The Boxes Experiment : Learning an Optimal Monitoring Interval

Sameen S. Fatima

Computer Science Technical Report 92-38

Experimental Knowledge Systems Laboratory
Department of Computer Science
University of Massachusetts
Amherst, Massachusetts 01003

Abstract

Monitoring is a kind of resource management, in which the resources are the events that bring about changes in the environment, and the goal is to minimize the costs incurred. In this project we were interested in performing an experiment to check whether humans would converge on optimal monitoring frequencies. The Boxes Experiment was designed and implemented for this very purpose. Data was collected and analysed, and statistically it was found that human subjects performed suboptimally in the monitoring task. We also attempted to provide reasons for this behavior.

This research is supported by DARPA under contracts #F49620-89-C-00113 and #F30602-91-C-0076 and by AFOSR under the Intelligent Real-time Problem Solving Initiative, contract #AFOSR-91-0067. The US Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation hereon.

1 Introduction

1.1 Aim of the project

1. To design and implement an experiment to test whether human subjects learn to monitor at the optimal rate for an event that occurs probabilistically.

2. To analyze results obtained from running the experiment on human subjects and to compare them with the optimal performance predicted by the mathematical model.

1.2 The monitoring problem

Monitoring is a kind of resource management, in which the resources are the events which bring about changes in the environment, and the goal is to minimize the costs incurred. Changes in the environment can be found by periodically checking for an event (polling model). But how frequently should checking be done? Intuitively, if the rate of arrival of events is very low then the frequency of polling should also be low and if the rate of arrival of events is high then the frequency of polling should also be high. On the other hand, if the cost of checking is very high, then one is tempted to check less often. Hence, in order to determine the optimum monitoring frequency, that is, the frequency at which the cost incurred is a minimum, it is necessary to know the probability of the occurence of events and the costs involved. For example in a fire-monitoring task, the event is a fire, and one needs to know how frequently fires are set, what is the cost of monitoring, and what is the cost of damage due to the fire, in order to determine the optimum monitoring frequency. A mathematical model for monitoring frequencies is already available. We were interested in performing an experiment to check whether humans would converge on optimal monitoring frequencies. If so we could say that the model was true in general. If not we could learn from the deviations from optimality about how humans set monitoring frequencies. Before I explain the design of my experiment I will state the underlying assumptions, identify the parameters and describe the model.

1.3 The mathematical model

The assumption is that events happening in the environment are stationary, independent bernoulli processes. Stationary means that the probability of the occurence of events on any given trial is the same and independent means each event is independent of anything that happened on a previous trial. Further, only one event can happen at any time instant, never more than one. The parameters are:
p = the probability that the event will happen in any given time unit
H = the cost of monitoring the event (monitoring is done once per r time units)
F = the cost of the damage caused by the event per time unit
r = the monitoring frequency i.e., the number of time units that elapse between monitoring.

For example in the fire-monitoring task, the event is a fire, and monitoring involves sending up a helicopter. Then, p is the probability that a fire starts in any given time unit, F is the damage done by the fire in each time unit, and H is the cost of sending up the helicopter. Only one new fire can start at any given instant. Mathematically the question now is, for given values of p, H, and F, what value of r minimizes total cost?

An event could happen at time r-1 with a probability p and continue for 1 time unit, an event could happen at time r-2 with a probability p and continue for 2 time units, etc., until an event that happens at the beginning with a probability p would continue for r time units. The total cost incurred during a monitoring interval r, is the sum of H and the cost of damage caused by all the events that happened in interval r. Hence,
Expected number of time units for which the events have been causing damage is

$$\sum_{i=1}^{r} pi = p\sum_{i=1}^{r} i = \frac{pr(r+1)}{2}$$

Expected cost of damage caused by the events during interval, r is $F\frac{pr(r+1)}{2}$

Expected cost during the monitoring interval, r is $H + F\frac{pr(r+1)}{2}$

Expected cost per unit time is C = $\frac{H+F\frac{pr(r+1)}{2}}{r}$

In order to find the value of r, which minimizes C, we differentiate C with respect to r, and equate it to 0, and solve for r.

