Metrics for Differential Privacy in Concurrent Systems

Lili Xu[1,3,4] Konstantinos Chatzikokolakis[2,3]
Huimin Lin[4] Catuscia Palamidessi[1,3]

[1]INRIA [2]CNRS [3]Ecole Polytechnique
[4]Institute of Software, Chinese Academy of Sciences

HotSpot 2014

Outline

1. Introduction
   - Concurrent Systems
   - Differential Privacy
   - The Verification Framework

2. Three Pseudometrics
   - The Accumulative Bijection Pseudometric
   - The Amortised Bijection Pseudometric
   - A Multiplicative Variant of the Kantorovich Pseudometric
   - Comparison

Motivation

- The model: Concurrent systems modeled as probabilistic automata.
- The measure of the level of privacy: Differential privacy

Motivation

- The model: Concurrent systems modeled as probabilistic automata.
- The measure of the level of privacy: Differential privacy

Goal:

To verify differential privacy properties for concurrent systems

Outline

1. Introduction
   - Concurrent Systems
   - Differential Privacy
   - The Verification Framework

2. Three Pseudometrics
   - The Accumulative Bijection Pseudometric
   - The Amortised Bijection Pseudometric
   - A Multiplicative Variant of the Kantorovich Pseudometric
   - Comparison

Our Model

A probabilistic automaton is a tuple $(S, \bar{s}, A, D)$

- S: a finite set of states;
- $\bar{s} \in S$: the start state;
- A: a finite set of action labels;
- $D \subseteq S \times A \times Disc(S)$: a transition relation. We also write $s \xrightarrow{a} \mu$.

Definition (Concurrent Systems with Secret Information)

Let U be a set of secrets. A concurrent system with secret information $\mathcal{A}$ is a mapping of secrets to probabilistic automata, where $\mathcal{A}(u), u \in U$ is the automaton modelling the behavior of the system when running on u.

How to Reason about Probabilistic Observations?

- A scheduler ζ resolves the non-determinism based on the history of a computation, inducing a probability measure over traces.

How to Reason about Probabilistic Observations?

- A scheduler ζ resolves the non-determinism based on the history of a computation, inducing a probability measure over traces.

Probabilities of finite traces

Let α be the history up to the current state s. The probability of observing a finite trace $\vec{t}$ starting from α, denoted by $\Pr_\zeta[\alpha \rhd \vec{t}\,]$, is defined recursively as follows.

$$\Pr_\zeta[\alpha \rhd \vec{t}\,] = \begin{cases} 1 & \text{if } \vec{t} \text{ is empty,} \\ 0 & \text{if } \vec{t} = a^\frown \vec{t}\,', \zeta(\alpha) = s \xrightarrow{b} \mu \text{ and } b \neq a, \\ \sum_{s_i} \mu(s_i) \Pr_\zeta[\alpha a s_i \rhd \vec{t}\,'] & \text{if } \vec{t} = a^\frown \vec{t}\,' \text{ and } \zeta(\alpha) = s \xrightarrow{a} \mu. \end{cases}$$

An example: A PIN-Checking System

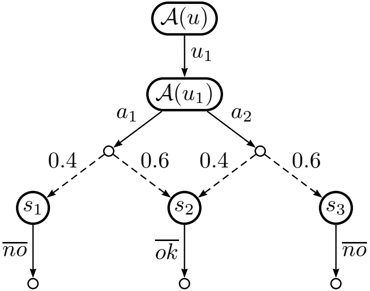

Example: The scheduler executes the a_1-branch.

$$\begin{aligned}
Pr_\varsigma[\mathcal{A}(u_1) \triangleright a_1\overline{ok}] &= 0.6 & \quad Pr_\varsigma[\mathcal{A}(u_2) \triangleright a_1\overline{ok}] &= 0.4\\
Pr_\varsigma[\mathcal{A}(u_1) \triangleright a_1\overline{no}] &= 0.4 & \quad Pr_\varsigma[\mathcal{A}(u_2) \triangleright a_1\overline{no}] &= 0.6\\
Pr_\varsigma[\mathcal{A}(u_1) \triangleright a_2\overline{ok}] &= 0 & \quad Pr_\varsigma[\mathcal{A}(u_2) \triangleright a_2\overline{ok}] &= 0\\
Pr_\varsigma[\mathcal{A}(u_1) \triangleright a_2\overline{no}] &= 0 & \quad Pr_\varsigma[\mathcal{A}(u_2) \triangleright a_2\overline{no}] &= 0
\end{aligned}$$

