurs in the rule which propagates a bound action through a ν -binder (rule (\dagger)) and has no counterpart outside of the de Bruijn setting. As a propagating binder passes through another binder, their relative position in the syntax is exchanged, and so to preserve naming R is rewired with a “braid” that swaps 0 and 1.

Although its use in the operational semantics is unique to the de Bruijn setting, swap will also play an important role when we consider the relationship between slices of causally

equivalent traces (§3 below), where it captures how the relative position of binders changes between different (but causally equivalent) interleavings.

2.2 Galois connections for slicing

We now compositionally assemble a Galois connection for each component of execution, starting with renamings, and then proceeding to individual transitions and entire traces, which relates forward and backward slices of the initial and terminal state.

Slicing renamings. The application ρx of a renaming to a name, and the lifting ρ^*P of that operation to a process give rise to the Galois connections defined here.

► **Definition 3** (Galois connection for ρx). Suppose $\rho : \Gamma \longrightarrow \Gamma'$ and $x \in \Gamma$. Define the following pair of monotone functions between $\downarrow(\rho, x)$ and $\downarrow(\rho x)$.

$$\begin{array}{ll} \text{app}_{\rho,x} & : \downarrow(\rho, x) \longrightarrow \downarrow(\rho x) \\ \text{app}_{\rho,x}(\sigma, \square) & = \square \\ \text{app}_{\rho,x}(\sigma, x) & = \sigma x \end{array} \qquad \begin{array}{ll} \text{unapp}_{\rho,x} & : \downarrow(\rho x) \longrightarrow \downarrow(\rho, x) \\ \text{unapp}_{\rho,x} u & = (x \mapsto_{\rho} u, \rho_x^{-1} u) \end{array}$$

$$\begin{array}{ll} \text{where } x \mapsto_{\rho} \cdot & : \downarrow(\rho x) \longrightarrow \downarrow \rho \\ (x \mapsto_{\rho} u) x & = u \\ (x \mapsto_{\rho} u) y & = \square \text{ (if } y \neq x) \end{array} \qquad \begin{array}{ll} \rho_x^{-1} & : \downarrow(\rho x) \longrightarrow \downarrow x \\ \rho_x^{-1} \square & = \square \\ \rho_x^{-1} u & = x \text{ (if } u \neq \square) \end{array}$$

It is convenient to decompose $\text{unapp}_{\rho,x}$ into two components: $x \mapsto_{\rho} u$ denotes the least slice of ρ which maps x to u , and $\rho_x^{-1} u$ denotes the least slice of x such that $\rho x = u$.

► **Lemma 4.** ($\text{app}_{\rho,x}, \text{unapp}_{\rho,x}$) is a Galois connection.

1. $\text{app}_{\rho,x} \circ \text{unapp}_{\rho,x} \geq \text{id}_{\rho x}$
2. $\text{unapp}_{\rho,x} \circ \text{app}_{\rho,x} \leq \text{id}_{\rho,x}$

► **Definition 5** (Galois connection for a renaming ρ^*P).

Suppose $\rho : \Gamma \longrightarrow \Gamma'$ and $\Gamma \vdash P$. Define monotone functions between $\downarrow(\rho, P)$ and $\downarrow(\rho^*P)$ by structural recursion on $\downarrow P$, using the following equations. Here \square_{ρ} denotes the least slice of ρ , namely the renaming which maps every $x \in \Gamma$ to \square .

$$\begin{array}{ll} \text{ren}_{\rho,P} & : \downarrow(\rho, P) \longrightarrow \downarrow(\rho^*P) \\ \text{ren}_{\rho,P}(\sigma, \square) & = \square \\ \text{ren}_{\rho,0}(\sigma, \mathbf{0}) & = \mathbf{0} \\ \text{ren}_{\rho,\underline{x}.P}(\sigma, \underline{x}.R) & = \underline{x}'.\text{ren}_{\rho+1,P}(\sigma, R) \text{ where } \underline{x}' = \text{app}_{\rho,x}(\sigma, u) \\ \text{ren}_{\rho,\overline{x}(y).P}(\sigma, \overline{x}(y).R) & = \overline{x}'(y').\text{ren}_{\rho,P}(\sigma, R) \text{ where } \underline{x}' = \text{app}_{\rho,x}(\sigma, u) \text{ and } y' = \text{app}_{\rho,y}(\sigma, v) \\ \text{ren}_{\rho,P+Q}(\sigma, R + S) & = \text{ren}_{\rho,P}(\sigma, R) + \text{ren}_{\rho,Q}(\sigma, S) \\ \text{ren}_{\rho,P|Q}(\sigma, R | S) & = \text{ren}_{\rho,P}(\sigma, R) | \text{ren}_{\rho,Q}(\sigma, S) \\ \text{ren}_{\rho,\nu P}(\sigma, \nu R) & = \nu \text{ren}_{\rho+1,P}(\sigma + 1, R) \\ \text{ren}_{\rho,!P}(\sigma, !R) & = !(\text{ren}_{\rho,P}(\sigma, R)) \end{array}$$

