



A Hierarchical Approach to Semiconductor Mass Storage

Peter Gillingham, CTO

CMOS ET Workshop, Vancouver, 24 Sep. 2009

NAND Flash Applications



USB Drives & Memory Card

1-4 NAND die

1 channel

40MB/s



Handheld Device

4-8 NAND die

1 channel

40MB/s



Consumer SSD

16-64 NAND die

8 channels

320MB/s



Enterprise SSD

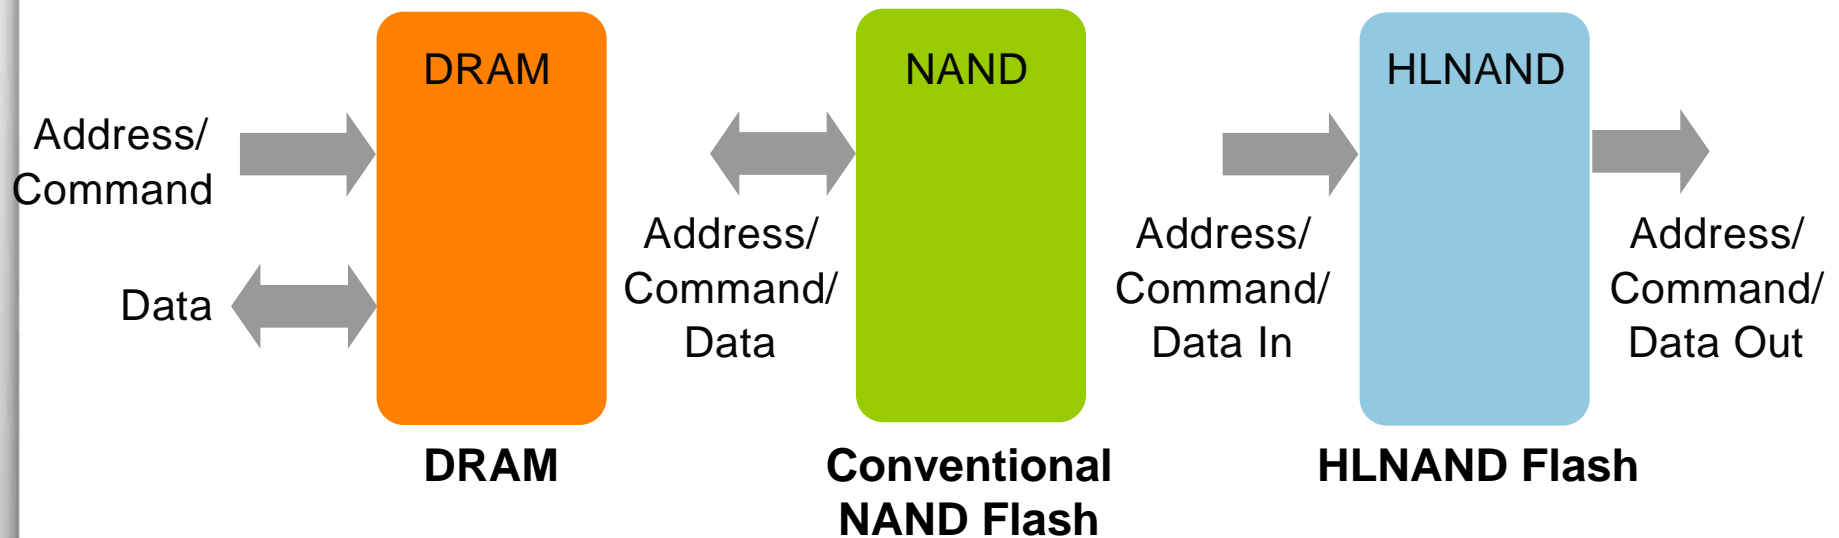
128-512 NAND die

~32 channels

~1TB/s

- Conventional NAND Flash operates at 40MB/s with up to 4 packages per channel – insufficient for emerging high performance, high capacity applications

