

A SHORT PROOF OF THE GITTINS INDEX THEOREM¹

BY JOHN N. TSITSIKLIS

Massachusetts Institute of Technology

We provide a short and elementary proof of the Gittins index theorem for the multi-armed bandit problem, for the case where each bandit is modeled as a finite-state semi-Markov process. We also indicate how this proof can be extended to the branching bandits and Klimov problems.

1. Introduction. There is a long history of alternative proofs of the Gittins index theorem for the multi-armed bandit problem. The original proof of Gittins and Jones [3] and a later proof by Gittins [2] relied on an interchange argument. A different interchange argument was provided by Varaiya, Walrand and Buyukkoc [9] and was simplified further in [10]. The simplest interchange argument available seems to be the one by Weiss [12], which in fact establishes an index theorem for the more general branching bandits model. A different proof, based on dynamic programming, was provided by Whittle [13] and subsequently simplified by Tsitsiklis [7]. Weber [11] has outlined a new proof that avoids any calculations and rests on more qualitative reasoning. Finally, a proof based on a polyhedral characterization of a suitably defined “performance region” has been provided by Tsoucas [8], for the average-cost Klimov problem, and by Bertsimas and Nino-Mora [1] for many other classes of multi-armed bandit problems.

This paper presents yet another proof of the same result. This proof has some common elements with the proof in [12] but appears to be simpler in that it is based on a simple inductive argument and uses only trivial calculations. The induction is in terms of the cardinality of the state spaces of the bandits involved and, for this reason, the proof is valid only for the case of finite-state bandits.

The rest of the paper is organized as follows. Section 2 presents the model to be employed and the proof of our main result. Section 3 contains some discussion on how to handle the more general branching bandits and Klimov problems.

2. The multi-armed bandit model. There are n bandit processes. The i th such process is a semi-Markov process with a finite state space \mathcal{X}_i . We assume, for simplicity, that the state spaces of the different bandits are disjoint and we let $\mathcal{X} = \mathcal{X}_1 \cup \dots \cup \mathcal{X}_n$. If the i th bandit is at some state $x \in \mathcal{X}_i$

Received February 1993; revised June 1993.

¹Supported by ARO Grant DAAL03-92-G-0115 and by Siemens A. G.

AMS 1991 subject classifications. Primary 60G40, 90B35; secondary 62L05, 90C40.

Key words and phrases. Multi-armed bandit problem, branching bandits, Klimov’s problem, priority scheduling.

and is selected to be “played,” then a random reward $R(x)$ is received and the bandit remains active over a time period of random length $T(x)$. After $T(x)$ time units, the play is completed and the bandit moves to a random new state $Y(x)$. At that point, we are free to choose the same or another bandit to be played.

We assume that the joint probability distribution of the random vector $(T(x), R(x), Y(x))$ is known and is the same for every play of bandit i for which bandit i is at the same state $x \in \mathcal{X}_i$. In addition, the random vectors corresponding to different plays of the same or of different bandits are assumed to be statistically independent.

A *policy* for the multi-armed bandit problem is defined as a mapping $\pi: \mathcal{X}_1 \times \cdots \times \mathcal{X}_n \mapsto \{1, \dots, n\}$ which at time zero and at any play completion time chooses a bandit to be played next, as a function of the current states of the n bandits. Given a particular policy, the time t_i at which the i th play starts and the reward R_i received at that time are well-defined random variables. Let $\beta > 0$ be a discount rate. We are interested in the problem of finding a policy that maximizes the expected discounted reward

$$E \left[\sum_{i=1}^{\infty} R_i e^{-\beta t_i} \right],$$

for every initial state.

We now comment on some consequences of the general fact that it is only $E[R(x)]$ and not the entire probability distribution of $R(x)$ that matters. Let us fix a particular policy. Let x_i be the state of the bandit that is played at the i th play and let \mathcal{F}_i stand for all random variables realized during the first $i - 1$ plays. In particular, t_i and x_i are determined by the outcomes of the first $i - 1$ plays and are contained in \mathcal{F}_i . Conditioned on \mathcal{F}_i , the expected discounted reward resulting from the i th play is given by

$$e^{-\beta t_i} E[R(x_i) | x_i].$$

Let us now consider the following alternative reward structure: Whenever a bandit at some state x is played, then rewards are received throughout the duration of that play at a constant rate $r(x)$, where

$$(2.1) \quad r(x) = \frac{E[R(x)]}{E \left[\int_0^{T(x)} e^{-\beta t} dt \right]}.$$

