

Security and Trust I:

4. Flow Security

Dusko Pavlovic

UHM ICS 355
Fall 2014

Outline

Covert channels and flows

Possibilistic models

Probabilistic models

Quantifying noninterference

What did we learn?

Outline

Covert channels and flows

Interference

Definition of covert channel

Examples

Possibilistic models

Probabilistic models

Quantifying noninterference

What did we learn?

Covert

Interference

Definition

Examples

Possibilistic

Probabilistic


Quantifying

Lesson

The elevator example again

Elevator model

- ▶ $Q = \{\text{floor}_0, \text{floor}_1\}$
- ▶ $I_k = \{k:\text{call}_0, k:\text{call}_1\}$, $k \in \mathbb{L} = \{\text{Alice}, \text{Bob}\}$
- ▶ $O = \{\text{go}_0, \text{go}_1, \text{stay}\}$
- ▶ θ :



The elevator example again

Elevator interference

The histories

(A:call0 B:call1) and (A:call1 B:call1)

are for Bob

- ▶ indistinguishable by the inputs, since he only sees Bob:call1 in both of them, yet they are
- ▶ distinguishable by the outputs, since Bob's channel outputs are
 - ▶ (A:call0 B:call1) \mapsto go1
 - ▶ (A:call1 B:call1) \mapsto stay

The elevator example again

Covert

Interference

Definition

Examples

Possibilistic

Probabilistic

Quantifying

Lesson

Question

How does Bob really use the interference?

The elevator example again

Answer

He derives *another* channel

$$\frac{\{A:\text{call}0, A:\text{call}1, B:\text{call}0, B:\text{call}1\}^+ \rightarrow \{\text{stay}, \text{go}\}}{\{B:\text{call}0, B:\text{call}1\}^+ \rightarrow \{A_home, A_out\}}$$

The elevator example again

Answer

He derives *another* channel

$$\frac{\{A:\text{call}0, A:\text{call}1, B:\text{call}0, B:\text{call}1\}^+ \rightarrow \{\text{stay}, \text{go}\}}{\{B:\text{call}0, B:\text{call}1\}^+ \rightarrow \{A_home, A_out\}}$$

This is a *covert channel*.

The elevator example again

Different flows

- ▶ $\{A:\text{call}0, A:\text{call}1, B:\text{call}0, B:\text{call}1\}^+ \rightarrow \{\text{stay}, \text{go}\}$
makes Alice and Bob flow through the elevator
- ▶ $\{B:\text{call}0, B:\text{call}1\}^+ \rightarrow \{A_home, A_out\}$
makes the information about Alice flow to Bob

What is *flow*?

Intuition

The *flow* of a channel is the observed traffic that flows through it

- ▶ (water flow, information flow, traffic flow...)

Covert

Interference

Definition

Examples

Possibilistic

Probabilistic

Quantifying

Lesson

What is *flow*?

Flow vs channel

- ▶ A deterministic unshared channel implements a single flow. There are two usages
 - ▶ either the channel $I^+ \xrightarrow{f} O$ induces the flow $I^* \xrightarrow{\vec{f}} O^*$
 - ▶ or the history \vec{x} induces the flow $\vec{f}(\vec{x})$ along the channel $I^+ \xrightarrow{f} O$

Covert

Interference

Definition

Examples

Possibilistic

Probabilistic

Quantifying

Lesson

What is *flow*?

Flow vs channel

- ▶ A deterministic unshared channel implements a single flow. There are two usages
 - ▶ either the channel $I^+ \xrightarrow{f} O$ induces the flow $I^* \xrightarrow{\vec{f}} O^*$
 - ▶ or the history \vec{x} induces the flow $\vec{f}(\vec{x})$ along the channel $I^+ \xrightarrow{f} O$
- ▶ A deterministic *shared* channel $I^+ \xrightarrow{\vec{f}} O$ contains the flows $I_k^* \xrightarrow{\vec{f}_k} O^*$.
 - ▶ The mapping $I^* \xrightarrow{\vec{f}} O^*$ is a flow only if there is a global observer.

What is *flow*?

Flow vs channel

- ▶ A deterministic unshared channel implements a single flow. There are two usages
 - ▶ either the channel $I^+ \xrightarrow{f} O$ induces the flow $I^* \xrightarrow{\vec{f}} O^*$
 - ▶ or the history \vec{x} induces the flow $\vec{f}(\vec{x})$ along the channel $I^+ \xrightarrow{f} O$
- ▶ A deterministic *shared* channel $I^+ \xrightarrow{\vec{f}} O$ contains the flows $I_k^* \xrightarrow{\vec{f}_k} O^*$.
 - ▶ The mapping $I^* \xrightarrow{\vec{f}} O^*$ is a flow only if there is a global observer.
- ▶ A possibilistic channel $I^+ \xrightarrow{f} \wp O$ contains multiple deterministic channels which induce the possible flows

Channeling interference

In general, any user k who seeks the interferences in a shared channel \vec{f} builds a derived *interference channel* \widehat{f}_k

$$\frac{I^* \xrightarrow{\vec{f}} O^*}{\begin{array}{l} I_k^* \quad \widehat{f}_k \\ \vec{x}_k \quad \mapsto \quad \{ \vec{f}_k(\vec{y}) \mid \vec{y} \upharpoonright_k = \vec{x}_k \} \end{array}}$$

Channeling interference

In general, any user k who seeks the interferences in a shared channel \vec{f} builds a derived *interference channel* \widehat{f}_k

$$\frac{I^* \xrightarrow{\vec{f}} O^*}{\begin{array}{l} I_k^* \xrightarrow{\widehat{f}_k} \emptyset O \\ \vec{x}_k \mapsto \{\vec{f}_k(\vec{y}) \mid \vec{y} \upharpoonright_k = \vec{x}_k\} \end{array}}$$

On the input \vec{x}_k the interference channel \widehat{f}_k outputs a *possible* output $\vec{f}_k(\vec{y})$, where $\vec{y} \upharpoonright_k = \vec{x}_k$, i.e. \vec{y} is a *possible world* for \vec{x}_k .

