THE MINIMAL HIDDEN COMPUTER NEEDED TO IMPLEMENT A VISIBLE COMPUTATION

David H. Wolpert

Santa Fe Institute, MIT, ASU [*http://davidwolpert.weebly.com*](http://davidwolpert.weebly.com)

Joint work with

Artemy Kolchinsky, *Santa Fe Institute*; Jeremy Owen, *MIT*

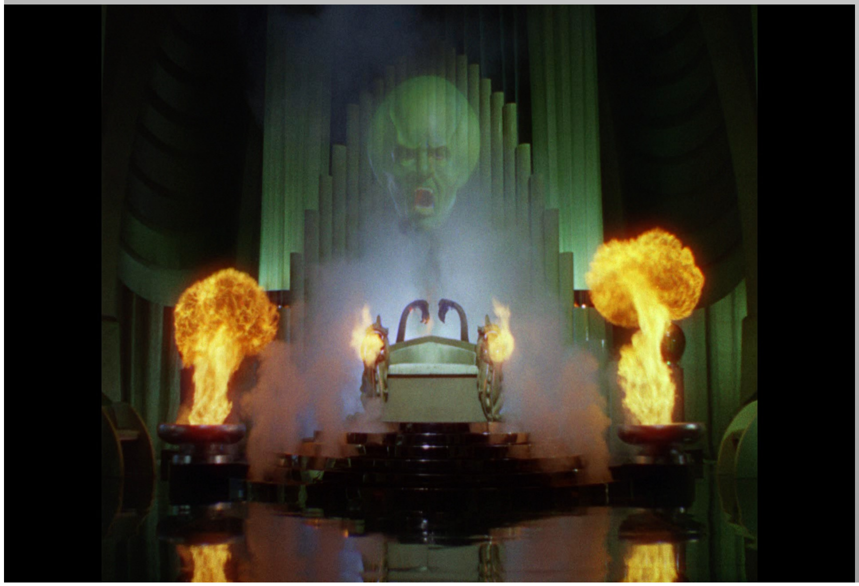

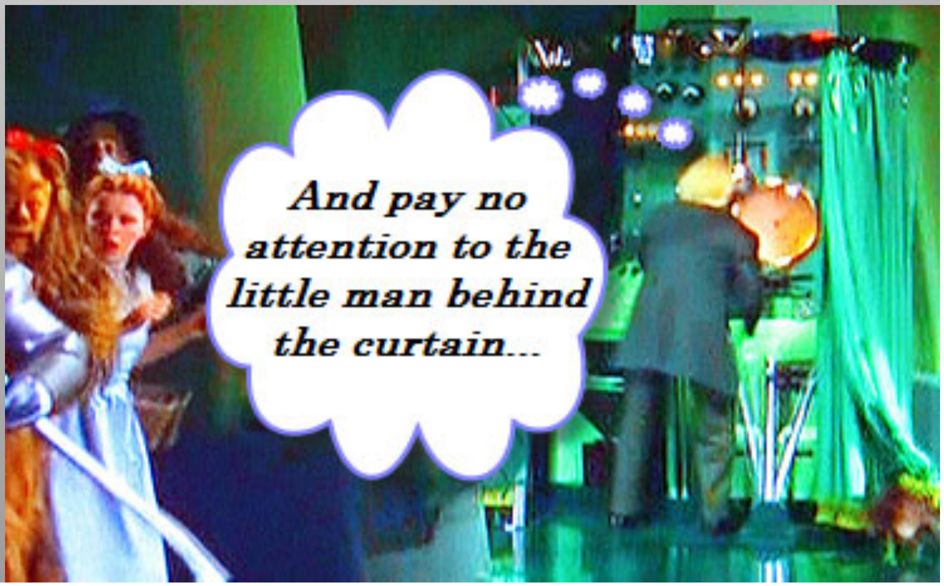

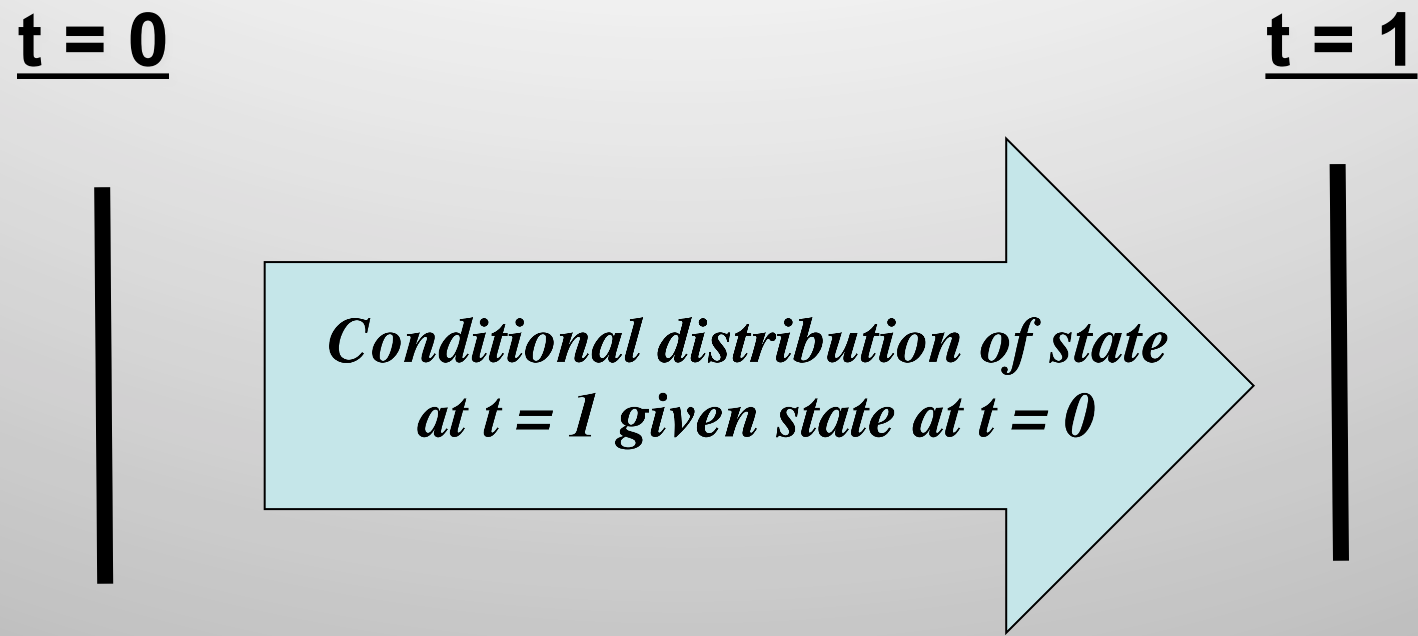