Outline

1. Introduction
   - Concurrent Systems
   - Differential Privacy
   - The Verification Framework
2. Three Pseudometrics
   - The Accumulative Bijection Pseudometric
   - The Amortised Bijection Pseudometric
   - A Multiplicative Variant of the Kantorovich Pseudometric
   - Comparison

How To Quantify the Amount of Privacy?

Definition (Standard Definition of Differential Privacy)

A query mechanism $\mathcal{A}$ is ϵ-differentially private if for any two adjacent databases u_1 and u_2, i.e. which differ only for one individual, and any property Z, the probability distributions of $\mathcal{A}(u_1), \mathcal{A}(u_2)$ differ on Z at most by e^ϵ, namely,

$$\Pr[\mathcal{A}(u_1) \in Z] \leq e^\epsilon \cdot \Pr[\mathcal{A}(u_2) \in Z].$$

The lower the value ϵ is, the better the privacy is protected.

How To Quantify the Amount of Privacy?

Definition (Standard Definition of Differential Privacy)

A query mechanism $\mathcal{A}$ is ϵ-differentially private if for any two adjacent databases u_1 and u_2, i.e. which differ only for one individual, and any property Z, the probability distributions of $\mathcal{A}(u_1), \mathcal{A}(u_2)$ differ on Z at most by e^ϵ, namely,

$$\Pr[\mathcal{A}(u_1) \in Z] \leq e^\epsilon \cdot \Pr[\mathcal{A}(u_2) \in Z].$$

The lower the value ϵ is, the better the privacy is protected.

Some Merits of Differential Privacy

- Strong notion of privacy.
- Independence from side knowledge.
- Robustness to attacks based on combining various sources of information.
- Looser restrictions between non-adjacent secrets.

Differential Privacy in the Context of Concurrent Systems

- The scheduler can easily break many security and privacy properties.
- We consider a restricted class of schedulers, called admissible schedulers.
  - make them unable to distinguish between secrets in the histories.

Definition (Differential Privacy in Our Setting)

A concurrent system $\mathcal{A}$ satisfies ϵ-*differential privacy* (DP) iff for any two adjacent secrets u, u', all finite traces $\vec{t}$ and all admissible schedulers ζ:

$$\Pr_{\zeta}[\mathcal{A}(u) \triangleright \vec{t}\,] \leq e^{\epsilon} \cdot \Pr_{\zeta}[\mathcal{A}(u') \triangleright \vec{t}\,]$$

The PIN-Checking System Revisited

Definition (Differential Privacy in Our Setting)

A concurrent system $\mathcal{A}$ satisfies ϵ-*differential privacy* (DP) iff for any two adjacent secrets u, u', all finite traces $\vec{t}$ and all admissible schedulers ζ:

$$\Pr_{\zeta}[\mathcal{A}(u) \triangleright \vec{t}\,] \leq e^{\epsilon} \cdot \Pr_{\zeta}[\mathcal{A}(u') \triangleright \vec{t}\,]$$

