


I've looked at the power consumption formula within your 3rd critique, and it appears that the factors involved would be beyond our scope at this time. For instance, we would not be able to switch power transitions that I know of, we may or may not be able to do loading capacities, measure voltage supply, and not able to measure clock frequency, unless they're talking about processor speed in MHz/GHz, etc. Those are the factors for the formula. Nonetheless, while we're both not electrical engineers, I know that, at least, I wouldn't be able to set up such an environment within the 1.5 weeks left.

On a different note, I looked further into scheduling, and (correct me if I'm wrong) scheduling with respect to operating systems are for many processes utilizing a resource, such as the CPU, at the same time. It's kind of hard to imagine something like that with migrating processes. Looking back at our documentation, we tried to intend scheduling processes that need to migrate.

Although, our resource is shared memory, I don't know how else to schedule these processes. Perhaps via some kind of priority, which might be too much complexity to worry about this late in the ball game. Additionally, we had mentioned to not worry about scheduling right now.

Finally, I'm able to come up with the following for our measurements:

1. Our chart template is comparing logging block size to time (probably milliseconds). Each logging block size interval can be increments of 5k. So, ideally we would need to be able to easily adjust the max log size.



Of course, for zero there is no logging going on for the same set of processes and that is what we can measure against. Hopefully, while crossing our fingers our times for logging are faster than for not logging. I think the wonderful thing about this is if we can hypothesize that there is an optimal point (i.e. log block size) where beyond that, the time would start increasing then we have shown something significant, actually. In other words I'm imagining the curve/line would be slopping downward, but if the sending node spends a lot of time logging then we may have find the optimal point.

2. Have the chart in 1 above, for the sending node.
3. Have the chart in 1 above, for the receiving node, with respect to both cached and non-cached processes. In other words, at the receiving end we can measure the amount of time the receiving node strips off the processes from the process block and fires them up, etc.
4. Have the chart in 1 above, for the total time in of both 2 and 3.

5. We need to use the same set of processes for each run above, i.e. 2, 3, 4.

Finally, please insert this info into documentation for the web site. Along with all of the other stuff, I think it may be important for Dr. Wang, when he assesses our work.