Channeling interference

Remark

- ▶ \widehat{f}_k is not a deterministic channel.
- ▶ Nondeterministic channels may be
 - ▶ *possibilistic* $I^+ \rightarrow \wp_* O \subset \{0, 1\}^O$
 - ▶ *probabilistic* $I^+ \rightarrow \Upsilon O \subset [0, 1]^O$
 - ▶ *quantum* $I_+ \rightarrow \Theta O \subset \{z \in \mathbb{C} \mid |z| \leq 1\}^O$

Channeling interference

Remark


- ▶ \widehat{f}_k is not a deterministic channel.
- ▶ Nondeterministic channels may be
 - ▶ *possibilistic* $I^+ \rightarrow \wp_* O \subset \{0, 1\}^O$
 - ▶ *probabilistic* $I^+ \rightarrow \Upsilon O \subset [0, 1]^O$
 - ▶ *quantum* $I_+ \rightarrow \Theta O \subset \{z \in \mathbb{C} \mid |z| \leq 1\}^O$

(We define the possibilistic and the probabilistic versions later, and do not study the quantum channels here.)

Channeling interference

Lemma

A channel $I^* \xrightarrow{\vec{f}} O^*$ satisfies the noninterference requirement for k if and only if the induced interference channel $I_k^+ \xrightarrow{\widehat{f}_k} \wp O$ is deterministic, i.e. emits at most one output for every input.



Covert

Interference

Definition

Examples

Possibilistic

Probabilistic

Quantifying

Lesson

Covert channel

Covert

Interference

Definition

Examples

Possibilistic

Probabilistic

Quantifying

Lesson

Definition

Given a shared channel f , a *covert channel* \hat{f} is derived from f by one or more subjects in order to implement different flows from those specified for f .

Covert channel

Remarks

- ▶ The covert channels in the literature usually extract the *information* about the interference.
- ▶ If channels model any resource use in general, then covert channels model any covert resource use, or abuse.
- ▶ Many familiar information flow attack patterns apply to other resources besides information.
- ▶ Modeling the information flows in a broader context of resource flows seems beneficial both for information security and for resource security.

Example 1

TSA liquid requirement



No more than 3.4oz of liquid carried by passengers.

Covert

Interference

Definition

Examples

Possibilistic

Probabilistic

Quantifying

Lesson

Example 1

TSA checkpoint breach

A group of passengers can form a covert channel by adding

- ▶ a new security level for **bombers**
- ▶ a new state **bomb** and
- ▶ a new transition where the bombers pool their resources

Example 1

TSA checkpoint breach

A group of passengers can form a covert channel by adding

- ▶ a new security level for **bombers**
- ▶ a new state **bomb** and
- ▶ a new transition where the bombers pool their resources

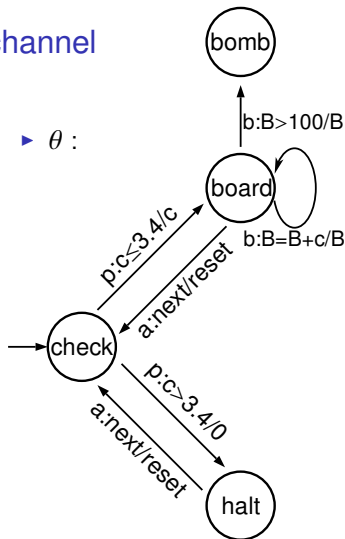
Attack: n subjects with a clearance **b** join their liquids together into a container **B** to get up to $n \times 3.4$ oz of liquid.

Example 1

TSA checkpoint with covert channel

- ▶ $Q = \{\text{check, board, halt, bomb}\}$
- ▶ $L = \{\mathbf{p}:\text{passenger} < \mathbf{a}:\text{agent}, \mathbf{p}:\text{passenger} < \mathbf{b}:\text{bomber}\}$
- ▶ $I_p = \{p:c \leq 3.4, p:c > 3.4\}$
- ▶ $I_a = \{a:\text{next}\}$
- ▶ $I_b = \{b:B = B+c\}$
- ▶ $O = \{c, B, 0, \text{reset}\}$

▶ $\theta :$



Example 2

Fortress gate

- ▶ The fortress wall prevents entry into the city.
- ▶ The fortress gate is an entry channel which
 - ▶ stops soldiers with **weapons**
 - ▶ lets merchants with **merchandise**

Example 2

Fortress gate breach

The attackers form a covert channel by adding

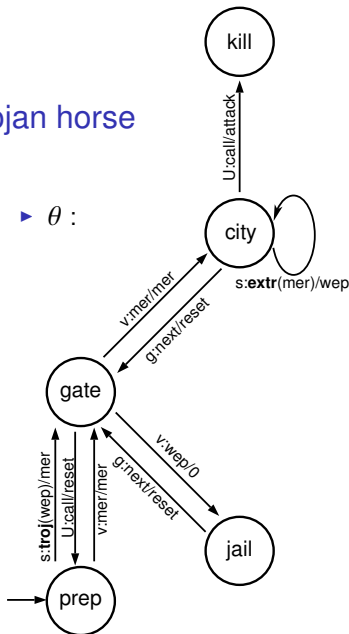
- ▶ new security classes **soldier** and **Ulysses**
- ▶ new actions
 - ▶ **troj**(wep): hide a weapon into a merchandise
 - ▶ **extr**(mer): extract a hidden weapon
 - ▶ **call**: call soldiers to kill
- ▶ new states to
 - ▶ **prep**are for the attack
 - ▶ **kill** the inhabitants
- ▶ new transitions
 - ▶ **prep**→**gate**
 - ▶ **gate**→**prep**
 - ▶ **city**→**kill**

Example 2

Fortress gate breach with Trojan horse


- ▶ $Q = \{\text{gate, city, jail, prep, kill}\}$
- ▶ $\mathbb{L} = \{\text{visitor} < \mathbf{guard}, \text{visitor} < \mathbf{soldier} < \mathbf{Ulysses}\}$
- ▶ $I_v = \{v:\text{mer}, v:\text{wep}\}$
- ▶ $I_g = \{g:\text{next}\}$
- ▶ $I_s = \{s:\text{mer}, s:\mathbf{extr}(\text{mer}), s:\text{wep}, s:\mathbf{troj}(\text{wep})\}$
- ▶ $I_U = \{U:\text{call}\}$
- ▶ $O = \{\text{mer, wep, 0, reset,}$

▶ $\theta :$



Trojan horse

The same attack pattern applies for most channel types



The authentication is often realized through
social engineering.