Example

$$\begin{aligned}
\Pr_{\zeta}[\mathcal{A}(u_1) \triangleright a_1\overline{ok}\,] &= 0.6 & \qquad \Pr_{\zeta}[\mathcal{A}(u_2) \triangleright a_1\overline{ok}\,] &= 0.4\\
\Pr_{\zeta}[\mathcal{A}(u_1) \triangleright a_1\overline{no}\,] &= 0.4 & \qquad \Pr_{\zeta}[\mathcal{A}(u_2) \triangleright a_1\overline{no}\,] &= 0.6\\
\Pr_{\zeta}[\mathcal{A}(u_1) \triangleright a_2\overline{ok}\,] &= 0 & \qquad \Pr_{\zeta}[\mathcal{A}(u_2) \triangleright a_2\overline{ok}\,] &= 0\\
\Pr_{\zeta}[\mathcal{A}(u_1) \triangleright a_2\overline{no}\,] &= 0 & \qquad \Pr_{\zeta}[\mathcal{A}(u_2) \triangleright a_2\overline{no}\,] &= 0
\end{aligned}$$

In this case, the level of differential privacy $\epsilon = \ln \frac{3}{2}$.

Outline

1. Introduction
   - Concurrent Systems
   - Differential Privacy
   - The Verification Framework
2. Three Pseudometrics
   - The Accumulative Bijection Pseudometric
   - The Amortised Bijection Pseudometric
   - A Multiplicative Variant of the Kantorovich Pseudometric
   - Comparison

Neighboring processes have neighboring behaviors.

- For example: behavioural equivalences
  - $\mathcal{A}(u) \simeq \mathcal{A}(u') \Longrightarrow$ Secrecy [Abadi and Gordon, the Spi-calculus]

The property of differential privacy requires that the observations generated by two adjacent secrets are probabilistically close.

Neighboring processes have neighboring behaviors.

- For example: behavioural equivalences
  - $\mathcal{A}(u) \simeq \mathcal{A}(u') \Longrightarrow$ Secrecy [Abadi and Gordon, the Spi-calculus]

The property of differential privacy requires that the observations generated by two adjacent secrets are probabilistically close.

Verification Technique

- Behavioural approximation:Pseudometrics on processes.
- Find a pseudometric m on states of a concurrent system for two adjacent secrets u, u', such that:

$$m(\mathcal{A}(u), \mathcal{A}(u')) \leq \epsilon \Longrightarrow \mathcal{A}(u) \text{ and } \mathcal{A}(u') \text{ are } \epsilon\text{-differentially private.}$$

Outline

The Accumulative Bijection Pseudometric

It stems from the work of

- Michael C. Tschantz, Dilsun Kaynar, and Anupam Datta. Formal verification of differential privacy for interactive systems. ENTCS 2011.

We reformulate the notion of approximate similarity proposed in the above work in terms of a pseudometric, and exhibit its properties as a distance relation.

Definitions

We define an approximate bisimulation relation:

Definition (Accumulative Bisimulation)

A relation $\mathcal{R} \subseteq S \times S \times [0, \epsilon]$ is an ϵ-accumulative bisimulation iff for all $(s, t, c) \in \mathcal{R}$:

- $s \xrightarrow{a} \mu$ implies $t \xrightarrow{a} \mu'$ with $\mu \mathcal{L}^D(\mathcal{R}, c)\mu'$
- $t \xrightarrow{a} \mu'$ implies $s \xrightarrow{a} \mu$ with $\mu \mathcal{L}^D(\mathcal{R}, c)\mu'$

Definitions

First, lift a relation over states to a relation over distributions.

Definition (D-Approximate Lifting)

$\mu\mathcal{L}^D(\mathcal{R},c)\mu'$ iff $\exists$ bijection $\beta : supp(\mu) \to supp(\mu')$ such that

$\forall s \in supp(\mu) : (s, \beta(s), c+\sigma) \in \mathcal{R}$ where $\sigma = \max_{s \in supp(\mu)} |\ln \frac{\mu(s)}{\mu'(\beta(s))}|$

We define an approximate bisimulation relation:

Definition (Accumulative Bisimulation)

A relation $\mathcal{R} \subseteq S \times S \times [0,\epsilon]$ is an ϵ-accumulative bisimulation iff for all $(s,t,c) \in \mathcal{R}$:

- $s \xrightarrow{a} \mu$ implies $t \xrightarrow{a} \mu'$ with $\mu\mathcal{L}^D(\mathcal{R},c)\mu'$
- $t \xrightarrow{a} \mu'$ implies $s \xrightarrow{a} \mu$ with $\mu\mathcal{L}^D(\mathcal{R},c)\mu'$