Memory Interface Comparison



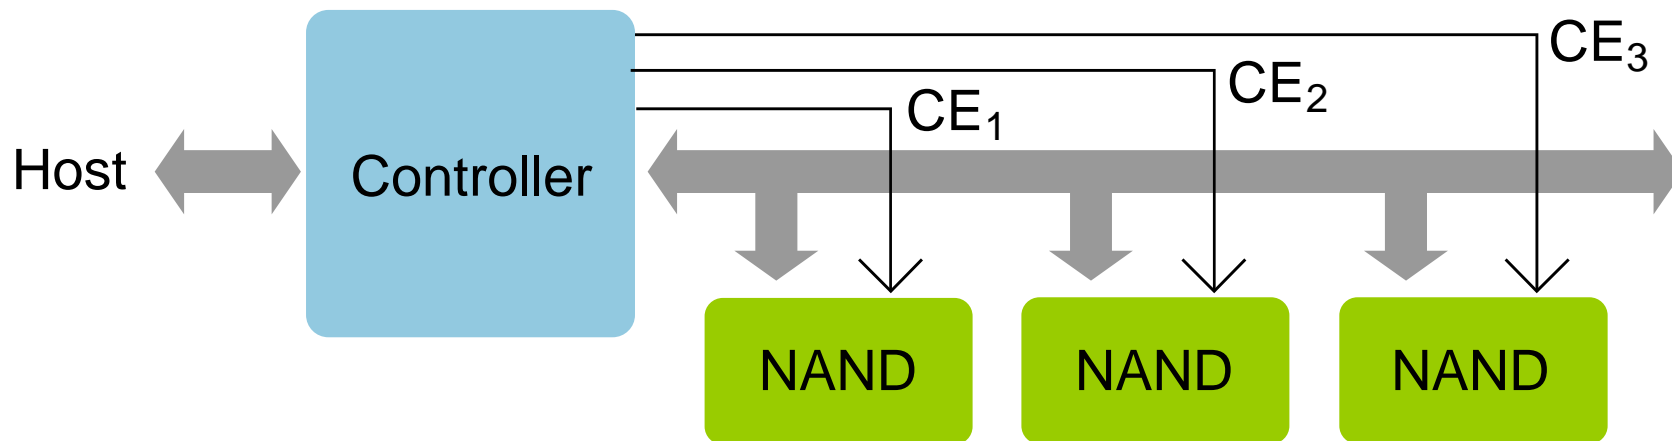
- code storage
- low latency operation
- issue commands during data transfers
- multi-drop bus
- limited number of loads

- high latency core
- long data transfers
- interrupt data transfer to issue commands
- multi-drop bus
- limited number of loads

- high latency core
- long data transfers
- interrupt data transfer to issue commands
- point-to-point daisy chain
- unlimited number of loads

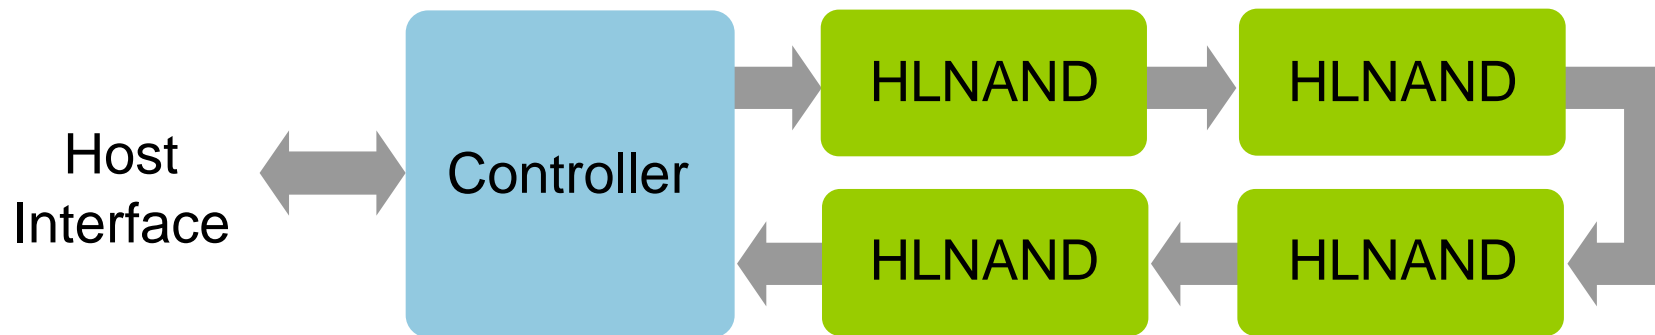
Conventional NAND Flash Interface

- 8 bit, bidirectional, multi-drop bus
- Asynchronous LVTTL signaling up to 40Mb/s/pin
- Speed degradation with more than 4 devices on bus
- Chip Enable (CE) signal required for each device
- Power hungry 3.3V I/O



