Updates on Research and Collaborations

Matthew Hale
Department of Mechanical and Aerospace Engineering
University of Florida

Center of Excellence for Assured Autonomy in Contested Environments
Fall 2021 Review
November 9th, 2021

Collaborations with CoE PIs

- Hybrid multi-agent optimization with Dawn Hustig-Schultz, Ricardo Sanfelice (UCSC)
  - "Exponentially Converging Distributed Gradient Descent with Intermittent Communication via Hybrid Methods" to appear at CDC '21
  - Ricardo visited UF, made plans for next steps/journal version

- New privacy mechanism on the unit simplex with Parham Gohari, Bo Wu, Ufuk Topcu (UT-A)
  - P. Gohari, B. Wu, C. Hawkins, M. Hale and U. Topcu, "Differential Privacy on the Unit Simplex via the Dirichlet Mechanism," *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 2326-2340, 2021.
  - Parham visited UF again, worked out basis for private policy synthesis in MDPs

- Resilient multi-agent control with Fred Zegers (UF & AFRL) and John Shea, Warren Dixon (UF)
  - F. M. Zegers, M. T. Hale, J. M. Shea and W. E. Dixon, "Event-Triggered Formation Control and Leader Tracking With Resilience to Byzantine Adversaries: A Reputation-Based Approach," *IEEE Transactions on Control of Network Systems*, vol. 8, no. 3, pp. 1417-1429, Sept. 2021.
  - Discussions on incorporating privacy into event-triggered communication

Collaborations with Air Force Colleagues

- Applied optimization work to weapon-target assignment (WTA) problems
  - K. Hendrickson, P. Ganesh, K. Volle, P. Buzaud, K. Brink, and M.T. Hale, "Decentralized Weapon-Target Assignment under Asynchronous Communications", Under review.
  - Kat and Kyle are full-time at UF REEF, collaborations continue

- Developed order-optimal algorithm for anomaly detection in multi-armed bandits with switching costs
  - With Ben Robinson and Beth Morrison at AFRL/RY
  - Publication forthcoming

- Engaging with AFRL every summer
  - William Warke was a Summer Scholar in 2018, 2019 at RW
  - I was a Summer Faculty Fellow at RW in 2020
  - Matthew Ubl was a Summer Scholar in 2021 at RY
  - William Warke applying to RW for 2022,
    Gabriel Behrendt applying to RV for 2022

UF UNIVERSITY of FLORIDA · Duke UNIVERSITY · TEXAS The University of Texas at Austin · UC SANTA CRUZ

Department of Mechanical & Aerospace Engineering
UNIVERSITY of FLORIDA

Differential Privacy for Symbolic Systems with applications to Markov Chains

Bo Chen[a], Kevin Leahy[b], Austin Jones[c], Matthew Hale[a]

Program Review for Center of Excellence on Assured Autonomy in Contested Environments

November 9th

[a] University of Florida, [b] MIT Lincoln Laboratory, [c] Arbor Biotechnologies.

Data driven systems create privacy threats

- Modern systems use more data than ever.

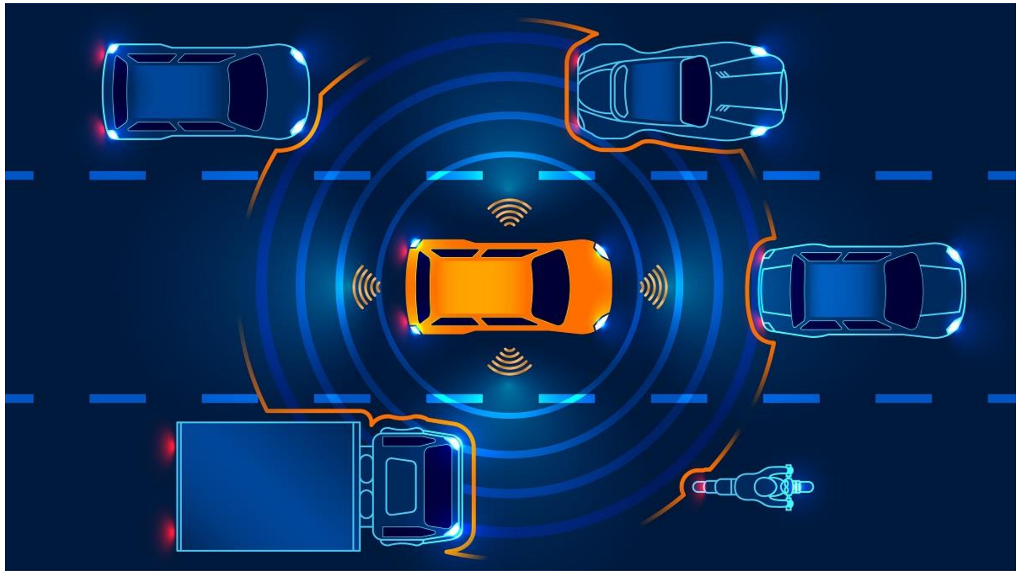


- In controls, sensitive information might be agents' dynamics, control inputs, state trajectories, etc.
- Agents might reveal sensitive information while collaborating.
- In this talk, we focus on state trajectories for symbolic systems.

Strong data protections are difficult to get

- Simply making data anonymous does not work, e.g. Netflix was subject to a linkage attack.

NETFLIX

IMDb

- Takeaway: we don't know what else an adversary might know about us.
- Key question: how can we safeguard information against these threats in symbolic systems?

Solution? Differential Privacy!

- Formal definition of privacy from computer science literature.

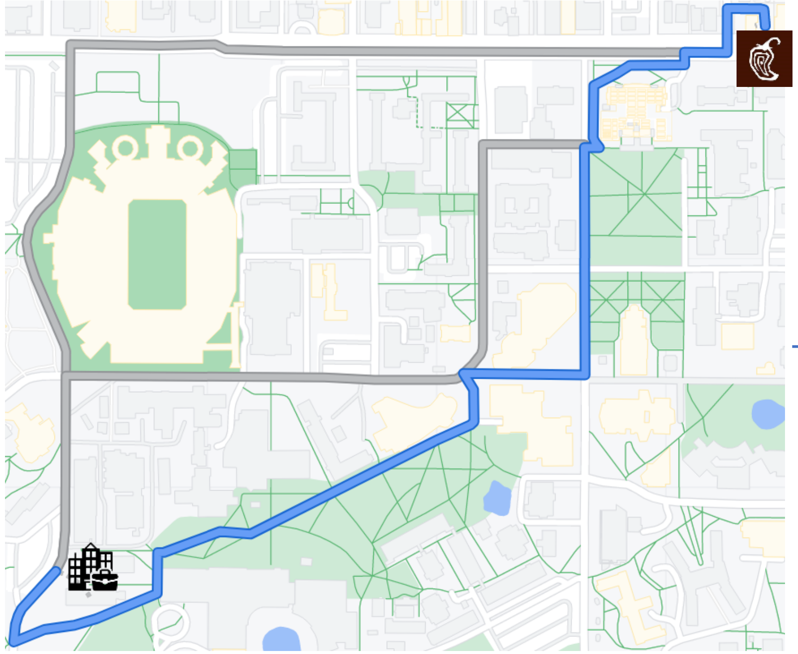

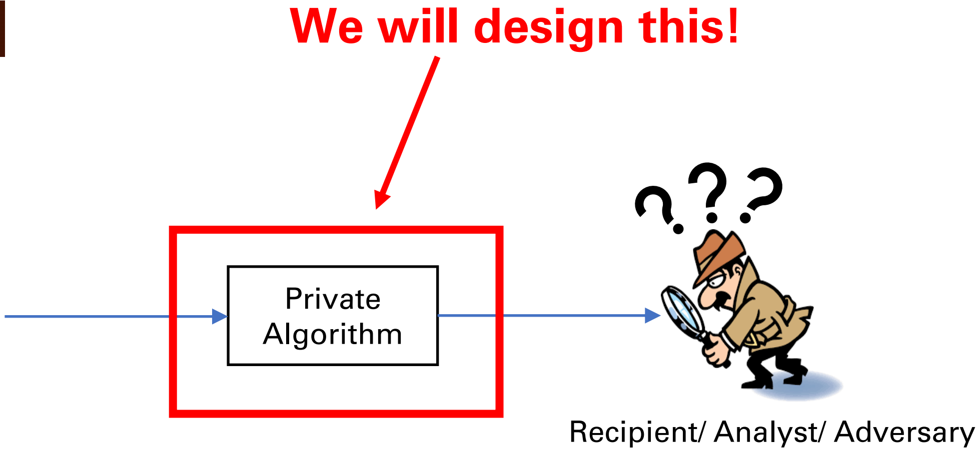

- In short, randomize data to protect it. (Details later.)
- Immune to post-processing: x is private implies f(x) is private.
- No need to anticipate types of privacy attacks.
- Used by Google, Apple, Uber, and the 2020 Census.

