

USQCD Jeopardy Policy for Institutional Clusters

(Dated: June 3, 2019)

I. PREAMBLE

USQCD purchases access to a certain number of nodes for a year at a time on the institutional clusters at BNL and Fermilab.¹ If a fiscal year ends with total usage below the purchased computing power, it is simple to conclude that our user community has wasted money and does not need as much computing as claimed. The point of this policy is to prevent such an outcome, by rewarding projects that consume allocations above pace and penalizing those that lag.

In principal, there can be episodes in which USQCD as a whole uses a smaller or greater number of nodes. In practice, however, low usage at the beginning of a fiscal year (or allocation year) is difficult to recoup.

II. NOTATION

Let \mathcal{Q} be the set of all projects running on the institutional clusters. It can be decomposed into

$$\mathcal{Q} = \mathcal{P} \cup \mathcal{R} \cup \mathcal{O}, \quad (2.1)$$

where

$$\mathcal{P} = \text{set of projects being } \textit{penalized}, \quad (2.2)$$

$$\mathcal{R} = \text{set of projects being } \textit{rewarded}, \quad (2.3)$$

$$\mathcal{O} = \text{other projects.} \quad (2.4)$$

In the policy spelled out in this document, these sets are disjoint. The policy defines the set \mathcal{O} to consist of projects with less than 1% of the total institutional cluster resource; such projects will be neither penalized nor rewarded for lackluster or exuberant running.

Let $A_q(m)$ be the allocation of project $q \in \mathcal{Q}$ remaining at the end of month m ; further, let $m + \epsilon$ denote the first day of the next month. The rules will specify $A_q(m + \epsilon)$ based on $A_q(m)$ and various criteria. The months are counted from the beginning of an allocation year: July \mapsto 1, August \mapsto 2, September \mapsto 3, \dots , June \mapsto 12.

III. PENALTIES

The jeopardy starts with a formula also in use at JLab:

$$A_p(m + \epsilon) = \min\left(\frac{1}{12}(14 - m)A_p(0), A_p(m)\right) \quad (3.1)$$

with the set \mathcal{P} for month m defined to be the set of projects for which the first choice is smaller. This formula imposes no penalty during the first two 2 months: $\frac{1}{12}(14 - m)A_p(0) \geq A_p(0) \geq A_p(m)$ for $m \leq 2$.

¹ Jefferson Lab still operates under the dedicated-cluster model and has its own jeopardy policy, based on the same basic principals.

IV. REWARDS

The time taken from projects in \mathcal{P} has to be allocated to other projects; otherwise it won't be used. We also set aside a certain number of node-years

$$R_0 = \frac{1}{10} \sum_{q \in \mathcal{Q}} A_q(0) = 10\% \times (\text{total resource}) \quad (4.1)$$

to be distributed during the first four months. Thus, the amount to redistribute is

$$r(m) = r_0(m) + \sum_{p \in \mathcal{P}} (A_p(m) - A_p(m + \epsilon)), \quad (4.2)$$

where

$$r_0(m) = \begin{cases} \frac{1}{4} R_0, & 1 \leq m \leq 4, \\ 0, & \text{otherwise.} \end{cases} \quad (4.3)$$

Then, for $r \in \mathcal{R}$,

$$A_r(m + \epsilon) = A_r(m) + r(m) \frac{A_r(m - 1) - A_r(m)}{\sum_r [A_r(m - 1) - A_r(m)]}, \quad (4.4)$$

that is the reward is distributed in proportion to last month's running.

The set \mathcal{R} consists of every project that is not in \mathcal{P} and that ran during the past month at 110% or more of the annualized allocation, i.e.,

$$A_r(m - 1) - A_r(m) \geq \frac{11}{120} A'_r(m - 1), \quad (4.5)$$

where the annualized allocation is defined to be

$$A'_r(m - 1) = \frac{12}{m - 1} A_r(m - 1 + \epsilon). \quad (4.6)$$

Defining $r(m)$ and \mathcal{R} in this way restores the allocation that the SPC might have awarded, had it not set aside R_0 for the reward pool, to those who can demonstrate they can use the resource from day one. In other words, if everyone runs this way, the reward pool is simply returned evenly to everyone.