1. Can you provide an FTE breakdown of the Collaboration effort that is available for Operations (data taking, calibrations, reconstruction, processing etc.)? for Analysis?

Breakdown of several different tasks, and estimate of how much of collaboration time is spent on those tasks:

Data Taking: Howard+Nur (2.0) + Shifters (1.4) + Detector Experts (1.2)   4.6
Calibrations: (MINERvA plus MINOS both)   2
Processing 0.7FTE   0.7
Keepup Processing: 0.45 FTE   0.45
MINOS Data Quality and Validation: 0.2FTE   0.2
Software Releases: 0.2FTE   0.2
Algorithm Development   2
Common Tools   2
Generator Work   2
Flux Studies   1
Testbeam Analyses (3 students, 0.5FTE)   1.5
Analysis Subgroup Coordination: 6 groups, 0.2FTE per group   1.2
Paper Review (2 days per committee member per paper) 8 days * 8 papers 0.26
Collaboration-wide Meeting (20 people 2 hours per week)   1
Week-long Collaboration-wide plus 2 3-day meetings per year   2.2
Collaboration Management (spokes)   (1.5)

  Remaining work to be done on Individual Analyses   18.2

We conclude that slightly less than 50% of the collaboration effort is available for work on analysis of data. The spokespersons "guess" would have generated a similar number, based on what we see people working on around the collaboration.

2. Can you provide an estimate of the FTE available from the Collaboration as a function of time (FY17 and beyond)? When do you anticipate you might experience a significant shortage of person-power?

We only do a census once a year, so based on the past experience and the information we have about new groups, we expect to stay at the current level. We just had two new groups join that expects to devote 1FTE of students, 1FTE of postdocs, ramping up over the next few years. Total PhD's: 2,8,6,7 for years 2013,2014,2015,2016.

While we could arguably add more data analysis output if we enhanced the collaboration, we don't see either any trend in the next several years of a reduction or enhancement in personnel at either the student or postdoc level. There is some trend of senior collaborators spending more time on future projects, mostly DUNE, but this cross-pollination is also important for MINERvA's mission.

We do expect that once we stop operating the detector clearly people will go elsewhere to give their students hardware and operations experience. We have already had 1 institution join with "limited authorship", where they provide students to do analyses, they sign only the papers they contributed to directly, and then they have a reduced operations burden.   That has resulted in a postdoc and student expending roughly half effort on MINERvA, mostly on data analysis.

3. Please comment on your plan to finish the LE publications.

Here is a list of the LE publications we envisage broken down by status
- Publications in internal review, students have graduated
    a. NC K+ production, including application to water Cereknov backgrounds to $p \to K^{+}\upsilon$
    b. Anti-neutrino/neutrino total CC cross-section ratio
    c. Anti-neutrino double differential CCQE cross-section
    d. Neutrino CCQE muon+proton ratios on nuclear targets
- Publications in preparation or nearly so (if student led, students are preparing theses)
    a. CC coherent pion production
    b. Anti-neutrino "low recoil" nuclear effects (uses neutron tagging, under development)
    c. CC π^0 production
    d. Neutrino double- differential CCQE cross-section
- Publications being led by PhDs (postdocs, junior faculty) in advanced stages
    a. Antineutrino CC π^- production
    b. "Low recoil" nuclear effects on Pb compared to CH
- Analyses still being explored or in early stages
    a. Antineutrino muon+neutron CCQE [relies on new neutron tagging technique]
    b. Antineutrino CC exclusive π^+ production [very high background, rare process probing nuclear effects]
    c. Transverse variables for exclusive nuclear CCQE muon+proton [relies on sample whose inclusive cross-sections were just approved]

The overall picture is that most LE analyses are in advanced stages, with the exception of a few where new techniques are being developed. If there continue to be new ideas for analyses or techniques where the low energy data is the right data set, right now because either because they need lower neutrino energies or because they require antineutrinos, there is no reason to discourage this.

Our recent publication record has demonstrated that we can get analyses out the door once the infrastructure for the analyses is completed.

4. Please comment on which analyses are highest priority wrt to the oscillation neutrino community for the ME program, both neutrino and antineutrino running? How many POT do you need for these analyses to become interesting?

A table like this would be helpful: For each high priority analysis topic tell us -
  A. Precision at which it becomes interesting
  B. #POT req'd to achieve a)
  C. when you estimate you'll achieve b)
  D. Ultimate precision goal
  E. #POT req'd to achieve d)
  F. when you estimate you'll achieve e)
  G. Comments

This is a difficult question to answer in a concise way for this review because we have such a variety of analyses with different physics goals. An initial reference document that starts to address this is the December 2014 document to the PAC from MINERvA. We note that at the time we wrote this, we had not yet developed several of the analyses described below.