We can now define a pseudometric based on accumulative bisimulation as:

$$m^D(s,t) = \min\{\epsilon \mid (s,t,0) \in \mathcal{R} \text{ for some } \epsilon\text{-accumulative bisimulation } \mathcal{R}\}$$

Proposition

m^D is a pseudometric, that is:

- *(reflexivity)* $m^D(s,s) = 0$
- *(symmetry)* $m^D(s_1,s_2) = m^D(s_2,s_1)$
- *(triangle inequality)* $m^D(s_1,s_3) \leq m^D(s_1,s_2) + m^D(s_2,s_3)$

Verification of differential privacy using m^D

Theorem

A concurrent system $\mathcal{A}$ is ϵ-differentially private if $m^D(\mathcal{A}(u), \mathcal{A}(u')) \leq \epsilon$ for any two adjacent secrets u and u'.

The PIN-Checking System Revisited

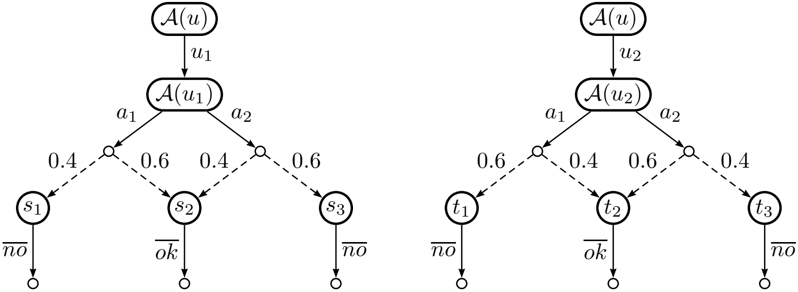

Example

The following relation is a $\ln \frac{3}{2}$-accumulative bisimulation between $\mathcal{A}(u_1)$ and $\mathcal{A}(u_2)$.

$$\begin{aligned} \mathcal{R} = \{ & (\mathcal{A}(u_1), \mathcal{A}(u_2), 0), & (s_1, t_1, \ln \tfrac{3}{2}) \\ & (s_2, t_2, \ln \tfrac{3}{2}), & (s_3, t_3, \ln \tfrac{3}{2}) \} \end{aligned}$$

Thus $m^D(\mathcal{A}(u_1), \mathcal{A}(u_2)) = \ln \frac{3}{2}$, system $\mathcal{A}$ is $\ln \frac{3}{2}$-differentially private.

The Use of the Privacy Budget May Be a bit Wasteful?

m^D is useful for verifying differential privacy. However,

- the amount of leakage is only accumulated.
- the accumulation is the same for all branches, and equal to the worst branch.

The Use of the Privacy Budget May Be a bit Wasteful?

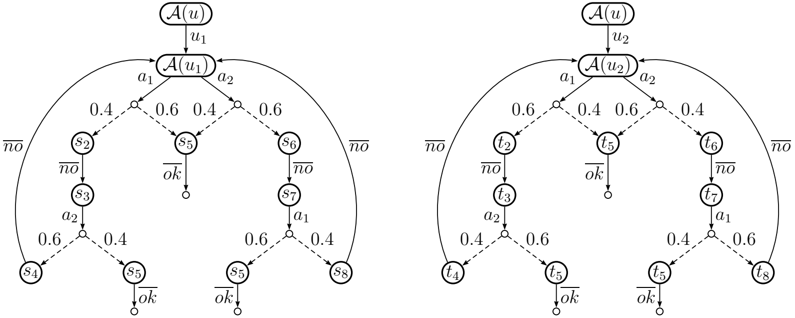

Consider the above example. m^D gives ∞ for the distance between $\mathcal{A}(u_1)$ and $\mathcal{A}(u_2)$.

The Use of the Privacy Budget May Be a bit Wasteful?