$$\begin{array}{ll} \text{unren}_{\rho,P} & : \downarrow(\rho^*P) \longrightarrow \downarrow(\rho, P) \\ \text{unren}_{\rho,P} \square & = (\square_{\rho}, \square) \\ \text{unren}_{\rho,0} \mathbf{0} & = (\square_{\rho}, \mathbf{0}) \\ \text{unren}_{\rho,\underline{x}.P} \underline{x}.R & = (\rho' \sqcup (x \mapsto_{\rho} u), \rho_x^{-1} u.P') \text{ where } \text{unren}_{\rho+1,P} R = (\rho' + 1, P') \\ \text{unren}_{\rho,\overline{x}(y).P} \overline{x}(y).R & = (\rho' \sqcup (x \mapsto_{\rho} u) \sqcup (y \mapsto_{\rho} v), \overline{\rho_x^{-1} u}(\rho_y^{-1} v).P') \text{ where } \text{unren}_{\rho,P} R = (\rho', P') \\ \text{unren}_{\rho,P+Q} (R + S) & = (\rho_1 \sqcup \rho_2, P' + Q') \text{ where } \text{unren}_{\rho,P} R = (\rho_1, P') \text{ and } \text{unren}_{\rho,Q} S = (\rho_2, Q') \\ \text{unren}_{\rho,P|Q} (R | S) & = (\rho_1 \sqcup \rho_2, P' | Q') \text{ where } \text{unren}_{\rho,P} R = (\rho_1, P') \text{ and } \text{unren}_{\rho,Q} S = (\rho_2, Q') \\ \text{unren}_{\rho,\nu P} \nu R & = (\rho', \nu P') \text{ where } \text{unren}_{\rho+1,P} R = (\rho' + 1, P') \\ \text{unren}_{\rho,!P} !R & = (\rho', !P') \text{ where } \text{unren}_{\rho,P} R = (\rho', P') \end{array}$$

► **Lemma 6.** ($\text{ren}_{\rho,P}, \text{unren}_{\rho,P}$) is a Galois connection.

$$\begin{array}{c}
\frac{}{\Box_P \xrightarrow{\Box_a} \Box_{P'}} \quad \frac{}{u_x.R_P \xrightarrow{u_x} R_P} \quad \frac{}{u_x \langle v_y \rangle . R_P \xrightarrow{u_x \langle v_y \rangle} R_P} \quad \frac{R_P \xrightarrow{a'_a} R'_{P'}}{R_P + S_Q \xrightarrow{a'_a} R'_{P'}} \\
\frac{R_P \xrightarrow{c'_c} R'_{P'}}{R_P | S_Q \xrightarrow{c'_c} R'_{P'} | S_Q} \quad \frac{R_P \xrightarrow{b'_b} R'_{P'}}{R_P | S_Q \xrightarrow{b'_b} R'_{P'} | \text{push}^* S_Q} \quad \frac{R_P \xrightarrow{x} R'_{P'} \quad S_Q \xrightarrow{\bar{x}(y_z)} S'_{Q'}}{R_P | S_Q \xrightarrow{\tau} (\text{pop } y_z)^* R'_{P'} | S'_{Q'}} \\
\frac{R_P \xrightarrow{a_x} R'_{P'} \quad S_Q \xrightarrow{a'_x(z)} S'_{Q'}}{R_P | S_Q \xrightarrow{\Box_\tau} \Box_{(\text{pop } z)^* P' | Q'}} \quad \frac{R_P \xrightarrow{(x+1)(0)} R'_{P'}}{vR_P \xrightarrow{x} R'_{P'}} \quad \frac{R_P \xrightarrow{a_{(x+1)(0)}} R'_{P'}}{vR_P \xrightarrow{\Box_{\bar{x}}} \Box_{P'}} \\
\frac{R_P \xrightarrow{x} R'_{P'} \quad S_Q \xrightarrow{\bar{x}} S'_{Q'}}{R_P | S_Q \xrightarrow{\tau} v(R'_{P'} | S'_{Q'})} \quad \frac{R_P \xrightarrow{a_x} R'_{P'} \quad S_Q \xrightarrow{a'_x} S'_{Q'}}{R_P | S_Q \xrightarrow{\Box_\tau} \Box_{v(P' | Q')}} \quad \frac{R_P \xrightarrow{\text{push}^* c'_c} R'_{P'}}{vR_P \xrightarrow{c'_c} vR'_{P'}} \\
\frac{R_P \xrightarrow{c'_{\text{push}^* c}} R'_{P'}}{vR_P \xrightarrow{\Box_c} \Box_{vP'}} \quad \frac{R_P \xrightarrow{\text{push}^* b'_b} R'_{P'}}{vR_P \xrightarrow{b'_b} v(\text{swap}^* R'_{P'})} \quad \frac{R_P \xrightarrow{b'_{\text{push}^* b}} R'_{P'}}{vR_P \xrightarrow{\Box_b} \Box_{v(\text{swap}^* P')}} \quad \frac{R_P | !R_P \xrightarrow{a'_a} R'_{P'}}{!R_P \xrightarrow{a'_a} R'_{P'}}
\end{array}$$

■ **Figure 4** Forward slicing judgement $R_P \xrightarrow{a'_a} R'_{P'}$

1. $\text{ren}_{\rho, P} \circ \text{unren}_{\rho, P} \geq \text{id}_{\rho^* P}$
2. $\text{unren}_{\rho, P} \circ \text{ren}_{\rho, P} \leq \text{id}_{\rho, P}$

Proof. In each case by induction on P , using Lemma 4 and the invertibility of $\cdot + 1$. ◀

Slicing transitions. Transitions also lift to the lattice setting, in the form of Galois connections defined by structural recursion over the proof that $t : P \xrightarrow{a} P'$. Figures 4 and 5 define the forward and backward slicing judgements. As presented the rules are non-deterministic, but the reader should consider there to be implicit determinising side-conditions which prefer earlier rules over subsequent rules.