Covert

Interference

Definition

Examples

Possibilistic

Probabilistic

Quantifying

Lesson

Resource security beyond policies

Covert

Interference

Definition

Examples

Possibilistic

Probabilistic

Quantifying

Lesson

- ▶ Norms and policies are established to assure the behaviors of the *specified* subjects participating in a *specified* process
 - ▶ Access control limits the interactions through *specified* channels.
 - ▶ Noninterference also limits the interactions through *unspecified* channels.

Resource security beyond policies

Covert

Interference

Definition

Examples

Possibilistic

Probabilistic

Quantifying

Lesson

- ▶ But sometimes (in networks) you don't know
 - ▶ who you are sharing a resource with, or
 - ▶ what exactly is the process of sharing

Resource security beyond policies

Covert

Interference

Definition

Examples

Possibilistic

Probabilistic

Quantifying

Lesson

- ▶ But sometimes (in networks) you don't know
 - ▶ who you are sharing a resource with, or
 - ▶ what exactly is the process of sharing
- ▶ The external influences of *unspecified subjects* in *unknown roles* can only be observed as *nondeterminism*:
 - ▶ possibilistic, or
 - ▶ probabilistic

Outline

Covert channels and flows

Possibilistic models

Probabilistic models

Quantifying noninterference

What did we learn?

Recall interference channel

- Shared deterministic flows induce possibilistic channels

$$I^* \xrightarrow{\vec{f}} O^*$$

$$I_k^* \xrightarrow{\widehat{f}_k} \emptyset O$$

$$\vec{x}_k \mapsto \{\vec{f}_k(\vec{y}) \mid \vec{y} \upharpoonright_k = \vec{x}_k\}$$

Recall interference channel

- Shared deterministic flows induce possibilistic channels


$$\frac{I^* \xrightarrow{\vec{f}} O^*}{I_k^* \xrightarrow{\widehat{f}_k} \mathcal{O} \quad \vec{x}_k \mapsto \{\vec{f}_k(\vec{y}) \mid \vec{y}|_k = \vec{x}_k\}}$$

- The interferences at the level k of the deterministic channel Q are observed as the possibility of multiple different outputs on the same local input.
 - A deterministic channel f satisfies the noninterference requirement at the level k if and only if the interference channel \widehat{f}_k is deterministic.

Possibilistic channels

Example: Car rental process

- ▶ $Q = \wp(\text{Cars})$
- ▶ $I_k = \{k:\text{get}, k:\text{ret}\}$, $k \in \mathbb{L} = \text{Customers}$
- ▶ $O = \text{Cars} \cup \text{Invoices} \cup \{\text{Out}\}$
- ▶ θ :



Possibilistic channels

Example: Car rental channel

When a subject k requests a car, the cars that she may *possibly* get depend on the other subjects' requests:

$$\begin{aligned} \{k:\text{get}, k:\text{ret} \mid k \in \mathbb{L}\}^+ &\rightarrow \wp(\text{Cars}) \\ \vec{x} @ k:\text{get} &\mapsto Y_{\vec{x}} \subseteq \text{Cars} \end{aligned}$$

where $Y_{\vec{x}} = \text{Cars} \setminus (\text{gotten out in } \vec{x} \setminus \text{returned back in } \vec{x})$

Possibilistic channels

Example: Car rental channel

When a subject k requests a car, the cars that she may *possibly* get depend on the other subjects' requests:

$$\{k:\text{get}, k:\text{ret} \mid k \in \mathbb{L}\}^+ \rightarrow \wp(\text{Cars})$$

$$\vec{x} @ k:\text{get} \mapsto Y_{\vec{x}} \subseteq \text{Cars}$$

where $Y_{\vec{x}} = \text{Cars} \setminus (\text{gotten out in } \vec{x} \setminus \text{returned back in } \vec{x})$

The interference is unavoidable.

What is a possibilistic channel?

Definition

A *possibilistic channel* with

- ▶ the *inputs* (or *actions*) from A
- ▶ the *outputs* (or *observations*) from B

is a relation

$$f : A^+ \rightarrow \wp B$$

which is prefix closed, in the sense that

$$f(\vec{x}@a) \neq \emptyset \implies f(\vec{x}) \neq \emptyset$$

holds for all $\vec{x} \in A^+$ and $a \in A$.

What is a possibilistic channel?

Notation

For a possibilistic channel $I^+ \xrightarrow{f} \wp O$, we write

$$\vec{x} \vdash_f y \quad \text{when} \quad y \in f(\vec{x})$$

What is a possibilistic channel?

Notation

For a possibilistic channel $I^+ \xrightarrow{f} \wp O$, we write

$$\vec{x} \vdash_f y \quad \text{when} \quad y \in f(\vec{x})$$

When there is just one channel, or f is clear from the context, we elide the subscript and write

$$\vec{x} \vdash y \quad \text{when} \quad y \in f(\vec{x})$$

What is a possibilistic channel?