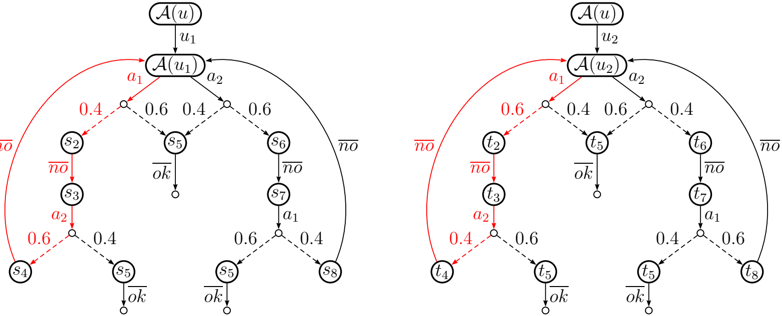

Assume that the scheduler executes the a_1-branch. The ratios of probabilities for $\mathcal{A}(u_1)$ and $\mathcal{A}(u_2)$ producing the same finite sequences:

$$(a_1\overline{no}\,a_2\overline{no})^* \quad : \quad = \left(\frac{0.4\times0.6}{0.6\times0.4}\right)^* = 1$$

$$(a_1\overline{no}\,a_2\overline{no})^* a_1\overline{ok} \quad : \quad = \frac{3}{2}$$

$$(a_1\overline{no}\,a_2\overline{no})^* a_1\overline{no}\,a_2\overline{ok} \quad : \quad = \frac{9}{4}$$

Outline

1. Introduction
   - Concurrent Systems
   - Differential Privacy
   - The Verification Framework

2. Three Pseudometrics
   - The Accumulative Bijection Pseudometric
   - The Amortised Bijection Pseudometric
   - A Multiplicative Variant of the Kantorovich Pseudometric
   - Comparison

The Amortised Bijection Pseudometric

We employ the amortised bisimulation relation from:

- Astrid Kiehn and S. Arun-Kumar.
  Amortised bisimulations. In *FORTE*, 2005.
- Gerald Lüttgen and Walter Vogler.
  Bisimulation on speed: A unified approach. *Theor. Comuput. Sci.*, 2006.

Intuition

The privacy budget in each simulation step may be either reduced due to a negative difference of probabilities, or increased due to a positive difference. Hence, the long-term budget might get amortised.

Definitions

We define amortised bisimulation:

Definition (Amortised bisimulation)

A relation $\mathcal{R} \subseteq S \times S \times [-\epsilon, \epsilon]$ is an ϵ-amortised bisimulation iff for all $(s, t, c) \in \mathcal{R}$:

- $s \xrightarrow{a} \mu$ implies $t \xrightarrow{a} \mu'$ with $\mu \mathcal{L}^A(\mathcal{R}, c)\mu'$
- $t \xrightarrow{a} \mu'$ implies $s \xrightarrow{a} \mu$ with $\mu \mathcal{L}^A(\mathcal{R}, c)\mu'$

Definitions

First, define the corresponding lifting:

Definition (A-Approximate Lifting)

$$\mu \mathcal{L}^A(\mathcal{R}, c)\mu' \quad \text{iff} \quad \exists \text{ bijection } \beta : supp(\mu) \to supp(\mu') \text{ such that}$$

$$\forall s \in supp(\mu) : (s, \beta(s), c + \ln \frac{\mu(s)}{\mu'(\beta(s))}) \in \mathcal{R}$$

We define amortised bisimulation:

Definition (Amortised bisimulation)

A relation $\mathcal{R} \subseteq S \times S \times [-\epsilon, \epsilon]$ is an ϵ-amortised bisimulation iff for all $(s, t, c) \in \mathcal{R}$:

- $s \xrightarrow{a} \mu$ implies $t \xrightarrow{a} \mu'$ with $\mu \mathcal{L}^A(\mathcal{R}, c)\mu'$
- $t \xrightarrow{a} \mu'$ implies $s \xrightarrow{a} \mu$ with $\mu \mathcal{L}^A(\mathcal{R}, c)\mu'$

Verification of differential privacy using m^A

Similarly to the previous section, we can finally define a pseudometric on states as:

$$m^A(s,t) = \min\{\epsilon \mid (s,t,0) \in \mathcal{R} \text{ for some } \epsilon\text{-amortised bisimulation } \mathcal{R}\}$$

Proposition

m^A is a pseudometric.