$$\frac{\partial C}{\partial r} = \frac{Fp}{2} - \frac{H}{r^2} = 0$$

$$r_{opt} = \sqrt{\frac{2H}{Fp}}$$

2 Description of the boxes experiment

The boxes experiment was designed to test whether human subjects learn to monitor at the optimal rate as predicted by the above model. In this experiment, 1200 boxes appear one after another on the screen. Each box is analogous to a time unit. The boxes contain dots, with each box containing at least as many dots as in the box immediately before it, and possibly one more, depending on chance. So the number of dots gradually increases from box to box. The way to control this is by monitoring. The subject is repeatedly asked to choose a monitoring interval, which is the same as the number of boxes. For every monitoring interval that the subject chooses, the next box in the sequence starts from zero dots again. So monitoring is a way of preventing too many dots from accumulating in the boxes.

There is a cost for monitoring, as well as a cost for the dots. The subject has to weigh the cost of monitoring against the cost of the dots in order to determine how often to monitor. The subject is asked to choose a monitoring interval in the range of 5 to 60, with the aim to bring the score as low possible, where the score for the entire experiment is:

(total number of dots)*(cost per dot) + (total times monitored)*(cost for monitoring)

Each time the subject monitors, we show a display. The display consists of two parts.

1. In the top part of the display appear some boxes, one for each second that has elapsed since the last monitoring. Each box might contain a colored dot. If a box contains a dot, then every box after it contains this same dot. The subject is charged for the number of dots that appear on the screen, plus a constant charge for monitoring.

2. At the bottom of the display is the message area, which displays the monitoring interval, the dots cost, the monitoring cost, the total cost (which is the total number of dots plus the cost of monitoring), as well as the average cost per second for the current monitoring interval, and the cumulative average cost per second since the beginning of the session. The subject adjusts the monitoring interval to keep these average costs as low as possible.

Below is a diagramatic representation of the experiment along with the formulae for the costs in general terms:

F = the cost of a dot
D_i = total number of dots on the screen in interval i
H = Monitoring or uncovering cost for interval i
T_i = Time interval for the ith click

```
A 1200 second session:

 Interval1    Interval2    Interval3    Interval4
o----------o-----------o----------o-------------o----.........----o
```

Dot cost for interval i = FD_i
Total cost for interval i = $FD_i + H$
Average cost per second for interval i = $A_i = \frac{FD_i+H}{T_i}$
Average cumulative cost per second since the beginning of the session = $C_i = \frac{\sum_{i=1}^{r}(FD_i+H)}{\sum_{i=1}^{r} T_i}$

This experiment has bearing on the fire-monitoring task. An analogy can be drawn between the scenario in this experiment and fires starting and being detected in the fire-monitoring task. Analogous to each box is a time unit and analogous to each dot is a fire. As time progresses fires get added, analogous to dots increasing from box to box. In order to prevent damage caused by the fire, periodic monitoring is done, analogous to monitoring to prevent dots from accumulating. Monitoring is done with a view to minimize costs, given that F is the cost of damage done by a fire, and H is the cost of monitoring by sending up a helicopter. This is the same scenario presented to the subjects in the boxes experiment.

3 Results of the experiment

3.1 Input parameters

The experiment was conducted with the following input parameters:

1. Set1 : H = 100, F = 1, p = 0.3
2. Set2 : H = 200, F = 1, p = 0.3

The session consists of 1200 boxes, so that the subject had enough time to learn the optimal monitoring rate. The optimal monitoring rate, r_{opt}, and the optimal cost, C_{opt}, in each case were:

1. Set1 : $r_{opt} = 26$, $C_{opt} = 7.89$
2. Set2 : $r_{opt} = 37$, $C_{opt} = 11.105$

In order to see the variation in the expected cost per second with respect to monitoring interval for each of the above sets of parameters, we plotted expected cost per second vs. monitoring interval, r (Figure 1), using the formula

$$\text{Expected cost per second} = \frac{H + Fp\frac{r^2+r}{2}}{r}$$

From the plots the following relationships were observed:

Table 1:

Set1		Set2	
Range of monitoring interval, r	Range of expected cost per sec	Range of monitoring interval, r	Range of expected cost per sec
16-43	7.89 - 9	25-54	11.105 - 12
18-38	7.89 - 8.5	28-47	11.105 - 11.5
22-30	7.89 - 8		

Note from line 1 of the above table, that for nearly the same range of r, the range of expected cost per second for set 1 is greater than for set 2. This means that set 2 is more sensitive to changes in monitoring interval as compared to set 1.

