Massachusetts Institute of Technology | Lecture 24
6.892: Theory of Parallel Systems | May 15, 1997
Professor Charles E. Leiserson | Scribe: Cotton Seed, Mehul Shah, Michael Taylor

Lecture 24: Resource Allocation in the Face of Uncertainty

1 Introduction

This lecture motivates and considers schemes for scheduling workstation resources in the face of uncertainty. Initially, we describe the single player "Pick-a-Winner" game in which one attempts to maximize "profit". This game maps directly onto the scheduling problem. There are several obvious, yet misleading observations pointing to a simple scheme for maximizing profit. We discuss those and present a counterintuitive optimal strategy. Intuitions and essential facts are delineated to support the optimal strategy. Finally, we summarize and show that the optimal strategy can be easily extended to the multi-player version of our game.

2 Pick-a-Winner Game

We present the rules of this game and give a simple example. Next, we discuss an on-line version by first describing the scenario and then asking the question of maximizing profit.

Rules: One player has $n = 8$ options (A-H) to choose from. These can be anything from stocks to workstation resources. The options pay dividends of at most a dollar each day. There are $d = 10$ days from which the player can choose an option. The player may only make a single selection at any time, and no changes are allowed. The player's profit, P, accrues only after he has made the selection.

Example: If D is chosen on the fourth day, then the players profit is $P = \$2$. The optimal profit in this example is $P_{OPT} = \$6$ achieved by choosing option B on the first day.

	1	2	3	4	5	6	7	8	9	10
A	\$1	\$1								\$1
B	\$1		\$1		\$1	\$1		\$1	\$1	
C	\$1			\$1						
D		\$1	\$1	\$1	\$1					
E	\$1					\$1				
F	\$1					\$1		\$1		
G					\$1					
H										

Now consider an on-line version of this game, where the player sees the dividends by each option day-by-day, with the following conditions. The past dividend payments are known but the future payment schedule is *not* known. The past and future payments are not necessarily correlated. Finally, the future is not necessarily random. The future is pre-determined by an adversary who knows the player's strategy. The fundamental question is how much money the player make?

3 Maximizing Profit

Keeping in mind that we are considering the worst case scenario under which an adversary knows the player's strategy, we make a few observations leading to a simple scheme. The optimal solution, however, is counterintuitive.

Clearly, the player makes no profit for any deterministic strategy even if the optimal profit is positive. The adversary can simply pay a dividend on each day until the player makes a selection. Then the dividends are stopped for the selected option. However, with a random strategy the player obtains a better profit. The adversary only knows the player's strategy, which is choosing randomly. Thus, he cannot plot deliberately against the player since he has no knowledge of the random numbers.

The obvious strategy is to choose a random option at the start. In this case, the player can get an expected profit of P_{OPT}/n because he has $1/n$ chance of selecting the optimal option. This scheme also seems to be the optimal for two reasons. Since we do not the future, a random choice is the best we can do. Also, the past has no bearing on the future, so it cannot help to delay the selection.

However, that strategy is not the best randomized strategy possible. In fact, with the right scheme, a player can expect a profit of $c\,P_{OPT}/log\,n$ with 99% probability if he knows P_{OPT}. Without knowledge of P_{OPT}, the expected return is only a little worse, $c\,P_{OPT}/(log\,n\;\;log\,d)$. The constant c depends on the desired probability of success. The probability of success depends only on the player's random choices and holds for *all* payment schedules. Furthermore, these are the optimal bounds for maximizing profit.

4 Optimal Strategy

The player can do much better than making a random choice at the start with the right randomize strategy.

Theorem 1 *The player can get profit $cP_{opt}/\log n$ with probability 0.99 if he knows P_{opt}*

c is a constant that depends on the desired probability of success. The probability of success depends only on the player's random numbers. The probabilty of success holds for *all* payment schedules. These bounds are the best possible.

Without knowledge of P_{opt}, the expected return is $cP_{opt}/(\log n \log d)$.

Player selects target profit P and confidence level 1 - q where $P_{opt} > 2P \log n/q$. The following algorithm returns a profit greater than P with probabilty greater than the confidence $1 - q$.

for t = 1 to d (days)
for S = 1 to n (options)
if a selection has not yet been made, and
 S pays its ith dividend on day t,
 then we select S with probabilty $2^{qi/P}/Pn^2$

5 Intuition and Key Facts

The selection probabilities start small and grow quickly (but not too quickly).

The **slow start** is so that the adversary can't mislead the player with false starts – e.g., options that pay P-1 dividends early and then nothing later.

$$Prob[MisleadPlayer] < Pn\left(\frac{2^q}{Pn^2}\right) = O\left(\frac{1}{n}\right)$$

We need **fast growth** to insure that we make some selection with high probability.

$$Prob[NoSelection] < e^{-\frac{n}{2}}$$

But, the growth must be **slow enough** to insure that whatever option is selected, the a posteriori probability is high that the option was selected at least P payments before the final payment (see Figure 1).

6 Summary

Combining the preceding 3 conditions yields a proof:

If P_{opt} is known, choose target P = q $P_{opt}/(3logn)$ and confidence parameter 1 - q as desired.

– Profit P will be obtained with prob. 1 - q.

If P_{opt} is not known, guess it. I.e. choose target $P = q2^R/(3logn)$ where R is uniformly selected from [1,logd].

– 2^R will approximate P_{opt} with probability 1/logd

– Profit $P_{opt}/(clogn)$ will be obtained with prob. 1/(clogd).

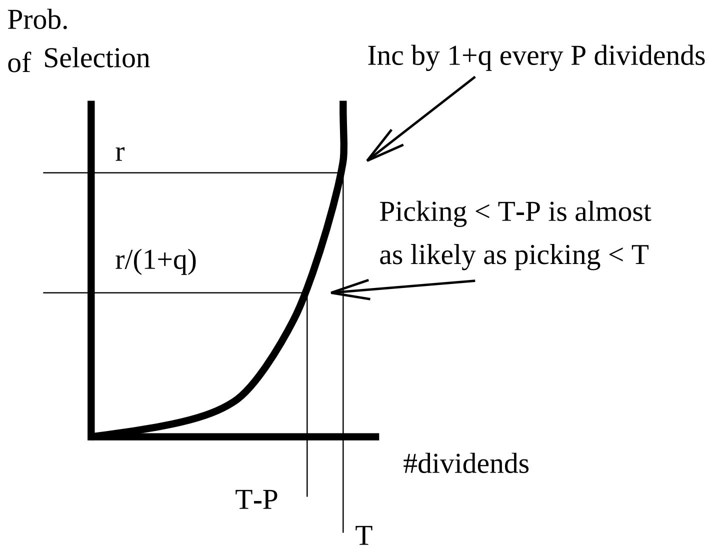

Figure 1: Approximate growth of the selection probability with time

– c is a fixed constant.

Even if the past is not correlated with the future, the past can be combined with randomness to provide a very good prediction of future performance.

The intuition that past winners are likely to be future winners can be used as the basis for a succesful strategy, even though the intuition is not precisely correct.