

A Cut Principle for Information Flow[★]

Joshua D. Guttman and Paul D. Rowe

The MITRE Corporation
{guttman,prowe}@mitre.org

Abstract. We view a distributed system as a graph of active locations with unidirectional channels between them, through which they pass messages. In this context, the graph structure of a system constrains the propagation of information through it.

Suppose a set of channels is a cut set between an information source and a potential sink. We prove that, if there is no disclosure from the source to the cut set, then there can be no disclosure to the sink. We introduce a new formalization of partial disclosure, called *blur operators*, and show that the same cut property is preserved for disclosure to within a blur operator. A related compositional principle ensures limited disclosure for a class of systems that differ only beyond the cut. **February 11, 2015**

1 Introduction

In this paper, we consider information flow in a true-concurrency, distributed model. Events in an execution may be only partially ordered, and locations communicate via synchronous message-passing. Each message traverses a channel. The locations and channels form a directed graph.

Evidently, the structure of this graph will constrain the flow of information. Distant locations may have considerable information about each other's actions, but only if the information in intermediate regions accounts for this. If a kind of information does not traverse the boundary of some portion of the graph (a *cut set*), then it can never be available beyond that. We represent these limits on disclosure, i.e. kinds of information that do not escape, using *blur operators*. A blur operator is a specialized closure operator; it returns a set of behaviors local to the information source that should be indistinguishable to the observer. When disclosure from a source to a cut set is limited to within a blur operator, then disclosure to a more distant region is limited to within the same blur operator (see Thm. 26, the *cut-blur* principle).

Thus, we combine *what*-dimension declassification with a *where*-dimension perspective [47]. The blur operators formalize the semantic content of limited disclosures. The cut-blur principle gives a criterion localizing those disclosure limitations within a system architecture.

A related result, Thm. 30, supports *compositional* security. Consider any other system that differs from a given one only in its structure beyond the

[★] Copyright © 2015 The MITRE Corporation. All rights reserved.

cut. That system will preserve the flow limitations of the first, assuming that it has the same local behaviors as the first in the cut set. We illustrate this (Examples 31–32) to show that secrecy and anonymity properties of a firewall and a voting system are preserved under some environmental changes. Flow properties of a simple system remain true for more complex systems, if the latter do not distort the behavior of the simple system.

Our model is intended for many types of systems, including networks, software architectures, virtualized systems, and distributed protocols such as voting systems. Our (somewhat informal) examples are meant to be suggestive. Network examples, which involve little local state, are easy to describe, and rely heavily on the directed graph structure. Blur operators pay off by allowing many of their security goals now to be regarded as information-flow properties. Voting systems offer an interesting notion of limited disclosure, since they must disclose the result but not the choices of the individual voters. Their granularity encourages composition, since votes are aggregated from multiple precincts (Example 32).

Motivation. A treatment of information flow—relying on the graph structure of distributed systems—facilitates compositional security design and analysis.

Many systems have a natural graph structure, which is determined early in the design process. Some are distributed systems where the components are on separate platforms, and the communication patterns are a key part of the specifications. In other cases, the components may be software, such as processes or virtual machines, and the security architecture is largely concerned with their communication patterns. The designers may want to validate that these communication patterns support the information flow goals of the design early in the life cycle. Thm. 30 justifies the designers in concluding that a set of eventual systems all satisfy these security goals, when those systems all agree on “the part that really matters.”

We have also established a refinement result, which characterizes a class of system refinements that will preserve a security goal [22]. We will study it in future work.

Contributions of this paper. Our main result is the cut-blur principle, Thm. 26. The definition of *blur operator* (Def. 20) is a supplementary contribution. It identifies the logical form, so to speak, of any notion of limited disclosure supporting the proof of the cut-blur principle. Thm. 30 brings the ideas to a compositional form.

Structure of this paper. After discussing some related work (Section 2), we introduce our systems, called *frames*, and their execution model in Section 3. This is a static model, in the sense that the channels connecting different locations do not change during execution. Section 4 proves the cut-blur principle for the simple case of no disclosure of information at all across the boundary.

We introduce blur operators (Section 5) to formalize partial disclosure. In Section 6, we show that the cut idea is also applicable to blurs (Thms. 26, 30).

Section 7 provides rigorous results to relate our model to the literature. We end by indicating some future directions. Appendix A contains longer proofs, and additional lemmas.

2 Some related work

We return to related work also in Sections 7–8.

Noninterference and nondeducibility. There is a massive literature on information-flow security; Goguen and Meseguer were key early contributors [20]. Sutherland introduced the non-deducibility idea [48] as a way to formalize lack of information flow, which we have adopted in our “non-disclosure” (Def. 10). Subsequent work has explored a wide range of formalisms, including state machines [42]; process algebras such as CSP [39,38,43] and CCS [18,19,7]; and bespoke formalisms [32,26].

Irvine, Smith, and Volpano initiated a language-based approach [51], in which systems are programs. Typing ensures that their behaviors satisfy information-flow goals. Sabelfeld and Myers [45] reviewed the early years of this approach. Distributed execution has been considered in language-based work, e.g. [54,10,6]. Our work here is not specifically language-based, since the behaviors of our locations are defined only as sets of traces. If these behaviors are specified as programs, however, our results may be usefully combined with language-based methods.

Declassification. Declassification is a major concern for us. A blur operator (Def. 20) determines an upper bound on what a system may declassify. It may declassify information unless its blur operators require those details to be blurred out. Like escape-hatches [46], this is disclosure along the *what*-dimension, in the Sabelfeld-Sands classification [47]. The cut-blur principle connects this *what* declassification to *where* the processing responsible for the declassification will occur in a system architecture. In this regard, it combines a semantic view of what information is declassified with an architectural view related to intransitive noninterference [42,49].

Composability and refinement. McCullough first raised the questions of non-determinism and composability of information-flow properties [31,32]. This was a major focus of work through much of the period since, persisting until today [25,53,27,28,41,36]. Mantel, Sands, and Sudbrock [29] use a rely/guarantee method for compositional reasoning about flow in the context of imperative programs. Roscoe [40,39] offers a definition based on determinism, which is intrinsically composable. Morgan’s [35] programming language treatment clarifies the refinements that preserve security. Van der Meyden [50] provides an architectural treatment designed to achieve preservation under refinement. Our work is distinguished from these in offering a new notion of composition (see Thm. 30); in focusing on declassification; and in applying uniformly to a range of declassification policies, defined by the blur operators.

Van der Meyden’s work with Chong [11,12] is most closely related to ours. They consider “architectures,” i.e. directed graphs that express an intransitive noninterference style of *what*-dimension flow policy. The nodes of an architecture are security domains, intended to represent levels of information sensitivity. The authors define when a (monolithic) deterministic state machine, whose transitions are annotated by domains, *complies* with an architecture. The main result in [11] is a cut-like epistemic property on the architecture graph: Roughly, any knowledge acquired by a recipient about a source implies that the same knowledge is available at every cut set in the architecture graph.

A primary contrast between this paper and [11] is our distributed execution model. We consider it a sharper, more localized link to a development model, since components are likely to be designed, implemented, and upgraded piecemeal. Chong and van der Meyden focus instead on the specification level, in which security specifications constitute the directed graph. This is a new and unfamiliar artifact the developer must generate before analysis can occur. Their epistemic logic allows nested occurrences of the *knowledge* modality K_G , or occurrences of K_G in the hypothesis of an implication. Thus, it may be more expressive than our blur properties. However, this surplus extra expressiveness may not be useful. Their examples do not have nested K_G operators, or occurrences of it in the hypothesis of an implication. Indeed, our clean proof methods suggest that our model may have the right degree of generality, and be easy to understand, apply, and enrich.

Recently [12], they label the arrows by functions f , where f filters information from its source, bounding visibility to its target. They appear not to have re-established their cut-like epistemic property in the richer model, however.

Like us, Van der Meyden and Chong propose a refinement method [50,12]. Theirs relates to one possible use of homomorphisms [22], namely those in which the refined system has a homomorphism *onto* the less refined one. A homomorphism may also be an embedding of one frame *into* a richer one. In this latter case, the refinement relation appears to run in the same direction as the homomorphism. Viewed this way, our result there entails that security properties are preserved by certain refinements. In this it does not run afoul of the refinement paradox [24,35]: our theorem identifies additional assumptions that ensure the blurs are preserved.

3 Frames and Executions

We represent systems by *frames*. Each frame is a directed graph. Each node, called a *location*, is equipped with a labeled transition system that defines its possible local behaviors. The arrows are called *channels*, and allow the synchronous transmission of a message from the location at the arrow tail to the location at the arrow head. Each message also carries some *data*. For instance, we can regard a network as determining a frame in various ways:

Example 1 (Network with filtering). A corporate network uses two routers r_1, r_2 to control the connection to the public network. Router r_1 is connected




Fig. 1. A simple network in two representations

to the external internet i ; r_2 is connected to two separate internal networks n_1, n_2 . In the graph on the left in Fig. 1(a), we regard the routers and network regions as the locations. In Fig. 1(b) regions are displayed as \bullet ; routers, as \square ; and interfaces, as \triangle . When a router has an interface onto a segment, a pair of locations—representing that interface as used in each direction—lie between this router and each peer router [21].

Let $\text{Dir} = \{\text{inb}, \text{outb}\}$ represent the inbound direction and the outbound directions from routers, respectively. Suppose Rt is a set of routers r , each with a set of interfaces $\text{intf}(r)$, and a set of network regions Rg containing end hosts.

Each member of Rt, Rg is a separate location. Each interface-direction pair $(i, r) \in (\bigcup_{r \in \text{Rt}} \text{intf}(r)) \times \text{Dir}$ is also a location. Essentially, we are introducing one interface between each router/region adjacent pair, and two interfaces between each router/router adjacent pair. See the right graph in Fig. 1(b).

The channels consist of those shown. In particular, each interface has a pair of channels that allow datagrams to pass between the router and the interface, and between the interface and an adjacent entity. We also include a self-loop channel at each network region i, n_1, n_2 ; it represents transmissions and receptions among the hosts and network infrastructure coalesced into the region.

The behaviors of the locations are easily specified. The router routes, i.e. receives packets from inbound interfaces, and chooses an outbound interface for each one. The directed interfaces enforce the filtering rules, discarding packets or delivering them. If the router is executing other sorts of processing, for instance Network Address Translation or the IP Security Protocols, then the behavior is slightly more complex [2,21], but sharply localized. ///

3.1 The Static Model

In this paper, we will be concerned with a static version of the model, in which channel endpoints are never transmitted from one location to another, although

we will briefly describe a dynamic alternative in Section 8, in which channels may also be delivered over other channels.

Definition 2 (Frame). *A (static) frame \mathcal{F} has three disjoint domains:*

Locations \mathcal{LO} : *Each location $\ell \in \mathcal{LO}$ is equipped with a set of traces, $\text{traces}(\ell)$ and other information, further constrained below.*

Channels \mathcal{CH} : *Each channel $c \in \mathcal{CH}$ is equipped with two endpoints, $\text{entry}(c)$ and $\text{exit}(c)$. It is intended as a one-directional conduit of data values between the endpoints.*

Data values \mathcal{D} : *Data values $v \in \mathcal{D}$ may be delivered through channels.*

We will write \mathcal{EP} for the set of channel endpoints, which we formalize as $\mathcal{EP} = \{\text{entry}, \text{exit}\} \times \mathcal{CH}$, although we will generally continue to type $\text{entry}(c)$ and $\text{exit}(c)$ to stand for $\langle \text{entry}, c \rangle$ and $\langle \text{exit}, c \rangle$. For each $\ell \in \mathcal{LO}$, there is:

1. *a set of endpoints $\text{ends}(\ell) \subseteq \mathcal{EP}$ such that*
 - (a) *$\langle e, c \rangle \in \text{ends}(\ell)$ and $\langle e, c \rangle \in \text{ends}(\ell')$ implies $\ell = \ell'$; and*
 - (b) *there is an ℓ such that $\text{entry}(c) \in \text{ends}(\ell)$ iff there is an ℓ' such that $\text{exit}(c) \in \text{ends}(\ell')$;*

When $\text{entry}(c) \in \text{ends}(\ell)$ we will write $\text{sender}(c) = \ell$; when $\text{exit}(c) \in \text{ends}(\ell)$ we will write $\text{rcpt}(c) = \ell$. Thus, $\text{sender}(c)$ is the location that can send messages on c , while $\text{rcpt}(c)$ is the recipient location that can receive them. We write $\text{chans}(\ell)$ for $\{c : \text{sender}(c) = \ell \text{ or } \text{rcpt}(c) = \ell\}$.

We say that λ is a label for ℓ if $\lambda = (c, v)$ where $c \in \text{chans}(\ell)$ and $v \in \mathcal{D}$. We divide labels into local, transmission, and reception labels, where c, v is:

local to ℓ if $\text{sender}(c) = \ell = \text{rcpt}(c)$;
a transmission for ℓ if $\text{sender}(c) = \ell \neq \text{rcpt}(c)$;
a reception for ℓ if $\text{sender}(c) \neq \ell = \text{rcpt}(c)$.

2. *a set $\text{traces}(\ell)$, each trace being a finite or infinite sequence of labels λ . ///*

In this definition, we do not require that the local behaviors $\text{traces}(\ell)$ should be determined in any particular way. They could be specified by associating a program to each location, or a term in a process algebra, or a labeled transition system, or a mixture of these for the different locations.

