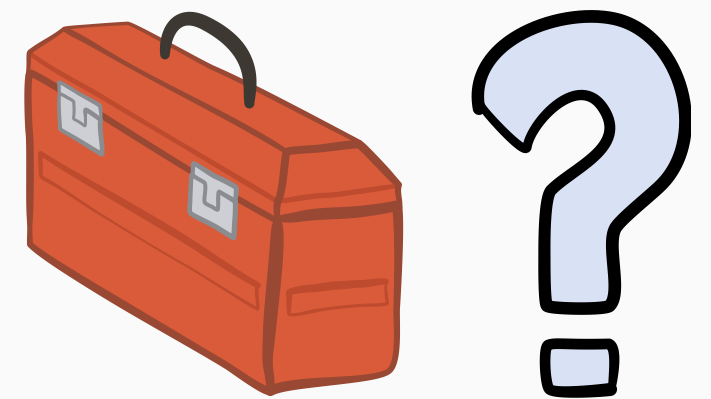


The Role of **Advanced Math**
in Teaching **Performance Modeling**

Ziv Scully
Cornell ORIE

TeaPACS 2023



Goals of performance modeling

Goals of performance modeling

Analyzing systems

Goals of performance modeling

Analyzing systems



Designing and optimizing systems

Goals of performance modeling

Analyzing systems

Describing systems



Designing and optimizing systems

Goals of performance modeling

Analyzing systems

Describing systems



Designing and optimizing systems



What role does
math play?

Goals of performance modeling

- exact model analysis
- guide simulation
- what to measure?

Analyzing systems

Describing systems



Designing and optimizing systems



What role does
math play?

Goals of performance modeling

- exact model analysis
- guide simulation
- what to measure?


Analyzing systems

Describing systems



Designing and optimizing systems

- find optimal policies
- evaluate heuristics
- what to optimize?

 What role does math play?

Goals of performance modeling

- exact model analysis
- guide simulation
- what to measure?

Analyzing systems




- stochastic modeling
- define load, stability
- what is predictable?

Describing systems



Designing and optimizing systems

- find optimal policies
- evaluate heuristics
- what to optimize?

 What role does math play?

Goals of performance modeling

- exact model analysis
- guide simulation
- what to measure?

Analyzing systems

- stochastic modeling
- define load, stability
- what is predictable?

Describing systems



Designing and optimizing systems

- find optimal policies
- evaluate heuristics
- what to optimize?

How much math?



What role does math play?

**Performance modeling
needs advanced math**

**Performance modeling
needs advanced math**

**We can teach advanced
math accessibly**

Part 1

**Performance modeling
needs advanced math**

Part 2

**We can teach advanced
math accessibly**

Part 1


**Performance modeling
needs advanced math**

Part 2

**We can teach advanced
math accessibly**

What do jobs look like?


What do jobs look like?




[Tirmazi et al., 2020]

Heavy tails are ubiquitous

What do jobs look like?



[Tirmazi et al., 2020]




[Zhao et al., 2022]

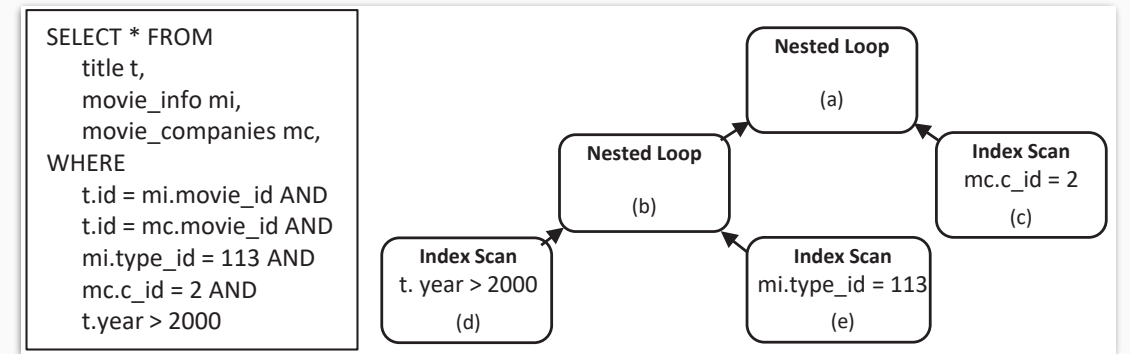
Jobs have complex structure

Heavy tails are ubiquitous

What do jobs look like?




[Tirmazi et al., 2020]



[Zhao et al., 2022]


Jobs have complex structure

Heavy tails are ubiquitous

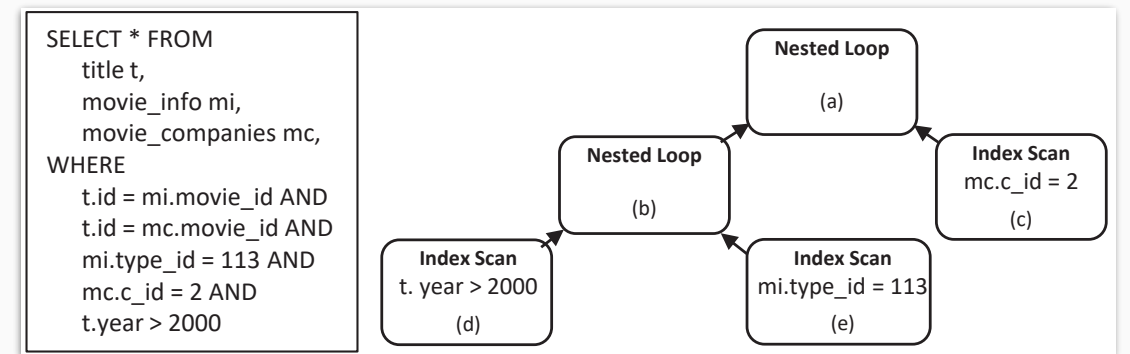


Finite-state Markov chains aren't enough

What do jobs look like?




[Tirmazi et al., 2020]




[Zhao et al., 2022]

Jobs have complex structure




Age and remaining work
aren't enough

Heavy tails are ubiquitous




Finite-state Markov chains
aren't enough

What do jobs look like?




[Tirmazi et al., 2020]




[Zhao et al., 2022]

Jobs have complex structure



Age and remaining work aren't enough

Heavy tails are ubiquitous



Finite-state Markov chains aren't enough

Need: general Markov processes

Stability in complex systems

Stability in complex systems




Metastable
failures

Stability in complex systems



Metastable failures




[Huang et al., 2020]


Stability in complex systems



Metastable failures



[Huang et al., 2020]




[Bronson et al., 2020]


Stability in complex systems



Metastable failures



[Huang et al., 2020]



[Bronson et al., 2020]

Need: drift methods, mean field methods

Scheduling practicalities

Scheduling practicalities

Theory: SRPT

Scheduling practicalities



SRPT in
networks

Theory: SRPT

Practice: Homa

[Montazeri et al., 2020]


Scheduling practicalities

SRPT in
networks

Theory: SRPT

Practice: Homa

[Montazeri et al., 2020]




Continuous priority,
no overhead/delay

Scheduling practicalities

SRPT in networks


Theory: SRPT



Continuous priority,
no overhead/delay

Practice: Homa

[Montazeri et al., 2020]