- Visible states arrayed along black lines

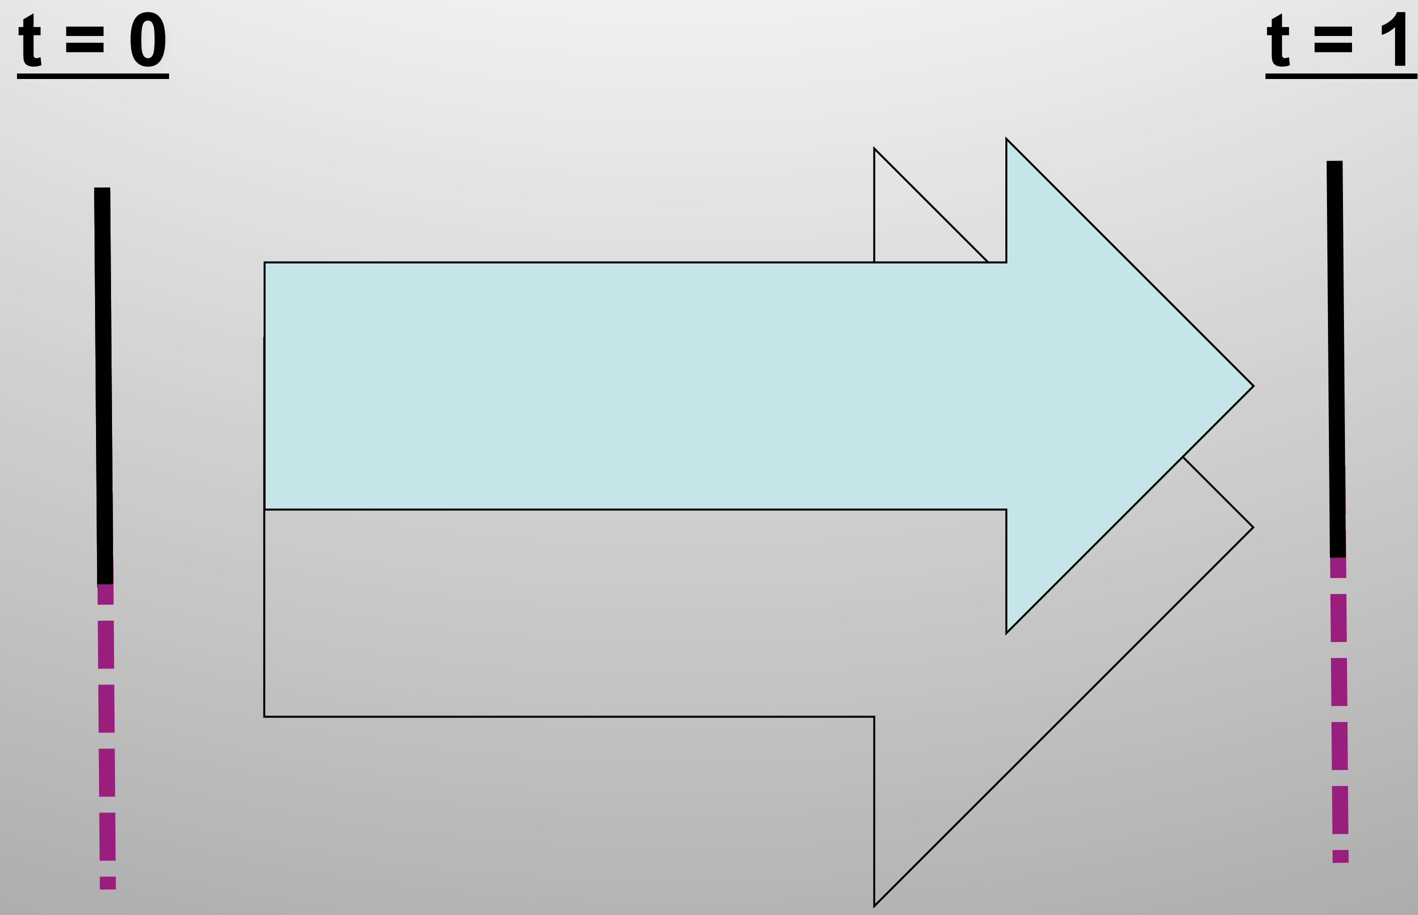

- Visible states arrayed along black lines
- Hidden states arrayed along dotted purple lines

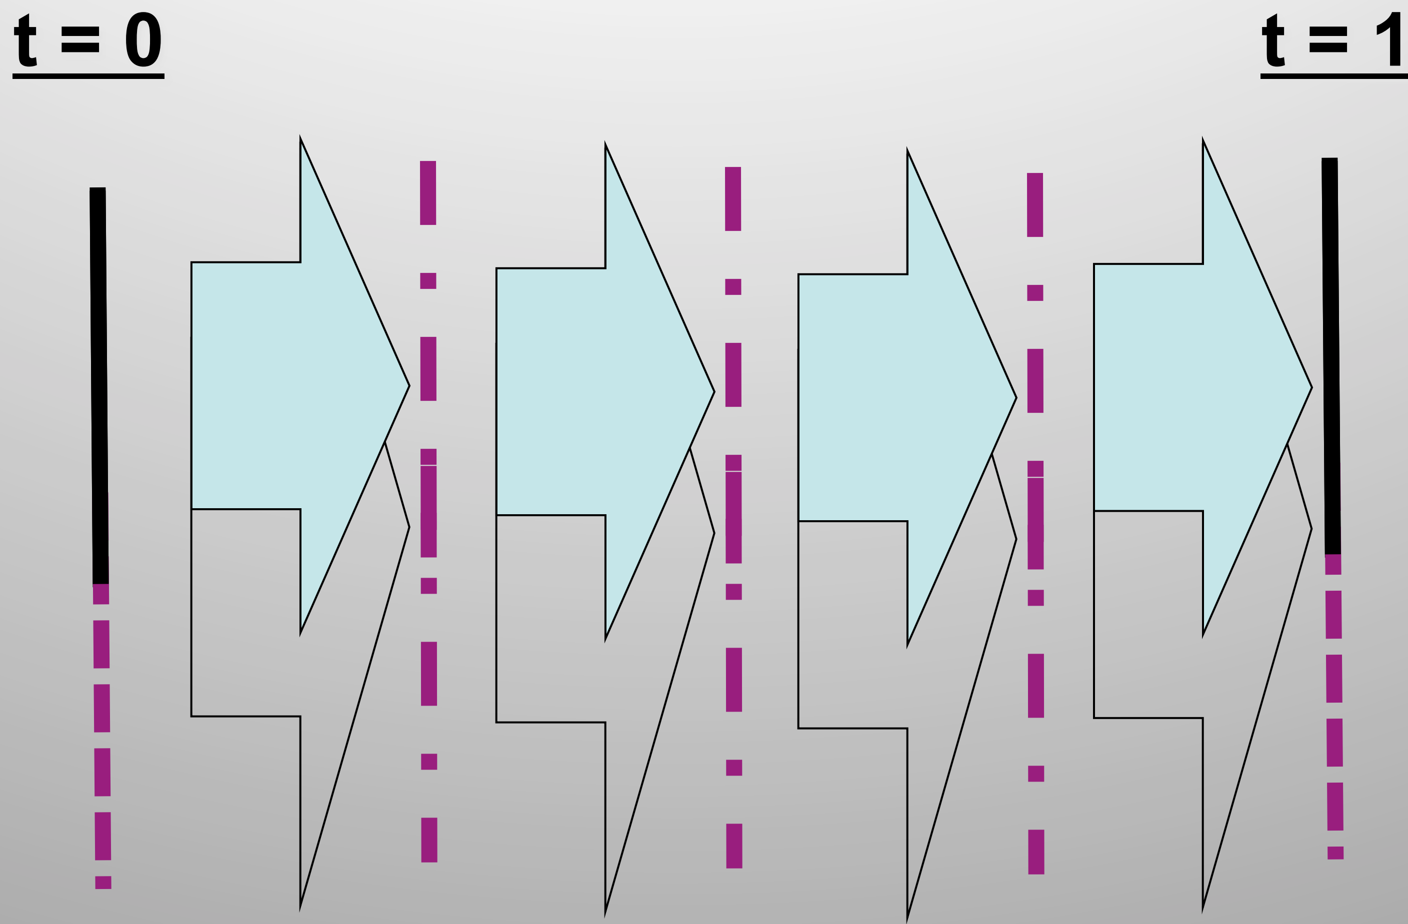

- Visible states arrayed along black lines
- Hidden states arrayed along dotted purple lines
- Each "column" of arrows = one timestep

