# 6.231 DYNAMIC PROGRAMMING 

## LECTURE 16

## LECTURE OUTLINE

- Review of computational theory of discounted problems
- Value iteration (VI), policy iteration (PI)
- Optimistic PI
- Computational methods for generalized discounted DP
- Asynchronous algorithms


## DISCOUNTED PROBLEMS

- Stationary system with arbitrary state space

$$
x_{k+1}=f\left(x_{k}, u_{k}, w_{k}\right), \quad k=0,1, \ldots
$$

- Bounded $g$. Cost of a policy $\pi=\left\{\mu_{0}, \mu_{1}, \ldots\right\}$
$J_{\pi}\left(x_{0}\right)=\lim _{N \rightarrow \infty} \underset{\substack{w_{k} \\ k=0,1, \ldots}}{E}\left\{\sum_{k=0}^{N-1} \alpha^{k} g\left(x_{k}, \mu_{k}\left(x_{k}\right), w_{k}\right)\right\}$
- Shorthand notation for DP mappings ( $n$-state Markov chain case)

$$
(T J)(x)=\min _{u \in U(x)} E\{g(x, u, w)+\alpha J(f(x, u, w))\}, \forall x
$$

$T J$ is the optimal cost function for the one-stage problem with stage cost $g$ and terminal cost $\alpha J$.

- For any stationary policy $\mu$

$$
\left(T_{\mu} J\right)(x)=E\{g(x, \mu(x), w)+\alpha J(f(x, \mu(x), w))\}, \forall x
$$

Note: $T_{\mu}$ is linear [in short $T_{\mu} J=P_{\mu}\left(g_{\mu}+\alpha J\right)$ ].

## "SHORTHAND" THEORY - A SUMMARY

- Cost function expressions (with $J_{0} \equiv 0$ )

$$
J_{\pi}=\lim _{k \rightarrow \infty} T_{\mu_{0}} T_{\mu_{1}} \cdots T_{\mu_{k}} J_{0}, \quad J_{\mu}=\lim _{k \rightarrow \infty} T_{\mu}^{k} J_{0}
$$

- Bellman's equation: $J^{*}=T J^{*}, J_{\mu}=T_{\mu} J_{\mu}$
- Optimality condition:

$$
\mu: \text { optimal }<==>\quad T_{\mu} J^{*}=T J^{*}
$$

- Contraction: $\left\|T J_{1}-T J_{2}\right\| \leq \alpha\left\|J_{1}-J_{2}\right\|$
- Value iteration: For any (bounded) $J$

$$
J^{*}=\lim _{k \rightarrow \infty} T^{k} J
$$

- Policy iteration: Given $\mu^{k}$,
- Policy evaluation: Find $J_{\mu^{k}}$ by solving

$$
J_{\mu^{k}}=T_{\mu^{k}} J_{\mu^{k}}
$$

- Policy improvement: Find $\mu^{k+1}$ such that

$$
T_{\mu^{k+1}} J_{\mu^{k}}=T J_{\mu^{k}}
$$

## INTERPRETATION OF VI AND PI



## VI AND PI METHODS FOR Q-LEARNING

- We can write Bellman's equation as

$$
J^{*}(i)=\min _{u \in U(i)} Q^{*}(i, u) \quad i=1, \ldots, n
$$

where $Q^{*}$ is the vector of optimal Q -factors

$$
Q^{*}(i, u)=\sum_{j=1}^{n} p_{i j}(u)\left(g(i, u, j)+\alpha J^{*}(j)\right)
$$

- VI and PI for Q-factors are mathematically equivalent to VI and PI for costs.
- They require equal amount of computation ... they just need more storage.
- For example, we can write the VI method as

$$
J_{k+1}(i)=\min _{u \in U(i)} Q_{k+1}(i, u), \quad i=1, \ldots, n
$$

where $Q_{k+1}$ is generated for all $i$ and $u \in U(i)$ by

$$
Q_{k+1}(i, u)=\sum_{j=1}^{n} p_{i j}(u)\left(g(i, u, j)+\alpha \min _{v \in U(j)} Q_{k}(j, v)\right)
$$

## APPROXIMATE PI

- Suppose that the policy evaluation is approximate, according to,

$$
\max _{x}\left|J_{k}(x)-J_{\mu^{k}}(x)\right| \leq \delta, \quad k=0,1, \ldots
$$

and policy improvement is approximate, according to,
$\max _{x}\left|\left(T_{\mu^{k+1}} J_{k}\right)(x)-\left(T J_{k}\right)(x)\right| \leq \epsilon, \quad k=0,1, \ldots$ where $\delta$ and $\epsilon$ are some positive scalars.