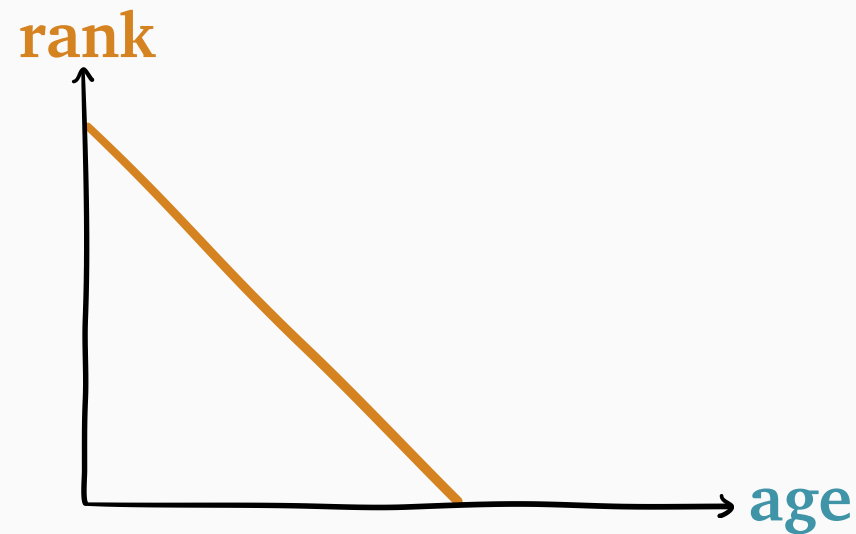


Discrete priorities,
overheads/delays

Scheduling practicalities

SRPT in networks


Theory: SRPT



Continuous priority,
no overhead/delay

Practice: Homa

[Montazeri et al., 2020]

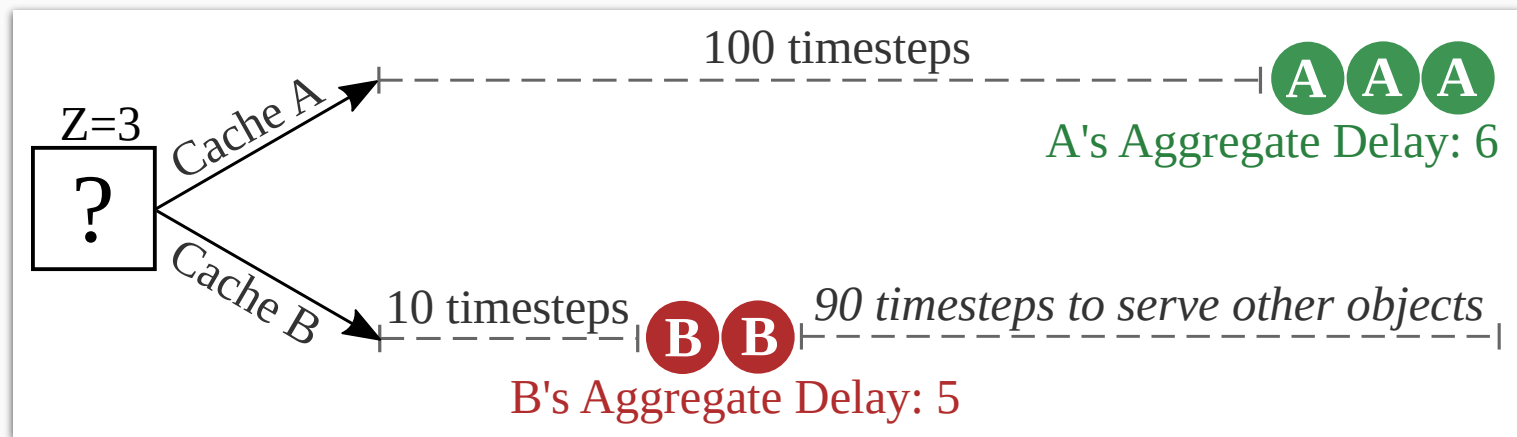


Discrete priorities,
overheads/delays

**Need: analyze variety
of scheduling policies**

What should we measure?

What should we measure?



[Atre et al., 2020]

What should we measure?



[Atre et al., 2020]

Algorithm 1 Estimating AggregateDelay

```
1: struct OBJECTMETADATA
2:   NumWindows = 0
3:   CumulativeDelay = 0
4:   WindowStartIdx =  $-\infty$ 
5:
6: function ESTIMATEAGGREGATEDELAY(X: OBJECTMETADATA)
7:   return  $\frac{X.CumulativeDelay}{X.NumWindows}$ 
8: end function
9:
10: function ONACCESS(TimeIdx, X: OBJECTMETADATA)
11:   // Time since start of the previous miss window
12:   TSSW = (TimeIdx - X.WindowStartIdx)
13:
14:   if TSSW  $\geq$  Z then
15:     // This access commences a new miss window
16:     X.NumWindows += 1
17:     X.CumulativeDelay += Z
18:     X.WindowStartIdx = TimeIdx
19:   else
20:     // This access is part of the previous miss window
21:     X.CumulativeDelay += (Z - TSSW)
22:   end if
23: end function
```

[Atre et al., 2020]

What should we measure?



[Atre et al., 2020]

Algorithm 1 Estimating AggregateDelay

```
1: struct OBJECTMETADATA
2:   NumWindows = 0
3:   CumulativeDelay = 0
4:   WindowStartIdx =  $-\infty$ 
5:
6: function ESTIMATEAGGREGATEDELAY(X: OBJECTMETADATA)
7:   return  $\frac{X.CumulativeDelay}{X.NumWindows}$ 
8: end function
9:
10: function ONACCESS(TimeIdx, X: OBJECTMETADATA)
11:   // Time since start of the previous miss window
12:   TSSW = (TimeIdx - X.WindowStartIdx)
13:
14:   if TSSW  $\geq$  Z then
15:     // This access commences a new miss window
16:     X.NumWindows += 1
17:     X.CumulativeDelay += Z
18:     X.WindowStartIdx = TimeIdx
19:   else
20:     // This access is part of the previous miss window
21:     X.CumulativeDelay += (Z - TSSW)
22:   end if
23: end function
```

[Atre et al., 2020]

Need: expectations from different perspectives

Part 1

**Performance modeling
needs advanced math**

Part 2

**We can teach advanced
math accessibly**

Part 1

Performance modeling
needs advanced math

Part 2


**We can teach advanced
math accessibly**

Part 1

Performance modeling
needs advanced math

Part 2

**We can teach advanced
math accessibly**




Simplify core
foundations

Part 1


Performance modeling
needs advanced math

Part 2

**We can teach advanced
math accessibly**



Simplify core
foundations




Prioritize very
flexible tools



Problem: many students lack math background




Problem: many students lack math background



Solution: hand-wave



Problem: many students lack math background




Solution: hand-wave



Problem: how to know when to hand-wave?




Problem: many students lack math background



Solution: hand-wave




Problem: how to know when to hand-wave?



Solution: clear rules for hand-waving




Problem: many students lack math background



Solution: hand-wave



Problem: how to know when to hand-wave?




Solution: clear rules for hand-waving

- *Principles:* rules that work most of the time




Problem: many students lack math background



Solution: hand-wave



Problem: how to know when to hand-wave?




Solution: clear rules for hand-waving

- *Principles:* rules that work most of the time
- *Recipes:* common patterns for using principles




Problem: many students lack math background



Solution: hand-wave



Problem: how to know when to hand-wave?



Solution: clear rules for hand-waving


- *Principles:* rules that work most of the time
- *Recipes:* common patterns for using principles



Problem: each topic needs many principles




Problem: many students lack math background



Solution: hand-wave



Problem: how to know when to hand-wave?