Example 3 (Network with filtering). Using the notation of Example 1:

$\mathcal{LO} = \text{Rt} \cup \text{Rg} \cup ((\bigcup_{r \in \text{Rt}} \text{intf}(r)) \times \text{Dir})$;
 $\mathcal{CH} = \{(\ell_1, \ell_2) \in \mathcal{LO} \times \mathcal{LO} : \ell_1 \text{ delivers datagrams directly to } \ell_2\}$;
 $\mathcal{D} =$ the set of IP datagrams;
 $\text{ends}(\ell) = \{\text{entry}(\ell, \ell_2) : (\ell, \ell_2) \in \mathcal{CH}\} \cup \{\text{exit}(\ell_1, \ell) : (\ell_1, \ell) \in \mathcal{CH}\}$,
for each $\ell \in \mathcal{LO}$.

The traces of the locations are generated by labeled state transition systems. For a router $r \in \text{Rt}$, its state is a set of received but not yet routed datagrams, and the sole initial state is \emptyset . The transition relation, when receiving a datagram, adds it to this set. When transmitting a datagram d in the current set, it removes d from the next state and selects an outbound channel as determined by the routing table. For simplicity, we regard the routing table as an unchanging part of determining the transition relation.

For an interface/direction pair $(i, d) \in (\text{intf}(r) \times \text{Dir})$, the state again consists of a set of received but not-yet-processed datagrams. The transition relation depends on an unchanging filter function, which determines for each datagram pair whether to discard it or retransmit it.

If $n \in \text{Rg}$ is a region, its state is the set of datagrams it has received and not yet retransmitted. It can receive a datagram; transmit one from its state; or else initiate a new datagram. If assumed to be well-configured, these all have source address in a given range of IP addresses. Otherwise, they may be arbitrary. ///

Many problems can be viewed as frames. For instance, a voting scheme such as ThreeBallot [37] may be viewed as defining a family of frames; the n voters are some of the locations (see Example 24). An attestation architecture for secure virtualized systems [14] is, at one level, a set of virtual machines communicating through one-directional channels.

Each \mathcal{F} determines directed and undirected graphs, $\text{gr}(\mathcal{F})$ and $\text{ungr}(\mathcal{F})$:

Definition 4. *If \mathcal{F} is a frame, then the graph of \mathcal{F} , written $\text{gr}(\mathcal{F})$, is the directed graph (V, E) whose vertices V are the locations \mathcal{LO} , and such that there is an edge $(\ell_1, \ell_2) \in E$ iff, for some $c \in \mathcal{CH}$, $\text{sender}(c) = \ell_1$ and $\text{rcpt}(c) = \ell_2$.*

The undirected graph $\text{ungr}(\mathcal{F})$ has those vertices, and an undirected edge (ℓ_1, ℓ_2) whenever either (ℓ_1, ℓ_2) or (ℓ_2, ℓ_1) is in the edges of $\text{gr}(\mathcal{F})$. ///

3.2 Execution semantics

We give an execution model for a frame via partially ordered sets of events. The key property is that the events that involve any single location ℓ should be in $\text{traces}(\ell)$.

Definition 5 (Events; Executions). *Let \mathcal{F} be a frame, and let \mathcal{E} be a structure $\langle E, \text{chan}, \text{msg} \rangle$. The members of E are events, equipped with the functions:*

$\text{chan}: E \rightarrow \mathcal{CH}$ returns the channel of each event; and
 $\text{msg}: E \rightarrow \mathcal{D}$ returns the message communicated in each event.

1. $\mathcal{B} = (B, \preceq)$ is a system of events, written $\mathcal{B} \in \text{ES}(\mathcal{E})$, iff (i) $B \subseteq E$; (ii) \preceq is a partial ordering on B ; and (iii), for every $e_1 \in B$, $\{e_0 \in B: e_0 \preceq e_1\}$ is finite.

Hence, \mathcal{B} is well-founded. If $\mathcal{B} = (B, \preceq)$, we refer to B as $\text{ev}(\mathcal{B})$ and to \preceq as $\preceq_{\mathcal{B}}$. Now let $\mathcal{B} = (B, \preceq) \in \text{ES}(\mathcal{E})$:

2. Define $\text{proj}(B, \ell) =$

$$\{e \in B : \text{sender}(\text{chan}(e)) = \ell \text{ or } \text{rcpt}(\text{chan}(e)) = \ell\}.$$

3. \mathcal{B} is an execution, written $\mathcal{B} \in \text{Exc}(\mathcal{F})$ iff, for every $\ell \in \mathcal{LO}$,

(a) $\text{proj}(B, \ell)$ is linearly ordered by \preceq , and

(b) $\text{proj}(B, \ell)$ yields an element of $\text{traces}(\ell)$ when viewed as a sequence. ///

By Def. 5, Clauses 1 and 3a, $\text{proj}(B, \ell)$ is a finite or infinite sequence, ordered by \preceq . Thus Clause 3b is well-defined. We will often write $\mathcal{A}, \mathcal{A}'$, etc., when $\mathcal{A}, \mathcal{A}' \in \text{Exc}(\mathcal{F})$. For a frame \mathcal{F} , the choice between two different structures $\mathcal{E}_1, \mathcal{E}_2$ makes little difference: If $\mathcal{E}_1, \mathcal{E}_2$ have the same cardinality, then to within isomorphism they lead to the same systems of events and hence also executions. Thus, we will suppress the parameter \mathcal{E} , henceforth.

Definition 6. Let $\mathcal{B}_1 = (B_1, \preceq_1), \mathcal{B}_2 = (B_2, \preceq_2) \in \text{ES}(\mathcal{F})$.

1. \mathcal{B}_1 is a substructure of \mathcal{B}_2 iff $B_1 \subseteq B_2$ and $\preceq_1 = (\preceq_2 \cap B_1 \times B_1)$.
2. \mathcal{B}_1 is an initial substructure of \mathcal{B}_2 iff it is a substructure of the latter, and for all $y \in B_1$, if $x \preceq_2 y$, then $x \in B_1$. ///

Lemma 7. 1. If \mathcal{B}_1 is a substructure of $\mathcal{B}_2 \in \text{ES}(\mathcal{F})$, then $\mathcal{B}_1 \in \text{ES}(\mathcal{F})$.

2. If \mathcal{B}_1 is an initial substructure of $\mathcal{B}_2 \in \text{Exc}(\mathcal{F})$, then $\mathcal{B}_1 \in \text{Exc}(\mathcal{F})$.

3. Being an execution is preserved under chains of initial substructures: Suppose that $\langle \mathcal{B}_i \rangle_{i \in \mathbb{N}}$ is a sequence where each $\mathcal{B}_i \in \text{Exc}(\mathcal{F})$, such that $i \leq j$ implies \mathcal{B}_i is an initial substructure of \mathcal{B}_j . Then $(\bigcup_{i \in \mathbb{N}} \mathcal{B}_i) \in \text{Exc}(\mathcal{F})$. ///

Partially vs. totally ordered executions. Def. 5 does not require the ordering \preceq of “occurring before” to be total. When events occur on different channels, neither has to precede the other. Thus, our executions need not be sequential.

This has three advantages. First, it is more inclusive, since executions with total orders satisfy our definition as do those with (properly) partial orders. Indeed, the main claims of this paper remain true when restricted to executions that are totally ordered. Second, reasoning is simplified. We do not need to interleave events when combining two local executions to construct a global one, as encapsulated in the proof of Lemma 39. Nor do we need to “compact” events, when splitting off a local execution, as we would if we used a particular index set for sequences. This was probably an advantage to us in developing these results. Third, the minimal partial order is a reflection of causality, which can be used also to reason about independence. We expect this to be useful in future work.

There is also a disadvantage: unfamiliarity. It requires some caution. Moreover, mechanized theorem provers have much better support for induction over sequences than over well-founded orders. This was an inconvenience when a colleague used PVS to formalize parts of this work, eventually choosing to use totally ordered executions. With that difference, Thms. 26 and 30 have been confirmed in PVS, as have the basic properties of Example 32.

4 Non-disclosure

Following Sutherland [48], we think of information flow in terms of *deducibility* or *disclosure*. A participant observes part of the system behavior, trying to draw conclusions about a different part. If his observations exclude some possible behaviors of that part, then he can *deduce* that those behaviors did not occur. His observations have *disclosed* something.

These observations occur on a set of channels $C_o \subseteq \mathcal{CH}$, and the deductions constrain the events on a set of channels $C_s \subseteq \mathcal{CH}$. C_o is the set of *observed* channels, and C_s is the set of *source* channels. The observer has access to the events on the channels in C_o in an execution, using these events to learn about what happened at the source. The observed events may rule out some behaviors on the channels C_s .

Definition 8. Let $C \subseteq \mathcal{CH}$, and $\mathcal{B} \in \text{ES}(\mathcal{F})$.

1. The restriction $\mathcal{B} \upharpoonright C$ of \mathcal{B} to C is (B_0, R) , where
 $B_0 = \{e \in \mathcal{B} : \text{chan}(e) \in C\}$, and
 $R = (\preceq \cap B_0 \times B_0)$.
2. $\mathcal{B} \in \text{ES}(\mathcal{F})$ is a C -run iff for some $\mathcal{A} \in \text{Exc}(\mathcal{F})$, $\mathcal{B} = \mathcal{A} \upharpoonright C$. We write $C\text{-runs}(\mathcal{F})$, or sometimes $C\text{-runs}$, for the set of C -runs of \mathcal{F} . A local run is a member of $C\text{-runs}$ for the relevant C .
3. $J_{C' \triangleleft C}(\mathcal{B})$ gives the C' -runs compatible with a C -run \mathcal{B} :

$$J_{C' \triangleleft C}(\mathcal{B}) = \{\mathcal{A} \upharpoonright C' : \mathcal{A} \in \text{Exc}(\mathcal{F}) \text{ and } \mathcal{A} \upharpoonright C = \mathcal{B}\}. \quad ///$$

$\mathcal{B} \upharpoonright C \in \text{ES}(\mathcal{F})$ by Lemma 7. In $J_{C' \triangleleft C}(\mathcal{B})$, the lower right index C indicates what type of local run \mathcal{B} is. The lower left index C' indicates the type of local runs in the resulting set. J stands for “joint.” $J_{C' \triangleleft C}(\mathcal{B})$ makes sense even if C and C' overlap, though behavior on $C \cap C'$ is not hidden from observations at C .

Lemma 9. 1. $\mathcal{B} \notin C\text{-runs}$ implies $J_{C' \triangleleft C}(\mathcal{B}) = \emptyset$.

2. $\mathcal{B} \in C\text{-runs}$ implies $J_{C \triangleleft C}(\mathcal{B}) = \{\mathcal{B}\}$.

3. $J_{C' \triangleleft C}(\mathcal{B}) \subseteq C'\text{-runs}$. ///

No disclosure means that any observation \mathcal{B} at C is compatible with everything that could have occurred at C' , where *compatible* means that there is some execution that combines the local C -run with the desired C' -run.

We summarize “no disclosure” by the Leibnizian slogan: *Everything possible is compossible*, “compossible” being his coinage meaning possible together. If $\mathcal{B}, \mathcal{B}'$ are each separately possible—being C, C' -runs respectively—then there’s an execution \mathcal{A} combining them, and restricting to each of them.

Definition 10. \mathcal{F} has no disclosure from C to C' iff, for all C -runs \mathcal{B} , $J_{C' \triangleleft C}(\mathcal{B}) = C'\text{-runs}$. ///

4.1 Symmetry of disclosure

Like Shannon's mutual information and Sutherland's non-deducibility [48], “no disclosure” is symmetric:

Lemma 11. 1. $\mathcal{B}' \in J_{C' \triangleleft C}(\mathcal{B})$ iff $\mathcal{B} \in J_{C \triangleleft C'}(\mathcal{B}')$.
 2. \mathcal{F} has no disclosure from C to C' iff \mathcal{F} has no disclosure from C' to C .

Proof. 1. By the definition, $\mathcal{B}' \in J_{C' \triangleleft C}(\mathcal{B})$ iff there exists an execution \mathcal{B}_1 such that $\mathcal{B}_1 \upharpoonright C = \mathcal{B}$ and $\mathcal{B}_1 \upharpoonright C' = \mathcal{B}'$. Which is equivalent to $\mathcal{B} \in J_{C \triangleleft C'}(\mathcal{B}')$.

2. There is no disclosure from C' to C if for every C -run \mathcal{B} , for every C' -run \mathcal{B}' , $\mathcal{B}' \in J_{C' \triangleleft C}(\mathcal{B})$. By Clause 1, this is the same as $\mathcal{B} \in J_{C \triangleleft C'}(\mathcal{B}')$. \square

Because of this symmetry, we speak of no disclosure *between* C and C' .

\mathcal{A} witnesses for $\mathcal{B}' \in J_{C' \triangleleft C}(\mathcal{B})$ iff $\mathcal{A} \in \text{Exc}(\mathcal{F})$, $\mathcal{B} = \mathcal{A} \upharpoonright C$, and $\mathcal{B}' = \mathcal{A} \upharpoonright C'$. We also sometimes say that \mathcal{B}_1 witnesses for $\mathcal{B}' \in J_{C' \triangleleft C}(\mathcal{B})$ if $\mathcal{B}_1 \in C_1$ -runs, where $C \cup C' \subseteq C_1$, when $\mathcal{B}_1 \upharpoonright C = \mathcal{B}$ and $\mathcal{B}_1 \upharpoonright C' = \mathcal{B}'$.

