CS630 Homework 2

Qilin Ye

February 11, 2025

Solution to problem 1. If we define X_i to be the processing time of job i, then $X_i = 1 + (k-1)Y_i$ where Y_i is a Bernoulli random variable with parameter $1-p$. Fix a random processor. With abuse of notation let $X_1, \cdots, X_{n/m}$ be its scheduled jobs. The runtime T can be written as

$$T = \sum_{i=1}^{n/m} X_i = \frac{n}{m} + (k-1)\sum_{i=1}^{n/m} Y_i.$$

To make computations easier, we use Hoeffding's inequality instead on $Z = \sum Y_i$, which gives

$$\mathbb{P}(T - \mathbb{E}T \geqslant (k-1)\epsilon) = \mathbb{P}(Z - \mathbb{E}Z \geqslant \epsilon) \leqslant \exp\left(-\frac{2\epsilon^2}{n/m}\right)$$

and union bound implies

$$\mathbb{P}(\text{latest completion time} \geqslant \mathbb{E}T + (k-1)\epsilon) \leqslant m\exp\left(-\frac{2\epsilon^2}{n/m}\right).$$

Likewise for the lower tail bound. We want above to happen with low probability of form for example $1/n$. Solving $m\exp(-2m\epsilon^2/n) \leqslant 1/n$ gives $\epsilon = \mathcal{O}(\sqrt{n/m \cdot \log n})$. Clearly, $\mathbb{E}T = n/m \cdot (p + k(1-p))$. Therefore, we conclude that with high probability, the completion time is bounded by

$$\left[\frac{n}{m}(p + k(1-p)) - \mathcal{O}(\sqrt{n/m \cdot \log n}),\ \frac{n}{m}(p + k(1-p)) + \mathcal{O}(\sqrt{n/m \cdot \log n})\right].$$

Solution to problem 2. Yannan Bai tossed me some links[1][2] which turned out to be massively helpful, so I cite them here. Consider the following send-receive $(n+2)$-bit pairs

$$(X)\ 00\ \underbrace{(00\cdots0)}_{n/2 \text{ bits}} \mapsto \underbrace{(00\cdots0)}_{n/2 \text{ bits}}\ 01\ (X). \tag{*}$$

where $X \in \{0,1\}^{n/2}$ with k ones. We claim that by routing all such pairs, the edge

$$\tau = \underbrace{(00\cdots00)}_{n/2 \text{ bits}}\ 00\ \underbrace{(00\cdots00)}_{n/2 \text{ bits}} \to \underbrace{(00\cdots00)}_{n/2 \text{ bits}}\ 01\ \underbrace{(00\cdots00)}_{n/2 \text{ bits}}$$

will except exponential congestion. The analysis will be performed w.r.t. n, but the offset of 2 bits should not affect the outcome.

Fix k to be some value to be determined later. Any such send-receive pair must be routed across $2k$ edges, k of which turn ones into zeros, and the remaining k vice versa. Thus, a randomly permuted bit-fixing scheme traverses through τ if and only if it fixes *all* k ones in the front-half X before flipping any zeroes to ones in the second half of the packet. There are a total of $\binom{n/2}{k}$ send-receive pairs of form (*), and each of them independently crosses τ with probability $\binom{2k}{k}^{-1}$. Using the well-known two-sided Sterling approximation

$$\left(\frac{n}{k}\right)^k \leqslant \binom{n}{k} \leqslant \left(\frac{ne}{k}\right)^k$$

we obtain

$$\mathbb{E}X := \mathbb{E}[\#\text{ of pairs crossing }\tau] = \binom{n/2}{k}\binom{2k}{k}^{-1} \geqslant \left(\frac{n/2}{k}\right)^k \left(\frac{2ke}{k}\right)^{-k} = \left(\frac{n}{4ke}\right)^k.$$

Taking logarithm gives $k(\log n - \log(4ke))$ which implies a local maximum is obtained at $n/4e^2$. We set our k to be at this value, and it follows that

$$\mathbb{E}X \geqslant \left(\frac{n}{4e\cdot n/(4e^2)}\right)^k = e^k = \exp(n/4e^2) > 2^{n/4e^2}.$$

Finally, all send-receive pairs of form (*) that cross τ must simultaneously reach τ after k flips, so the expected congestion at τ is at least $\mathbb{E}X$. Proof complete!

Solution to problem 3. Consider the event that none of the first $k = c\sqrt{n}$ balls collide, which evidently has probability $\prod_{j=0}^{k-1}\left(1-\frac{j}{n}\right)$. Heuristically, this quantity roughly equals $\prod_{j=0}^{k-1}\exp(-j/n) = \exp(-k^2/2n) = \exp(-c^2/2)$. To make use of this identity:

- If we let $c_1 = \sqrt{2}$, then $\sum_{j=0}^{c_1\sqrt{n}} j \approx n$ [the exact value may be off by $o(1)$ which is negligible as n gets large], and so

$$\mathbb{P}(\text{no collision with } k = c_1\sqrt{n}) = \prod_{j=0}^{k-1}\left(1-\frac{j}{n}\right) \leqslant \exp\left(-\sum_{j=0}^{\sqrt{2n}}\frac{j}{n}\right) \sim \exp(-2n/2n) = \frac{1}{e}.$$