Under this new reward structure, the expected reward resulting from the i th play, conditioned on \mathcal{F}_i , is given by

$$\begin{aligned} E \left[\int_{t_i}^{t_i + T(x_i)} e^{-\beta t} r(x_i) dt \middle| \mathcal{F}_i \right] &= e^{-\beta t_i} r(x_i) E \left[\int_0^{T(x_i)} e^{-\beta t} dt \middle| x_i \right] \\ &= e^{-\beta t_i} E[R(x_i) | x_i], \end{aligned}$$

where the last equality follows from (2.1). It follows that, under either reward structure, the infinite-horizon expected discounted reward of any policy is the same. We will be using this fact later in our proof.

We say that a policy is a *priority rule* if there is an ordering of the elements of $\mathcal{X} = \mathcal{X}_1 \cup \dots \cup \mathcal{X}_n$ such that, at each decision point, the bandit whose state is ordered highest is chosen. Our basic result is the following.

THEOREM 2.1. *If each \mathcal{X}_i , $i = 1, \dots, n$, is finite, then there exists a priority rule which is optimal.*

PROOF. Let N be the cardinality of the set \mathcal{X} . The proof proceeds by induction on N . If $N = 1$, we have a single bandit and the only available policy is trivially a priority rule.

Let us now assume that the result is true for all multi-armed bandit problems for which $N = K$, where K is some positive integer. We consider a multi-armed bandit problem for which $N = K + 1$, and we will show that there exists an optimal policy which is a priority rule. This will complete the induction and the proof of the theorem.

Let us pick some state $s^* \in \mathcal{X}$ such that $r(s^*) = \max_{x \in \mathcal{X}} r(x)$. Let i^* be such that $s^* \in \mathcal{X}_{i^*}$. The following lemma states that s^* can be chosen as a top priority state.

LEMMA 2.1. *There exists an optimal policy that obeys the following rule: Whenever bandit i^* is at state s^* , then bandit i^* is played.*

PROOF. Consider an optimal policy π . Suppose that at time 0, bandit i^* is at state s^* . If policy π chooses to play bandit i^* , then there is nothing to prove. Suppose now that π chooses some other bandit to play. Define the random variable τ as the first time at which bandit i^* is played under policy π . (We let $\tau = \infty$ if bandit i^* is never played.)

We now define a new policy, call it π' , which plays bandit i^* once and from then on mimics the actions of policy π . However, when (and if) policy π plays bandit i^* for the first time, policy π' skips that play of bandit i^* . Let $\bar{r}(t)$ be the reward rate, as a function of time, under policy π . [That is, if $x(t)$ is the state of the bandit that is played at time t , then $\bar{r}(t) = r(x(t))$.] Using the definition of s^* , we have $\bar{r}(t) \leq r(s^*)$ for all t .

The expected discounted reward $J(\pi)$ under policy π is given by

$$J(\pi) = E \left[\int_0^\tau \bar{r}(t) e^{-\beta t} dt + e^{-\beta \tau} \int_0^{T(s^*)} r(s^*) e^{-\beta t} dt + \int_{\tau+T(s^*)}^\infty \bar{r}(t) e^{-\beta t} dt \right].$$

Similarly, the expected discounted reward $J(\pi')$ under policy π' is given by

$$J(\pi') = E \left[\int_0^{T(s^*)} r(s^*) e^{-\beta t} dt + e^{-\beta T(s^*)} \int_0^\tau \bar{r}(t) e^{-\beta t} dt + \int_{\tau+T(s^*)}^\infty \bar{r}(t) e^{-\beta t} dt \right].$$

We wish to show that $J(\pi') \geq J(\pi)$. Equivalently, we wish to show that

$$(2.2) \quad E \left[(1 - e^{-\beta \tau}) \int_0^{T(s^*)} r(s^*) e^{-\beta t} dt \right] \geq E \left[(1 - e^{-\beta T(s^*)}) \int_0^\tau \bar{r}(t) e^{-\beta t} dt \right].$$

We note that if $\bar{r}(t)$ were equal to $r(s^*)$ for all t , then the two sides of (2.2) would be equal. This observation and the fact $\bar{r}(t) \leq r(s^*)$ show that (2.2) is valid and, therefore, $J(\pi') \geq J(\pi)$. Since π was assumed optimal, π' is also optimal. However, if it is optimal to give top priority to state s^* at time 0, then (by the optimality of stationary policies) it is also optimal to give top priority to state s^* at every decision time. \square

Lemma 2.1 states that there exists an optimal policy within the set of policies that give top priority to state s^* ; call this set of policies $\Pi(s^*)$. We will now consider the problem of finding a policy which is optimal within the set $\Pi^*(s)$.