Solution: clear rules for hand-waving

- *Principles:* rules that work most of the time
- *Recipes:* common patterns for using principles



Problem: each topic needs many principles



Solution: focus on a few very powerful topics

Proposed toolbox



Proposed toolbox



Description: model with *Markov processes*

Proposed toolbox



Description: model with *Markov processes*

Metrics: define using *long-run averages*

Proposed toolbox



Description: model with *Markov processes*

Metrics: define using *long-run averages*

Analysis: reduce to questions about *drift*

Description via Markov processes

Description via Markov processes

State: all info we need to describe evolution

Description via Markov processes

State: all info we need to describe evolution

current state

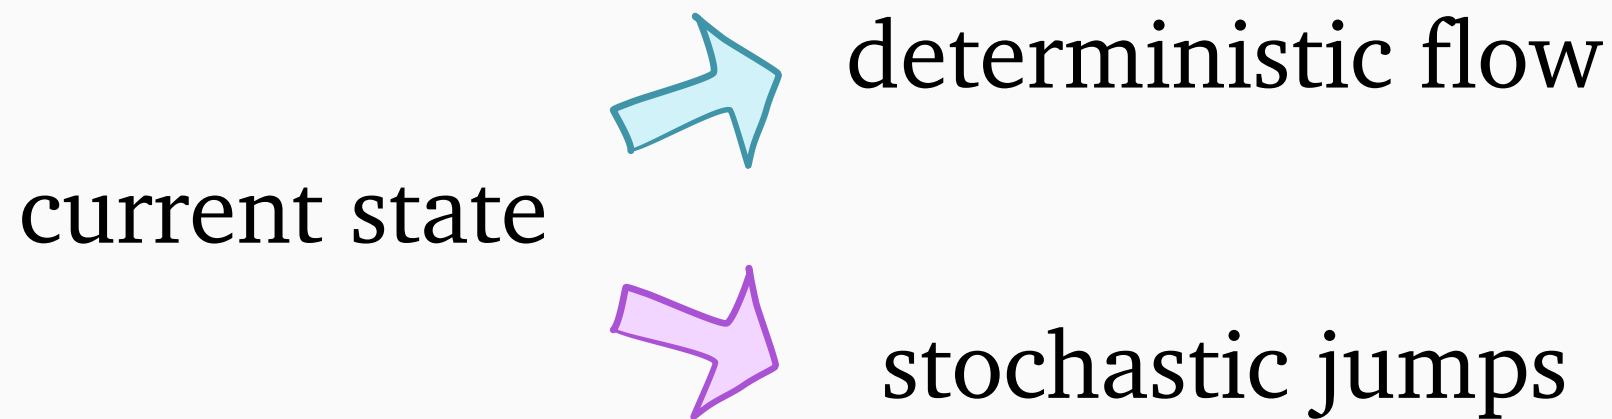
Description via Markov processes

State: all info we need to describe evolution

current state  deterministic flow

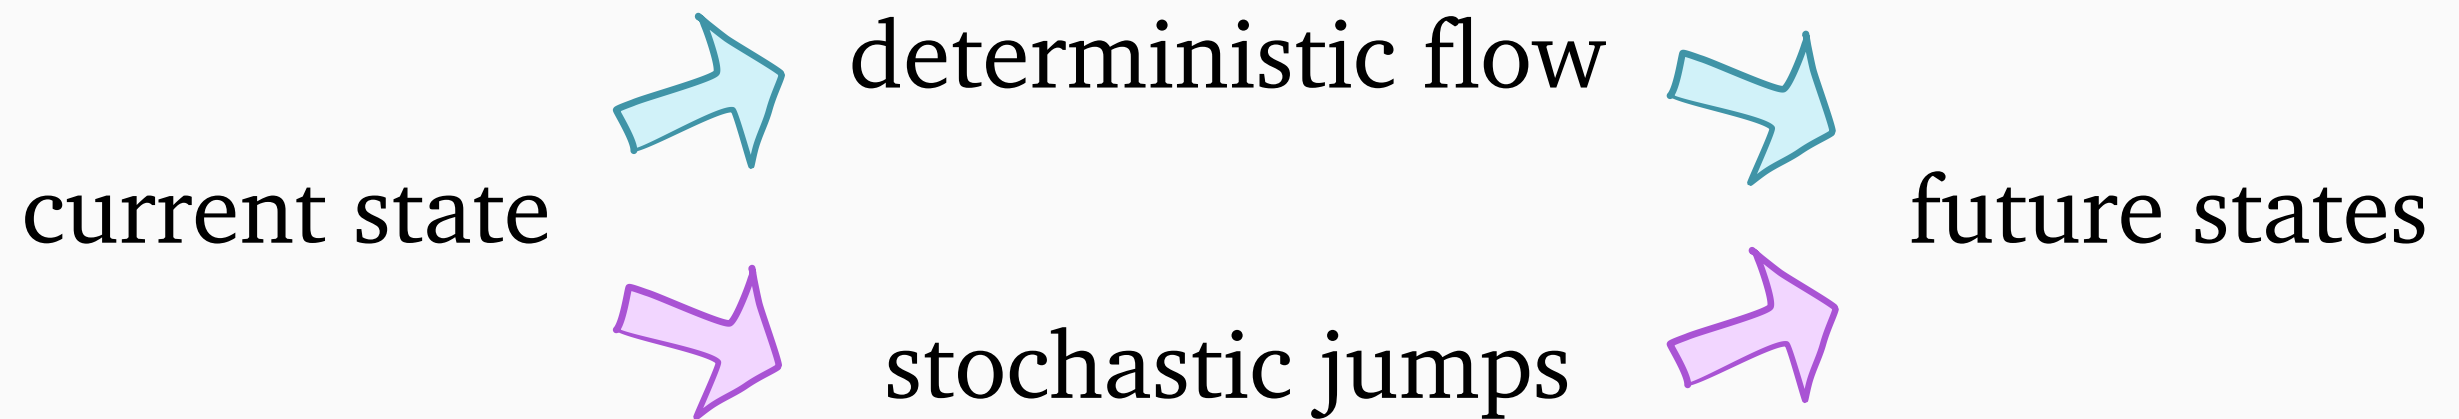
Description via Markov processes

State: all info we need to describe evolution



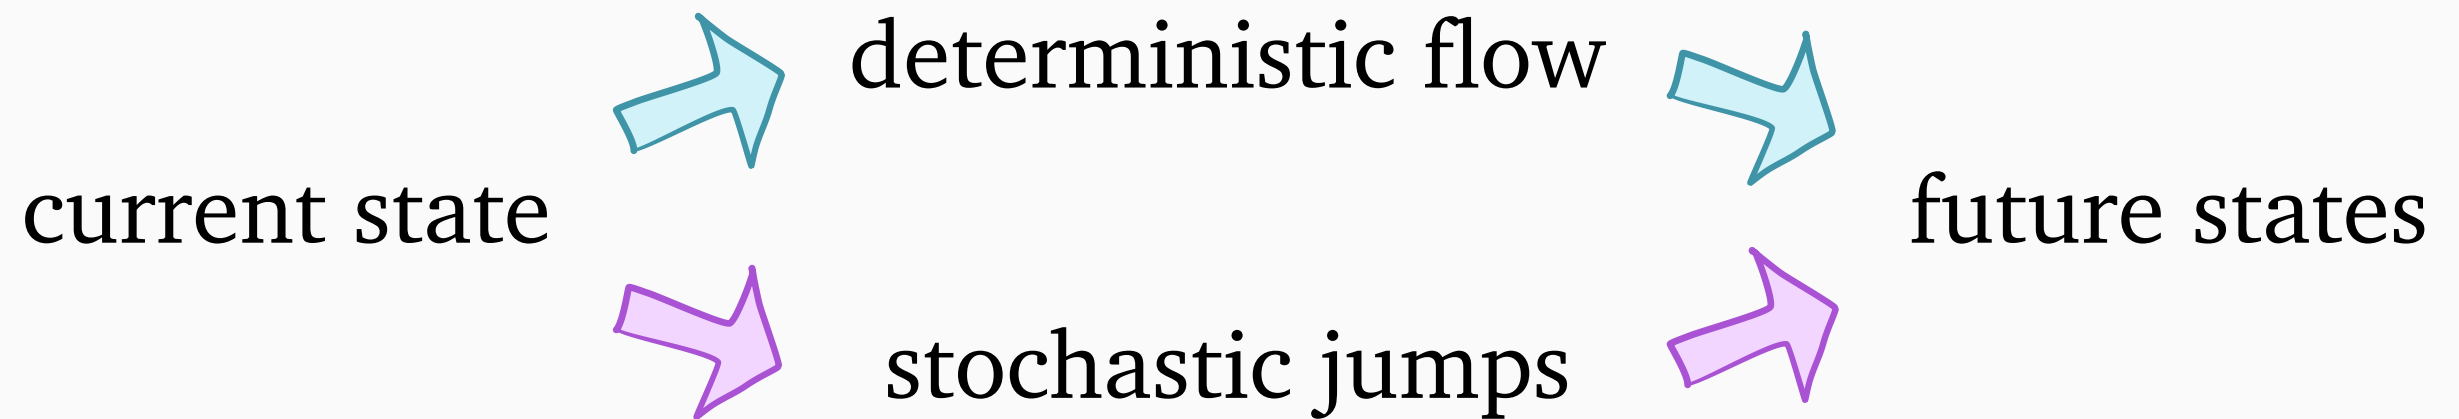
Description via Markov processes

State: all info we need to describe evolution



Description via Markov processes

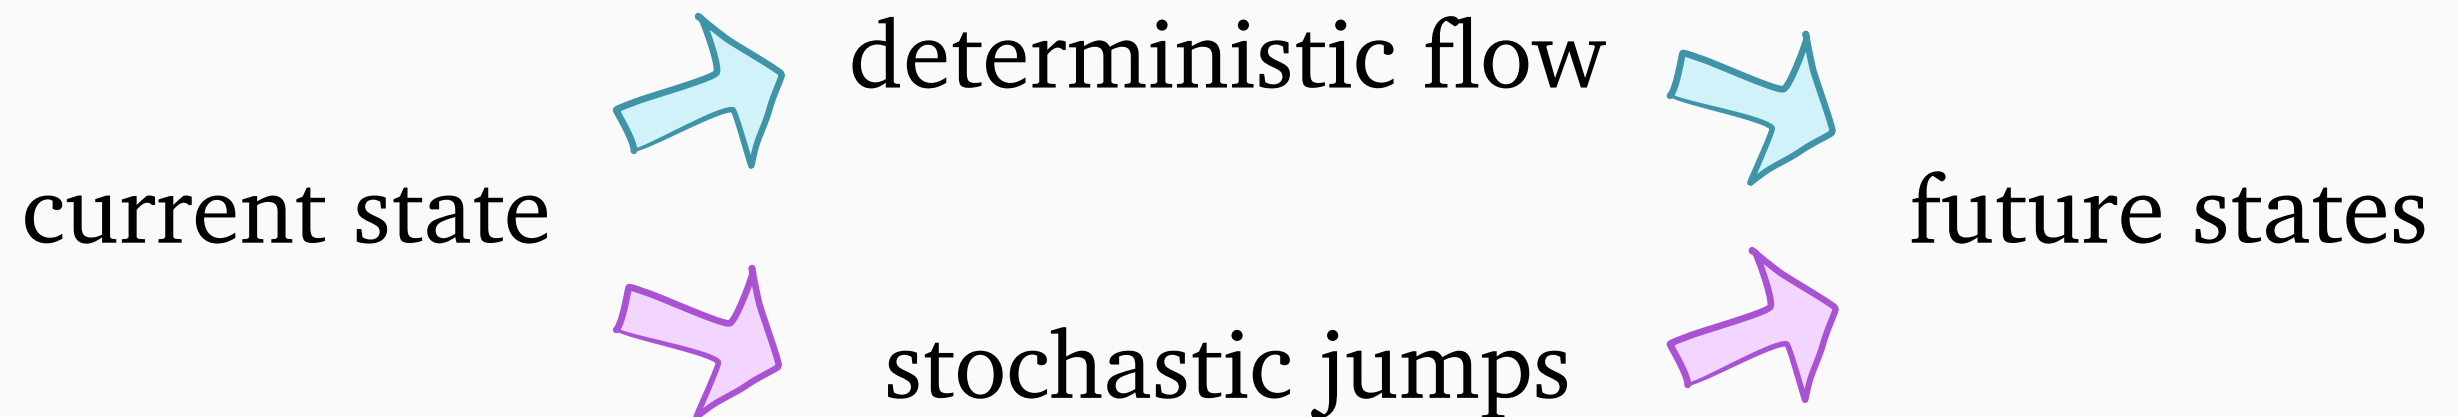
State: all info we need to describe evolution



