



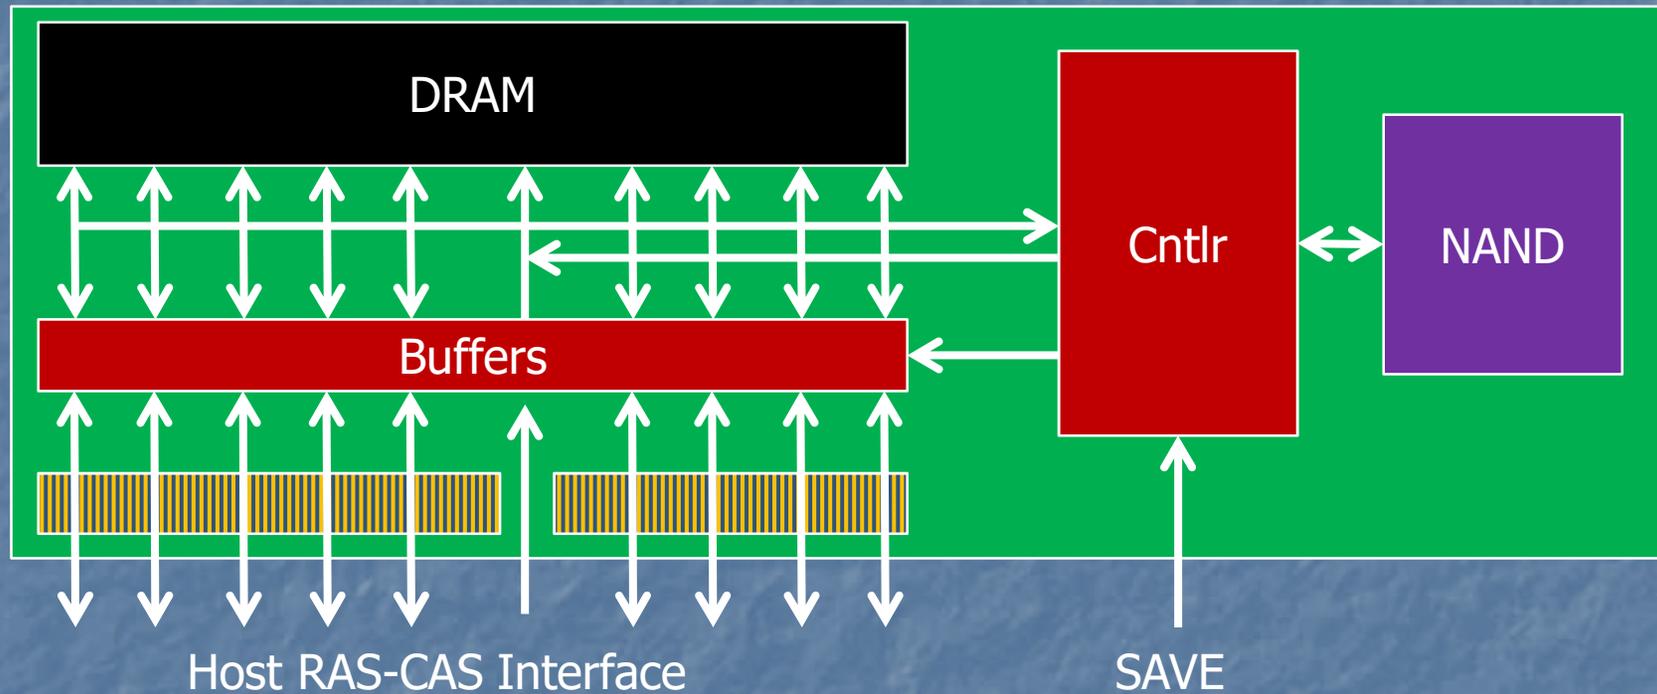
DISCOBOLUS
DESIGNS

NVDIMM-P

A New Hybrid Architecture

14 April 2016

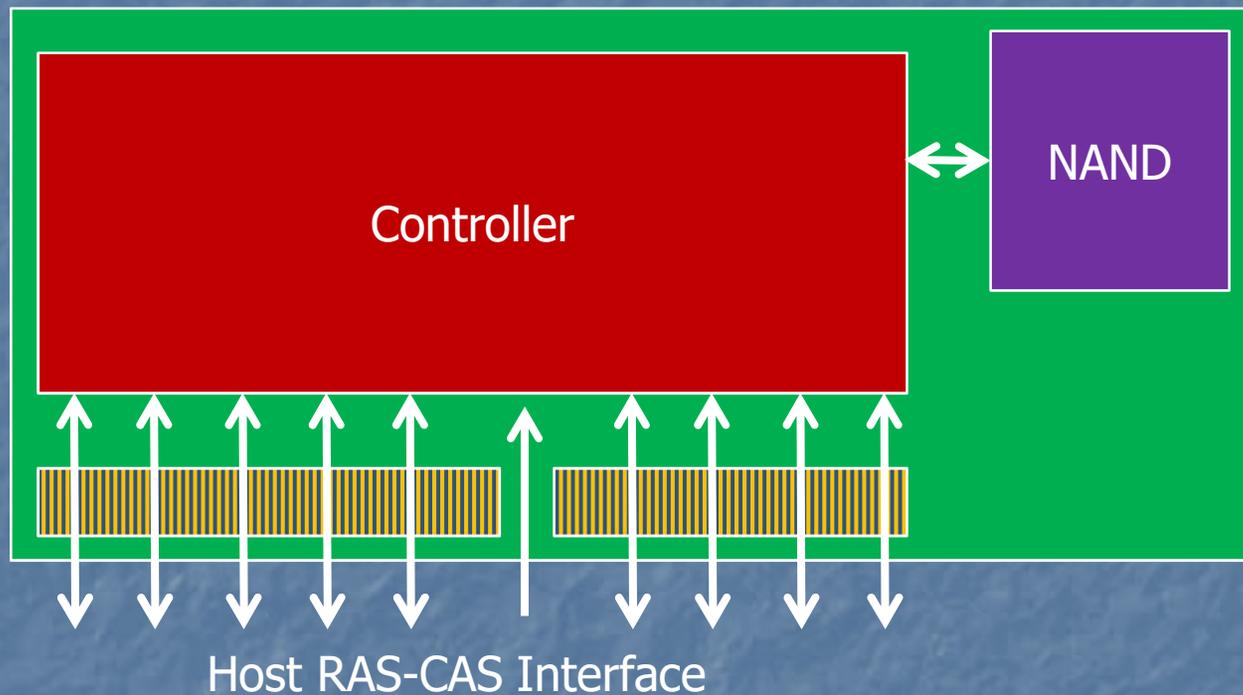
NVDIMM-N: DRAM Persistence



DRAM accessed at DRAM speeds
Contents saved to NAND on power fail
Restored to DRAM when power resumes