3.2 Optimum statistical learner (OSL)

An optimal statistical learner is one who knows the model, but nothing more about the input parameters than the human subject knows. The value of p is estimated by the optimal statistical learner by using the formula

$$p_{est} = \frac{Number-of-new-dots-since-the-session-began}{Time-elapsed-since-the-session-began}$$

For every monitoring the optimal statistical learner estimates the value of p, and based on that calculates the value of the optimal monitoring frequency,

$$r_{osl} = \sqrt{\frac{2H}{Fp_{est}}}$$

The optimal statistical learner program was written in order to compare the results obtained from the subjects with the optimally expected results.

3.3 Data collection

The boxes experiment was carried out on 30 subjects and the results were tabulated. A sample table for a single subject is shown below:

Table 2:

Click no	Monitoring interval	Cost of dots	Cost of monitoring	Total cost	Average cost	Cumulative average cost
1	30	166	100	266	8.87	8.87
2	15	48	100	148	9.87	9.20
.	.	.	.	.	.	.
.	.	.	.	.	.	.

In order to counterbalance the input parameter settings, half of my subjects got set1 first, and the other half got set2 first. Data was also collected for 30 optimal statistical learners. The number 30 was chosen in order to make statistical comparisons of data obtained from subjects vs. data obtained from the optimal statistical learner. The cumulative average cost was calculated for the entire 1200 box session for each of the 30 subjects and each of the 30 optimal statistical learners. The distributions of the cumulative average costs for the subjects (Figure 2) and the OSL (Figure 3) were plotted, and the mean cumulative average cost, $\bar{x}$, and the standard deviation, σ were calculated in each case.

Table 3:

Set Number	$\bar{x}_{subjects}$	$\sigma_{subjects}$	$\bar{x}_{osl}$	σ_{osl}
1	8.657	0.939	7.8	0.174
2	11.996	1.678	10.961	0.290

From the above table the question to be answered is "How does the performance of the subjects compare with the performance of the OSL?"

3.4 Data analysis

The one-tailed z-test was performed to find out how significant the above differences in performance were between the subjects and the OSL. The z-score was calculated using the formula:

$$z = \frac{\overline{x}_{subjects} - \overline{x}_{osl}}{\sqrt{\frac{\sigma^2_{subjects}}{n_{subjects}} + \frac{\sigma^2_{osl}}{n_{osl}}}}$$

- Set 1 : $z_{subjects-osl} = 4.92 > 2.33$
- Set 2 : $z_{subjects-osl} = 3.33 > 2.33$

Both the above z scores are greater than 2.33, and hence we can say with greater than 99% confidence that $\overline{x}_{subjects}$ is statistically different from $\overline{x}_{osl}$. This means that the subjects do not monitor optimally.

4 Why do subjects monitor suboptimally?

In the earlier section we saw that the subjects do not monitor optimally. But it's hard to find an answer to why they do not. It could be for any of the following reasons:

1. Subject may not have a well thought monitoring strategy at first, and may develop one during the experiment. By the time the subject developed a strategy the game would have ended. Probably if the subjects were allowed to practice, strategies could be learned and the performance might improve.

2. Subjects may not have understood the underlying assumptions, that is the events happening in the environment are stationary, independent bernoulli processes. For example, few of my subjects had asked me if the probability of dots getting added was changing with time.

3. The strategy used (learned) may be sub-optimal (i.e., it may rely on heuristics that are sub-optimal)

4. The strategy may be confused by noise. Noise is due to the probabilistic nature of the events. Heuristics that work without noise may not work with noise. This is because noise may reinforce incorrect relationships between moitoring interval and the average cost per second.

By talking to the subjects after the experiment, I found there were several different strategies they used. Two of these strategies were implemented to see if they performed in a similar way as the subjects.