Another difficulty is that analyses using the same sample have different observables. For example, our current result on CCQE on the nuclear targets looks at a single differential cross-section ratio in a variable that measures Q^2 of the reaction. However, one of the most interesting things to do with this sample, which we barely have statistics in the ME beam to do, is to look at a variety of cross-section variables correlating muon and proton kinematics because these separate different nuclear effects. In this case, the ultimate precision for the more inclusive measurement would be reached before the more "differential" measurement in event kinematics.

For the oscillation motivated analyses, there are three main types of analyses that are the most important output of the medium energy program.
  1. Analyses where the neutrino energy dependence of some other feature of the ME beam (e.g., we expect the flux to be better known from neutrino-electron scattering constraints) is of interest. Examples of these are CCQE differential cross-sections on CH, "low recoil" 2p2h analyses on CH, etc.
  2. Analyses that measure exclusive processes on the passive nuclear targets, e.g., CCQE muon+proton in the nuclear targets. These targets are of order 10% of the scintillator mass, and so almost all these results are statistics limited in the LE beam.
  3. Analyses of rare reactions, ~1% of the rate or less which are statistics limited in the LE beam, e.g., electron neutrino CCQE or coherent pion production.

Aside from infrastructure issues that we described in the review (increased activity and the flux problem), analyses of type 1 could be carried out today for the neutrino and early for anti-neutrino beams and would be of interest.

Analyses of type 2 are almost all statistics limited in such a way that the physics goal requires full statistics. A non-oscillation goal in the PAC document that illustrates this are the studies of the EMC effect that will be one of the physics legacy measurements of MINERvA. We have some more information about extrapolating the CCQE muon+proton result, which clearly benefits from all the available statistics, especially if we wish to break it down in more differential variables. A higher statistics measurement which might arguably reach diminishing returns with a subset of the ME neutrino is the "low recoil" 2p2h measurement comparing Pb targets and CH.

Analyses of type 3 almost all benefit from increased statistics. A good example is the electron neutrino CCQE measurement. There are models that suggest there may be nuclear dependence that leads to differences in electron neutrino and muon neutrino CCQE cross sections. Those difference show up in either the overall rate at the level of a few percent, or might show up as a larger difference in the lowest Q^2 bins. Either way, we expect the uncertainties in this analysis to reduce with statistics (many of the largest systematic effects are constrained by sidebands). For example, in the ratio of electron to muon neutrino CCQE, the full statistics of the neutrino mode ME run will provide a sample roughly 6 times that of the current LE analysis, which in turn should reduce the uncertainty on the cross-section ratio for Q^2 < 0.2 GeV^2 from ~10% to ~4%.

5. p21 of Debbie's first talk mentions the possibility that you may need to change the way you "slice" events. Please comment on what criteria you'll use to decide whether or not you need to do this. Please comment on what that would imply in terms of changes to your data processing, etc. Would that significantly increase the required FTE to calibrate, process, analyze, and publish the data in this instance?

We showed that our event rate per POT drops as intensity increases. There are two criteria that determine how we proceed with the slicing retuning:

1. How much does the event rate per POT (especially in our lowest-statistics analyses) drop at 700 kW?
2. How well can we simulate that inefficiency?

If the answer to point 1 is "not too much" and to point 2 is "well enough", then we may decide to eat the inefficiency and not pursue improvements in the slicing algorithm.

If that's not the case, then we need to retune. The plan would then be to tune the (currently very conservative) parameters of the slicing algorithm to be as aggressive as possible without splitting apart individual neutrino interactions. This retuning is probably an FTE-month or two. Retuning would be based primarily on MC whose timing distributions have been very well tuned. (The accuracy in the MC timing distribution was a necessary ingredient for the K+ identification in our series of papers.)

We would then have to re-run both data and MC with the new tuning, which is a small amount of FTE to submit the jobs, but a few months of wall clock time.

Analyzers would then have to consider re-tuning cuts, but mostly this should be minimal: if they've already tuned their analysis on lower-intensity data, then aggressively-sliced higher-intensity data should look similar.

6. Can you provide a more detailed description of the plan to migrate to SL6 for Nearline processing? for MINOS data processing?