The judgement $R_P \xrightarrow{a'_a} R'_{P'}$ asserts that there is a “replay” transition from $R \leq P$ to $(a', R') \leq (a, P)$. In this case R is the input to the judgement and (a', R') is the output. The judgement $R'_P \xleftarrow{a'_a} R_P$ asserts that there is an “rewind” transition from $(a', R) \leq (a, P')$ to $R' \leq P$. In this case (a', R) is the input and R' the output. When writing R_P where $R \leq P$ we exploit the preservation and reflection of \leq by all constructors in the notation, writing $v(R_P | S_Q)$ for $v(R | S)_{v(P|Q)}$, for example.

As a further notational convenience, we permit the renaming application operator $*$ to appear in a pattern-matching (input) position of the backward-slicing judgement, indicating a use of the lower adjoint unren . Specifically, given a renaming application $\rho^* P$, the pattern $\sigma^* P'$ matches any slice R of $\rho^* P$ such that $\text{unren}_{\rho, P}(R) = (\sigma, P')$.

► **Definition 7** (Galois connection for a transition). Suppose $t : P \xrightarrow{a} P'$. Define the following pair of monotone functions between $\downarrow P$ to $\downarrow(a, P')$.

$$\begin{array}{ll}
\text{step}_t & : \downarrow P \longrightarrow \downarrow(a, P') \\
\text{step}_t \ R & = (a', R') \text{ where } R_P \xrightarrow{a'_a} R'_{P'} \\
\text{unstep}_t & : \downarrow(a, P') \longrightarrow \downarrow P \\
\text{unstep}_t \ (R, a') & = R' \text{ where } R'_P \xleftarrow{a'_a} R_P
\end{array}$$

We omit the proofs that these equations indeed define total, deterministic, monotone relations.

► **Theorem 8** ($(\text{step}_t, \text{unstep}_t)$ is a Galois connection).

1. $\text{step}_t \circ \text{unstep}_t \geq \text{id}_{\downarrow(a, P')}$
2. $\text{unstep}_t \circ \text{step}_t \leq \text{id}_{\downarrow P}$

$$\begin{array}{c}
\frac{}{\square_P \xleftarrow{\square_a} \square_{P'}} \quad \frac{}{\underline{u}_x.R_P \xleftarrow{\underline{u}_x} R_P} \quad \frac{}{\square_x.R_P \xleftarrow{\square_x} R_P} \quad \frac{}{\overline{u}_x \langle v_y \rangle . R_P \xleftarrow{\overline{u}_x \langle v_y \rangle} R_P} \\
\frac{}{\square_x \langle \square_y \rangle . R_P \xleftarrow{\square_x \langle \square_y \rangle} R_P} \quad \frac{R'_P \xleftarrow{a'_a} R_{P'}}{R'_P + \square_Q \xleftarrow{a'_a} R_{P'}} \quad \frac{R'_P \xleftarrow{c'_c} R_{P'}}{R'_P | S_Q \xleftarrow{c'_c} R_{P'} | S_Q} \quad \frac{R'_P \xleftarrow{c'_c} \square_{P'}}{R'_P | \square_Q \xleftarrow{c'_c} \square_{P'} | \square_Q} \\
\frac{R'_P \xleftarrow{b'_b} R_{P'}}{R'_P | S_Q \xleftarrow{b'_b} R_{P'} | \rho_{\text{push}^*} S_Q} \quad \frac{R'_P \xleftarrow{b'_b} \square_{P'}}{R'_P | \square_Q \xleftarrow{b'_b} \square_{P'} | \rho_{\text{push}^*} S_Q} \quad \frac{R'_P \xleftarrow{x} R_{P'} \quad S'_Q \xleftarrow{\bar{x}(\rho 0_z)} S_{Q'}}{R'_P | S'_Q \xleftarrow{a_\tau} \rho_{\text{pop } z}^* R_{P'} | S_{Q'}} \\
\frac{R_P \xleftarrow{x} \square_{P'} \quad S_Q \xleftarrow{\bar{x}(\square_z)} \square_{Q'}}{R_P | S_Q \xleftarrow{\tau} \square_{(\text{pop } z)^* P'} | \square_{Q'}} \quad \frac{R'_P \xleftarrow{\overline{(x+1)}(0)} R_{P'}}{\nu R'_P \xleftarrow{a_{\bar{x}}} R_{P'}} \quad \frac{R'_P \xleftarrow{x} R_{P'} \quad S'_Q \xleftarrow{\bar{x}} S_{Q'}}{R'_P | S'_Q \xleftarrow{a_\tau} \nu(R_{P'} | S_{Q'})} \\
\frac{R_P \xleftarrow{x} \square_{P'} \quad S_Q \xleftarrow{\bar{x}} \square_{Q'}}{R_P | S_Q \xleftarrow{a_\tau} \nu \square_{P'} | \square_{Q'}} \quad \frac{R_P \xleftarrow{x} \square_{P'} \quad S_Q \xleftarrow{\bar{x}} \square_{Q'}}{R_P | S_Q \xleftarrow{\tau} \square_{\nu(P' | Q')}} \quad \frac{R'_P \xleftarrow{\text{push}^* c} R_{P'}}{\nu R'_P \xleftarrow{c'_c} \nu R_{P'}} \\
\frac{R'_P \xleftarrow{\text{push}^* c} \square_{P'}}{\nu R'_P \xleftarrow{c'_c} \square_{\nu P'}} \quad \frac{R'_P \xleftarrow{\text{push}^* b} R_{P'}}{\nu R'_P \xleftarrow{b'_b} \nu(\rho_{\text{swap}}^* R_{P'})} \quad \frac{R'_P \xleftarrow{\text{push}^* b} \square_{P'}}{\nu R'_P \xleftarrow{b'_b} \square_{\nu(\text{swap}^* P')}} \quad \frac{R'_P | R''_P \xleftarrow{a'_a} R_{P'}}{(!R' \sqcup R'')_{!P} \xleftarrow{a'_a} R_{P'}}
\end{array}$$