NVDIMM-F: Block Accessed NAND Flash

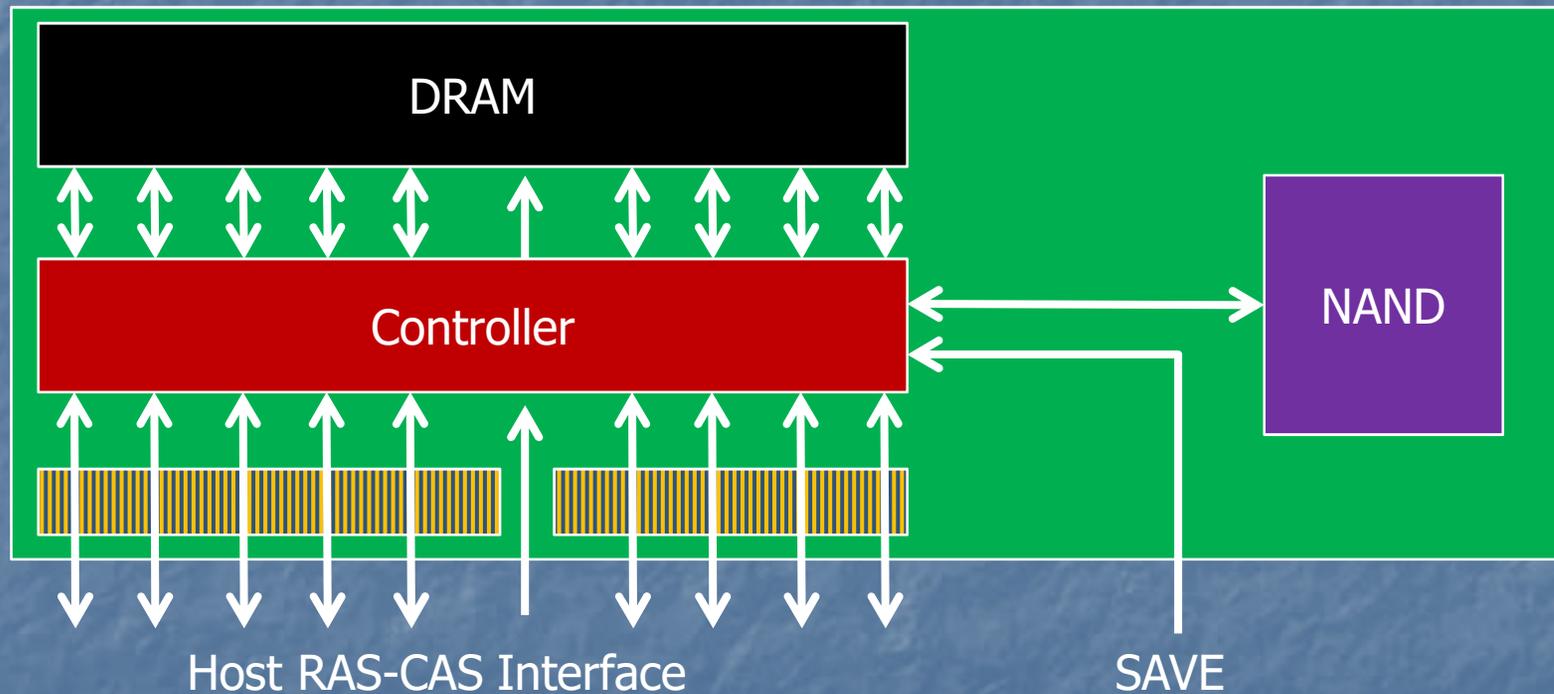


No DRAM
Flash accessed in native block format



NVDIMM-P

Combines DRAM & Flash



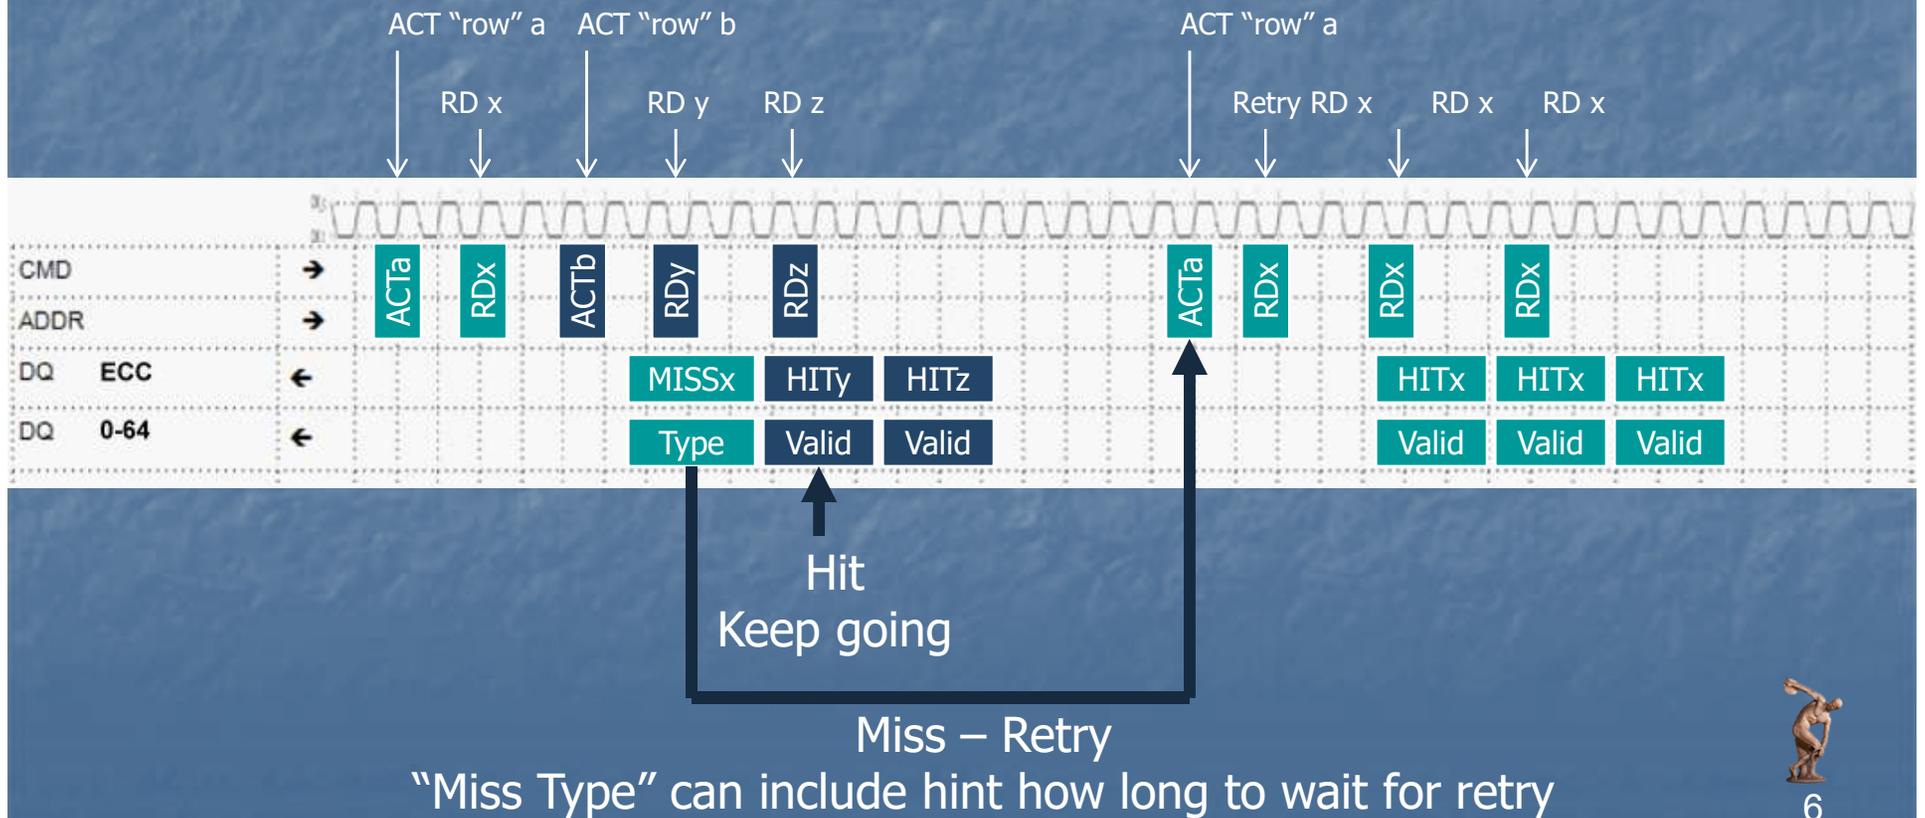
Variations Under Discussion

- NVDIMM-controlled cached interface
 - Host blindly requests access
 - On cache hit, request satisfied immediately
 - On cache miss, request postponed + retried
- Host-directed cache interface
 - Host memory tracks NVDIMM cache
 - Forces cache fill & replacement
 - No cache misses on DDR4/DDR5 bus



