Memory Address Decoding and Fault Analysis

Analyzing memory error logs can help us to identify and fix previously unknown faults in memory. Using Cray's memory error log tool, xthwerrlog, we're able to see when an error occurred and what address it's associated with but we aren't able to dive deeper into DRAM for fault identification. For instance, all the ECC is done in the integrated memory controller (iMC) so there's no direct way to tell if the fault was in the iMC, the DIMM, or the channel between them. So, how do we determine where the error occurred? The address bits contain the location of the row, bank, bank group, etc. of the bad bit(s) when we are looking at a physical address. If those bad bits are consistently the same ones going bad we can say something about the fault being in the same physical location with good probability. For instance, we can statistically determine what sort of fault we are seeing. If we see lots of faults on a DIMM that vary only in the column value, then we can call that a column fault, while that same logic applies to the DRAM row, bank group, and bank. Using a set of Intel documentation, we were able to decode physical addresses given to us by xthwerrlog and statistically determine the location of those faults in memory. In this talk, we look at some of the statistics we have been able to gather from this work.

LA-UR: LA-UR-20-25757