
Budgeted Optimization with Concurrent Stochastic-Duration Experiments

Javad Azimi, Alan Fern, Xiaoli Z. Fern
 School of EECS, Oregon State University
 {azimi, afern, xfern}@eeecs.oregonstate.edu

1 Proof

Proposition 1. *For any number of experiments n and labs l , let \mathcal{S}_N be the set of corresponding N stage schedules, where N must be at least $\lceil n/l \rceil$. For any $S \in \mathcal{S}_N$, $CPE(S) = \max_{S' \in \mathcal{S}_N} CPE(S')$ if and only if S is uniform.*

Proof. First, we rewrite the CPE objective as follows.

$$CPE(S) = \sum_{i=2}^N n_i \sum_{j=1}^{i-1} n_j = \frac{1}{2} \sum_{i=1}^N \sum_{j \neq i} n_i n_j$$

From this it is clear that all uniform schedules achieve the same CPE value since the multi-set of n_i values in any uniform schedule are identical. It remains to show that if a schedule S is not uniform, then it is not optimal. Since S is not uniform we know that there exists two stages i and j with $|n_i - n_j| > 1$. Without loss of generality, we assume that $n_i < n_j$. Consider the summation of terms in the above expression that involve either n_i and n_j , denoting this sum as $f(S, i, j)$, we have:

$$f(S, i, j) = n_i \sum_{l \neq i, j} n_l + n_j \sum_{l \neq i, j} n_l + n_i n_j$$

Now consider a new schedule S' with $n_i' = n_i + 1$ and $n_j' = n_j - 1$, but otherwise identical to S , we have:

$$\begin{aligned} f(S', i, j) &= (n_i + 1) \sum_{l \neq i, j} n_l + (n_j - 1) \sum_{l \neq i, j} n_l + (n_i + 1)(n_j - 1) \\ &= n_i \sum_{l \neq i, j} n_l + \sum_{l \neq i, j} n_l + n_j \sum_{l \neq i, j} n_l - \sum_{l \neq i, j} n_l + n_i n_j + n_j - n_i - 1 \\ &= f(S, i, j) + n_j - n_i - 1 \\ &> f(S, i, j) \end{aligned}$$

where the last inequality follows from our assumption that $|n_i - n_j| > 1$. This shows that $CPE(S') > CPE(S)$ and thus S is not optimal in terms of CPE . \square

Lemma 1. For any duration distribution P_d that is log-concave, if an N stage schedule $S = \langle (n_i, d_i) \rangle$ is p -safe, then there is a p -safe N stage schedule $S' = \langle (n_i, d'_i) \rangle$ such that if $n_i = n_j$ then $d'_i = d'_j$.

Proof. Consider an N stage schedule $S = \langle (n_i, d_i) \rangle$. Based on the n_i values, we group the N stages into subsets that contain stages with the same n_i values, denoted by $\Pi_j, j = 1, \dots, k$, where k is the total number of distinct n_i values in S . Due to the IID nature of the stages, we can consider each group separately. Focusing on a specific group Π_j , the probability that an execution of the stages in Π_j is safe is

$$\prod_{i \in \Pi_j} P_d(d_i)^{n_j}$$

Taking the log of the probability, we have:

$$n_j \sum_{i \in \Pi_j} f(d_i)$$

where $f(d_i) = \log P_d(d_i)$. Since P_d is log-concave, f is a concave function. Based on Jensen's inequality, we have:

$$\begin{aligned} \frac{1}{|\Pi_j|} \sum_{i \in \Pi_j} f(d_i) &\leq f\left(\frac{1}{|\Pi_j|} \sum_{i \in \Pi_j} d_i\right) = f(\bar{d}_j) \Rightarrow \\ \sum_{i \in \Pi_j} f(d_i) &\leq |\Pi_j| f(\bar{d}_j) \Rightarrow \\ n_j \sum_{i \in \Pi_j} \log P_d(d_i) &\leq n_j |\Pi_j| \log P_d(\bar{d}_j) \Rightarrow \\ \prod_{i \in \Pi_j} P_d(d_i)^{n_j} &\leq P_d(\bar{d}_j)^{n_j |\Pi_j|} = \prod_{i \in \Pi_j} P_d(\bar{d}_j)^{n_j} \end{aligned}$$

This shows that the probability of a safe execution is at least as large if we replace all d_i 's within group j with their average \bar{d}_j . Since this can be done for all groups, this completes our proof. \square

Theorem 2. Let $\Pi(s, t)$ be a policy generator and $\bar{\pi}$ be the switching policy computed with ϵ -accurate estimates. For any state s , stages-to-go t , and base policy π , $C_t^{\bar{\pi}}(s, \pi) \geq \max_{\pi' \in \Pi(s, t) \cup \{\pi\}} C_t^{\pi'}(s) - 2t\epsilon$.

