# Sequential Decision Processes, Master MICAS, Part I 

Michèle Wigger

Telecom Paris, 27 November 2020

## Outline of the Course: Part I

Michèle Wigger (3C58) and Mustapha Hamad (3C54)

- Markov Chains
- Dynamic Programming for Finite Horizon and Shortest-Paths Problems
- Dynamic Programming for Infinite Horizon Problems with Discounted and Average Cost Functions
- Constrained Markov Decision Processes: Solutions and Suboptimal Policies
- 2 TDs and 1 TP


## Outline of the Course: Part II

Mireille Sarkiss, Telecom SudParis, 3C56

- Markov Decision Processes without known transition probabilities
- Reinforcement Learning: exploration/exploitation tradeoff
- Epsilon Greedy, Boltzman Algorithm
- Deep reinforcement learning


# Lecture 1 - Finite-State Markov Chains 

## Definitions and Types of Markov Chains

## Definition (First-order Markov Chain)

A stochastic process $\left\{X_{k}\right\}_{k \geq 0}=\left\{X_{0}, X_{1}, X_{2}, \ldots,\right\}$ over an alphabet $\mathcal{X}$ is called a (first-order) Markov chain if for all $k=1,2, \ldots$, :

$$
P_{X_{k} \mid X_{k-1}, X_{k-2}, \ldots, x_{0}}(a \mid b, c, \ldots, z)=P_{X_{k} \mid X_{k-1}}(a \mid b), \quad \forall a, b, c, \ldots, z \in \mathcal{X}
$$

- Examples: Random walk, memoryless process, ...
- Statistics of the stochastic process $\left\{X_{k}\right\}_{k \geq 0}$ is determined by $P_{X_{0}}$ and $\left\{P_{X_{k} \mid X_{k-1}}\right\}_{k \geq 1}$. In fact:

$$
P_{X_{0}, X_{1}, \ldots, X_{K}}(a, b, c, \ldots, z)=P_{X_{0}}(a) \cdot P_{X_{1} \mid X_{0}}(b \mid a) \cdot P_{X_{2} \mid X_{1}}(c \mid b) \cdots P_{X_{K} \mid X_{K_{1}}}(z \mid y)
$$

## Homogeneous/Time-Invariant Markov Chains

## Definition (Homogeneous Markov Chains)

A Markov chain $\left\{X_{k}\right\}_{k \geq 0}$ over an alphabet $\mathcal{X}$ is called homogeneous or time-invariant if the transition probability $P_{X_{k} \mid X_{k-1}}$ does not depend on the index $k$. That means, there exists a conditional probability mass function $W(\cdot \mid \cdot)$ such that:

$$
P_{X_{k} \mid X_{k-1}}(a \mid b)=W(a \mid b), \quad \forall k=1,2, \ldots, \text { and } a, b \in \mathcal{X}
$$

- The alphabet $\mathcal{X}$ is typically called the state space and $W$ the transition law of the homogeneous Markov chain.


## State-Transition Diagramme for Homogeneous Markov Chains

- A node for all possible states $a \in \mathcal{X}$ and an arrow from state $b$ to state $a$ labelled by the probability $W(a \mid b)>0$. (If $W(a \mid b)=0$ there is no arrow.)
- Each outgoing edge from state $b$ represents a probability $W(\cdot \mid b)$ $\Rightarrow$ the labels of all outgoing edges from a given node have to sum to 1 !

Life in Lockdown:


## Describing a Homogeneous Markov Chain with its Transition Matrix

- Transition matrix W: each row and each column is associated with a state $\rightarrow \mathrm{W}$ is square of dimension $|\mathcal{X}| \times|\mathcal{X}|$

$$
\mathrm{W}=\left(\begin{array}{ccccc}
W(a \mid a) & W(b \mid a) & W(c \mid a) & \cdots & W(z \mid a) \\
W(a \mid b) & W(b \mid b) & W(c \mid b) & \cdots & W(z \mid b) \\
\vdots & \cdots & \ddots & \cdots & \\
W(a \mid z) & \underbrace{W(b \mid z)}_{W_{i, b}} & \cdots & \cdots & W(z \mid z)
\end{array}\right)
$$

- Each row of W sums to $1 \rightarrow$ a (right) stochastic matrix
- For any state $b$ :

$$
P_{X_{1}}(b)=\sum_{x \in \mathcal{X}} P_{X_{0}}(x) W(b \mid x)=\pi_{0} \cdot W_{:, b}
$$

where $\pi_{k}=\left(P_{x_{k}}(a), P_{x_{k}}(b), \ldots, P_{x_{k}}(z)\right)$.

- Summary for all $b \in \mathcal{X}$ :

$$
\pi_{1}=\pi_{0} \mathrm{~W}
$$

## The Markov Process in Matrix Notation

- Let $\pi_{k}=\left(P_{x_{k}}(a), P_{x_{k}}(b), \ldots, P_{x_{k}}(z)\right)$. Then:

$$
\begin{aligned}
\boldsymbol{\pi}_{1} & =\boldsymbol{\pi}_{0} \cdot \mathrm{~W} \\
\boldsymbol{\pi}_{2} & =\boldsymbol{\pi}_{1} \cdot \mathrm{~W}=\boldsymbol{\pi}_{0} \cdot \mathrm{~W} \cdot \mathrm{~W} \\
& \vdots \\
\boldsymbol{\pi}_{k} & =\boldsymbol{\pi}_{0} \cdot \mathrm{~W}^{k} .
\end{aligned}
$$

$\rightarrow$ the statistics is determined by $\pi_{0}$ and W

## Transient and Recurrent States

## Definition (Recurrent State Class)

Consider a homogeneous Markov process. A class of states $\mathcal{S} \subseteq \mathcal{X}$ is called recurrent, if the following two conditions hold:
(1) For any two states $a, b \in \mathcal{S}$ there are positive integers $k, i, j$ such that

$$
\operatorname{Pr}\left[X_{k+i}=b \mid X_{k}=a\right]>0 \quad \text { and } \quad \operatorname{Pr}\left[X_{k+j}=a \mid X_{k}=b\right]>0
$$

(We say that states $a$ and $b$ communicate.)
(2) For any states $a \in \mathcal{S}$ and $b \in \mathcal{X} \backslash \mathcal{S}$ and for all $k, i>0$ :

$$
\operatorname{Pr}\left[X_{k+i}=b \mid X_{k}=a\right]>0
$$

If $\mathcal{X}$ is a recurrent class, the Markov process $\left\{X_{k}\right\}_{k \geq 0}$ is said irreducible.

## Definition (Recurrent and Transient States)

A state $a \in \mathcal{X}$ that belongs to some recurrent class is called recurrent. A state that does not belong to any recurrent class is called transient. For any transient state $a$ :

$$
\lim _{i \rightarrow \infty} \operatorname{Pr}\left[X_{k+i}=a \mid X_{k}=a\right]=0
$$

## Periodicity of States And Aperiodic Chains

## Definition (Periods of a states)

The period $d(x)$ of a state $x$ is the smallest positive integer such that irrespective of the starting distribution $\operatorname{Pr}\left[X_{\ell+k}=x \mid X_{k}=x\right]=0$ if $\ell$ is not a multiple of $d(x)$.

period of states:

## Definition (Aperiodic Markov Chains)

A Markov chain $\left\{X_{k}\right\}$ is said aperiodic if $d(x)=1$ for all states $x \in \mathcal{X}$.

## A Stationary Process

## Definition (Stationary Process)

A stochastic process $\left\{X_{k}\right\}_{k \geq 0}$ is called stationary, if for all integers $k, n \geq 0$ :

$$
P_{x_{k}, x_{k+1}, \ldots, x_{k+n}}(a, b, \ldots, z)=P_{X_{0}, x_{1}, \ldots, x_{n}}(a, b, \ldots, z), \quad \forall a, b, \ldots, z \in \mathcal{X}
$$

## Theorem

A Markov process $\left\{X_{k}\right\}_{k \geq 0}$ with transition matrix W and initial distribution $\pi_{0}$ is stationary if, and only if,

$$
\boldsymbol{\pi}_{0}=\boldsymbol{\pi}_{0} \cdot \mathrm{~W}
$$

Proof: The "only if" direction is trivial because $\pi_{1}=\pi_{0} \cdot \mathrm{~W}$.
To see the "if"-direction, notice that for any $k \geq 1$ :

$$
\pi_{k}=\pi_{0} \cdot \mathrm{~W}^{k}=\underbrace{\pi_{0} \cdot \mathrm{~W}}_{=\pi_{0}} \cdot \mathrm{~W}^{k-1}=\pi_{0} \cdot \mathrm{~W}^{k-1}=\underbrace{\pi_{0} \cdot \mathrm{~W}}_{=\pi_{0}} \cdot \mathrm{~W}^{k-2}=\cdots=\pi_{0} \cdot \mathrm{~W}=\pi_{0}
$$

and thus by Bayes' rule and the Markov property:

$$
\begin{aligned}
& P_{X_{k}, X_{k+1}, \ldots, X_{k+n}}(a, b, \ldots, z)=P_{X_{k}}(a) P_{X_{k+1} \mid X_{k}}(b \mid a) \cdots P_{X_{k+n} \mid X_{k+n-1}}(z \mid y) \\
& =\pi_{0}(a) \cdot W(b \mid a) \cdot W(c \mid b) \cdots W(z \mid y)=P_{X_{0}}(a) P_{X_{1} \mid X_{0}}(b \mid a) \cdots P_{X_{n} \mid X_{n-1}}(z \mid y) \\
& =P_{X_{0}, X_{1}, \ldots, x_{n}}(a, b, \ldots, z)
\end{aligned}
$$

## More on Stationary Distributions

Consider a Markov chain $\left\{X_{k}\right\}_{k \geq 0}$ with transition matrix W .

- Any distribution $\pi$ satisfying the fix-point equation

$$
\boldsymbol{\pi}=\boldsymbol{\pi} \cdot \mathrm{W}
$$

is called a stationary distribution of this Markov chain.

- Any such $\pi$ is an eigenvector of W corresponding to eigenvalue 1 .
- Aperiodic and irreducible Markov chains have a unique stationary distribution $\pi^{*}$.
- Transient states have 0 probability under $\pi^{*}$.


## Convergence of the Transition Matrix

## Theorem

The following limit exists

$$
\mathrm{W}^{*}:=\lim _{N \rightarrow \infty} \mathrm{~W}^{N}
$$

and $\mathrm{W}^{*}$ is a stochastic matrix.
For an irreducibile and aperiodic Markov chain:

$$
\mathrm{W}^{*}=\mathbf{1}^{\top} \boldsymbol{\pi}^{*}
$$

where $\pi^{*}$ is the unique stationary distribution.

## Proof.

Omitted.

## Convergence to A Stationary Process

## Theorem

If the Markov chain $\left\{X_{k}\right\}_{k \geq 0}$ is aperiodic and irreducible, then for any initial distribution $\pi_{0}$ :

$$
\lim _{N \rightarrow \infty} \boldsymbol{\pi}_{N} \rightarrow \boldsymbol{\pi}^{*}
$$

where $\pi^{*}$ is the only stationary distribution of the Markov chain.
Proof:

$$
\lim _{N \rightarrow \infty} \boldsymbol{\pi}_{N}=\lim _{N \rightarrow \infty}\left(\boldsymbol{\pi}_{0} \cdot W^{N}\right)=\boldsymbol{\pi}_{0} \cdot \lim _{N \rightarrow \infty} W^{N}=\underbrace{\boldsymbol{\pi}_{0} \cdot \mathbf{1}^{\top}}_{=1} \boldsymbol{\pi}^{*}
$$

# Sequential Decision Processes，Master MICAS，Part I 

Michèle Wigger

Telecom Paris， 27 November 2020

## Lecture 2 - Markov Decision Processes and Dynamic Programming over a Finite Horizon

## A Discrete-Time Dynamic System Model

- State evolution

$$
X_{k+1}=f_{k}\left(X_{k}, U_{k}, W_{k}\right), \quad k=0,1,2, \ldots
$$

- $X_{k}$ is the time- $k$ state over a state space $\mathcal{X}$
- $U_{k}$ is the time- $k$ (control) action over a space $\mathcal{U}$
- $W_{k}$ the random disturbance


## Markov Decision Process (MDP) -A Markov Chain with Actions

The discrete-time dynamic system is a Markov decision process if

- the sequence $\left\{W_{k}\right\}$ is memoryless; and
- a reward $R_{u}\left(x, x^{\prime}\right)$ is associated to each action $u$ and pair of states $x, x^{\prime} \in \mathcal{X}$
$\rightarrow$ Generalization of a Markov chain to incorporate actions and where the transition law depends on these actions:

$$
\begin{aligned}
P_{X_{k+1} \mid X_{k}, \ldots, X_{0}, U_{k}, \ldots, u_{0}}(a \mid b, \ldots, z, u, \ldots, v)= & P_{X_{k+1} \mid X_{k}, u_{k}}(a \mid b, u), \\
& \forall a, b, \ldots, z \in \mathcal{X}, u, v \in \mathcal{U} .
\end{aligned}
$$

## An MDP Example with Graph Representation



- Boxes are states; labels on arrows designate actions and transition probabilities. E.g.:

$$
\operatorname{Pr}\left[X_{k+1}=" \mathrm{I} " \mid X_{k}=" \mathrm{U} ", U_{k}=" \mathrm{i} "\right]=0.6 .
$$

## Finite-Horizon Dynamic Programming Problem Setup

(Slightly more general than introduced for MDPs)

- Discrete-time dynamic system:

$$
X_{k+1}=f_{k}\left(X_{k}, U_{k}, W_{k}\right), \quad k=0,1,2, \ldots, N-1
$$

where given $\left(X_{k}, U_{k}\right)$ the noise $W_{k}$ is conditionally independent of $\left(X_{0}, \ldots, X_{k-1}, U_{1}, \ldots, U_{k-1}, W_{1}, \ldots, W_{k-1}\right)$

- $N$ is called the horizon of the control problem
- Admissible control sets $\left\{\mathcal{U}_{k}(a)\right\}_{a \in \mathcal{X}}$ for action $U_{k}=\mu_{k}\left(X_{k}\right)$
$\rightarrow$ The set of functions $\mu_{0}, \ldots, \mu_{N-1}$ is called a policy $\pi$
- Additive expected cost
$\mathbb{E}\left[g_{N}\left(X_{N}\right)+\sum_{k=0}^{N-1} g_{k}\left(X_{k}, U_{k}, W_{k}\right)\right]=\mathbb{E}_{\left\{W_{k}\right\}}\left[g_{N}\left(X_{N}\right)+\sum_{k=0}^{N-1} g_{k}\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right)\right]$
where $g_{N}\left(X_{N}\right)$ denotes a terminal cost


## Decomposition of Expected Cost

- Expected time $i$-to- $j$ cost starting from state $a \in \mathcal{X}$ :

$$
J_{i \rightarrow j, \pi}(a)=\mathbb{E}\left[\sum_{k=i}^{j} g_{k}\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right) \mid X_{i}=a\right], \quad 0 \leq i<j \leq N
$$

where $g_{N}\left(X_{N}, \mu_{N}\left(X_{N}\right), W_{N}\right):=g_{N}\left(X_{N}\right)$.

- Decomposition of finite-horizon expected cost for $i<j \leq N$ :

$$
\begin{aligned}
& J_{i \rightarrow N, \pi}(a)=\mathbb{E}\left[g_{N}\left(X_{N}\right)+\sum_{k=i}^{N-1} g_{k}\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right) \mid X_{j}=b, X_{i}=a\right] \\
& =\sum_{b \in \mathcal{X}} \operatorname{Pr}\left[X_{j}=b \mid X_{i}=a\right] \mathbb{E}\left[g_{N}\left(X_{N}\right)+\sum_{k=i}^{N-1} g_{k}\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right) \mid X_{j}=b, X_{i}=a\right] \\
& =\sum_{b \in \mathcal{X}} \operatorname{Pr}\left[X_{j}=b \mid X_{i}=a\right] \mathbb{E}\left[g_{N}\left(X_{N}\right)+\sum_{k=j}^{N-1} g_{k}\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right) \mid X_{j}=b, X_{i}=a\right] \\
& \quad+\sum_{b \in \mathcal{X}} \operatorname{Pr}\left[X_{j}=b \mid X_{i}=a\right] \mathbb{E}\left[\sum_{k=i}^{j-1} g_{k}\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right) \mid X_{j}=b, X_{i}=a\right] \\
& = \\
& \sum_{b \in \mathcal{X}} \operatorname{Pr}\left[X_{j}=b \mid X_{i}=a\right] J_{j \rightarrow N, \pi}(b)+J_{i \rightarrow j-1, \pi}(a)
\end{aligned}
$$

## Minimizing the Expected Finite-Horizon Cost

- Minimize expected cost for $a \in \mathcal{X}: \quad J_{0 \rightarrow N}^{*}(a)=\min _{\pi} J_{0 \rightarrow N, \pi}(a)$
- Decomposition of optimization problem:

$$
\begin{aligned}
\min _{\pi} J_{0 \rightarrow N, \pi}(a) & =\min _{\mu_{0}}\left[J_{0 \rightarrow 0, \mu_{0}}(a)+\min _{\mu_{1}, \ldots, \mu_{N-1}} \sum_{b \in \mathcal{X}} \operatorname{Pr}\left[X_{1}=b \mid X_{0}=a\right] J_{1 \rightarrow N, \pi}(b)\right] \\
& \geq \min _{\mu_{0}}\left[J_{0 \rightarrow 0, \mu_{0}}(a)+\sum_{b \in \mathcal{X}} \operatorname{Pr}\left[X_{1}=b \mid X_{0}=a\right] \min _{\mu_{b, 1}, \ldots, \mu_{b, N-1}} J_{1 \rightarrow N, \pi_{b}}(b)\right]
\end{aligned}
$$

where equality holds when optimal policies $\mu_{b, 1}, \ldots, \mu_{b, N-1}$ don't depend on $b$.

$$
\begin{gathered}
\min _{\pi} J_{1 \rightarrow N, \pi}(b) \geq \min _{\mu_{1}}\left[J_{1 \rightarrow 1, \mu_{1}}(b)+\sum_{c \in \mathcal{X}} \operatorname{Pr}\left[X_{2}=c \mid X_{1}=b\right] \min _{\mu_{c, 2}, \ldots, \mu_{c, N-1}} J_{2 \rightarrow N, \pi_{c}}(c)\right] \\
\vdots \\
\min _{\pi} J_{N-1 \rightarrow N, \pi}(x) \geq \min _{\mu_{N-2}}\left[J_{N-1 \rightarrow N-1, \mu_{N-1}}(x)\right. \\
\left.\quad+\sum_{y \in \mathcal{X}} \operatorname{Pr}\left[X_{N}=y \mid X_{N-1}=x\right] \min _{\mu_{y, N-1}} J_{N \rightarrow N, \pi_{y}}(y)\right]
\end{gathered}
$$

- Will see: optimal $\mu_{a, i}, \ldots, \mu_{a, N-1}$ don't depend on $a \Rightarrow$ Ineq. are equalities
- Find the optimal solution starting backwards!!