■ **Figure 5** Backward slicing judgement $R'_P \xleftarrow{a'_a} R_{P'}$

Proof. By induction on $t : P \xrightarrow{a'} P'$, using Lemma 6 for the cases involving renaming. ◀

Slicing traces. Finally we extend slicing to entire runs of a π -calculus program. A sequence of transitions \tilde{t} is called a *trace*; the empty trace at P is written ε_P , and the composition of a transition $t : P \xrightarrow{a} R$ and trace $\tilde{t} : R \xrightarrow{\tilde{a}} S$ is written $t \cdot \tilde{t} : P \xrightarrow{a \cdot \tilde{a}} S$ where actions are composable whenever their source and target contexts match.

► **Definition 9** (Galois connection for a trace \tilde{t}). Suppose $\tilde{t} : P \xrightarrow{\tilde{a}} P'$. Define the following pair of monotone functions between $\downarrow P$ and $\downarrow P'$, which use variants of step_t and unstep_t which discard the action slice (going forward) and which use \square as the action slice (going backward).

$$\begin{array}{l}
\text{fwd}_{\tilde{t}} \quad : \quad \downarrow P \longrightarrow \downarrow P' \\
\text{fwd}_{\varepsilon_P} \quad = \quad \text{id}_{\downarrow P} \\
\text{fwd}_{t \cdot \tilde{t} \square} \quad = \quad \square \\
\text{fwd}_{t \cdot \tilde{t} R} \quad = \quad \text{fwd}_{\tilde{t}}(\text{step}'_t R) \quad (R \neq \square) \\
\text{step}'_t \quad : \quad \downarrow P \longrightarrow \downarrow P' \\
\text{step}'_t R \quad = \quad R' \text{ where } \text{step}_t R = (a', R') \\
\text{bwd}_{\tilde{t}} \quad : \quad \downarrow P' \longrightarrow \downarrow P \\
\text{bwd}_{\varepsilon_{P'}} \quad = \quad \text{id}_{\downarrow P'} \\
\text{bwd}_{t \cdot \tilde{t} \square} \quad = \quad \square \\
\text{bwd}_{t \cdot \tilde{t} R} \quad = \quad \text{unstep}'_t(\text{bwd}_{\tilde{t}} R) \quad (R \neq \square) \\
\text{unstep}'_t \quad : \quad \downarrow P' \longrightarrow \downarrow P \\
\text{unstep}'_t R' \quad = \quad \text{unstep}_t(\square, R')
\end{array}$$

At the empty trace ε_P the Galois connection is simply the identity on $\downarrow P$. Otherwise, we recurse into the structure of the trace $t \cdot \tilde{t}$, composing the Galois connection for the single transition t with the Galois connection for the tail of the computation \tilde{t} .

► **Theorem 10** ($(\text{fwd}_{\tilde{t}}, \text{bwd}_{\tilde{t}})$ is a Galois connection).

1. $\text{fwd}_{\tilde{t}} \circ \text{bwd}_{\tilde{t}} \geq \text{id}_{P'}$
2. $\text{bwd}_{\tilde{t}} \circ \text{fwd}_{\tilde{t}} \leq \text{id}_P$

Note that the trace used to define forward and backward slicing for a computation is not an auxiliary data structure recording the computation, such as a redex trail or memory, but simply the proof term witnessing $P \xrightarrow{\tilde{a}} P'$.

3 Slicing and causal equivalence

In this section, we show that when dynamic slicing a π -calculus program, slicing with respect to any causally equivalent execution yields essentially the same slice. “Essentially the same” here means modulo lattice isomorphism. In other words slicing discards the same quantity of information regardless of which interleaving is chosen to do the slicing.

Proof-relevant causal equivalence. Causally equivalent computations are generated by transitions which share a start state, but which are independent. Following Lévy [13], we call such transitions *concurrent*, written $t \smile t'$. We illustrate this idea, and the non-trivial relationship that it induces between terminal states, by way of example. For the full definition of concurrency for π -calculus, we refer the interested reader to [17] or to the Agda definition¹. For the sake of familiarity the example uses regular names instead of de Bruijn indices.