Verification of differential privacy using m^A

Similarly to the previous section, we can finally define a pseudometric on states as:

$$m^A(s,t) = \min\{\epsilon \mid (s,t,0) \in \mathcal{R} \text{ for some } \epsilon\text{-amortised bisimulation } \mathcal{R}\}$$

Proposition

m^A is a pseudometric.

Theorem

A concurrent system $\mathcal{A}$ is ϵ-differentially private if $m^A(\mathcal{A}(u), \mathcal{A}(u')) \leq \epsilon$ for any two adjacent secrets u and u'.

Indeed, a Thriftier Use of the Privacy Leakage Budget

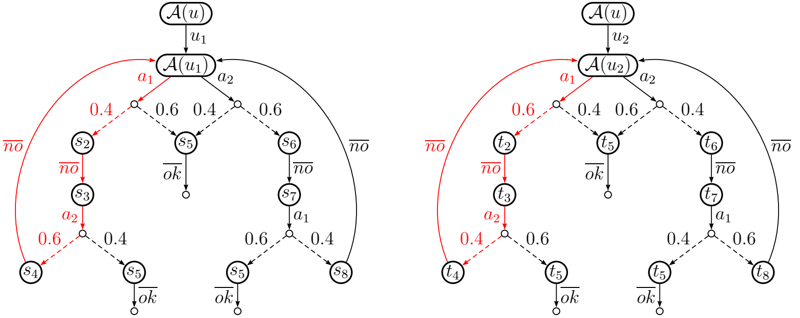

The following relation is an amortised bisimulation between $\mathcal{A}(u_1)$ and $\mathcal{A}(u_2)$.

$$\mathcal{R} = \{ (\mathcal{A}(u_1), \mathcal{A}(u_2), 0),\ (s_2, t_2, \ln\tfrac{2}{3}),\ (s_5, t_5, \ln\tfrac{3}{2}),\ (s_3, t_3, \ln\tfrac{2}{3}),\ (s_4, t_4, 0),\ (s_5, t_5, \ln\tfrac{4}{9}),\ (s_6, t_6, \ln\tfrac{3}{2}),\ (s_5, t_5, \ln\tfrac{2}{3}),\ (s_7, t_7, \ln\tfrac{3}{2}),\ (s_8, t_8, 0),\ (s_5, t_5, \ln\tfrac{9}{4}) \}$$

Thus $m^A(\mathcal{A}(u_1), \mathcal{A}(u_2)) = \ln\frac{9}{4}$, system $\mathcal{A}$ is $\ln\frac{9}{4}$-differentially private.

Outline

How can we get rid of the bijection requirement?

- The second pseudometric is an improvement of the first pseudometric.
- But, both of them are too restrictive! (Bijections between states.)

How can we get rid of the bijection requirement?

- The second pseudometric is an improvement of the first pseudometric.
- But, both of them are too restrictive! (Bijections between states.)

Try to use:

A conventional bisimulation metric: based on the Kantorovich metric.

- Josee Desharnais, Radha Jagadeesan, Vineet Gupta, and Prakash Panangaden.
  The metric analogue of weak bisimulation for probabilistic processes. 2002.
- The Kantorovich metric is a measure of the distance between two probabilistic distributions.

The Standard Definition of Kantorovich Metric.

- Consider a metric m on states, also referred to as the ground distance.
- We lift metric on states to metric on probabilistic distributions, using the Kantorovich metric.
  - Let μ, μ' be distributions on states, the metric $m(\mu, \mu')$ is given by the optimal value of the following primal (dual) program.