## Optimal Dynamic Programming Algorithm

- For each $x_{N} \in \mathcal{X}$ initialize $J_{N \rightarrow N}^{*}\left(x_{N}\right)=g_{N}\left(x_{N}\right)$
$\rightarrow$ trivially the same $\mu_{N}$ achieves optimal $J_{N \rightarrow N}^{*}\left(x_{N}\right)$ for all $x_{N} \in \mathcal{X}$
- For each $i=N-1, \ldots, 0$ calculcate for each $x_{i} \in \mathcal{X}$ :

$$
\begin{aligned}
& J_{i \rightarrow N}^{*}\left(x_{i}\right) \\
&\left.:=\min _{\mu_{i}}\left[J_{i \rightarrow i, \mu_{i}}\left(x_{i}\right)+\sum_{x_{i+1} \in \mathcal{X}} \operatorname{Pr}\left[X_{i+1}=x_{i+1} \mid X_{i}=x_{i}\right]\right]_{i+1 \rightarrow N}^{*}\left(x_{i+1}\right)\right] \\
&=\min _{\mu_{i}}\left[\mathbb{E}_{w_{i}}\left[g_{i}\left(x_{i}, \mu_{i}\left(x_{i}\right), W_{i}\right)+J_{i+1 \rightarrow N}^{*}\left(X_{i+1}\right) \mid X_{i}=x_{i}\right]\right]
\end{aligned}
$$

$\rightarrow$ If optimal policies $\mu_{i+1}^{*}, \ldots, \mu_{N}^{*}$ for $J_{i+1 \rightarrow N}^{*}\left(x_{i+1}\right)$ don't depend on $x_{i+1} \in \mathcal{X}$, then optimal policies $\mu_{i}^{*}, \mu_{i+1}, \ldots, \mu_{N}^{*}$ for $J_{i \rightarrow N}^{*}\left(x_{i}\right)$ don't depend on $x_{i}$ !

## Optimality Principle for Finite-Horizon Dynamic Programming

## Theorem (Optimality Principle)

Let $\pi^{*}=\left(\mu_{0}^{*}, \mu_{1}^{*}, \mu_{2}^{*}, \ldots, \mu_{N-1}^{*}\right)$ be an optimal policy for $J_{0 \rightarrow N, \pi}$ :

$$
J_{0 \rightarrow N, \pi^{*}}(a)=\min _{\pi} J_{0 \rightarrow N, \pi}(a)=: J_{0 \rightarrow N}^{*}(a), \quad \forall a \in \mathcal{X}
$$

Then $\forall b \in \mathcal{X}$ the truncated policy $\pi_{i \rightarrow N}^{*}:=\left(\mu_{i}^{*}, \ldots, \mu_{N-1}^{*}\right)$ minimizes the sub-problem $J_{i \rightarrow N, \pi}$ :

$$
J_{i \rightarrow N, \pi_{i \rightarrow N}^{*}}(b)=\min _{\pi} J_{i \rightarrow N, \pi}(b)=: J_{i \rightarrow N}^{*}(b), \quad \forall b \in \mathcal{X}
$$

Proof by Contradiction: Given policy $\pi_{i \rightarrow N}=\left(\mu_{0}, \mu_{1}, \ldots, \mu_{N-1}\right)$ satisfying

$$
J_{i \rightarrow N, \pi}(b)<J_{i \rightarrow N, \pi^{*}}(b), \quad \forall b \in \mathcal{X}
$$

Then for all $a \in \mathcal{X}$ and policy $\tilde{\pi}=\left(\mu_{0}^{*}, \mu_{1}^{*}, \ldots, \mu_{i-1}^{*}, \mu_{i}, \ldots, \mu_{N-1}\right)$ :

$$
\begin{aligned}
J_{0 \rightarrow N, \pi^{*}}(a) & =\sum_{b \in \mathcal{X}} \operatorname{Pr}\left[X_{i}=b \mid X_{0}=a\right] J_{i \rightarrow N, \pi^{*}}(b)+J_{0 \rightarrow i-1, \pi^{*}}(a) \\
& >\sum_{b \in \mathcal{X}} \operatorname{Pr}\left[X_{i}=b \mid X_{0}=a\right] J_{i \rightarrow N, \pi}(b)+J_{0 \rightarrow i-1, \pi^{*}}(a) \\
& =J_{0 \rightarrow N, \tilde{\pi}}(a)
\end{aligned}
$$

## Example: Inventory Control

- state $x_{k}$ : stock at the beginning of period $k$
- action $u_{k}$ : stock order (and delivery) at the beginning of period $k$
- disturbance $w_{k}$ : random demand during period $k$
- state evolution:

$$
x_{k+1}=f\left(x_{k}, u_{k}, w_{k}\right)=x_{k}+u_{k}-w_{k} .
$$

- cost $g_{k}\left(x_{k}, u_{k}, w_{k}\right)$ in period $k$ consists of inventory cost/penalty $r\left(x_{k}\right)$ and purchase cost $c u_{k}$ :

$$
g_{k}\left(x_{k}, u_{k}, w_{k}\right)=r\left(x_{k}\right)+c \cdot u_{k}
$$

- Wish to minimize total expected cost over horizon $N$ :

$$
J_{0 \rightarrow N, \pi}=\mathbb{E}\left[\sum_{k=0}^{N} r\left(x_{k}\right)+\sum_{k=0}^{N-1} c \cdot u_{k} \mid X_{0}=a\right], \quad a \geq 0
$$

## Optimal DP Algorithm for the Inventory Control Example

- Initialize $J_{N \rightarrow N}^{*}\left(x_{N}\right)=r\left(x_{N}\right)$
- First iteration:

$$
\begin{aligned}
J_{N-1 \rightarrow N}^{*}\left(x_{N-1}\right) & =\min _{u_{N-1}}\left\{r\left(x_{N-1}\right)+c u_{N-1}+\mathbb{E}\left[r\left(X_{N}\right)\right]\right\} \\
& =r\left(x_{N-1}\right)+\min _{u_{N-1}}\left\{c u_{N-1}+\mathbb{E}_{W_{N-1}}\left[r\left(x_{N-1}+u_{N-1}+W_{N-1}\right)\right]\right\}
\end{aligned}
$$

- Second iteration:

$$
\begin{aligned}
J_{N-2 \rightarrow N}^{*} & =\min _{u_{N-2}}\left\{r\left(x_{N-2}\right)+c u_{N-2}+\mathbb{E}\left[J_{N-1 \rightarrow N}^{*}\left(X_{N-1}\right)\right]\right\} \\
& =r\left(x_{N-2}\right)+\min _{u_{N-2}}\left\{c u_{N-2}+\mathbb{E}_{W_{N-2}}\left[J_{N-1 \rightarrow N}^{*}\left(x_{N-2}+u_{N-2}+W_{N-2}\right)\right]\right\}
\end{aligned}
$$

- $i$-th iteration:

$$
J_{N-i \rightarrow N}^{*}=r\left(x_{N-i}\right)+\min _{u_{N-i}}\left\{c u_{N-i}+\mathbb{E}_{W_{N-i}}\left[J_{N-i-1 \rightarrow N}^{*}\left(x_{N-i}+u_{N-i}+W_{N-i}\right)\right]\right\}
$$

- Solution obtained after $N$ iterations: $J_{0 \rightarrow N}^{*}$


## Deterministic MDPs and Shortest-Path Problems

- No disturbance $\rightarrow$ state evolution $x_{k+1}=f\left(x_{k}, u_{k}\right)$ and cost $g_{k}\left(x_{k}, u_{k}\right)$
- Graph representation:

- At each stage $k=1,2, \ldots, N$ there is a node for each $x_{k} \in \mathcal{X}$
- Arrows indicate transitions for different actions $\rightarrow$ label arrows with actions $u_{k}$ and costs $g_{k}\left(x_{k}, u_{k}\right)$
- Total cost $J_{0 \rightarrow N, \pi}$ is the sum of the costs on the path indicated by $\pi$

Finding minimum total cost $J_{0 \rightarrow N, \pi}$ equivalent to finding "shortest path" $\rightarrow$ DP algorithm can be run in reverse order

## Travelling Salesman Problem and Label Correcting Method

Initialize $d_{1}=0$ and
$d_{2}=\cdots=d_{t}=\infty$

## Label Correcting Algorithm

Step 1: Remove a node $i$ from OPEN and for each child $j$ of $i$, execute step 2.
Step 2: If $d_{i}+a_{i j}<\min \left\{d_{j}\right.$, UPPER $\}$, set $d_{j}=d_{i}+a_{i j}$ and set $i$ to be the parent of $j$. In addition, if $j \neq t$, place $j$ in OPEN if it is not already in OPEN, while if $j=t$, set UPPER to the new value $d_{i}+a_{i t}$ of $d_{t}$.

Step 3: If OPEN is empty, terminate; else go to step 1.

| Iter. No. | Node Exiting OPEN | OPEN at the End of Iteration | UPPER |
| :---: | :---: | :---: | :---: |
| 0 | - | 1 | $\infty$ |
| 1 | 1 | $2,7,10$ | $\infty$ |
| 2 | 2 | $3,5,7,10$ | $\infty$ |
| 3 | 3 | $4,5,7,10$ | $\infty$ |
| 4 | 4 | $5,7,10$ | 43 |
| 5 | 5 | $6,7,10$ | 43 |
| 6 | 6 | 7,10 | 13 |
| 7 | 7 | 8,10 | 13 |
| 8 | 8 | 9,10 | 13 |
| 9 | 9 | 10 | 13 |
| 10 | 10 | Empty | 13 |

- Dijkstra's method always chooses the node in OPEN with smallest $d_{i}$.


## Dynamic Programming in a Hidden Markov Model

- In a Hidden Markov Model (HMM) or Partially Observable Markov Process (POMP), an observer does not observe the state sequences $X_{0}, X_{1}, \ldots, X_{N}$ directly but a related sequence $Z_{1}, \ldots, Z_{N}$, where

$$
P_{X_{0}, x_{1}, \ldots, x_{N}, z_{1}, \ldots, z_{N}}=P_{X_{0}} \cdot \prod_{k=1}^{N} P_{X_{k} \mid X_{k-1}} \cdot P_{Z_{k} \mid X_{k}, X_{k-1}}
$$

- Observe $z_{1}, \ldots, z_{N}$ and solve

$$
\begin{aligned}
& \min _{x_{0}, x_{1}, \ldots, x_{N}}-\log P_{x_{0}, x_{1}, \ldots, x_{N}, z_{1}, \ldots, z_{N}}\left(x_{0}, x_{1}, \ldots, x_{N}, z_{1}, \ldots, z_{N}\right) \\
& =\min _{x_{0}, x_{1}, \ldots, x_{N}}\left[-\log P_{x_{0}}\left(x_{0}\right)-\sum_{k=1}^{N} \log P_{X_{k} \mid x_{k-1}}\left(x_{k} \mid x_{k-1}\right) P_{z_{k} \mid x_{k}, x_{k-1}}\left(z_{k} \mid x_{k}, x_{k-1}\right)\right]
\end{aligned}
$$

$\rightarrow$ Apply Forward DP algorithm on a Trellis

## The Viterbi Algorithm

- Trellis:


Edges from $s$ to $x_{0}$ are labeled with $P_{x_{0}}$, edges from $x_{N}$ to $t$ by 0 and edges from $x_{k-1}$ to $x_{k}$ by $-\log P_{X_{k} \mid x_{k-1}}\left(x_{k} \mid x_{k-1}\right) P_{Z_{k} \mid x_{k}, x_{k-1}}\left(z_{k} \mid x_{k}, x_{k-1}\right)$

- Shortest Path from $s$ to $t$ solves minimization problem
- Apply forward DP algorithm and cut the branches that are suboptimal


# Sequential Decision Processes, Master MICAS, Part I 

Michèle Wigger

Telecom Paris, 8 December 2020

## Lecture 3 - Dynamic Programming over an Infinite Horizon: The Discounted Case

## Review of Lecture 2: Finite Horizon and Decomposition of the Cost

- Discrete-time dynamic system:

$$
X_{k+1}=f_{k}\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right), \quad k=0,1,2, \ldots, N-1
$$

$\left\{W_{k}\right\}$ is independent and identically distributed (i.i.d.)

- Minimize total cost for given initial state $a \in \mathcal{X}$ :

$$
J_{0 \rightarrow N}^{*}(a):=\min _{\pi} \underbrace{\mathbb{E}\left[\sum_{k=0}^{N-1} g_{k}\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right)+g_{N}\left(X_{N}\right) \mid X_{0}=a\right]}_{=: J_{0 \rightarrow N, \pi}(a)}
$$

- Optimal Backward DP Algorithm: Initialize $J_{N \rightarrow N}^{*}\left(x_{N}\right):=g_{N}\left(x_{N}\right)$ and compute for $i=N-1, \ldots, 0$

$$
\begin{aligned}
J_{i \rightarrow N}^{*}\left(x_{i}\right) & =\min _{\mu_{i}}\left(\mathbb{E}\left[g_{i}\left(x_{i}, \mu_{i}\left(x_{i}\right), W_{i}\right)+\sum_{x_{i+1} \in \mathcal{X}} \operatorname{Pr}\left[X_{i+1}=x_{i+1} \mid X_{i}=x_{i}\right] J_{i+1 \rightarrow N}^{*}\left(x_{i+1}\right)\right)\right. \\
& =\min _{\mu_{i}} \mathbb{E}_{W_{i}}\left[g_{i}\left(x_{i}, \mu_{i}\left(x_{i}\right), W_{i}\right)+J_{i+1 \rightarrow N}^{*}\left(f_{i}\left(x_{i}, \mu_{i}\left(x_{i}\right), W_{i}\right)\right)\right]
\end{aligned}
$$

- For deterministic problems optimal DP algorithm can be run forwards


## Optimality of Memoryless Policies

- Restriction to memoryless policies $u_{i}=\mu_{i}\left(x_{i}\right)$ is without loss of optimality. (I.e., there is no need to consider policies of the form $\left.u_{i}=\mu_{i}\left(x_{0}, \ldots, x_{i}, u_{1}, \ldots, u_{i-1}\right).\right)$
- Recall

$$
J_{i \rightarrow N}^{*}\left(x_{i}\right)=\min _{\mu_{i}}\left(\mathbb{E}\left[g_{i}\left(x_{i}, \mu_{i}\left(x_{i}\right), W_{i}\right)+\sum_{x_{i+1} \in \mathcal{X}} \operatorname{Pr}\left[X_{i+1}=x_{i+1} \mid X_{i}=x_{i}\right] J_{i+1 \rightarrow N}^{*}\left(x_{i+1}\right)\right)\right.
$$

- $J_{i \rightarrow N}^{*}\left(x_{i}\right)$ only depends on $P_{x_{i+1} \mid} \mid X_{i}$ and $P_{x_{i} U_{i}} \rightarrow$ introducing memory would have no effect at all on the value of $J_{i \rightarrow N}^{*}\left(x_{i}\right)$.
- Deterministic policies suffice because the minimum has a deterministic solution


## Infinite-Horizon Dynamic Programming with Discounted Costs

- Time-invariant discrete-time dynamic system:

$$
X_{k+1}=f\left(X_{k}, U_{k}, W_{k}\right), \quad k=0,1,2, \ldots,
$$

- Bounded time-invariant cost function $g(x, u, w) \in[-M, M]$


## Definition (Optimal Discounted Cost)

Given a discounting factor $\gamma>0$, the discounted expected cost for policy $\pi=\left(\mu_{0}, \mu_{1}, \ldots,\right)$ is:

$$
J_{\pi}(a):=\mathbb{E}_{\left\{W_{k}\right\}}\left[\sum_{k=0}^{\infty} \gamma^{k} g\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right) \mid X_{0}=a\right]
$$

The optimal infinite-horizon discounted cost is $J^{*}(a):=\min _{\pi} J_{\pi}(a)$

## A Closer Look at the Finite-Horizon Discounted Cost Problem

- The finite-horizon cost for our problem and policy $\pi . \forall L<N$ :

$$
\begin{aligned}
& J_{0 \rightarrow N, \pi}(a) \\
& =\mathbb{E}_{\mid X_{0}=a}\left[\sum_{k=0}^{L-1} \gamma^{k} g\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right)+\sum_{k=L}^{N-1} \gamma^{k} g\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right)+\gamma^{N} g_{N}\left(X_{N}\right)\right] \\
& \leq \mathbb{E}_{\mid X_{0}=a}\left[\sum_{k=0}^{L-1} \gamma^{k} g\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right)\right]+\gamma^{L} g_{L}\left(X_{L}\right)+\sum_{k=L}^{N} \gamma^{k} M-\gamma^{L} g_{L}\left(X_{L}\right) \\
& \leq J_{0 \rightarrow L, \pi}(a)+M \gamma^{L}\left(1+\frac{1-\gamma^{N-L+1}}{1-\gamma}\right)
\end{aligned}
$$

- Let $N \rightarrow \infty$ and take $\min _{\pi}$ on both sides:

$$
J^{*}(a):=\min _{\pi} \lim _{N \rightarrow \infty} J_{0 \rightarrow N, \pi}(a) \leq \min _{\pi} J_{0 \rightarrow L, \pi}(a)+M \gamma^{L} \frac{2-\gamma}{1-\gamma}
$$

Similarly, we obtain

$$
J^{*}(a) \geq J_{0 \rightarrow L}^{*}(a)-M \gamma^{L} \frac{2-\gamma}{1-\gamma}
$$

## Optimal Infinite-Horizon Discounted Cost as a Limit

By a sandwiching argument and $L \rightarrow \infty$ :

## Theorem

The Optimal Infinite-Horizon Discounted Cost can be obtained as:

$$
J^{*}(a)=\lim _{L \rightarrow \infty} J_{0 \rightarrow L}^{*}(a), \quad \forall a \in \mathcal{X}
$$

irrespective of the termination costs $\left\{\gamma^{L} g_{L}\left(X_{L}\right)\right\}$.

