Dendou Oni Prototyping Hardware Ver. 1.0 Floorplan Note:

There would be at least two flavors: BGA and TSOP for the SDRAM memory. That would eventually be four in total for two separate version RAMs, in two footprint versions – for both Master and Slave processor boards. Therefore, I also elected to put in few headers for peoples wanting more than they would settle on, for few options: Write Disable for EEPROM / Ferroelectric RAM write protect to prevent the firmware ROM image from being overwritten, if desired. And, Oscillator Disable header lets you pick one out of two electrically live oscillators, both in DIP-8 form and 2.5 x 2.0 mm SMD form, to easily change out the frequency quickly and easily. There is a RS232 header for PropClip too.

And, the SMD components were chosen to minimize the real-estate board size, while stuffing lot of features at the same times, even though there would be few DIP-style parts (like that optional blister Quartz oscillator in DIP-8 form), that form alone would allow you to stack as much as you want up to 256 P8X32A processors callable by their IDs (which is also callable by FTDI USB to RS232 converter which would also be included here, along with Vinculum II USB 2.0 host controller) in small space as possible rather than using DIP-40 processors.

Even though unfinished at this time of writing of this sentence, lemme show you the Write Disable pin, Oscillator Disable / Select pins and RS232 header, along with DIP-8 plastic oscillator (yes, there is such few DIP-8 plastic oscillator, and I am even sure they COULD be either Quartz or MEMS, depending on whom you buy from) – under the DIP-8 plastic oscillator, there is a QFN-4 footprint for the tiny oscillator (which would allow you to solder DIP-8 socket atop the oscillator).