Definition

A *possibilistic channel* with

- ▶ the *inputs* (or *actions*) from A
- ▶ the *outputs* (or *observations*) from B

is a relation

$$\vdash \subseteq A^+ \times B$$

which is prefix closed, in the sense that

$$\exists z. \vec{x}@a \vdash z \implies \exists y. \vec{x} \vdash y$$

holds for all $\vec{x} \in A^+$ and $a \in A$.

(Possibilistic state machines and processes)

Definition

A *possibilistic state machine* is a map

$$Q \times I \xrightarrow{Nx} \wp(Q \times O)$$

where Q, I, O are finite sets.

(Possibilistic state machines and processes)

Definition

A *possibilistic state machine* is a map

$$Q \times I \xrightarrow{Nx} \wp(Q \times O)$$

where Q, I, O are finite sets.

A *possibilistic process* is a possibilistic state machine with a chosen initial state.

(Possibilistic state machines and processes)

Remark

Possibilistic processes do not in general induce possibilistic channels.

Possibilistic output machines and processes

Definition

A *possibilistic output machine* is a map

$$Q \times I \xrightarrow{\theta} Q \times \wp O$$

where Q, I, O are finite sets.

A *possibilistic output process* is a possibilistic output machine with a chosen initial state.

Possibilistic output machines and processes)

Remark

Possibilistic output processes induce possibilistic channels.

Trace representation

$$\frac{q \in Q \quad Q \times I \xrightarrow{\theta} Q \times \emptyset O}{\frac{I^* \rightarrow \emptyset O}{I^* \times I \xrightarrow{\theta^*} I^* \times \emptyset O}}$$

Memory

- ▶ A possibilistic channel *with no memory* is a binary relation $A \rightarrow \wp B$.

Flows through a possibilistic channel

Definition

The *flow* through a channel $f : A^* \rightarrow \wp B$ is a partial function

$$\vec{f}_\bullet : A^* \rightarrow B^*$$

such that

$$\vec{f}_\bullet() = () \text{ and}$$

$$\vec{f}_\bullet(\vec{x}) \downarrow \wedge \exists b. \vec{x} @ a \vdash_f b \iff \vec{f}_\bullet(\vec{x} @ a) = \vec{f}_\bullet(\vec{x}) @ b$$

holds for all $\vec{x} \in A^*$ and $a \in A$.

Remark

- ▶ Specifying a deterministic channel was equivalent to specifying a deterministic flow.
- ▶ Every possibilistic channel induces many flows.

Possibilistic channels in computation

- ▶ Bob and Charlie using the same network at the same clearance level may enter the same inputs in parallel, and observe several outputs at once.
- ▶ The possible multiple outputs may be observed by entering the same inputs
 - ▶ *sequentially* or
 - ▶ *in parallel*.
- ▶ The actual computations are abstracted away from the channels.

Possibilistic channels in computation

- ▶ Bob enters his inputs into the channel, and observes the interferences with Alice's inputs as the multiple possible outputs.
 - ▶ He observes the interference as the different results of the same local actions.

Possibilistic channels in computation

- ▶ Bob enters his inputs into the channel, and observes the interferences with Alice's inputs as the multiple possible outputs.
 - ▶ He observes the interference as the different results of the same local actions.
- ▶ In network computation, the subjects usually don't even know each other.
 - ▶ The different possibilities are viewed as the *external choices* made by the unobservable environment.

Security consequence

- ▶ A user of a deterministic channel could recognize interference by observing different outputs on the same input:

$$\frac{I^+ \xrightarrow{\vec{f}} O}{I_k^* \xrightarrow{\widehat{f}_k} \emptyset O}$$

Security consequence

- ▶ A user of a deterministic channel could recognize interference by observing different outputs on the same input:

$$\frac{I^+ \xrightarrow{\vec{f}} O}{I_k^* \xrightarrow{\widehat{f}_k} \emptyset O}$$

- ▶ A user of a possibilistic channel can always expect different outputs of the same input:

$$\frac{I^+ \xrightarrow{\vec{f}} \emptyset O}{I_k^* \xrightarrow{\widehat{f}_k} \emptyset O}$$

Security consequence

- ▶ A user of a deterministic channel could recognize interference by observing different outputs on the same input:

$$\frac{I^+ \xrightarrow{\vec{f}} O}{I_k^* \xrightarrow{\widehat{f}_k} \emptyset O}$$

- ▶ A user of a possibilistic channel can always expect different outputs of the same input:

$$\frac{I^+ \xrightarrow{\vec{f}} \emptyset O}{I_k^* \xrightarrow{\widehat{f}_k} \emptyset O}$$

- ▶ The user does not even know who she interferes with
- ▶ The environment makes the "*external choices*"

Security consequence

- ▶ Possibilistic channels arise in nature
- ▶ Possibilistic models are too crude for security.

Outline

Covert channels and flows

Possibilistic models

Probabilistic models

Quantifying noninterference

What did we learn?

Probabilistic channels

Example: Car rental channel

When a subject k requests to rent a car, the cars that she will *probably* get depend on the other subjects' requests, *and* on the habits of the channel

$$\{k:\text{get}, k:\text{ret} \mid k \in \mathbb{L}\}^+ \rightarrow \Upsilon(\text{Cars})$$

$$\vec{x} @ k:\text{get} \mapsto Y_{\vec{x}}$$

where $Y_{\vec{x}}$ is a random selection from


$$\text{Cars} \setminus (\text{Taken in } \vec{x} \setminus \text{Returned in } \vec{x})$$

.

Probabilistic channels

Example: Car rental process

- ▶ $Q = \emptyset(\text{Cars})$
- ▶ $I_k = \{k:\text{get}, k:\text{ret}\}$, $k \in \mathbb{L} = \text{Customers}$
- ▶ $O = \text{Cars} \cup \text{Invoices} \cup \{\text{Out}\}$
- ▶ θ :



What is a probabilistic channel?