Proof. We use induction on t for which the base case of $t = 1$ is easily verified. For the inductive case, let $T(s, d)$ be a random variable that is distributed over next states given that a decision d is made in state s . It is easily verified that $C_{t+1}^{\pi}(s) = E[C_t^{\pi}(T(s, \pi(s, t)))]$ for any base policy π , s , and t . Also, we have that $C_{t+1}^{\bar{\pi}}(s, \pi) = E[C_t^{\bar{\pi}}(T(s, \pi^*(s, t)))]$, where π^* is the base policy selected at $t + 1$, which returns the decision $d^* = \pi^*(s, t + 1)$.

$$\begin{aligned} C_{t+1}^{\bar{\pi}}(s, \pi) &= E[C_t^{\bar{\pi}}(T(s, d^*), \pi^*)] \\ &\geq E\left[\max_{\pi' \in \Pi(s, t) \cup \{\pi^*\}} C_t^{\pi'}(T(s, d^*)) - 2t\epsilon\right] \\ &\geq \max_{\pi' \in \Pi(s, t) \cup \{\pi^*\}} E[C_t^{\pi'}(T(s, d^*))] - 2t\epsilon \\ &\geq E[C_t^{\pi^*}(T(s, d^*))] - 2t\epsilon \\ &= C_{t+1}^{\pi^*}(s) - 2t\epsilon \\ &\geq \max_{\pi' \in \Pi(s, t) \cup \{\pi\}} C_{t+1}^{\pi'}(s) - 2(t+1)\epsilon \end{aligned}$$

The first inequality follows from the inductive hypothesis and the estimation error bound. The third and fourth steps follow from the definition of expectation and maximization. Finally, the last step follows from our assumption of ϵ -accurate loss estimates. \square

2 Results

In this section we present the results of the proposed policies with their corresponding variances. In general, each number in Tables 1, 2, 3 represents the average regret of each policy \pm its variance for different horizons. It can be seen that the variance is very small for all of the proposed policies.

Table 1: Horizon=4

Function	$h = \infty$	OnFCP	OfStaged	OfIL	OnMEL	PS
Cosines	.142 \pm .00	.339 \pm .04	.181 \pm .01	.195 \pm .01	.275 \pm .03	.205 \pm .01
FuelCell	.160 \pm .01	.240 \pm .01	.182 \pm .01	.191 \pm .01	.258 \pm .01	.206 \pm .01
Hydro	.025 \pm .00	.115 \pm .01	.069 \pm .00	.070 \pm .00	.123 \pm .02	.059 \pm .00
Rosen	.008 \pm .00	.013 \pm .00	.010 \pm .00	.009 \pm .00	.013 \pm .00	.008 \pm .00
Hart(3)	.037 \pm .00	.095 \pm .00	.070 \pm .00	.069 \pm .00	.096 \pm .00	.067 \pm .00
Michal	.465 \pm .01	.545 \pm .01	.509 \pm .01	.508 \pm .01	.525 \pm .01	.502 \pm .01
Shekel	.427 \pm .03	.660 \pm .04	.630 \pm .03	.648 \pm .03	.688 \pm .03	.623 \pm .03
Hart(6)	.265 \pm .01	.348 \pm .01	.338 \pm .01	.340 \pm .01	.354 \pm .01	.347 \pm .01
CPE	190	55	100	100	66	100

Table 2: Horizon=5

Function	$h = \infty$	OnFCP	OfStaged	OfIL	OnMEL	PS
Cosines	.142 \pm .00	.339 \pm .04	.181 \pm .01	.194 \pm .01	.274 \pm .04	.150 \pm .01
FuelCell	.160 \pm .01	.240 \pm .01	.167 \pm .01	.190 \pm .01	.239 \pm .01	.185 \pm .01
Hydro	.025 \pm .00	.115 \pm .01	.071 \pm .01	.069 \pm .01	.086 \pm .01	.042 \pm .00
Rosen	.008 \pm .00	.013 \pm .00	.009 \pm .00	.008 \pm .00	.011 \pm .00	.008 \pm .00
Hart(3)	.037 \pm .00	.095 \pm .00	.055 \pm .00	.064 \pm .00	.081 \pm .01	.045 \pm .00
Michal	.465 \pm .01	.545 \pm .01	.500 \pm .01	.510 \pm .01	.521 \pm .01	.494 \pm .01
Shekel	.427 \pm .04	.660 \pm .04	.635 \pm .04	.645 \pm .03	.682 \pm .03	.540 \pm .04
Hart(6)	.265 \pm .01	.348 \pm .01	.334 \pm .01	.330 \pm .01	.333 \pm .01	.297 \pm .01
CPE	190	55	100	100	91	118

Table 3: Horizon=6

Function	$h = \infty$	OnFCP	OfStaged	OfIL	OnMEL	PS
Cosines	.142 \pm .00	.339 \pm .04	.167 \pm .01	.147 \pm .00	.270 \pm .03	.156 \pm .00
FuelCell	.160 \pm .01	.240 \pm .01	.154 \pm .01	.163 \pm .01	.230 \pm .01	.153 \pm .01
Hydro	.025 \pm .00	.115 \pm .01	.036 \pm .00	.035 \pm .00	.064 \pm .00	.025 \pm .00
Rosen	.008 \pm .00	.013 \pm .00	.007 \pm .00	.009 \pm .00	.010 \pm .00	.009 \pm .00
Hart(3)	.037 \pm .00	.095 \pm .00	.045 \pm .00	.050 \pm .00	.070 \pm .00	.038 \pm .00
Michal	.465 \pm .01	.545 \pm .01	.477 \pm .01	.460 \pm .01	.502 \pm .02	.480 \pm .01
Shekel	.427 \pm .03	.660 \pm .03	.530 \pm .04	.564 \pm .03	.576 \pm .05	.510 \pm .03
Hart(6)	.265 \pm .01	.348 \pm .01	.304 \pm .01	.266 \pm .01	.301 \pm .01	.262 \pm .01
CPE	190	55	133	137	120	138