New privacy notions are needed here

- Differential privacy is often implemented on numerical system
  - Numerical system: state trajectories can be represented by numbers.
  - For data x, we have a private data $\tilde{x} = x + z$, z is Gaussian or Laplace noise.

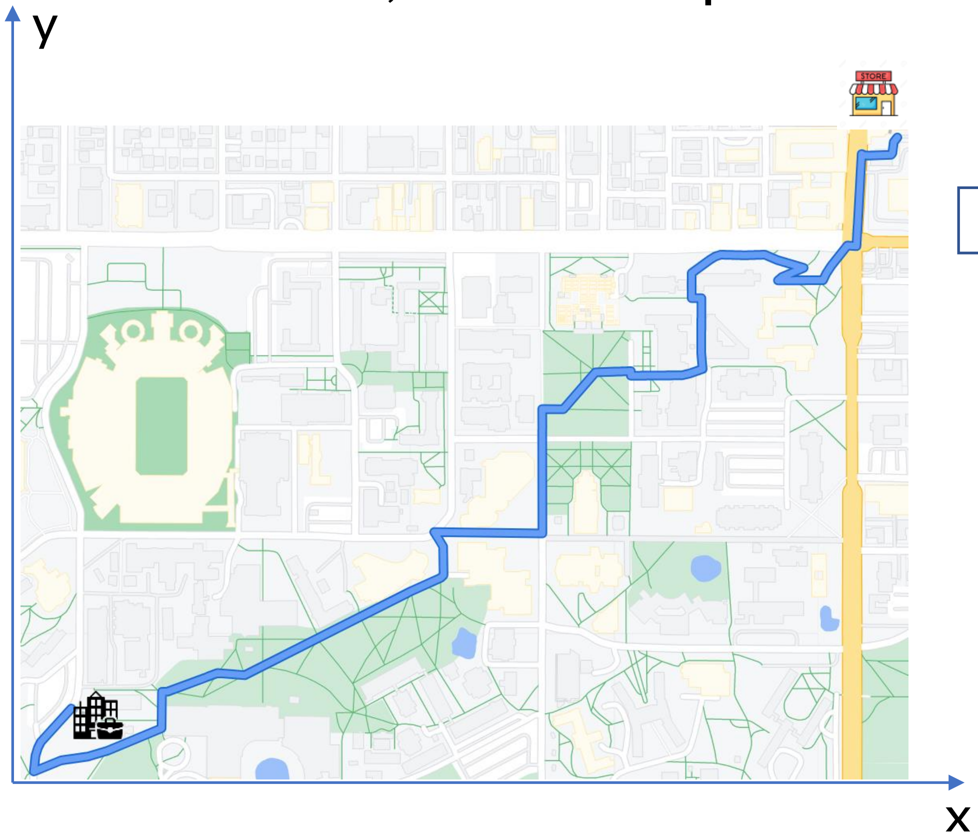

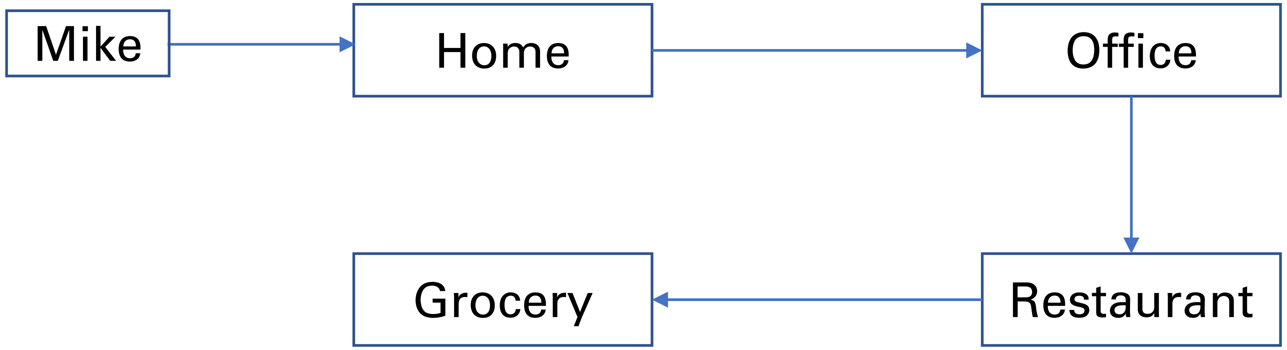

Mike's daily trajectory: ("Home", "Office", "Restaurant", "Grocery")

- How about symbolic systems?

Markov chain, a special symbolic system

- In this talk we focus on Markov chains.
- A Markov chain is a stochastic model describing a sequence of random variables $S_1, S_2 \ldots S_n$ such that

$$\Pr[S_{t+1} | S_t, S_{t-1}, \ldots, S_1] = \Pr[S_{t+1} | S_t]$$

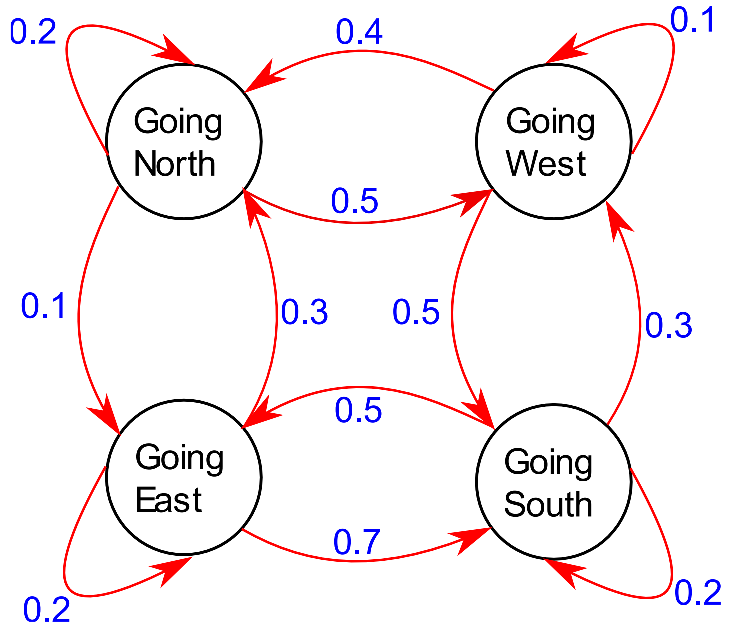

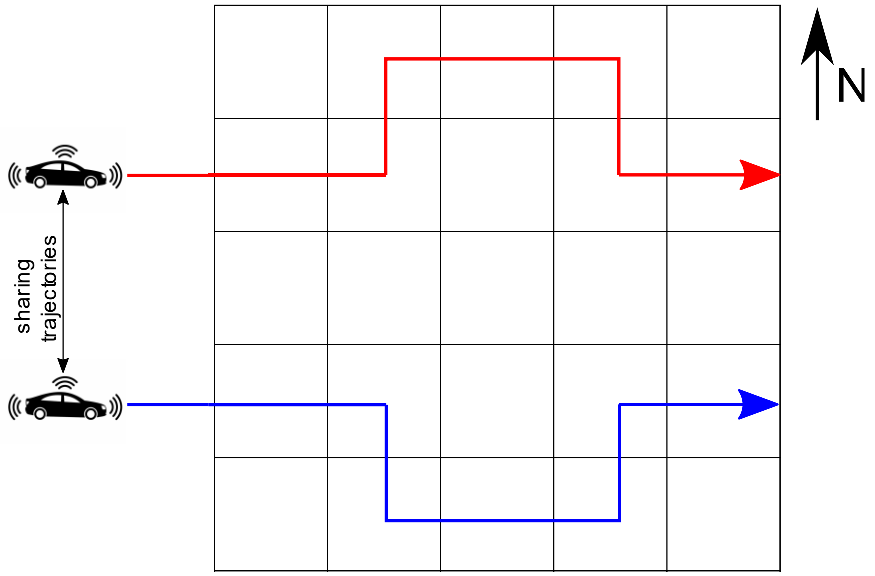

- States can be non-numerical.