Example. Consider the process $P_0 \stackrel{\text{def}}{=} (\nu yz) (\bar{x}(y).P) \mid \bar{x}(z).Q$ for some unspecified processes P and Q . This process can take *two* transitions, which we will call t and t' . Transition $t : P_0 \xrightarrow{\bar{x}(y)} P_1$ extrudes y on the channel x :

$$P_0 \xrightarrow{\bar{x}(y)} (\nu z) P \mid \bar{x}(z).Q \stackrel{\text{def}}{=} P_1$$

whereas transition $t' : P_0 \xrightarrow{\bar{x}(z)} P'_1$ extrudes z , also on the channel x :

$$P_0 \xrightarrow{\bar{x}(z)} (\nu y) (\bar{x}(y).P) \mid Q \stackrel{\text{def}}{=} P'_1$$

In both cases the output actions are bound, representing the extruding binder. Moreover, t and t' are *concurrent*, written $t \smile t'$, meaning they can be executed in either order. Having taken t , one can *mutatis mutandis* take t' , and vice versa. Concurrency is an irreflexive and symmetric relation defined over transitions which are *coinitial* (have the same source state).

The qualification is needed because t' will need to be adjusted to operate on the target state of t , if t is the transition which happens first. If t' happens first then t will need to be adjusted to operate on the target state of t' . The adjusted version of t' is called the *residual* of t' after t , and is written t'/t . In this case t'/t can still extrude z :

$$P_1 = (\nu z) P \mid \bar{x}(z).Q \xrightarrow{\bar{x}(z)} P \mid Q \stackrel{\text{def}}{=} P'_0$$

whereas the residual t/t' can still extrude y :

$$P'_1 = (\nu y) (\bar{x}(y).P) \mid Q \xrightarrow{\bar{x}(y)} P \mid Q = P'_0$$

The independence of t and t' is confirmed by the fact that $t \cdot t'/t$ and $t' \cdot t/t'$ are *cofinal* (share a target state), as shown on the left below.



We say that the traces $\tilde{t} \stackrel{\text{def}}{=} t \cdot t'/t$ and $\tilde{u} \stackrel{\text{def}}{=} t' \cdot t/t'$ are *causally equivalent*, written $\tilde{t} \simeq \tilde{u}$. The two interleavings are also equivalent from a slicing point of view, in the sense that the

¹ <https://github.com/rolyp/proof-relevant-pi/blob/master/Transition/Concur.agda>

XX:10 Causally consistent dynamic slicing

right-hand square commutes (Theorem 16 below). Here $\overline{\text{step}}_t$ denotes the Galois connection $(\text{step}_t, \text{unstep}_t)$.

However [17], which formalised causal equivalence for π -calculus, showed that causally equivalence traces do not always reach exactly the same state, but only the same state up to some permutation of the binders in the resulting processes. This will become clear if we consider another process $Q_0 \stackrel{\text{def}}{=} (x(y').R) \mid x(z').S$ able to synchronise with both of the extrusions raised by P_0 and consider the two different ways that $P_0 \mid Q_0$ can evolve.

First note that Q_0 can also take two independent transitions: $u : Q_0 \xrightarrow{x(y')} R \mid x(z').S \stackrel{\text{def}}{=} Q_1$ inputs on x and binds the received name to y' ; and $u' : Q_0 \xrightarrow{x(z')} (x(y').R) \mid S \stackrel{\text{def}}{=} Q'_1$ also inputs on x and binds the received name to z' . (Assume z is not free in the left-hand side of Q_0 and that y is not free in the right-hand side.) The respective residuals $Q_1 = R \mid x(z').S \xrightarrow{x(z')} R \mid S \stackrel{\text{def}}{=} Q'_0$ and $Q'_1 = (x(y').R) \mid S \xrightarrow{x(y')} R \mid S = Q'_0$ again converge on the same state Q'_0 , leading to a diamond for Q_0 similar to the one for P_0 above.

The subtlety arises when we put P_0 and Q_0 into parallel composition, since now we have two concurrent synchronisation possibilities. For clarity we give the derivations, which we call s and s' :

$$\frac{t : P_0 \xrightarrow{\bar{x}(y)} P_1 \quad u : Q_0 \xrightarrow{x(y')} Q_1}{s : P_0 \mid Q_0 \xrightarrow{\tau} (\nu y) P_1 \mid Q_1\{y/y'\}} \quad \frac{t' : P_0 \xrightarrow{\bar{x}(z)} P'_1 \quad u' : Q_0 \xrightarrow{x(z')} Q'_1}{s' : P_0 \mid Q_0 \xrightarrow{\tau} (\nu z) P'_1 \mid Q'_1\{z/z'\}}$$


The labelled transition system is closed under renamings; thus the residual u'/u has an image in the renaming $\cdot\{y/y'\}$, and u/u' has an image in the renaming $\cdot\{z/z'\}$, allowing us to derive composite residual s'/s :

$$\frac{t'/t : P_1 \xrightarrow{\bar{x}(z)} P'_0 \quad \frac{u'/u : Q_1 \xrightarrow{x(z')} Q'_0}{(u'/u)\{y/y'\} : Q_1\{y/y'\} \xrightarrow{x(z')} Q'_0\{y/y'\}}}{s'/s : (\nu y) P_1 \mid Q_1\{y/y'\} \xrightarrow{\tau} (\nu yz) P'_0 \mid Q'_0\{y/y'\}\{z/z'\}}$$

By similar reasoning we can derive s/s' :

$$s/s' : (\nu z) P'_1 \mid Q'_1\{y/y'\} \xrightarrow{\tau} (\nu zy) P'_0 \mid Q'_0\{z/z'\}\{y/y'\}$$