Lemma 12. 1. Suppose $C_0 \subseteq C_1$ and $C'_0 \subseteq C'_1$. If \mathcal{F} has no disclosure from C_1 to C'_1 , then \mathcal{F} has no disclosure from C_0 to C'_0 .
 2. When $C_1, C_2, C_3 \subseteq \mathcal{CH}$,

$$J_{C_3 \triangleleft C_1}(\mathcal{B}_1) \subseteq \bigcup_{\mathcal{B}_2 \in J_{C_2 \triangleleft C_1}(\mathcal{B}_1)} J_{C_3 \triangleleft C_2}(\mathcal{B}_2). \quad ///$$

Lemma 14 gives a condition under which this inclusion is an equality. See Appendix A for this proof, and the longer subsequent proofs.

4.2 The Cut Principle for Non-disclosure

Our key observation is that non-disclosure respects the graph structure of a frame \mathcal{F} . If $\text{cut} \subseteq \mathcal{CH}$ is a cut set in the undirected graph $\text{ungr}(\mathcal{F})$, then disclosure from a source set $\text{src} \subseteq \mathcal{CH}$ to a sink $\text{obs} \subseteq \mathcal{CH}$ is controlled by disclosure to cut . If there is no disclosure from src to cut , there can be no disclosure from src to obs . As we will see in Section 6, this property extends to limited disclosure in the sense of disclosure to within a blur operator.

We view a cut as separating one set of channels as source from another set of channels as sink. Although it is more usual to take a cut to separate sets of nodes than sets of channels, it is easy to transfer between the channels and the relevant nodes. If $C \subseteq \mathcal{CH}$, we let $\text{ends}(C) = \{\ell: \exists c \in C. \text{sender}(c) = \ell \text{ or } \text{rcpt}(c) = \ell\}$; conversely, $\text{chans}(L) = \{c: \text{sender}(c) \in L \text{ or } \text{rcpt}(c) \in L\}$. For a singleton set $\{\ell\}$ we suppress the curly braces and write $\text{chans}(\ell)$.

Definition 13. Let $\text{src}, \text{cut}, \text{obs} \subseteq \mathcal{CH}$ be sets of channels; cut is an undirected cut (or simply a cut) between src, obs iff

1. $\text{src}, \text{cut}, \text{obs}$ are pairwise disjoint; and
2. every undirected path p_1 in $\text{ungr}(\mathcal{F})$ from any $\ell_1 \in \text{ends}(\text{obs})$ to any $\ell_2 \in \text{ends}(\text{src})$ traverses some member of cut . ///

For instance, in Fig. 1(b), $\{c_1, c_2\}$ is a cut between $\text{chans}(i)$ and $\text{chans}(\{n_1, n_2\})$. Lemma 14 serves as the heart of the proofs of the two main theorems about cuts, Thms. 15 and 26.

Lemma 14. *Let cut be an undirected cut between src, obs, and let $\mathcal{B}_o \in \text{src-runs}$. Then*

$$J_{\text{src} \triangleleft \text{obs}}(\mathcal{B}_o) = \bigcup_{\mathcal{B}_c \in J_{\text{cut} \triangleleft \text{obs}}(\mathcal{B}_o)} J_{\text{src} \triangleleft \text{cut}}(\mathcal{B}_c). \quad ///$$

Lemma 14 is in fact a corollary of Lemma 29, which makes an analogous assertion about a pair of overlapping frames.

Theorem 15. *Let cut be an undirected cut between src, obs in \mathcal{F} . If there is no disclosure between src and cut, then there is no disclosure between src and obs.*

Proof. Suppose that $\mathcal{B}_s \in \text{src-runs}$ and $\mathcal{B}_o \in \text{obs-runs}$. We must show $\mathcal{B}_s \in J_{\text{src} \triangleleft \text{obs}}(\mathcal{B}_o)$. To apply Lemma 14, let $\mathcal{A} \in \text{Exc}(\mathcal{F})$ such that $\mathcal{B}_o = \mathcal{A} \upharpoonright \text{obs}$; \mathcal{A} exists by the definition of obs-run. Letting $\mathcal{B}_c = \mathcal{A} \upharpoonright \text{cut}$, we have $\mathcal{B}_c \in J_{\text{cut} \triangleleft \text{obs}}(\mathcal{B}_o)$.

Since there is no disclosure between cut and src, $\mathcal{B}_s \in J_{\text{src} \triangleleft \text{cut}}(\mathcal{B}_c)$, and Lemma 14 applies. \square

Example 16. In Fig. 1(b) let r_1 be configured to discard all inbound packets, and r_2 to discard all outbound packets. Then the empty event system is the only member of $\{c_1, c_2\}$ -runs. Hence there is no disclosure between $\text{chans}(i)$ and $\{c_1, c_2\}$. By Thm. 15, there is no disclosure to $\text{chans}(\{n_1, n_2\})$. $///$

So disconnected portions of a frame cannot interfere:

Corollary 17. *If there is no path between src and obs in $\text{ungr}(\mathcal{F})$, then there is no disclosure between them.*

Proof. Then $\text{cut} = \emptyset$ is an undirected cut set, and there is only one cut-run, namely the empty system of events. It is thus compatible with all src-runs. \square

Thm. 15 and its analogue Thm. 26, while reminiscent of the max flow/min cut principle (cf. e.g. [15, Sec. 26.2]), are however quite distinct from it, as the latter depends essentially on the quantitative structure of network flows. Our results may also seem reminiscent of the Data Processing Inequality, stating that when three random variables X, Y, Z form a Markov chain, the mutual information $I(X; Z) \leq I(X; Y)$. Indeed, Thm. 15 entails the special case where $I(X; Y) = 0$, choosing $\text{gr}(\mathcal{F})$ to be a single path $X \rightarrow Y \rightarrow Z$. For more on quantitative information flow and our work, see the conclusion (Sec. 8).

5 Blur operators

Turning from no disclosure to partial disclosure, for every frame and regions of interest $\text{src} \subseteq \mathcal{CH}$ and $\text{obs} \subseteq \mathcal{CH}$, there is an equivalence relation on src-runs:

Definition 18. Let $\text{src}, \text{obs} \subseteq \mathcal{CH}$. If $\mathcal{B}_1, \mathcal{B}_2$ are src -runs, we say that they are obs -equivalent, and write $\mathcal{B}_1 \approx_{\text{obs}} \mathcal{B}_2$, iff, for all obs -runs \mathcal{B}_o ,

$$\mathcal{B}_1 \in J_{\text{src} \triangleleft \text{obs}}(\mathcal{B}_o) \text{ iff } \mathcal{B}_2 \in J_{\text{src} \triangleleft \text{obs}}(\mathcal{B}_o) \quad ///$$

Lemma 19. For each obs , \approx_{obs} is an equivalence relation. ///

No disclosure means that all src -runs are obs -equivalent. Any notion of partial disclosure must respect obs -equivalence, since no observations can possibly “split” apart obs -equivalent src -runs. Partial disclosures always respect unions of obs -equivalence classes. Hence, every notion of partial equivalence has three properties. We will group them together in the following notion of a *blur operator*, because they are the technical conditions that matter in our proofs:

Definition 20. A function f on sets is a blur operator iff it satisfies:

Inclusion: For all sets S , $S \subseteq f(S)$;

Idempotence: f is idempotent, i.e. for all sets S , $f(f(S)) = f(S)$; and

Union: f commutes with unions: If $S_{a \in I}$ is a family indexed by the set I , then

$$f\left(\bigcup_{a \in I} S_a\right) = \bigcup_{a \in I} f(S_a). \quad (1)$$

S is an f -blurred set iff f is a blur operator and $S = f(S)$. ///

By Idempotence, S is f -blurred iff it is in the range of the blur operator f . Since $S = \bigcup_{a \in S} \{a\}$, the Union property says that f is determined by its action on the singleton subsets of S . Thus, Inclusion could have said $a \in f(\{a\})$.

Monotonicity also follows from the Union property; if $S_1 \subseteq S_2$, then $S_2 = S_0 \cup S_1$, where $S_0 = S_2 \setminus S_1$. Thus $f(S_2) = f(S_0) \cup f(S_1)$, so $f(S_1) \subseteq f(S_2)$.

For intuition about blurs, think of blurring an image: The viewer no longer knows the details of the scene. The viewer knows only that it was some scene which, when blurred, would look like this. In this interpretation of the knowledge available to the observer, we follow the tradition of epistemic logic [17]; see also [5].

Example 21. Imaginary Weather Forecasting Inc. (IWF) sells tailored, high-resolution weather data and forecasts to airlines, airports, and commodity investment firms, and it sells low-resolution weather data more cheaply to TV and radio stations. IWF’s low-tier subscribers should not learn higher resolution data than they have paid for. There is some disclosure about high resolution data because (e.g.) when low-tier subscribers see warm temperatures, they know that the high-resolution data is inconsistent with snow. We can formalize this partial disclosure as a blur.

Suppose IWF creates its low-resolution data d by applying a lossy compression function comp to high-resolution data D . Then when low-tier subscribers receive d , they know that the high-resolution data IWF measured from the environment is some element of $\text{comp}^{-1}(d) = \{D : \text{comp}(D) = d\}$. These sets are f -blurred where $f(\{D\}) = \{D' : \text{comp}(D') = \text{comp}(D)\}$. ///

Lemma 22. Suppose that A is a set, and \mathcal{R} is a partition of the elements of A . There is a unique function $f_{\mathcal{R}}$ on sets $S \subseteq A$ such that

1. $f_{\mathcal{R}}(\{a\}) = S$ iff $a \in S$ and $S \in \mathcal{R}$;
2. $f_{\mathcal{R}}$ commutes with unions (Eqn. 1).

Moreover, $f_{\mathcal{R}}$ is a blur operator.

Proof. Since $S = \bigcup_{a \in S} \{a\}$, $f_{\mathcal{R}}(S)$ is defined by the union principle (Eqn. 1).

Inclusion and *Union* are immediate from the form of the definition. *Idempotence* holds because being in the same \mathcal{R} -equivalence class is transitive. \square

Although every equivalence relation determines a blur operator, the converse is not true: Not every blur operator is of this form. For instance, consider the set $A = \{a, b\}$, and let $f(\{a\}) = \{a\}$, $f(\{b\}) = f(\{a, b\}) = \{a, b\}$.

We will study information disclosure to within blur operators f , which we interpret as meaning that $J_{C' \triangleleft C}(\mathcal{B}_c)$ is f -blurred. Essentially, this is an “upper bound” on how much information may be disclosed when \mathcal{B}_c is observed. The observer will know an f -blurred set that the behavior at C' belongs to. However, the observer cannot infer anything finer than the f -blurred sets.

Definition 23. Let $\text{obs}, \text{src} \subseteq \mathcal{CH}$, and let f be a function on sets of src -runs.

\mathcal{F} restricts disclosure from src to obs to within f iff f is a blur operator and, for every obs -run \mathcal{B}_o , $J_{\text{src} \triangleleft \text{obs}}(\mathcal{B}_o)$ is an f -blurred set.

We also say that \mathcal{F} f -limits src -to- obs flow. ///

At one extreme, the no-disclosure condition is also disclosure to within a blur operator, namely the blur operator that ignores S and adds all C' -runs:

$$f_{\text{all}}(S) = \{\mathcal{A} \upharpoonright C' : \mathcal{A} \in \text{Exc}(\mathcal{F})\}.$$

At the other extreme, the maximally permissive security policy is represented as disclosure to within the identity function $f_{\text{id}}(S) = S$. The blur operator f_{id} also shows that every frame restricts disclosure to within *some* blur operator. The equivalence classes generated by compatibility are always f_{id} -blurred.

\mathcal{F} may f -limit src -to- obs flow even when the intersection $\text{obs} \cap \text{src}$ is non-empty, as long as f is not too fine-grained; see below e.g. in Def. 36.

Example 24. Suppose that \mathcal{F} is an electronic voting system such as ThreeBallot or its siblings [37]. Some locations L_{EC} are run by the election commission. We will regard the voters themselves as a set of locations L_V . Each voter delivers a message containing, in some form, his vote for some candidate.

The election officials observe the channels connected to L_{EC} , i.e. $\text{chans}(L_{EC})$. To determine the correct outcome, they must infer a property of the local run at $\text{chans}(L_V)$, namely, how many votes for each candidate occurred. However, they should not find out which voter voted for which candidate [16].

We formalize this via a blur operator. Suppose $\mathcal{B}' \in \text{chans}(L_V)$ -runs is a possible behavior of all voters in L_V . Suppose that π is a permutation of L_V . Let $\pi \cdot \mathcal{B}'$ be the behavior in which each voter $\ell \in L_V$ casts not his own actual vote,

but the vote actually cast by $\pi(\ell)$. That is, π represents one way of reallocating the actual votes among different voters. Now for any $S \subseteq \text{chans}(L_V)\text{-runs}$ let

$$f_0(S) = \{\pi \cdot \mathcal{B}' : \mathcal{B}' \in S \wedge \pi \text{ is a permutation of } L_V\}. \quad (2)$$

This is a blur operator: (i) the identity is a permutation; (ii) permutations are closed under composition; and (iii) Eqn. 2 implies commutation with unions. The election commission should learn nothing about the votes of individuals, meaning that, for any $\mathcal{B} \in \text{chans}(L_{EC})\text{-runs}$ the commission could observe, $J_{\text{chans}(L_V) \triangleleft \text{chans}(L_{EC})}(\mathcal{B})$ is f_0 -blurred. Permutations of compatible voting patterns are also compatible.

This example is easily adapted to other considerations. For instance, the commissioners of elections are also voters, and they know how they voted themselves. Thus, we could define a (narrower) blur operator f_1 that only uses the permutations that leave commissioners' votes fixed.