1. Strategy 1 :
The idea behind this strategy was to find an interval which would make the dots cost equal to the monitoring cost. This can also be worked out from the model by making approximations. Events happening in the model are the same as the dots appearing in

the boxes experiment. Therefore,

Expected cost of dots getting added in interval r, is $F\frac{pr(r+1)}{2}$

For large r, r+1 $\approx$ r, and
Expected cost of dots getting added in interval r, is $\frac{Fpr^2}{2}$

For optimal performance, r = $r_{opt} = \sqrt{\frac{2H}{Fp}}$. Hence
Minimum expected dots cost = $\frac{Fpr_{opt}^2}{2} = \frac{Fp}{2}.\frac{2H}{Fp}$ = H

This shows that for large r, the minimum expected dots cost is equal to the monitoring cost.

The program for simulating this strategy, increases the monitoring interval if the dots cost is lesser than the monitoring cost, and decreases the monitoring interval if the dots cost is greater than the monitoring cost, until it finds a monitoring interval such that the dots cost is nearly equal to the monitoring cost.

2. Strategy 2 :
   The basic idea is hill-climbing with gradient descent. The aim is to find a cumulative cost scale with three points on it such that the middle one has the lowest cost associated with it. This would form a v-shaped scale. The program moves the 3 points about, trying to find the v-shaped costs. If the weights on the scale make a slope, roll down the slope. It first finds a v-shaped scale with the points 10 units aparts, then 5, 2, and finally 1 unit apart.

Again a program was written to generate 30 data sets using each of the above strategies. This was done in order to compare their performance with the subjects and with the optimal statistical learner. The distributions of the cumulative average costs for strategy 1 (Figure 4)and strategy 2 (Figure 5) were plotted , and the mean cumulative average cost, $\overline{x}$, and the standard deviation, σ were calculated in each case.

Table 4:

Set Number	$\overline{x}_{strategy1}$	$\sigma_{strategy1}$	$\overline{x}_{strategy2}$	$\sigma_{strategy2}$
1	8.005	0.190	8.285	0.358
2	11.2334	0.218	11.299	0.360

Based on the above statistics the following z-scores were calculated:

1. Strategy 1:

   - Set 1: $z_{subject-strategy1} = 3.73$, $z_{osl-strategy1} = 4.36$
   - Set 2: $z_{subject-strategy1} = 2.47$, $z_{osl-strategy1} = 4.14$

2. Strategy 2:

    - Set 1: $z_{subject-strategy2} = 2.03$, $z_{osl-strategy2} = 6.68$
    - Set 2: $z_{subject-strategy2} = 2.22$, $z_{osl-strategy2} = 4.00$

The above z-scores are greater than 2.33 in most cases and are greater than 1.65 in some. Hence we can say with greater than 99% and 95% confidence in each of these two cases respectively that $\overline{x}_{s1}$ and $\overline{x}_{s2}$ are statistically different from $\overline{x}_{subjects}$ and from the $\overline{x}_{osl}$. From this we see that there isn't really one strategy that is being used. there are several different strategies. Further to see how the two strategies compare with each other, the following z-scores were computed:

- Set 1: $z_{strategy1-strategy2} = 3.78$
- Set 2: $z_{strategy1-strategy2} = 0.84$

The z-score for set 1 is greater than 2.33 and hence we can say with greater than 99% confidence that strategy 1 is statistically different from strategy 2. Whereas for set 2, we cannot say that there is a statistical difference between the two strategies.

5 Conclusion

In this project the boxes experiment was carried out on 30 human subjects and 30 optimum statistical learners, in order to be able to make statistical comparisons of their performances on the monitoring task. From the data collected the cumulative average costs were computed for each of the 30 human subjects and each of the 30 OSLs. It was found that the mean of the cumulative average costs of the 30 subjects, $\overline{x}_{subjects}$ was greater than the mean of the cumulative average costs of the 30 OSLs, $\overline{x}_{osl}$. The one-tailed z-test was performed, and it was found that with greater than 99% confidence, $\overline{x}_{subjects}$ was statistically different from $\overline{x}_{osl}$. Based on this, we reached the conclusion that human subjects performed suboptimally in the monitoring task.