If s^* is the only possible state of bandit i^* , then the policy that always plays bandit i^* is evidently optimal and is a priority rule. We henceforth assume that \mathcal{X}_{i^*} is not a singleton. Suppose that bandit i^* is in some state $x \neq s^*$ and that this bandit is played. If this play causes a transition to state s^* , bandit i^* will be played again and again until eventually a transition to some state different from s^* results. We can view this succession of plays as a single (composite) play which cannot be interrupted due to our restriction to $\Pi(s^*)$. This single play has a random duration $\hat{T}(x)$ equal to the total time elapsed until a transition to a state different than s^* . Furthermore, by the discussion preceding the statement of Theorem 2.1, the reward of every policy remains the same if the discounted reward $\int_0^{\hat{T}(x)} e^{-\beta t} \bar{r}(t) dt$ received during this composite play is replaced by a constant reward rate equal to

$$(2.3) \quad \hat{r}(x) = \frac{E \left[\int_0^{\hat{T}(x)} e^{-\beta t} \bar{r}(t) dt \right]}{E \left[\int_0^{\hat{T}(x)} e^{-\beta t} dt \right]},$$

to be received throughout the duration of this composite play. We may thus replace bandit i^* by a new bandit in which state s^* is absent, $T(x)$ and $r(x)$ are replaced by $\hat{T}(x)$ and $\hat{r}(x)$, respectively, and the transition probabilities are suitably modified. We call this procedure "reducing bandit i^* by removing state s^* ."

The preceding argument shows that the problem of finding an optimal policy within the class $\Pi(s^*)$ is a new multi-armed bandit problem for which the sum of the cardinalities of the state spaces of the different bandits is equal to K . The induction hypothesis shows that there exists a priority rule $\hat{\pi}$ which is optimal for the latter problem. It follows that there exists a priority rule which is optimal for the original problem: Give top priority to state s^* and follow the priority rule $\hat{\pi}$ for the remaining states. \square

To every state $x \in \mathcal{X}$, we associate a number $\gamma(x)$, which we will call an *index*, using the following procedure:

INDEX ALGORITHM. (a) Pick a state s^* such that $r(s^*) = \max_{x \in \mathcal{X}} r(x)$ and let $\gamma(s^*) = r(s^*)$. Let i^* be such that $s^* \in \mathcal{X}_{i^*}$.

(b) If \mathcal{X}_{i^*} is a singleton, remove bandit i^* . Else, reduce bandit i^* by removing state s^* . Go back to (a).

From the proof of Theorem 2.1, it is apparent that the statistics of the random variables $\hat{T}(x)$ and $\hat{r}(x)$, as well as the transition probabilities of the reduced bandit i^* are completely determined by the corresponding statistics and transition probabilities of the original bandit i^* . This shows that the indices of the various states of a particular bandit are completely determined by the statistics associated with that bandit. In other words, the index algorithm can be carried out separately for each different bandit, still yielding the same index values.

The Gittins index theorem establishes something more than Theorem 2.1. In particular, not only does it show that there exists a priority policy which is optimal, but also that an optimal priority ordering can be found by ordering the states according to the numerical values of a certain index which can be computed separately for each bandit. We can also get this stronger result as follows:

THEOREM 2.2. *Let the index of each state be determined according to our index algorithm. Then any priority policy in which states with a higher index have higher priority is optimal.*

PROOF. The proof of Theorem 2.1 shows that any priority policy that orders the states in the same order as they are picked by the index algorithm is optimal. Therefore, it only remains to show that x can be picked before y by the index algorithm if and only if $\gamma(x) \geq \gamma(y)$. Given the recursive nature of the algorithm, it suffices to show that if state s^* is the first one picked by the index algorithm and q^* is the next state to be picked, then $\gamma(s^*) \geq \gamma(q^*)$. Let i^* be such that $s^* \in \mathcal{X}_{i^*}$. If $q^* \in \mathcal{X}_{i^*}$, then, using (2.3), we have

$$\gamma(q^*) = \hat{r}(q^*) \leq \max_{x \in \mathcal{X}} r(x) = r(s^*) = \gamma(s^*).$$

If on the other hand $q^* \notin \mathcal{X}_{i^*}$, then $\gamma(q^*) = r(q^*) \leq r(s^*) = \gamma(s^*)$. \square

For the case of discrete-time Markov bandits, our index algorithm is the same as the one in [6] and [9]. These references, as well as [10], provide some more detail on how the needed calculations can be carried out. Our algorithm is also a special case of the algorithm in [12].