In fact, voters are divided among different precincts in many cases, and tallies are reported on a per-precinct basis. Thus, we have sets V_1, \dots, V_k of voters registered at the precincts P_1, \dots, P_k respectively. The relevant blur function says that we can permute the votes of any two voters $v_1, v_2 \in V_i$ within the same precinct. One cannot permute votes between different precincts, since that could change the tallies in the individual precincts. ///

Example 25. Suppose in Fig. 1(b): The inbound interface from i to router r_1 discards downward-flowing packets unless their source is an address in i and the destination is an address in n_1, n_2 . The inbound interface for downward-flowing to router r_2 discards packets unless the destination address is the IP for a web server **www** in n_1 , and the destination port is 80 or 443, or else their source port is 80 or 443 and their destination port is ≥ 1024 .

We filter outbound (upward-flowing) packets symmetrically.

A packet is *importable* iff its source address is in i and either its destination is **www** and its destination port is 80 or 443; or else its destination address is in n_1, n_2 , its source port is 80 or 443, and its destination port is ≥ 1024 .

It is *exportable* iff, symmetrically, its destination address is in i and either its source is **www** and its source port is 80 or 443; or else its source address is in n_1, n_2 , its destination port is 80 or 443, and its source port is ≥ 1024 .

We will write $\text{select } \mathcal{B} p$ for the result of selecting those events $e \in \text{ev}(\mathcal{B})$ that satisfy the predicate $p(e)$, restricting \preceq to the selected events. Now consider the operator f_i on $\text{chans}(i)\text{-runs}$ generated as in Lemma 22 from the equivalence relation:

$\mathcal{B}_1 \approx_i \mathcal{B}_2$ iff they agree on all *importable* events, i.e.:

$$\begin{aligned} \text{select } \mathcal{B}_1 (\lambda e. \text{msg}(e) \text{ is importable}) &\cong \\ \text{select } \mathcal{B}_2 (\lambda e. \text{msg}(e) \text{ is importable}). \end{aligned}$$

The router configurations mentioned above are intended to ensure that there is f_i -limited flow from $\text{chans}(i)$ to $\text{chans}(\{n_1, n_2\})$.

This is an *integrity* condition; it is meant to ensure that systems in n_1, n_2 cannot be affected by bad (i.e. non-importable) packets from i .

Outbound, the blur f_e on $\text{chans}(\{n_1, n_2\})$ -runs is generated from the equivalence relation:

$\mathcal{B}_1 \approx_e \mathcal{B}_2$ iff they agree on all *exportable* events, i.e.:

$$\begin{aligned} \text{select } \mathcal{B}_1 (\lambda e . \text{msg}(e) \text{ is exportable}) &\cong \\ \text{select } \mathcal{B}_2 (\lambda e . \text{msg}(e) \text{ is exportable}). \end{aligned}$$

The router configurations are also intended to ensure that there is f_e -limited flow from $\text{chans}(\{n_1, n_2\})$ to $\text{chans}(i)$.

This is a *confidentiality* condition; it is meant to ensure that external observers learn nothing about the non-exportable traffic, which was not intended to exit the organization.

In this example, it is helpful that the exportable and importable packets are disjoint, and the transmission of one of these packets is never dependent on the reception of a packet of the other kind. In applications lacking this property, proving flow limitations is harder. ///

6 The Cut-Blur Principle

The symmetry of non-disclosure (Lemma 11) no longer holds for disclosure to within a blur. We have, however, the natural extension of Thm 15:

Theorem 26 (Cut-Blur Principle). *Let cut be an undirected cut between src, obs in \mathcal{F} . If \mathcal{F} f -limits src-to-cut flow, then \mathcal{F} f -limits src-to-obs flow.*

Proof. By the hypothesis, f is a blur operator. Let \mathcal{B}_o be a obs-run. We want to show that $J_{\text{src} \triangleleft \text{obs}}(\mathcal{B}_o)$ is an f -blurred set, i.e. $J_{\text{src} \triangleleft \text{obs}}(\mathcal{B}_o) = f(J_{\text{src} \triangleleft \text{obs}}(\mathcal{B}_o))$.

For convenience, let $S_c = J_{\text{cut} \triangleleft \text{obs}}(\mathcal{B}_o)$.

By Lemma 14, $J_{\text{src} \triangleleft \text{obs}}(\mathcal{B}_o) = \bigcup_{\mathcal{B}_c \in S_c} J_{\text{src} \triangleleft \text{cut}}(\mathcal{B}_c)$. Thus, we must show that the latter is f -blurred.

By the assumption that each $J_{\text{src} \triangleleft \text{cut}}(\mathcal{B}_c)$ is f -blurred and by idempotence, $J_{\text{src} \triangleleft \text{cut}}(\mathcal{B}_c) = f(J_{\text{src} \triangleleft \text{cut}}(\mathcal{B}_c))$. Now:

$$\begin{aligned} \bigcup_{\mathcal{B}_c \in S_c} J_{\text{src} \triangleleft \text{cut}}(\mathcal{B}_c) &= \bigcup_{\mathcal{B}_c \in S_c} f(J_{\text{src} \triangleleft \text{cut}}(\mathcal{B}_c)) \\ &= f\left(\bigcup_{\mathcal{B}_c \in S_c} J_{\text{src} \triangleleft \text{cut}}(\mathcal{B}_c)\right), \end{aligned}$$

applying the union property (Eqn. 1). Hence, $\bigcup_{\mathcal{B}_c \in S_c} J_{\text{src} \triangleleft \text{cut}}(\mathcal{B}_c)$ is f -blurred. \square

This proof is the reason we introduced the *Union* principle, rather than simply considering all closure operators [33]. Eqn. 1 distinguishes those closure operators that allow the “long distance reasoning” summarized in the proof.

Example 27. The frame of Example 25 has f_i -limited flow from $\text{chans}(i)$ to the cut $\{c_1, c_2\}$. Thus, it has f_i -limited flow from $\text{chans}(i)$ to $\text{chans}(\{n_1, n_2\})$.

It also has f_e -limited flow from $\text{chans}(\{n_1, n_2\})$ to the cut $\{c_1, c_2\}$. This implies f_e -limited flow to $\text{chans}(i)$. ///

6.1 A Compositional Relation between Frames

Our next technical result gives us a way to “transport” a blur security property from one frame \mathcal{F}_1 to another frame \mathcal{F}_2 . It assumes that the two frames share a common core, some set of locations L_0 . These locations should hold the same channel endpoints in each of $\mathcal{F}_1, \mathcal{F}_2$, and should engage in the same traces. The boundary separating L_0 from the remainder of $\mathcal{F}_1, \mathcal{F}_2$ necessarily forms a cut set cut. Assuming that the local runs at cut are respected, blur properties are preserved from \mathcal{F}_1 to \mathcal{F}_2 .

Definition 28. A set L_0 of locations is shared between \mathcal{F}_1 and \mathcal{F}_2 iff $\mathcal{F}_1, \mathcal{F}_2$ are frames with locations $\mathcal{LO}_1, \mathcal{LO}_2$, endpoints $\text{ends}_1, \text{ends}_2$ and traces $\text{traces}_1, \text{traces}_2$, resp., where $L_0 \subseteq \mathcal{LO}_1 \cap \mathcal{LO}_2$, and for all $\ell \in L_0$, $\text{ends}_1(\ell) = \text{ends}_2(\ell)$ and $\text{traces}_1(\ell) = \text{traces}_2(\ell)$.

When L_0 is shared between \mathcal{F}_1 and \mathcal{F}_2 , let:

$\text{left}_0 = \{c \in \mathcal{CH}_1 : \text{ both endpoints of } c \text{ are locations } \ell \in L_0\};$
 $\text{cut}_0 = \{c \in \mathcal{CH}_1 : \text{ exactly one endpoint of } c \text{ is a location } \ell \in L_0\}; \text{ and}$
 $\text{right}_i = \{c \in \mathcal{CH}_i : \text{ neither endpoint of } c \text{ is a location } \ell \in L_0\}, \text{ for } i = 1, 2.$

We will also use $C\text{-runs}_1$ and $C\text{-runs}_2$ to refer to the local runs of C within \mathcal{F}_1 and \mathcal{F}_2 , resp.; and $J_{C' \triangleleft C}^1(\mathcal{B})$ and $J_{C' \triangleleft C}^2(\mathcal{B})$ will refer to the compatible C' runs in the frames \mathcal{F}_1 and \mathcal{F}_2 , resp. ///

Indeed, cut_0 is an undirected cut between left_0 and right_i in \mathcal{F}_i , for $i = 1$ and 2 . In an undirected path that starts in left_0 and never traverses cut_0 , each arc always has both ends in L_0 . We next prove a two-frame analog of Lemma 14.

Lemma 29. Let L_0 be shared between frames $\mathcal{F}_1, \mathcal{F}_2$. Let $\text{src} \subseteq \text{left}_0$, and $\mathcal{B}_c \in \text{cut}_0\text{-runs}_1 \cap \text{cut}_0\text{-runs}_2$.

1. $J_{\text{src} \triangleleft \text{cut}_0}^1(\mathcal{B}_c) = J_{\text{src} \triangleleft \text{cut}_0}^2(\mathcal{B}_c).$
2. Assume $\text{cut}_0\text{-runs}(\mathcal{F}_2) \subseteq \text{cut}_0\text{-runs}(\mathcal{F}_1)$. Let $\text{obs} \subseteq \text{right}_2$, and $\mathcal{B}_o \in \text{obs-runs}_2$.
Then

$$J_{\text{src} \triangleleft \text{obs}}^2(\mathcal{B}_o) = \bigcup_{\mathcal{B}_c \in J_{\text{cut}_0 \triangleleft \text{obs}}^2(\mathcal{B}_o)} J_{\text{src} \triangleleft \text{cut}_0}^1(\mathcal{B}_c).$$

Part 1 states that causality acts locally. The variable portions $\text{right}_1, \text{right}_2$ of \mathcal{F}_1 and \mathcal{F}_2 can affect what happens in their shared part left. But it does so only by changing which $\text{cut}_0\text{-runs}$ are possible. Whenever both frames agree on any $\mathcal{B}_c \in \text{cut}_0\text{-runs}_1 \cap \text{cut}_0\text{-runs}_2$, then the left-runs runs compatible with \mathcal{B}_c are the same. The variation between $\text{right}_1, \text{right}_2$ affects the distant behavior within left only by changing the possible local runs at the boundary cut_0 .

The assumption $\text{cut}_0\text{-runs}(\mathcal{F}_2) \subseteq \text{cut}_0\text{-runs}(\mathcal{F}_1)$ in Part 2 and Thm. 30 is meant to limit this variability in one direction.

Theorem 30. Suppose that L_0 is shared between frames $\mathcal{F}_1, \mathcal{F}_2$, and assume $\text{cut}_0\text{-runs}(\mathcal{F}_2) \subseteq \text{cut}_0\text{-runs}(\mathcal{F}_1)$. Consider any $\text{src} \subseteq \text{left}_0$ and $\text{obs} \subseteq \text{right}_2$. If \mathcal{F}_1 f -limits src-to-cut_0 flow, then \mathcal{F}_2 f -limits src-to-obs flow.

The proof is similar to the proof of the cut-blur principle, which effectively results from it by replacing Lemma 29 by Lemma 14, and omitting the subscripts on frames and their local runs. The cut-blur principle is in fact the corollary of Thm. 30 for $\mathcal{F}_1 = \mathcal{F}_2$.

6.2 Two Applications

Thm. 30 is a useful compositional principle. It implies, for instance in connection with Example 27, that non-exportable traffic in n_1, n_2 remains unobservable even as we vary the top part of Fig. 1(b).

Example 31. Regarding Fig. 1(b) as the frame \mathcal{F}_1 , let L_0 be the locations below $\{c_1, c_2\}$, and let $\text{cut} = \{c_1, c_2\}$. Let \mathcal{F}_2 contain L_0, cut as shown, and have any graph structure above cut such that cut remains a cut between the new structure and \mathcal{F}_0 . Let the new locations have any transition systems such that the local runs agree, i.e. $\text{cut-runs}(\mathcal{F}_2) = \text{cut-runs}(\mathcal{F}_1)$. Then by Thm. 30, external observations of $\text{chans}(\{n_1, n_2\})$ are guaranteed to blur out non-exportable events.

///

It is appealing that our security goal is independent of changes in the structure of the internet that we do not control. A similar property holds for the integrity goal of Example 27 as we alter the internal network. The converse questions—preserving the confidentiality property as the internal network changes, and the integrity property as the internet changes—appear to require a different theorem. Thm. 42 in Appendix B is intended to handle these questions.

Example 32. Consider a frame \mathcal{F}_1 representing a precinct, as shown in Fig. 2. It consists of a set of voters $\bar{v} = \{v_1, \dots, v_k\}$, a ballot box BB_1 , and a channel c_1 connecting that to the election commission EC . The EC publishes the results over the channel p to the public Pub .

We have proved that a particular implementation of BB_1 ensures that \mathcal{F}_1 blurs the votes; we formalized this within the theorem prover PVS. That is, if a pattern of voting in precinct i is compatible with an observation at c_1 , then any permutation of the votes at \bar{v} is also compatible.