Definitions we'll need

A *partial random element* X over a countable set A is given by a subprobability distribution ν_X over A , i.e. a function

$$\nu_X : A \rightarrow [0, 1]$$

such that $\sum_{x \in A} \nu(x) \leq 1$.

What is a probabilistic channel?

Definitions we'll need

A *partial random element* X over a countable set A is given by a subprobability distribution ν_X over A , i.e. a function

$$\nu_X : A \rightarrow [0, 1]$$

such that $\sum_{x \in A} \nu(x) \leq 1$.

We usually write

$$\nu_X(x) = \nu(X = x)$$

What is a probabilistic channel?

Definitions we'll need

The set of all partial random elements over the set X is

$$\mathcal{r}A = \left\{ \nu(X = -) : A \rightarrow [0, 1] \mid \sum_{x \in A} \nu(X = x) \leq 1 \right\}$$

What is a probabilistic channel?

Definitions we'll need

A *partial random function* is a function $f : A \rightarrow \Upsilon B$.

What is a probabilistic channel?

Definition

A *probabilistic channel* with

- ▶ the *inputs* (or *actions*) from A
- ▶ the *outputs* (or *observations*) from B

is partial random function

$$f : A^+ \rightarrow \Upsilon B$$

which is prefix closed, in the sense that

$$\sum_{z \in B} \nu(f(\vec{x}@a) = z) \leq \sum_{y \in B} \nu(f(\vec{x}) = y)$$

for all $\vec{x} \in A^+$ and $a \in A$.

What is a probabilistic channel?

Notation

For a probabilistic channel $I^+ \xrightarrow{f} \Upsilon O$, we write

$$[\vec{x} \vdash_f y] = v(f(\vec{x}) = y)$$

What is a probabilistic channel?

Notation

For a probabilistic channel $I^+ \xrightarrow{f} \Upsilon O$, we write

$$[\vec{x} \vdash_f y] = v(f(\vec{x}) = y)$$

When there is just one channel, or f is clear from the context, we elide the subscript and write

$$[\vec{x} \vdash y] = v(f(\vec{x}) = y)$$

What is a probabilistic channel?

Notation

For a probabilistic channel $I^+ \xrightarrow{f} \Upsilon O$, we write $[\vec{x} \vdash Y]$ and view Y as the source where

$$v(Y = y) = v(f(\vec{x}) = y)$$

for the given history $\vec{x} \in I^+$

What is a probabilistic channel?

Definition

A *probabilistic channel* with

- ▶ the *inputs* (or *actions*) from A
- ▶ the *outputs* (or *observations*) from B

is a partial random element

$$[- \vdash -] \in \Upsilon(A^+ \times B)$$

which is prefix closed, in the sense that

$$\sum_{z \in B} [\vec{x} @ a \vdash z] \leq \sum_{y \in B} [\vec{x} \vdash y]$$

holds for all $\vec{x} \in A^+$ and $a \in A$.

Memory

- ▶ A probabilistic channel *with no memory* is a partial random function $A \rightarrow \Upsilon B$.

Information theoretic channel

Any probabilistic channel can be extended

$$\frac{I^+ \xrightarrow{f} \mathcal{O}}{\begin{array}{ccc} \Upsilon(I^+) & \xrightarrow{\vec{f}} & \mathcal{O} \\ \vec{X} & \mapsto & Y \end{array}}$$

where

$$v(Y = y) = \sum_{\vec{x} \in I^+} v(\vec{X} = \vec{x}) \cdot v(f(\vec{x}) = y)$$

Information theoretic channel

Notation

The extensions align with the usual information theoretic channel notation

$$[X_1, X_2, \dots, X_n \vdash Y] = v(\bar{f}(X_1, X_2, \dots, X_n) = Y)$$

Probabilistic interference channel

Shared channels induce interference channels

$$\frac{I^+ \xrightarrow{[\vdash]} \gamma O}{I_k^+ \xrightarrow{[\vdash]_k} \gamma O}$$

where

$$[\vec{x}_k \vdash y]_k = \sum_{\vec{x} \in I^+} \nu(\vec{x}_k = \vec{x} \upharpoonright_k) \cdot [\vec{x} \vdash y]$$

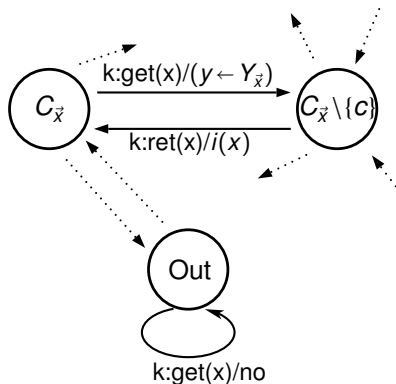
Probabilistic interference channel

Probabilistic interference is exploited through Bayesian inference.

Has Alice rented a car?

Example: Car rental process

- ▶ $Q = \emptyset(\text{Cars})$, $\text{Cars} = \{9 \text{ toyotas}, 1 \text{ porsche}\}$
- ▶ $I_k = \{k:\text{get}(x), k:\text{ret}(x)\}$, $k \in \{\text{Alice}, \text{Bob}\} \cup \text{Others}$, $x \in \text{Cars}$
- ▶ $O = \text{Cars} \cup \text{Invoices} \cup \{\text{Out}\}$
- ▶ θ :



Has Alice rented a car?

Covert channel

- ▶ Bob wonders whether Alice is in town.
 - ▶ She always rents a car.
- ▶ Bob knows that Alice likes to rent the porsche.
 - ▶ She does not get it one in 5 times.
- ▶ Bob requests a rental and gets the porsche.
 - ▶ How likely is it that Alice is in town?

Has Alice rented a car?

Bob considers the following events

a: Alice has rented a car

- ▶ Alice:get(car) occurs in \vec{x}

m: The porsche is available

- ▶ Bob:get(porsche) results in $\text{porsche} \leftarrow Y_{\vec{x}}$

Has Alice rented a car?