We suggested four possible reasons for this, and did further experiments to investigate one of these possibilities, that humans do not monitor optimally because they rely on strategies that are not optimal. We programmed two of the strategies that subjects claimed to use and found that although they performed quite well, in fact better than the subjects on average, they still performed sub-optimally. This provides some evidence for the conclusion that humans monitor sub-optimally because the strategies they rely on, consist of heuristics that give good, but sub-optimal, performance. To support the other two reasons, additional studies can be carried out to see the effect of practice and the effect of noise.

It would also be interesting to decide what level of performance can be considered to represent deviation from the model. This is because humans may follow the model, but differ from optimal performance due to factors like fatigue, cognitive distraction, and learning. Further work can be done to take these factors into account.

Acknowledgements

Special thanks to my advisor Prof. Paul R. Cohen for introducing me to the subject of experimental methods in AI, and for his valuable suggestions and encouragement all through this project. I thank Eric Hansen for the discussions that brought my project to a completion, and to David Westbrook for his help in developing the graphics for the experiment. I would also thank all my subjects, specially my lab mates who patiently gave their time. I express my gratitude to Prof. Carole Beal for her suggestions as my second reader.

References

1. Cohen, P. R. 1992. Experimental Methods in AI. Manuscript in preparation.
2. Cohen, P. R., Greenberg M. L., Hart, D. M., Howe, A. E. 1989. Trial by Fire : Understanding the Design Requirements for Agents in Complex Environments. AI Magazine, 10(3), Fall 1989.
3. Gupta, S. C., Kapoor, V. K. Fundamentals of Mathematical Statistics. Sultan-Chand and Sons Publishers.
4. Phillips, John L Jr. How to Think about Statistics. W.H. Freeman and Company.