The cut-blur principle implies this blur also applies to observations at channel p to the public. Other implementations of BB_1 also achieve this property. Three-Ballot and VAV [37] appear to have this effect; they involve some additional data delivered to Pub , namely the receipts for the ballots.¹

However, elections generally concern many precincts. Frame \mathcal{F}_2 contains i precincts, all connected to the election commission EC (Fig. 3). Taking $L_0 = \bar{v} \cup \{BB_1\}$, we may apply Thm. 30. We now have $\text{cut} = \{c_1\}$. Thus, to infer that \mathcal{F}_2 blurs observations of the voters in precinct 1, we need only check that $\{c_1\}$ has no new local runs in \mathcal{F}_2 .

By symmetry, each precinct in \mathcal{F}_2 enjoys the same blur.

¹ Our claim is non-probabilistic. For quantitative conclusions, this no longer holds: Some permutations are more likely than others, given the receipts [13,34].




Fig. 2. A single precinct: \mathcal{F}_1




Fig. 3. Multiple precincts report to EC : \mathcal{F}_2

Thus—for a given local run at p —any permutation of the votes at \bar{v} preserves compatibility in \mathcal{F}_2 , and any permutation of the votes at \bar{w} preserves compatibility in \mathcal{F}_2 . However, Thm. 30 does not say that any pair of permutations at \bar{v} and \bar{w} must be jointly compatible. That is, does every permutation on $\bar{v} \cup \bar{w}$ that respects the division between the precinct of the \bar{v} s and the precinct of the \bar{w} s preserve compatibility? Although Thm. 30 does not answer this question, the answer is yes, as we can see by applying Lemma 39 to \mathcal{F}_2 . ///

Thm. 30 is a tool to justify abstractions. Fig. 1(b) is a sound abstraction of a variety of networks, and Fig. 2 is a sound abstraction of the various multiple precinct instances of Fig. 3.

7 Relating Blurs to Noninterference and Nondeducibility

If we specialize frames to state machines (see Fig. 4), we can reproduce some of the traditional definitions. Let $D = \{d_1, \dots, d_k\}$ be a finite set of *domains*, i.e. sensitivity labels; $\hookrightarrow \subseteq D \times D$ specifies which domains *may influence* each other. We assume \hookrightarrow is reflexive, though not necessarily transitive. A is a set of *actions*, and $\text{dom}: A \rightarrow D$ assigns a domain to each action; O is a set of outputs.

$M = \langle S, s_0, A, \delta, \text{obs} \rangle$ is a (possibly non-deterministic) state machine with states S , initial state s_0 , transition relation $\delta \subseteq S \times A \times S$, and observation function $\text{obs}: S \times D \rightarrow O$. M has a set of traces α , and each trace α determines a sequence of observations for each domain [42,49,50].




Fig. 4. Machine M , domains $\{d_1, \dots, d_k\}$

M accepts commands from A along the incoming channels c_i^{in} from the d_i ; each command $a \in A$ received from d_i has sensitivity $\text{dom}(a) = d_i$. M delivers observations along the outgoing channels c_i^{out} . The frame requires a little extra memory, in addition to the states of M , to deliver outputs over the channels c_i^{out} .

\mathcal{F} being star-like, each $\mathcal{A} \in \text{Exc}(\mathcal{F})$ has a linearly ordered $\preceq_{\mathcal{A}}$, since all events are in the linearly ordered $\text{proj}(\mathcal{A}, M)$. Let us write:

$$\begin{aligned}
 C_i &= \{c_i^{\text{in}}, c_i^{\text{out}}\} \text{ for } d_i\text{'s input and output for } M; \\
 \text{vis}(d_i) &= \{c_j^{\text{in}} : d_j \hookrightarrow d_i\} \text{ for inputs visible to } d_i; \\
 \text{IN} &= \{c_x^{\text{in}} : 1 \leq x \leq k\} \text{ for the input channels;} \\
 \text{input}(\mathcal{A}) &= \mathcal{A} \upharpoonright \text{IN} \text{ for all input behavior in } \mathcal{A}.
 \end{aligned}$$

Noninterference and nondeducibility. *Noninterference* [20] and its variants are defined by *purge* functions p for each target domain d_i , defined by recursion on input behaviors $\text{input}(\mathcal{A})$. The original Goguen-Meseguer (GM) purge function p_o for d_i [20] retains the events $e \in \text{input}(\mathcal{A})$ satisfying the predicate

$$\text{chan}(e) \in \text{vis}(d_i).$$

A purge function for intransitive \hookrightarrow relations was subsequently proposed by Haigh and Young [23]. In the purge function for domain d_i , any input event $e_0 \in \text{input}(\mathcal{A})$ is retained if $\text{input}(\mathcal{A})$ has an increasing subsequence $e_0 \preceq e_1 \preceq \dots \preceq e_j$ where $\text{dom}(\text{chan}(e_j)) = d_i$ and, for each k with $0 \leq k < j$,

$$\text{chan}(e_k) \in \text{vis}(\text{dom}(e_{k+1})).$$

In [49], van der Meyden's purge functions yield tree structures instead of subsequences; every path from a leaf to the root in these trees is a subsequence consisting of permissible effects $\text{chan}(e_k) \in \text{vis}(\text{dom}(e_{k+1}))$. This tightens the notion of security, because the trees do not include ordering information between events that lie on different branches to the root.

We formalize a *purge function* for a domain $d_i \in D$ as being a function from executions \mathcal{A} to some range set A . It should be sensitive only to *input* events in \mathcal{A} (condition 1), and it should certainly reflect *all* the inputs *visible* to level d_i (condition 2). In most existing definitions, the range A consists of sequences of input events, though in van der Meyden's [49], they are trees of input events. In [12], the range depends on how declassification conditions are defined.

Definition 33. Let \mathcal{F} be as in Fig. 4, and A any set. A function $p: \text{Exc}(\mathcal{F}) \rightarrow A$ is a d_i -purge function, where $d_i \in D$, iff

1. $\text{input}(\mathcal{A}) = \text{input}(\mathcal{A}')$ implies $p(\mathcal{A}) = p(\mathcal{A}')$;
2. $p(\mathcal{A}) = p(\mathcal{A}')$ implies $\mathcal{A} \upharpoonright \text{vis}(d_i) = \mathcal{A}' \upharpoonright \text{vis}(d_i)$.

If p is a d_i -purge, $\mathcal{A} \approx^p \mathcal{A}'$ means $p(\mathcal{A}) = p(\mathcal{A}')$. ///

Each purge p determines notions of noninterference and nondeducibility.

Definition 34. Let p be a purge function for $d_i \in D$. \mathcal{F} is p -noninterfering, written $\mathcal{F} \in \text{NI}^p$, iff, for all $\mathcal{A}, \mathcal{A}' \in \text{Exc}(\mathcal{F})$,

$$\mathcal{A} \approx^p \mathcal{A}' \text{ implies } \mathcal{A} \upharpoonright C_i = \mathcal{A}' \upharpoonright C_i.$$

\mathcal{F} is p -nondeducible ($\mathcal{F} \in \text{ND}^p$), iff, for all $\mathcal{A}, \mathcal{A}' \in \text{Exc}(\mathcal{F})$,

$$\mathcal{A} \approx^p \mathcal{A}' \text{ implies } \mathcal{A}' \upharpoonright \text{IN} \in J_{\text{IN} \triangleleft C_i}(\mathcal{A} \upharpoonright C_i). \quad ///$$

Here we take non-deducibility to mean that d_i 's observations provide no more information about all inputs than the purge p preserves. Thus, $\mathcal{A} \upharpoonright C_i$ is akin to Sutherland's *view* [48, Sec. 5.2]. We have adapted it slightly. Sutherland assumes that d_i also observes the outputs on the channels c_j^{out} where $d_j \hookrightarrow d_i$; we instead let the definition of machine M replicate them onto c_i^{out} when desired.

Sutherland's *hidden-from* could be interpreted as $\mathcal{A}' \upharpoonright \{c_j^{\text{in}} : d_j \not\rightarrow d_i\}$, i.e. the inputs that would not be visible to d_i . This agrees with our proposed definition in the case Sutherland considered, namely the classic GM purge for noninterference. The assumption $\mathcal{A} \approx^p \mathcal{A}'$ is meant to extend nondeducibility for other purges. As expected, noninterference is tighter than nondeducibility [48, Sec. 7]:

Lemma 35. Let p be a purge function for domain d_i . $\mathcal{F} \in \text{NI}^p$ implies $\mathcal{F} \in \text{ND}^p$.

Proof. Assume that $\mathcal{F} \in \text{NI}^p$ and $\mathcal{A}, \mathcal{A}' \in \text{Exc}(\mathcal{F})$, where $\mathcal{A} \approx^p \mathcal{A}'$.

By the definition, $\mathcal{A} \upharpoonright C_i = \mathcal{A}' \upharpoonright C_i$. Thus, $J_{\text{IN} \triangleleft C_i}(\mathcal{A} \upharpoonright C_i) = J_{\text{IN} \triangleleft C_i}(\mathcal{A}' \upharpoonright C_i)$.

But $\mathcal{A}' \upharpoonright \text{IN} \in J_{\text{IN} \triangleleft C_i}(\mathcal{A}' \upharpoonright C_i)$ because \mathcal{A}' is itself a witness. □

NI^p and ND^p are not equivalent, as ND^p has an additional (implicit) existential quantifier. The witness execution showing that $\mathcal{A}' \upharpoonright \text{IN} \in J_{\text{IN} \triangleleft C_i}(\mathcal{A} \upharpoonright C_i)$ may differ from \mathcal{A}' on channels $c \notin \text{IN} \cup C_i$, namely the output channels c_j^{out} for $j \neq i$.

The symmetry of nondisclosure (Lemma 11) does not hold for NI^p and ND^p . For instance, relative to the GM purge for flow to d_i , there may be noninterference for inputs at d_j , while there is interference for flow from d_i to d_j . The asymmetry arises because the events to be concealed are only inputs at the source, while the observed events are both inputs and outputs [48].

The idea of p -noninterference is useful only when M is deterministic, since otherwise the outputs observed on c_i^{out} may differ even when $\text{input}(\mathcal{A}) = \text{input}(\mathcal{A}')$. For non-deterministic M , nondeducibility is more natural.

Purges and blurs. We can associate a blur operator f^p with each purge function p , such that ND^p amounts to respecting the blur operator f^p . We regard ND^p as saying that the input/output events on C_i tell d_i no more about all the inputs than the purged input $p(\mathcal{A})$ would disclose. We use a compatibility relation where the observed channels and the source channels overlap on c_i^{in} .

Definition 36. Let p be a purge function for d_i , and define the equivalence relation $\mathcal{R} \subseteq (\text{IN-runs} \times \text{IN-runs})$ by the condition: $\mathcal{R}(\mathcal{B}_1, \mathcal{B}_2)$ iff

$$\exists \mathcal{A}_1, \mathcal{A}_2 \in \text{Exc}(\mathcal{F}) . \bigwedge_{j=1,2} \mathcal{B}_j = \mathcal{A}_j \upharpoonright \text{IN} \wedge \mathcal{A}_1 \approx^p \mathcal{A}_2. \quad (3)$$

Define $f^p: \mathcal{P}(\text{IN-runs}) \rightarrow \mathcal{P}(\text{IN-runs})$ to close under the \mathcal{R} -equivalence classes as in Lemma 22. ///

In fact, ND is a form of disclosure limited to within a blur:

Lemma 37. Let p be a purge function for domain d_i . For all \mathcal{F} , $\mathcal{F} \in \text{ND}^p$ iff \mathcal{F} f^p -limits IN-to- C_i flow.

Proof. **1. ND^p implies f^p -limited flow.** Suppose that $\mathcal{F} \in \text{ND}^p$; $\mathcal{B}_i \in C_i$ -runs; and $\mathcal{B}_1 \in J_{\text{IN} \triangleleft C_i}(\mathcal{B}_i)$. If $\mathcal{B}_2 \in f^p(\mathcal{B}_1)$, we must show that $\mathcal{B}_2 \in J_{\text{IN} \triangleleft C_i}(\mathcal{B}_i)$.

By Def. 36 there are $\mathcal{A}_1, \mathcal{A}_2$ so that $\mathcal{A}_1 \upharpoonright \text{IN} = \mathcal{B}_1$ and $\mathcal{A}_2 \upharpoonright \text{IN} = \mathcal{B}_2$ and $\mathcal{A}_1 \approx^p \mathcal{A}_2$. Furthermore, let \mathcal{A} witness $\mathcal{B}_1 \in J_{\text{IN} \triangleleft C_i}(\mathcal{B}_i)$. Then $\mathcal{A} \upharpoonright \text{IN} = \mathcal{B}_1 = \mathcal{A}_1 \upharpoonright \text{IN}$. So Def. 33, Clause 1 says $\mathcal{A} \approx^p \mathcal{A}_1$ and so also $\mathcal{A} \approx^p \mathcal{A}_2$. Since $\mathcal{F} \in \text{ND}^p$, $\mathcal{A}_2 \upharpoonright \text{IN} \in J_{\text{IN} \triangleleft C_i}(\mathcal{A} \upharpoonright C_i)$. That is, $\mathcal{B}_2 \in J_{\text{IN} \triangleleft C_i}(\mathcal{B}_i)$ as required.

2. f^p -limited flow implies ND^p . Assume $J_{\text{IN} \triangleleft C_i}(\mathcal{B}_i)$ is f^p -blurred for all \mathcal{B}_i . We must show that, for all $\mathcal{A}_1, \mathcal{A}_2$ such that $\mathcal{A}_1 \approx^p \mathcal{A}_2$, $(\mathcal{A}_2 \upharpoonright \text{IN}) \in J_{\text{IN} \triangleleft C_i}(\mathcal{A}_1 \upharpoonright C_i)$.