Goal: clear process definition

Description via Markov processes

State: all info we need to describe evolution

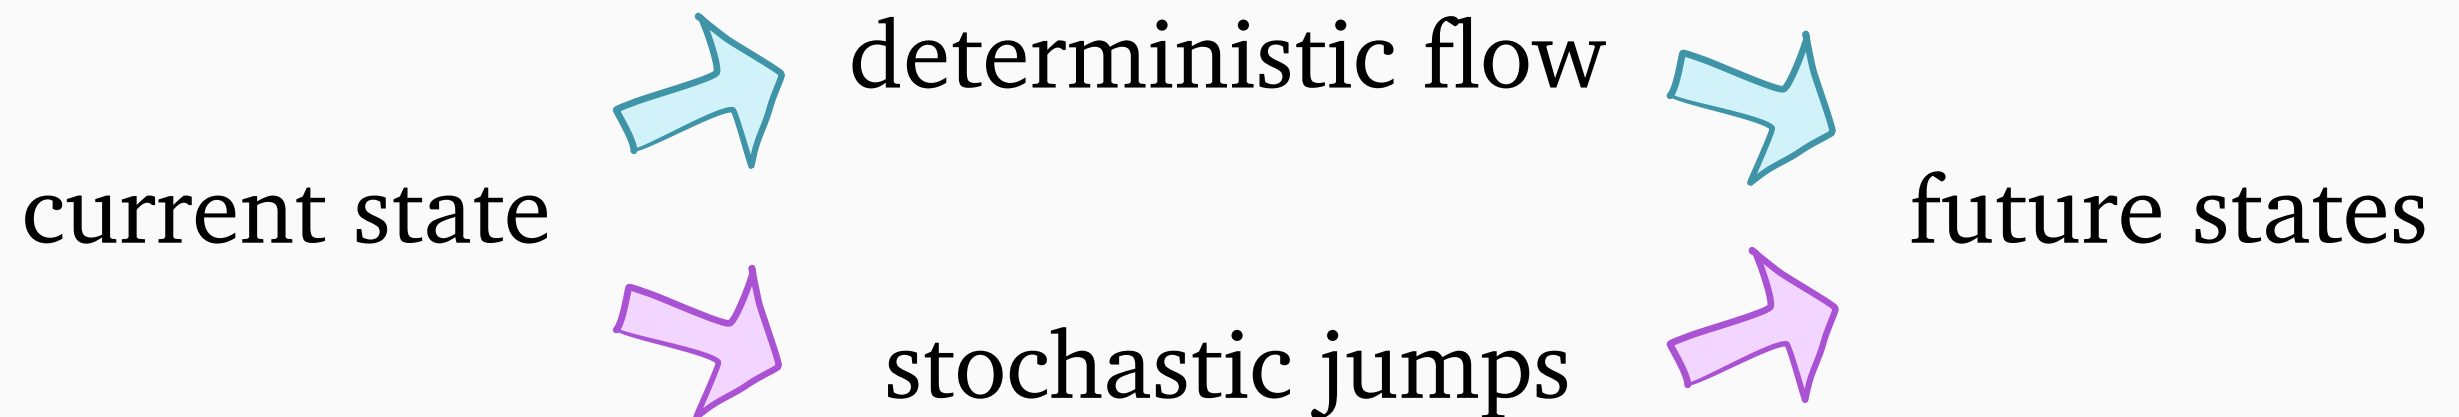


Goal: clear process definition

Non-goal (yet): tractable analysis

Description via Markov processes

State: all info we need to describe evolution



Goal: clear process definition

Non-goal (yet): tractable analysis

Non-goal: verifying Markov property

Example: M/G/1

Example: M/G/1

State: list with remaining work of each job

$$[r_1, \dots, r_n]$$

Example: M/G/1

State: list with remaining work of each job

$$[r_1, \dots, r_n]$$

Dynamics:

Example: M/G/1

State: list with remaining work of each job

$$[r_1, \dots, r_n]$$

Dynamics:

- If list nonempty: decrease r_1 at rate 1

Example: M/G/1

State: list with remaining work of each job

$$[r_1, \dots, r_n]$$

Dynamics:

- If list nonempty: decrease r_1 at rate 1
- When $r_1 = 0$: remove r_1 from list

Example: M/G/1

State: list with remaining work of each job

$$[r_1, \dots, r_n]$$

Dynamics:

- If list nonempty: decrease r_1 at rate 1
- When $r_1 = 0$: remove r_1 from list
- Poisson(λ): draw from S , append it to list

Example: M/G/1

State: list with remaining work of each job

$$[r_1, \dots, r_n]$$

Dynamics:

- If list nonempty: decrease r_1 at rate 1
- When $r_1 = 0$: remove r_1 from list
- Poisson(λ): draw from S , append it to list

$$\text{Work: } w([r_1, \dots, r_n]) = r_1 + \dots + r_n$$

Example: M/G/1

State: list with remaining work of each job

$$[r_1, \dots, r_n]$$

Dynamics:

- If list nonempty: decrease r_1 at rate 1
- When $r_1 = 0$: remove r_1 from list
- Poisson(λ): draw from S , append it to list

$$\text{Work: } w([r_1, \dots, r_n]) = r_1 + \dots + r_n$$

$$\text{Queue length: } q([r_1, \dots, r_n]) = (n - 1)^+$$

Metrics via long-run averages

$X(t)$ = state at time t

mean waiting time = $\mathbf{E}_{\text{arrival}}[w(X)]$



mean number in queue = $\mathbf{E}_{\text{time}}[q(X)]$

Metrics via long-run averages

$X(t)$ = state at time t

mean waiting time = $\mathbf{E}_{\text{arrival}}[w(X)]$

mean number in queue = $\mathbf{E}_{\text{time}}[q(X)]$





Metrics via long-run averages

$X(t)$ = state at time t

mean waiting time = $\mathbf{E}_{\text{arrival}}[w(X)]$

mean number in queue = $\mathbf{E}_{\text{time}}[q(X)]$




$$\mathbf{E}_{\text{arrival}}[f(X)] = \frac{\sum_{t \text{ arrival}} f(X(t))}{\# \text{ arrivals}}$$

Metrics via long-run averages


$X(t)$ = state at time t

mean waiting time = $\mathbf{E}_{\text{arrival}}[w(X)]$

mean number in queue = $\mathbf{E}_{\text{time}}[q(X)]$



$$\mathbf{E}_{\text{arrival}}[f(X)] = \frac{\sum_{t \text{ arrival}} f(X(t))}{\# \text{ arrivals}}$$



$$\mathbf{E}_{\text{time}}[f(X)] = \frac{\int_0^{\text{long time}} f(X(t)) dt}{\text{long time}}$$


Principles for long-run averages

Principles for long-run averages

Base principle: when averaging over entire timeline,
ignore edge effects

Principles for long-run averages

Base principle: when averaging over entire timeline,
ignore edge effects




Little's law

Principles for long-run averages

Base principle: when averaging over entire timeline,
ignore edge effects




Little's law




Renewal-reward

Principles for long-run averages


Base principle: when averaging over entire timeline,
ignore edge effects



Little's law




Renewal-reward




Palm inversion

Principles for long-run averages


Base principle: when averaging over entire timeline,
ignore edge effects



Little's law



Renewal-reward



Palm inversion

Rate conservation law:




for any f , average rate of change in $f(X)$ is 0

Principles for long-run averages

Base principle: when averaging over entire timeline,

ignore edge effects

requires
stability!



Little's law

Renewal-reward

Palm inversion

Rate conservation law:

for any f , average rate of change in $f(X)$ is 0

Analysis via drift

Dynamics:

- If list nonempty: decrease head at rate 1
- When head = 0: remove head of list
- Poisson(λ): draw from S , append it to list

Analysis via drift

Dynamics:

- If list nonempty: decrease head at rate 1
- When head = 0: remove head of list
- Poisson(λ): draw from S , append it to list



$$\begin{aligned} 0 = & \mathbf{E}_{\text{time}} \left[\frac{\partial}{\partial r_1} f(X) \right] \\ & + \lambda \mathbf{E}_{\text{departure}} \left[f(\text{tail}(X)) - f(X) \right] \\ & + \lambda \mathbf{E}_{\text{arrival}} \left[f(\text{join}(X, [S])) - f(X) \right] \end{aligned}$$

Analysis via drift

Dynamics:

- If list nonempty: decrease head at rate 1
- When head = 0: remove head of list
- Poisson(λ): draw from S , append it to list



$$\begin{aligned} 0 = & \mathbf{E}_{\text{time}} \left[\frac{\partial}{\partial r_1} f(X) \right] \\ & + \lambda \mathbf{E}_{\text{departure}} \left[f(\text{tail}(X)) - f(X) \right] \\ & + \lambda \mathbf{E}_{\text{arrival}} \left[f(\text{join}(X, [S])) - f(X) \right] \end{aligned}$$

Analysis via drift

Dynamics:

- If list nonempty: decrease head at rate 1
- When head = 0: remove head of list
- Poisson(λ): draw from S , append it to list



$$\begin{aligned} 0 = & \mathbf{E}_{\text{time}} \left[\frac{\partial}{\partial r_1} f(X) \right] \\ & + \lambda \mathbf{E}_{\text{departure}} \left[f(\text{tail}(X)) - f(X) \right] \\ & + \lambda \mathbf{E}_{\text{arrival}} \left[f(\text{join}(X, [S])) - f(X) \right] \end{aligned}$$

Analysis via drift

Dynamics:

- If list nonempty: decrease head at rate 1
- When head = 0: remove head of list
- Poisson(λ): draw from S , append it to list



$$\begin{aligned} 0 = & \mathbf{E}_{\text{time}} \left[\frac{\partial}{\partial r_1} f(X) \right] \\ & + \lambda \mathbf{E}_{\text{departure}} \left[f(\text{tail}(X)) - f(X) \right] \\ & + \lambda \mathbf{E}_{\text{arrival}} \left[f(\text{join}(X, [S])) - f(X) \right] \end{aligned}$$

Analysis via drift

Dynamics:

- If list nonempty: decrease head at rate 1
- When head = 0: remove head of list
- Poisson(λ): draw from S , append it to list



$$\begin{aligned} 0 = & \mathbf{E}_{\text{time}} \left[\frac{\partial}{\partial r_1} f(X) \right] \\ & + \lambda \mathbf{E}_{\text{departure}} \left[f(\text{tail}(X)) - f(X) \right] \\ & + \lambda \mathbf{E}_{\text{arrival}} \left[f(\text{join}(X, [S])) - f(X) \right] \end{aligned}$$

$$f(x) = w(x)$$

Analysis via drift


Dynamics:

- If list nonempty: decrease head at rate 1
- When head = 0: remove head of list
- Poisson(λ): draw from S , append it to list



$$\begin{aligned} 0 = & \mathbf{E}_{\text{time}} \left[\frac{\partial}{\partial r_1} f(X) \right] \\ & + \lambda \mathbf{E}_{\text{departure}} \left[f(\text{tail}(X)) - f(X) \right] \\ & + \lambda \mathbf{E}_{\text{arrival}} \left[f(\text{join}(X, [S])) - f(X) \right] \end{aligned}$$

$$f(x) = w(x)$$



$$\mathbf{P}_{\text{time}}[X \text{ empty}] = 1 - \lambda \mathbf{E}[S]$$

Analysis via drift


Dynamics:

- If list nonempty: decrease head at rate 1
- When head = 0: remove head of list
- Poisson(λ): draw from S , append it to list



$$\begin{aligned} 0 = & \mathbf{E}_{\text{time}} \left[\frac{\partial}{\partial r_1} f(X) \right] \\ & + \lambda \mathbf{E}_{\text{departure}} \left[f(\text{tail}(X)) - f(X) \right] \\ & + \lambda \mathbf{E}_{\text{arrival}} \left[f(\text{join}(X, [S])) - f(X) \right] \end{aligned}$$

$$f(x) = w(x)$$



$$\mathbf{P}_{\text{time}}[X \text{ empty}] = 1 - \lambda \mathbf{E}[S]$$

$$f(x) = w(x)^2$$

Analysis via drift


Dynamics:

- If list nonempty: decrease head at rate 1
- When head = 0: remove head of list
- Poisson(λ): draw from S , append it to list




$$\begin{aligned} 0 = & \mathbf{E}_{\text{time}} \left[\frac{\partial}{\partial r_1} f(X) \right] \\ & + \lambda \mathbf{E}_{\text{departure}} [f(\text{tail}(X)) - f(X)] \\ & + \lambda \mathbf{E}_{\text{arrival}} [f(\text{join}(X, [S])) - f(X)] \end{aligned}$$

$$f(x) = w(x)$$



$$\mathbf{P}_{\text{time}}[X \text{ empty}] = 1 - \lambda \mathbf{E}[S]$$

$$f(x) = w(x)^2$$



$$\mathbf{E}_{\text{time}}[w(X)] = \frac{\frac{\lambda}{2} \mathbf{E}[S^2]}{1 - \lambda \mathbf{E}[S]}$$

Analysis via drift


Dynamics:

- If list nonempty: decrease head at rate 1
- When head = 0: remove head of list
- Poisson(λ): draw from S , append it to list




$$\begin{aligned} 0 = & \mathbf{E}_{\text{time}} \left[\frac{\partial}{\partial r_1} f(X) \right] \\ & + \lambda \mathbf{E}_{\text{departure}} [f(\text{tail}(X)) - f(X)] \\ & + \lambda \mathbf{E}_{\text{arrival}} [f(\text{join}(X, [S])) - f(X)] \end{aligned}$$

$$f(x) = w(x)$$



$$\mathbf{P}_{\text{time}}[X \text{ empty}] = 1 - \lambda \mathbf{E}[S]$$

$$f(x) = w(x)^2$$



$$\mathbf{E}_{\text{time}}[w(X)] = \frac{\frac{\lambda}{2} \mathbf{E}[S^2]}{1 - \lambda \mathbf{E}[S]}$$

Principle: PASTA
 $\mathbf{E}_{\text{arrival}}[\cdot] = \mathbf{E}_{\text{time}}[\cdot]$

Analysis via drift

Dynamics:


- If list nonempty: decrease head at rate 1
- When head = 0: remove head of list
- Poisson(λ): draw from S , append it to list



$$\begin{aligned}
 0 = & \mathbf{E}_{\text{time}} \left[\frac{\partial}{\partial r_1} f(X) \right] \\
 & + \lambda \mathbf{E}_{\text{departure}} [f(\text{tail}(X)) - f(X)] \\
 & + \lambda \mathbf{E}_{\text{arrival}} [f(\text{join}(X, [S])) - f(X)]
 \end{aligned}$$


Recipe: to get n th-order info, use $(n+1)$ th-order function f

$$f(x) = w(x)$$



$$\mathbf{P}_{\text{time}}[X \text{ empty}] = 1 - \lambda \mathbf{E}[S]$$

$$f(x) = w(x)^2$$



$$\mathbf{E}_{\text{time}}[w(X)] = \frac{\frac{\lambda}{2} \mathbf{E}[S^2]}{1 - \lambda \mathbf{E}[S]}$$

Principle: PASTA
 $\mathbf{E}_{\text{arrival}}[\cdot] = \mathbf{E}_{\text{time}}[\cdot]$

Beyond the M/G/1

Beyond the M/G/1

Work decomposition law: under M/G arrivals,

$$\mathbf{E}_{\text{time}}[w(X)] = \frac{\frac{\lambda}{2}\mathbf{E}[S^2] + \mathbf{E}_{\text{time}}[u(X)w(X)]}{1 - \lambda\mathbf{E}[S]}$$

Beyond the M/G/1

Work decomposition law: under M/G arrivals,

$$\mathbf{E}_{\text{time}}[w(X)] = \frac{\frac{\lambda}{2}\mathbf{E}[S^2] + \mathbf{E}_{\text{time}}[u(X)w(X)]}{1 - \lambda\mathbf{E}[S]}$$

unused service

Beyond the M/G/1

Work decomposition law: under M/G arrivals,

$$\mathbf{E}_{\text{time}}[w(X)] = \frac{\frac{\lambda}{2}\mathbf{E}[S^2] + \mathbf{E}_{\text{time}}[u(X)w(X)]}{1 - \lambda\mathbf{E}[S]}$$

unused service



M/G/1

$$\mathbf{E}_{\text{time}}[u(X)w(X)] = 0$$

Beyond the M/G/1

Work decomposition law: under M/G arrivals,

$$\mathbf{E}_{\text{time}}[w(X)] = \frac{\frac{\lambda}{2}\mathbf{E}[S^2] + \mathbf{E}_{\text{time}}[u(X)w(X)]}{1 - \lambda\mathbf{E}[S]}$$

unused service



M/G/1


$$\mathbf{E}_{\text{time}}[u(X)w(X)] = 0$$




M/G/k

$$\frac{\mathbf{E}_{\text{time}}[u(X)w(X)]}{1 - \lambda\mathbf{E}[S]} \text{ is work of } \leq k-1 \text{ jobs}$$