1. Bo Chen, Kevin Leahy, Austin Jones, Matthew Hale, "Differential Privacy for Symbolic Systems with applications to Markov Chains"

Private outputs can't be nonsense

- Goal for privacy of Markov chains: privatize sequence of states.[1]
- For a private sequence $w = s_1 s_2 \ldots s_n$ and any t, we must enforce

$$\Pr[s_{t+1} | s_t] > 0$$

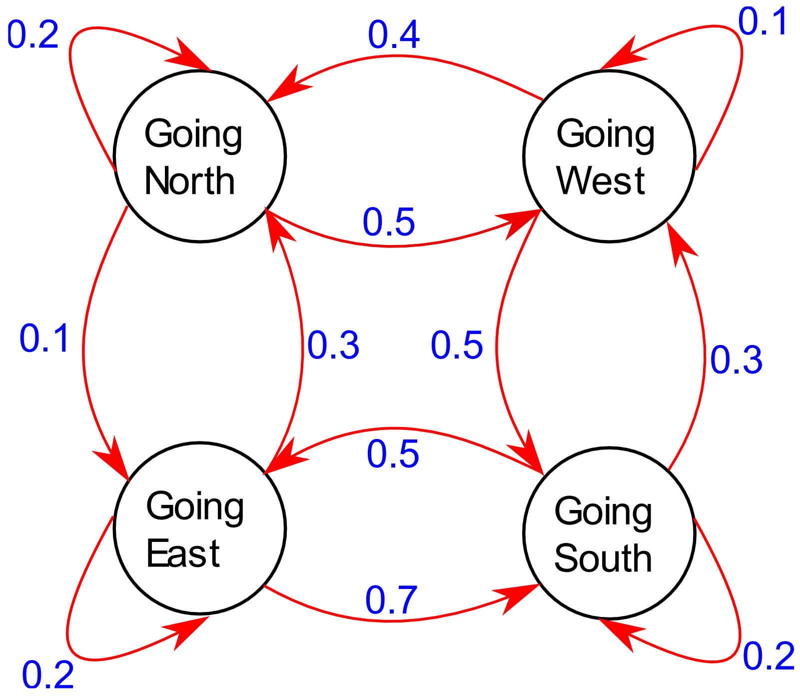

(N,W,N,E) → Private Algorithm → ($\color{red}{N}$,$\color{red}{S}$,E,E) Infeasible!

Output 1

(N,W,N,E) → Private Algorithm → (N,E,E,N) Feasible!

Output 2

1. Bo Chen, Kevin Leahy, Austin Jones, Matthew Hale, "Differential Privacy for Symbolic Systems with applications to Markov Chains"

Differential Privacy on symbolic systems

- Goal of differential privacy: generate randomized outputs in order to "mask" differences between "similar" sequences.
- "Similar" sequences are defined by adjacency relationship.

Definition 1 (Word Adjacency): For a positive integer n and k, the word adjacency relation between two words w_1, w_2 is $Adj_{n,k} = \{(w_1, w_2) | d(w_1, w_2) \leq k\}$.

Definition 2 (Hamming Distance): The Hamming distance between sequences w_1, w_2 denoted by $d(w_1, w_2)$, is the minimum number of substitutions that can be applied to w_1 to convert it to w_2.

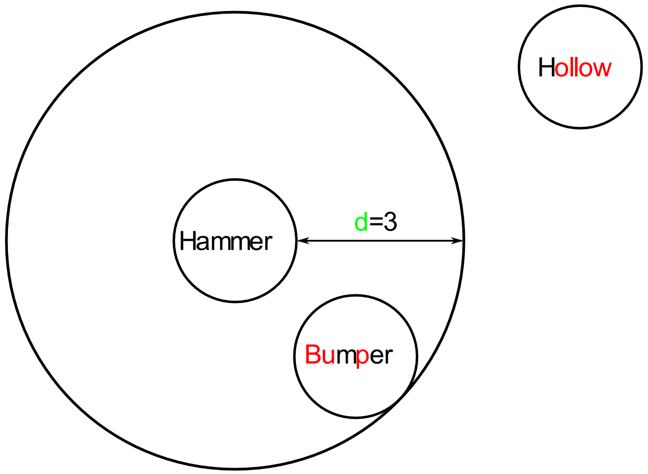

Differential Privacy on symbolic systems

- Goal of differential privacy: generate randomized outputs in order to "mask" differences between "similar" sequences.
- "Mask" means an adversary can not reliably tell if an output sequence is generated by an individual sequence or any adjacent sequence.

Definition 3 (Word Differential Privacy):
Let $\varepsilon > 0$. A randomized algorithm M is ε-differential private if for all $S \subseteq \mathrm{Range(M)}$ and for all word adjacent sequence $(w_1, w_2) \in Adj_{n,k}$ we have

$$\Pr[M(w_1) \in S] \leq \exp(\varepsilon)\Pr[M(w_2) \in S]$$

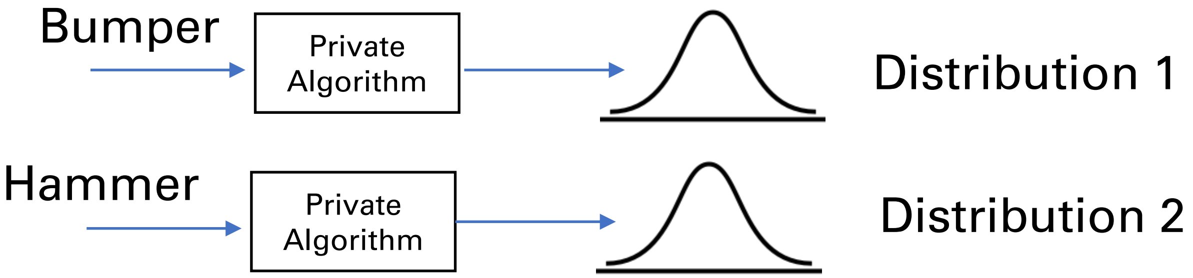

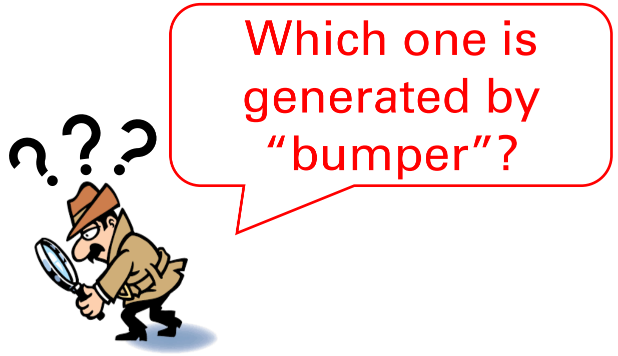

Recipient/ Analyst/ Adversary

- Smaller epsilon implies stronger privacy.
- In literatures, ε is ranging from [0.01,10].[2]

Department of Mechanical & Aerospace Engineering
UNIVERSITY of FLORIDA

2. Hsu, Justin and Gaboardi, Marco and Haeberlen, Andreas and Khanna, Sanjeev and Narayan, Arjun and Pierce, Benjamin C. and Roth, Aaron, "Differential Privacy: An Economic Method for Choosing Epsilon".

We need two types of mechanisms

i) Offline Mechanism (batch privacy)

- Privatize the whole sensitive sequence $w = \sigma_1\sigma_2 \dots \sigma_n$ at once.

ii) Online Mechanism (real-time privacy)

- Generate differentially private outputs but future states are unknown.

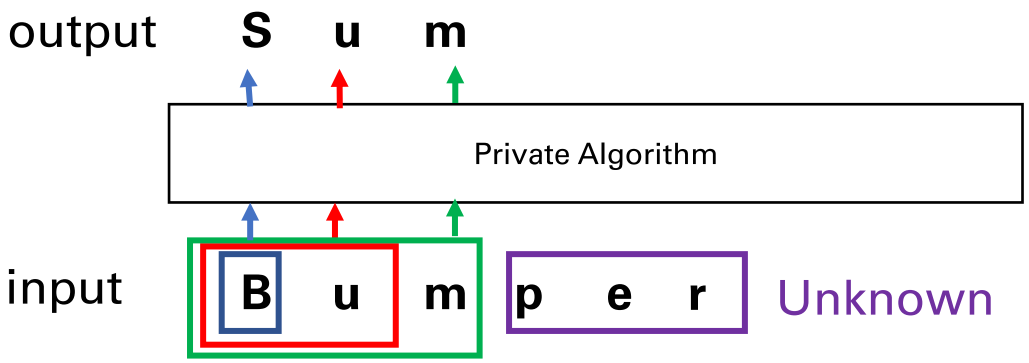

iii) For both

- The private outputs of states are feasible.
- Quantify tradeoffs between privacy and accuracy.