So choose executions with $\mathcal{A}_1 \approx^p \mathcal{A}_2$. By Def. 36, $\mathcal{R}(\mathcal{A}_1 \upharpoonright \text{IN}, \mathcal{A}_2 \upharpoonright \text{IN})$, since $\mathcal{A}_1, \mathcal{A}_2$ satisfy the condition. Thus, since $J_{\text{IN} \triangleleft C_i}(\mathcal{A}_1 \upharpoonright C_i)$ is f^p -blurred and contains $\mathcal{A}_1 \upharpoonright \text{IN}$, $\mathcal{A}_2 \upharpoonright \text{IN} \in J_{\text{IN} \triangleleft C_i}(\mathcal{A}_1 \upharpoonright C_i)$. □

Semantic sensitivity. Blur operators provide an explicit semantic representation of the information that will not be disclosed when flow is limited. This is in contrast to intransitive non-interference [42,23,49], which considers only whether the “ \hookrightarrow plumbing” among domains is correct.

Example 38. We represent Imaginary Weather Forecasting (IWF, see Example 21) as a state machine frame as in Fig. 4. It has domains $\{ws, \ell, p, cmp\}$ for the weather service, low-tier customer, premium-tier customer and compression service respectively. Let \hookrightarrow be the smallest reflexive (but intransitive) relation extending Eqn. 4, where all reports must flow through the compression service:

$$ws \hookrightarrow cmp \hookrightarrow p \text{ and } cmp \hookrightarrow \ell. \quad (4)$$

The *cmp* service should compress reports lossily before sending them to ℓ and compress them losslessly for p . However, a faulty *cmp* may compress losslessly for both ℓ and p . Purge functions [23,49] do not distinguish between correct and faulty *cmps*. In both cases, all information from *ws* does indeed pass through *cmp*. The blur of Example 21, however, defines the desired goal semantically. With the faulty *cmp*, the high-resolution data compatible with the observation of ℓ is more sharply defined than an f -blurred set. ///

8 Future Work

We have explored how the graph structure of a distributed system helps to constrain information flow. We have established the cut-blur principle. It allows us to propagate conclusions about limited disclosure from a cut set cut to more remote parts of the graph. We have also showed a sufficient condition for limitations on disclosure to be preserved under homomorphisms. Most of our examples concern networks and filtering, but the ideas are much more widely applicable.

Quantitative treatment. It should be possible to equip frames with a quantitative information flow semantics. One obstacle here is that our execution model mixes some choices which are natural to view probabilistically—for instance, selection between different outputs when both are permitted by an LTS—with others that seem non-deterministic. The choice between receiving an input and emitting an output is an example of this, as is the choice between receiving inputs on different channels. This problem has been studied (e.g. [8,9]), but a tractable semantics may require new ideas.

A Dynamic Model. Instead of building $\text{ends}(\ell)$ into the frame, so that it remains fixed through execution, we may alternatively regard it as a component of the states of the individual locations. Let us regard $\text{traces}(\ell)$ as generated by a labeled transition system $\text{lts}(\ell)$. Then we may enrich the labels c, v so that they also involve a sequence of endpoints $\bar{p} \subseteq \mathcal{EP}$:

$$(c, v, \bar{p}).$$

The transition relation of $\text{lts}(\ell)$ is then constrained to allow a transmission (c, v, \bar{p}) in a state only if $p \subseteq \text{ends}(\ell)$ holds in that state, in which case \bar{p} is omitted in the next state. A reception (c, v, \bar{p}) causes \bar{p} to be added to the next state of the receiving location.

The cut-blur principle remains true in an important case: A set cut is an *invariant cut* between src and obs if it is an undirected cut, and moreover the execution of the frame preserves this property. Then the cut-blur principle holds in the dynamic model for invariant cuts. Our definition of homomorphism (Appendix B) continues to apply.

This dynamic model suggests an analysis of security-aware software using object capabilities. Object capabilities may be viewed as transmission endpoints. To use a capability, one transmits a message down the channel to the object itself, which holds the reception endpoint. To transfer a capability, one transmits its transmission endpoint down another channel.

McCamant and Ernst [30] generate a directed graph of this sort in memory at runtime. They use the program source text to compute quantify flow from sensitive source to public sink for this run. One would like to provide a maximum over all possible runs, and this would appear to depend on inferring some invariants on the structure of the graphs. Our methods might be helpful for this.

Cryptographic Masking. Encryption is not a blur. Encrypting messages makes their contents unavailable in locations lacking the decryption keys. In particular, locations lacking the decryption key may form a cut set between the source and destination of the encrypted message. However, at the destinations, where the keys are available, the messages can be decrypted and their contents observed. Thus, the cut-blur theorem says it would be wrong to view encryption as a blur in this set-up: Its effects can be undone beyond the cut.

Instead, we enrich the model. The messages on a channel are not observed in their specificity. Instead, the observation is a projection of messages on the channel. One could represent this projection using the Abadi-Rogaway [1] pattern operator \square . Alternatively, one could rely directly on the cryptographic indistinguishability relation [4].

In general, this means applying the compatibility function $J_{C' \triangleleft C}(\mathcal{B})$ not only to local C -runs \mathcal{B} , but also to their projections $J_{C' \triangleleft C}(\text{proj}(\mathcal{B}))$ for projection functions of interest. This change is quite compatible with many of our proof methods, but the statements of theorems will need adaptation. We would like to use the resulting set-up to reason about cryptographic voting systems, such as Helios and Prêt-à-Voter [3,44].

We also intend to provide tool support for defining relevant blurs and establishing that they limit disclosure in several application areas.

Acknowledgments. We are grateful to Megumi Ando, Aslan Askarov, Stephen Chong, John Ramsdell, and Mitchell Wand. In particular, John Ramsdell formalized Thms. 26 and 30 in PVS, as well as Example 24; this suggested Example 32.

References

1. Martín Abadi and Phillip Rogaway. Reconciling two views of cryptography (the computational soundness of formal encryption). *Journal of Cryptology*, 15(2):103–127, 2002.
2. P Adao, Claudio Bozzato, R Focardi Gian-Luca Dei Rossi, and F Luccio. Mignis: A semantic based tool for firewall configuration. In *IEEE CSF*, 2014.
3. Ben Adida. Helios: Web-based open-audit voting. In *Usenix Security Symposium*. Usenix Association, 2008.
4. Aslan Askarov, Daniel Hedin, and Andrei Sabelfeld. Cryptographically-masked flows. *Theoretical Computer Science*, 402(2):82–101, 2008.
5. Aslan Askarov and Andrei Sabelfeld. Gradual release: Unifying declassification, encryption and key release policies. In *IEEE Symp. Security and Privacy*, pages 207–221. IEEE, 2007.
6. Aaron Bohannon, Benjamin C Pierce, Vilhelm Sjöberg, Stephanie Weirich, and Steve Zdancewic. Reactive noninterference. In *ACM conference on Computer and Communications Security*, pages 79–90. ACM, 2009.
7. Annalisa Bossi, Riccardo Focardi, Carla Piazza, and Sabina Rossi. Verifying persistent security properties. *Computer Languages, Systems & Structures*, 30(3):231–258, 2004.
8. Ran Canetti, Ling Cheung, Dilsun Kirli Kaynar, Moses Liskov, Nancy A. Lynch, Olivier Pereira, and Roberto Segala. Analyzing security protocols using time-bounded task-PIOAs. *Discrete Event Dynamic Systems*, 18(1):111–159, 2008.

9. Konstantinos Chatzikokolakis and Catuscia Palamidessi. Making random choices invisible to the scheduler. In *CONCUR*, pages 42–58. Springer, 2007.
10. Stephen Chong, Jed Liu, Andrew C Myers, Xin Qi, Krishnaprasad Vikram, Lantian Zheng, and Xin Zheng. Secure web applications via automatic partitioning. In *ACM SIGOPS Operating Systems Review*, volume 41, pages 31–44. ACM, 2007.
11. Stephen Chong and Ron van der Meyden. Deriving epistemic conclusions from agent architecture. In *Theoretical Aspects of Rationality and Knowledge*, pages 61–70. ACM, 2009.
12. Stephen Chong and Ron van der Meyden. Using architecture to reason about information security. Technical Report arXiv:1409.0309, Arxiv, Sept. 2014. <http://arxiv.org/abs/1409.0309v1>.
13. Jeremy Clark, Aleks Essex, and Carlisle Adams. On the security of ballot receipts in e2e voting systems. In *Workshop on Trustworthy Elections (WOTE)*, 2007.
14. George Coker, Joshua Guttman, Peter Loscocco, Amy Herzog, Jonathan Millen, Brian O’Hanlon, John Ramsdell, Ariel Segall, Justin Sheehy, and Brian Sniffen. Principles of remote attestation. *International Journal of Information Security*, 10(2):63–81, 2011.
15. Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein. *Introduction to Algorithms*. MIT Press, Cambridge, MA, second edition, 2003.
16. Stéphanie Delaune, Steve Kremer, and Mark Ryan. Verifying privacy-type properties of electronic voting protocols. *Journal of Computer Security*, 17(4):435–487, 2009.
17. Ronald Fagin, Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Reasoning about Knowledge*. MIT Press, Cambridge, MA, 1995.
18. Riccardo Focardi and Roberto Gorrieri. The compositional security checker: A tool for the verification of information flow security properties. *IEEE Transactions on Software Engineering*, 23(9), September 1997.
19. Riccardo Focardi and Roberto Gorrieri. Classification of security properties. In *Foundations of Security Analysis and Design*, pages 331–396. Springer, 2001.
20. Joseph A. Goguen and José Meseguer. Security policies and security models. In *IEEE Symposium on Security and Privacy*, 1982.
21. Joshua D. Guttman and Amy L. Herzog. Rigorous automated network security management. *International Journal for Information Security*, 5(1–2):29–48, 2005.
22. Joshua D. Guttman and Paul D. Rowe. A cut principle for information flow. Arxiv, October 2014. <http://arxiv.org/abs/1410.4617>.
23. J Thomas Haigh and William D Young. Extending the non-interference version of mls for sat. *IEEE Transactions on Software Engineering*, 13(2), 1987.
24. Jeremy Jacob. Basic theorems about security. *Journal of Computer Security*, 1(3):385–411, 1992.
25. Dale M Johnson and F Javier Thayer. Security and the composition of machines. In *CSFW*, volume 88, pages 72–89, 1988.
26. Heiko Mantel. Possibilistic definitions of security—an assembly kit. In *IEEE Computer Security Foundations*, pages 185–199. IEEE, 2000.
27. Heiko Mantel. Preserving information flow properties under refinement. In *IEEE Computer Security Foundations*, pages 78–91, 2001.
28. Heiko Mantel. On the composition of secure systems. In *IEEE Security and Privacy*, pages 88–101, 2002.
29. Heiko Mantel, David Sands, and Henning Sudbrock. Assumptions and guarantees for compositional noninterference. In *IEEE Computer Security Foundations Symposium*, pages 218–232, 2011.

30. Stephen McCamant and Michael D Ernst. Quantitative information flow as network flow capacity. *ACM SIGPLAN Notices (PLDI)*, 43(6):193–205, 2008.
31. Daryl McCullough. Specifications for multi-level security and a hook-up property. In *IEEE Symposium on Security and Privacy*, pages 161–161, 1987.
32. Daryl McCullough. Noninterference and the composability of security properties. In *IEEE Symposium on Security and Privacy*, pages 177–186, 1988.
33. John McLean. A general theory of composition for trace sets closed under selective interleaving functions. In *IEEE Symp. Security and Privacy*, pages 79–93, 1994.
34. Murat Moran, James Heather, and Steve Schneider. Automated anonymity verification of the ThreeBallot and VAV voting systems. *Software and Systems Modeling*, 2015. Forthcoming.
35. Carroll Morgan. The shadow knows: Refinement of ignorance in sequential programs. In *Mathematics of program construction*, pages 359–378. Springer, 2006.
36. Willard Rafnsson and Andrei Sabelfeld. Compositional information-flow security for interactive systems. In *IEEE Symp. CSF*, pages 277–292, July 2014.
37. Ronald L Rivest and Warren D Smith. Three voting protocols: ThreeBallot, VAV, and Twin. <http://people.csail.mit.edu/rivest/pubs/RS07.pdf>, 2007.
38. A. W. Roscoe and M. H. Goldsmith. What is intransitive noninterference? In *12th IEEE Computer Security Foundations Workshop*, pages 228–238. IEEE CS Press, June 1999.
39. A William Roscoe. CSP and determinism in security modelling. In *IEEE Security and Privacy*, pages 114–127. IEEE, 1995.
40. A.W. Roscoe, J.C.P. Woodcock, and L. Wulf. Non-interference through determinism. In *Computer Security—ESORICS 94*, pages 31–53. Springer, 1994.
41. Sabina Rossi and Damiano Macedonio. Information flow security for service compositions. In *ICUMT*, pages 1–8, 2009.
42. John Rushby. *Noninterference, transitivity, and channel-control security policies*. SRI International, Computer Science Laboratory, 1992.
43. Peter Ryan and Steve Schneider. Process algebra and non-interference. *J. Comput. Secur.*, 9(1-2):75–103, January 2001.
44. Peter YA Ryan, David Bismark, James Heather, Steve Schneider, and Zhe Xia. Prêt à voter: a voter-verifiable voting system. *Information Forensics and Security, IEEE Transactions on*, 4(4):662–673, 2009.
45. Andrei Sabelfeld and Andrew C. Myers. Language-based information-flow security. *IEEE Journal on Selected Areas in Communication*, 21(1):5–19, January 2003.
46. Andrei Sabelfeld and Andrew C Myers. A model for delimited information release. In *Software Security-Theories and Systems*, pages 174–191. Springer, 2004.
47. Andrei Sabelfeld and David Sands. Declassification: Dimensions and principles. *Journal of Computer Security*, 17(5):517–548, 2009.
48. David Sutherland. A model of information. In *9th National Computer Security Conference*. National Institute of Standards and Technology, 1986.
49. Ron van der Meyden. What, indeed, is intransitive noninterference? In *Computer Security—ESORICS 2007*, pages 235–250. Springer, 2007.
50. Ron van der Meyden. Architectural refinement and notions of intransitive noninterference. *Formal Aspects of Computing*, 24(4-6):769–792, 2012.
51. Dennis Volpano, Geoffrey Smith, and Cynthia Irvine. A sound type system for secure flow analysis. *Journal of Computer Security*, 4(3):1–21, 1996.
52. J Todd Wittbold and Dale M Johnson. Information flow in nondeterministic systems. In *IEEE Symp. Security and Privacy*, pages 144–144, 1990.
53. Aris Zakynthinos and E Stewart Lee. The composability of non-interference. *Journal of Computer Security*, 3(4):269–281, 1995.