- Is there a way to efficiently compute this limit?
$\rightarrow$ Yes, because of time-invariance and since the starting point does not matter!


## Rephrasing the Finite-Horizon Cost

- Finite-horizon Optimal DP algorithm:

$$
J_{i \rightarrow N}^{*}(a):=\min _{\mu} \mathbb{E}_{W_{i}}\left[\gamma^{i} g\left(a, \mu(a), W_{i}\right)+J_{i+1 \rightarrow N}^{*}\left(f\left(a, \mu(a), W_{i}\right)\right)\right]
$$

for starting condition $J_{N \rightarrow N}^{*}(a):=\gamma^{N} g_{N}(a)$ for all $a \in \mathcal{X}$.

- For $i<N$ define $V_{N-i}(a):=\frac{1}{\gamma^{i}} J_{i \rightarrow N}^{*}(a)$ and $W_{N-i}^{\prime}:=W_{i}$, and $k=N-i$ :

$$
\begin{aligned}
& V_{0}(a)=J_{N \rightarrow N}^{*}(a) \\
& V_{k}(a)=\min _{\mu} \mathbb{E}_{W_{k}^{\prime}}\left[g\left(a, \mu(a), W_{k}^{\prime}\right)+\gamma V_{k-1}\left(f\left(a, \mu(a), W_{k}^{\prime}\right)\right)\right], \quad k=1, \ldots, N
\end{aligned}
$$

- Recursion independent of $N$ and $\forall N: V_{N}(a)=J_{0 \rightarrow N}^{*}(a)$ ! (with same $g_{N}$.)


## Lemma

$$
J^{*}(a)=\lim _{N \rightarrow \infty} V_{N}(a)
$$

where

$$
V_{k}=\min _{\mu} \mathbb{E}\left[g+\gamma V_{k-1}\right], \quad k=1,2, \ldots,
$$

and starting vector $V_{0}$ can be arbitrary.

## The Value-Iteration Algorithm for Dynamic Programming

- Finds an approximation to the solution vector $J^{*}$ for an infinite-horizon DP problem with discounted and bounded costs
- Algorithm:
- Select an arbitrary starting vector $V_{0} \in \mathbb{R}^{|\mathcal{X}|}$
- For $k=1,2, \ldots$, calculate for each $a \in \mathcal{X}$ :

$$
V_{k}(a)=\min _{\mu} \mathbb{E}_{W}\left[g(a, \mu(a), W)+\gamma V_{k-1}(f(a, \mu(a), W)] .\right.
$$

- Stop according to some convergence criterion, for example when the value on each component does not change more than a given value $\epsilon$.
- How fast does it converge? Error bounds?
- Attention: In the literature $V$ is often also called $J$


## Exponential Decay on Difference of Iterations

## Lemma

Given two bounded initial vectors $V_{0}$ and $V_{0}^{\prime}$ such that

$$
\max _{a \in \mathcal{X}}\left|V_{0}(a)-V_{0}^{\prime}(a)\right| \leq c .
$$

If $V_{1}, \ldots, V_{k}$ and $V_{1}^{\prime}, \ldots, V_{k}^{\prime}$ are obtained from the DP recursion for $V_{0}$ and $V_{0}^{\prime}$, respectively:

$$
\max _{a \in \mathcal{X}}\left|V_{k}(a)-V_{k}^{\prime}(a)\right| \leq \alpha^{k} \max _{a \in \mathcal{X}}\left|V_{0}(a)-V_{0}^{\prime}(a)\right| .
$$

Proof: By induction:

$$
\begin{aligned}
V_{1}(a) & =\min _{\mu} \mathbb{E}_{W}\left[g(a, \mu(a), W)+\gamma V_{0}(f(a, \mu(a), W))\right] \\
& \leq \min _{\mu} \mathbb{E}_{W}\left[g(a, \mu(a), W)+\gamma V_{0}^{\prime}(f(a, \mu(a), W))\right]+\gamma c=V_{1}^{\prime}(a)+\gamma c \\
V_{k}(a) & =\min _{\mu} \mathbb{E}_{W}\left[g(a, \mu(a), W)+\gamma V_{k-1}(f(a, \mu(a), W))\right] \\
& \leq \min _{\mu} \mathbb{E}_{W}\left[g(a, \mu(a), W)+\gamma V_{k-1}^{\prime}(f(a, \mu(a), W))\right]+\gamma \gamma^{k-1} c=V_{k}^{\prime}(a)+\gamma^{k} c
\end{aligned}
$$

Similarly, $V_{1}(a) \geq V_{1}^{\prime}(a)-\gamma c$ and $V_{k}(a) \geq V_{k}^{\prime}(a)-\gamma^{k} c$

## Error Bounds on the Value-Iteration Algorithm

- By Bellman's equation ahead, $V_{0}^{\prime}=J^{*}$ implies $V_{1}^{\prime}=\cdots V_{k}^{\prime}=J^{*}$ and thus

$$
\max _{a \in \mathcal{X}}\left|V_{k}(a)-J^{*}(a)\right| \leq \alpha^{k} \max _{a \in \mathcal{X}}\left|V_{0}(a)-J^{*}(a)\right| .
$$

- The error in the value-iteration algorithm vanishes exponentially fast with each iteration


## The Operator Interpretation

- Operator $\mathbb{T}$ (or $\mathbb{T}_{f, g, \gamma}$ ) acts on vector $V \in \mathcal{R}^{|\mathcal{X}|}$ componentwise as:

$$
(\mathbb{T} V)(a)=\min _{\mu} \mathbb{E}_{W}[g(a, \mu(a), W)+\gamma V(f(a, \mu(a), W))], \quad \forall a \in \mathcal{X} .
$$

- Optimal DP iteration is described as: $V_{k+1}=\mathbb{T} V_{k}$.
- The operator $\mathbb{T}$ is contracting since $\exists \rho \in(0,1)$ :

$$
\left\|\mathbb{T}(J)-\mathbb{T}\left(J^{\prime}\right)\right\| \leq \rho\left\|J-J^{\prime}\right\|, \quad \forall J, J^{\prime},
$$

where here $\|\cdot\|$ denotes the infinity norm (i.e., the maximum component)

- Irrespective of $V$, as $k \rightarrow \infty$ the operator $\mathbb{T}^{k} V=\underbrace{\mathbb{T}(\mathbb{T}(\cdots \mathbb{T}}_{k \text { applications of } \mathbb{T}}(V)))$
converges to a unique $J^{*}$ that satisfies the fix-point equation

$$
J^{*}=\mathbb{T} J^{*}
$$

## Bellman's Equation

## Theorem

The cost vector $J^{*}$ is optimal if, and only if, it satisfies

$$
J^{*}(a)=\min _{\mu} \mathbb{E}_{W}\left[g(a, \mu(a), W)+\gamma J^{*}(f(a, \mu(a), W))\right], \quad \forall a \in \mathcal{X}
$$

There is a unique finite cost-vector $J^{*}$ satisfying above equation.
Proof: "If"-direction: Set $J^{*}$ as starting vector in iteration. "Only if'-direction uses the previous bounds. $\forall a \in \mathcal{X}$ :

$$
\begin{aligned}
J^{*}(a)-M \gamma^{L+1} \frac{2-\gamma}{1-\gamma} & \leq V_{L+1}=\min _{\mu} \mathbb{E}_{W}\left[g(a, \mu(a), W)+\gamma V_{L}(f(a, \mu(a), W))\right] \\
& \leq \min _{\mu} \mathbb{E}_{W}\left[g(a, \mu(a), W)+\gamma J^{*}(f(a, \mu(a), W))\right]+M \gamma^{L} \frac{2-\gamma}{1-\gamma}
\end{aligned}
$$

Similarly:

$$
J^{*}(a)+M \gamma^{L+1} \frac{2-\gamma}{1-\gamma} \geq \min _{\mu} \mathbb{E}_{W}\left[g(a, \mu(a), W)+\gamma J^{*}(f(a, \mu(a), W))\right]-M \gamma^{L} \frac{2-\gamma}{1-\gamma} .
$$

Taking $L \rightarrow \infty$ by sandwiching argument proves "only-if" direction. Uniqueness follows by convergence of $\left\{V_{k}\right\}_{k \geq 0}$ irrespective of $V_{0}$.

## About Stationary Policies

- A policy of the form $\pi=(\mu, \mu, \mu, \ldots)$ is called stationary.
- For any stationary policy $\mu$ and arbitrary initial vector $V_{0}$ :

$$
V_{k, \mu}(a)=\mathbb{E}_{W}\left[g(a, \mu(a), W)+\gamma V_{k-1, \mu}(f(a, \mu(a), W))\right]
$$

converges for each $a \in \mathcal{X}$. Call the convergence point $J_{\mu}(a)$.

- If $V_{1, \mu}(a) \leq V_{0, \mu}(a)$ for all $a \in \mathcal{X}$, then $V_{k, \mu}$ is a decreasing sequence


## Lemma (Optimality of Stationary Policies)

A stationary policy $\mu^{*}$ is optimal if, and only if,

$$
\begin{aligned}
& \mathbb{E}_{W}\left[g\left(a, \mu^{*}(a), W\right)+\gamma J^{*}\left(f\left(a, \mu^{*}(a), W\right)\right)\right] \\
& \quad=\min _{\mu} \mathbb{E}_{W}\left[g(a, \mu(a), W)+\gamma J^{*}(f(a, \mu(a), W))\right], \quad \forall a \in \mathcal{X} .
\end{aligned}
$$

Proof: Follows essentially from Bellman's equation and the uniqueness of the solution $J^{*}$.

## Finding an Improved Stationary Policy

## Theorem

Let $\mu$ and $\bar{\mu}$ be stationary policies satisfying $\forall a \in \mathcal{X}$ :

$$
\mathbb{E}_{W}\left[g(a, \bar{\mu}(a), W)+\gamma J_{\mu}(a, \bar{\mu}(a), W)\right]=\min _{u} \mathbb{E}_{W}\left[g(a, u)+\gamma J_{\mu}(f(a, u, W))\right] .
$$

Then,

$$
J_{\bar{\mu}}(a) \leq J_{\mu}(a), \quad \forall a \in \mathcal{X}
$$

where inequality is strict for at least one $a \in \mathcal{X}$ whenever $\mu$ is not optimal.
Proof:

$$
\begin{aligned}
\underbrace{J_{\mu}(a)}_{V_{0, \bar{\mu}}} & =\mathbb{E}\left[g(a, \mu(a), W)+\gamma J_{\mu}(a)(f(a, \mu(a), W))\right] \\
& \geq \underbrace{\mathbb{E}\left[g(a, \bar{\mu}(a), W)+\gamma J_{\mu}(a)(f(a, \bar{\mu}(a), W))\right]}_{V_{1, \bar{\mu}}} \\
& \geq V_{2, \bar{\mu}} \geq V_{3, \bar{\mu}} \geq \ldots \\
& \geq J_{\bar{\mu}}(a) .
\end{aligned}
$$

## Policy Iteration Algorithm

- Finds the exact solution vector $J^{*}$ for an infinite-horizon DP problem with discounted and bounded costs
- Algorithm:
- Select an arbitrary policy $\mu_{0}$ and find $J_{\mu_{0}}$ by solving the linear system of equations:

$$
J_{\mu_{0}}(a)=\mathbb{E}\left[g\left(a, \mu_{0}(a), W\right)\right]+\gamma \mathbb{E}\left[J_{\mu_{0}}\left(f\left(a, \mu_{0}(a), W\right)\right)\right], \quad a \in \mathcal{X} .
$$

- For $k=1,2, \ldots$ solve the minimization problem

$$
\mu_{k}(a):=\operatorname{argmin}_{u \in \mathcal{U}} \mathbb{E}_{W}\left[g(a, u, W)+\gamma J_{\mu_{k-1}}(f(a, u, W)], \quad a \in \mathcal{X} .\right.
$$

and find $J_{\mu_{k}}$ by solving the linear system of equations:

$$
J_{\mu_{k}}(a)=\mathbb{E}\left[g\left(a, \mu_{k}(a), W\right)\right]+\gamma \mathbb{E}\left[J_{\mu_{k}}\left(f\left(a, \mu_{k}(a), W\right)\right)\right], \quad a \in \mathcal{X} .
$$

- Stop when $\mu_{k}=\mu_{k-1}$ and produce $J^{*}=J_{\mu_{k-1}}$
- Advantage: There is only a finite number of stationary policies and thus the algorithm finds the exact optimal discounted cost $J^{*}$.


## A Simple Binary Example

- Let $\mathcal{X}=\{a, b\}$ and $\mathcal{U}=\{1,2\}$. Moreover, $W_{i} \sim \mathcal{B}(1 / 4)$ and $\gamma=0.9$.
- Transition function: $f(x, u, w)=a$ if $(u=1, w=1)$ or $\left(u_{2}=2, w=0\right)$, and $f(x, u, w)=b$ else
- Cost function: $\mathbb{E}_{w}[g(a, 1, W)]=2, \mathbb{E}_{w}[g(a, 2, W)]=0.5$, $\mathbb{E}_{w}[g(b, 1, W)]=1, \quad \mathbb{E}_{w}[g(b, 2, W)]=3$.
- Value iteration algorithm with starting point $V_{0}=(0,0)^{\top}$ :

$$
\begin{aligned}
V_{1}(a) & =\min _{\mu}\left(\mathbb{E}[g(a, \mu(a), W)]+\mathbb{E}\left[\gamma V_{0}(f(a, \mu(a), W))\right]\right) \\
& =\min _{u \in\{1,2\}} \mathbb{E}[g(a, u, W)]=\min \{2,0.5\}=0.5 . \\
V_{1}(b) & =\min _{u \in\{1,2\}} \mathbb{E}[g(a, u, W)]=\min \{1,3\}=1 . \\
V_{2}(a) & =\min \left\{\mathbb{E}\left[g(a, 1, W)+\gamma V_{1}(f(a, 1, W))\right], \mathbb{E}\left[g(a, 2, W)+\gamma V_{1}(f(a, 2, W))\right]\right\} \\
& =\min \{2+0.9 \cdot(0.5 \cdot 3 / 4+1 \cdot 1 / 4), 0.5+0.9 \cdot(0.5 \cdot 1 / 4+1 \cdot 3 / 4)\} \\
& =\min \{2+0.9 \cdot 5 / 8,0.5+0.9 \cdot 7 / 8\}=0.5+0.9 \cdot 7 / 8=1.2875 \\
V_{2}(b) & =\min \{1+0.9 \cdot 5 / 8,3+0.9 \cdot 7 / 8\}=1+0.9 \cdot 5 / 8=1.5625
\end{aligned}
$$

## Example Continued

- Value iteration algorithm continued:

$$
\begin{array}{cc}
V_{3}(a)=1.844 & V_{3}(b)=2.220 \\
V_{4}(a)=2.414 & V_{4}(b)=2.745 \\
V_{5}(a)=2.896 & V_{5}(b)=3.247 \\
\vdots & \vdots \\
V_{15}(a)=5.783 & V_{15}(b)=6.128
\end{array}
$$