UF Department of Mechanical & Aerospace Engineering
UNIVERSITY of FLORIDA

Construct offline mechanism

- Main idea: each feasible sequence can be selected based on Hamming distance. → Time complexity: $O(|S|^n)$.
- Challenge: We need to make sure this is efficient!
- Procedure of constructing offline mechanism.
  - Step 1: Select a Hamming distance l.
  - Step 2: Select a private output from only sequences that have Hamming distance l to the input sequence.

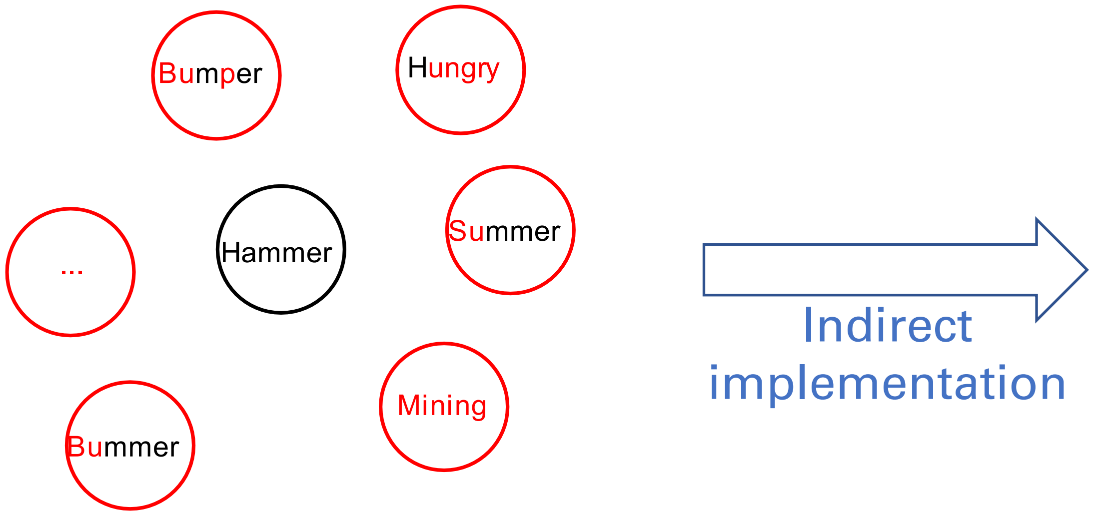

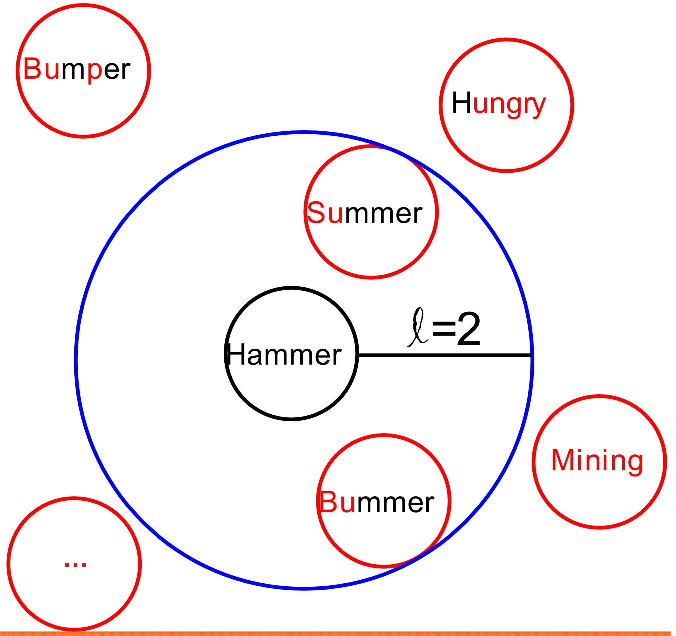

Step 1: Select a Hamming distance

- For an input sequence x, Adjacency $Adj_{|x|,k}$, and privacy parameter ε, select a Hamming distance using the distribution

$$\rho(l;|x|,k)=\frac{m_l \exp\left(-\frac{\varepsilon l}{2k}\right)}{\sum_{i=0}^{|x|} m_i \exp\left(-\frac{\varepsilon i}{2k}\right)}$$

Length of sensitive input word

Number of possible sequences that are distance i from input

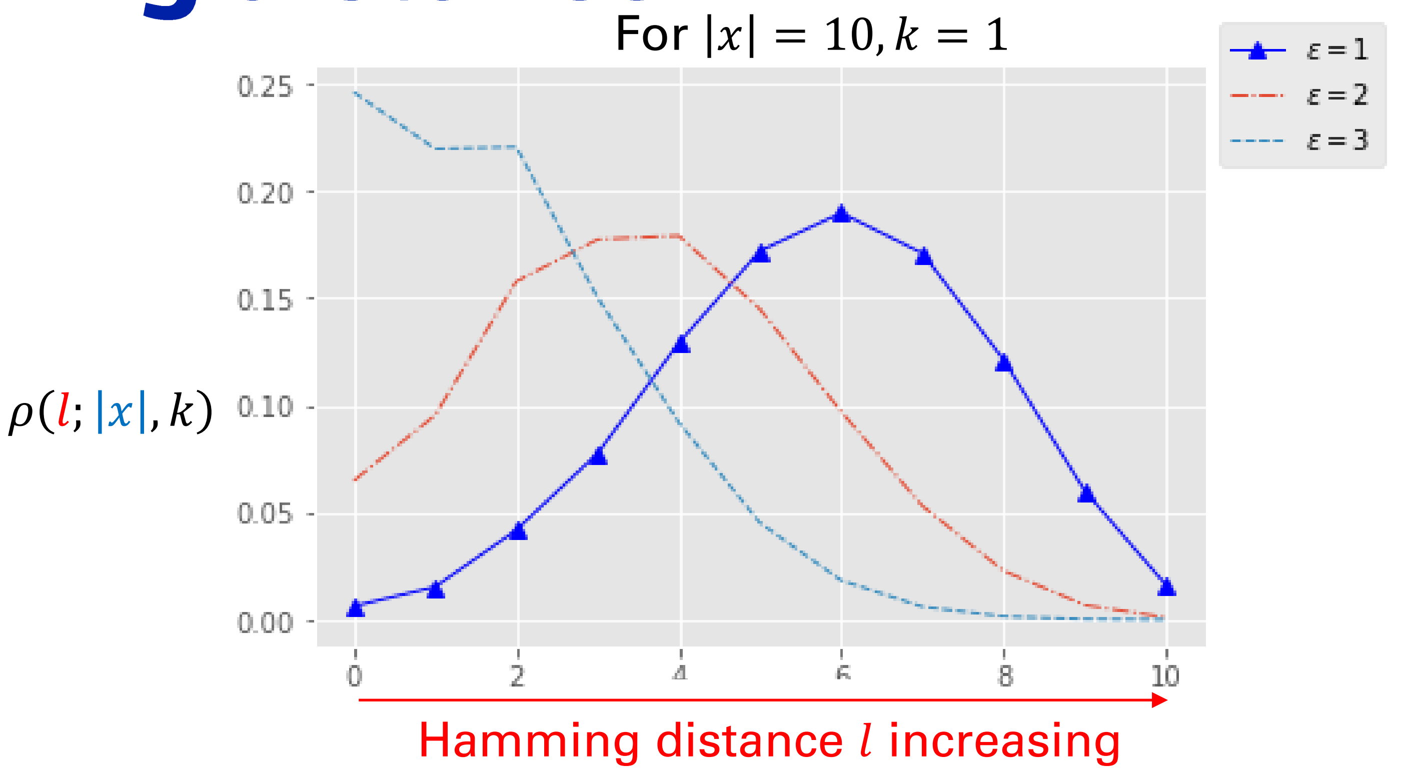

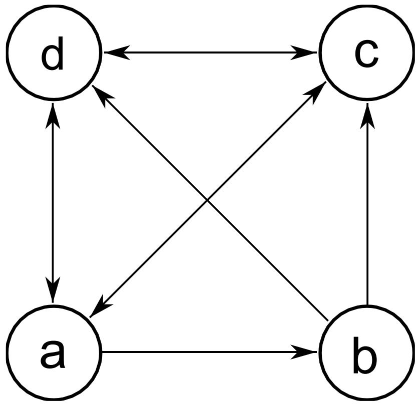

Step 2: Select a private output

- Find sequences which are with a Hamming distance l efficiently.
- For a non-Markov symbolic system, we can do this using modified Hamming distance automaton. (Will make it Markov in next slide)
- **Takeaway**: This automaton is efficient and generate output sequence.