MINERvA Nearline

- The overall plan is to use current SLF5 based system when beam comes back. Run SLF6 system in parallel, testing components of the workflow
  - (SCD, Both SCD and MINERvA, MINERvA)
- SLF6 Migration:
  - Install SLF6 on new machines (Done)
    - Logger, and 2 worker nodes
  - Configure worker nodes (In progress)
    - Install MINERvA framework
    - Configure to accept Condor jobs (shifter plots)
  - Configure logger (In progress)
    - Install MINERvA framework for GMBrowser (live processing of data)
    - Condor submission routine
  - Direct job output (histograms) to web viewable area
    - Parallel area on our echecklist instances (minerva-exp.fnal.gov, minerva-wbm.fnal.gov)
- Once it is verified new system produces GMBrowser and shift plots reliably will redirect shifters to new areas and GMBrowser instance
  - Keep old system running as backup
- Repurpose old system machines as backups
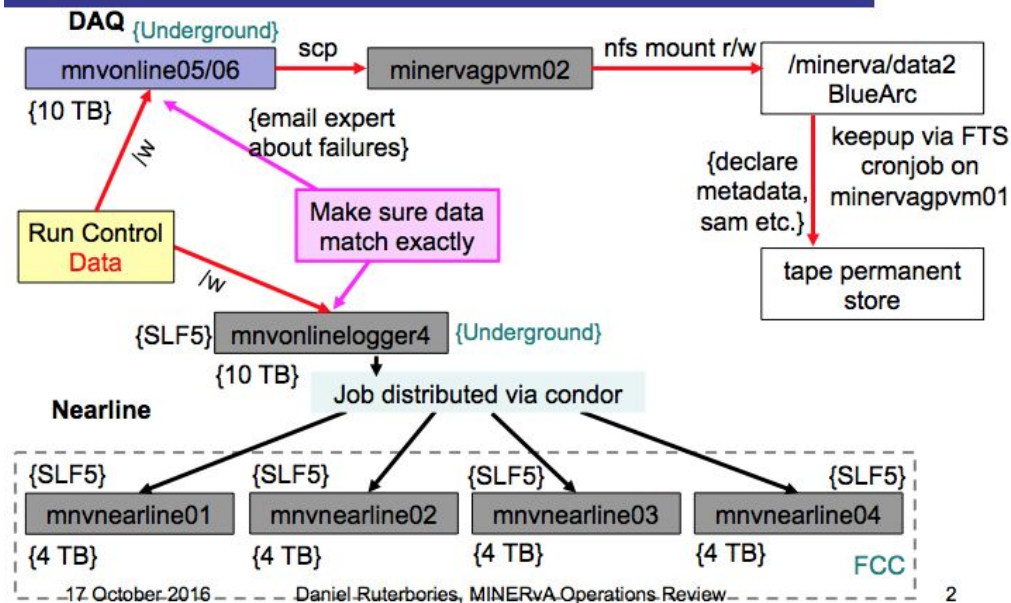- Note: All necessary SCD experts have been consulted and are helping work on this migration

MINOS processing

- MINOS data processing has been tested on SLF6 machines, but not fully migrated
- Keepup list is produced on an SLF5 machine
  - Samproject file, daily list, machine OS independent
- Job submission is done from an SLF6 machine
  - Already had to be done with new DCAFI authentication
- Will migrate Keepup list to slf6 machine after BlueArc dismount tests are completed

MINERvA Personnel

•Operations group: Ozgur Altinok, current nearline processing expert; Barbara Yaeggy, next nearline processing expert; Nur, to oversee and assist

•Offline Computing Experts: Daniel Ruterbories, will help with condor and FTS

Nearline Current Data Flow Structure

DAQ {Underground}

mnvonline05/06 — scp → minervagpvm02 — nfs mount r/w → /minerva/data2 BlueArc

{10 TB}

/w

{email expert about failures}

Run Control **Data**

Make sure data match exactly

/w

{declare metadata, sam etc.}

keepup via FTS cronjob on minervagpvm01

tape permanent store

{SLF5} mnvonlinelogger4 {Underground}

{10 TB}

Job distributed via condor

Nearline

{SLF5} mnvnearline01 {4 TB}

{SLF5} mnvnearline02 {4 TB}

{SLF5} mnvnearline03 {4 TB}

{SLF5} mnvnearline04 {4 TB}

FCC

Future Nearline

DAQ {Underground}

mnvonline05/06

{10 TB}

/w

Run Control **Data**

Make sure data match exactly

/w

{declare metadata, sam etc.}

FTS online copies

dCache + tape permanent store

{SLF6} mnvonlinelogger6 {Underground}

{10 TB}

Job distributed via fifebatch

Nearline

{SLF6} mnvonlineworker1 {10 TB}

{SLF6} mnvonlineworker2 {10 TB}

{SLF6} mnvonlineworker3 {10 TB}

FCC