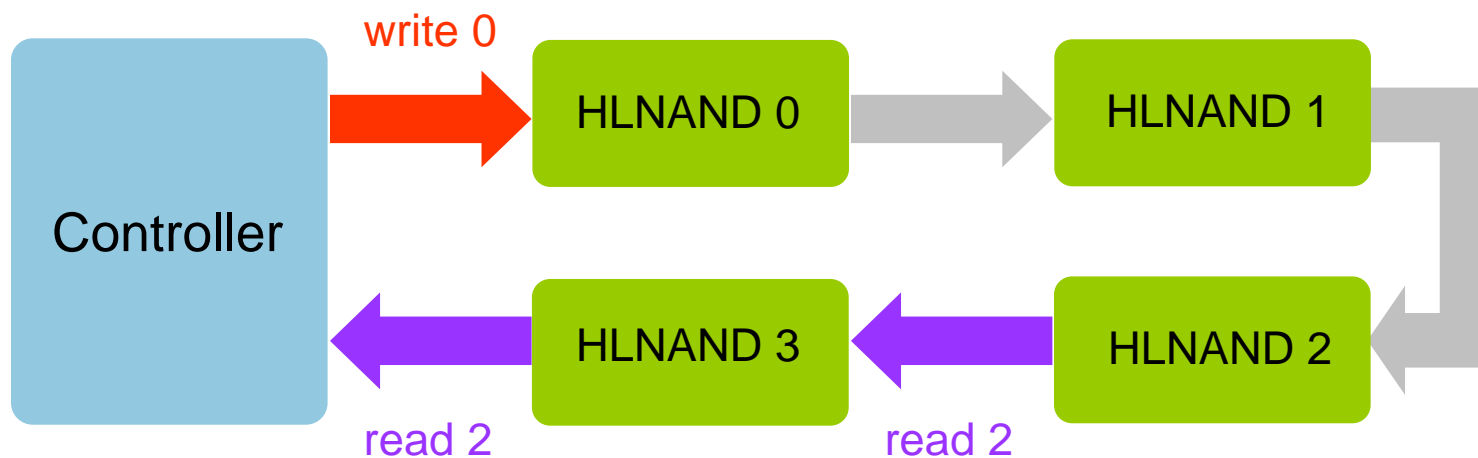
HyperLink Interface

- Unidirectional, point-to-point, daisy-chain cascade supporting as many as 255 devices in a ring
- No bandwidth degradation with additional devices
- Device address assigned on initialization
- High speed DDR signaling up to 800Mb/s/pin
- Dynamic link width programmable from 1 to 8 bits
- Low Power 1.8V I/O



