## 8: Gittins index

Concordia April 5, 2016

In this lecture, we consider a special type of MDP often encountered in supply chains and an efficient "divide-and-conquer" solution for it.

## 1 Multiarmed bandit



Figure 1: By Hagen, 2001.
Consider the following sequential decision problem (special case of an MDP). The time steps are $t=1,2, \ldots$ At every time step, the decision maker must choose one of $n$ alternative actions (imagine the case $n=2$ ). The state of the system is $n$-dimensional random variable:

$$
X_{t}=\left(X_{t}^{1}, \ldots, X_{t}^{n}\right) \in \mathbb{X}^{n}
$$

where $\mathbb{X}$ is a finite or countable set for simplicity.
Let $f^{1}, \ldots, f^{n}$ denote fixed deterministic functions, and let $\left\{W_{1}^{1}, W_{2}^{1}, \ldots\right\}, \ldots,\left\{W_{1}^{n}, W_{2}^{n}, \ldots\right\}$ denote i.i.d. sequences that are also mutually independent. If the $i$-th action is chosen,
i.e., $A_{t}=i$, then the state transition is as follows

$$
\begin{aligned}
X_{t+1}^{1} & =X_{t}^{1} \\
& \ldots \\
X_{t+1}^{i} & =f^{i}\left(X_{t}^{i}, W_{t}^{i}\right), \\
& \ldots \\
X_{t+1}^{n} & =X_{t}^{n}
\end{aligned}
$$

Observe that these transitions are Markovian, and only the $i$-th element changes.
The reward corresponding to action $A_{t}$ at time $t$ is

$$
\sum_{k=1}^{n} r^{k}\left(X_{t}^{k}\right) 1_{\left[A_{t}=k\right]}
$$

which depends only on the reward function for the arm $A_{t}$ that is chosen. The expected total discounted reward with discount factor $\beta \in(0,1)$ and given initial condition $X_{0}$ is

$$
\mathbb{E} \sum_{t=0}^{\infty} \beta^{t} \sum_{k=1}^{n} r^{k}\left(X_{t}^{k}\right) 1_{\left[A_{t}=k\right]},
$$

where the expectation is over the $\left\{W_{t}^{i}\right\}$ random variables.
Example 1.1 (Advertising and pricing). There are $n$ possible ads (prices). The state is the success frequency of each ad (price).

Due to the curse of dimensionality, when $n$ is large, typical dynamic programming solutions are not appropriate (complexity proportional to $\mathbb{X}^{n}$ ).

## 2 Gittins index

One approach to solve the above bandit problem is to assign a number to each pair $(x, i)$ of state $x \in \mathbb{X}$ and bandit $i$ :

$$
\nu^{i}(x)=\sup _{\tau} \frac{\mathbb{E}\left[\sum_{t=0}^{\tau} \beta^{t} r^{i}\left(X_{t}^{i}\right) \mid X_{0}^{i}=x\right]}{\mathbb{E}\left[\sum_{t=0}^{\tau} \beta^{t} \mid X_{0}^{i}=x\right]},
$$

where the sup is taken over all stopping times $\mathbb{1}^{1}$, or alternatively:

$$
\nu^{i}(x)=\sup \left\{\lambda: \sup _{\ell>0} \mathbb{E}\left[\sum_{t=0}^{\ell} \beta^{t}\left(r^{i}\left(X_{t}^{i}\right)-\lambda\right) \mid X_{0}^{i}=x\right]\right\} .
$$

This is the Gittins index of bandit $i$ at state $x$. Then, at every time step, the optimal action is to choose an action with the highest Gittins index.

[^0]There are various ways to compute the Gittins index. Consider a single bandit (with another action giving 0 reward). Instead of stopping rules, the Gittins index can also be defined in terms of stopping sets:

$$
S(x)=\left\{x^{\prime} \in \mathbb{X}: \nu\left(x^{\prime}\right) \leq \nu(x)\right\} .
$$

The Gittins index is then:

$$
\nu^{i}(x)=\sup _{S(x) \subseteq \mathbb{X}} \frac{\mathbb{E}\left[\sum_{t=0}^{\tau(S(x))} \beta^{t} r^{i}\left(X_{t}^{i}\right) \mid X_{0}^{i}=x\right]}{\mathbb{E}\left[\sum_{t=0}^{\tau(S(x))} \beta^{t} \mid X_{0}^{i}=x\right]},
$$

where $\tau(S(x))=\inf \left\{t>0: X_{t} \in S(x)\right\}$.

## 3 Computing the Gittins index

Consider a single bandit (with another action giving 0 reward).

- Input: $r, f$, distribution of $\left\{W_{t}\right\}$.
- Find the state with the highest index (with stopping set $S\left(x_{(1)}\right)=\mathbb{X}$ ): $x_{(1)}=$ $\arg \max _{x \in \mathbb{X}} r(x)$.
- Output $\nu\left(x_{(1)}\right)=r\left(x_{(1)}\right)$.
- Find $x_{(2)}$ with the second highest index, where the stopping set is $S\left(x_{(2)}\right)=\mathbb{X} \backslash$ $x_{(1)}$. Compute the probability distribution of $\tau\left(S\left(x_{(2)}\right)\right)=\inf \left\{t>0: X_{t} \neq x_{(1)}\right\}$.
- Output $\nu\left(x_{(2)}\right)$.
- Continue for $x_{(3)}, x_{(4)}, \ldots$


## 4 References

- "Multi-armed bandits, Gittins index, and its calculation." J. Chakravorty and A. Mahajan, Methods and Applications of Statistics in Clinical Trials, 2014.


[^0]:    ${ }^{1}$ For all $t$, the event $\tau=t$ depends only on $X_{0}, X_{1}, \ldots, X_{t}$.