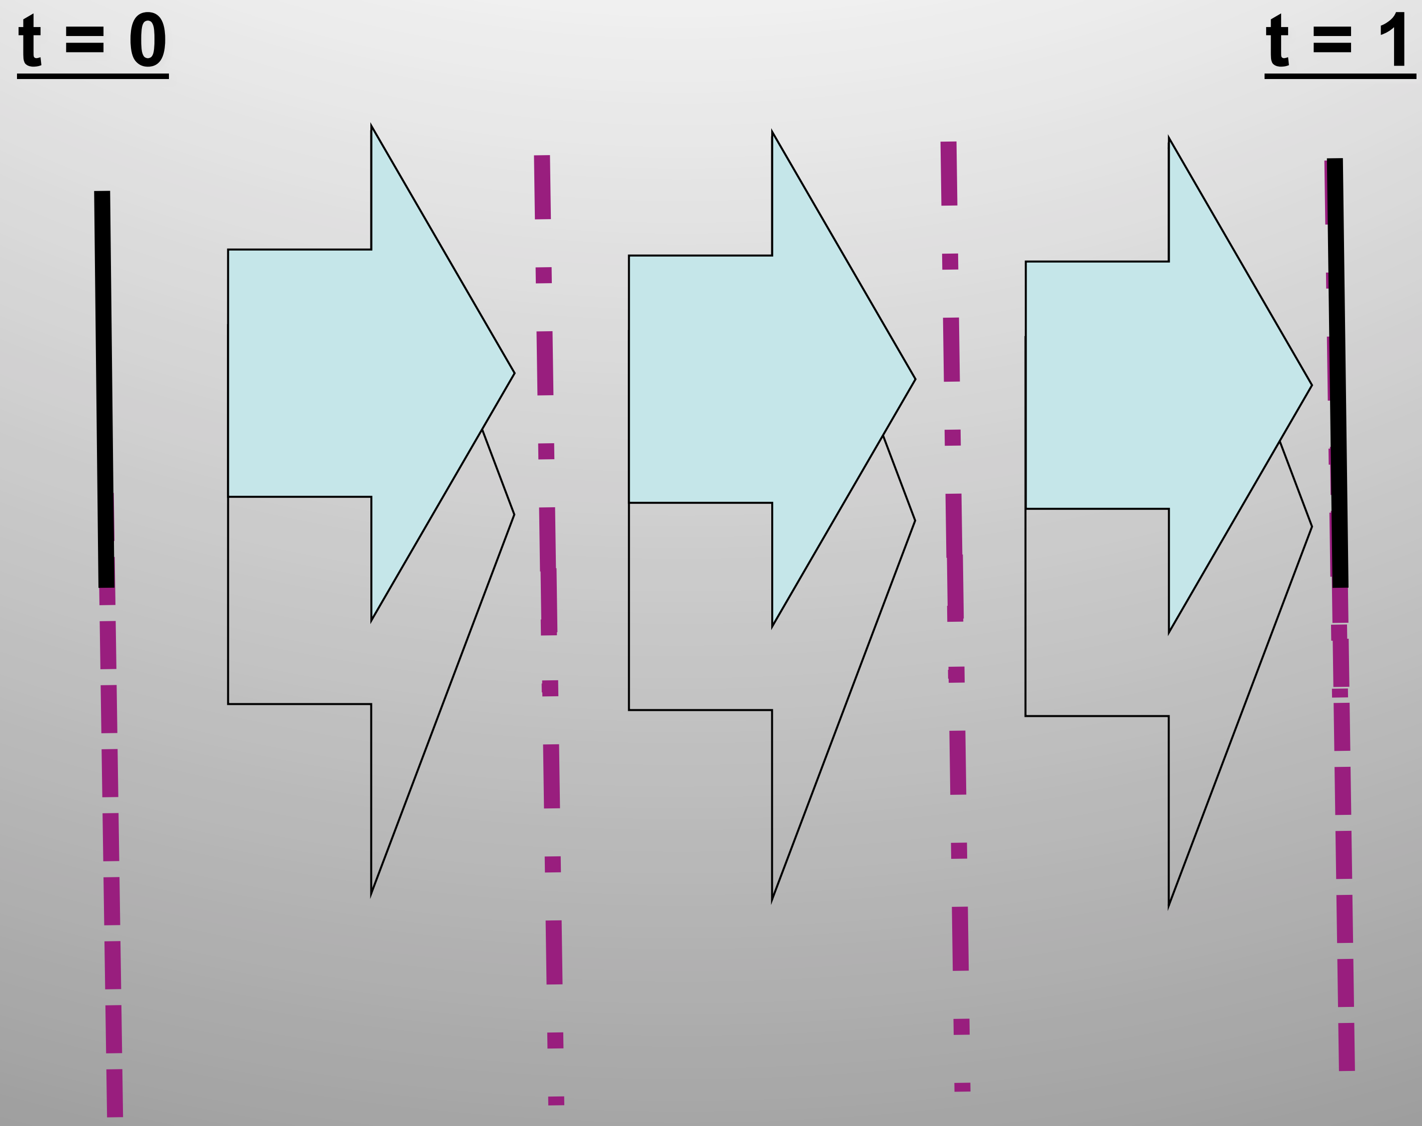

More hidden states allows fewer timesteps

ROADMAP

Overview

How to erase a bit: continually-embeddable matrices

How to flip a bit - hidden states

How to flip a bit - hidden timesteps

Tradeoff of numbers of hidden states and timesteps

PHYSICALLY IMPLEMENTING CONDITIONAL DISTRIBUTIONS

- Consider a physical system implementing a conditional distribution $\pi(x_1 \mid x_0)$ over **"visible" state space X** (X finite, time units arbitrary)

  - π observed to govern some naturally occurring system

  - π constructed by human engineers

    Example: Update function of a (re-usable) gate in a digital circuit (a single-valued conditional distribution)

    Example: Update function of an entire (re-usable) digital computer

Will mostly focus here on π that are functions (like gates)

PHYSICALLY IMPLEMENTING CONDITIONAL DISTRIBUTIONS

- Consider a physical system implementing a conditional distribution $\pi(x_1 \mid x_0)$ over a visible state space X

- "Implement" here means the dynamics for $t \in [0, 1]$ is described by a ***(time-inhomogeneous) continuous-time Markov chain (CTMC)***

  - This encompasses stochastic thermodynamics
  - However our results are even more general

- **Example**: Stochastic thermodynamic analysis of flipping a bit
  stored as state of a quantum dot.