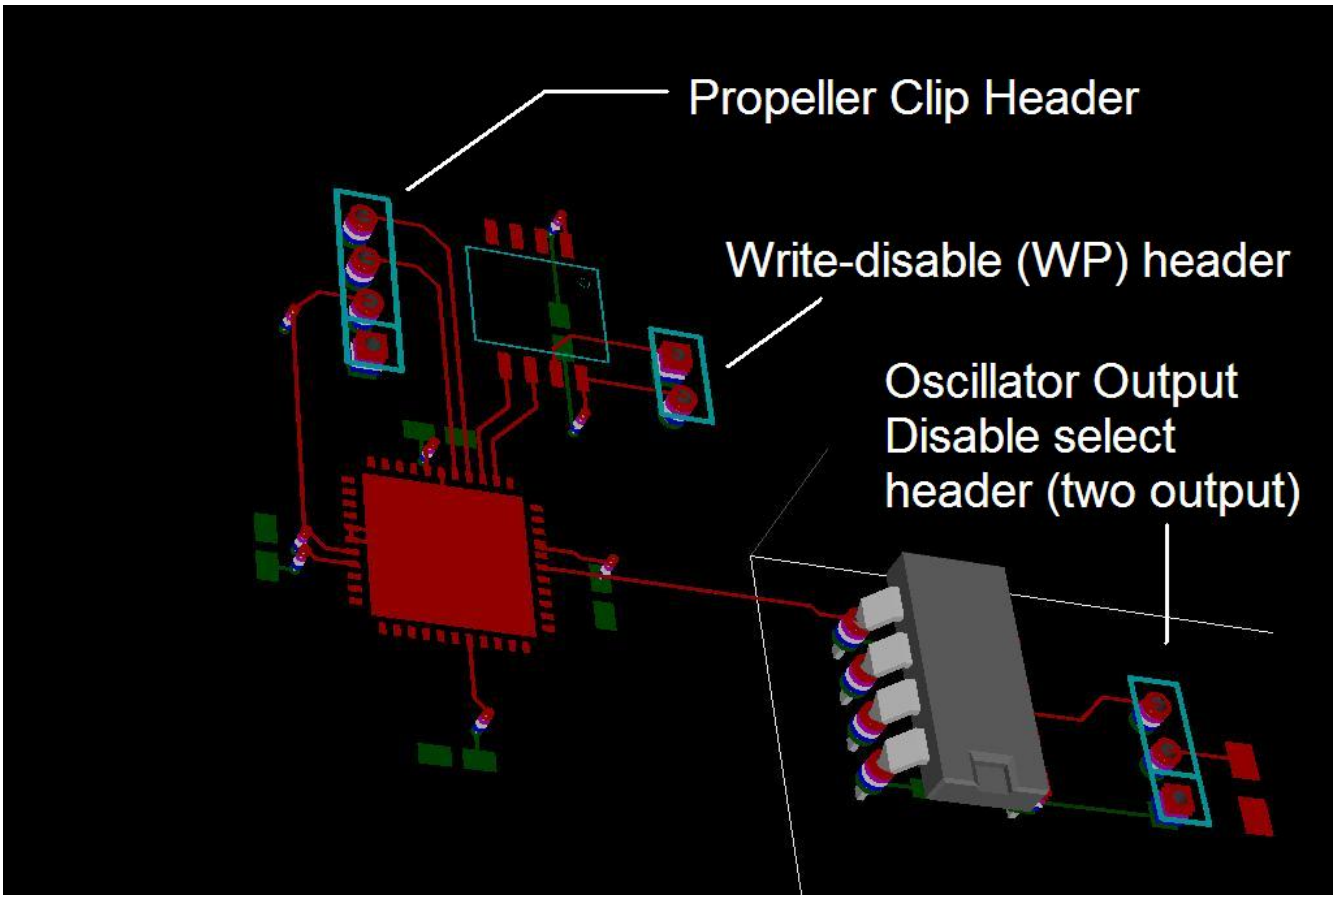

Also, at that time, I have been thinking what kind of power supply option I may go with, regular LDO (Low-Dropout) regulator or switchmode regulator, although I may go with switchmode one as the more boards get connected, the more power the total hardwares start to draw. Consider the currents being drawn:

A Parallax Semiconductors Propeller – 0.3 A.
A Cypress CY62177DV30 4 MB SRAM – 0.015 A.
A typical SDRAM – 0.3 A.
An blister oscillator - ~ 0.004 – 0.05 A (depending on technology and resonator type).
A SiTime SiT8003 MEMS Oscillator (on my board, at least) – 0.003 A.
A typical I2C EEPROM - ~ 0.004 – 0.2 A (depending on I2C speed and size).

Now, to total up all currents being sucked up by a single-processor setup: Maximum current of 1.05 Amps. That's nothing if it's just by itself, a single-processor, but it will just get worse as you expand the whole system – a board that I am designing, the whole bells and whistles there, would draw 2.2 Amps, but that's without a portable hard drive or jumpdrive being plugged in to the Vinculum II chipset which can easily add up to 0.6 A each USB ports. 2 Amps Max. for each boards, and if you add up to 8 boards, you better allow up to 20A as it can easily melt the FDC (Floppy Disk Controller) style cable's wiring insulators. I would add in power headers for the whole boards to allow monstrous currents to be carried among high-voltage current wires, up to 12 volts DC as allowed by switchmode voltage regulator chips on all those boards. What about the signal currents being carried throughout the cables? “It's all digital, matey!” Actually, for eight boards, you would really see up to 450mA on it, no big deal, and they're terminated via 300 ohms resistors mostly as an insurance for sucking up unwanted IO noises, like all that resistors at end of the DIMM on your motherboards, soaking up the end-of-line currents to keep the processor's integrated RAM controller happy (of course, if any of you guys end up getting the PC motherboards with Rambus XDR-DRAM memory bus on it, all four XDIMM slots on it, you won't find any termination resistors at the end of the XDIMM circuitry, as they're all etched on both processor's IMC and the XDR-DRAM chips on the modules and I am thoroughly aware that AMD Athlon X2 and Phenom II CPUs have the termination resistors built-in, with three resistance values: 75, 150 and 300 ohms), and that really serves the purpose here, for the IO pins.

Also, as far as thermal management is concerned, I included the thermal viases onto underneath of the die attach pad on P8X32A-M44 QFN-44 package, it's there as an insurance policy for the extended period of time the whole hardware is to be left operating – that way even when overclocked, it won't overheat. And the thermal viases also carry the current (Vss = GND) to the processor through this point. Also, the bypass capacitors are connected via the thermal mass plane (from the thermal viases of P8X32A to that capacitors) which is also connected to – 3.3V (GND) of the voltage regulator, that way the P8X32A processor's PLL won't get upset. Note that the P8X32A is operated off by the external oscillator(s) as they're pretty much away from P8X32A processor, if it does get warm in operation after few hours or even days, the oscillators won't get affected thermally. I chose to go with external oscillator rather than crystal, it's mostly stability reason – on-die PLLs are very easily affected by heat, even can be exacerbated by crappy filtering on the power supply rail, so I had to put on several capacitors and implement the external oscillators. But, when the tiny oscillator is disabled by pin header, or isn't used, the clock pin can be linked to the clock generator (if more than 1 - 3 feet, RF coaxial cable, like the one used for wireless LAN, should be used: outer mesh to GND and core to CLK, that way you don't mess the clock pulse shape up, or even radiate it into air). Also, note the solder pad for tiny oscillator under the plastic DIP-8 oscillator – you would see why I decided to include Disable / Selection header for two oscillators. They're optional: I made it that way for you guys.

