From: Smith-Tone, Daniel C. (Fed)
To: internal-pqc
Subject: RAM model
Date: Friday, March 11, 2022 10:03:54 AM

Hi, everyone,

I want to come to the defense of the RAM model, since I believe that there is a lot of RAM model bashing going on. I am going to restrict myself in such a way that I may only say things that are true, but I will start by placing my tongue firmly in my cheek.

Here is one method for making something close to the RAM model a realistic model for *some* large memory computations. It only requires a small bit of engineering of new technology, and none of this stuff is as technically challenging as building a quantum computer.

First, we must fully automate the process of building computing and robotic equipment so that the only thing required from humans is the design of a machine. We are not so far from this state right now. This task requires that machines mine the appropriate ores, isolate the mineral components in an unsupervised way, and perform miniature fabrication on their own. We have a great deal of automation in all of these tasks today, we merely need to replace the repetitive tasks humans currently perform with another level of automation.

Second, we must launch some small number of robots to Mercury to build a few of these automation stations. Once there, the process will speed up rapidly. Some robots will build components for orbital power stations while other robots build facilities for constructing more robots.

After disassembling a certain percentage of Mercury, the power stations are towed via a combination of rockets and solar sails to form a Dyson swarm around the sun in the orbit of the somewhat smaller planet Mercury. Now, there is no real limit on power consumption, the limits are heat dissipation and propagation of information at the speed of light.

To overcome heat dissipation, every processing unit on a spacecraft will need to have a (sometimes extremely massive) heat sink radiating heat into space. It is difficult to deal with heat dissipation in space, but to put things in perspective, the heat sink that the ISS uses to dissipate its heat is only about 2 meters long. I haven't computed this yet, but the heat sinks for some of the more critical processing nodes will be LARGER. Still, we can do this because we have all of Mercury to work with.

To overcome the speed of light and three-dimensional space barrier, (this is where the *some* comes in), we distribute the memory across the swarm and have them constantly beaming their memory contents to several processing nodes with special purpose hardware to handle the incoming data stream and relay it to data processing nodes. The reason this can be faster than calling on physical memory is that it depends on the density of information in the data stream, and not on the physical properties of a medium through which the data are being transmitted through wires. This step requires some significant engineering, but is not impossible.

With these steps, memory access can be made to have an essentially constant cost, and it can be fairly small. I believe that the biggest difference this method has to others is that it is essentially agnostic to energy costs, since the energy is free to every piece of equipment.

Of course, this could only work effectively when the algorithm does not need to make a lot of write queries to memory. Otherwise you need to wait until the memory is written, and that might be a while.

Cheers,
Daniel