HIDDEN STATES

- Consider a system governed by a CTMC that implements a conditional distribution $\pi(x_1 \mid x_0)$ over a "visible" state space X

We prove that for many $\boldsymbol{\pi}$'s, the CTMC must actually evolve over a space including "**hidden**" states Z, in addition to the visible states X

- More precisely:
  - For many $\pi(x_1 \mid x_0)$, ***any*** CTMC implementing $\boldsymbol{\pi}$ over X must actually evolve across some space $\mathbf{X \cup Z}$

  - $\pi(x_1 \mid x_0)$ is the restriction to X of the CTMC over $X \cup Z$

HIDDEN STATES

For many $\boldsymbol{\pi}$'s, the CTMC must actually evolve over a space including hidden states Z, in addition to the visible states X

- **Bit flip example**:
  - X = {0, 1}
  - Start in either state 0 or 1

1 → 0

0
1

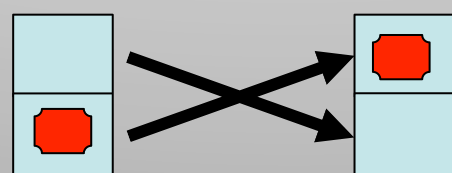

0 → 1

0
1

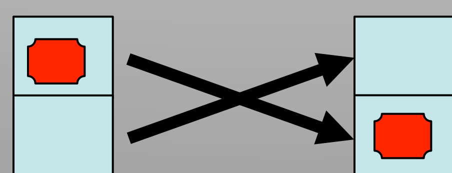

HIDDEN STATES

For many π's, the CTMC must actually evolve over a space including hidden states Z, in addition to the visible states X

- **Bit flip example**:
  - X = {0, 1}
  - Start in either state 0 or 1

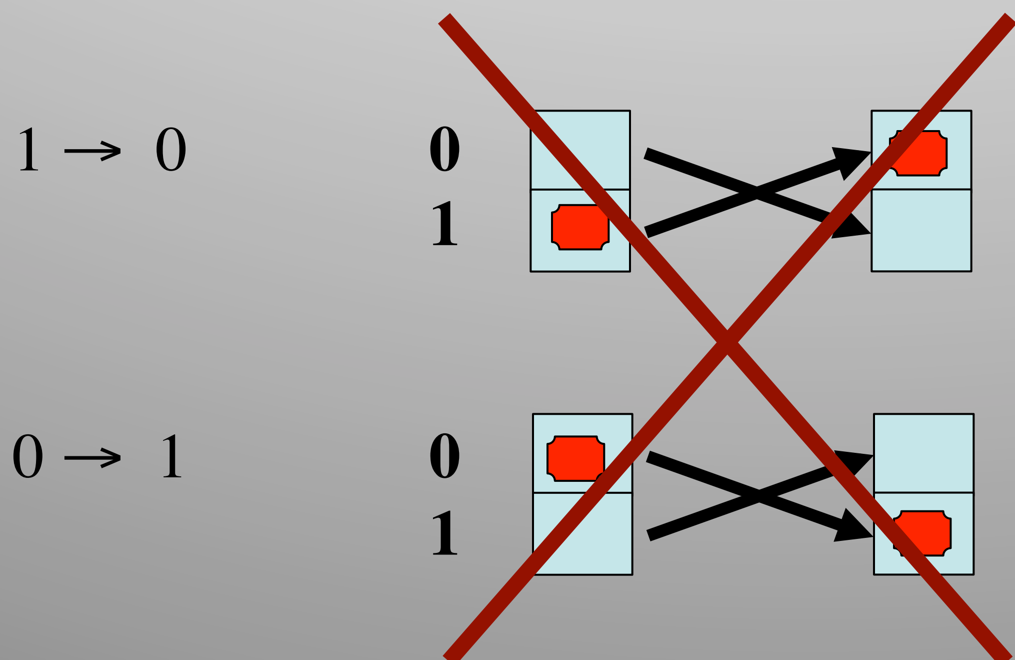

HIDDEN STATES

For many π's, the CTMC must actually evolve over a space including hidden states Z, in addition to the visible states X

- **Bit flip example**:
  - X = {0, 1}, Z = {2}
  - Start in either state 0 or 1

$1 \rightarrow 0$

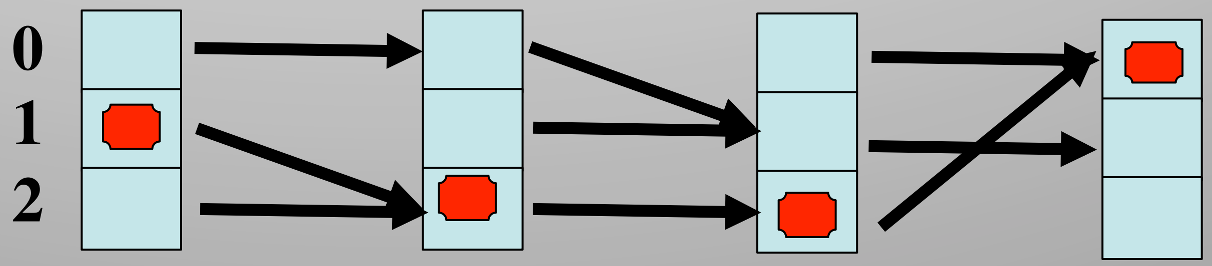

HIDDEN STATES

For many π's, the CTMC must actually evolve over a space including hidden states Z, in addition to the visible states X

- **Bit flip example**:
  - X = {0, 1}, Z = {2}
  - Start in either state 0 or 1

$1 \rightarrow 0$

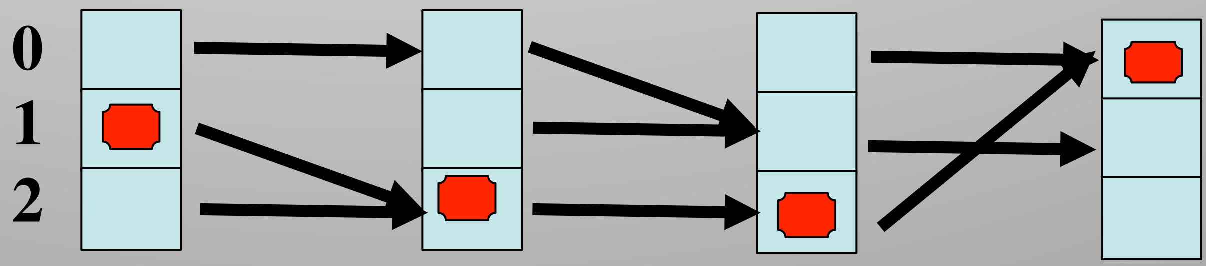

$0 \rightarrow 1$

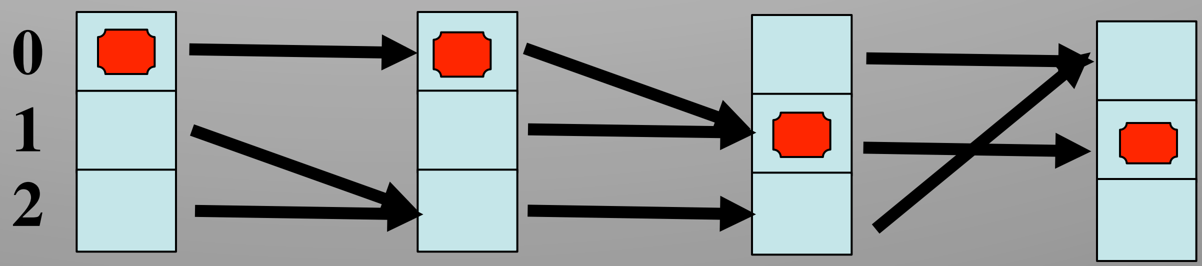

HIDDEN TIMESTEPS

- Consider a system governed by a CTMC that implements a conditional distribution $\pi(x_1 \mid x_0)$ over a visible state space X