FIG 1: EXPECTED COST PER SECOND VERSUS MONITORING INTERVAL

Set 1 : F = 1, H = 100, p = 0.3

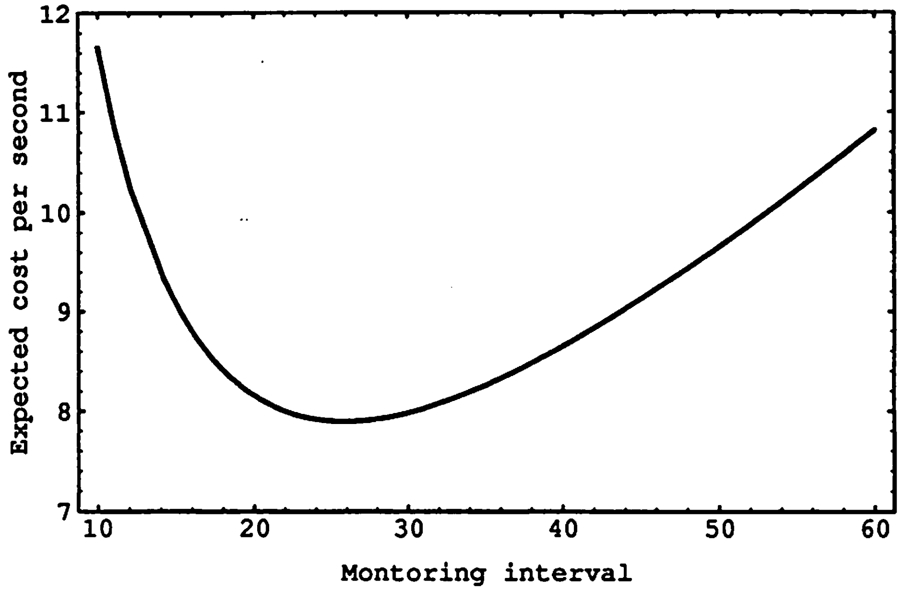

Set 2 : F = 1, H = 200, p = 0.3

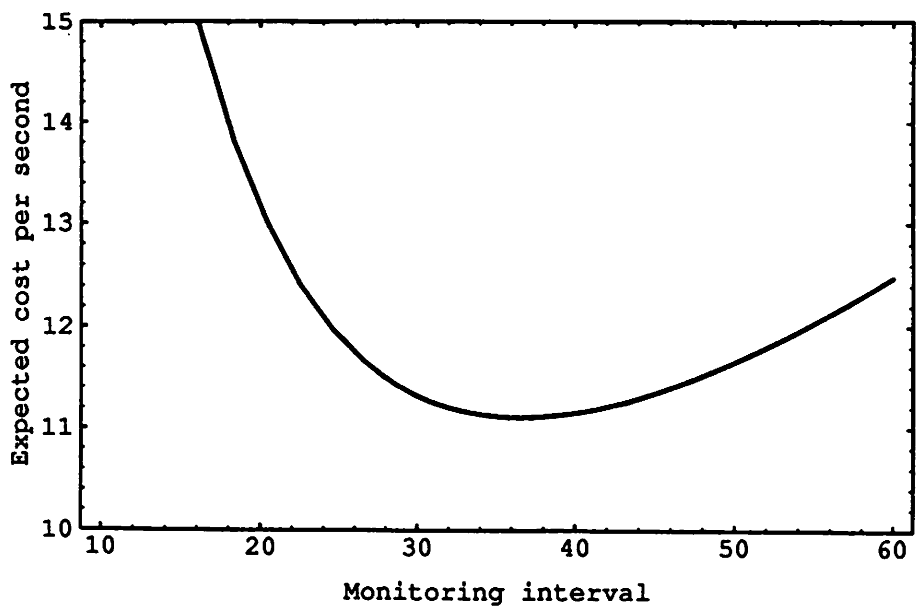

FIG 2: HISTOGRAMS FOR THE SUBJECTS

Set 1 : F = 1, H = 100, p = 0.3
$\bar{x}$ = 8.657, σ = 0.939

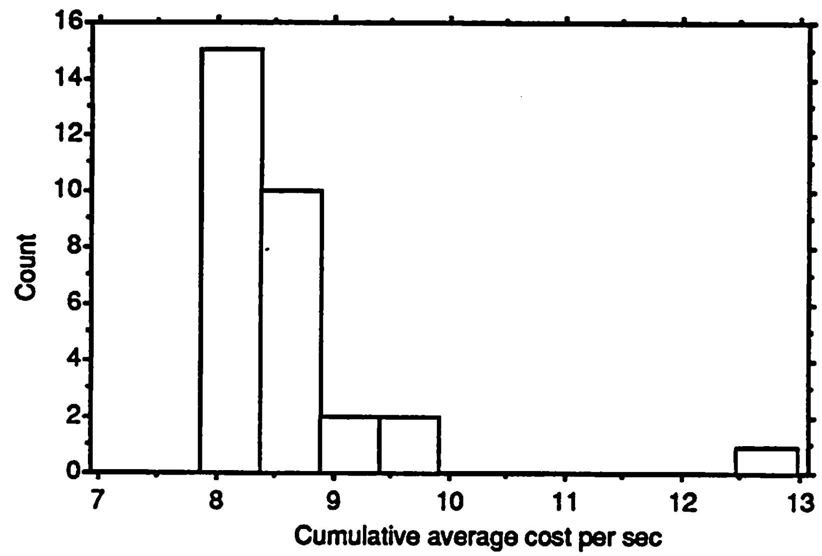

Set 2 : F = 1, H = 200, p = 0.3