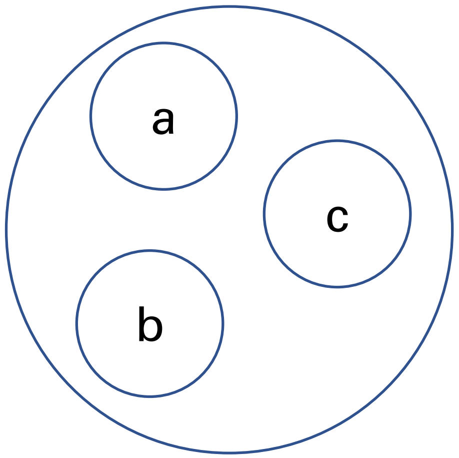

Set of possible output characters which are independent with each other.

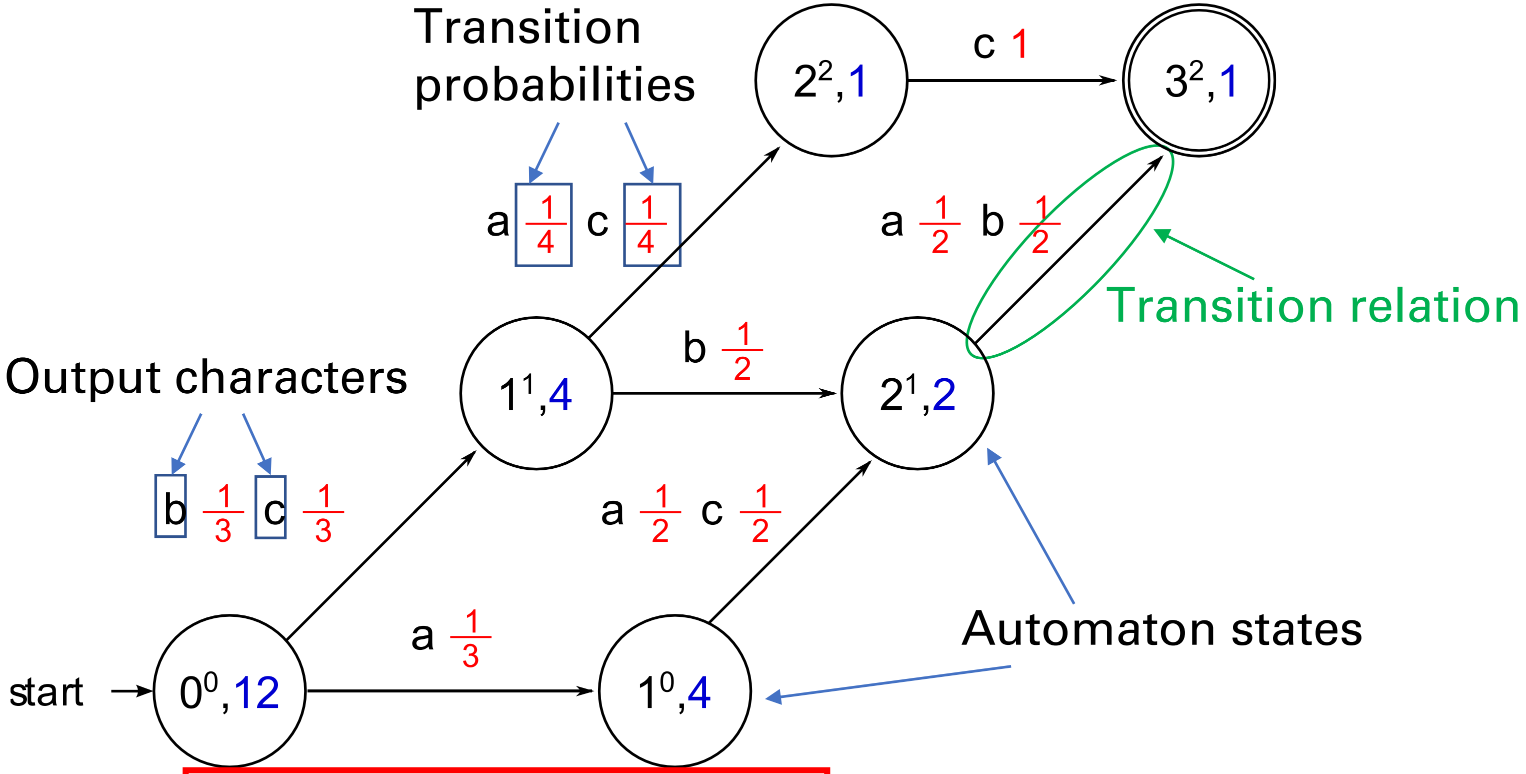

For an input "abc" and l=2

Step 2: Select a private output

- Main idea: to extend to a Markov chain, we make synchronous product of the modified Hamming distance automaton and the Markov chain.

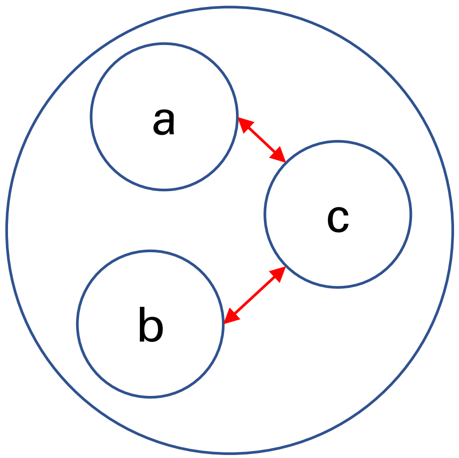

State space of a Markov chain

Synchronous state

Construct it in this way

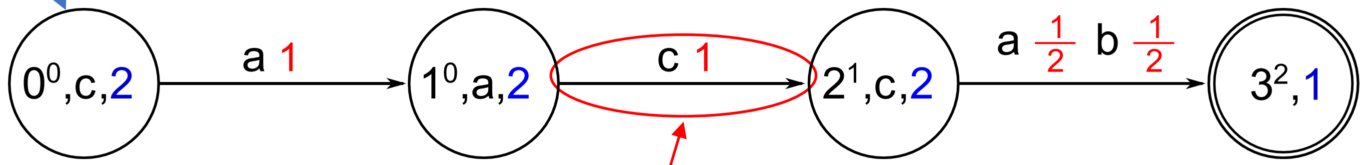

Exist when it satisfies the transition relation of both the modified Hamming distance automaton and the Markov chain

- The offline mechanism first selects a Hamming distance l, then selects an output sequence by running a product modified Hamming distance automaton.

Key Result: The offline mechanism is ε-differentially private.

Concentration bounds

Theorem 2 (differential privacy and concentration bounds): For an input sequence w_i, let w_o be an output sequence generated by the offline mechanism, then w_o is ε-differentially private and the expectation and variance of distance is bounded by

Depends on ε and k

$$\frac{n(N_{min}-1)B_{\varepsilon,k}[(N_{min}-1)B_{\varepsilon,k}+1]^{n-1}}{\sum_{i=0}^{|x|} m_i \exp\left(-\frac{\varepsilon i}{2k}\right)} \leq E[d(w_i,w_o)] \leq \frac{nN_{max}B_{\varepsilon,k}[N_{max}B_{\varepsilon,k}+1]^{n-1}}{\sum_{i=0}^{|x|} m_i \exp\left(-\frac{\varepsilon i}{2k}\right)}$$