There is a natural way to view any CTMC as dividing $t \in [0, 1]$ into a countable number of contiguous intervals

- I.e., any CTMC taking x_0 to $\pi(x_1 \mid x_0)$ runs through a sequence of "**hidden timesteps**" within [0, 1]

- Often there is a **cost** to any engineer constructing a CTMC to implement $\boldsymbol{\pi}$, which increases with the number of hidden timesteps

HIDDEN STATES

- **Example:** To implement a bit flip requires at least ***three hidden timesteps***

- **Bit flip example**:
  - $X = \{0, 1\}, Z = \{2\}$
  - Start in either state 0 or 1

$1 \rightarrow 0$

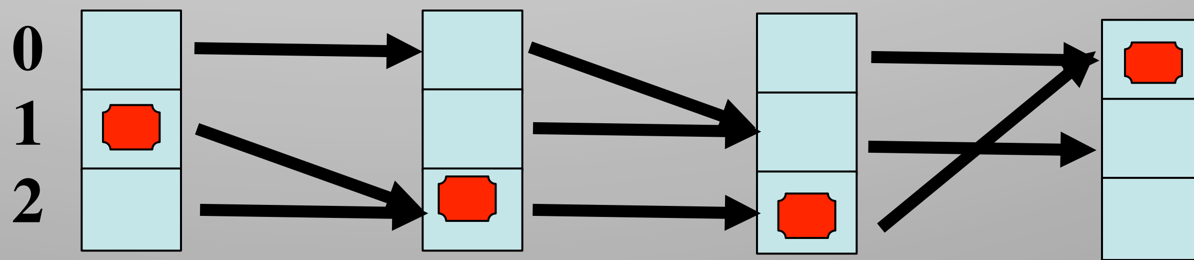

$0 \rightarrow 1$

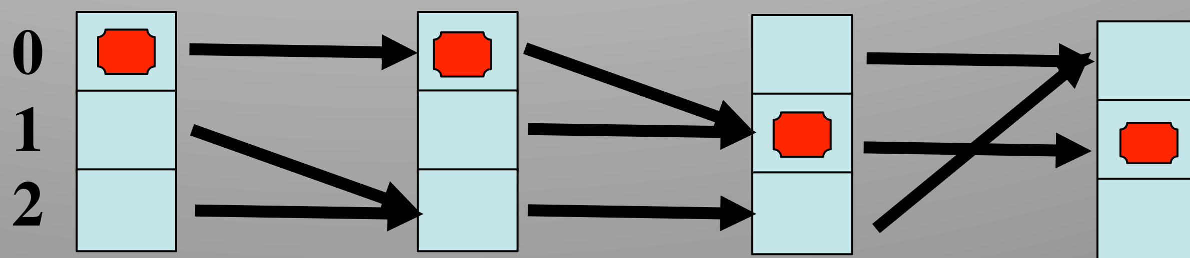

HIDDEN TIMESTEPS

- Consider a system governed by a CTMC that implements a conditional distribution $\pi(x_1 \mid x_0)$ over a visible state space X

There is a natural way to view any CTMC as dividing $t \in [0, 1]$ into a countable number of contiguous intervals

- In general, for any $\pi(x_1 \mid x_0)$, ***the more hidden states*** a CTMC can use, ***the fewer hidden timesteps*** it needs to implement π.

- I.e., a ***tradeoff*** between number of hidden states and minimal number of hidden timesteps needed to implement π with a CTMC

- This tradeoff depends on the details of π

ROADMAP

Overview

↓

How to erase a bit: continually-embeddable matrices

↓

How to flip a bit - hidden states

↓

How to flip a bit - hidden timesteps

↓

Tradeoff of numbers of hidden states and timesteps

PAST WORK ON EMBEDDING

- Given any $\pi(x_1 \mid x_0)$, the ***embedding problem*** is to determine if there is a CTMC with rate matrix $R_{x,x'}(t)$ such that $OE(R)[1] = \pi(x_1 \mid x_0)$
- First studied by Kingman (1962) who derived necessary and sufficient conditions for any $\pi(x_1 \mid x_0)$ to be embeddable ***for binary X***, by a time-homogeneous CTMC
- We still do not know necessary and sufficient conditions for larger X, even for time-homogeneous CTMCs

PAST WORK ON EMBEDDING

- Given any $\pi(x_1 \mid x_0)$, the embedding problem is to determine if there is a CTMC with rate matrix $R_{x,x'}(t)$ such that $OE(R)[1] = \pi(x_1 \mid x_0)$
- Goodman (1970) derived necessary conditions for π to be embeddable by a (time-*inhomogeneous*) CTMC for arbitrary finite X
- In particular,

If det $\pi \leq 0$, π cannot be embedded by any CTMC

- Intuition: For any time-varying $R_{x,x'}(t)$,

$$\det\ \pi = e^{\int_0^1 dt\ \mathrm{Tr} R(t)} > 0$$

- Lencastre et al. (2016) is a nice review.

CONTINUALLY-EMBEDDABLE π

If det $\pi \leq 0$, π cannot be embedded by any CTMC

- But ... bit erasure is the stochastic matrix $\begin{pmatrix} 0 & 0 \\ 1 & 1 \end{pmatrix}$

- This has determinant 0, and yet many physical systems erase bits.

???

- ***Intuitive Solution***: A "quasi-static" CTMC, that is arbitrarily close to bit erasure.

  (So the determinant of the matrix π that it implements is infinitesimal - but positive)

CONTINUALLY-EMBEDDABLE π

If $\det \pi \le 0$, π cannot be embedded by any CTMC

- Intuitive Solution: A "quasi-static" CTMC, that is arbitrarily close to bit erasure (so determinant is infinitesimal − but positive)

- Formally: π is ***continually-embeddable*** if $\exists$ sequence of CTMCs with transition matrices $\{\mathbf{T}^{(n)}(\boldsymbol{t}, \boldsymbol{t'}) : n = 1,2,\ldots\}$ such that

  1) $T(t, t')$ is continuous in t and t' for all $t, t' \in [0, 1] : t < t'$
  2) $\pi = T(0, 1)$

where for all $t, t' \in [0, 1]$, $\mathbf{T}(\boldsymbol{t}, \boldsymbol{t'}) = \mathbf{lim}_{n\to\infty}\, \mathbf{T}^{(n)}(\boldsymbol{t}, \boldsymbol{t'})$

ROADMAP

Overview

How to erase a bit: continually-embeddable matrices

How to flip a bit - hidden states

How to flip a bit - hidden timesteps

Tradeoff of numbers of hidden states and timesteps

HIDDEN STATES

Theorem: *Any* noninvertible function (like bit erasure) is continually embeddable

Theorem: *No* invertible function (except the identity) is continually-embeddable

- So ***bit-flip is not continually-embeddable***

  - Intuition: bit-flip is the matrix $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ with determinant -1.