By side-conditions on the transition rules the renamings $\cdot\{y/y'\}$ and $\cdot\{z/z'\}$ commute and so $Q'_0\{y/y'\}\{z/z'\} \stackrel{\text{def}}{=} Q''_0 = Q'_0\{z/z'\}\{y/y'\}$. However, the positions of binders y and z are transposed in the terminal states of s'/s and s/s' . Instead of the usual diamond shape, we have the pentagon on the left below, where ϕ is a *braid* representing the transposition of the binders. Lifted to slices, ϕ becomes the unique isomorphism $\overline{\text{braid}}_\phi$ relating slices of the terminal states, as shown in the commutative diagram on the right:



In the de Bruijn setting, a braid like ϕ does not relate two processes of the form $(\nu yz) R$ and $(\nu zy) R$ but rather two processes of the form $\nu\nu R$ and $\nu\nu(\text{swap}^* R)$: the transposition of the (nameless) binders is represented by the transposition of the roles of indices 0 and 1 in the body of the innermost binder.

► **Definition 11** (Bound braid $P \times R$). Inductively define the symmetric relation $P \times R$ using the rules below.

$$\begin{array}{l} \text{vv-swap}_P \frac{}{vvP \times vvP'} P = \text{swap}^* P' \quad \cdot + Q \frac{P \times R}{P + Q \times R + Q} \quad P + \cdot \frac{Q \times S}{P + Q \times P + S} \\ \cdot | Q \frac{P \times R}{P | Q \times R | Q} \quad P | \cdot \frac{Q \times S}{P | Q \times P | S} \quad \cdot v \cdot \frac{P \times R}{vP \times vR} \quad ! \cdot \frac{P \times R}{!P \times !R} \end{array}$$

Following [17], we adopt a compact term-like notation for \times proofs, using the rule names which occur to the left of each rule in Definition 11. For the extrusion example above, ϕ (in de Bruijn indices notation) would be a leaf case of the form $\text{vv-swap}_{\cdot|}$.

► **Definition 12** (Lattice isomorphism for bound braid). Suppose $\phi : Q \times Q'$. Define the following pair of monotone functions between $\downarrow Q$ and $\downarrow Q'$ by structural recursion on ϕ .

$$\begin{array}{ll} \text{braid}_{\phi} : \downarrow Q \longrightarrow \downarrow Q' & \text{unbraid}_{\phi} : \downarrow Q' \longrightarrow \downarrow Q \\ \text{braid}_{\text{vv-swap}_P} (\text{vv}R) = \text{vv}(\text{ren}_{\text{swap},P}(R)) & \text{unbraid}_{\text{vv-swap}_P} (\text{vv}R) = \text{vv}(\text{ren}_{\text{swap},P}(R)) \\ \text{braid}_{\phi+S} (R+S) = \text{braid}_{\phi} R + S & \text{unbraid}_{\phi+S} (R+S) = \text{unbraid}_{\phi} R + S \\ \text{braid}_{R+\psi} (R+S) = R + \text{braid}_{\psi} S & \text{unbraid}_{R+\psi} (R+S) = R + \text{unbraid}_{\psi} S \\ \text{braid}_{\phi|S} (R|S) = \text{braid}_{\phi} R | S & \text{unbraid}_{\phi|S} (R|S) = \text{unbraid}_{\phi} R | S \\ \text{braid}_{R|\psi} (R|S) = R | \text{braid}_{\psi} S & \text{unbraid}_{R|\psi} (R|S) = R | \text{unbraid}_{\psi} S \\ \text{braid}_{v\phi} (vR) = v(\text{braid}_{\phi} R) & \text{unbraid}_{v\phi} (vR) = v(\text{unbraid}_{\phi} R) \\ \text{braid}_{!\phi} (!R) = !(\text{braid}_{\phi} R) & \text{unbraid}_{!\phi} (!R) = !(\text{unbraid}_{\phi} R) \end{array}$$

► **Lemma 13.**

1. $\text{braid}_{\phi} \circ \text{unbraid}_{\phi} = \text{id}_{\downarrow Q'}$
2. $\text{unbraid}_{\phi} \circ \text{braid}_{\phi} = \text{id}_{\downarrow Q}$

Proof. Induction on ϕ . In the base case use the idempotence of swap lifted to lattices. ◀


► **Definition 14** (Lattice isomorphism for cofinality map). Suppose $t \smile t'$ with $\text{tgt}(t'/t) = Q$ and $\text{tgt}(t/t') = Q'$. By Theorem 1 of [17], there exists a unique $\gamma_{t,t'}$ witnessing $Q = Q'$, $Q \times Q'$ or $Q \times Q'$. Define the following pair of monotone functions between $\downarrow Q$ and $\downarrow Q'$.

$$\begin{array}{ll} \text{map}_{\gamma_{t,t'}} : \downarrow Q \longrightarrow \downarrow Q' & \text{unmap}_{\gamma_{t,t'}} : \downarrow Q' \longrightarrow \downarrow Q \\ \text{map}_{Q=Q'} = \text{id}_{\downarrow Q} & \text{unmap}_{Q=Q'} = \text{id}_{\downarrow Q} \\ \text{map}_{Q \times Q'} = \text{ren}_{\text{swap},Q} & \text{unmap}_{Q \times Q'} = \text{unren}_{\text{swap},Q} \\ \text{map}_{\phi:Q \times Q'} = \text{braid}_{\phi} & \text{unmap}_{\phi:Q \times Q'} = \text{unbraid}_{\phi} \end{array}$$

► **Lemma 15.**

1. $\text{map}_{\gamma_{t,t'}} \circ \text{unmap}_{\gamma_{t,t'}} = \text{id}_{\downarrow Q'}$
2. $\text{unmap}_{\gamma_{t,t'}} \circ \text{map}_{\gamma_{t,t'}} = \text{id}_{\downarrow Q}$

► **Theorem 16.** Suppose $t \smile t'$ as on the left. Then the pentagon on the right commutes.