Dispatching systems

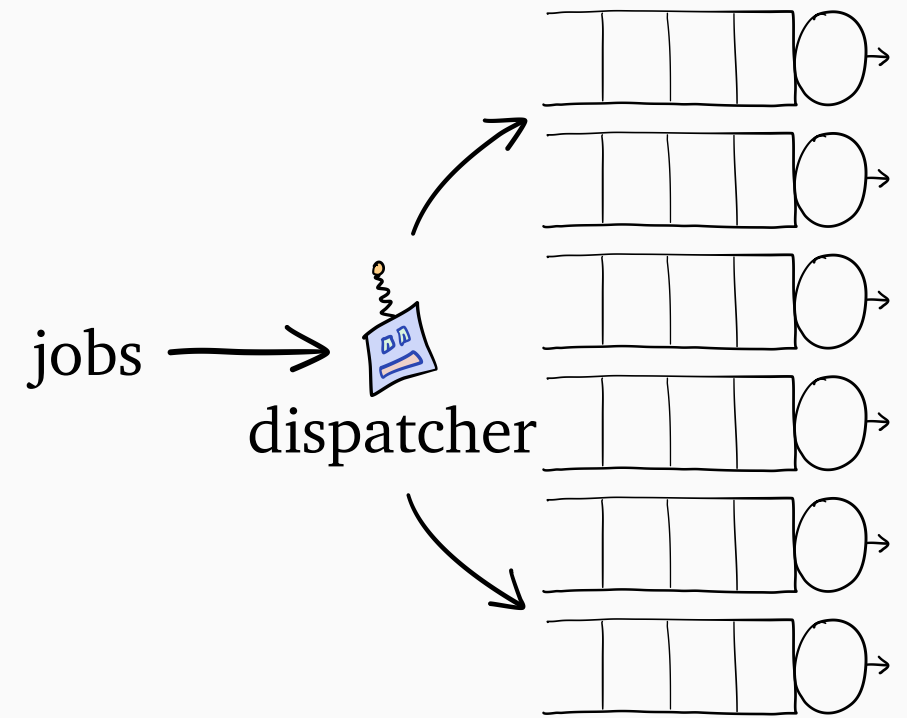


Dispatching systems

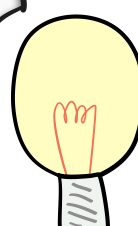


Key: $\mathbf{E}_{\text{time}}[u(X) w(X)]$

Dispatching systems




Key: $\mathbf{E}_{\text{time}}[u(X)w(X)]$

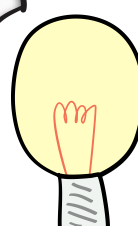
 If lots of work,
want servers busy

Dispatching systems

Possible policy:
dispatch to server
with less work





Key: $E_{\text{time}}[u(X)w(X)]$

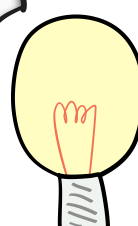
 If lots of work,
want servers busy

Dispatching systems

Possible policy:
dispatch to server
with less work





Key: $\mathbf{E}_{\text{time}}[u(X) w(X)]$

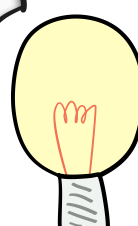
 If lots of work,
want servers busy

Dispatching systems

Possible policy:
dispatch to server
with less work


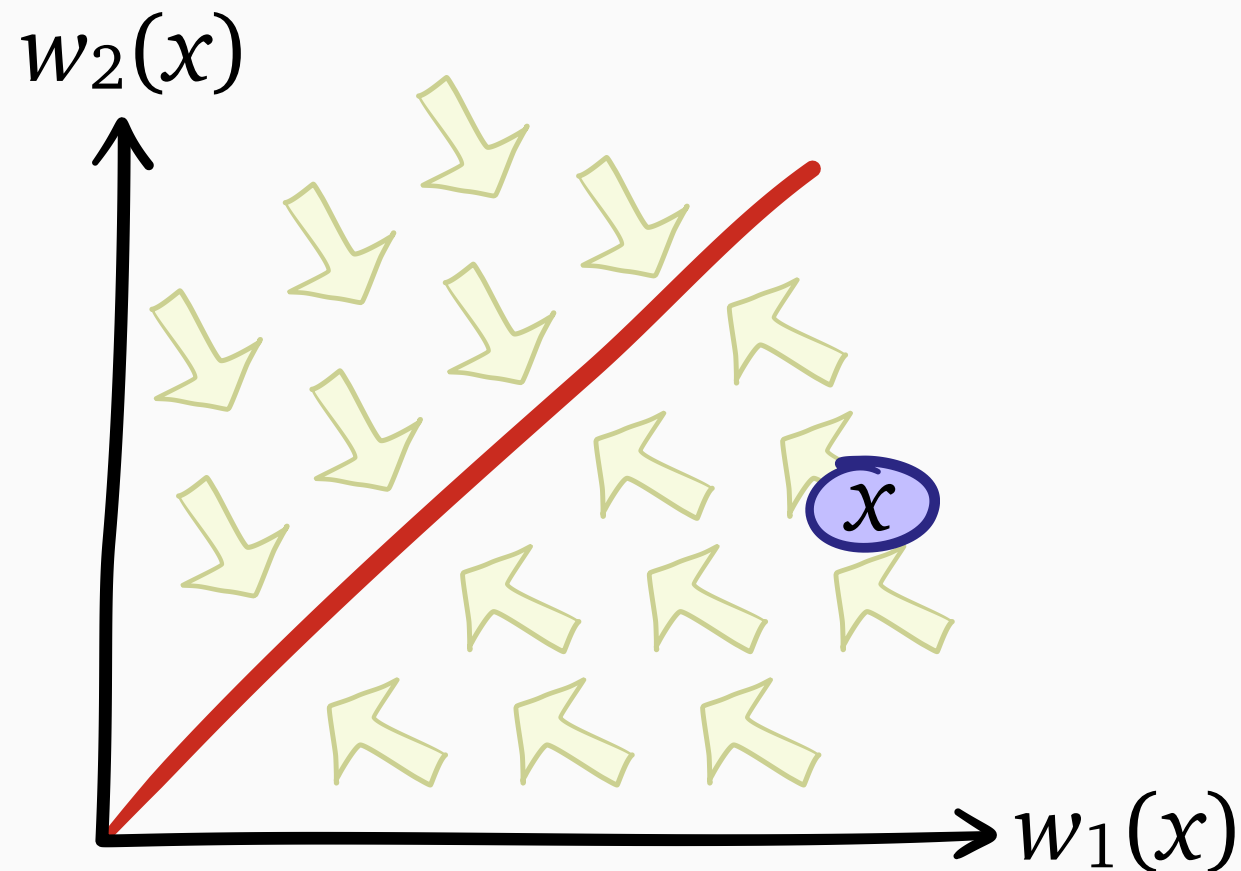


Key: $\mathbf{E}_{\text{time}}[u(X) w(X)]$

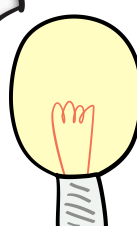
 If lots of work,
want servers busy

Dispatching systems

Possible policy:
dispatch to server
with less work


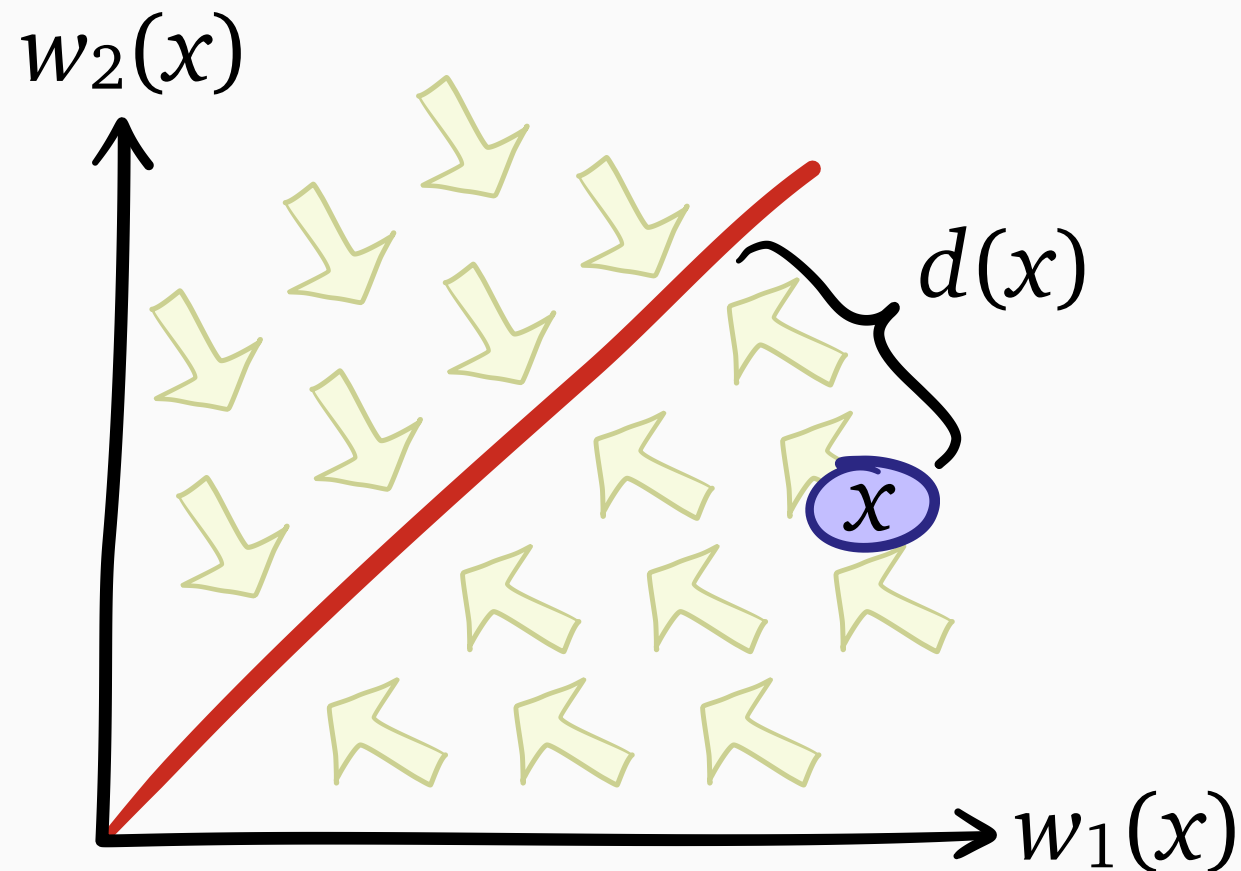


Key: $\mathbf{E}_{\text{time}}[u(X)w(X)]$

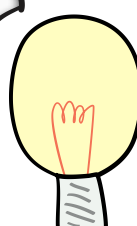
 If lots of work,
want servers busy

Dispatching systems

Possible policy:
dispatch to server
with less work





Key: $\mathbf{E}_{\text{time}}[u(X)w(X)]$

 If lots of work,
want servers busy

Dispatching systems

Possible policy:
dispatch to server
with less work



Key: $\mathbf{E}_{\text{time}}[u(X)w(X)]$

If lots of work,
want servers busy

$f(x) = \exp(\theta d(x))$ **RCL**  state space collapse

What principles do we need?

Dynamics:

- If list nonempty: decrease head at rate 1
- When head = 0: remove head of list
- Poisson(λ): draw from S , append it to list



$$\begin{aligned} 0 = & \mathbf{E}_{\text{time}} \left[\frac{\partial}{\partial r_1} f(X) \right] \\ & + \lambda \mathbf{E}_{\text{departure}} \left[f(\text{tail}(X)) - f(X) \right] \\ & + \lambda \mathbf{E}_{\text{arrival}} \left[f(\text{join}(X, [S])) - f(X) \right] \end{aligned}$$

What principles do we need?

Dynamics:

- If list nonempty: decrease head at rate 1
- When head = 0: remove head of list
- Poisson(λ): draw from S , append it to list