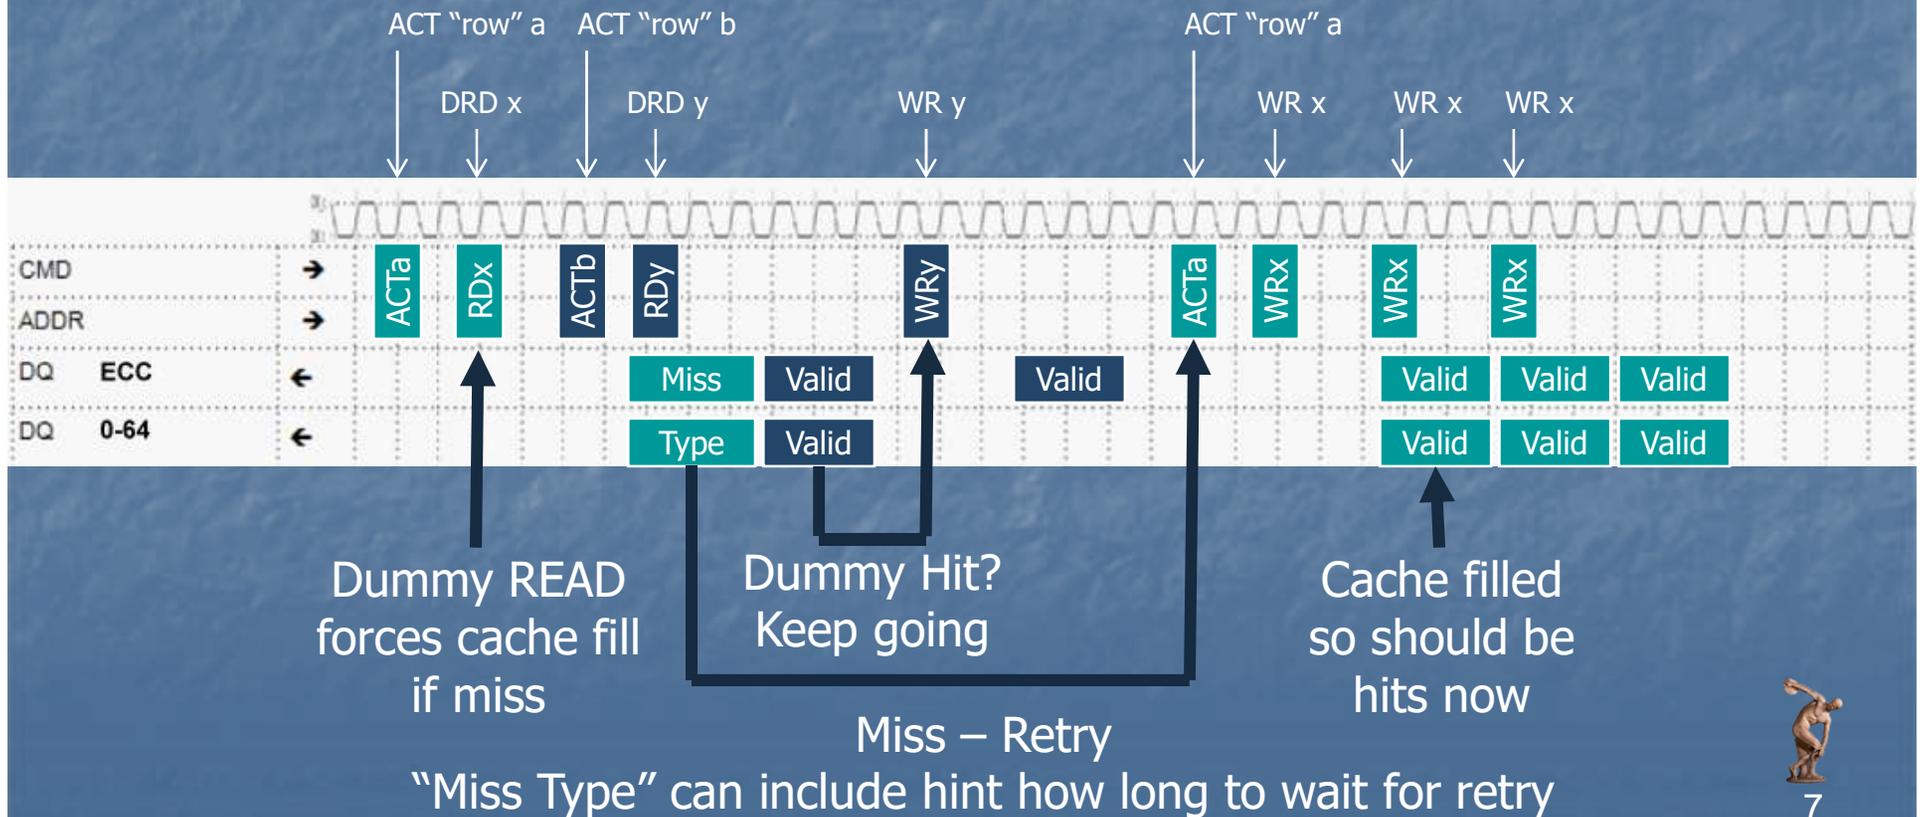
NVDIMM-P Cache Control – READ

- Data response for READ includes cache hit/miss in ECC
- Retry only needed on miss
- O-o-O transactions allowed by simple try-retry