$\bar{x}$ = 11.996, σ = 1.678

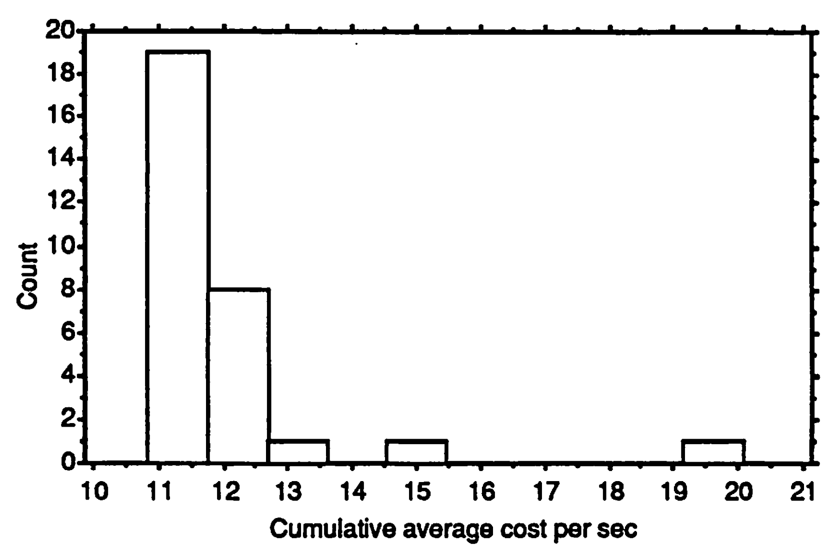

FIG 3: HISTOGRAMS FOR THE OSL

Set 1 : F = 1, H = 100, p = 0.3
$\bar{x} = 7.8$, $\sigma = 0.174$

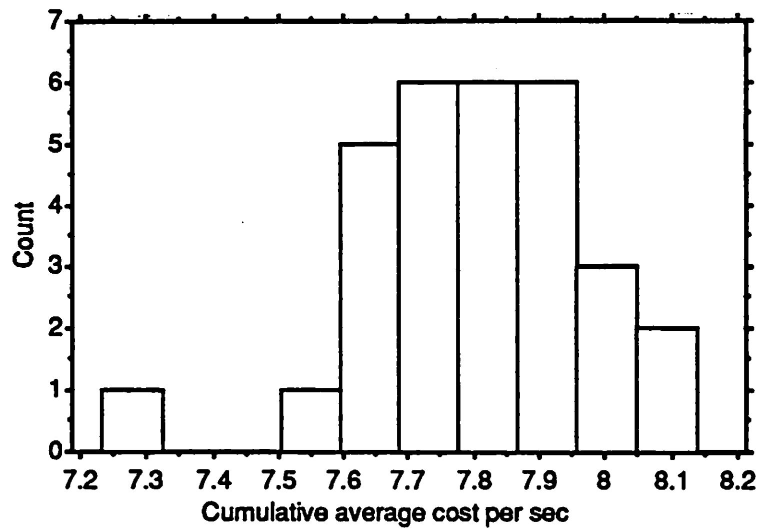

Set 2 : F = 1, H = 200, p = 0.3
$\bar{x} = 10.961$, $\sigma = 0.290$

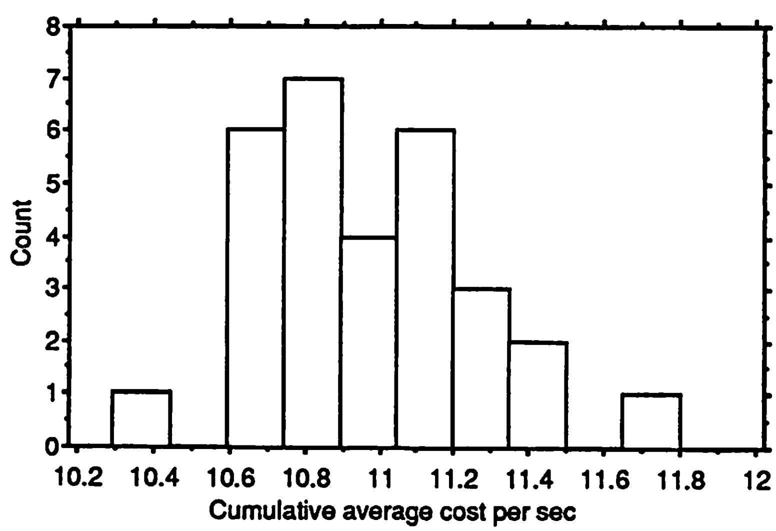

FIG 4: HISTOGRAMS FOR STRATEGY 1

Set 1 : F = 1, H = 100, p = 0.3