Lattice isomorphism for arbitrary causal equivalence. Concurrent transitions $t \smile t'$ induce an “atom” of causal equivalence, $t \cdot t'/t \simeq t' \cdot t/t'$. The full relation is generated by closing under the trace constructors (for horizontal composition) and transitivity (for vertical composition). In [17] this yields a composite form of cofinality map γ_α where $\alpha : \tilde{t} \simeq \tilde{u}$ is an arbitrary causal equivalence. We omit further discussion for reasons of space, but note that γ_α is built by composing and translating (by contexts) atomic cofinality maps, and so gives rise, by composition of isomorphisms, to a lattice isomorphism between $\downarrow\text{tgt}(\tilde{t})$ and $\downarrow\text{tgt}(\tilde{u})$.

4 Related work

Reversible process calculi. Reversible process calculi have recently been used for speculative execution, debugging, transactions, and distributed protocols that require backtracking. A key challenge is to permit backwards execution to leverage concurrency whilst ensuring causal consistency. In contrast to our work, reversible calculi focus on mechanisms for reversibility, such as the thread-local memories used by Danos and Krivine’s reversible CCS [6], Lanese et al’s $\rho\pi$ [12], and Cristescu et al’s reversible π -calculus [5]. We intentionally remain agnostic about implementation strategy, whilst providing a formal guarantee that causally consistent rewind and replay are suitable foundation for any implementation.

Concurrent program slicing. An early example of concurrent dynamic slicing is Duesterwald et al, who consider a language with synchronous message-passing [9]. They give a notion of correctness with respect to a slicing criterion, but find that computing least slices is undecidable. Following Cheng’s [4], most subsequent work has recast dynamic slicing as a dependency-graph reachability problem; we take a different approach, namely to slice with respect to a particular execution but show how to map the result to the slice that would be obtained from another execution with the same dependency structure.

Moreover, there is a notable lack of correctness and minimality properties for systems for concurrent dynamic slicing. Goswami and Mall consider a language with shared-memory concurrency [10], and Mohapatra et al tackle slicing for concurrent Java [15], but both present only algorithms, with no formal guarantees. Tallam et al develop an approach based on dependency graphs, but again offer only algorithms and empirical results [19]. Finally, most prior work restricts the slicing criteria to the (entire) values of particular variables, rather than arbitrary parts of configurations.

Provenance and slicing. Our interest in slicing [16, 3] arises in part due to connections with provenance and dependency-tracking [2], and recent applications of provenance to security [1]. Others have also considered provenance models in concurrency calculi, including Souliah et al [18] and Dezani-Ciancaglini et al [8]. We have focused on formalising slicing for π -calculus using Galois connections; further study is needed to relate this model to provenance and security.

5 Conclusion

The main contribution of this paper is to extend our previous approach to slicing based on Galois connections to a concurrent setting, and show that the resulting notion of slicing for the π -calculus is invariant under causal equivalence of traces. Because of the complexity of the underlying definitions, we have built upon a previous formalization of causal equivalence for the π -calculus in Agda [17]. Although this makes the resulting definitions less transparent (as usual with de Bruijn indices), this is a small price to pay for definitions that are correct by construction. This paper provides a foundation for future development of rigorous provenance

tracing or dynamic slicing techniques for practical concurrent programs, which we plan to investigate in future work.