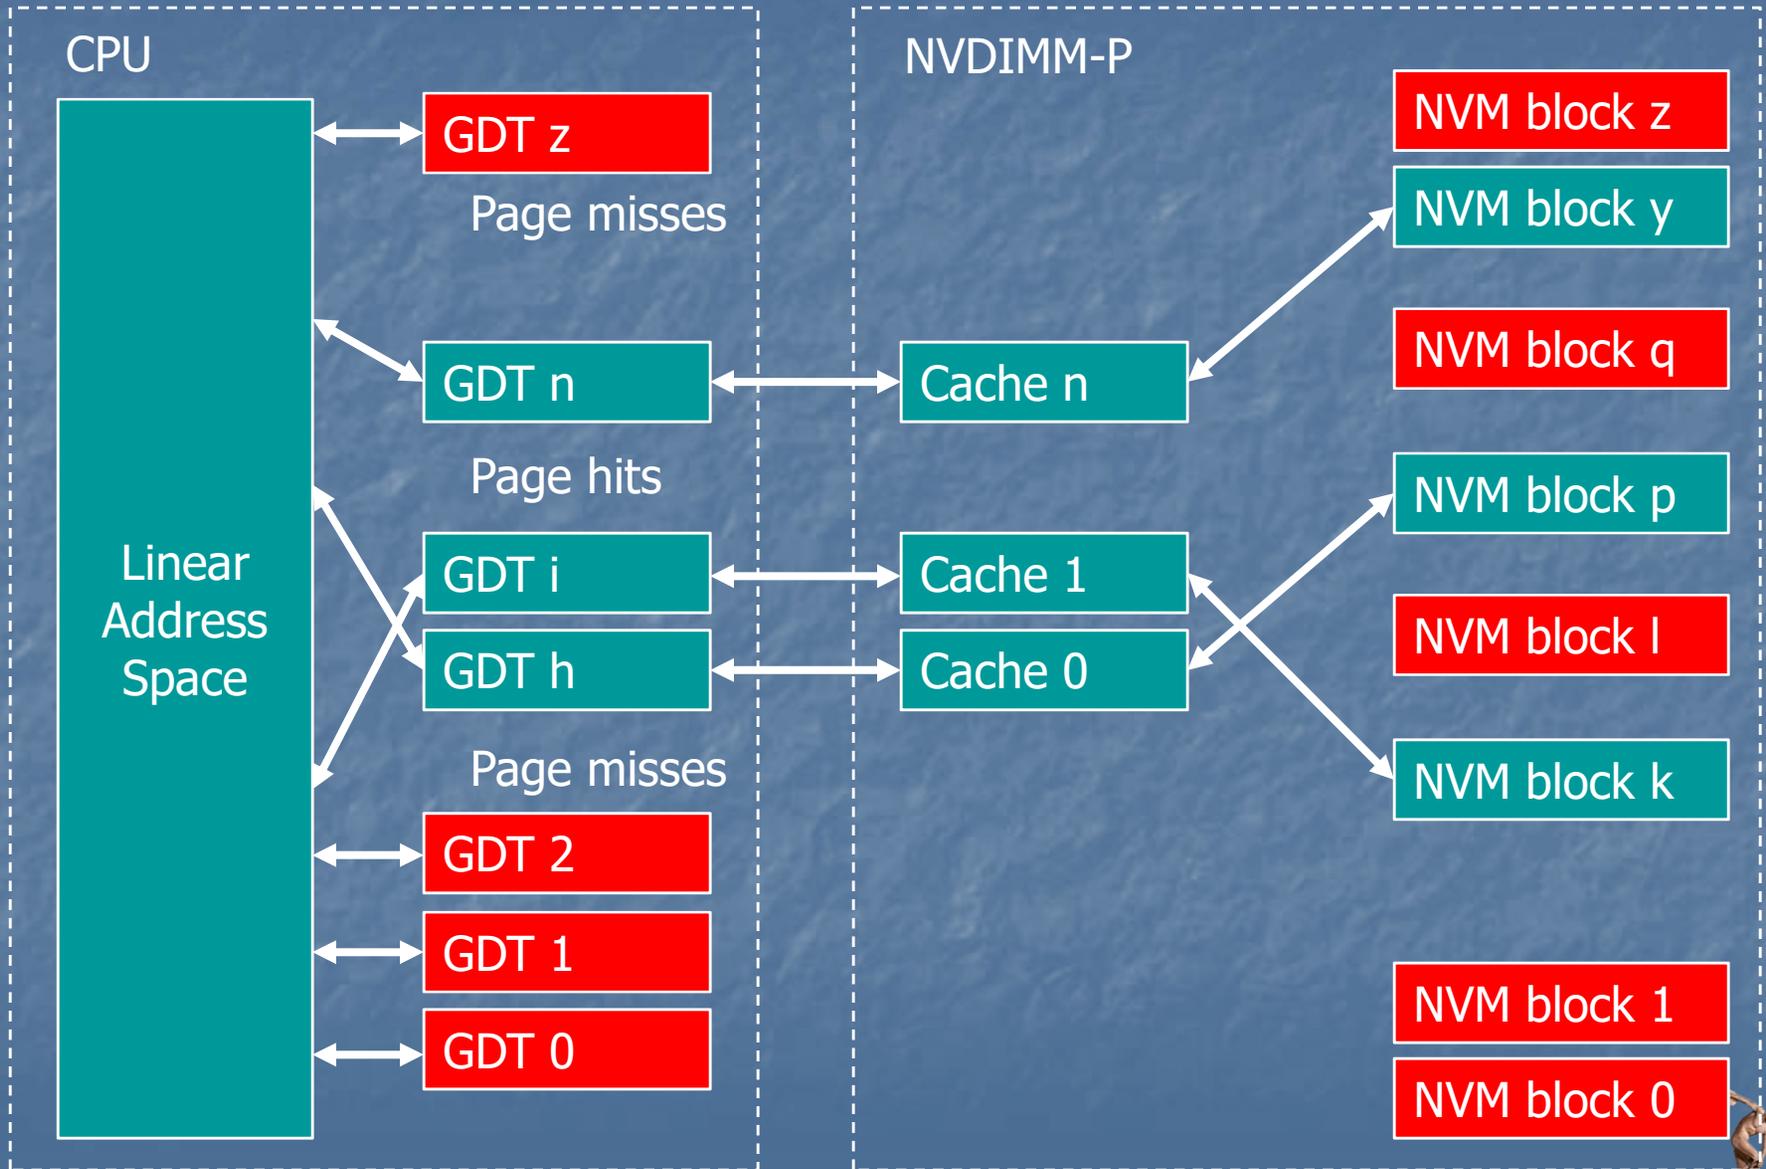


NVDIMM-P Cache Control – WRITE

- Dummy READ to check cache status, force cache fill
- Dummy READ gets valid reply? Cache hit, go ahead
- Subsequent WRITES do not require ACK