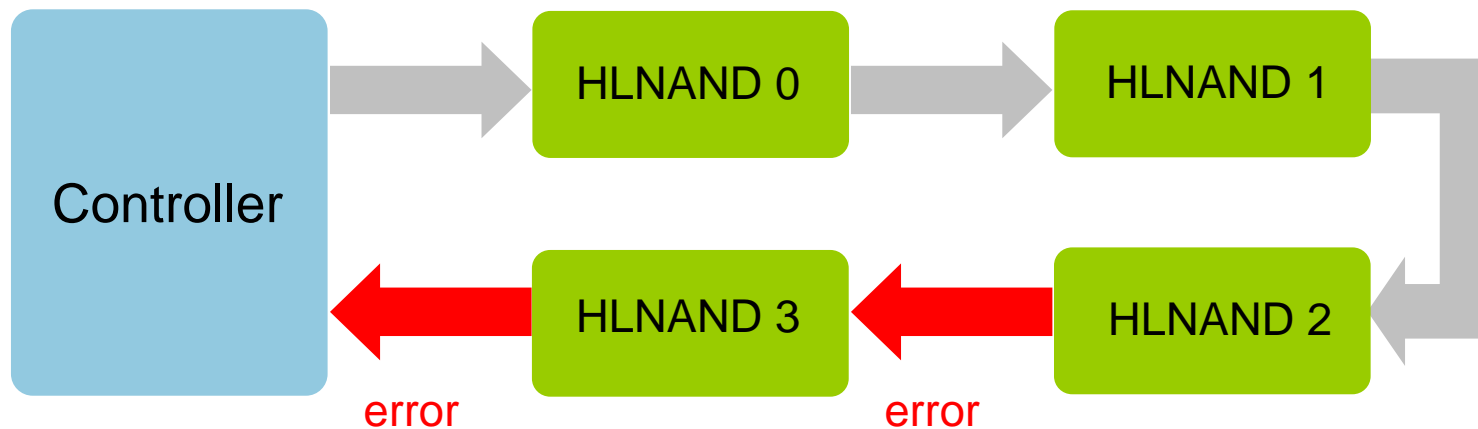
HyperLink Packet Truncation

- Once packet reaches addressed device the write data payload is truncated
- Simultaneous data transfer possible if write device is upstream of read device or the same device



Command Packet Error Recovery

- HLNAND checks Hamming code byte within command packet – if error is detected command is not executed
- Controller recognizes error in returned command packet
- Controller reads status register of HLNAND to determine whether command was executed
- Controller re-issues command if necessary