54. Steve Zdancewic, Lantian Zheng, Nathaniel Nystrom, and Andrew C Myers. Secure program partitioning. *ACM Transactions on Computer Systems (TOCS)*, 20(3):283–328, 2002.

A Some Lemmas and Proofs

Lemma 12.

1. Suppose $C_0 \subseteq C_1$ and $C'_0 \subseteq C'_1$. If \mathcal{F} has no disclosure from C_1 to C'_1 , then \mathcal{F} has no disclosure from C_0 to C'_0 .
2. When $C_1, C_2, C_3 \subseteq \mathcal{CH}$,

$$J_{C_3 \triangleleft C_1}(\mathcal{B}_1) \subseteq \bigcup_{\mathcal{B}_2 \in J_{C_2 \triangleleft C_1}(\mathcal{B}_1)} J_{C_3 \triangleleft C_2}(\mathcal{B}_2).$$

Proof. 1. Suppose \mathcal{B}_0 is a C_0 -run, and \mathcal{B}'_0 is a C'_0 -run. We want to show that $\mathcal{B}'_0 \in J_{C'_0 \triangleleft C_0}(\mathcal{B}_0)$.

Since they are local runs, there exist $\mathcal{A}_0, \mathcal{A}'_0 \in \text{Exc}(\mathcal{F})$ such that $\mathcal{B}_0 = \mathcal{A}_0 \upharpoonright C_0$ and $\mathcal{B}'_0 = \mathcal{A}'_0 \upharpoonright C'_0$. But let $\mathcal{B}_1 = \mathcal{A}_0 \upharpoonright C_1$ and let $\mathcal{B}'_1 = \mathcal{A}'_0 \upharpoonright C'_1$. By no-disclosure, $\mathcal{B}'_1 \in J_{C'_1 \triangleleft C_1}(\mathcal{B}_1)$. So there is an $\mathcal{A} \in \text{Exc}(\mathcal{F})$ such that $\mathcal{B}_1 = \mathcal{A} \upharpoonright C_1$ and $\mathcal{B}'_1 = \mathcal{A} \upharpoonright C'_1$.

However, then \mathcal{A} witnesses for $\mathcal{B}'_0 \in J_{C'_0 \triangleleft C_0}(\mathcal{B}_0)$: After all, since $C_0 \subseteq C_1$, $\mathcal{A} \upharpoonright C_0 = (\mathcal{A} \upharpoonright C_1) \upharpoonright C_0$. Similarly for the primed versions.

2. Suppose that $\mathcal{B}_3 \in J_{C_3 \triangleleft C_1}(\mathcal{B}_1)$, so that there exists an $\mathcal{A} \in \text{Exc}(\mathcal{F})$ such that $\mathcal{B}_1 = \mathcal{A} \upharpoonright C_1$ and $\mathcal{B}_3 = \mathcal{A} \upharpoonright C_3$. Letting $\mathcal{B}_2 = \mathcal{A} \upharpoonright C_2$, the execution \mathcal{A} ensures that $\mathcal{B}_2 \in J_{C_2 \triangleleft C_1}(\mathcal{B}_1)$ and $\mathcal{B}_3 \in J_{C_3 \triangleleft C_2}(\mathcal{B}_2)$. \square

We now consider different frames $\mathcal{F}_1, \mathcal{F}_2$ that overlap on a common subset L_0 , and show how local runs in the two can be pieced together. In this context, we use the notation of Def. 28.

Lemma 39. *Let L_0 be shared between frames $\mathcal{F}_1, \mathcal{F}_2$. Let*

$$\mathcal{B}_{lc} \in (\text{left} \cup \text{cut})\text{-runs}_1 \text{ and } \mathcal{B}_{rc} \in (\text{right}_2 \cup \text{cut})\text{-runs}_2$$

agree on cut, i.e. $\mathcal{B}_{lc} \upharpoonright \text{cut} = \mathcal{B}_{rc} \upharpoonright \text{cut}$. Then there is an $\mathcal{A} \in \text{Exc}(\mathcal{F}_2)$ such that

$$\mathcal{B}_{lc} = \mathcal{A} \upharpoonright (\text{left} \cup \text{cut}) \text{ and } \mathcal{B}_{rc} = \mathcal{A} \upharpoonright (\text{right}_2 \cup \text{cut}).$$

Proof. Since \mathcal{B}_{lc} and \mathcal{B}_{rc} are local runs of $\mathcal{F}_1, \mathcal{F}_2$ resp., they are restrictions of executions, so choose $\mathcal{A}_1 \in \text{Exc}(\mathcal{F}_1)$ and $\mathcal{A}_2 \in \text{Exc}(\mathcal{F}_2)$ so that $\mathcal{B}_{lc} = \mathcal{A}_1 \upharpoonright (\text{left} \cup \text{cut})$ and $\mathcal{B}_{rc} = \mathcal{A}_2 \upharpoonright (\text{right}_2 \cup \text{cut})$. Now define \mathcal{A} by stipulating:

$$\text{ev}(\mathcal{A}) = \text{ev}(\mathcal{B}_{lc}) \cup \text{ev}(\mathcal{B}_{rc}) \tag{5}$$

$$\preceq_{\mathcal{A}} = \text{the least partial order extending } \preceq_{\mathcal{B}_{lc}} \cup \preceq_{\mathcal{B}_{rc}}. \tag{6}$$

Since $\mathcal{A}_1, \mathcal{A}_2$ agree on cut , $\text{ev}(\mathcal{A}) = \text{ev}(\mathcal{B}_{lc} \upharpoonright \text{left}) \cup \text{ev}(\mathcal{B}_{rc})$, and we could have used the latter as an alternate definition of $\text{ev}(\mathcal{A})$, as well as the symmetric restriction of \mathcal{B}_{rc} to right_2 leaving \mathcal{B}_{lc} whole.

The definition of $\preceq_{\mathcal{A}}$ as a partial order is sound, because there are no cycles in the union (6). Cycles would require \mathcal{A}_1 and \mathcal{A}_2 to disagree on the order of events in their restrictions to cut , contrary to assumption. Likewise, the finite-predecessor property is preserved: $x_0 \preceq_{\mathcal{A}} x_1$ iff x_0, x_1 belong to the same $\mathcal{B}_{?c}$ and are ordered there, or else there is an event in $\mathcal{B}_{?c} \upharpoonright \text{cut}$ which comes between them. So the events preceding x_1 form the finite union of finite sets. Thus, $\mathcal{A} \in \text{ES}(\mathcal{F}_2)$.

Moreover, \mathcal{A} is an execution $\mathcal{A} \in \text{Exc}(\mathcal{F}_2)$: If $\ell \in L_0$, then $\text{proj}(\mathcal{A}, \ell) = \text{proj}(\mathcal{B}_{lc}, \ell)$, and the latter is a trace in $\text{traces}_1(\ell) = \text{traces}_2(\ell)$. If $\ell \notin L_0$, then $\text{proj}(\mathcal{A}, \ell) = \text{proj}(\mathcal{B}_{rc}, \ell)$, and the latter is a trace in $\text{traces}_2(\ell)$.

There is no ℓ with channels in both left and right_2 . \square

What makes this proof work? Any one location either has all of its channels lying in $\text{left}_0 \cup \text{cut}_0$ or else all of them lying in $\text{right}_i \cup \text{cut}$. When piecing together the two executions $\mathcal{A}_1, \mathcal{A}_2$ into a single execution \mathcal{A} , no location needs to be able to execute a trace that comes partly from \mathcal{A}_1 and partly from \mathcal{A}_2 . This is what determines our definition of cuts using the undirected graph $\text{ungr}(\mathcal{F})$.

We next prove the two-frame analog of Lemma 14.

Lemma 29. *Let L_0 be shared between frames $\mathcal{F}_1, \mathcal{F}_2$. Let $\text{src} \subseteq \text{left}$, and $\mathcal{B}_c \in \text{cut}_0\text{-runs}_1 \cap \text{cut}_0\text{-runs}_2$.*

1. $J_{\text{src} \triangleleft \text{cut}_0}^1(\mathcal{B}_c) = J_{\text{src} \triangleleft \text{cut}_0}^2(\mathcal{B}_c)$.
2. Assume $\text{cut}_0\text{-runs}(\mathcal{F}_2) \subseteq \text{cut}_0\text{-runs}(\mathcal{F}_1)$. Let $\text{obs} \subseteq \text{right}_2$, and $\mathcal{B}_o \in \text{obs-runs}_2$. Then

$$J_{\text{src} \triangleleft \text{obs}}^2(\mathcal{B}_o) = \bigcup_{\mathcal{B}_c \in J_{\text{cut}_0 \triangleleft \text{obs}}^2(\mathcal{B}_o)} J_{\text{src} \triangleleft \text{cut}_0}^1(\mathcal{B}_c).$$

Proof. **1.** First, we show that $\mathcal{B}_s \in J_{\text{src} \triangleleft \text{cut}_0}^1(\mathcal{B}_c)$ implies $\mathcal{B}_s \in J_{\text{src} \triangleleft \text{cut}_0}^2(\mathcal{B}_c)$.

Let \mathcal{A}_1 witness for $\mathcal{B}_s \in J_{\text{src} \triangleleft \text{cut}_0}^1(\mathcal{B}_c)$, and let \mathcal{A}_2 witness for $\mathcal{B}_c \in \text{cut-runs}_2$. Define

$$\mathcal{B}_{lc} = \mathcal{A}_1 \upharpoonright (\text{left} \cup \text{cut}) \text{ and } \mathcal{B}_{rc} = \mathcal{A}_2 \upharpoonright (\text{right}_2 \cup \text{cut}).$$

Now the assumptions for Lemma 39 are satisfied. So let $\mathcal{A} \in \text{Exc}(\mathcal{F}_2)$ restrict to \mathcal{B}_{lc} and \mathcal{B}_{rc} as in the conclusion. Thus, $\mathcal{A} \upharpoonright \text{src} = \mathcal{B}_s$.

For the converse, we rely on the symmetry of “ L_0 is shared between frames $\mathcal{F}_1, \mathcal{F}_2$.”

- 2.** By the assumption, whenever $\mathcal{B}_c \in J_{\text{cut}_0 \triangleleft \text{obs}}^2(\mathcal{B}_o)$, then also $\mathcal{B}_c \in \text{cut-runs}_1$. Thus, we can apply part 1 after using Lemma 12:

$$\begin{aligned} J_{\text{src} \triangleleft \text{obs}}^2(\mathcal{B}_o) &\subseteq \bigcup_{\mathcal{B}_c \in J_{\text{cut}_0 \triangleleft \text{obs}}^2(\mathcal{B}_o)} J_{\text{src} \triangleleft \text{cut}_0}^2(\mathcal{B}_c) \\ &\subseteq \bigcup_{\mathcal{B}_c \in J_{\text{cut}_0 \triangleleft \text{obs}}^2(\mathcal{B}_o)} J_{\text{src} \triangleleft \text{cut}_0}^1(\mathcal{B}_c). \end{aligned}$$

For the reverse inclusion, assume that $\mathcal{B}_s \in J_{\text{src} \triangleleft \text{cut}_0}^1(\mathcal{B}_c)$, where $\mathcal{B}_c \in J_{\text{cut}_0 \triangleleft \text{obs}}^2(\mathcal{B}_o)$. Thus, we can apply Lemma 39, obtaining $\mathcal{A} \in \text{Exc}(\mathcal{F}_2)$ which agrees with \mathcal{B}_s , \mathcal{B}_c , and \mathcal{B}_o . So \mathcal{A} witnesses for $\mathcal{B}_s \in J_{\text{src} \triangleleft \text{obs}}^2(\mathcal{B}_o)$. \square

We now turn to the one-frame corollary, which we presented earlier as Lemma 14.

Lemma 14. *Let cut be an undirected cut between src , obs , and let $\mathcal{B}_o \in \text{src-runs}$. Then*

$$J_{\text{src} \triangleleft \text{obs}}(\mathcal{B}_o) = \bigcup_{\mathcal{B}_c \in J_{\text{cut} \triangleleft \text{obs}}(\mathcal{B}_o)} J_{\text{src} \triangleleft \text{cut}}(\mathcal{B}_c).$$

Proof. Define L_0 to be the smallest set of locations such that

1. $\ell \in L_0$ if $\text{chans}(\ell) \cap \text{src} \neq \emptyset$;
2. L_0 is closed under reachability by paths that do not traverse cut .