3. Discussion. The proof given here is very simple and it is quite surprising that it was not known earlier. Perhaps a reason is that for the proof to go through, we have to consider semi-Markov bandits rather than the usual discrete-time Markov bandits.

We remark that our proof is easily extended to cover the case of arm-acquiring bandits [14] and the even more general case of branching bandits, thus recovering the results of [12]. We assume that the reader is familiar

with the framework of [12] and we only point out a few minor modifications of the proof of Theorem 2.1 that are needed. Instead of assuming that the different bandits have disjoint state spaces, we now assume that all bandits share the same state space. We then use induction on the cardinality of this common state space. As in the proof of Theorem 2.1, we pick a top priority state s^* whose reward rate is maximal. We then “eliminate” state s^* and form a reduced bandit as follows: If a bandit at some state $x \neq s^*$ is played, then the play lasts until all type s^* descendants of that bandit have been eliminated; the reward rate during this composite play is also suitably defined, similarly to (2.3). The resulting index algorithm is identical to the algorithm in [12].

We finally comment that the same approach goes through for the average-cost Klimov problem [4, 10] as well. As in Lemma 2.1, an interchange argument establishes that highest priority can be given to a particular customer class, and the problem is then reduced to one with a smaller number of customer classes. The resulting algorithm is the same as the one established in [5] by means of a more complicated interchange argument.

REFERENCES

- [1] BERTSIMAS, D. and NINO-MORA, J. (1993). Conservation laws, extended polymatroids and multi-armed bandit problems: A unified approach to indexable systems. Unpublished manuscript.
- [2] GITTINS, J. C. (1989). *Multi-Armed Bandit Allocation Indices*. Wiley, New York.
- [3] GITTINS, J. C. and JONES, D. M. (1974). A dynamic allocation index for the design of experiments. In *Progress in Statistics* (J. Gani, K. Sarkadi and I. Vince, eds.) *Colloq. Math. Soc. János Bolyai* **9**. North-Holland, Amsterdam.
- [4] KLIMOV, G. P. (1974). Time-sharing service systems I. *Theory Probab. Appl.* **19** 532–551.
- [5] NAIN, P., TSUCAS, P. and WALRAND, J. (1989). Interchange arguments in stochastic scheduling. *J. Appl. Probab.* **27** 815–826.
- [6] ROBINSON, D. R. (1981). Algorithms for evaluating the dynamic allocation index. Research Report 80/DRR/4, Manchester–Sheffield School of Probability and Statistics.
- [7] TSITSIKLIS, J. N. (1986). A lemma on the multi-armed bandit problem. *IEEE Trans. Automat. Control* **31** 576–577.
- [8] TSUCAS, P. (1991). The region of achievable performance in a model of Klimov. Research report, IBM.
- [9] VARAIYA, P., WALRAND, J. and BUYUKKOC, C. (1985). Extensions of the multi-armed bandit problem: The discounted case. *IEEE Trans. Automat. Control* **AC-30** 426–439.
- [10] WALRAND, J. (1988). *An Introduction to Queueing Networks*. Prentice-Hall, Englewood Cliffs, NJ.
- [11] WEBER, R. (1992). On the Gittins index for multiarmed bandits. *Ann. Appl. Probab.* **2** 1024–1033.
- [12] WEISS, G. (1988). Branching bandit processes. *Probab. Engng. Inform. Sci.* **2** 269–278.
- [13] WHITTLE, P. (1980). Multi-armed bandits and the Gittins index. *J. Royal Statist. Soc. Ser. B* **42** 143–149.
- [14] WHITTLE, P. (1981). Arm acquiring bandits. *Ann. Probab.* **9** 284–292.

LABORATORY FOR INFORMATION
AND DECISION SYSTEMS
ROOM 35-214
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
CAMBRIDGE, MASSACHUSETTS 02139