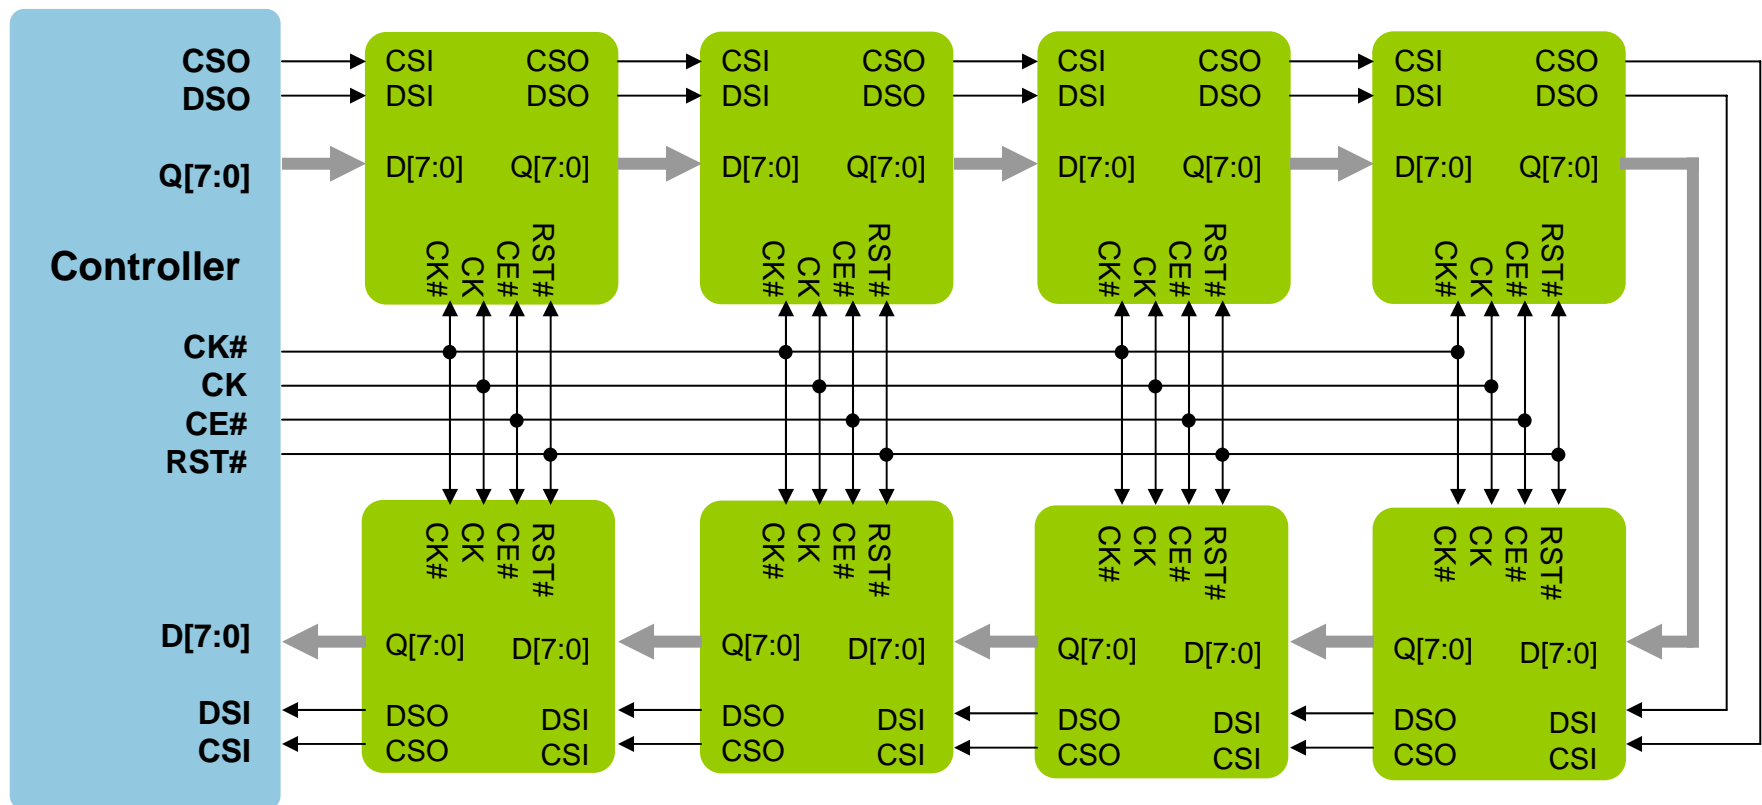


HyperLink Features

- HyperLink interface can be applied to any device
 - SLC NAND, MLC NAND, NOR, PRAM, DRAM etc.
 - self identified via configuration registers on power up
- Two modes of operation
 - HL1 supporting speeds up to 266Mb/s/pin
 - parallel distributed clock – no PLL required
 - LVCMOS signaling
 - HL2 supporting speeds up to 800Mb/s/pin
 - Source synchronous clocking – a PLL is required
 - HSTL Class1 signaling – matched output driver - no termination required - zero static power
 - Single device supports both modes by sensing Vref pin
- Optional Vpp supply pin for die cost/power reduction