- Policy iteration algorithm with initial policy $\mu_{0}(a)=1$ and $\mu_{0}(b)=2$ :
- Policy evaluation to determine $J_{\mu_{0}}$ :

$$
\begin{aligned}
& J_{\mu_{0}}(a)=2+0.9 \cdot\left(J_{\mu_{0}}(a) \cdot 3 / 4+J_{\mu_{0}}(b) \cdot 1 / 4\right) \\
& J_{\mu_{0}}(b)=3+0.9 \cdot\left(J_{\mu_{0}}(a) \cdot 1 / 4+J_{\mu_{0}}(b) \cdot 3 / 4\right) \\
& \Rightarrow J_{\mu_{0}}=\binom{2}{3}+\underbrace{\left(\begin{array}{cc}
0.9 \cdot 3 / 4 & 0.9 \cdot 1 / 4 \\
0.9 \cdot 1 / 4 & 0.9 \cdot 3 / 4
\end{array}\right)}_{\substack{\text { state transition matrix } \\
P_{\mu_{0}} \text { from } X_{0} \text { to } X_{1}}} J_{\mu_{0}}=\binom{24.091}{25.909}
\end{aligned}
$$

## Example Continued II

- Policy improvement to determine $\mu_{1}$ :

$$
\begin{aligned}
& \mu_{1}(a)= 1+\mathbb{1}\left\{\mathbb { E } _ { W } \left[g(a, 1, W)+\gamma J_{\mu_{0}}(f(a, 1, W)]\right.\right. \\
&>\mathbb{E}_{W}\left[g(a, 2, W)+\gamma J_{\mu_{0}}(f(a, 2, W)]\right\} \\
&= 1+\mathbb{1}\{2+0.9 \cdot 3 / 4 \cdot 24.091+0.9 \cdot 1 / 4 \cdot 25.909 \\
&>0.5+0.9 \cdot 1 / 4 \cdot 24.091+0.9 \cdot 3 / 4 \cdot 25.909\} \\
&= 1+\mathbb{1}\{24.909>23.409\}=2 \\
& \mu_{1}(b)=1+\mathbb{1}\{1+0.9 \cdot 3 / 4 \cdot 24.091+0.9 \cdot 1 / 4 \cdot 25.909 \\
&>3+0.9 \cdot 1 / 4 \cdot 24.091+0.9 \cdot 3 / 4 \cdot 25.909\} \\
&= 1+\mathbb{1}\{22.909>25.909\}=1
\end{aligned}
$$

- Policy evaluation to determine $J_{\mu_{1}}$ :

$$
\begin{aligned}
J_{\mu_{1}}(a)= & 0.5+0.9 \cdot\left(J_{\mu_{1}}(a) \cdot 1 / 4+J_{\mu_{1}}(b) \cdot 3 / 4\right) \\
J_{\mu_{1}}(b)= & 1+0.9 \cdot\left(J_{\mu_{1}}(a) \cdot 3 / 4+J_{\mu_{1}}(b) \cdot 1 / 4\right) \\
& \Rightarrow J_{\mu_{1}}=\underbrace{\binom{0.5}{1}+\left(\begin{array}{cc}
0.9 \cdot 1 / 4 & 0.9 \cdot 3 / 4 \\
0.9 \cdot 3 / 4 & 0.9 \cdot 1 / 4
\end{array}\right)}_{\begin{array}{c}
\text { state transition matrix } \\
P_{\mu_{1}} \text { from } X_{1} \text { to } X_{2}
\end{array}} J_{\mu_{1}}=\binom{7.3276}{7.6724}
\end{aligned}
$$

## Example Continued III

- Policy improvement to determine $\mu_{2}$ :

$$
\begin{aligned}
& \mu_{2}(a)= 1+\mathbb{1}\left\{\mathbb { E } _ { W } \left[g(a, 1, W)+\gamma J_{\mu_{1}}(f(a, 1, W)]\right.\right. \\
&>\mathbb{E}_{W}\left[g(a, 2, W)+\gamma J_{\mu_{1}}(f(a, 2, W)]\right\} \\
&= 1+\mathbb{1}\{2+0.9 \cdot 3 / 4 \cdot 27.3276+0.9 \cdot 1 / 4 \cdot 7.6724 \\
&>0.5+0.9 \cdot 1 / 4 \cdot 7.3276+0.9 \cdot 3 / 4 \cdot 7.6724\} \\
&= 1+\mathbb{1}\{8,6724>7.3276\}=2 \\
& \mu_{2}(b)=1+\mathbb{1}\{1+0.9 \cdot 3 / 4 \cdot 7.3276+0.9 \cdot 1 / 4 \cdot 7.6724 \\
&>3+0.9 \cdot 1 / 4 \cdot 7.3276+0.9 \cdot 3 / 4 \cdot 7.6724\} \\
&= 1+\mathbb{1}\{7.6724>9.8276\}=1
\end{aligned}
$$

- Notice that policy $\mu_{2}=\mu_{1}$ ! So, we terminate.
- $\mu_{1}, \mu_{2}$ are optimal policies and $J^{*}=J_{\mu_{1}}$


# Sequential Decision Processes, Master MICAS, Part I 

Michèle Wigger

Telecom Paris, 18 December 2020


## Lecture 4- LP Approach to Discounted Infinite-Horizon Dynamic Programming

## Review of Lecture 3: The Discounted Case

- Time-invariant discrete-time dynamic system:

$$
X_{k+1}=f\left(X_{k}, U_{k}, W_{k}\right), \quad k=0,1,2, \ldots,
$$

- Bounded time-invariant cost function $g(x, u, w) \in[-M, M]$
- Optimal discounted infinite-horizon cost:

$$
J^{*}(a):=\min _{\pi} \mathbb{E}_{\left\{w_{k}\right\}}\left[\sum_{k=0}^{\infty} \gamma^{k} g\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right) \mid X_{0}=a\right]
$$

- Bellman's Equation: Optimal cost function $J^{*}(a)$ satisfies

$$
J^{*}(a)=\min _{\mu} \mathbb{E} W\left[g(a, \mu(a), W)+\gamma J^{*}(f(a, \mu(a), W))\right], \quad \forall a \in \mathcal{X} .
$$

## Review of Lecture 3, continued

- Value iteration algorithm based on the fact:

$$
\lim _{k \rightarrow \infty} V_{k}(a)=J^{*}(a)
$$

for any starting vector $V_{0}$ and

$$
\begin{equation*}
V_{k+1}(a)=\min _{\mu} \mathbb{E}_{W}\left[g(a, \mu(a), W)+\gamma V_{k}(f(a, \mu(a), W))\right], \quad k=0,1,2, \ldots \tag{1}
\end{equation*}
$$

$\rightarrow$ Start with $V_{0}=\mathbf{0}$ and apply iteration (1) until satisfied with precision

- Policy iteration algorithm based on the following fact:

$$
\begin{align*}
& \mathbb{E}_{W}\left[g\left(a, \mu_{k+1}(a), W\right)+\gamma J_{\mu_{k}}\left(a, \mu_{k+1}(a), W\right)\right]=\min _{u} \mathbb{E}_{W}\left[g(a, u)+\gamma J_{\mu_{k}}(f(a, u, W))\right] \\
& \quad \text { then } J_{\mu_{k+1}}(a) \leq J_{\mu_{k}}(a), \quad \forall a \in \mathcal{X} \tag{2}
\end{align*}
$$

$\rightarrow$ Start with any policy $\mu_{0}$, and apply policy iteration in (2)

## Dynamic Programming Operator and Monotonicity

## Definition (Dynamic Programming Operator)

Operator $\mathbb{T}$ (or $\mathbb{T}_{f, g, \gamma}$ ) acts on vector $V \in \mathcal{R}^{|\mathcal{X}|}$ componentwise as:

$$
(\mathbb{T} V)(a):=\min _{\mu} \mathbb{E}_{W}[g(a, \mu(a), W)+\gamma V(f(a, \mu(a), W))], \quad \forall a \in \mathcal{X} .
$$

- Monotonicity of $\mathbb{T}$ : If $V(a) \leq(\mathbb{T} V(a)$ for all $a \in \mathcal{X}$, then

$$
\begin{equation*}
V(a) \leq(\mathbb{T} V)(a) \leq\left(\mathbb{T}^{2} V\right)(a) \leq \cdots J^{*}(a) \tag{3}
\end{equation*}
$$

- The optimal cost vector $J^{*}$ satisfies (3) by Bellman's equation: $\left(\mathbb{T} J^{*}\right)=J^{*}$
- Thus $J^{*}$ is the largest vector satisfying $V(a) \leq(\mathbb{T} V)(a)$ for all $a \in \mathcal{X}$.
- Since $\mathbb{T}$ contains a $\min , V(a) \leq(\mathbb{T} V)(a)$ is equivalent to:

$$
V(a) \leq \mathbb{E}_{W}[g(a, \mu(a), W)+\gamma V(f(a, \mu(a), W))], \quad \forall a \in \mathcal{X}, \text { and } \forall \mu .
$$

## Linear Programming Approach to find Vector J*

- Let $\mathcal{X}=\{1, \ldots, m\}$ and $J(i)=J_{i}$.
- Pick positive weights $p_{0}(1), \ldots, p_{0}(m)$ summing to 1 and solve


## Linear Programming Optimization Problem

$$
\max _{J_{1}, \ldots, J_{m}}(1-\gamma) \sum_{i=1}^{m} p_{0}(i) J_{i}
$$

subject to:

$$
J_{i} \leq \mathbb{E}_{W}[g(i, u, W)]+\gamma \cdot \sum_{j=1}^{m} P_{u, i j} J_{j}, \quad \forall i, u
$$

where $P_{u, i j}:=\operatorname{Pr}[f(i, u, W)=j]$
(Indices $i$ and $j$ were mixed up in the previous version of the slides! Also, we used policy $\mu$ instead of action $u$. We can use a single action $u$ because for each $i$ the constraint only depends on the single action in state i)

- Problem: the number of constraints can be huge.


## Basic Optimization Theory: Primal-Dual LP Problems

## Primal Problem

$$
\max _{x_{1}, \ldots, x_{n}} \sum_{j=1}^{n} c_{j} x_{j}
$$

subject to

$$
\sum_{j=1}^{n} a_{i, j} x_{j} \leq b_{i}, \quad i=1, \ldots, m
$$

## Dual Problem

$$
\min _{\lambda_{1}, \ldots, \lambda_{m}} \sum_{i=1}^{m} b_{i} \lambda_{i}
$$

subject to

$$
\begin{aligned}
& \sum_{i=1}^{m} a_{i, j} \lambda_{i}=c_{j}, \quad j=1, \ldots, n \\
& \lambda_{i} \geq 0, \quad i=1, \ldots, m
\end{aligned}
$$

- Solution has at most $L$ non-degenerate components (i.e., components satisfying the constraints with strict inequalities)


## The Dual Optimization Problem to the LP on the Previous Slide

## Dual Problem

$$
\min _{\{\rho(i, u)\}} \sum_{i=1}^{m} \sum_{u} \mathbb{E}_{W}[g(i, u, W)] \cdot \rho(i, u)
$$

subject to:

$$
\begin{equation*}
\sum_{u} \rho(i, u)-\sum_{j=1}^{m} \sum_{u} \gamma P_{u, i j} \cdot \rho(j, u)=(1-\gamma) p_{0}(i), \quad \forall i=1, \ldots, m \tag{4}
\end{equation*}
$$

where $P_{u, i j}:=\operatorname{Pr}[f(i, u, W)=j]$ and $\rho(i, u) \geq 0$ for all $i, u$.

- Solutions of linear programs are at the extreme points (corner points) of the intersection plane defined by the $m$ constraints (4)
$\rightarrow \exists$ an optimal solution $\rho^{*}(i, u)$ with only $m$ components $\rho^{*}(i, u)>0$
- If $\rho(i, u)=0 \forall u$ for a specific $i$, then (4) cannot be satisfied for this $i$ (the two sides (4) have different signs for constraint $i$ )
$\Rightarrow$ For each $i=1, \ldots, m$ there is exactly one $\rho^{*}(i, u)>0$
There exists an optimal stationary deterministic policy $\mu^{*}(u \mid i)=\frac{\rho^{*}(i, u)}{\sum_{v} \rho^{*}(i, v)}$


## The Dual Optimization Problem to the LP on the Previous Slide

## Dual Problem

$$
\min _{\{\rho(i, u)\}} \sum_{i=1}^{m} \sum_{u} \mathbb{E}_{W}[g(i, u, W)] \cdot \rho(i, u)
$$

subject to:

$$
\begin{equation*}
\sum_{u} \rho(i, u)-\sum_{j=1}^{m} \sum_{u} \gamma P_{u, i j} \cdot \rho(j, u)=(1-\gamma) p_{0}(i), \quad \forall i=1, \ldots, m \tag{4}
\end{equation*}
$$

where $P_{u, i j}:=\operatorname{Pr}[f(i, u, W)=j]$ and $\rho(i, u) \geq 0$ for all $i, u$.

- Summing both sides of (4) over $i=1, \ldots, m$ shows that for any feasible $\rho(i, u)$ :

$$
\sum_{i=1}^{m} \sum_{u} \rho(i, u)=\sum_{i=1}^{m} p_{0}(i)=1
$$

So any feasible $\rho(i, u)$ can be a probability distribution over the states and actions.

## Randomized Policies

- A stationary randomized policy $\mu$ chooses action $U_{k}=u$ with probability $\mu(u \mid i)$ when $X_{k}=i$
- We start with a random initial state $X_{0} \sim p_{0}$ and calculate the expected discounted cost of this randomized policy

$$
\begin{aligned}
J_{\mu}\left(p_{0}\right) & :=\lim _{N \rightarrow \infty} \sum_{k=0}^{N} \gamma^{k} \mathbb{E}\left[g\left(X_{k}, \mu\left(X_{k}\right), W\right)\right] \\
& =\lim _{N \rightarrow \infty} \sum_{k=0}^{N} \sum_{w} \sum_{i=1}^{m} \sum_{u} \gamma^{k} g(i, u, w) \mu(u \mid i) P_{X_{k}}(i) P_{W}(w)
\end{aligned}
$$

whre $P_{X_{k}}(i)$ depends on the initial distribution $p_{0}$, and of course the stationary randomized policy $\mu$ and the state-transition function $f(\cdot, \cdot, \cdot)$.

## State-Action Frequencies (also called Occupation Measures)

- Given an infinite-horizon policy $\pi$ and initial state-distribution $p_{0}(i)=\operatorname{Pr}\left[X_{0}=i\right]$, define the state-action frequency:

$$
\rho_{p_{0}}^{\pi}(i, u):=(1-\gamma) \sum_{k=0}^{\infty} \gamma^{k} P_{p_{0}, k}^{\pi}(i, u), \quad i=1, \ldots, m
$$

where $P_{P_{0}, k}^{\pi}(i, u)=\operatorname{Pr}\left[X_{k}=i, U_{k}=u\right]$ under policy $\pi$ and initial state-distribution $p_{0}$.

- Define the state-frequency

$$
\rho_{P_{0}}^{\pi}(i):=\sum_{u} \rho_{p_{0}}^{\pi}(i, u)=(1-\gamma) \sum_{k=0}^{\infty} \gamma^{k} P_{p_{0}, k}^{\pi}(i), \quad i=1, \ldots, m
$$

- Under policy $\pi$ and initial state-distribution $p_{0}$ :

$$
\begin{aligned}
& =(1-\gamma) J_{\pi}\left(p_{0}\right) \quad=(1-\gamma) \sum_{k=0}^{\infty} \gamma^{k} \mathbb{E}\left[g\left(X_{k}, U_{k}, W_{k}\right)\right] \\
& =(1-\gamma) \sum_{k=0}^{\infty} \gamma^{k} \sum_{i, u} \mathbb{E}\left[g\left(i, u, W_{k}\right)\right] P_{p_{0}, k}^{\pi}(i, u) \\
& =(1-\gamma) \sum_{i, u} \mathbb{E}\left[g\left(i, u, W_{k}\right)\right] \sum_{k=0}^{\infty} \gamma^{k} P_{p_{0}, k}^{\pi}(i, u)=\sum_{i, u} \mathbb{E}\left[g\left(i, u, W_{k}\right)\right] \rho_{p_{0}}^{\pi}(i, u)
\end{aligned}
$$

## Stationary Randomized Policy Deduced from State-Action Frequencies

- Given $\pi$, define a stationary randomized policy $\tilde{\pi}=\left(\mu_{p_{0}}^{\pi}, \mu_{p_{0}}^{\pi}, \ldots,\right)$ as

$$
\mu_{P_{0}}^{\pi}(u \mid i):=\frac{\rho_{P_{0}}^{\pi}(i, u)}{\rho_{P_{0}}^{\pi}(i)}, \quad \text { if } \rho_{p_{0}}^{\pi}(i)>0
$$

and $\mu_{p_{0}}^{\pi}(u \mid i)$ arbitrary if $\rho_{\rho_{0}}^{\pi}(i)=0$. (From any state-action frequencies $\rho(i, u)>0$ one can derive a stationary policy.)

