Jeopardy Policies for USQCD

(Dated: February 1, 2019)

I. NOTATION

Let $\mathcal{Q}$ be the set of all projects running on the cluster in question. It can be decomposed into

$$\mathcal{Q} = \mathcal{P} \cup \mathcal{R} \cup \mathcal{O}, \tag{1.1}$$

where

$$\mathcal{P} = \text{set of projects being } \textit{penalized}, \tag{1.2}$$

$$\mathcal{R} = \text{set of projects being } \textit{rewarded}, \tag{1.3}$$

$$\mathcal{O} = \text{other projects.} \tag{1.4}$$

The sets need not be disjoint, but in the proposals below they are. It is necessary to define $\mathcal{P}$ and $\mathcal{R}$ explicitly; then $\mathcal{O}$ is simply shorthand for the complement $\overline{\mathcal{P} \cup \mathcal{R}}$. We can also put "small" projects, with less than, say, 12M core-hours, into $\mathcal{O}$.

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.

II. PENALTIES

The jeopardy committee proposes to adopt a formula in use at JLab:

$$A_p(m + \epsilon) = \min\left(\tfrac{1}{12}(14 - m)A_p(0), A_p(m)\right) \tag{2.1}$$

with the set $\mathcal{P}$ defined to be the set of projects for which the first choice is smaller. This formula gives a grace period of 2 months: $\frac{1}{12}(14 - m)A_p(0) \geq A_p(0) \geq A_p(m)$ for $m \leq 2$.

III. REWARDS

The time taken from projects in $\mathcal{P}$ has to be allocated to other projects; otherwise it won't be used. We also can set aside a certain number of node-years R_0 to be distributed throughout the year. It is sensible to divide this into ten portions, to be distribution at the end of the first through tenth months (for the second through eleventh). Thus, the amount to redistribute is

$$r(m) = r_0(m) + \sum_{p \in \mathcal{P}} A_p(m) - A_p(m + \epsilon), \tag{3.1}$$

where

$$r_0(m) = \begin{cases} \frac{1}{4}R_0, & 1 \leq m \leq 4 \\ 0, & \text{otherwise} \end{cases} \tag{3.2}$$

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

$$A_r(m+\epsilon) = A_r(m) + r(m)\frac{A_r(m-1) - A_r(m)}{\sum_r \left[A_r(m-1) - A_r(m)\right]}, \tag{3.3}$$

that is the reward is distributed in proportion to last month's running. Some suggest replacing $m-1$ with 0 in the above formula, i.e., base the reward on accumulated running since the beginning of the year.

We propose that $\mathcal{R}$ consists of every project not in $\mathcal{P}$ and that ran 110% or more of the annualized allocation, i.e.,

$$A_r(m-1) - A_r(m) \geq \tfrac{11}{120} A'_r(m-1), \tag{3.4}$$

where the annualized allocation is defined to be

$$A'_r(m-1) = \frac{12}{m-1} A_r(m-1+\epsilon). \tag{3.5}$$

Defining $\mathcal{R}$ in this way restores the allocation that the SPC might have awarded, had it not set aside 10% of the total computing resource 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 to everyone.