N_{min}, N_{max}: min/max outdegree

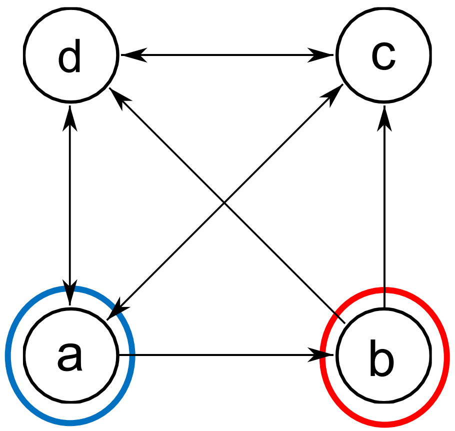

$N_{min} = 2, N_{max} = 3$

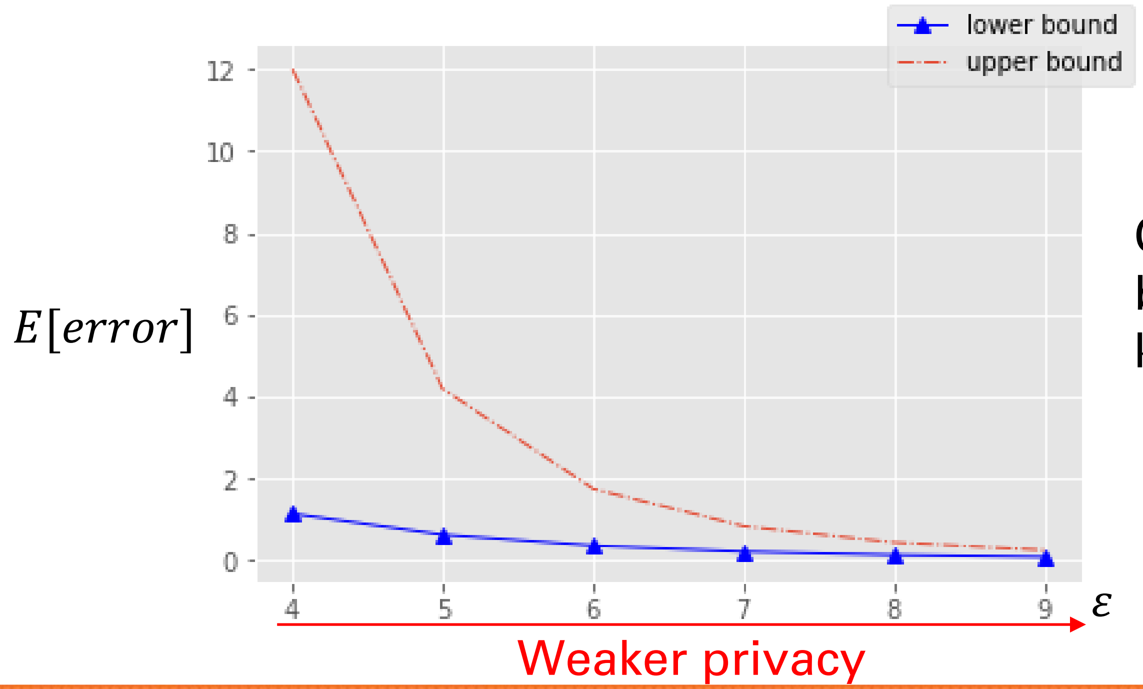

Concentration bounds when k=1 and n=10.

Online mechanism for Markov chains

- Main idea: each output state is generated based on the most recently generated private state.

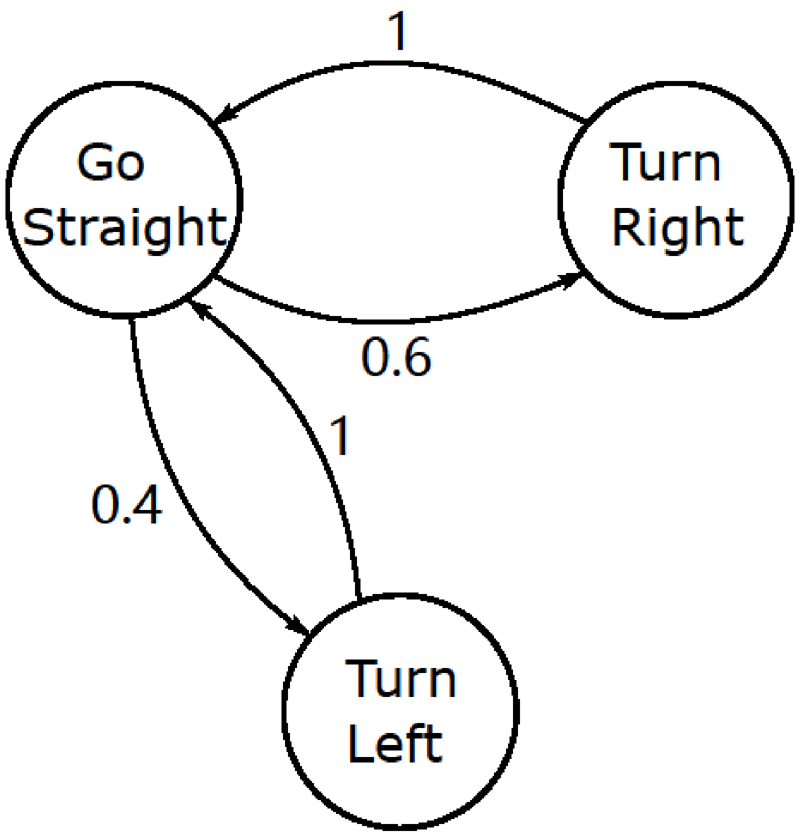

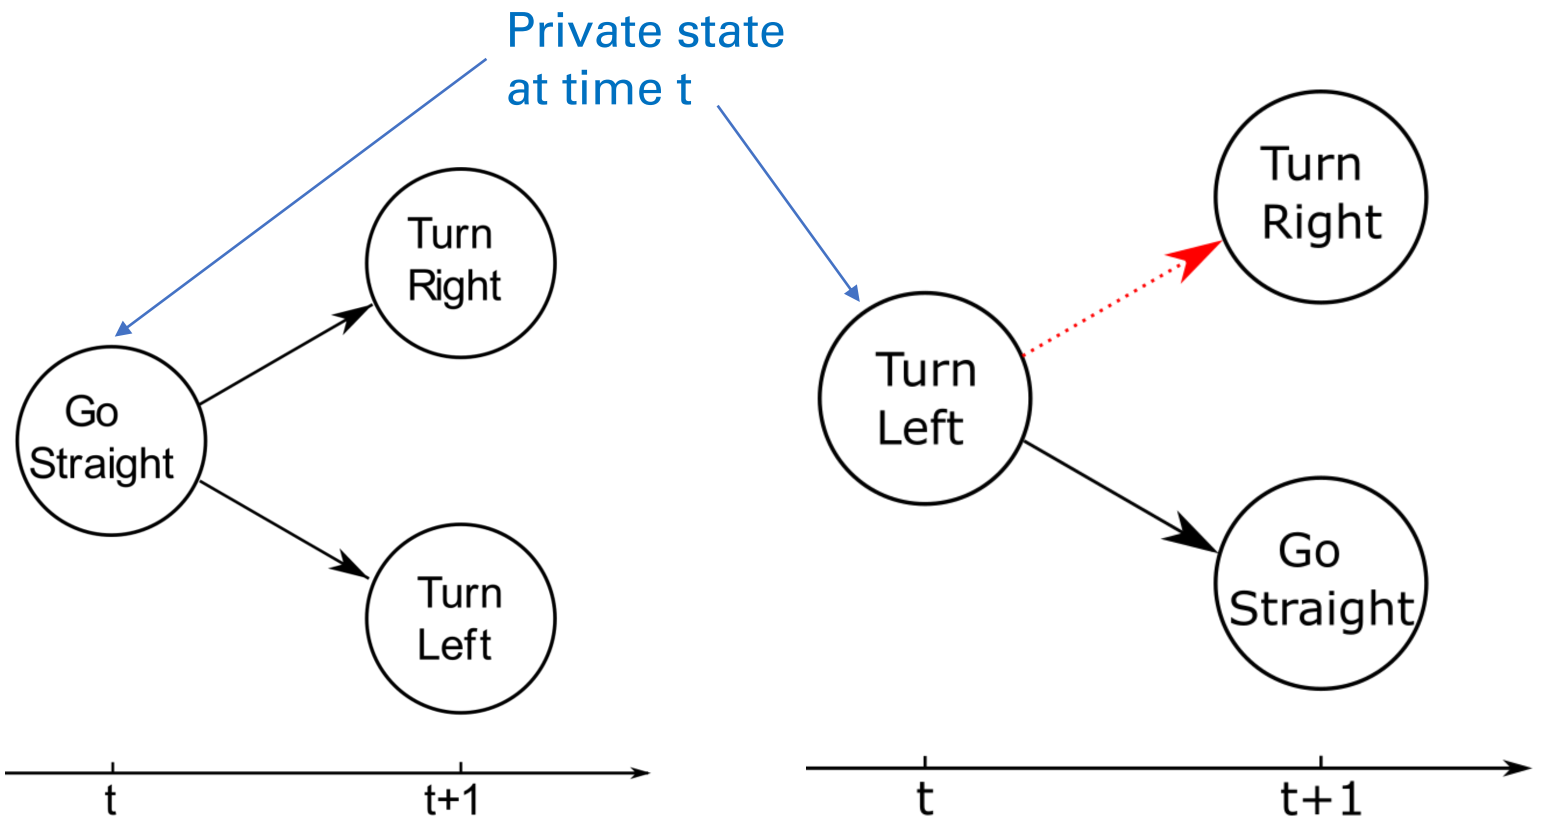

Markov chain application 1

Markov chain application 2

Online mechanism for Markov chains

- For a sensitive input state s_t,
  - If s_t is feasible from the most recently private state, then $\Pr[s_t]$ is set to τ and other states will have identical probability whose sum is $1 - \tau$.
  - If s_t is not feasible then all feasible states will have identical probability whose sum is 1.