- Under policy $\mu=\mu_{p_{0}}^{\pi}$ (proof on next slide):

$$
\rho_{p_{0}}^{\mu}(i, u)=\rho_{p_{0}}^{\pi}(i, u), \quad \forall i, u
$$

- Therefore:

$$
\begin{aligned}
(1-\gamma) J_{\mu}\left(p_{0}\right) & =\sum_{i, u} \mathbb{E}\left[g\left(i, u, W_{k}\right)\right] \rho_{\rho_{0}}^{\mu}(i, u) \\
& =\sum_{i, u} \mathbb{E}\left[g\left(i, u, W_{k}\right)\right] \rho_{p_{0}}^{\pi}(i, u)=(1-\gamma) J_{\pi}\left(p_{0}\right)
\end{aligned}
$$

$\Rightarrow$ For any $\pi$ there is an equally-good stationary randomized policy $\mu$ $\Rightarrow$ Without loss in performance one can restrict to stationary policies

Proof that $\rho_{p_{0}}^{\mu}(i, u)=\rho_{p_{0}}^{\pi}(i, u)$

$$
\begin{align*}
& (1-\gamma)^{-1} \rho_{\rho_{0}}^{\pi}(i) \\
& =\sum_{k=0}^{\infty} \gamma^{k} P_{P_{0}, k}^{\pi}(i)=p_{0}(i)+\sum_{k=1}^{\infty} \gamma^{k} P_{P_{0}, k}^{\pi}(i) \\
& \stackrel{k^{\prime}}{ }=k-1 p_{0}(i)+\gamma \sum_{k^{\prime}=0}^{\infty} \gamma^{k^{\prime}} P_{p_{0}, k^{\prime}+1}^{\pi}(i) \\
& =p_{0}(i)+\gamma \sum_{k^{\prime}=0}^{\infty} \gamma^{k^{\prime}} \operatorname{Pr}\left[X_{k^{\prime}+1}=i\right] \\
& =p_{0}(i)+\gamma \sum_{k^{\prime}=0}^{\infty} \gamma^{k^{\prime}} \sum_{j, u} \operatorname{Pr}_{\pi}\left[X_{k^{\prime}}=j, U_{k^{\prime}}=u\right] \cdot \operatorname{Pr}\left[X_{k^{\prime}+1}=i \mid X_{k^{\prime}}=j, U_{k^{\prime}}=u\right] \\
& =p_{0}(i)+\gamma \sum_{j, u} \sum_{k^{\prime}=0}^{\infty} \gamma^{k^{\prime}} \operatorname{Pr}_{\pi}\left[X_{k^{\prime}}=j, U_{k^{\prime}}=u\right] \cdot P_{u, j i} \\
& =p_{0}(i)+\frac{\gamma}{1-\gamma} \sum_{j, u} \rho_{\rho_{0}}^{\pi}(j, u) \cdot P_{u, j i}  \tag{5}\\
& =p_{0}(i)+\frac{\gamma}{1-\gamma} \sum_{j} \rho_{\rho_{0}}^{\pi}(j) \cdot \underbrace{\sum_{u} \mu(u \mid j) \cdot P_{u, j i}}_{=P_{\mu, j i}}=p_{0}(i)+\frac{\gamma}{1-\gamma} \sum_{j} \rho_{p_{0}}^{\pi}(j) \cdot P_{\mu, j i}
\end{align*}
$$

Proof that $\rho_{p_{0}}^{\mu}(i, u)=\rho_{p_{0}}^{\pi}(i, u)$ continued

- Vectors $\boldsymbol{\rho}_{\rho_{0}}^{\pi}:=\left(\rho_{\rho_{0}}^{\pi}(1), \ldots, \rho_{\rho_{0}}^{\pi}(m)\right)$ and $\mathbf{p}_{0}:=\left(p_{0}(1), \ldots, p_{0}(m)\right)$ (Attention: changed to row-vectors for simplicity.)
- $\mathrm{P}_{\mu}$ the matrix with row- $j$ and column- $i$ entry equal to $P_{\mu, j i}$
- Then:

$$
\boldsymbol{\rho}_{\rho_{0}}^{\pi}=(1-\gamma) \mathbf{p}_{0}+\gamma \boldsymbol{\rho}_{p_{0}}^{\pi} \mathrm{P}_{\mu}
$$

- Therefore:
$\boldsymbol{\rho}_{p_{0}}^{\pi}=(1-\gamma) \mathbf{p}_{0}\left(\mathrm{I}-\gamma \mathrm{P}_{\mu}\right)^{-1}=(1-\gamma) \mathbf{p}_{0} \cdot \sum_{k=0}^{\infty} \gamma^{k} \mathrm{P}_{\mu}^{k}=(1-\gamma) \sum_{k=0}^{\infty} \gamma^{k} \mathbf{P}_{p_{0}, k}^{\mu}=\boldsymbol{\rho}_{p_{0}}^{\mu}$,
where $\mathbf{P}_{p_{0}, k}^{\mu}$ is the vector with $i$-th entry equal to $P_{p_{0}, k}^{\mu}(i)$.


## Proof that $\rho_{p_{0}}^{\mu}(i, u)=\rho_{p_{0}}^{\pi}(i, u)$ continued II

- At the end of the previous slide we proved that the policies $\pi$ and $\mu$ have same state-frequencies:

$$
\rho_{P_{0}}^{\pi}(i)=\rho_{P_{0}}^{\mu}(i), \quad \forall i .
$$

- We now prove that the two policies also have same state-action frequencies:

$$
\begin{aligned}
\rho_{p_{0}}^{\pi}(i, u) & =\rho_{P_{0}}^{\pi}(i) \mu(u \mid i)=\rho_{p_{0}}^{\mu}(i) \mu(u \mid i) \\
& =(1-\gamma) \sum_{k=0}^{\infty} \gamma^{k} \operatorname{Pr}_{\mu}\left[X_{k}=i\right] \mu(u \mid i) \\
& =(1-\gamma) \sum_{k=0}^{\infty} \gamma^{k} \operatorname{Pr}_{\mu}\left[X_{k}=i, U_{k}=u\right]=\rho_{p_{0}}^{\mu}(i, u)
\end{aligned}
$$

## State-Action Frequencies are the Variables in the Dual Problem, Slide 7

For any stationary policy $\mu$, the state-action frequencies are feasible variables for the dual problem on slide 7 because $\rho_{\rho_{0}}^{\mu}(i, u)>0$ and by eq. (5) on slide 11 :

$$
\begin{equation*}
\underbrace{\sum_{u} \rho_{P_{0}}^{\mu}(i, u)}_{=\rho_{P_{0}}^{\mu}(i)}-\sum_{j=1}^{m} \sum_{u} \gamma \rho_{P_{0}}^{\mu}(j, u) P_{u, j i}=(1-\gamma) p_{0}(i), \quad \forall i \tag{6}
\end{equation*}
$$

Moreover,

$$
(1-\gamma) J_{\mu}\left(p_{0}\right)=\sum_{i, u} \mathbb{E}[g(i, u, W)] \rho_{p_{0}}^{\mu}(i, u)
$$

and thus minimizing above right-hand side over all $\rho(i, u)$ satisfying (6) yields the minimum discounted infinite-horizon cost $J^{*}\left(p_{0}\right)$. (Recall that for any $\rho(i, u)>0$ satisfying (6), it is possible to find a corresponding stationary policy $\mu$ s.t., $\rho(i, u)$ are the state-action frequencies of $\mu$.)

Dual variables can be interpreted as the state-action frequencies!

## Adding Constraints

- Can add a constraints on the cost to the linear programme on slide 6 !
- Determininistic policies might not be optimal anymore, but randomized policies can have better performances.


# Sequential Decision Processes, Master MICAS, Part I 

Michèle Wigger

Telecom Paris, 18 December 2020


## Lecture 5 - Multi-Armed Bandits and Unbounded Costs

## Problems with Retirement Option

- Consider an infinite-horizon problem with bounded cost-per-stage $|g(a, u, w)| \leq M$, where at each stage $k$ one can retire at cost $\gamma^{k} \cdot M_{\infty}$.
- Let $J_{\text {ret }}^{*}\left(a, M_{\infty}\right)$ be the optimal cost function for this problem. It satisfies the modified Bellman equation:
$J_{\text {ret }}^{*}\left(a, M_{\infty}\right)=\min \left\{M_{\infty}, \min _{\mu} \mathbb{E}_{W}\left[g(a, \mu(a), W)+\gamma J_{\text {ret }}^{*}\left(f(a, \mu(a), W), M_{\infty}\right)\right]\right\}$.
- If $M_{\infty} \geq \frac{1}{1-\gamma} M$, then never retire
- If $M_{\infty} \leq-\frac{1}{1-\gamma} M$, then retire immediately


## Optimal Policy under a Retirement Option

$$
\xrightarrow{J_{\text {ret }}^{*}\left(a, M_{\infty}\right)=M_{\infty}}
$$

- Define

$$
m(a):=\max \left\{M^{\prime}: J_{\text {ret }}^{*}\left(a, M^{\prime}\right)=M^{\prime}\right\}
$$

## Optimal Policy

Assume at stage $k$ we have $X_{k}=a$.

- Retire if

$$
m(a) \geq M_{\infty}
$$

- If $m(a)<M_{\infty}$, then play the optimal policy from Bellman's equation


## Multi-Armed Bandits with Known Behaviours/Scheduling Projects

- Consider now $L$ different DP problems $X_{0}^{\ell}, X_{1}^{\ell}, X_{2}^{\ell}, \ldots$ with different state evolution and cost functions $f^{\ell}(a, u, w)$ and $g^{\ell}(a, u, w)$, for $\ell=1, \ldots, L$
- At each stage $k$ one can retire at cost $\gamma^{k} \cdot M_{\infty}$
- Initial state $\mathbf{x}_{0}=\left(x_{0}^{1}, x_{0}^{2}, \ldots, x_{0}^{L}\right)$
- At each stage $k$, retire or choose a project $\ell_{k}^{*} \in\{1, \ldots, L\}$ and an action $u$. If you don't retire:

$$
X_{k+1}^{\ell_{k}^{*}}=f^{\ell_{k}^{*}}\left(X_{k}^{\ell^{*}}, u, W\right) \quad \text { and } \quad X_{k+1}^{\ell}=X_{k}^{\ell}, \forall \ell \in\{1, \ldots, L\} \backslash\left\{\ell_{k}^{*}\right\}
$$

and the stage- $k$ cost is given by

$$
g\left(x_{1}, \ldots, x_{L},\left(u, \ell_{k}^{*}\right), W\right)=g^{\ell_{k}^{*}}\left(x_{\ell^{*}}, u, W\right)
$$

- Wish to maximize the infinite-horizon discounted cost until retirement (if the player retires at all)


## Optimal Scheduling Policy for Multi-Armed Bandit Problems

- Calculate the retirement threshold $m^{\ell}($ a) for each project $\ell=1, \ldots, L$ and state $a \in \mathcal{X}$ as explained before


## Optimal Policy

Assume that at time $k$ the states of the $L$ projects are $x_{1}, \ldots, x_{L}$.

- Retire if

$$
m^{\ell}\left(x_{\ell}\right) \geq M_{\infty}, \quad \forall \ell \in\{1, \ldots, L\} .
$$

- Otherwise choose (ties can be split arbitrary)

$$
\ell_{k}^{*}=\operatorname{argmin}_{\ell} m^{\ell}\left(x_{\ell}\right)
$$

and play the optimal policy for this project $\ell_{k}^{*}$ according to Bellman's equation.

## Unbounded but Positive Costs

- Positive (possibly unbounded) costs $g(x, u, w) \in[0, \infty)$
- Discount factor $\gamma<1$
- Bellman's equation remains valid:

$$
J^{*}=T J^{*} .
$$

But the solution might not be unique.
The optimal cost function is given by the smallest fix-point!

- Value-iteration algorithm still works and provides optimal cost and optimal stationary policy!
$\rightarrow$ finite-horizon solutions converge to the infinite-horizon solutions
- Policy iteration algorithm does not necessarily converge to optimal solution


## The Quadratic Gaussian Case

- Vector states $\mathbf{x}_{0}, \mathbf{x}_{1}, \mathbf{x}_{2}, \ldots \in \mathbb{R}^{n}$ and actions $\mathbf{u}_{0}, \mathbf{u}_{1}, \mathbf{u}_{2}, \ldots \in \mathbb{R}^{m}$
- i.i.d. Gaussian noise vectors $\mathbf{W}_{k}$ of covariance matrix $\mathrm{K}_{w}$
- State evolution when noise $\mathbf{W}_{k}=\mathbf{w}_{k}$ and controls $\mathbf{u}_{0}, \mathbf{u}_{1}, \mathbf{u}_{2}, \ldots$,

$$
\mathbf{x}_{k+1}=f\left(\mathbf{x}_{k}, \mathbf{u}_{k}, \mathbf{w}_{k}\right)=\mathrm{A} \mathbf{x}_{k}+\mathrm{B} \mathbf{u}_{k}+\mathbf{w}_{k}, \quad k=0,1,2, \ldots
$$

for given matrices $A$ and $B$.

- Deterministic cost function

$$
\sum_{k=0}^{\infty} \gamma^{k} g\left(\mathbf{x}_{k}, \mathbf{u}_{k}, \mathbf{w}_{k}\right)=\sum_{k=0}^{\infty} \gamma^{k}\left(\mathbf{x}_{k}^{\top} \mathrm{Q} \mathbf{x}_{k}+\mathbf{u}_{k}^{\top} \mathrm{R} \mathbf{u}_{k}\right) .
$$

- Let R and Q be positive semi-definite.


## Value-Iteration Algorithm on the Quadratic Gaussian Case

- Value-Iteration update rule for $k=1,2, \ldots$

$$
\begin{aligned}
V_{k}(\mathbf{x}) & =\min _{\mu} \mathbb{E}_{\mathbf{W}}\left[g(\mathbf{x}, \mu(\mathbf{x}), \mathbf{W})+\gamma \boldsymbol{V}_{k-1}(f(\mathbf{x}, \mu(\mathbf{x}), \mathbf{W}))\right] \\
& =\min _{\mathbf{u}}\left[\mathbf{x}^{\top} \mathbf{Q} \mathbf{x}+\mathbf{u}^{\top} \mathbf{R} \mathbf{u}+\gamma \mathbb{E}\left[V_{k-1}(\mathrm{~A} \mathbf{x}+\mathrm{Bu}+\mathbf{W})\right]\right.
\end{aligned}
$$

- Start with $\mathbf{V}_{0}(\mathbf{x})=0$, for all vectors $\mathbf{x}$
- Notice that because $R$ is positive semi-definite, $\mathbf{u}^{\top} R \mathbf{u} \geq 0$ with equality for $\mathbf{u}=\mathbf{0}$. Thus:

$$
\mathbf{V}_{1}(\mathbf{x})=\min _{\mathbf{u}} \mathbf{x}^{\top} \mathbf{Q} \mathbf{x}+\mathbf{u}^{\top} R \mathbf{u}=\mathbf{x}^{\top} \mathbf{Q} \mathbf{x}
$$

- For $k=2$ :

$$
\begin{aligned}
& \mathbf{V}_{2}(\mathbf{x})=\min _{\mathbf{u}}\left[\mathbf{x}^{\top} \mathbf{Q} \mathbf{x}+\mathbf{u}^{\top} \mathbf{R} \mathbf{u}+\gamma \mathbb{E}_{W}\left[\left(\mathbf{x}^{\top} \mathrm{A}^{\top}+\mathbf{u} \mathbf{B}^{\top}+\mathbf{W}^{\top}\right) \mathrm{Q}(\mathbf{W}+\mathrm{B} \mathbf{u}+\mathrm{A} \mathbf{x})\right]\right] \\
&\left.=\mathbf{x}^{\top} \mathbf{Q} \mathbf{x}+\gamma \mathbb{E}\left[\mathbf{W}^{\top} \mathbf{Q} \mathbf{W}\right]+\min _{\mathbf{u}}\left[\mathbf{u}^{\top} \mathbf{R} \mathbf{u}+\gamma\left(\mathbf{x}^{\top} \mathrm{A}^{\top}+\mathbf{u} \mathrm{B}^{\top}\right) \mathbf{Q}(\mathrm{B} \mathbf{u}+\mathrm{A} \mathbf{x})\right]\right] \\
&=\mathbf{x}^{\top} \underbrace{\left(\mathbf{Q}+\mathrm{A}^{\top} \mathbf{Q A}\right)}_{\text {positive semidefinite }} \mathbf{x}+\gamma \mathbb{E}\left[\mathbf{W}^{\top} \mathbf{Q} \mathbf{W}\right] \\
&+\min _{\mathbf{u}}[\mathbf{u}^{\top} \underbrace{\left(\mathrm{R}+\gamma \mathbf{B}^{\top} \mathbf{Q B}\right)}_{\text {positive semidefinite }} \mathbf{u}+2 \gamma \mathbf{x}^{\top} \mathbf{A}^{\top} \mathbf{Q B} \mathbf{u}]
\end{aligned}
$$

## Minimizing Quadratic Forms

- Consider the quadratic form in $\mathbf{u}$ :

$$
f(\mathbf{u})=\frac{1}{2} \mathbf{u}^{\top} \mathbf{M} \mathbf{u}+\mathbf{c}^{\top} \mathbf{u}
$$

where $\mathbf{c}$ is an arbitrary vector and $\mathbf{M}$ is a positive semidefinite matrix. (This latter assumption is need to ensure convexity of the function $f$.)