$$\begin{aligned} 0 = & \mathbf{E}_{\text{time}} \left[\frac{\partial}{\partial r_1} f(X) \right] \\ & + \lambda \mathbf{E}_{\text{departure}} \left[f(\text{tail}(X)) - f(X) \right] \\ & + \lambda \mathbf{E}_{\text{arrival}} \left[f(\text{join}(X, [S])) - f(X) \right] \end{aligned}$$

Principle: translating dynamics to mean rate

What principles do we need?

Dynamics:

- If list nonempty: decrease head at rate 1
- When head = 0: remove head of list
- Poisson(λ): draw from S , append it to list



$$\begin{aligned} 0 = \mathbf{E}_{\text{time}} & \left[\frac{\partial}{\partial r_1} f(X) \right] \\ & + \lambda \mathbf{E}_{\text{departure}} \left[f(\text{tail}(X)) - f(X) \right] \\ & + \lambda \mathbf{E}_{\text{arrival}} \left[f(\text{join}(X, [S])) - f(X) \right] \end{aligned}$$

Principle: translating dynamics to mean rate

? Principle for stability?

What principles do we need?

Dynamics:

- If list nonempty: decrease head at rate 1
- When head = 0: remove head of list
- Poisson(λ): draw from S , append it to list



$$\begin{aligned} 0 = & \mathbf{E}_{\text{time}} \left[\frac{\partial}{\partial r_1} f(X) \right] \\ & + \lambda \mathbf{E}_{\text{departure}} \left[f(\text{tail}(X)) - f(X) \right] \\ & + \lambda \mathbf{E}_{\text{arrival}} \left[f(\text{join}(X, [S])) - f(X) \right] \end{aligned}$$

Principle: translating dynamics to mean rate



Principle for stability?



Principle for mean field?

What principles do we need?

Dynamics:

- If list nonempty: decrease head at rate 1
- When head = 0: remove head of list
- Poisson(λ): draw from S , append it to list



$$\begin{aligned} 0 = \mathbf{E}_{\text{time}} & \left[\frac{\partial}{\partial r_1} f(X) \right] \\ & + \lambda \mathbf{E}_{\text{departure}} \left[f(\text{tail}(X)) - f(X) \right] \\ & + \lambda \mathbf{E}_{\text{arrival}} \left[f(\text{join}(X, [S])) - f(X) \right] \end{aligned}$$

Principle: translating dynamics to mean rate



Principle for stability?



Principle for mean field?




Principles for composition?

Part 1


**Performance modeling
needs advanced math**

Part 2

**We can teach advanced
math accessibly**



Simplify core
foundations



Prioritize very
flexible tools