Dear Colleagues,

Thank you for the very careful reading of the proposal, and for your questions. First, please note that the time requested on the "19g" nodes can indeed be implemented on the "21g" nodes, that is the 19g-resources requested in Table 3 can be allocated to the 21g nodes; we have not had the time to do a detailed costing, but the exchange rate of 1 Mi100 GPU-hour = 130/21.6 RTX GPU-hours is probably reasonable. Our answers to the remaining questions are as follows.

> *Which deliverables would you prioritize in case the allocation is reduced?*

The calculation of the disconnected contributions to hadron structure can only be performed on CPUs, and therefore the $8,014K$ KNL hours, 2 times that listed in Table 2 (up to a typo!) shown is the highest priority for KNL resources. This then requires that we exploit the RTX/Mi100 resources of Table 3 in order to compute the needed perambulators.

Regarding the second project, namely that of isovector PDFs and GPDs, the calculation of the generalized perambulators on our E2 and E3 ensembles detailed in Table 4 are our highest priority; the ensemble E4 we believe to be of too large a volume to implement on the Mi100, and cannot be implemented on non-error-correcting GPUs such as the RTX nodes.

> *Can one comment on the cost/gain ratio for the use of the distillation method?*

A detailed analysis of the cost/benefit of the distillation framework vs a more usual sequential-source approach is difficult since the most computationally demanding part of our computation, namely that of the so-called *Generalized Perambulators* corresponding to an operator insertion in a three-point function are "multi-use". The same generalized perambulators can in general be used for different source/sink particles such as the pion and the nucleon, and indeed for different momenta of the particles at source and sink.

For the case of the nucleon charges, we did attempt to compare the number of inversions using the distillation with vanilla Jacobi smearing in Egerer *et al*, Phys.Rev.D 99 (2019) 3, 034506, but emphasised that comparison was for a single operator at a single momentum and therefore did not take account of the multi-use feature of the distillation framework.

Perhaps the most important feature to emphasise is that we believe the distillation framework enabled us to do calculations that we would not otherwise have attempted. Thus in the recent calculation of the helicity distribution of the gluon, our most statistically demanding calculation, the use of distillation allowed us both to use the sGEV method with a non-trival basis of operators, and moreover sample the ensembles more completely by performing a time-slice sum at both source and sink nucleon, rather than only at the sink as would be the case in the sequential-source approach; it was still necessary to compute correlators from every time slice to get reasonable statistics.

> *What is the range of t and ξ values expected to get good signal within the obtained statistics?*

In the pseudo-PDF approach to the calculation of the GPDs, the GPDs are computed for discrete values of the momentum transfer $-t$, and of the skewness ξ; the range of t depends on the value of the skewness. For the ensemble E1 of Table 1, we believe the range of momentum transfers $-t$ will be to around 1.2 Gev^2, for around five values of the skewness between 0 and 1. The extent to which we get a good signal must await the completion of our analysis.