- The gradient of $f$ with respect to $\mathbf{u}$ is:

$$
\nabla f(\mathbf{u})=\mathrm{M} \mathbf{x}+\mathbf{c}
$$

- The function $f$ is minimized for

$$
\mathbf{u}^{*}=-\mathrm{M}^{-1} \mathbf{c}
$$

and the minimum value of $f$ is

$$
f_{\min }:=\min _{\mathbf{u}} f(\mathbf{u})=-\frac{1}{2} \mathbf{c}^{\top} \mathrm{M}^{-1} \mathbf{c}
$$

## Quadratic Gaussian Example continued

- We obtain for $k=2$ :

$$
\begin{aligned}
\mathbf{V}_{2}(\mathbf{x}) & =\mathbf{x}^{\top}\left(\mathrm{Q}+\gamma \mathrm{A}^{\top} \mathrm{QA}\right) \mathbf{x}+\gamma \mathbb{E}\left[\mathbf{W}^{\top} \mathrm{Q} \mathbf{W}\right]-\gamma^{2} \mathbf{x}^{\top} \mathrm{A}^{\top} \mathrm{QB}\left(\mathrm{R}+\gamma \mathrm{B}^{\top} \mathrm{QB}\right)^{-1} \mathrm{~B}^{\top} \mathrm{QA} \mathbf{x} \\
& =\gamma \mathbb{E}\left[\mathbf{W}^{\top} \mathrm{QW}\right]+\mathbf{x}^{\top} \underbrace{\left(\mathrm{Q}+\gamma \mathrm{A}^{\top} \mathrm{QA}-\gamma^{2} \mathrm{~A}^{\top} \mathrm{QB}\left(\mathrm{R}+\gamma \mathrm{B}^{\top} \mathrm{QB}\right)^{-1} \mathrm{~B}^{\top} \mathrm{QA}\right)}_{=: M_{2}} \mathbf{x}
\end{aligned}
$$

- The optimal control is linear:

$$
\mathbf{u}^{*}=-\gamma\left(\mathrm{R}+\gamma \mathbf{B}^{\top} \mathbf{Q B}\right)^{-1} \mathbf{B}^{\top} \mathbf{Q A} \mathbf{x}
$$

- $\mathbf{V}_{2}$ has a similar form to $\mathbf{V}_{1}$ but with $M_{2}$ (which is positive semi-definite, see slide 12) instead of $\mathbf{Q}$, and there is an additional summand $\gamma \operatorname{tr}\left(\mathrm{K}_{W} \mathrm{Q}\right)$
- Can obtain $\mathbf{V}_{3}$ following the same reasoning, but exchanging Q with $\mathrm{M}_{2}$ and adding $\gamma \cdot \gamma \mathbb{E}\left[\mathbf{W}^{T} \mathbf{Q} \mathbf{W}\right]$ to the cost


## Semi-positivity of matrix $\mathrm{M}_{2}$

- By standard manipulations on matrices:

$$
\begin{aligned}
\Gamma: & =\gamma \mathrm{A}^{\top} \mathrm{QA}-\gamma^{2} \mathrm{~A}^{\top} \mathrm{QB}\left(\mathrm{R}+\gamma \mathrm{B}^{\top} \mathrm{QB}\right)^{-1} \mathrm{~B}^{\top} \mathrm{QA} \\
= & \gamma \mathrm{A}^{\top}\left(\mathrm{Q}-\gamma \mathrm{QB}\left(\mathrm{R}+\gamma \mathrm{B}^{\top} \mathrm{QB}\right)^{-1} \mathrm{~B}^{\top} \mathrm{Q}\right) \mathrm{A} \\
= & \gamma \mathrm{A}^{\top}\left(\mathrm{QB}\left(\mathrm{~B}^{\top} \mathrm{QB}\right)^{-1} \mathrm{~B}^{\top} \mathrm{Q}-\gamma \mathrm{QB}\left(\mathrm{R}+\gamma \mathrm{B}^{\top} \mathrm{QB}\right)^{-1} \mathrm{~B}^{\top} \mathrm{Q}\right) \mathrm{A} \\
= & \gamma \mathrm{A}^{\top} \mathrm{QB}\left(\left(\mathrm{~B}^{\top} \mathrm{QB}\right)^{-1}-\gamma\left(\mathrm{R}+\gamma \mathrm{B}^{\top} \mathrm{QB}\right)^{-1}\right) \mathrm{B}^{\top} \mathrm{QA} \\
= & \gamma \mathrm{A}^{\top} \mathrm{QB}\left(\left(\mathrm{~B}^{\top} \mathrm{QB}\right)^{-1}\left(\mathrm{R}+\gamma \mathrm{B}^{\top} \mathrm{QB}\right)\left(\mathrm{R}+\gamma \mathrm{B}^{\top} \mathrm{QB}\right)^{-1}\right. \\
& \left.\quad \quad\left(\mathrm{B}^{\top} \mathrm{QB}\right)^{-1}\left(\mathrm{~B}^{\top} \mathrm{QB}\right) \gamma\left(\mathrm{R}+\gamma \mathrm{B}^{\top} \mathrm{QB}\right)^{-1}\right) \mathrm{B}^{\top} \mathrm{QA} \\
= & \gamma \mathrm{A}^{\top} \mathrm{QB}\left(\mathrm{~B}^{\top} \mathrm{QB}\right)^{-1}\left(\left(\mathrm{R}+\gamma \mathrm{B}^{\top} \mathrm{QB}\right)-\gamma\left(\mathrm{B}^{\top} \mathrm{QB}\right)\right)\left(\mathrm{R}+\gamma \mathrm{B}^{\top} \mathrm{QB}\right)^{-1} \mathrm{~B}^{\top} \mathrm{QA} \\
= & \gamma \mathrm{A}^{\top} \mathrm{QB}\left(\mathrm{~B}^{\top} \mathrm{QB}\right)^{-1} \mathrm{R}\left(\mathrm{R}+\gamma \mathrm{B}^{\top} \mathrm{QB}\right)^{-1} \mathrm{~B}^{\top} \mathrm{QA}
\end{aligned}
$$

- $\Gamma \succeq 0$ is positive semidefinite because: - $\mathrm{Q}, \mathrm{R}$ are positive semidefinite and for any positive semidefinite matrices $M, N$ and arbitrary matrix $S$ :
$M+N \succeq 0, M \cdot N \succeq 0, M^{-1} \succeq 0, S^{\top} M S \succeq 0$ are also positive semidefinite.
- By the same reasons, also $M_{2}=\Gamma+Q$ is positive semidefinite


## Quadratic Gaussian Example continued II

- We obtain for $k=3$ :

$$
\begin{aligned}
\mathbf{V}_{3}(\mathbf{x})= & \min _{\mathbf{u}}\left[\mathbf{x}^{\top} \mathbf{Q} \mathbf{x}+\mathbf{u}^{\top} \mathrm{R} \mathbf{u}+\gamma \mathbb{E} w\left[\left(\mathbf{x}^{\top} \mathrm{A}^{\top}+\mathbf{u} \mathrm{B}^{\top}+\mathbf{W}^{\top}\right) \mathrm{M}_{2}(\mathbf{W}+\mathrm{B} \mathbf{u}+\mathrm{A} \mathbf{x})\right]\right] \\
& +\gamma^{2} \mathbb{E}\left[\mathbf{W}^{\top} \mathbf{Q} \mathbf{W}\right] \\
= & \mathbf{x}^{\top}\left(\mathrm{Q}+\gamma \mathrm{A}^{\top} \mathrm{M}_{2} \mathrm{~A}\right) \mathbf{x}+\gamma^{2} \mathbb{E}\left[\mathbf{W}^{\top} \mathrm{Q} \mathbf{W}\right]+\gamma \mathbb{E}\left[\mathbf{W}^{\top} \mathrm{M}_{2} \mathbf{W}\right] \\
& +\min _{\mathbf{u}}\left[\mathbf{u}^{\top}\left(\mathrm{R}+\gamma \mathrm{B}^{\top} \mathrm{M}_{2} \mathrm{~B}\right) \mathbf{u}+2 \gamma \mathbf{x}^{\top} \mathrm{A}^{\top} \mathrm{M}_{2} \mathrm{~B} \mathbf{u}\right] \\
= & \gamma^{2} \mathbb{E}\left[\mathbf{W}^{\top} \mathbf{Q} \mathbf{W}\right]+\gamma \mathbb{E}\left[\mathbf{W}^{\top} \mathrm{M}_{2} \mathbf{W}\right] \\
& +\mathbf{x}^{\top} \underbrace{\left(\mathbf{Q}+\gamma \mathrm{A}^{\top} \mathrm{M}_{2} \mathrm{~A}-\gamma^{2} \mathrm{~A}^{\top} \mathrm{M}_{2} \mathrm{~B}\left(\mathrm{R}+\gamma \mathrm{B}^{\top} \mathrm{M}_{2} \mathrm{~B}\right)^{-1} \mathrm{~B}^{\top} \mathrm{M}_{2} \mathrm{~A}\right)}_{=: \mathrm{M}_{3}} \mathbf{x}
\end{aligned}
$$

- The optimal control is linear:

$$
\mathbf{u}^{*}=-\gamma\left(\mathrm{R}+\gamma \mathrm{B}^{\top} \mathrm{M}_{2} \mathrm{~B}\right)^{-1} \mathrm{~B}^{\top} \mathrm{M}_{2} \mathrm{~A} \mathbf{x}
$$

- Can obtain $\mathbf{V}_{4}$ following the same reasoning, but exchanging $\mathrm{M}_{2}$ with $\mathrm{M}_{3}$ and adding $\gamma \cdot\left(\gamma^{2} \mathbb{E}\left[\mathbf{W}^{T} \mathbf{Q} \mathbf{W}\right]+\gamma \mathbb{E}\left[\mathbf{W}^{T} \mathrm{M}_{2} \mathbf{W}\right]\right)$ to the cost. $\mathbb{E T C}$.


## Quadratic Gaussian Example continued III

- Continuing along the same lines, we observe:

$$
\mathbf{V}_{k}(\mathbf{x})=\sum_{\ell=1}^{k-1} \gamma^{k-\ell} \mathbb{E}\left[\mathbf{W}^{T} \mathrm{M}_{\ell} \mathbf{W}\right]+\mathbf{x}^{\top} \mathrm{M}_{k} \mathbf{x}
$$

where $M_{1}=Q$ and for $k=2,3, \ldots$ :

$$
\begin{aligned}
M_{k} & =Q+\gamma A^{\top} M_{k-1} A-\gamma^{2} A^{\top} M_{k-1} B\left(R+\gamma B^{\top} M_{k-1} B\right)^{-1} B^{\top} M_{k-1} A \\
& =Q+\tilde{A}^{\top} M_{k-1} \tilde{A}-\tilde{A}^{\top} M_{k-1} \tilde{B}\left(R+\tilde{B}^{\top} M_{k-1} \tilde{B}\right)^{-1} \tilde{B}^{\top} M_{k-1} \tilde{A},
\end{aligned}
$$

where $\tilde{A}:=\sqrt{\gamma} \mathrm{A}$ and $\mathrm{B}:=\sqrt{\gamma} \mathrm{B}$

- It can again be shown that $M_{k} \succeq 0$ is positive semidefinite.
- The sequence $M_{k}$ is known to converge to $M^{*}$ the solution of the Algebraic Riccatti Equation (important in control theory)

$$
M=Q+\tilde{A}^{\top} M \tilde{A}-\tilde{A}^{\top} M \tilde{B}\left(R+\tilde{B}^{\top} M \tilde{B}\right)^{-1} \tilde{B}^{\top} M \tilde{A}
$$

whenever the pair $(\tilde{A}, \tilde{B})$ is controllable and $(\tilde{A}, \tilde{C})$ is observable, where $Q=C^{\top} C$.

## Controllability and Observability

## Definition (Controllability)

A pair $(A, B)$, where $A$ is an $n \times n$ matrix and $B$ a $n \times m$ matrix, is said controllable if the $n \times n m$ matrix

$$
\left[B, A B, A^{2} B, \ldots A^{n-1} B\right]
$$

has full rank

## Definition (Observability)

A pair $(A, C)$ is said observable if the pair $\left(A^{\top}, C^{\top}\right)$ is controllable.

## The Solution of the Quadratic Gaussian Example

- Since $M_{\ell}$ converges, also the weighted sum of the noise-terms converges. Using the geometric sum formula:

$$
\lim _{k \rightarrow \infty} \sum_{\ell=1}^{k-1} \gamma^{k-\ell} \mathbb{E}\left[\mathbf{W}^{T} \mathbf{M}_{\ell} \mathbf{W}\right]=\frac{1}{1-\gamma} \mathbb{E}\left[\mathbf{W}^{T} \mathbf{M}^{*} \mathbf{W}\right]
$$

where $\mathrm{M}^{*}$ is the solution to the Algebraic Riccatti equation

$$
\begin{equation*}
M=Q+\tilde{A}^{\top} M \tilde{A}-\tilde{A}^{\top} M \tilde{B}\left(R+\tilde{B}^{\top} M \tilde{B}\right)^{-1} \tilde{B}^{\top} M \tilde{A} \tag{1}
\end{equation*}
$$

## Optimal Infinite cost $J^{*}(x)$

For any state vector $\mathbf{x}$ :

$$
J^{*}(\mathbf{x})=\frac{1}{1-\gamma} \mathbb{E}\left[\mathbf{W}^{\top} \mathbf{M}^{*} \mathbf{W}\right]+\mathbf{x}^{\top} \mathbf{M}^{*} \mathbf{x}
$$

where $\mathrm{M}^{*}$ is the solution to (1)

## english

# Sequential Decision Processes, Master MICAS, Part I 

Michèle Wigger

Telecom Paris, 8 Jan 2021


# Lecture 6- Constrained Discounted Problems and Average-Cost Problems 

## Outlook Today

- Time-invariant discrete-time dynamic system:

$$
X_{k+1}=f\left(X_{k}, U_{k}, W_{k}\right), \quad k=0,1,2, \ldots
$$

disturbance $\left\{W_{k}\right\}$ i.i.d.

- Bounded time-invariant cost function $g(x, u, w) \in[-M, M]$
- Optimal discounted infinite-horizon cost:

$$
\bar{J}^{*}\left(p_{0}\right):=\min _{\pi} \lim _{N \rightarrow \infty} \mathbb{E}_{X_{0},\left\{W_{k}\right\}}\left[\sum_{k=0}^{N-1} \gamma^{k} g\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right)\right]
$$

- Today we add cost constraints: A policy $\pi$ is admissible only if

$$
\mathbb{E}_{X_{0},\left\{W_{k}\right\}}^{\pi}\left[\sum_{k=0}^{\infty} \gamma^{k} d_{\ell}\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right)\right] \leq D_{\ell}, \quad \ell=1, \ldots, L
$$

## Outlook Today

- Time-invariant discrete-time dynamic system:

$$
X_{k+1}=f\left(X_{k}, U_{k}, W_{k}\right), \quad k=0,1,2, \ldots,
$$

disturbance $\left\{W_{k}\right\}$ i.i.d.

- Bounded time-invariant cost function $g(x, u, w) \in[-M, M]$
- Optimal average infinite-horizon cost:

$$
\bar{J}^{*}\left(p_{0}\right):=\min _{\pi} \lim _{N \rightarrow \infty} \mathbb{E}_{X_{0},\left\{W_{k}\right\}}\left[\sum_{k=0}^{N-1} \frac{1}{N} g\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right)\right]
$$

## Review of Lecture 4: LP Programming Approach

## Primal Problem

$$
\max _{J_{1}, \ldots, J_{m}}(1-\gamma) \sum_{i=1}^{m} p_{0}(i) J_{i}
$$

subject to:

$$
J_{i} \leq \mathbb{E}_{W}[g(i, u, W)]+\gamma \cdot \sum_{j=1}^{m} P_{u, i j} J_{j}, \quad \forall i, u
$$

where $P_{u, i j}:=\operatorname{Pr}[f(i, u, W)=j]$

## Review of Lecture 4: LP Programming Approach

## Dual Problem

$$
\min _{\{\rho(i, u)\}} \sum_{i=1}^{m} \sum_{u} \mathbb{E}_{W}[g(i, u, W)] \cdot \rho(i, u)
$$

subject to:

$$
\sum_{u} \rho(i, u)-\sum_{j=1}^{m} \sum_{u} \gamma P_{u, i j} \cdot \rho(j, u)=(1-\gamma) p_{0}(i), \quad i=1, \ldots, m
$$

where $P_{u, i j}:=\operatorname{Pr}[f(i, u, W)=j]$ and $\rho(i, u) \geq 0$ for all $i, u$.

- State-action frequencies/occupation measures $\rho(i, u)$ form a pmf and determine a randomized stationary policy $\mu(u \mid i)=\frac{\rho(i, u)}{\sum_{u} \rho(i, u)}$
- $\exists$ an optimal $\rho^{*}(i, u)>0$ with only $m$ components, one for each state $i$ $\Longrightarrow$ Deterministic stationary policies are optimal!