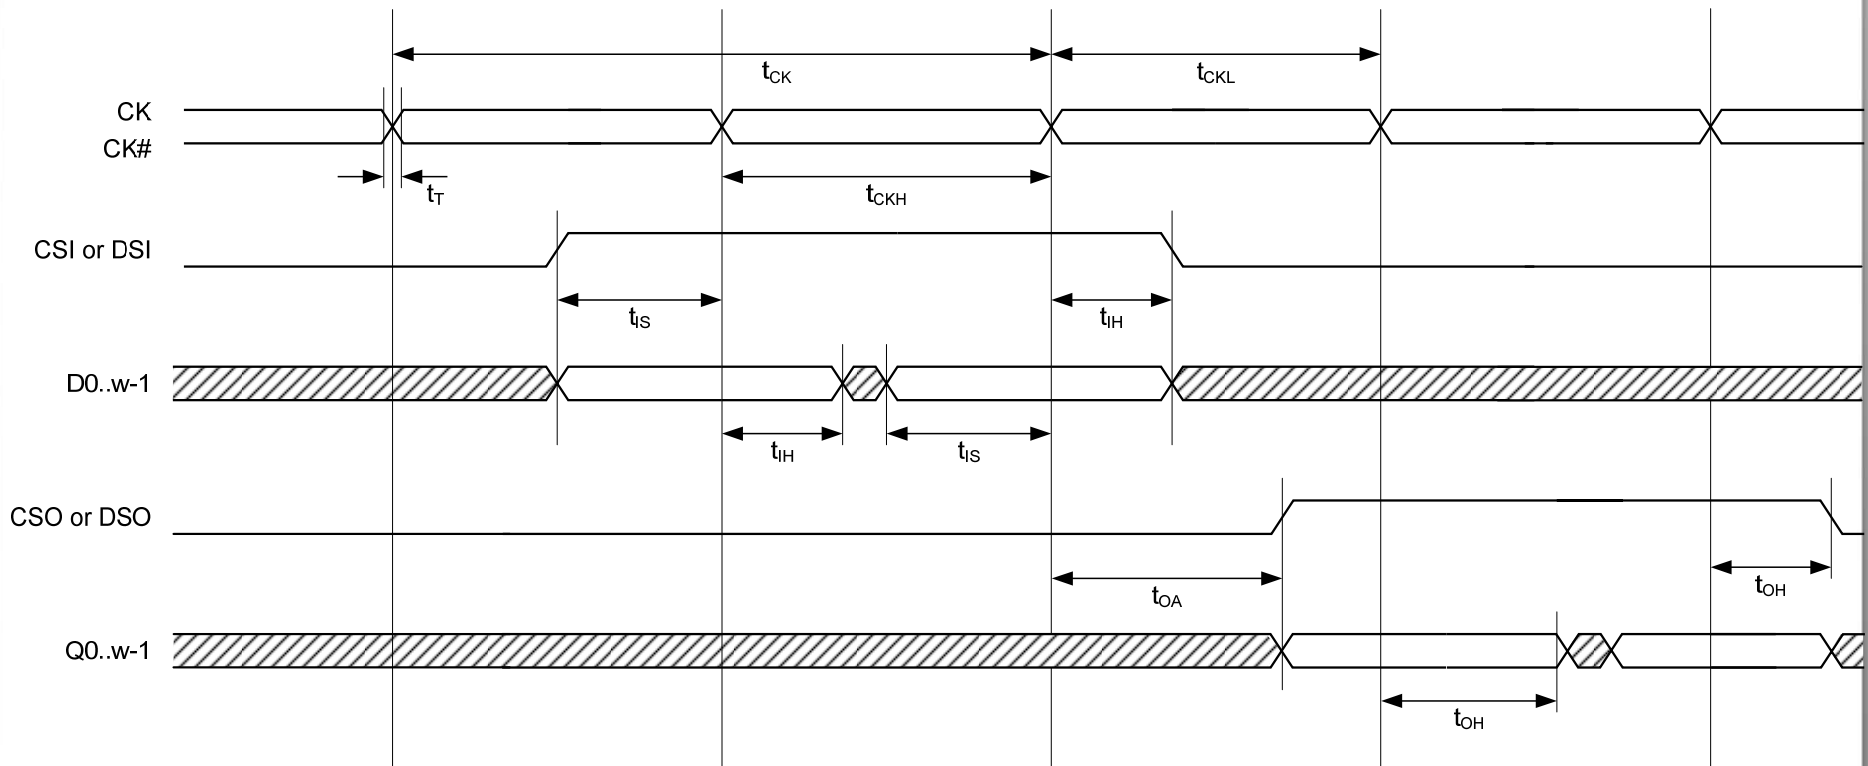
HL1 Interface

- DDR signaling up to 266MB/s, no PLL required
- No changes to existing Flash process