References

- 1 Umut A. Acar, Amal Ahmed, James Cheney, and Roly Perera. A core calculus for provenance. *Journal of Computer Security*, 21:919–969, 2013. Full version of a POST 2012 paper.
- 2 James Cheney, Amal Ahmed, and Umut A. Acar. Provenance as dependency analysis. *Mathematical Structures in Computer Science*, 21(6):1301–1337, 2011.
- 3 James Cheney, Amal Ahmed, and Umut A. Acar. Database queries that explain their work. In *PPDP 2014*, pages 271–282, 2014.
- 4 Jingde Cheng. Slicing concurrent programs: A graph-theoretical approach. In *Automated and Algorithmic Debugging*, number 749 in LNCS, pages 223–240. Springer-Verlag, 1993.
- 5 Ioana Cristescu, Jean Krivine, and Daniele Varacca. A compositional semantics for the reversible pi-calculus. In *LICS*, pages 388–397, June 2013.
- 6 Vincent Danos and Jean Krivine. Reversible communicating systems. In Philippa Gardner and Nobuko Yoshida, editors, *CONCUR*, volume 3170 of LNCS, pages 292–307. Springer, 2004. doi:10.1007/978-3-540-28644-8_19.
- 7 N.G. de Bruijn. Lambda-calculus notation with nameless dummies: a tool for automatic formula manipulation with application to the Church-Rosser theorem. *Indagationes Mathematicae*, 34(5):381–392, 1972. doi:10.1016/1385-7258(72)90034-0.
- 8 Mariangiola Dezani-Ciancaglini, Ross Horne, and Vladimiro Sassone. Tracing where and who provenance in linked data: A calculus. *Theor. Comput. Sci.*, 464:113–129, 2012.
- 9 Evelyn Duesterwald, Rajiv Gupta, and Mary Lou Soffa. Distributed slicing and partial re-execution for distributed programs. In *Proceedings of the 5th International Workshop on Languages and Compilers for Parallel Computing*, pages 497–511, London, UK, 1993. Springer.
- 10 D. Goswami and R. Mall. Dynamic slicing of concurrent programs. In Mateo Valero, Viktor K. Prasanna, and Sriram Vajapeyam, editors, *High Performance Computing – HiPC 2000*, volume 1970 of LNCS, pages 15–26. Springer, Berlin / Heidelberg, 2000. doi:10.1007/3-540-44467-X_2.
- 11 Daniel Hirschhoff. Handling substitutions explicitly in the pi-calculus. In *Proceedings of the Second International Workshop on Explicit Substitutions: Theory and Applications to Programs and Proofs*, 1999.
- 12 Ivan Lanese, Claudio Antares Mezzina, and Jean-Bernard Stefani. Reversing higher-order pi. In *CONCUR*, pages 478–493. Springer-Verlag, 2010.
- 13 Jean-Jacques Lévy. Optimal reductions in the lambda-calculus. In J. P. Seldin and J. R. Hindley, editors, *To H. B. Curry: Essays in Combinatory Logic, Lambda Calculus and Formalism*, pages 159–191. Academic Press, New York, NY, USA, 1980.
- 14 Robin Milner. *Communicating and mobile systems: the π calculus*. Cambridge University Press, Cambridge, UK, 1999.
- 15 D.P. Mohapatra, Rajib Mall, and Rajeev Kumar. An efficient technique for dynamic slicing of concurrent Java programs. In Suresh Manandhar, Jim Austin, Uday Desai, Yoshio Oyanagi, and AsokeK. Talukder, editors, *Applied Computing*, volume 3285 of LNCS, pages 255–262. Springer, Berlin / Heidelberg, 2004. doi:10.1007/978-3-540-30176-9_33.
- 16 Roly Perera, Umut A. Acar, James Cheney, and Paul Blain Levy. Functional programs that explain their work. In *ICFP*, pages 365–376, New York, NY, USA, 2012. ACM. doi:10.1145/2364527.2364579.
- 17 Roly Perera and James Cheney. Proof-relevant pi-calculus, 2016. Submitted. <http://arxiv.org/abs/1604.04575>.

XX:14 Causally consistent dynamic slicing

- 18 Issam Souilah, Adrian Francalanza, and Vladimiro Sassone. A formal model of provenance in distributed systems. In *TAPP 2009*, Berkeley, CA, USA, 2009. USENIX Association.
- 19 Sriraman Tallam, Chen Tian, and Rajiv Gupta. Dynamic slicing of multithreaded programs for race detection. In *24th IEEE International Conference on Software Maintenance (ICSM 2008), September 28 - October 4, 2008, Beijing, China*, pages 97–106. IEEE, 2008. doi: 10.1109/ICSM.2008.4658058.
- 20 Mark Weiser. Program slicing. In *Proceedings of the 5th International Conference on Software Engineering, ICSE '81*, pages 439–449, Piscataway, NJ, USA, 1981. IEEE Press.

A Agda module structure

Figure 6 summarises the module structure of the Agda formalisation. The module structure of the auxiliary libraries is described in [17]. All repositories can be found at the URL <https://github.com/rolyp>.

Auxiliary libraries

agda-stdlib-ext 0.0.3	Extensions to Agda library
proof-relevant-pi 0.3	Concurrent transitions, residuals and causal equivalence

Core modules

Action.Lattice	Action slices $a' \in \downarrow a$
Braiding.Proc.Lattice	Bound braid prefixes and slicing functions \mathbf{braid}_ϕ and $\mathbf{unbraid}_\phi$
ConcurrentSlicing	Include everything; compile to build project
ConcurrentSlicingCommon	Common imports from standard library
Example	Milner's scheduler example
Example.Helper	Utility functions for examples
Lattice	Lattice typeclass
Lattice.Product	Component-wise product of lattices
Name.Lattice	Name slices $y \in \downarrow x$
Proc.Lattice	Process slices $P' \in \downarrow P$
Proc.Ren.Lattice	Slicing functions $\mathbf{ren}_{\rho,P}$ and $\mathbf{unren}_{\rho,P}$
Ren.Lattice	Renaming slices $\sigma \in \downarrow \rho$ and slicing functions $\mathbf{app}_{\rho,x}$ and $\mathbf{unapp}_{\rho,x}$
Ren.Lattice.Properties	Additional properties relating to renaming slices
Transition.Lattice	Slicing functions \mathbf{step}_t and \mathbf{unstep}_{g_t}
Transition.Seq.Lattice	Slicing functions \mathbf{fwd}_i and \mathbf{bwd}_i

Common sub-modules

.GaloisConnection	Galois connection between lattices defined in parent module
-------------------	---

■ **Figure 6** concurrent-slicing module overview, release 0.1