L_0 is shared between \mathcal{F} and itself. Moreover, for the set of channels cut_0 defined in Def. 28, we have $\text{cut}_0 \subseteq \text{cut}$: cut_0 is the part of cut that actually lies on the boundary of L_0 .

By Lemma 29, we have

$$J_{\text{src} \triangleleft \text{obs}}(\mathcal{B}_o) = \bigcup_{\mathcal{B}_c \in J_{\text{cut}_0 \triangleleft \text{obs}}(\mathcal{B}_o)} J_{\text{src} \triangleleft \text{cut}}(\mathcal{B}_c).$$

Since $\text{cut}_0 \subseteq \text{cut}$,

$$\bigcup_{\mathcal{B}_c \in J_{\text{cut}_0 \triangleleft \text{obs}}(\mathcal{B}_o)} J_{\text{src} \triangleleft \text{cut}}(\mathcal{B}_c) \subseteq \bigcup_{\mathcal{B}_c \in J_{\text{cut} \triangleleft \text{obs}}(\mathcal{B}_o)} J_{\text{src} \triangleleft \text{cut}}(\mathcal{B}_c).$$

For the converse, suppose that $\mathcal{B}_s \in J_{\text{src} \triangleleft \text{cut}}(\mathcal{B}_c)$, for $\mathcal{B}_c \in J_{\text{cut} \triangleleft \text{obs}}(\mathcal{B}_o)$. Then there is \mathcal{A} such that $\mathcal{A} \upharpoonright \text{src} = \mathcal{B}_s$ and $\mathcal{A} \upharpoonright \text{obs} = \mathcal{B}_o$. Thus, $\mathcal{B}_s \in J_{\text{src} \triangleleft \text{cut}}(\mathcal{A} \upharpoonright \text{cut}_0)$ and $\mathcal{A} \upharpoonright \text{cut}_0 \in J_{\text{cut}_0 \triangleleft \text{obs}}(\mathcal{B}_o)$. \square

The cut-blur principle is also the one-frame corollary of Thm. 30. The proofs are very similar.

Theorem 30. *Suppose that L_0 is shared between frames $\mathcal{F}_1, \mathcal{F}_2$, and assume $\text{cut-runs}(\mathcal{F}_2) \subseteq \text{cut-runs}(\mathcal{F}_1)$. Consider any $\text{src} \subseteq \text{left}$ and $\text{obs} \subseteq \text{right}_2$. If \mathcal{F}_1 f -limits src-to-cut flow, then \mathcal{F}_2 f -limits src-to-obs flow.*

Proof. By the hypothesis, f is a blur operator. Letting $\mathcal{B}_o \in \text{obs-runs}_2$, we want to show that $J_{\text{src} \triangleleft \text{obs}}^2(\mathcal{B}_o)$ is an f -blurred set, i.e. $J_{\text{src} \triangleleft \text{obs}}^2(\mathcal{B}_o) = f(J_{\text{src} \triangleleft \text{obs}}^2(\mathcal{B}_o))$.

For convenience, let $S_c = J_{\text{cut} \triangleleft \text{obs}}^2(\mathcal{B}_o)$. By Lemma 29,

$$J_{\text{src} \triangleleft \text{obs}}^2(\mathcal{B}_o) = \bigcup_{\mathcal{B}_c \in S_c} J_{\text{src} \triangleleft \text{cut}}^1(\mathcal{B}_c);$$

thus, we must show that the latter is f -blurred. By the assumption that each $J_{\text{src} \triangleleft \text{cut}}^1(\mathcal{B}_c)$ is f -blurred, we have $J_{\text{src} \triangleleft \text{cut}}^1(\mathcal{B}_c) = f(J_{\text{src} \triangleleft \text{cut}}^1(\mathcal{B}_c))$. Using this and the union property (Eqn. 1):

$$\begin{aligned} \bigcup_{\mathcal{B}_c \in S_c} J_{\text{src} \triangleleft \text{cut}}^1(\mathcal{B}_c) &= \bigcup_{\mathcal{B}_c \in S_c} f(J_{\text{src} \triangleleft \text{cut}}^1(\mathcal{B}_c)) \\ &= f\left(\bigcup_{\mathcal{B}_c \in S_c} J_{\text{src} \triangleleft \text{cut}}^1(\mathcal{B}_c)\right), \end{aligned}$$

Hence, $J_{\text{src} \triangleleft \text{obs}}^2(\mathcal{B}_o)$ is f -blurred. \square

B Frame homomorphisms

This appendix is included only for the convenience of the referees. See the technical report version on arxiv.org, [22].

Homomorphisms in general are structure-preserving maps, and frame homomorphisms are maps that preserve both graph structure and also executions. Homomorphisms are useful for representing refinement steps.

Definition 40. Suppose given frames $\mathcal{F} = (\mathcal{LO}, \mathcal{CH}, \mathcal{D}, \text{ends}, \dots)$ and $\mathcal{F}' = (\mathcal{LO}', \mathcal{CH}', \mathcal{D}', \text{ends}', \dots)$. Let $h = (h_{\mathcal{LO}}, h_{\mathcal{CH}}, h_{\mathcal{D}}, h_{\text{runs}})$, where

$$h_L: \mathcal{LO} \rightarrow \mathcal{LO}', \quad h_C: \mathcal{CH} \rightarrow \mathcal{CH}', \quad h_D: \mathcal{D} \rightarrow \mathcal{D}',$$

and $h_{\text{runs}}: \text{ES}(\mathcal{F}) \rightarrow \text{ES}(\mathcal{F}')$. We apply h to values using the relevant components, and lift h to sets, writing $h(C)$ for $\{h(c): c \in C\} \subseteq \mathcal{CH}'$, etc.

1. h is a graph homomorphism $h: \mathcal{F} \rightarrow \mathcal{F}'$ iff, for all $c \in \mathcal{CH}$,

$$h(\text{sender}(c)) = \text{sender}(h(c)) \quad \text{and} \quad h(\text{rcpt}(c)) = \text{rcpt}(h(c)).$$

2. A graph homomorphism h is a frame homomorphism, written $h: \mathcal{F} \rightarrow \mathcal{F}'$, iff, for all $C \subseteq \mathcal{CH}$ and $\mathcal{B} \in C\text{-runs}(\mathcal{F})$,
 - (a) $h(\mathcal{B}) \in h(C)\text{-runs}(\mathcal{F}')$;
 - (b) for some \preceq -preserving $\phi: \text{ev}(\mathcal{B}) \rightarrow \text{ev}(h(\mathcal{B}))$, for all $e \in \text{ev}(\mathcal{B})$,

$$\text{chan}(\phi(e)) = h(\text{chan}(e)) \quad \text{and} \quad \text{msg}(\phi(e)) = h(\text{msg}(e));$$

- (c) if $C_0 \subseteq h^{-1}(C') \subseteq C$ and $h(C_0) \subseteq C'$, then

$$h(\mathcal{B} \upharpoonright h^{-1}(C')) = h(\mathcal{B}) \upharpoonright C'. \quad ///$$

A homomorphism $h: \mathcal{F} \rightarrow \mathcal{F}'$ is *onto* if the functions $h_{\mathcal{LO}}, h_{\mathcal{CH}}, h_{\mathcal{D}}, h_{\text{runs}}$ are surjective. In this case, \mathcal{F} exhibits finer structure than \mathcal{F}' , which has “forgotten” structure from \mathcal{F} . So \mathcal{F} is a kind of refinement of \mathcal{F}' . When $h: \mathcal{F} \rightarrow \mathcal{F}'$ is an *embedding*, i.e. its functions are injective, there is more information and structure in \mathcal{F}' . Then we can consider \mathcal{F}' to be a different kind of refinement.

Definition 41. Let $h: \mathcal{F} \rightarrow \mathcal{F}'$, and $C_1, C_2 \subseteq \mathcal{CH}$. Then h extends pre-images from C_1 to C_2 iff, for all $\mathcal{B}_1 \in C_1\text{-runs}(\mathcal{F})$ and $\mathcal{B}' \in h(C_1) \cup h(C_2)\text{-runs}(\mathcal{F}')$, if $\mathcal{B}' \upharpoonright h(C_1) = h(\mathcal{B}_1)$, then there is a $\mathcal{B} \in (C_1 \cup C_2)\text{-runs}(\mathcal{F})$ such that $\mathcal{B} \upharpoonright C_1 = \mathcal{B}_1$ and $h(\mathcal{B}) = \mathcal{B}'$. $///$

Theorem 42. Let $h: \mathcal{F} \rightarrow \mathcal{F}'$ and $\text{src}, \text{obs} \subseteq \mathcal{CH}$; and assume

1. $h^{-1}(h(\text{src})) = \text{src}$ and $h^{-1}(h(\text{obs})) = \text{obs}$;
2. h is injective on src -runs;
3. h extends pre-images from src to obs .

If \mathcal{F} restricts disclosure from src to obs to within f , then \mathcal{F}' restricts disclosure from $h(\text{src})$ to $h(\text{obs})$ to within the function g such that:

$$g(S) = S \cup h(f(h^{-1}(S))).$$

Proof. We must show that g is a blur operator on $h(\text{src})$ -runs, and then that the compatible sets are g -blurred.

1. The form of the definition guarantees the *Inclusion* and *Union* properties, so we only need to check *Idempotence*. Clearly, the union with S does not affect idempotence. Thus, we want to show that $g \circ g = (h \circ f \circ h^{-1}) \circ (h \circ f \circ h^{-1}) = g$.

Since h is injective on src -runs, if $S \subseteq \text{src}$ -runs, then $h^{-1}(h(S)) = S$. Thus, by associativity and the idempotence of f :

$$\begin{aligned} g \circ g &= (h \circ f \circ h^{-1}) \circ (h \circ f \circ h^{-1}) \\ &= h \circ f \circ f \circ h^{-1} \\ &= h \circ f \circ h^{-1} = g. \end{aligned}$$

2. Suppose that $\mathcal{B}_{hs} \in J_{h(\text{src}) \triangleleft h(\text{obs})}(\mathcal{B}_{ho})$. By the compatibility assumption, there is a $\mathcal{B}_h \in (h(\text{src}) \cup h(\text{obs}))$ -runs such that $\mathcal{B}_h \upharpoonright h(\text{src}) = \mathcal{B}_{hs}$ and $\mathcal{B}_h \upharpoonright h(\text{obs}) = \mathcal{B}_{ho}$.

Whenever $\mathcal{B}'_{hs} \in g(\{\mathcal{B}_{hs}\})$, we must show that $\mathcal{B}'_{hs} \in J_{h(\text{src}) \triangleleft h(\text{obs})}(\mathcal{B}_{ho})$. By the definition of g , either $\mathcal{B}'_{hs} \in \{\mathcal{B}_{hs}\}$ or else $\mathcal{B}'_{hs} \in h(f(h^{-1}(\{\mathcal{B}_{hs}\})))$. In the first case, we are done.

So assume the latter; thus, there exist $\mathcal{B}_s, \mathcal{B}'_s$ such that $h(\mathcal{B}_s) = \mathcal{B}_{hs}$, $h(\mathcal{B}'_s) = \mathcal{B}'_{hs}$, and $\mathcal{B}'_s \in f(\{\mathcal{B}_s\})$.

By the assumption that h extends pre-images from src to obs , there is a $\mathcal{B} \in (\text{src} \cup \text{obs})$ -runs(\mathcal{F}) such that $\mathcal{B} \upharpoonright \text{src} = \mathcal{B}_s$ and $h(\mathcal{B}) = \mathcal{B}_h$. Since $h^{-1}(h(\text{obs})) = \text{obs}$, Def. 40, Clause 2c implies

$$h(\mathcal{B} \upharpoonright \text{obs}) = h(\mathcal{B}) \upharpoonright h(\text{obs}) = \mathcal{B}_{ho}.$$

Let \mathcal{B}_o be $\mathcal{B} \upharpoonright \text{obs}$. Since \mathcal{B} is a witness, $\mathcal{B}_s \in J_{\text{src} \triangleleft \text{obs}}(\mathcal{B}_o)$. By assumption, the compatible runs are f -blurred, so $\mathcal{B}'_s \in J_{\text{src} \triangleleft \text{obs}}(\mathcal{B}_o)$. Thus, there is a $\mathcal{B}' \in (\text{src} \cup \text{obs})$ -runs(\mathcal{F}) such that $\mathcal{B}' \upharpoonright \text{src} = \mathcal{B}'_s$ and $\mathcal{B}' \upharpoonright \text{obs} = \mathcal{B}_o$. By Def. 40, Clause 2a, $h(\mathcal{B}') \in h(\text{src} \cup \text{obs})$ -runs(\mathcal{F}'). Since $h^{-1}(h(\text{src})) = \text{src}$, by Clause 2c,

$$h(\mathcal{B}') \upharpoonright h(\text{src}) = h(\mathcal{B}') \upharpoonright \text{src} = \mathcal{B}'_{hs}.$$

Similarly,

$$h(\mathcal{B}') \upharpoonright h(\text{obs}) = h(\mathcal{B}') \upharpoonright \text{obs} = h(\mathcal{B}_o) = \mathcal{B}_{ho}.$$

So $h(\mathcal{B}')$ is a witness for $\mathcal{B}'_{hs} \in J_{h(\text{src}) \triangleleft h(\text{obs})}(\mathcal{B}_{ho})$. □