Bob's beliefs

- ▶ $v(m | a) = \frac{1}{5}$
 - ▶ If Alice is in town, then the chance that the porsche is available is $\frac{1}{5}$.
- ▶ $v(m | \neg a) = \frac{9}{10}$
 - ▶ If Alice is not in town, then the chance that the porsche is available is $\frac{9}{10}$.
- ▶ $v(a) = \frac{1}{2}$
 - ▶ A priori, the chance that Alice is in town is 50-50.

Has Alice rented a car?

Bob's reasoning

$$\blacktriangleright v(a | m) = \frac{v(a,m)}{v(m)}$$

Has Alice rented a car?

Bob's reasoning

- ▶ $v(a | m) = \frac{v(a,m)}{v(m)}$
 - ▶ $v(a, m) = v(m|a) \cdot v(a) = \frac{1}{5} \cdot \frac{1}{2} = \frac{1}{10}$
 - ▶ $v(m) = v(a, m) + v(\neg a, m)$

Has Alice rented a car?

Bob's reasoning

- ▶ $v(a | m) = \frac{v(a,m)}{v(m)}$
 - ▶ $v(a, m) = v(m|a) \cdot v(a) = \frac{1}{5} \cdot \frac{1}{2} = \frac{1}{10}$
 - ▶ $v(m) = v(a, m) + v(\neg a, m)$
 - ▶ $v(m, \neg a) = v(m|\neg a) \cdot v(\neg a) = \frac{9}{10} \cdot \frac{1}{2} = \frac{9}{20}$

Has Alice rented a car?

Bob's reasoning

- ▶ $v(a | m) = \frac{v(a,m)}{v(m)}$
 - ▶ $v(a, m) = v(m|a) \cdot v(a) = \frac{1}{5} \cdot \frac{1}{2} = \frac{1}{10}$
 - ▶ $v(m) = v(a, m) + v(\neg a, m)$
 - ▶ $v(m, \neg a) = v(m|\neg a) \cdot v(\neg a) = \frac{9}{10} \cdot \frac{1}{2} = \frac{9}{20}$
 - ▶ $v(m) = \frac{1}{10} + \frac{9}{20} = \frac{11}{20}$

Has Alice rented a car?

Bob's reasoning

- ▶ $v(a | m) = \frac{v(a, m)}{v(m)}$
 - ▶ $v(a, m) = v(m|a) \cdot v(a) = \frac{1}{5} \cdot \frac{1}{2} = \frac{1}{10}$
 - ▶ $v(m) = v(a, m) + v(\neg a, m)$
 - ▶ $v(m, \neg a) = v(m|\neg a) \cdot v(\neg a) = \frac{9}{10} \cdot \frac{1}{2} = \frac{9}{20}$
 - ▶ $v(m) = \frac{1}{10} + \frac{9}{20} = \frac{11}{20}$
- ▶ $v(a | m) = \frac{\frac{1}{10}}{\frac{11}{20}} = \frac{2}{11}$

Has Alice rented a car?

Bob's reasoning

- ▶ $v(a | m) = \frac{v(a, m)}{v(m)}$
 - ▶ $v(a, m) = v(m|a) \cdot v(a) = \frac{1}{5} \cdot \frac{1}{2} = \frac{1}{10}$
 - ▶ $v(m) = v(a, m) + v(\neg a, m)$
 - ▶ $v(m, \neg a) = v(m|\neg a) \cdot v(\neg a) = \frac{9}{10} \cdot \frac{1}{2} = \frac{9}{20}$
 - ▶ $v(m) = \frac{1}{10} + \frac{9}{20} = \frac{11}{20}$
- ▶ $v(a | m) = \frac{\frac{1}{10}}{\frac{11}{20}} = \frac{2}{11}$

If the porsche is available, then the chance that Alice is in town is 2 in 11.

Has Alice rented a car?

Bob's learning

- ▶ Bob's *input information* (or *prior belief*) before renting the car was that the chance that Alice was in town was $\frac{1}{2}$.
- ▶ Bob's *channel information* (or *posterior belief*) after renting the car was that the chance that Alice was in town was $\frac{2}{11}$.

Has Alice rented a car?

Quantifying noninterference

- ▶ A channel satisfies the k -noninterference requirement if k learns nothing from using it:

$$\text{channel information} = \text{input information}$$

Has Alice rented a car?

Quantifying noninterference

- ▶ A channel satisfies the k -noninterference requirement if k learns nothing from using it:

$$\text{posterior belief} = \text{prior belief}$$

Has Alice rented a car?

Quantifying noninterference

- ▶ A channel satisfies the k -noninterference requirement if k learns nothing from using it:

$$\text{posterior belief} = \text{prior belief}$$

- ▶ The degree of the channel noninterference is

$$\frac{\text{posterior belief}}{\text{prior belief}} \leq 1 \quad \text{or} \quad \frac{\text{prior belief}}{\text{posterior belief}} \leq 1$$

Has Alice rented a car?

Quantifying noninterference

- ▶ A channel satisfies the k -noninterference requirement if k learns nothing from using it:

$$\text{posterior belief} = \text{prior belief}$$

- ▶ The degree of the channel noninterference is

for the rental channel: $\frac{\frac{2}{11}}{\frac{1}{2}} = \frac{4}{11}$

Outline

Covert channels and flows

Possibilistic models

Probabilistic models

Quantifying noninterference

What did we learn?