- On the other hand, if we let $c_2 = \sqrt{2\log 2}$, then

$$\begin{aligned}\mathbb{P}(\text{no collision with } k = c_2\sqrt{n}) &= \prod_{j=0}^{k-1}\left(1-\frac{j}{n}\right)\\ &\geqslant \exp\left(-\sum_{j=0}^{k-1}\frac{j}{n} - \sum_{j=0}^{k-1}\frac{j^2}{n^2}\right)\\ &\sim \exp\left(-\log 2 - \mathcal{O}(n^{-1/2})\right) \sim \frac{1}{2}.\end{aligned}$$

Solution to problem 4. We use the well-known fact that for a geometric random variable X with parameter p, its MGF is given by

$$\mathbb{E}(e^{tX}) = \frac{pe^t}{1-(1-p)e^t} \qquad \text{for } t < -\log(1-p).$$

In our problem, it translates to $\mathbb{E}(e^{tX_i}) = e^t/(2-e^t)$ for $t \in (0, \log 2)$. The rest follows from the standard proof:

each $\mathbb{E}X_i = 2$ and $\mathbb{E}X = 2n$, and for $0 < t < \log 2$ we have

$$\begin{aligned}\mathbb{P}(X \geqslant (1+\delta)2n) &= \mathbb{P}(e^{tX} \geqslant e^{t(1+\delta)(2n)}) \\ &\leqslant \frac{\mathbb{E}(e^{tX})}{e^{t(1+\delta)(2n)}} = \left(\frac{e^t}{2-e^t}\right)^n \cdot \exp\left(-t(1+\delta)(2n)\right) \\ &= \exp\left(n\log\left(\frac{e^t}{2-e^t}\right) - 2tn(1+\delta)\right) \\ &= \exp\left(n[t - \log(2-e^t) - 2t(1+\delta)]\right) = \exp(n[-\log(2-e^t) - t(1+2\delta)])\end{aligned}$$

To minimize the last expression w.r.t. t, it suffices to minimize the exponent $f(t) = -\log(2-e^t) - t(1+2\delta)$. Since $f''(t) = 2e^t/(2-e^t)^2 \geqslant 0$ is concave on $(0, \log 2)$, and so solving $f'(t) = e^t/(2-e^t) - (1+2\delta)$ gives the optimum. This yields $t^* = \log((1+2\delta)/(1+\delta))$. The rest is algebra.

$$\begin{aligned}\mathbb{P}(X \geqslant (1+\delta)2n) &\leqslant \exp\left(n[-\log(2-e^{t^*}) - t^*(1+2\delta)]\right) \\ &= \exp\left(n\left[-\log\left(2 - \frac{1+2\delta}{1+\delta}\right) - (1+2\delta)\log\left(\frac{1+2\delta}{1+\delta}\right)\right]\right) \\ &= \exp\left(-n\left[-\log(1+\delta) + (1+2\delta)\log\left(\frac{1+2\delta}{1+\delta}\right)\right]\right).\end{aligned}$$

Solution to problem 5. We assume $\sigma < \infty$. Let $S_n = \sum_{i=1}^n X_i$ which has mean $n\mu$ and variance $n\sigma^2$. Therefore, S_n/n has mean μ and variance σ^2/n. Therefore,

$$\mathbb{P}(|S_n/n - \mu| > \epsilon) \leqslant \frac{\mathrm{var}(S_n/n)}{\epsilon^2} = \frac{\sigma^2}{n\epsilon^2}$$

which, as $n \to \infty$, tends to 0. Hence $S_n/n \to \mu$ in probability.

Solution to problem 6.

(1) Let $x^*(i)$ be the optimal fractional solution; we round them to $x(i)$ via $\mathbb{P}(x(i) = 1) = x^*(i)$ and 0 otherwise. Immediately we have $\mathbb{E}\sum_i x(i) = \sum_i \mathbb{E}x(i) = \sum_i \mathbb{P}(x(i) = 1) = \sum_i x^*(i)$ so the post-rounding objective equals the optimum in expectation.

(2) Each S_i containing p has probability $x^*(i)$ of being chosen, so the desired quantity is $\prod_{i:p\in S_i}(1 - x^*(i))$.

(3) Using $e^{-x} \geqslant 1 - x$ we bound (2) by

$$\prod_{i:p\in S_i}(1 - x^*(i)) \leqslant \exp\left(-\sum_{i:p\in S_i} x^*(i)\right) \leqslant \frac{1}{e}$$

where the second $\leqslant$ is because by assumption x^* is feasible, and so $\sum_{i:p\in S_i} x^*(i) \geqslant 1$.

(4) For each element p, the probability that it is not covered by any individual C_i is $\leqslant 1/e$ by (3). The probability that p is not covered by all $C_1, \cdots, C_{2\log n}$ is $e^{-2\log n} = n^{-2}$. By union bound, some element remains uncovered after taking the union of all C_i is n^{-1}, i.e., the probability of outputting a valid set cover is $\geqslant 1 - 1/n$.

(5) By (1) each C_i is in expectation equally large as OPT, so the final output is $\mathcal{O}(\log n)$-approximate.