Kantorovich Metric: $m(\mu, \mu')$

Primal	maximize $\sum_i (\mu(s_i) - \mu'(s_i))x_i$ subject to $\forall i.\ 0 \leq x_i \leq 1$ $\forall i, j.\ x_i - x_j \leq m(s_i, s_j)$
Dual	minimize $\sum_{i,j} l_{ij} m(s_i, s_j)$ subject to $\forall i.\ \sum_j l_{ij} = \mu(s_i)$ $\forall j.\ \sum_i l_{ij} = \mu'(s_j)$ $\forall i, j.\ l_{ij} \geq 0$

The Standard Definition of Kantorovich Metric.

- Consider a metric m on states, also referred to as the ground distance.
- We lift metric on states to metric on probabilistic distributions, using the Kantorovich metric.
  - Let μ, μ' be distributions on states, the metric $m(\mu, \mu')$ is given by the optimal value of the following primal (dual) program.

Kantorovich Metric: $m(\mu, \mu')$

Primal	maximize $\sum_i (\mu(s_i) - \mu'(s_i))x_i$ subject to $\forall i.\ 0 \leq x_i \leq 1$ $\forall i,j.\ x_i - x_j \leq m(s_i, s_j)$
Dual	minimize $\sum_{i,j} l_{ij} m(s_i, s_j)$ subject to $\forall i.\ \sum_j l_{ij} = \mu(s_i)$ $\forall j.\ \sum_i l_{ij} = \mu'(s_j)$ $\forall i,j.\ l_{ij} \geq 0$

Intuition

Transportation Problem

- l_{ij}: the amount of mass moved from location i of μ to location j of μ'.
- $m(s_i, s_j)$: the cost of moving one unit of mass from location i to location j.

The Standard Kantorovich Metric does not imply differential privacy.

Consider the following example, the value given by the standard Kantorovich metric will be:

Example

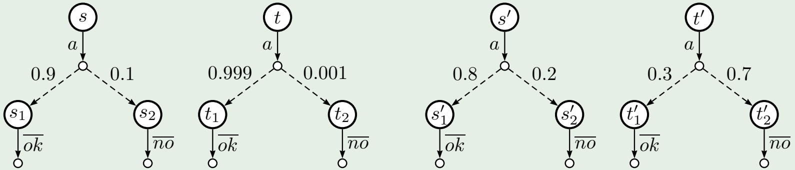

(g) $0.1 - 0.001 = 0.099$, while $\epsilon = \ln \frac{0.1}{0.001} = \ln 100$.

(h) $0.7 - 0.2 = 0.5$, while $\epsilon' = \ln \frac{0.7}{0.2} = \ln 3.5$.

- The standard Kantorovich metric exhibits an additive nature.
- That is inadequate for verifying a multiplicative property such as differential privacy.

A Multiplicative Variant of Kantorovich Metric

Adapting the Kantorovich Metric

	Kantorovich metric	A multiplicative variant
Primal	maximize $\sum_i (\mu(s_i) - \mu'(s_i))x_i$ subject to $\forall i.\ 0 \leq x_i \leq 1$ $\forall i,j.\ x_i - x_j \leq m(s_i, s_j)$	maximize $\ln \dfrac{\sum_i \mu(s_i)x_i}{\sum_i \mu'(s_i)x_i}$ subject to $\forall i.\ 0 \leq x_i \leq 1$ $\forall i,j.\ x_i \leq e^{m(s_i,s_j)}x_j$
Dual	minimize $\sum_{i,j} l_{ij} m(s_i, s_j)$ subject to $\forall i.\ \sum_j l_{ij} = \mu(s_i)$ $\forall j.\ \sum_i l_{ij} = \mu'(s_j)$ $\forall i,j.\ l_{ij} \geq 0$	minimize $\ln z$ subject to $\forall i. \sum_j l_{ij} - r_i = \mu(s_i)$ $\forall j.\ \sum_i l_{ij} e^{m(s_i,s_j)} - r_j \leq z \cdot \mu'(s_j)$ $\forall i,j.\ l_{ij}, r_i, z \geq 0$

This Multiplicative Variant is Well Defined.