- Error Bound: The sequence $\left\{\mu^{k}\right\}$ generated by approximate policy iteration satisfies

$$
\limsup _{k \rightarrow \infty} \max _{x \in S}\left(J_{\mu^{k}}(x)-J^{*}(x)\right) \leq \frac{\epsilon+2 \alpha \delta}{(1-\alpha)^{2}}
$$

- Typical practical behavior: The method makes steady progress up to a point and then the iterates $J_{\mu^{k}}$ oscillate within a neighborhood of $J^{*}$.


## OPTIMISTIC PI

- This is PI, where policy evaluation is carried out by a finite number of VI
- Shorthand definition: For some integers $m_{k}$

$$
\begin{aligned}
& T_{\mu^{k}} J_{k}=T J_{k}, \quad J_{k+1}=T_{\mu^{k}}^{m_{k}} J_{k}, \quad k=0,1, \ldots \\
& \quad-\text { If } m_{k} \equiv 1 \text { it becomes VI } \\
& \quad-\text { If } m_{k}=\infty \text { it becomes PI } \\
& \quad-\text { For intermediate values of } m_{k}, \text { it is generally } \\
& \quad \text { more efficient than either VI or PI }
\end{aligned}
$$



Approx. Policy Evaluation

## EXTENSIONS TO GENERALIZED DISC. DP

- All the preceding VI and PI methods extend to generalized/abstract discounted DP.
- Summary: For a mapping $H: X \times U \times R(X) \mapsto$ $\Re$, consider

$$
\begin{array}{ll}
(T J)(x)=\min _{u \in U(x)} H(x, u, J), & \forall x \in X . \\
\left(T_{\mu} J\right)(x)=H(x, \mu(x), J), & \forall x \in X .
\end{array}
$$

- We want to find $J^{*}$ such that

$$
J^{*}(x)=\min _{u \in U(x)} H\left(x, u, J^{*}\right), \quad \forall x \in X
$$

and a $\mu^{*}$ such that $T_{\mu^{*}} J^{*}=T J^{*}$.

- Discounted, Discounted Semi-Markov, Minimax

$$
\begin{gathered}
H(x, u, J)=E\{g(x, u, w)+\alpha J(f(x, u, w))\} \\
H(x, u, J)=G(x, u)+\sum_{y=1}^{n} m_{x y}(u) J(y) \\
H(x, u, J)=\max _{w \in W(x, u)}[g(x, u, w)+\alpha J(f(x, u, w))]
\end{gathered}
$$

## ASSUMPTIONS AND RESULTS

- Monotonicity assumption: If $J, J^{\prime} \in R(X)$ and $J \leq J^{\prime}$, then

$$
H(x, u, J) \leq H\left(x, u, J^{\prime}\right), \quad \forall x \in X, u \in U(x)
$$

- Contraction assumption:
- For every $J \in B(X)$, the functions $T_{\mu} J$ and $T J$ belong to $B(X)$.
- For some $\alpha \in(0,1)$ and all $J, J^{\prime} \in B(X), H$ satisfies

$$
\left|H(x, u, J)-H\left(x, u, J^{\prime}\right)\right| \leq \alpha \max _{y \in X} \mid J(y)-J^{\prime}(y)
$$

for all $x \in X$ and $u \in U(x)$.

- Standard algorithmic results extend:
- Generalized VI converges to $J^{*}$, the unique fixed point of $T$
- Generalized PI and optimistic PI generate $\left\{\mu^{k}\right\}$ such that

$$
\lim _{k \rightarrow \infty}\left\|J_{\mu^{k}}-J^{*}\right\|=0, \quad \lim _{k \rightarrow \infty}\left\|J_{k}-J^{*}\right\|=0
$$

- Analytical Approach: Start with a problem, match it with an $H$, invoke the general results.