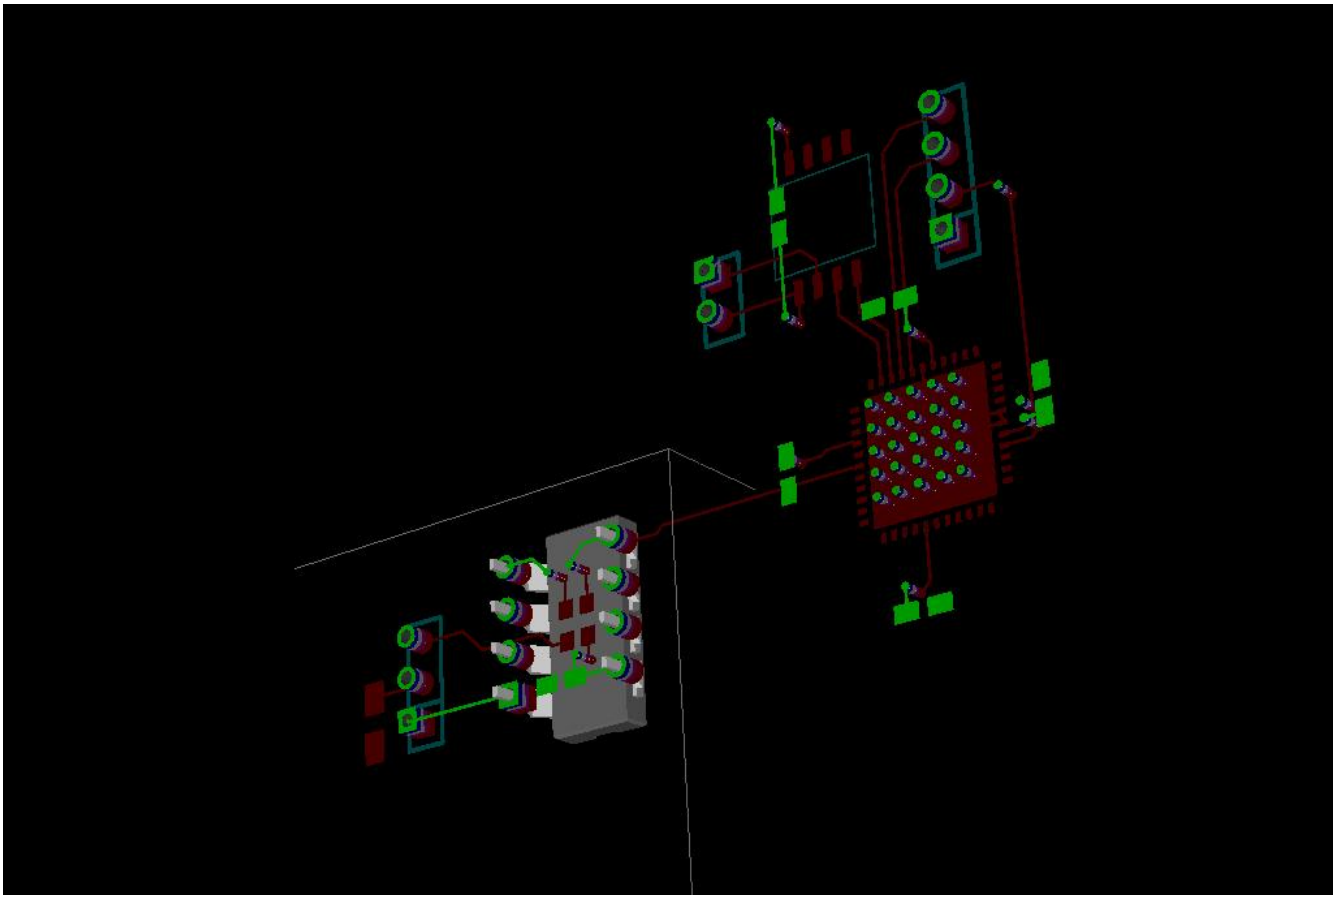

Yet, there's going to be a long way to go before I can finally clean it up and save the whole P8X32A circuitry as “generic” for few favors of boards, that way you get the best on every boards you get under my project (which have been going in and out a lot lately....) which was to give you the open-source supercomputing hardware - for you to mess with. I am going to finish up few favors, based on Parallax Semiconductors P8X32A – then research a bit more on Propeller II and do the same, only on larger board (with eight Prop II – which, funnily enough, reminds me of 68k-operated transputer board I saw out of very old (about 24 years old) DIY computing magazine BYTE).

And, at least as far as FBGA-54 and TSOP-54 SDRAM packages are concerned, I had to check both Micron Technologies, Inc. and ISSI SDRAM whitepapers to make sure the pinouts are the same to avoid disappointing few peoples willing to try out my designs. Fortunately, as supposed, they're of the same pinout so it's pretty standard as well. Also, I may try larger SDRAM on my board to be able to hold very complex data for supercomputing job while being very small, in FBGA-54 package, thus cutting in PCB real-estate. But I won't leave out TSOP version as well, since I am aware that not many peoples are experienced in QFN / BGA soldering, or just like to hold the soldering iron (no pun), so I may redo the board design as soon as I am finished with BGA / QFN version, to hold QFP and TSOP packaged chips, like P8X32A-Q44 TQFP-44 chip, instead of thinner and fancy package I am going to use on mine to keep the whole thing smaller. However, with both board version, it shouldn't matter what the firmware thinks, as it would contain some kind of training codes to take in account with the changes in trace length due to temperature changes and / or types of packages being used. Electrically, they both will be wired up exactly the same, that way the same firmware ROM images can be used on either board versions.