## Constrained Discounted Infinite-Horizon Problems

- Time-invariant discrete-time dynamic system:

$$
X_{k+1}=f\left(X_{k}, U_{k}, W_{k}\right), \quad k=0,1,2, \ldots,
$$

- Bounded time-invariant cost function $g(x, u, w) \in[-M, M]$ and constraint-cost functions $d_{\ell}(x, u, w)$, for $\ell=1, \ldots, L$, as well as maximum constraints $D_{1}, \ldots, D_{L}$
- Optimal discounted infinite-horizon cost:

$$
J^{*}(a):=\min _{\pi} \lim _{N \rightarrow \infty} \mathbb{E}_{X_{0},\left\{W_{k}\right\}}\left[\sum_{k=0}^{N} \gamma^{k} g\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right)\right]
$$

where minimum is over all policies $\pi=\left(\mu_{1}, \mu_{2}, \ldots\right)$ satisfying

$$
\lim _{N \rightarrow \infty} \mathbb{E}_{X_{0},\left\{W_{k}\right\}}\left[\sum_{k=0}^{N} \gamma^{k} d_{\ell}\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right)\right] \leq D_{\ell}, \quad \ell=1, \ldots, L .
$$

## Can express constraints using State-Action Frequencies

For all $\ell=1, \ldots, L$ :

$$
\begin{aligned}
& (1-\gamma) \mathbb{E}_{X_{0},\left\{W_{k}\right\}}\left[\sum_{k=0}^{\infty} \gamma^{k} d_{\ell}\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right)\right] \\
& =(1-\gamma) \sum_{k=0}^{\infty} \gamma^{k} \sum_{i, u} \mathbb{E}\left[d_{\ell}\left(i, u, W_{k}\right)\right] \operatorname{Pr}\left[X_{k}=i, \mu_{k}(i)=u\right] \\
& =\sum_{i, u} \mathbb{E}\left[d_{\ell}(i, u, W)\right](1-\gamma) \sum_{k=0}^{\infty} \gamma^{k} \operatorname{Pr}\left[X_{k}=i, \mu_{k}(i)=u\right] \\
& =\sum_{i, u} \mathbb{E}\left[d_{\ell}(i, u, W)\right] \rho(i, u) \\
& \leq(1-\gamma) D_{\ell}
\end{aligned}
$$

## Dual Linear Programming Problem with Constraints

Dual Linear Programming Problem

$$
J^{*}\left(p_{0}\right)=\min _{\rho(i, u) \geq 0} \sum_{i=1}^{m} \sum_{u} \mathbb{E}_{W}[g(i, u, W)] \cdot \rho(i, u)
$$

subject to:

$$
\sum_{u} \rho(i, u)-\sum_{j=1}^{m} \sum_{u} \gamma P_{u, i j} \cdot \rho(j, u)=(1-\gamma) p_{0}(i), \quad i=1, \ldots, m
$$

and

$$
\sum_{i, u} \mathbb{E}\left[d_{\ell}(i, u, W)\right] \rho(i, u) \leq(1-\gamma) D_{\ell}, \quad \ell=1, \ldots, L
$$

- Optimal policy is generally stationary with $\leq L$ randomized actions


## Dual Problem with Constraints $\rightarrow$ Lagrange Multipliers

## Dual Problem

$$
J^{*}\left(p_{0}\right)=\min _{\rho(i, u) \geq 0} \sum_{i=1}^{m} \sum_{u} \mathbb{E} w[\underbrace{g(i, u, W)}] \cdot \rho(i, u)
$$

subject to:

$$
\sum_{u} \rho(i, u)-\sum_{j=1}^{m} \sum_{u} \gamma P_{u, i j} \cdot \rho(j, u)=(1-\gamma) p_{0}(i), \quad i=1, \ldots, m
$$

and

$$
\sum_{i, u} \mathbb{E}\left[d_{\ell}(i, u, W)\right] \rho(i, u) \leq(1-\gamma) D_{\ell}, \quad \ell=1, \ldots, L
$$

- Add additional constraints using Lagrange Multipliers $\lambda_{1}, \ldots, \lambda_{L}$ !


## Dual Problem with Constraints $\rightarrow$ Lagrange Multipliers

## Dual Problem

$$
\begin{aligned}
J^{*}\left(p_{0}\right)=\min _{\rho(i, u) \geq 0} & \sup _{\lambda_{1}, \ldots, \lambda_{L} \geq 0} \sum_{i=1}^{m} \sum_{u} \mathbb{E}_{W}[\underbrace{g(i, u, W)+\sum_{\ell=1}^{L} \lambda_{\ell} d_{\ell}(i, u, W)}] \cdot \rho(i, u) \\
& -\sum_{\ell=1}^{L} \lambda_{\ell} D_{\ell}
\end{aligned}
$$

subject to:

$$
\sum_{u} \rho(i, u)-\sum_{j=1}^{m} \sum_{u} \gamma P_{u, i j} \cdot \rho(j, u)=(1-\gamma) p_{0}(i), \quad i=1, \ldots, m
$$

- Add additional constraints using Lagrange Multipliers $\lambda_{1}, \ldots, \lambda_{L}$ !


## Dual Problem with Constraints $\rightarrow$ Lagrange Multipliers

## Dual Problem

$$
\begin{aligned}
J^{*}\left(p_{0}\right)=\sup _{\lambda_{1}, \ldots, \lambda_{L} \geq 0} & \min _{\rho(i, u) \geq 0} \sum_{i=1}^{m} \sum_{u} \mathbb{E}_{W}[\underbrace{g(i, u, W)+\sum_{\ell=1}^{L} \lambda_{\ell} d_{\ell}(i, u, W)}] \cdot \rho(i, u) \\
& -\sum_{\ell=1}^{L} \lambda_{\ell} D_{\ell}
\end{aligned}
$$

subject to:

$$
\sum_{u} \rho(i, u)-\sum_{j=1}^{m} \sum_{u} \gamma P_{u, i j} \cdot \rho(j, u)=(1-\gamma) p_{0}(i), \quad i=1, \ldots, m
$$

- Add additional constraints using Lagrange Multipliers $\lambda_{1}, \ldots, \lambda_{L}$ !
- Strong duality holds by standard arguments


## Dual Problem with Constraints $\rightarrow$ Lagrange Multipliers

## Dual Problem

$$
\begin{aligned}
J^{*}\left(p_{0}\right)=\sup _{\lambda_{1}, \ldots, \lambda_{L} \geq 0} & \min _{\rho(i, u) \geq 0} \sum_{i=1}^{m} \sum_{u} \mathbb{E}_{W}[\underbrace{g(i, u, W)+\sum_{\ell=1}^{L} \lambda_{\ell} d_{\ell}(i, u, W)}_{\text {new cost function } \tilde{g}(i, u, W)}] \cdot \rho(i, u) \\
& -\sum_{\ell=1}^{L} \lambda_{\ell} D_{\ell}
\end{aligned}
$$

subject to:

$$
\sum_{u} \rho(i, u)-\sum_{j=1}^{m} \sum_{u} \gamma P_{u, i j} \cdot \rho(j, u)=(1-\gamma) p_{0}(i), \quad i=1, \ldots, m
$$

- Add additional constraints using Lagrange Multipliers $\lambda_{1}, \ldots, \lambda_{L}$ !
- Strong duality holds by standard arguments
- For each $\lambda_{1}, \ldots, \lambda_{L}$ : solve for the new cost function $\tilde{g}$
$\rightarrow$ minimum achieved by a deterministic stationary policy (proof as before)


## Optimal Average Cost Problems

- Optimal average infinite horizon cost:

$$
\bar{J}^{*}\left(p_{0}\right):=\min _{\pi} \bar{J}^{\pi}\left(p_{0}\right)
$$

where for a given policy $\pi$ :

$$
\bar{J}^{\pi}\left(p_{0}\right):=\varlimsup_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}_{X_{0},\left\{W_{k}\right\}}\left[\sum_{k=0}^{N-1} g\left(X_{k}, U_{k}, W_{k}\right)\right]
$$

- We can again restrict to Markov policies because objective function only depends on $\left\{P_{x_{k}}, v_{k}\right\}_{k \geq 0}$ as in the discounted case


## Unichain Assumption

- For a stationary policy $\mu$, the induced Markov chain has transition matrix

$$
\mathrm{P}_{\mu}(i, j):=\operatorname{Pr}\left[X_{k+1}=j \mid X_{k}=i\right]=\sum_{u} \mu(u \mid i) \operatorname{Pr}[f(i, u, W)=j]
$$

- Recall: If a Markov chain is irreducible (i.e., $\mathcal{X}$ is a recurrent class) and aperiodic, its state-distribution tends to the unique stationary distribution, irrespective of the $X_{0}$-distribution.
- If the Markov chain is periodic, the distribution can "toggle" between different distributions
- The same holds also when there is an additional set of transient states. (At some point the Markov chain will end in the recurrent class and converge (or toggle).)


## Definition (Unichain)

A Dynamic Programming Problem is called Unichain if the state space can be decomposed into $\mathcal{S} \cup \mathcal{T}=\mathcal{X}$, with $\mathcal{S} \cap \mathcal{T}=\emptyset$, so that for all stationary policies $\mu$, the set $\mathcal{S}$ forms a recurrent class and $\mathcal{T}$ is a set of transient states.

## Expressing the Cost-Function in State-Action Frequencies

- For a given policy $\pi$ :

$$
\begin{aligned}
\bar{J}^{\pi}\left(p_{0}\right) & :=\varlimsup_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}_{X_{0},\left\{W_{k}\right\}}\left[\sum_{k=0}^{N-1} g\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right)\right] \\
& =\varlimsup_{N \rightarrow \infty} \sum_{i, u} \mathbb{E}[g(i, u, W)] \cdot \frac{1}{N} \sum_{k=0}^{N-1} \operatorname{Pr}\left[X_{k}=i, \mu_{k}(i)=u\right] \\
& =\varlimsup_{N \rightarrow \infty} \sum_{i, u} \mathbb{E}[g(i, u, W)] \cdot \nu_{N}^{\pi}(i, u)
\end{aligned}
$$

- N-horizon state-action frequency

$$
\nu_{N}^{\pi}(i, u):=\frac{1}{N} \sum_{k=0}^{N-1} \operatorname{Pr}\left[X_{k}=i, \mu_{k}(i)=u\right]
$$

- $N$-horizon state-action frequency (occupation measure) $\nu_{N}^{\pi}(i, u)$ describes the probability of observing the state-action pair $(i, u)$ at a random time $T$ which is uniform over $\{0,1, \ldots, N-1\}$


## Convergence of $\nu_{N}^{\pi}(i, u)$

- Depending on the policy $\pi$, the sequences $\left\{\nu_{N}^{\pi}(i, u)\right\}_{N \geq 1}$ might diverge to various accumulation points! $\rightarrow$ therefore use limsup!
- Let $\nu^{\pi}$ be an accumulation point of $\left\{\nu_{N}^{\pi}(i, u)\right\}_{N \geq 1}$. Then (see next slide):

$$
\sum_{u} \nu^{\pi}(i, u)=\sum_{j, u} \nu^{\pi}(j, u) P_{u, j i}
$$

- Under the unichain assumption and stationary policy $\mu$, the sequences $\left\{\nu_{N}^{\mu}(i, u)\right\}_{N \geq 1}$ converge to the (infinite-horizon) state-action frequencies

$$
\nu_{\infty}^{\mu}(i, u):=\lim _{N \rightarrow \infty} \nu_{N}^{\mu}(i, u)=\xi^{\mu}(i) \cdot \mu(u \mid i),
$$

irrespective of $p_{0}$, and where $\xi^{\mu}=\left(\xi^{\mu}(1), \ldots, \xi^{\mu}(m)\right)$ is the stationary distribution of the Markov chain $\mathrm{P}_{\mu}$.
Proof: Apply Césaro's mean theorem and the limit $\operatorname{Pr}\left[X_{k}=i\right] \rightarrow \xi^{\mu}(i)$

## Proof that $\sum_{u} v^{\pi}(i, u)=\sum_{j, u} v^{\pi}(j, u) P_{u, j i}$

Consider any initial distribution $p(0)$ and increasing sequence $\left\{N_{/}\right\}_{\mid \geq 0}$ such that $\nu_{N_{l}}^{\pi}(i, u)$ converges to $v^{\pi}(i, u)$ as $I \rightarrow \infty$ for all $u, i$.
For any $\mathrm{l}>0$ :

$$
\begin{aligned}
& \sum_{v} \nu_{N_{l}}^{\pi}(i, v)-\frac{1}{N_{l}} p(0) \\
& =\sum_{v} \frac{1}{N_{l}} \sum_{k=1}^{N_{l}-1} \operatorname{Pr}\left[X_{k}=i, \mu_{k}(i)=v\right]=\frac{1}{N_{l}} \sum_{k=1}^{N_{l}-1} \operatorname{Pr}\left[X_{k}=i\right] \\
& =\frac{1}{N_{l}} \sum_{k=1}^{N_{l}-1} \sum_{j, u} \operatorname{Pr}\left[X_{k-1}=j, U_{k-1}=u\right] P_{u, j i} \\
& =\frac{1}{N_{l}} \sum_{k^{\prime}=0}^{N_{l}-2} \sum_{j, u} \operatorname{Pr}\left[X_{k^{\prime}}=j, U_{k^{\prime}}=u\right] P_{u, j i} \\
& =\frac{1}{N_{l}} \sum_{k^{\prime}=0}^{N_{l}-1} \sum_{j, u} \operatorname{Pr}\left[X_{k^{\prime}}=j, U_{k^{\prime}}=u\right] P_{u, j i}-\frac{1}{N_{l}} \operatorname{Pr}\left[X_{N_{l}-1}=j, U_{N_{l}-1}=u\right] P_{u, j i}
\end{aligned}
$$

Taking limits $I \rightarrow \infty$ and thus $N_{I} \rightarrow \infty$ on both sides, yields the desired expressions because the sums and the limit can be exchanged

## Can restrict to Stationary Policies

- Given any policy $\pi$ and accumulation point $\nu^{\pi}(i, u)$.
- Choose a stationary policy $\mu$ with

$$
\mu(u \mid i)=\frac{\nu^{\mu}(i, u)}{\sum_{v} \nu^{\mu}(i, v)}
$$

- $\pi$ and $\mu$ have same state-action frequencies:

$$
\nu^{\pi}(i, u)=\mu(u \mid i) \cdot\left(\sum_{v} \nu^{\mu}(i, v)\right)=\underbrace{\mu(u \mid i) \xi^{\mu}(i)}_{=\nu_{\infty}^{\mu}(i, u)} \cdot \underbrace{\frac{\sum_{v} \nu^{\mu}(i, v)}{\xi^{\mu}(i)}}_{=1, \text { see next slide }}=\nu_{\infty}^{\mu}(i, u)
$$

- Cost function of $\mu$ at least as good as for $\pi$ :

$$
\bar{J}^{\pi} \geq \sum_{i, u} \mathbb{E}[g(i, u, W)] \cdot \nu^{\pi}(i, u)=\sum_{i, u} \mathbb{E}[g(i, u, W)] \cdot \nu_{\infty}^{\mu}(i, u)=\bar{J}^{\mu}
$$

Can restrict to (randomized) stationary policies $\mu$

## Proof that $\sum_{v} \nu^{\mu}(i, v)=\xi^{\mu}(i)$

- We have

$$
\begin{aligned}
\nu^{\pi}(i) & :=\sum_{u} \nu^{\pi}(i, u)=\sum_{j, u} \nu^{\pi}(j, u) P_{u, j i}=\sum_{j} \nu^{\pi}(i) \sum_{u} \mu(u \mid j) P_{u, j i} \\
& =\sum_{j} \nu^{\pi}(j) \mathrm{P}_{\mu, j i},
\end{aligned}
$$

- Therefore $\boldsymbol{\nu}^{\pi}$ equals the unique stationary distribution $\boldsymbol{\xi}^{\mu}$ of the $\mathrm{MC} \mathrm{P}_{\mu}$ induced by action policy $\mu$.

Linear Programme Solution based on State-Action Frequencies

- Since we can restrict to stationary distributions:


## "Dual Problem" for Average Costs

$$
\bar{J}^{*}=\min _{\nu(i, u) \geq 0} \sum_{i=1}^{m} \sum_{u} \mathbb{E}_{W}[g(i, u, W)] \cdot \nu(i, u)
$$

subject to:

$$
\begin{gather*}
\sum_{v} \nu(i, v)=\sum_{j=1}^{m} \sum_{u} \nu(j, u) P_{u, j i} \quad i=1, \ldots, m  \tag{1}\\
\sum_{i, u} \nu(i, u)=1
\end{gather*}
$$

- $m$ constraints are linearly dependent because both sides of (1) sum to 1 . $\rightarrow$ Optimal $\nu^{*}(i, u)>0$ for at most $m$ pairs $(i, u)$ ( $m$ lin. indep. constr.)