## ASYNCHRONOUS ALGORITHMS

- Motivation for asynchronous algorithms
- Faster convergence
- Parallel and distributed computation
- Simulation-based implementations
- General framework: Partition $X$ into disjoint nonempty subsets $X_{1}, \ldots, X_{m}$, and use separate processor $\ell$ updating $J(x)$ for $x \in X_{\ell}$.
- Let $J$ be partitioned as $J=\left(J_{1}, \ldots, J_{m}\right)$, where $J_{\ell}$ is the restriction of $J$ on the set $X_{\ell}$.
- Synchronous algorithm: Processor $\ell$ updates $J$ for the states $x \in X_{\ell}$ at all times $t$,

$$
J_{\ell}^{t+1}(x)=T\left(J_{1}^{t}, \ldots, J_{m}^{t}\right)(x), \quad x \in X_{\ell}, \ell=1, \ldots, m
$$

- Asynchronous algorithm: Processor $\ell$ updates $J$ for the states $x \in X_{\ell}$ only at a subset of times $\mathcal{R}_{\ell}$,

$$
J_{\ell}^{t+1}(x)= \begin{cases}T\left(J_{1}^{\tau_{\ell 1}(t)}, \ldots, J_{m}^{\tau_{\ell}(t)}\right)(x) & \text { if } t \in \mathcal{R}_{\ell} \\ J_{\ell}^{t}(x) & \text { if } t \notin \mathcal{R}_{\ell}\end{cases}
$$

where $t-\tau_{\ell j}(t)$ are communication"delays"

## ONE-STATE-AT-A-TIME ITERATIONS

- Important special case: Assume $n$ "states", a separate processor for each state, and no delays - Generate a sequence of states $\left\{x^{0}, x^{1}, \ldots\right\}$, generated in some way, possibly by simulation (each state is generated infinitely often)
- Asynchronous VI: Change any one component of $J^{t}$ at time $t$, the one that corresponds to $x^{t}$ :

$$
J^{t+1}(\ell)= \begin{cases}T\left(J^{t}(1), \ldots, J^{t}(n)\right)(\ell) & \text { if } \ell=x^{t}, \\ J^{t}(\ell) & \text { if } \ell=x^{t},\end{cases}
$$

- The special case where

$$
\left\{x^{0}, x^{1}, \ldots\right\}=\{1, \ldots, n, 1, \ldots, n, 1, \ldots\}
$$

## is the Gauss-Seidel method

- More generally, the components used at time $t$ are delayed by $t-\tau_{\ell j}(t)$
- Flexible in terms of timing and "location" of the iterations
- We can show that $J^{t} \rightarrow J^{*}$ under assumptions typically satisfied in DP


## ASYNCHRONOUS CONV. THEOREM I

- Assume that for all $\ell, j=1, \ldots, m$, the set of times $\mathcal{R}_{\ell}$ is infinite and $\lim _{t \rightarrow \infty} \tau_{\ell j}(t)=\infty$
- Proposition: Let $T$ have a unique fixed point $J^{*}$, and assume that there is a sequence of nonempty subsets $\{S(k)\} \subset R(X)$ with $S(k+1) \subset S(k)$ for all $k$, and with the following properties:
(1) Synchronous Convergence Condition: Every sequence $\left\{J^{k}\right\}$ with $J^{k} \in S(k)$ for each $k$, converges pointwise to $J^{*}$. Moreover, we have

$$
T J \in S(k+1), \quad \forall J \in S(k), k=0,1, \ldots
$$

(2) Box Condition: For all $k, S(k)$ is a Cartesian product of the form

$$
S(k)=S_{1}(k) \times \cdots \times S_{m}(k),
$$

where $S_{\ell}(k)$ is a set of real-valued functions on $X_{\ell}, \ell=1, \ldots, m$.

Then for every $J \in S(0)$, the sequence $\left\{J^{t}\right\}$ generated by the asynchronous algorithm converges pointwise to $J^{*}$.

## ASYNCHRONOUS CONV. THEOREM II

- Interpretation of assumptions:


A synchronous iteration from any $J$ in $S(k)$ moves into $S(k+1)$ (component-by-component)

- Convergence mechanism:


Key: "Independent" component-wise improvement. An asynchronous component iteration from any $J$ in $S(k)$ moves into the corresponding component portion of $S(k+1)$ permanently!

## PRINCIPAL DP APPLICATIONS

- The assumptions of the asynchronous convergence theorem are satisfied in two principal cases:
- When $T$ is a (weighted) sup-norm contraction.
- When $T$ is monotone and the Bellman equation $J=T J$ has a unique solution.
- The theorem can be applied also to convergence of asynchronous optimistic PI for:
- Discounted problems (Section 2.6.2 of the text).
- SSP problems (Section 3.5 of the text).
- There are variants of the theorem that can be applied in the presence of special structure.
- Asynchronous convergence ideas also underlie stochastic VI algorithms like Q-learning.

MIT OpenCourseWare
http://ocw.mit.edu

### 6.231 Dynamic Programming and Stochastic Control

Fall 2015

For information about citing these materials or our Terms of Use, visit: http://ocw.mit.edu/terms.

