Lec fine 20 Note Title No class Thes April]. Optimistic Concurring control is to use version numbers on objects and detect conflicts at eact commit time $nead(d_1, v_1, t_1)$ read(dz, vz, tz) write (d3, V1+V2, t3) at end check that d, shell has version to, dzhas tz dy has Ez Why?

Local Counter. inc ()

commit Transaction () - dues tests like those I described for opt, conc. control.

Refer to Figures 1 and 2 in the paper for how to use the DSTM primitives.

Why does consistency have to be checked at each open()? Otherwise a transaction (that would have to be aborted) could read conflicting values committed by different transactions: this could lead to null pointer dereference or other nastiness: so the decision is made to abort as soon as the need is

ad this one Iread younite

Refer to Figure 3 for the representation of transactional objects and figures 4 and 5 for what happens when an object is opened for writing in the two cases that the most recent opener for write has committed or aborted. If the most recent opener is still active then we either need to abort it or wait awhile and try again (contention policy management decision).

To open an object for writing in transaction A O For read: A N B OL = start O. Start N1 if OL. B. active: B. abort OL = start else of OL.B. committed V= OL, new IF OL. B. committed & else v= OL. old NL. Ad = OL. new NL. new = OL. new. clone () currInead - vtable += CAS (start, OL, NL) (\circ, v) else if OL. B. aborted NL- dd = OL. old NL. new - OL. old. clone () CAS(stend, OL, NL) else // OL.B is active

abort B Validation at commit or open: - for read sysies: are all the (0, v) pairs shift current " for written objects - check if current hact have been aborted. Next Thursday Composable Menoy Transactions Harris, Marlow, Peyton-Jones, Herlihy Refer to Figure 6 for performance analysis. The top graph is not very interesting in my opinion. The most striking thing about the lower graph is how much, much, (much*) worse in ALL of the concurrent versions are than the single-threaded test case. The lesson that I take is that performance improvement from parallelism is doomed when operations on shared data represent too large a fraction of the total operations performed, regardless of whether one uses locking or STM methods. It is interesting that the STM methods do perform better than the locking methods beginning at a modest level of concurrency.