Deterministic stationary policy $\mu^{*}(u \mid i)=\frac{\nu^{*}(i, u)}{\sum_{v} \nu^{*}(i, v)}$ is optimal

## Value-Iteration Algorithm to Find Optimal Average Cost

- Modified update operator $\mathbb{T}_{\text {avg }}: \mathbf{V} \mapsto \min _{\mu}\left[\mathbb{E} w[g(i, \mu(i), W)]+\mathrm{P}_{\mu} \mathbf{V}\right]$
- A modified Bellman's equation holds
- For any initial vector $\mathbf{V}$ :

$$
\frac{1}{N} \mathbb{T}_{\text {avg }}^{N} \mathbf{V} \rightarrow \bar{J}^{*} \quad \text { as } N \rightarrow \infty
$$

- Value-iteration algorithm: Pick an arbitrary initial vector $\mathbf{J}_{0}$ and iterate until convergence:

$$
\mathbf{J}_{k+1}=\frac{k}{k+1} \mathbb{T}_{\text {avg }} \mathbf{J}_{k}, \quad k=0,1, \ldots,
$$

## Policy- Iteration Algorithm to Find Optimal Average Cost

- Modified operators $\mathbb{T}_{\text {avg }}$ and $\mathbb{T}_{\text {avg }, \mu}: \mathbf{V} \mapsto\left[\mathbb{E}_{w}[g(i, \mu(i), W)]+\mathrm{P}_{\mu} \mathbf{V}\right]$
- Policy-iteration algorithm: use above operators and slightly modified policy evaluation step.
- Start with arbitrary initial policy $\mu_{0}$ and iterate for $k=0,1, \ldots$ until $\mu_{k+1}=\mu_{k}$ :
(1) Policy evaluation: Find average and differential costs $J_{k} \in \mathbb{R}$ and $h_{k} \in \mathbb{R}^{m}$ satisfying for $i=1, \ldots, m$ :

$$
J_{k}+h_{k}(i)=\mathbb{E}\left[g\left(i, \mu_{k}(i), W\right)\right]+\sum_{j=1}^{m} P_{\mu_{k}, i j} h_{k}(j)
$$

$$
\left(J_{k}+h_{k}(i)=\mathbb{T}_{\text {avg }, \mu_{k}} \mathbf{h}_{k}\right)
$$

(2) Policy improvement: Find new policy $\mu_{k+1}$ satisfying for $i=1, \ldots, m$ :

$$
\begin{aligned}
& \quad \mu_{k+1}(i)+\sum_{j=1}^{m} P_{\mu_{k+1}, i j} h_{k}(j)=\min _{u \in \mathcal{U}}\left[\mathbb{E}_{W}[g(i, u, W)]+\sum_{j=1}^{m} P_{u, i j} h_{k}(j)\right] . \\
& \left(\mathbb{T}_{\text {avg }, \mu_{k+1}} \mathbf{h}_{k}=\mathbb{T}_{\mathrm{avg}} \mathbf{h}_{k}\right)
\end{aligned}
$$

## Average Infinite-Cost Case with L Cost-Constraints

- Optimal average infinite horizon cost:

$$
\bar{J}^{*}\left(p_{0}\right):=\min _{\pi} \bar{J}^{\pi}\left(p_{0}\right)
$$

where minimum is only over policies $\pi$ satisfying

$$
\varlimsup_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}_{X_{0},\left\{W_{k}\right\}}\left[\sum_{k=0}^{N-1} d_{\ell}\left(X_{k}, \mu_{k}\left(X_{k}\right), W_{k}\right)\right] \leq D_{\ell}, \quad \ell=1, \ldots, L .
$$

- Similar to before we can prove that we can restrict to stationary policies where the limsups are proper limits.
- Can express the average cost and the constraints with the state-action frequencies $\nu_{\infty}^{\mu}(i, u)$ of the stationary policies $\mu$


## Linear Programme for Optimal Average Cost with Constraints

## "Dual Problem" for Average Costs and Constraints

$$
\bar{J}^{*}=\min _{\nu(i, u) \geq 0} \sum_{i=1}^{m} \sum_{u} \mathbb{E}_{W}[g(i, u, W)] \cdot \nu(i, u)
$$

subject to:

$$
\begin{aligned}
\sum_{v} \nu(i, v) & =\sum_{j=1}^{m} \sum_{u} P_{u, i j} \cdot \nu(j, u), \quad i=1, \ldots, m \\
\sum_{i, u} \nu(i, u) & =1
\end{aligned}
$$

$$
\sum_{i=1}^{m} \sum_{u} \mathbb{E}_{W}\left[d_{\ell}(i, u, W)\right] \cdot \nu(i, u) \leq D_{\ell}, \quad \ell=1, \ldots, L
$$

- Optimal $\rho^{*}(i, u)>0$ for at most $m+L$ pairs $(i, u)$ (since there are $m+L$ lin. ind. constraints)

Maybe randomized actions in optimal policy $\mu^{*}=\frac{\nu^{*}(i, u)}{\sum_{\nu} \nu^{*}(i, v)}$

## Optimal Policy has L Randomization Points

- Randomized stationary policies with $L$ randomization points optimal
- Consider $L=1$ and optimal $\nu^{*}$ with $m+1$ positive entries:

$$
\nu^{*}\left(1, u_{1}\right), \nu^{*}\left(2, u_{2}\right), \nu^{*}\left(3, u_{3}\right), \ldots, \nu^{*}\left(m, u_{m}\right)>0
$$

and for some $j \in\{1, \ldots, m\}$ and $u_{j}^{\prime} \neq u_{j}$ :

$$
\nu^{*}\left(j, u_{j}^{\prime}\right)>0
$$

All other entries $\nu^{*}(i, u)=0$.

## Initial Randomization Suffices

- Idea: Randomize only at the beginning!
- Create the $m$-ary state-action frequencies

$$
\begin{aligned}
& \nu_{1}(i, u)= \begin{cases}\nu^{*}\left(j, u_{j}\right)+\nu^{*}\left(j, u_{j}^{\prime}\right) & i=j, u=u_{j} \\
0 & i=j, u=u_{j}^{\prime} \\
\mu^{*}(i, u), & \text { otherwise }\end{cases} \\
& \nu_{2}(i, u)= \begin{cases}0 & i=j, u=u_{j} \\
\nu^{*}\left(j, u_{j}\right)+\nu^{*}\left(j, u_{j}^{\prime}\right) & i=j, u=u_{j}^{\prime} \\
\mu^{*}(i, u), & \text { otherwise }\end{cases}
\end{aligned}
$$

- Construct the deterministic stationary policies

$$
\mu_{1}(u \mid i)=\frac{\nu_{1}(i, u)}{\sum_{v} \nu_{1}(i, v)} \quad \mu_{2}(u \mid i)=\frac{\nu_{2}(i, u)}{\sum_{v} \nu_{2}(i, v)}
$$

- At the beginning play each deterministic policy $\mu_{I}$ with prob. $q_{I}, I=1,2$,

$$
q_{1}:=\frac{\nu^{*}(j, u)}{\nu^{*}\left(j, u_{j}\right)+\nu^{*}\left(j, u_{j}^{\prime}\right)} \quad q_{2}:=\frac{\nu^{*}\left(j, u^{\prime}\right)}{\nu^{*}\left(j, u_{j}\right)+\nu^{*}\left(j, u_{j}^{\prime}\right)}
$$

## Initial Randomization Suffices, continued

- The expected cost of this mixed strategy is:

$$
\begin{aligned}
q_{1} \bar{J}^{\mu_{1}}+q_{2} \bar{J}^{\mu_{2}} & =q_{1} \sum_{i, u} \mathbb{E}[g(i, u, W)] \nu_{\infty}^{\mu_{1}}(i, u)+q_{2} \sum_{i, u} \mathbb{E}[g(i, u, W)] \nu_{\infty}^{\mu_{2}}(i, u) \\
& =q_{1} \sum_{i, u} \mathbb{E}[g(i, u, W)] \nu_{1}(i, u)+q_{2} \sum_{i, u} \mathbb{E}[g(i, u, W)] \nu_{2}(i, u) \\
& =\sum_{i, u} \mathbb{E}[g(i, u, W)]\left(q_{1} \cdot \nu_{1}(i, u)+q_{2} \cdot \nu_{2}(i, u)\right) \\
& =\sum_{i, u} \mathbb{E}[g(i, u, W)] \nu^{*}(i, u)=\bar{J}^{*}
\end{aligned}
$$

- The mixed strategy also satisfies the constraints for each $\ell=1, \ldots, L$ :

$$
\begin{aligned}
& q_{1} \sum_{i, u} \mathbb{E}\left[d_{\ell}(i, u, W)\right] \nu_{1}(i, u)+q_{2} \sum_{i, u} \mathbb{E}\left[d_{\ell}(i, u, W)\right] \nu_{2}(i, u) \\
&=\sum_{i, u} \mathbb{E}\left[d_{\ell}(i, u, W)\right]\left(q_{1} \cdot \nu_{1}(i, u)+q_{2} \cdot \nu_{2}(i, u)\right) \\
&=\sum_{i, u} \mathbb{E}\left[d_{\ell}(i, u, W)\right] \nu^{*}(i, u) \leq D_{l} l
\end{aligned}
$$

Optimal strategy: Randomly play one of $L$ deterministic policies

## Average Infinite-Cost Case with Constraints and Lagrange Multipliers

## "Dual Problem" for Average Costs and Constraints with Lagrange Multipliers

$$
\begin{aligned}
\bar{J}^{*}=\sup _{\lambda_{1}, \ldots, \lambda_{L} \geq 0} & \min _{\nu(i, u) \geq 0} \sum_{i=1}^{m} \sum_{u} \mathbb{E}_{W}\left[g(i, u, W)+\sum_{\ell} \lambda_{\ell} d_{\ell}(i, u, W)\right] \cdot \nu(i, u) \\
& -\sum_{\ell=1}^{L} \lambda_{\ell} D_{\ell}
\end{aligned}
$$

subject to:

$$
\begin{gathered}
\sum_{v} \nu(i, v)=\sum_{j=1}^{m} \sum_{u} P_{u, i j} \cdot \nu(j, u) \quad i=1, \ldots, m, \\
\sum_{i, u} \nu(i, u)=1 .
\end{gathered}
$$

- For each $\lambda_{1}, \ldots, \lambda_{L}$ a deterministic policy $\mu$ is optimal.


# Sequential Decision Processes, Master MICAS, Part I 

Michèle Wigger

Telecom Paris, 8 January 2021

## Lecture 7 - Algorithmic Dynamic Programming

## Algorithmic Paradigms

- Greedy Algorithm
- Construct solution incrementally
- Greedily choose the "right" subproblem by optimizing a local criterion
- Divide and Conquer
- Divide a problem into non-overlapping subproblems
- Solve each subproblem (in any order)
- Combine solutions of subproblems to obtain solution to initial problem
- Top-down approach



## Dynamic Programming (Bellman) Principle

- Breaking the problem into overlaping subproblems
- Calculate and store optimal solutions to subproblems
- Combine solutions to subproblems to solve the initial problem
- Solutions can be cached (stored) and reused

Top-down: Memoization


Bottom-up: Tabulation


Example: Binomial Coefficient $C_{n}^{k}=\binom{n}{k}=\frac{n!}{k!(n-k)!}$

## Recursive formula:

$$
C_{n}^{k}= \begin{cases}\binom{n-1}{k-1}+\binom{n-1}{k} & 0<k<n \\ 1 & \text { otherwise }\end{cases}
$$

Divide and Conquer Approach:

Function $C(n, k)$

1. if $(k=0)$ or $(k=n)$ return 1 ;
2. else return

$$
C(n-1, k-1)+C(n-1, k)
$$



- Time complexity:
- Exponential number of recursive calls: $O\left(\binom{n}{k}\right) \approx 2\binom{n}{k}$


## Example: Binomial Coefficient, continued

Pascal-triangle approach: Dynamic Programming with memoization based on 2-dimensional table

## Function C-mem(n, k)

```
1. for \((i=0 ; i \leq n ; i++)\)
2. for \((j=0 ; j \leq \min (i, k) ; j++)\)
3. if \((i=0)\) or \((j=i)\),
    \(T[i][j]=1\);
4. else
    \(T[i][j]=T[i-1][j-1]+T[i-1][j] ;\)
```

5. return $T[n][k]$;

|  | 0 | 1 | 2 | 3 | $\cdots$ | $n-1$ | $n$ |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0 | 1 |  |  |  |  |  |  |
| 1 | 1 | 1 |  |  |  |  |  |
| 2 | 1 | 2 | 1 |  |  |  |  |
| 3 | 1 | 3 | 3 | 1 |  |  |  |
| $\vdots$ | $\vdots$ | $\vdots$ | $\vdots$ |  | $\ddots$ |  |  |
| $n-1$ | 1 | $n-1$ | $\binom{n-1}{2}$ | $\binom{n-1}{3}$ | $\cdots$ | 1 |  |
| $n$ | 1 | $n$ | $\binom{n}{2}$ | $\binom{n}{3}$ | $\cdots$ | $n$ | 1 |

- Top -Down Approach
- Auxiliary space $O(n k)$ and time-complexity $O(n k)$.


## Example: Binomial Coefficient (3)

- Dynamic programming solution: Tabulation
- Create table with 1 dimension to compute small numbers
- Compute next row of pascal triangle using previous row Function C-dyn(n, k)

1. $T[0]=1$;
2. for $(i=0 ; i \leq n ; i++)$
3. $\quad$ for $(j=\min (i, k) ; j>0 ; j--)$ do $T[j]=T[j]+T[j-1]$;
4. return $T[k]$;

- Time complexity:
- Table of $k$ elements $\Rightarrow$ Auxiliary space $O(k)$
- Time complexity: $O(n k)$
- Optimized-space bottom-up DP approach


## How to design Dynamic Programming Solution

- Define subproblems
- Identify recursive relation between subproblems
- Avoid similar computation
- Resolve original problem by combining solutions of subproblems
- Tabulation approach:
- Recognize and solve the base cases
- Deduce dynamic programming algorithm in a bottom-up way
- Memoization approach:
- Deduce dynamic programming algorithm in a top-down way


# Sequential Decision Processes, Master MICAS, Part I 

Michèle Wigger

Telecom Paris, 8 January 2021

Lecture 7 - Some Shortest Paths Algorithms

## Deterministic MDPs and Shortest-Path Problems

- No disturbance $\rightarrow$ state evolution $x_{k+1}=f\left(x_{k}, u_{k}\right)$ and cost $g_{k}\left(x_{k}, u_{k}\right)$
- Graph representation:

- At each stage $k=1,2, \ldots, N$ there is a node for each $x_{k} \in \mathcal{X}$
- Arrows indicate transitions for different actions $\rightarrow$ label arrows with actions $u_{k}$ and costs $g_{k}\left(x_{k}, u_{k}\right)$
- Total cost $J_{0 \rightarrow N, \pi}$ is the sum of the costs on the path indicated by $\pi$

Finding minimum total cost $J_{0 \rightarrow N, \pi}$ equivalent to finding "shortest path" $\rightarrow$ DP algorithm can be run in reverse order

## Travelling Salesman Problem and Label Correcting Method

 Initialize $d_{s}=0$ and

- State space depends on stage $k$
$d_{2}=\cdots=d_{t}=$ upper $=\infty$


## Label Correcting Algorithm

Step 1: Remove a node $i$ from OPEN and for each child $j$ of $i$, execute step 2.
Step 2: If $d_{i}+a_{i j}<\min \left\{d_{j}\right.$, UPPER $\}$, set $d_{j}=d_{i}+a_{i j}$ and set $i$ to be the parent of $j$. In addition, if $j \neq t$, place $j$ in OPEN if it is not already in OPEN, while if $j=t$, set UPPER to the new value $d_{i}+a_{i t}$ of $d_{t}$.
Step 3: If OPEN is empty, terminate; else go to step 1.

| Iter. No. | Node Exiting OPEN | OPEN at the End of Iteration | UPPER |
| :---: | :---: | :---: | :---: |
| 0 | - | 1 | $\infty$ |
| 1 | 1 | $2,7,10$ | $\infty$ |
| 2 | 2 | $3,5,7,10$ | $\infty$ |
| 3 | 3 | $4,5,7,10$ | $\infty$ |
| 4 | 4 | $5,7,10$ | 43 |
| 5 | 5 | $6,7,10$ | 43 |
| 6 | 6 | 7,10 | 13 |
| 7 | 7 | 8,10 | 13 |
| 8 | 8 | 9,10 | 13 |
| 9 | 9 | 10 | 13 |
| 10 | 10 | Empty | 13 |

- Dijkstra's method always chooses the node in OPEN with smallest $d_{i}$.
- Bellman-Ford algorithm chooses the node in OPEN as first-in first-out.


## The Branch-and-Bound Algorithm

- Wish to minimize cost function $f(\cdot)$ over all elements of $\mathcal{X}$

Find functions $\bar{f}$ and $\underline{f}$ over subsets $\mathcal{Y} \subseteq \mathcal{X}$ such that :

$$
\underline{f}(\mathcal{Y}) \leq \min _{x \in \mathcal{Y}} f(x) \leq \bar{f}(\mathcal{Y}), \quad \forall \mathcal{Y} \subseteq \mathcal{X}
$$



- Construct a tree with subsets of $\mathcal{X}$ $\rightarrow$ including all singletons!
- If $\mathcal{Y}_{i} \subseteq \mathcal{Y} \Rightarrow \mathcal{Y}$ is a parent of $\mathcal{Y}_{i}$
- Label branch from $\mathcal{Y}$ to $\mathcal{Y}_{i}$ by $\underline{f}\left(\mathcal{Y}_{i}\right)-\underline{f}(\mathcal{Y}) \Rightarrow$ path length from $\mathcal{X}$ to $\mathcal{Y}$ equals $\underline{f}(\mathcal{Y})$


## Branch-and-Bound Algorithm

Step 1: Remove a node $Y$ from OPEN. For each child $Y_{j}$ of $Y$, do the following: If $\underline{f}_{Y j}<$ UPPER, then place $Y_{j}$ in OPEN. If in addition $\bar{f}_{Y j}<$ UPPER, then set UPPER $=\bar{f}_{Y j}$, and if $Y_{j}$ consists of a single solution, mark that solution as being the best solution found so far.
Step 2: (Termination Test) If OPEN is nonempty, go to step 1.
Otherwise, terminate; the best solution found so far is optimal.