  - This is not infinitesimally close to a positive determinant

- But many physical systems flip bits (not to mention perform more complicated invertible maps)

???

HIDDEN STATES

- **Illustration of flipping a bit:**
  - Visible states X = {0, 1}, hidden states Z = {2}
  - Start in either (visible) state 0 or 1

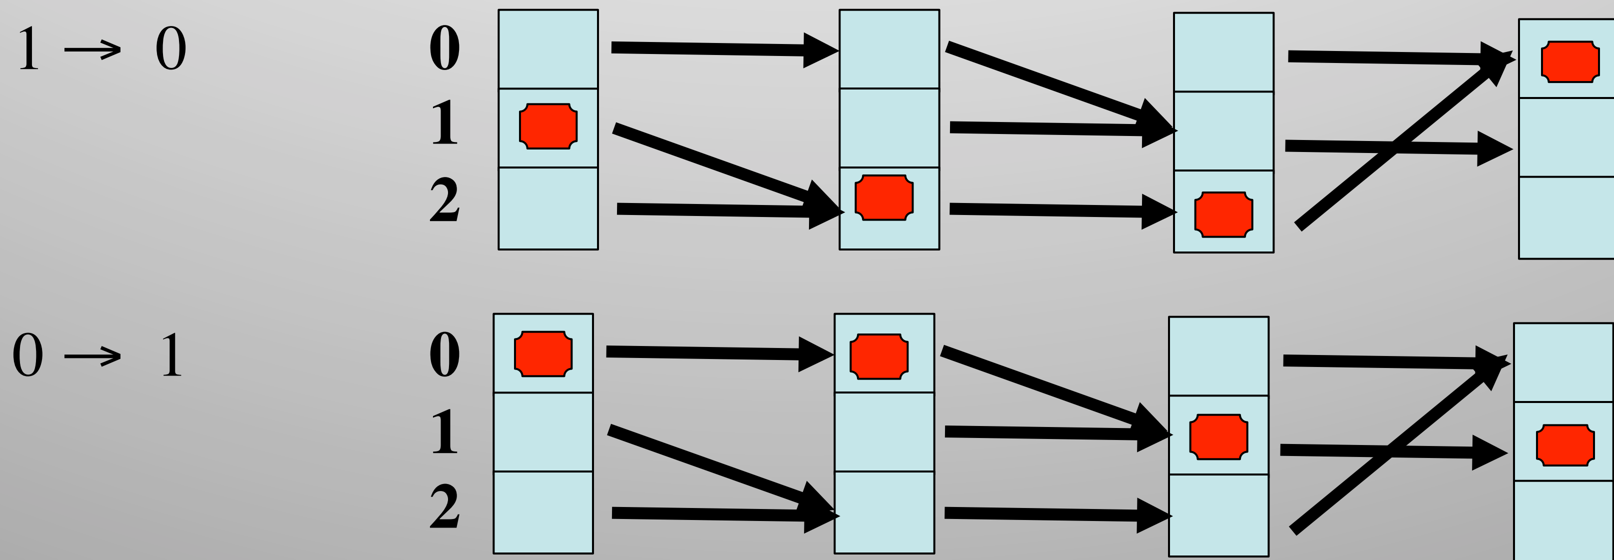

- Each step is noninvertible, and so continually-embeddable - **but over X ∪ Z**.
- ***Not*** continually-embeddable over X

ROADMAP

Overview

↓

How to erase a bit: continually-embeddable matrices

↓

How to flip a bit - hidden states

↓

How to flip a bit - hidden timesteps

↓

Tradeoff of numbers of hidden states and timesteps

HIDDEN TIMESTEPS

- Example of flipping a bit:

$1 \rightarrow 0$

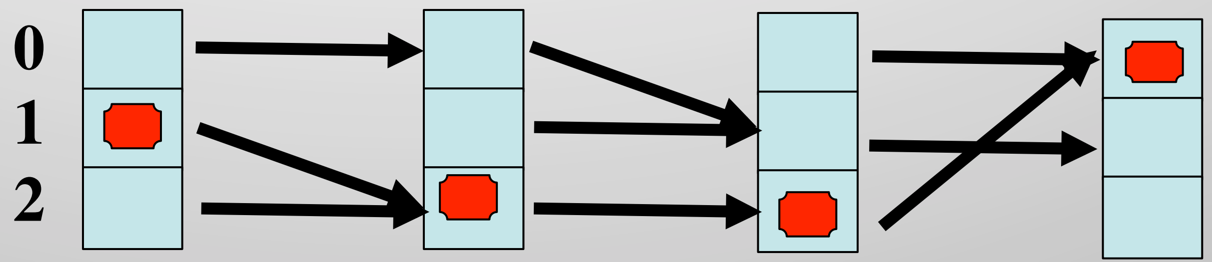

$0 \rightarrow 1$

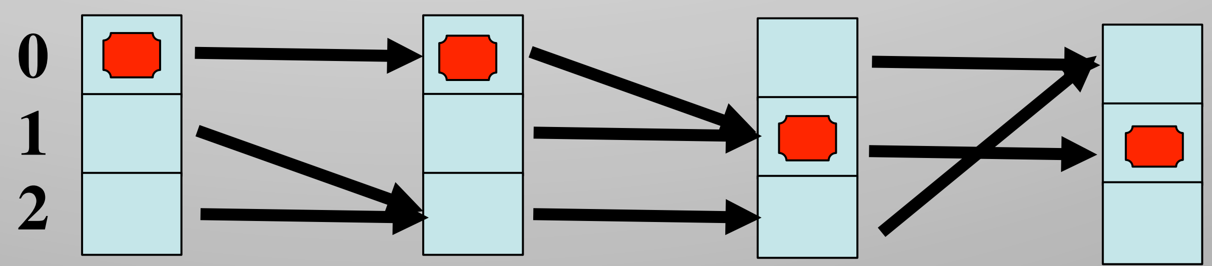

If construct a CTMC with transition matrix $T(t, t')$ to do this, you find that at transitions from column 1 to 2 and from column 2 to 3, ***a term in $T(t, t')$ changes from 0 to nonzero or vice-versa***

HIDDEN TIMESTEPS

- Example of flipping a bit:

$1 \rightarrow 0$

$0 \rightarrow 1$

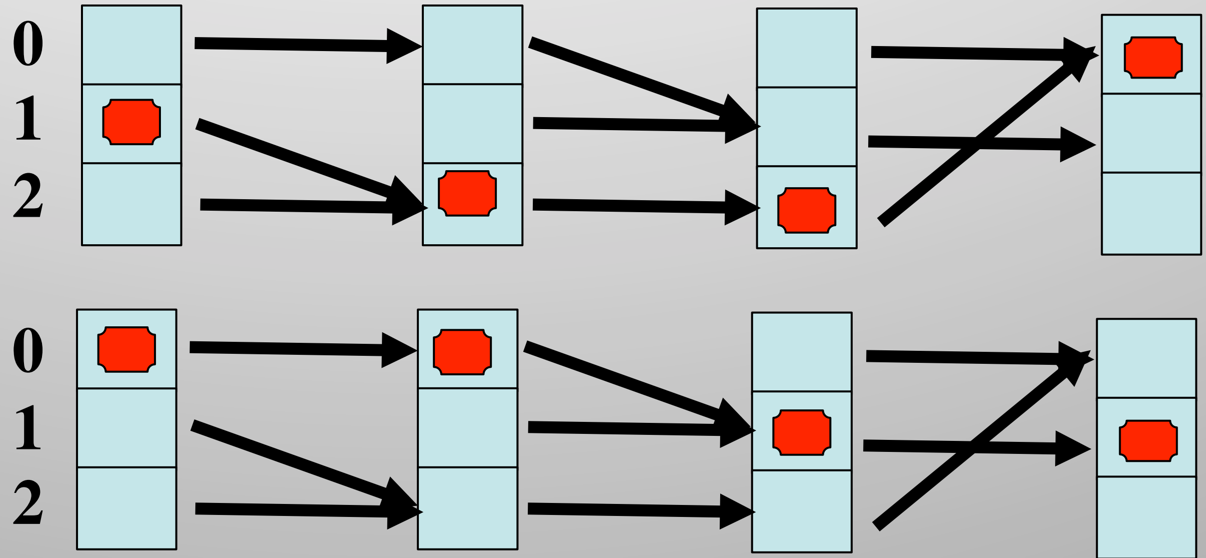

In conventional single heat-bath stochastic thermodynamics, such changes typically correspond to changing energy gaps from being ***infinite*** to being ***finite*** or vice-versa

HIDDEN TIMESTEPS

- Example of flipping a bit:

$1 \rightarrow 0$

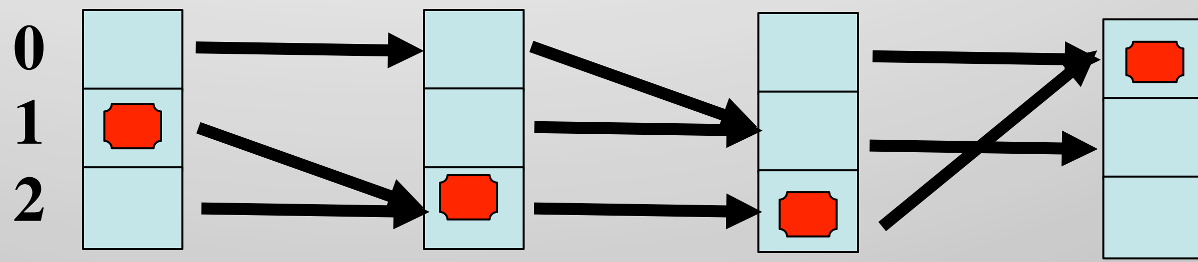

$0 \rightarrow 1$

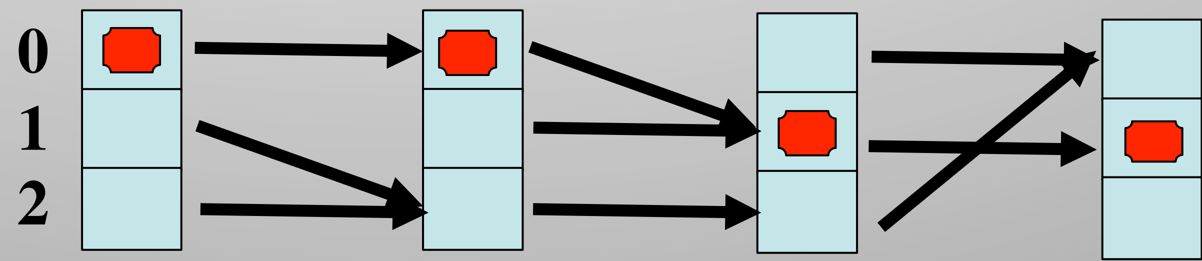

Often difficult for an engineer to construct a system whose transition matrix has terms that change from zero to nonzero or vice-versa.

Treat number of such changes as a cost

HIDDEN STATES

- **Example of flipping a bit:**
  - Visible states X = {0, 1}, hidden states Z = {2}
  - Start in either (visible) state 0 or 1

$1 \rightarrow 0$

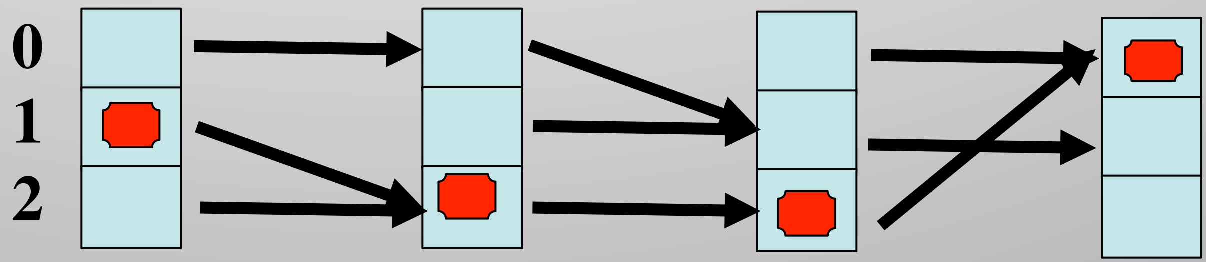

$0 \rightarrow 1$

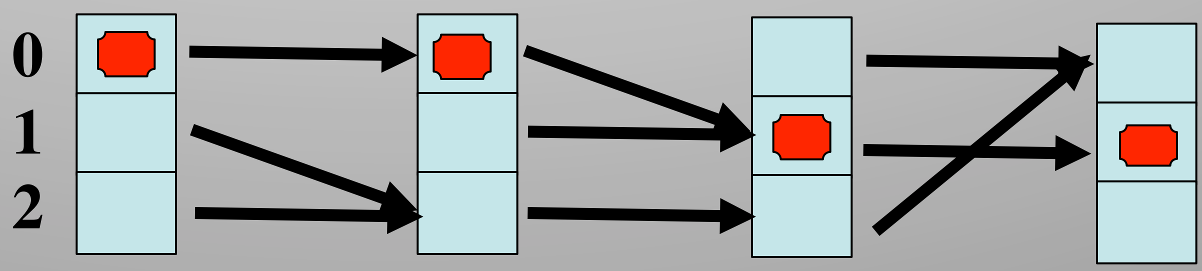

- 3 successive idempotent functions
  - Is that fewest possible? I.e., does ***timestep cost = 3?***

ROADMAP

Overview

How to erase a bit: continually-embeddable matrices

How to flip a bit - hidden states

How to flip a bit - hidden timesteps

Tradeoff of numbers of hidden states and timesteps

TIMESTEP COST

Theorem: Timestep cost of noninvertible π is minimal number of idempotent functions whose product is π

- Analyzing a formally identical semigroup theory question, Saito (1989) showed that this cost is either

$$\left\lceil \frac{|X| + \text{cycl}\ \pi - \text{fix}\ \pi}{|X| - |\text{img}\ \pi|} \right\rceil$$

or **1** more than this, where*:*

- cycl(π) is number of (invertible) cyclic orbits of π
- fix(π) is number of fixed points of π
- img(π) is size of image of X under π