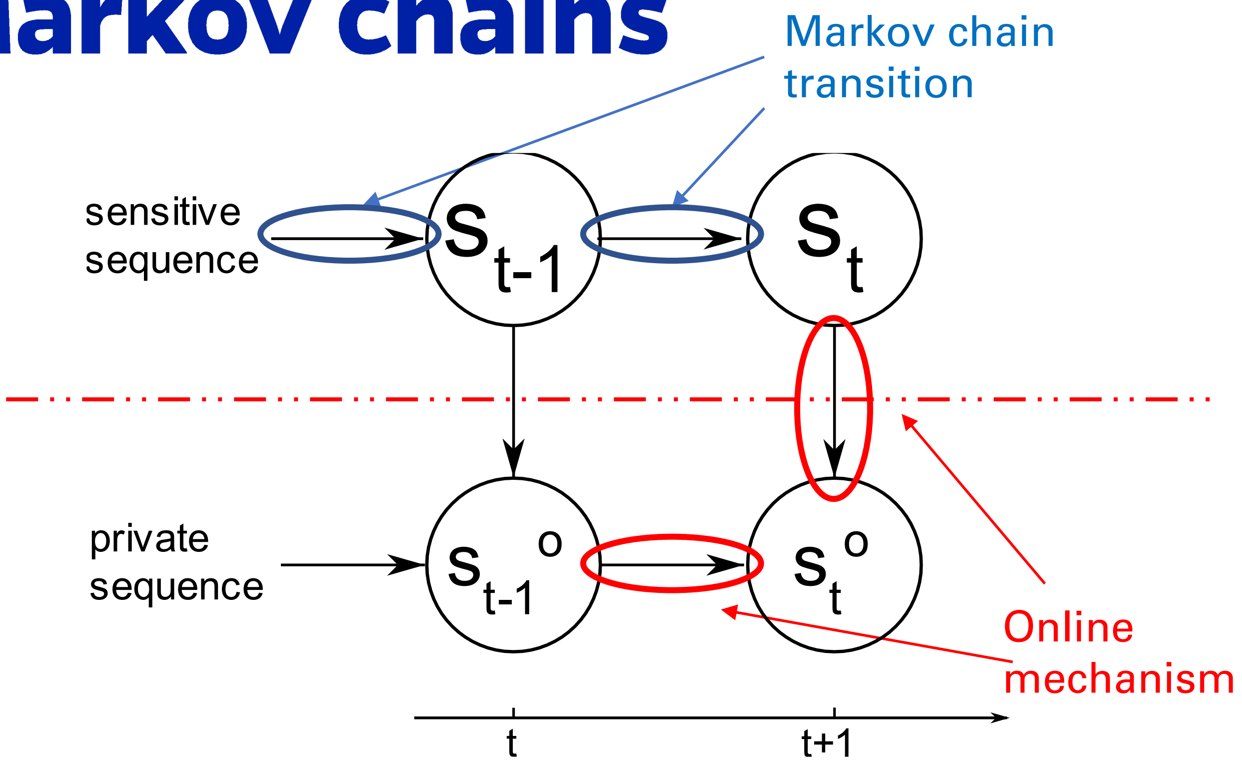

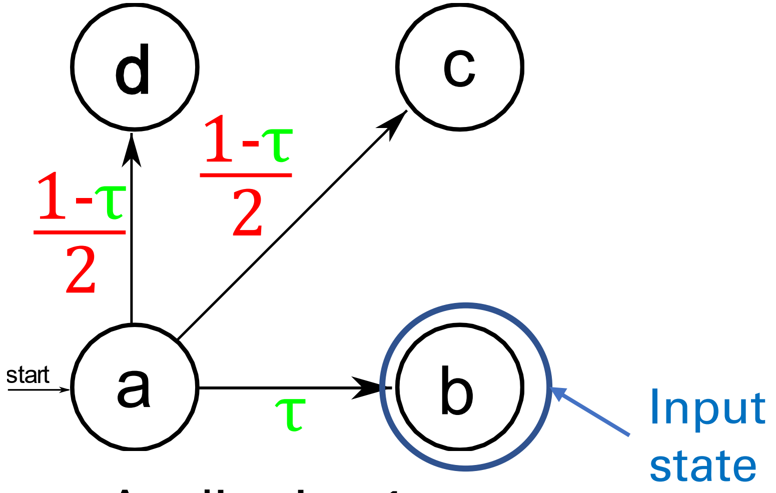

Application 1

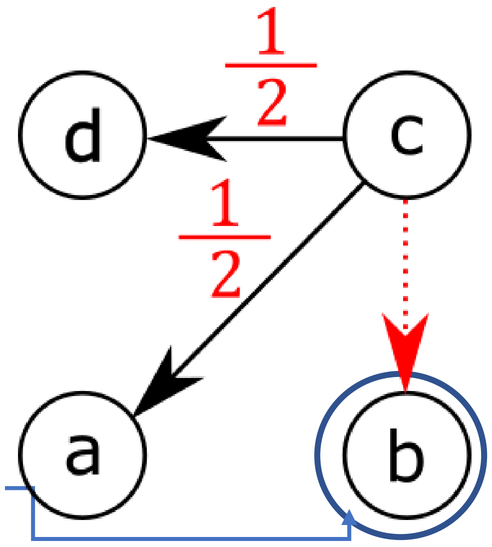

Application 2

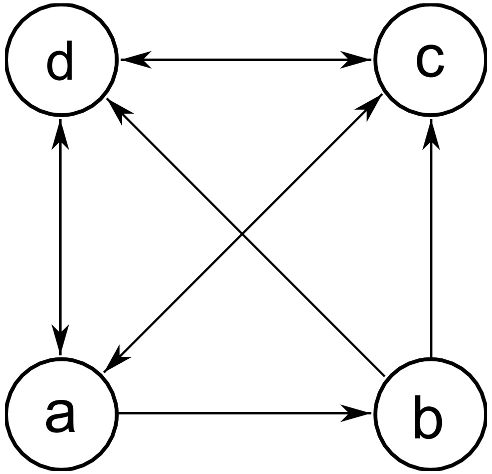

Online mechanism for Markov chains

Theorem 2 (Online Mechanism is differentially private): For a sensitive input sequence $w_o = s_1^o s_2^o \dots s_n^o$ and an initial private state s_0^o, the online mechanism is word ε-differentially private if

$$\tau(s_t^o) = \frac{1}{\left(N\left(s_t^o\right) - 1\right)\exp\left(-\frac{\varepsilon}{k}\right) + 1}$$