Definition (K-State-Metric)

A metric m is a K-state-metric if, for any ϵ, $m(s,t) \leq \epsilon$ implies that if $s \xrightarrow{a} \mu$ then there exists some μ' such that $t \xrightarrow{a} \mu'$ and $m(\mu, \mu') \leq \epsilon$.

We define m^K as the greatest K-state-metric:

$$m^K(s,t) = \min\{m(s,t) \mid m \text{ is a } K\text{-state-metric}\}.$$

This Multiplicative Variant is Well Defined.

Definition (K-State-Metric)

A metric m is a K-state-metric if, for any ϵ, $m(s,t) \leq \epsilon$ implies that if $s \xrightarrow{a} \mu$ then there exists some μ' such that $t \xrightarrow{a} \mu'$ and $m(\mu, \mu') \leq \epsilon$.

We define m^K as the greatest K-state-metric:

$$m^K(s,t) = \min\{m(s,t) \mid m \text{ is a } K\text{-state-metric}\}.$$

This multiplicative variant inherits good merits of the standard one:

Proposition

- *m^K is a pseudometric.*
- *m^K has a fixed-point characterization.*
- *m^K extends bimilarity.*

Verification of differential privacy using m^K

Similarly to the previous two pseudometrics, we can show that

Theorem

A concurrent system $\mathcal{A}$ is ϵ-differentially private if $m^K(\mathcal{A}(u), \mathcal{A}(u')) \leq \epsilon$ for any two adjacent secrets u and u'.

Outline

1. Introduction
   - Concurrent Systems
   - Differential Privacy
   - The Verification Framework
2. Three Pseudometrics
   - The Accumulative Bijection Pseudometric
   - The Amortised Bijection Pseudometric
   - A Multiplicative Variant of the Kantorovich Pseudometric
   - Comparison

Comparison of the Three Pseudometrics

The latter two pseudometrics are more liberal than the first one. Define $m_1 \preceq m_2$: $\forall s, t : m_1(s,t) \geq m_2(s,t)$.

Proposition

- $m^D \preceq m^A$
- $m^D \preceq m^K$

Although they are incomparable to each other. Consider the following toy example in which $m^K(s,t) > m^A(s,t)$:

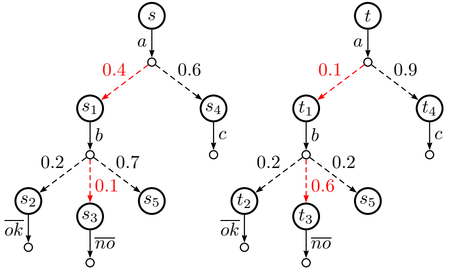

Summary

We have investigated three pseudometrics on states:

- The second pseudometric is designed so that the total privacy leakage bound gets amortised.
- The third one is built on a multiplicative variant of the Kantorovich metric.
- Each of the three pseudometrics establishes a framework for the formal verification of differential privacy for concurrent systems.

- Outlook
  - Whether and how can we define a new pseudometric that unifies the merits of the amortised pseudometric and the multiplicative variant of the Kantorovich metric
  - Compute the pseudometrics
  - Design more subtle approximation relation characterizing (ϵ, δ)-differential privacy

Related Work

Other formal methods on reasoning about differential privacy with programming languages

- type systems: linear types
  - Jason Reed and Benjamin C. Pierce. *Distance makes the types grow stronger: a calculus for differential privacy*. 2010.
  - Marco Gaboardi, Andreas Haeberlen, Justin Hsu, Arjun Narayan, and Benjamin C. Pierce. *Linear dependent types for differential privacy*. In POPL, 2013.
- logic formulations: a relational Hoare logic
  - Gilles Barthe and Boris Köpf and Federico Olmedo and Santiago Z. Béguelin. *Probabilistic relational reasoning for differential privacy*. In POPL. 2012.
  - Gilles Barthe, George Danezis, Benjamin Grégoire, César Kunz, and Santiago Zanella Béguelin. *Verified computational differential privacy with applications to smart metering*. In CSF, 2013.

The End

Thank you very much for your attention!

Questions?