Choice of frequency (1300 vs. 1207):

Summary opinion; I think it's probably true that there are enough advantages of 1300 MHz over 1207 MHz to make the change in design parameter. I.e., Piggy backing on the TESLA tube development, using their input coupler as a baseline (or at least as a starting point to minimize development time and cost), and copying their cryomodule. However, I have two caveats:

(1) One needs to be sure that there are no significant beam loss problems created by bringing the 1300 MHz cavities down to the 87 MeV point and by having the bigger step in frequency at the transition point. I'd do a few beam simulations to be sure that everything still looks good and there are not sizeable matching problems, beam loss and halo before I bricked up the opening.
(2) The other issue that should be considered is whether the existing 402.5-MHz linac provides a unique ability to test components with a real beam in an actual accelerator environment. I think it might be wise to just list the hardware items that you might consider testing in the existing linac and how much you'd lose by changing frequencies. There might well be a challenge in the future with regard to this issue and it would be good to have a defense in hand.

Detailed comments:

Advantages of using TESLA compatible frequencies

(1) I agree that the fewest different klystron types the better.
(2) I think it makes a lot of sense to take advantage of the TESLA's successful multi-beam klystron R&D development program. However, it's not clear what are the true benefits since you're trading off klystron development, which probably in going down to 1207 from 1300 isn't all that big or costly a deal, vs. cavity development. The costs quoted by JLAB are probably low and are not chiseled in stone. Things will come up at JLAB that will add to the cost burden, whereas, klystron manufacturers give fixed cost commitments.
(3) I also believe that fewer klystrons are better. I think a smaller body count means fewer chances for hardware failures. I really don't believe that there will be more failures with 10 Mw klystrons compared to 5 MW klystrons. In my experience anything above a few hundred kilowatts can, and usually does, have development problems and quirks that have to be ironed out, but once done will operate with equal reliability
(4) I doubt that SNS will ever replace the CCL. They have a lot of other improvement options that would be more effective use of the money. It would take something like a serious reliability problem (chronic water /vacuum leaks or arc downs and frequency detuning in the coupling cells) before replacement would ever get serious attention.

(5) I think it makes a lot of sense to use the TESLA input coupler since it seems to work, and that will result in a significant R&D savings. However, you still will want to have your own test stand for pre-conditioning of input couplers, and since that's the most expensive part of a coupler development program, I'd include an improvement program at some level to see if there are things that could easily done to get a better TESLA unit.

(6) I think using an existing cryomodule design, especially if you have access to the design drawings, would be a big cost and schedule factor for the construction project.

(7) There's no question that a lower tunnel height will save money. However I seem to remember that the structural engineers at SNS did claim that tunnel width was a big cost driver and a foot of tunnel width was more valuable to them than a foot of tunnel height. This could also be a site-specific issue since there was something about composition of the Chestnut Ridge soil that resulted in higher tunnel loading.

(8) WG elements are cheaper at 1300 MHz as long as you don't end up needing pressurized air or SF6 to hold the fields without arcing. The area of concern would be near the klystron and before the power splitters where you have to handle the full power.

(9) I don't think the size of the 805 KHz cavities is the problem, but rather, from what was presented at ASAC in February, the cleanliness of the rinse water system is. It seems that JLAB was using oil-sealed pumps and the QC wasn't always what it should have been. There was a SRF review committee that made a number of recommendations to SNS and JLAB on how to get more consistent performance, but a major recommendation seemed to be to just get someone dedicated to watching over production and making sure that the normal good handling practices are followed. They also recommended changing the oil-sealed pump.

Advantages of the SNS-Compatible 402.5 MHz/ 805 MHz/ 1207 MHz solution

(1) It's true that you would be able to use the SNS cavity designs at 0.61 and 0.81 and probably the RIA cavity design at 0.47.

(2) Some kind of change will be needed at 1300, either another beta section or maybe fewer cells in the early sections.

(3) I'm not sure how big an advantage the actual SNS cost data at 805 MHz is to the frequency decision. First, much of it extrapolates to other frequencies just as well. The modulators, water systems, interlocks and all that stuff are frequency independent. The bulk of the stuff that is in the waveguide system has very good scale factors that can be used over the frequency range of consideration. The cost estimates for the cold part of the linac will not be as valuable because you're going with a whole different cryomodule concept. It could also turn out that the SNS cost estimates could become a double-edged sword for whatever frequency is chosen. Suppose you come up with more cost effective solutions, or lower cost estimates because you expanded the range of manufacturers that could bid on the hardware, you'll find yourself in the

position of overcoming skepticism that will come with the comparison of what will be perceived as “reality”.

(4) I don’t think lower waveguide losses are a very strong argument for basing a decision on frequency. Losses in circulators, ferrite tuners and various mismatched components will probably dominant losses.

(5) The ability to exchange parts with SNS would be an advantage, but isn’t it a wash in the end because if you switch to 1300 MHz, you can trade parts with TESLA?

(6) The ability to test components, especially diagnostic equipment, with a beam in an accelerator environment is a big plus when it comes to commissioning. In my experience, half the problems with commissioning are caused by new hardware with a high rate of infant mortality that reduces beam physics time to next to nothing and prevents continuity in the commissioning program and the other half of the problems are diagnostics that don’t work properly and mislead you when you’re trying to interpret what’s going on with the beam. Setting up the beam properties according to the physics design is almost trivial when the other stuff is working. You’d also be able to test a cavity and maybe even a cryomodule with a beam to see how everything performs as a system.

(7) I don’t think having the documented SNS physic design is all that much of an advantage. Your “SNS-type” design has a different cryomodule design that results in different cavity and focusing element spacing and your cavity and real estate accelerating gradients are higher. Those differences are too big to directly use the SNS stuff. However, the computer codes that they are using are being validated as their commissioning progresses, and I believe that they are valuable for the FERMI design no matter what frequency is chosen.