$\bar{x} = 8.005$, $\sigma = 0.190$

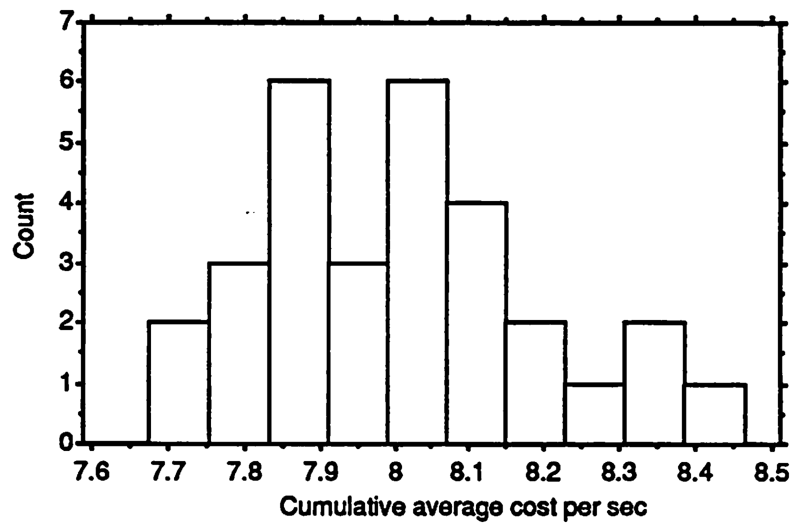

Set 2 : F = 1, H = 200, p = 0.3
$\bar{x} = 11.234$, $\sigma = 0.218$

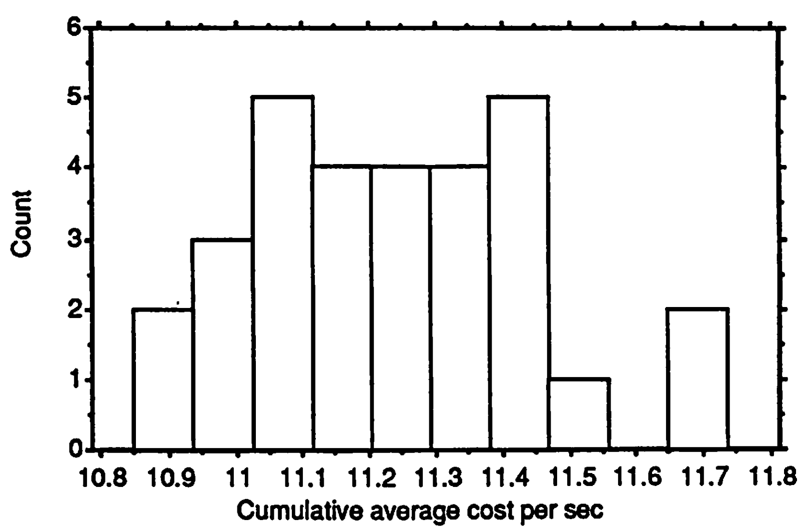

FIG 5: HISTOGRAMS FOR STRATEGY 2

Set 1 : F = 1, H = 100, p = 0.3
$\bar{x} = 8.285$, $\sigma = 0.358$

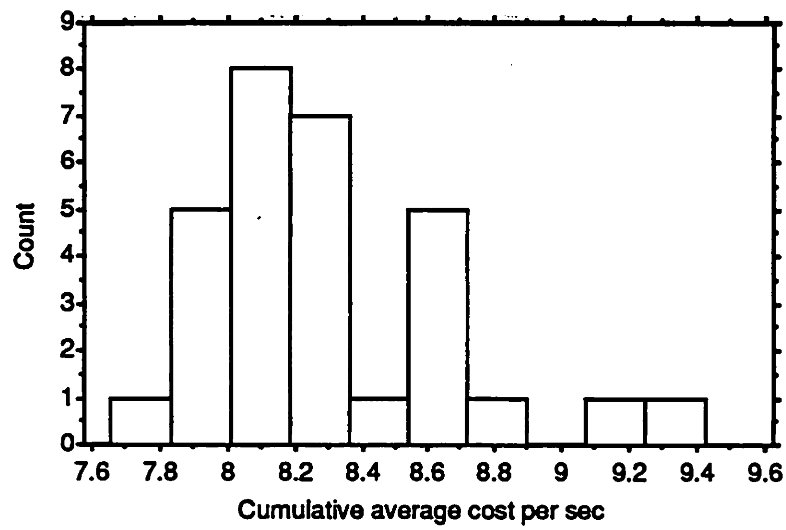

Set 2 : F = 1, H = 200, p = 0.3
$\bar{x} = 11.299$, $\sigma = 0.360$