the number of feasible states

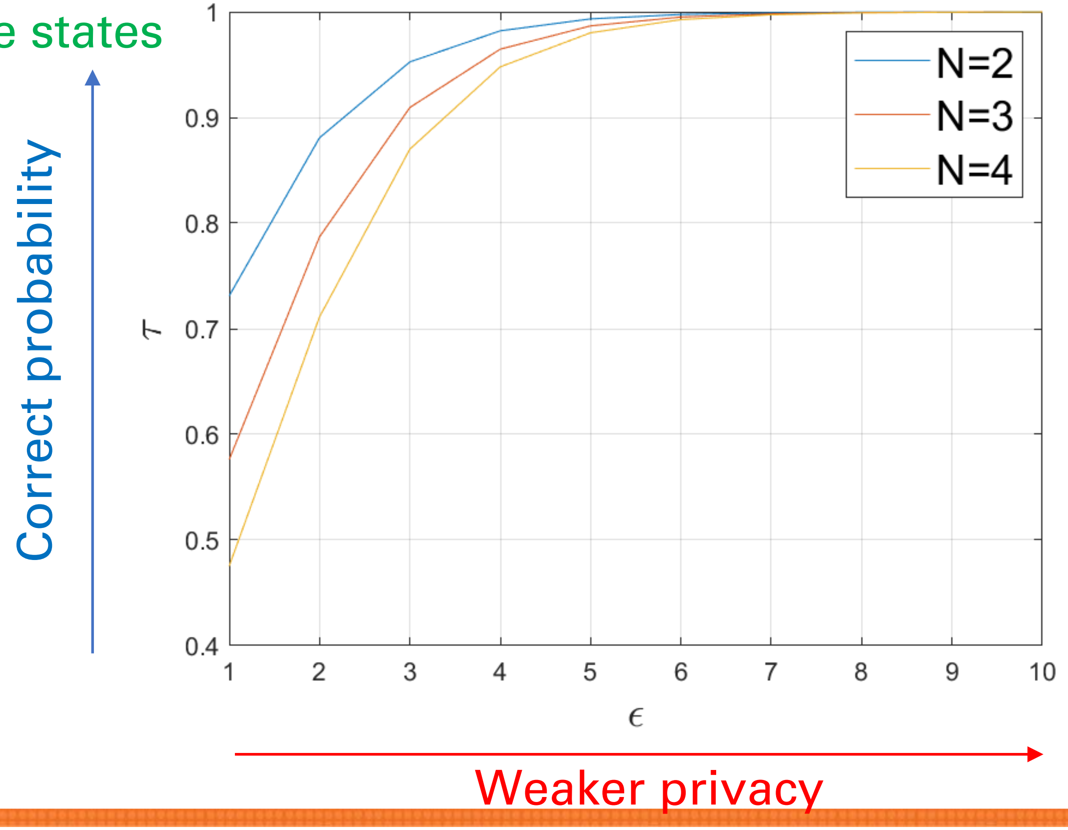

1. Bo Chen, Kevin Leahy, Austin Jones, Matthew Hale, "Differential Privacy for Symbolic Systems with applications to Markov Chains"

Experiment

- Example Markov chain is generated by the book “Green Eggs and Ham”.
- 50 unique words → 50 states in total.
- We generate differentially private versions of a sequence “I do so like green eggs and ham thank you thank you Sam I am”.

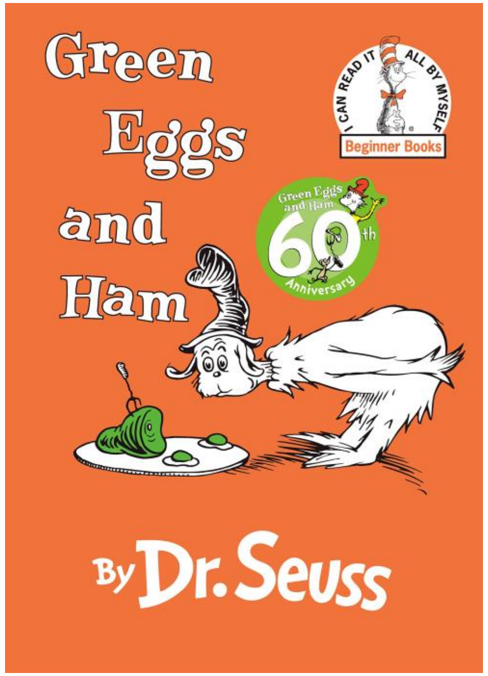

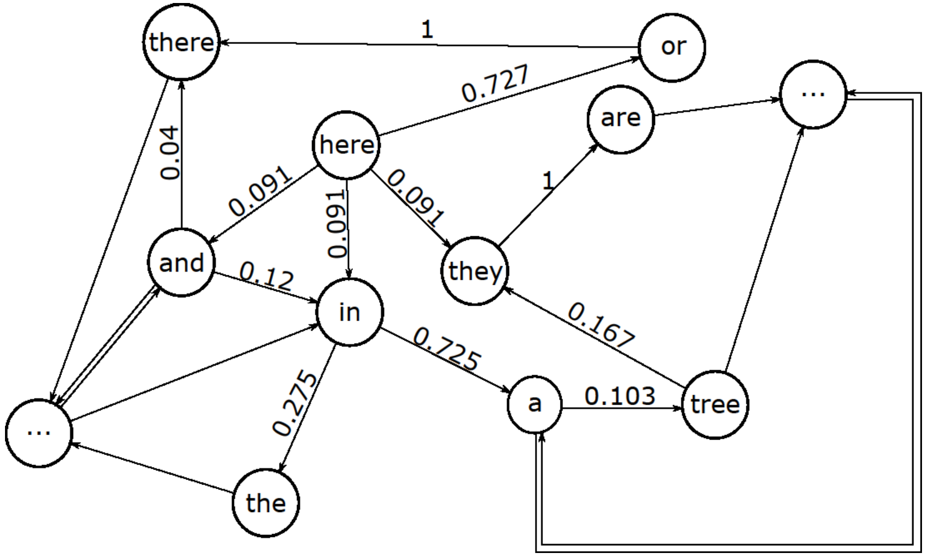

Results for different epsilon

Sensitive input: I do so like green eggs and ham thank you thank you Sam I am

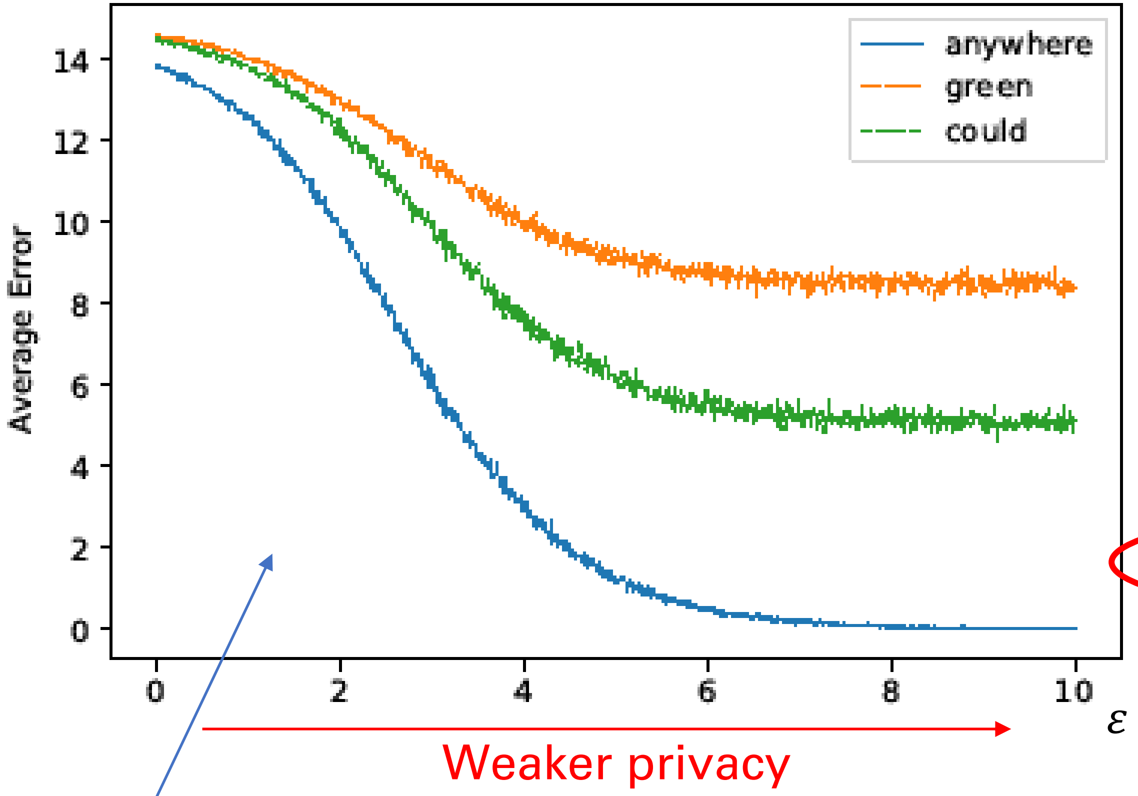

Table 1
Samples of differentially private strings with different values of the privacy parameter ϵ and initial state "anywhere".

ϵ	Start at "anywhere"	error
0.1	"you see you see so you in the dark I am say Sam Sam Sam"	15
1	"I would not eat green egg and ham thank you eat them eat them with	8
10	"I do so like green eggs and ham thank you thank you Sam I am"	0

This is random! We just got lucky!

Different starting word incurs different errors.

Department of Mechanical & Aerospace Engineering
UNIVERSITY of FLORIDA

1. Bo Chen, Kevin Leahy, Austin Jones, Matthew Hale, "Differential Privacy for Symbolic Systems with applications to Markov Chains"

Future works

- Generalize to Partial Observable MDPs.
- Using this work on multi-agent reinforcement learning.

Thanks for listening!

Result for different conditions.

Sensitive input: I do so like green eggs and ham thank you thank you Sam I am