TRADEOFF OF HIDDEN STATES AND TIMESTEPS

Theorem: Timestep cost of noninvertible π is minimal number of idempotent functions whose product is π

- Simple extension of this result to allow ***k hidden states,*** and include ***invertible π*** : Timestep cost is either

$$\left\lceil \frac{k + |X| + \max\big(\mathrm{cycl}(\pi) - k, 0\big) - \mathrm{fix}(\pi)}{k + |X| - |\mathrm{img}(\pi)|} \right\rceil$$

or 1 more than this

EXAMPLE: BIT FLIP

- Timestep cost of π with k hidden states:

$$\left\lceil \frac{k + |X| + \max\big(\mathrm{cycl}(\pi) - k, 0\big) - \mathrm{fix}(\pi)}{k + |X| - |\mathrm{img}(\pi)|} \right\rceil$$

$1 \rightarrow 0$

$0 \rightarrow 1$

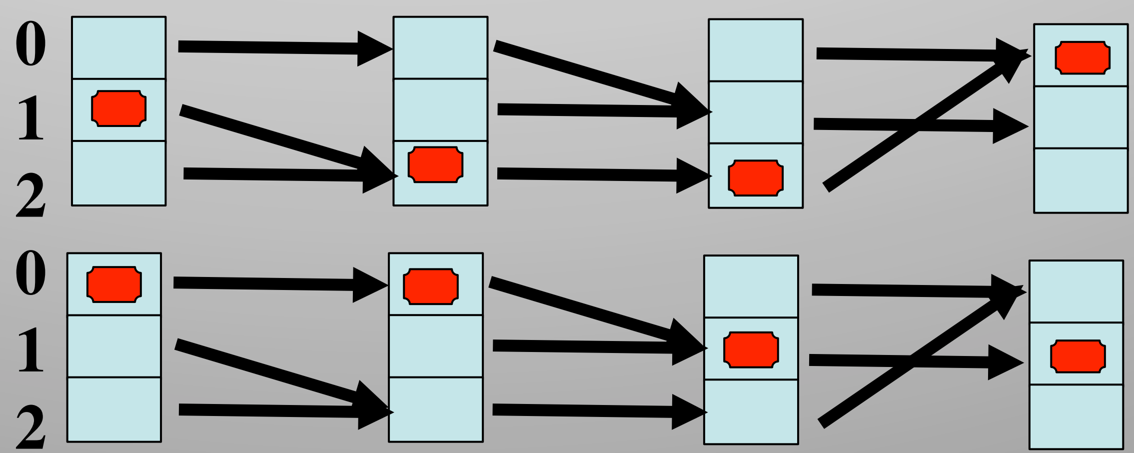

- |X| = 2, cycl(π) = 1, fix(π) = 0, img(π) = 2, k = 1
- So timestep cost = 3 – three successive steps is smallest possible.

EXAMPLE – Maps over 4 bits

- Space/time trade-off for two functions over X = {0, . . . , 15}: -
  - '**Cycle**' is x → x + 1 mod 16.
  - '**Complement**' represents each element of X as a four-bit string and then applies bitwise NOT.

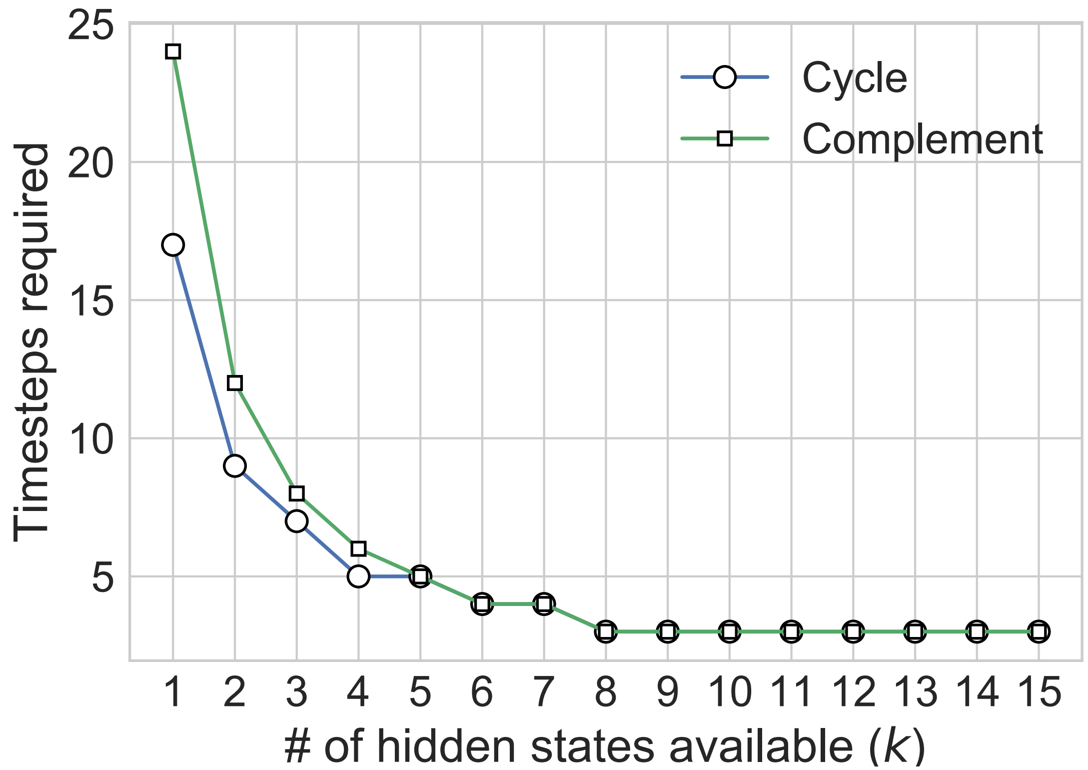

REALISTIC SETS OF IDEMPOTENT FUNCTIONS

- Analysis so far assumes can use arbitrary idempotent functions
- In real world, ***severe constraints on set of idempotent functions*** we can build into our devices
- **Ex**: X is all bit strings of length 128
  - Number of possible idempotent functions lower-bounded by the number of partitions of X, i.e., the Bell number of 2^{128}
  - This is ***huge*** − so results above, which assume we can use all those functions are not appropriate for such an X

How does analysis change with realistic constraints on set of idempotent functions we can use?

REALISTIC SETS OF IDEMPOTENT FUNCTIONS

Ex: X is all bit strings of length 128

- Suppose only ***two types of idempotent function*** we can use:
  - Functions that work on one spin (bit) at a time
  - Functions that work on two spins (bits) at a time

- The set of all such functions includes all logical ANDs, NOTs or ORs of individual bits

- So can implement any Boolean function of x by a sequence of such idempotent functions

- Calculating timestep cost with k hidden states similar to circuit complexity, ***but different***

CONCLUSIONS

- Derived a **novel "hidden" space/time trade-off** applicable to all continuous-time Markov chains

- **Physical meaningful** as minimal "costs" of any stochastic thermodynamic process that implements a given function

- Unlike traditional costs in thermodynamics of computation, these new ones involve ***state-space resources*** and ***timestep resources***

- Can extend to non-single-valued ("stochastic") π (another talk).

- Space / time tradeoffs of **a single gate** within an overall circuit of many gates... that is *itself* subject to space / time tradeoffs...

- Lots of future work!