Host Directed Cache Control – R/W

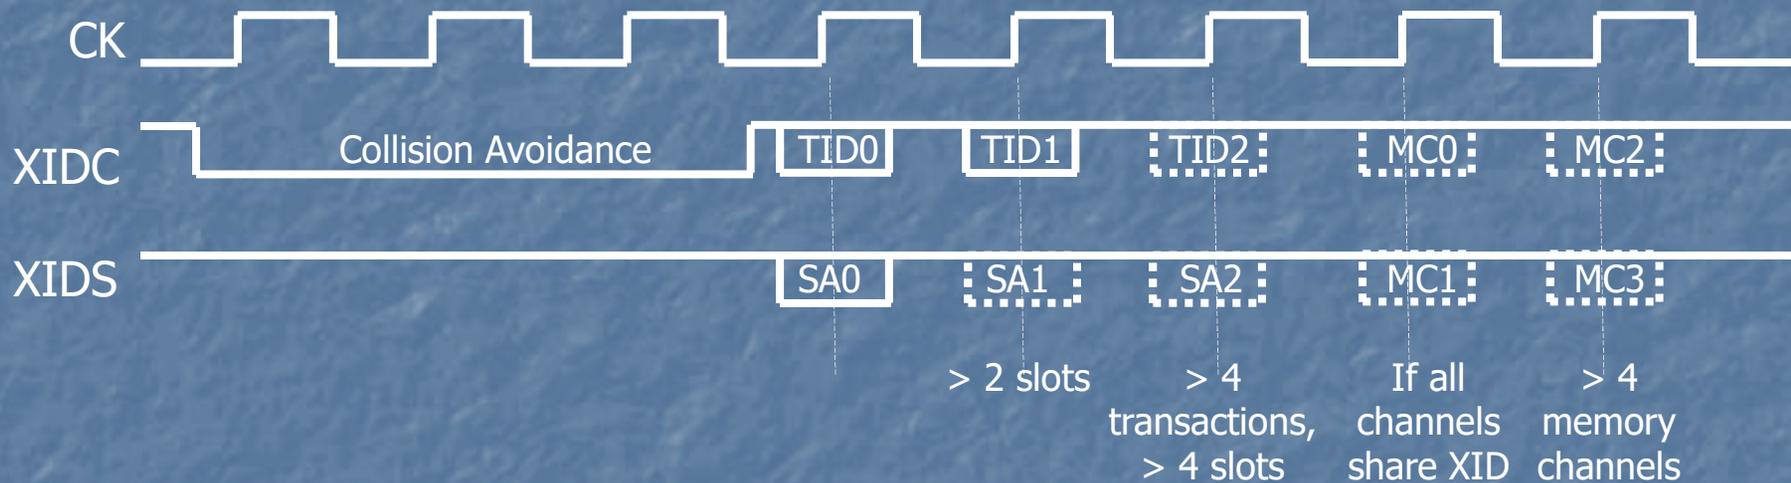


Optimizing for NVDIMM

- Polling on cache misses consumes bus bandwidth
- NVM activities are non-deterministic
 - Cache replacement
 - Error scrubbing
 - Wear leveling
- Sideband signals may be used to speed up response to miss
- Support for out-of-order needed



Considering Sideband ID Bus



- Host waits for XID on cache miss
- XID bus indicates completed O-o-O transaction



Linear Addressing

- Current DDR4 bus limits DIMM capacity to 128GB of linear mapped memory
- Desire to put multiple TB on an NVDIMM
- Transaction IDs for O-o-O included



Considering Extended Addresses

Function	CKE Previous Cycle	CKE Current Cycle	CS _n	ACT _n	RAS _n / A16	CAS _n / A15	WE _n / A14	BG0 - BG1	BA0 - BA1	C2-C0	BC _n / A12	A17	A13	A11	A10 / AP	A9 - A0
Bank Activate [ACT]	H	H	L	L	RA			BG	BA	V	RA					
Bank Activate Extension [EXT]	H	H	L	H	L	H	H	V	V	V	V	V	V	V	H	EA



- Increases total memory capacity by up to 2^{15}



Summary

- NVDIMM-P definition in discussion
- Existing DDR4 protocol supported
- Extensions to protocol under consideration
 - Sideband signals for transaction ID bus
 - Extended address for large linear addresses



Thank you

Bill Gervasi

bilge@discobolusdesigns.com



DISCOBOLUS
DESIGNS