Recall noninterference

Definition

A shared deterministic channel $I^+ \xrightarrow{f} O$ satisfies the *noninterference* requirement at the level k if for all states of the world $\vec{x}, \vec{y} \in I^*$ holds

$$\vec{x} \downarrow_k \vec{y} \implies \vec{x} \uparrow_{f_k} \vec{y}$$

where

$$\vec{x} \downarrow_k \vec{y} \iff \vec{x} \upharpoonright_k = \vec{y} \upharpoonright_k$$

$$\vec{x} \uparrow_{f_k} \vec{y} \iff f_k(\vec{x}) = f_k(\vec{y})$$

Recall noninterference

Definition

A shared deterministic channel $I^+ \xrightarrow{f} O$ satisfies the *noninterference* requirement at the level k if for all states of the world $\vec{x}, \vec{y} \in I^*$ holds

$$\vec{x} \upharpoonright_k \vec{y} \implies \vec{x} \upharpoonright_{f_k} \vec{y}$$

where

$$\vec{x} \upharpoonright_k \vec{y} \iff \vec{x} \upharpoonright_k = \vec{y} \upharpoonright_k \quad \leftarrow \text{input view}$$

$$\vec{x} \upharpoonright_{f_k} \vec{y} \iff f_k(\vec{x}) = f_k(\vec{y}) \quad \leftarrow \text{channel view}$$

Quantified noninterference

Definition

A shared probabilistic channel $I^+ \xrightarrow{f} \Upsilon O$ satisfies the *noninterference* requirement at the level k if for all states of the world $\vec{x}, \vec{y} \in I^*$ holds

$$\vec{x}[k]\vec{y} \leq \vec{x}[f_k]\vec{y}$$

where

$$\vec{x}[k]\vec{y} = \bigwedge_{\vec{x}_k \in I_k^+} \frac{v(\vec{x} \upharpoonright_k = \vec{x}_k)}{v(\vec{y} \upharpoonright_k = \vec{x}_k)}$$

$$\vec{x}[f_k]\vec{y} = \bigwedge_{z \in O} \frac{v(f_k(\vec{x}) = z)}{v(f_k(\vec{y}) = z)}$$

Quantified interference

Definition

The amount of interference that a user at the level k of the shared probabilistic channel $I^+ \xrightarrow{f} \Upsilon O$ can extract to distinguish the histories $\vec{x}, \vec{y} \in I^+$ is

$$\begin{aligned} \iota(\vec{x}, \vec{y}) &= -\log \left| \frac{\vec{x} \lfloor k \rfloor \vec{y}}{\vec{x} \lceil f_k \rceil \vec{y}} \right| \\ &= \left| \log(\vec{x} \lfloor k \rfloor \vec{y}) - \log(\vec{x} \lceil f_k \rceil \vec{y}) \right| \end{aligned}$$

where...

- ▶ The *normalized ratio* is defined

$$\frac{x}{y} = \begin{cases} \frac{x}{y} & \text{if } x \leq y \\ \frac{y}{x} & \text{if } x > y \end{cases}$$

- ▶ The *normalized ratio* is defined

$$\left| \frac{x}{y} \right| = \begin{cases} \frac{x}{y} & \text{if } x \leq y \\ \frac{y}{x} & \text{if } x > y \end{cases}$$

- ▶ It is the multiplicative version of the more familiar *absolute difference*

$$|x - y| = \begin{cases} y - x & \text{if } x \leq y \\ x - y & \text{if } x > y \end{cases}$$

Connection

- ▶ from absolute value to normalized ratio

$$\left| \frac{x}{y} \right| = 2^{|\log x - \log y|}$$

- ▶ from normalized ratio to absolute value

$$|x - y| = \log \left| \frac{2^x}{2^y} \right|$$

Question

- ▶ But why is this the right way to quantify noninterference?

Question

- ▶ But why is this the right way to quantify noninterference?
- ▶ In which sense do the numbers $\vec{x}[k]\vec{y}$ and $\vec{x}[f_k]\vec{y}$ quantify and generalize the relations $\vec{x}[k]\vec{y}$ and $\vec{x}[f_k]\vec{y}$

Quantified equivalences

Recall partial equivalence relations

An equivalence relation over a set A is a function

$$A \times A \xrightarrow{R} \{0, 1\}$$

such that

$$xRy = yRx \qquad xRy \wedge yRz \leq xRz$$

Quantified equivalences

Equivalence kernel

An *equivalence kernel* over a set A is a function

$$A \times A \xrightarrow{R} [0, 1]$$

such that

$$xRy = yRx \qquad xRy \cdot yRz \leq xRz$$

Quantified equivalences

Equivalence kernel over ΥA

Recall the set of partial random elements over A

$$\Upsilon A = \left\{ \nu(X = -) : A \rightarrow [0, 1] \mid \sum_{x \in A} \nu(X = x) \leq 1 \right\}$$

Quantified equivalences

Equivalence kernel over ΥA

Recall the set of partial random elements over A

$$\Upsilon A = \left\{ \nu(X = -) : A \rightarrow [0, 1] \mid \sum_{x \in A} \nu(X = x) \leq 1 \right\}$$

It comes equipped with the canonical equivalence kernel, defined

$$[X \sim Y] = \bigwedge_{a \in A} \left| \frac{\nu(X = a)}{\nu(Y = a)} \right|$$

Quantified equivalences

Exercise

Show that $[X \sim Y]$ is an equivalence kernel, i.e. that it satisfies the quantified symmetry and transitivity, as defined 3 slides ago.

Quantified equivalences

Input view is an equivalence kernel

k 's prior belief tells how likely is each $\vec{x}_k \in I_k^+$ to be the local view of any $\vec{y} \in I^+$, which is given by a partial random element

$$v(\vec{x}_k = \vec{x} \upharpoonright_k) : I^+ \rightarrow [0, 1]$$

Quantified equivalences

Input view is an equivalence kernel

k 's prior belief tells how likely is each $\vec{x}_k \in I_k^+$ to be the local view of any $\vec{y} \in I^+$, which is given by a partial random element

$$v(\vec{x}_k = \vec{x} \upharpoonright_k) : I^+ \rightarrow [0, 1]$$

Rearranging k 's beliefs into partial random elements over I_k^+

$$v(\vec{x} \upharpoonright_k = \vec{x}_k) : I_k^+ \rightarrow [0, 1]$$

we define the input view

$$\vec{x}[k]\vec{y} = \bigwedge_{\vec{x}_k \in I_k^+} \frac{v(\vec{x} \upharpoonright_k = \vec{x}_k)}{v(\vec{y} \upharpoonright_k = \vec{x}_k)}$$

Quantified equivalences

Remark

Note that for every $\vec{x}_k \in I^+$ and every $\vec{y} \in I^+$ holds

$$\vec{x}_k [k] \vec{y} = v(\vec{x}_k = \vec{y} \upharpoonright_k)$$

Quotients

Recall that every partial function $A \xrightarrow{f} B$ induces the partial equivalence relation on A

$$x(f)y \iff f(x) = f(y)$$

Recall that every partial function $A \xrightarrow{f} B$ induces the partial equivalence relation on A

$$x(f)y \iff f(x) = f(y)$$

Analogously, every partial stochastic function $A \xrightarrow{f} \Upsilon B$ induces the equivalence kernel

$$x(f)y = \bigwedge_{b \in B} \left| \frac{v(f(x) = b)}{v(f(y) = b)} \right|$$

Hence

$$\vec{x} [f_k] \vec{y} = \bigwedge_{z \in \mathcal{O}} \frac{v(f_k(\vec{x}) = z)}{v(f_k(\vec{y}) = z)}$$

... and hence noninterference

Definition

A shared probabilistic channel $I^+ \xrightarrow{f} \Upsilon O$ satisfies the *noninterference* requirement at the level k if for all states of the world $\vec{x}, \vec{y} \in I^*$ holds

$$\vec{x}[k]\vec{y} \leq \vec{x}[f_k]\vec{y}$$

where

$$\vec{x}[k]\vec{y} = \bigwedge_{\vec{x}_k \in I_k^+} \frac{v(\vec{x} \upharpoonright_k = \vec{x}_k)}{v(\vec{y} \upharpoonright_k = \vec{x}_k)}$$

$$\vec{x}[f_k]\vec{y} = \bigwedge_{z \in O} \frac{v(f_k(\vec{x}) = z)}{v(f_k(\vec{y}) = z)}$$

... and quantified interference

Definition


The amount of interference that a user at the level k of the shared probabilistic channel $I^+ \xrightarrow{f} \Upsilon O$ can extract to distinguish the histories $\vec{x}, \vec{y} \in I^+$ is

$$\begin{aligned} \iota(\vec{x}, \vec{y}) &= -\log \left| \frac{\vec{x}[k]\vec{y}}{\vec{x}[f_k]\vec{y}} \right| \\ &= \left| \log(\vec{x}[k]\vec{y}) - \log(\vec{x}[f_k]\vec{y}) \right| \end{aligned}$$

(An aside about the partitions)

The partition induced by the kernel of any function $A \xrightarrow{f} B$ or relation $A \xrightarrow{f} \wp B$ are obtained as the image of the composite with its inverse image

$$\begin{aligned} \wp B &\xrightarrow{f^*} \wp A \\ V &\mapsto \bigcup \{U \subseteq A \mid f(U) \subseteq V\} \end{aligned}$$



(An aside about the partitions)

The same construction lifts to *stochastic* functions, which are the partial random functions $A \xrightarrow{f} \Upsilon B$ such that for every $b \in B$ holds

$$f_{\bullet}(b) = \sum_{a \in A} f_a(b) < \infty$$

(An aside about the partitions)

The same construction lifts to *stochastic* functions, which are the partial random functions $A \xrightarrow{f} \Upsilon B$ such that for every $b \in B$ holds

$$f_{\bullet}(b) = \sum_{a \in A} f_a(b) < \infty$$

Hence


$$\frac{A \xrightarrow{f} \Upsilon B}{B \xrightarrow{\tilde{f}} \Upsilon A}$$

$$b \mapsto \frac{1}{f_{\bullet}(b)} \cdot \lambda a. f_a(b)$$

(An aside about the partitions)

The partition induced by the kernel of any stochastic function $A \xrightarrow{f} \Upsilon B$ are obtained as the image of the composite with its inverse image

$$\begin{aligned} \Upsilon B &\xrightarrow{f^*} \Upsilon A \\ \beta &\mapsto \sum_{b \in B} \beta(b) \cdot \tilde{f}_b \end{aligned}$$



Outline

Covert channels and flows

Possibilistic models

Probabilistic models

Quantifying noninterference

What did we learn?

What did we learn?

- ▶ Interference is exploited through a special family of *covert channels*.
- ▶ Other failures of channel security are realized through other types of covert channels.
- ▶ The external interferences¹ on the functioning of a channel manifest themselves through *many possible* outputs on the same input.
 - ▶ Hence *possibilistic processes*.
- ▶ Gathering information about the external interferences requires *quantifying* the *probabilities* of the various possible inputs.
 - ▶ *Possibilistic processes* allow quantifying interference.

¹by the environment, or by unobservable subject 

Statistical disclosure is a probabilistic channel

- ▶ Statistical disclosure outputs data from a family of databases randomized as to preserve privacy and anonymity.

Statistical disclosure is a probabilistic channel

- ▶ Statistical disclosure outputs data from a family of databases randomized as to preserve privacy and anonymity.
- ▶ A randomization method of statistical disclosure can be viewed as a shared probabilistic channel.

Differential privacy is a bound on interference

ICS 355:
Introduction

Dusko Pavlovic

Covert

Possibilistic

Probabilistic

Quantifying

Lesson

- ▶ Security of statistical disclosure is a difficult problem, recently solved in terms of *differential privacy*.

Differential privacy is a bound on interference

- ▶ Security of statistical disclosure is a difficult problem, recently solved in terms of *differential privacy*.
- ▶ Differential privacy turns out to be a method for **limiting the amount of interference**, as defined above.

Huh?

- ▶ But what is differential privacy?

Huh?

- ▶ But what is differential privacy?
- ▶ We first need to define privacy, don't we?