HL1 Timing

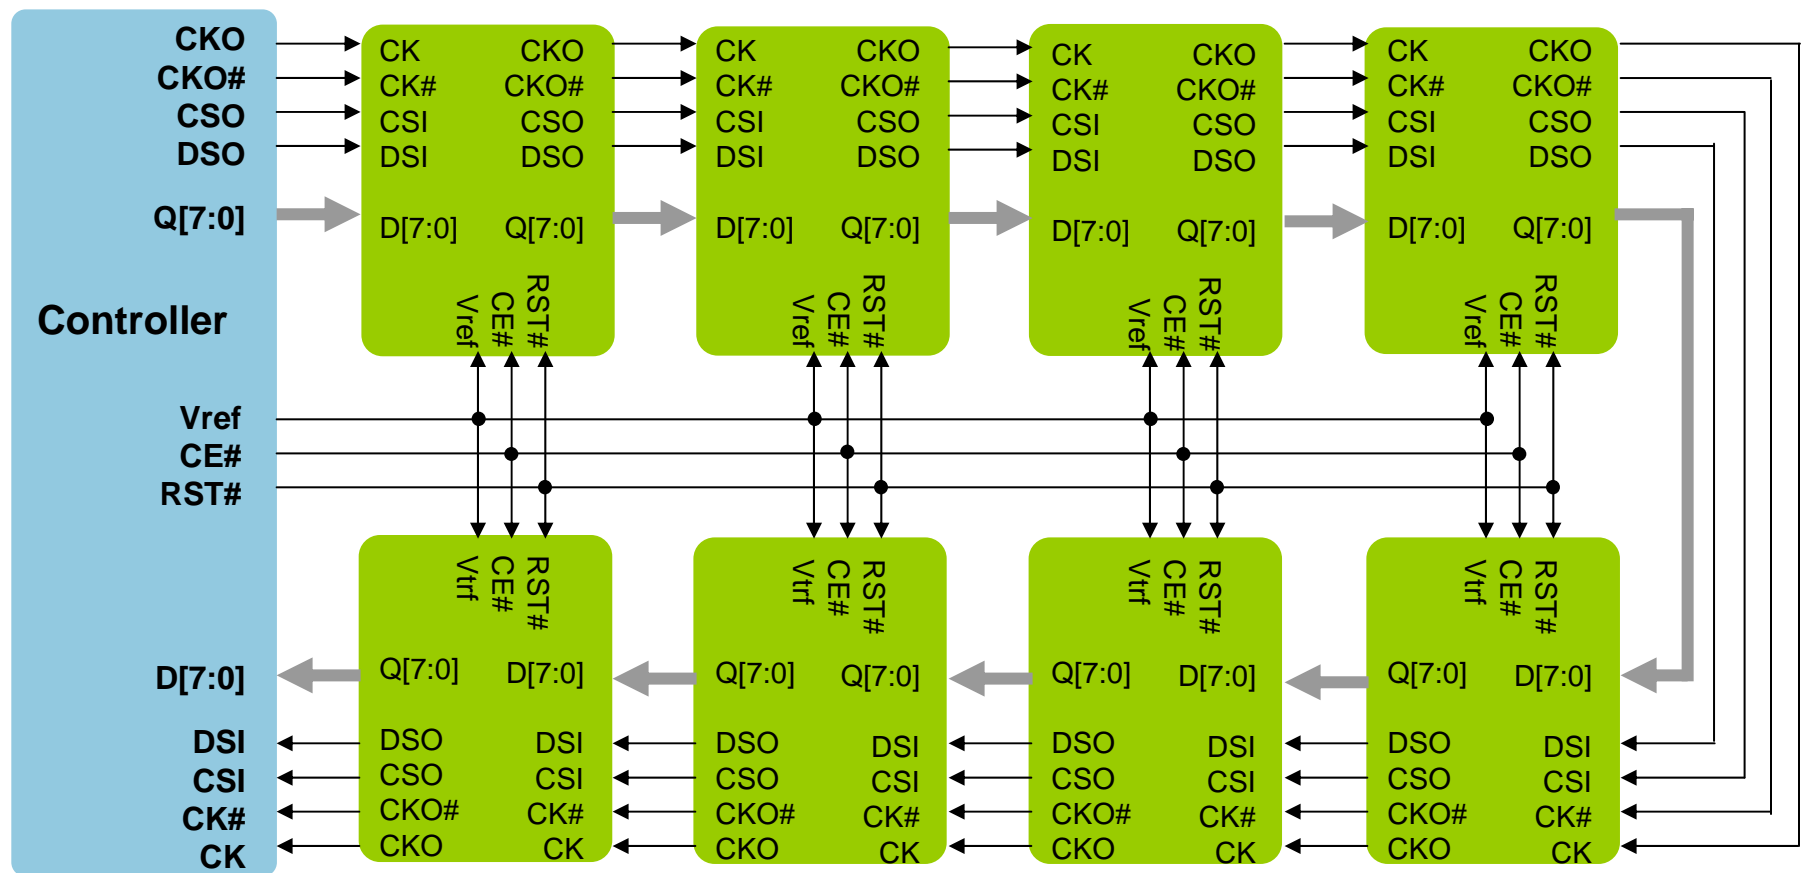
- CSI/CSO – Command and write data Strobe Input/Output
- DSI/DSO – read Data Strobe Input/Output
- Data bursts can be interrupted by terminating CSI/DSI



w = current link width

HL2 Interface

- Source synchronous DDR to 800MB/s with on-chip PLL
- May require improvement to I/O transistor performance
- Fully backward compatible to HL1



HLNAND Instructions

Operation	1 st Byte	2 nd Byte	3 rd Byte	4 th Byte	5 th Byte	6 th Byte	7 th Byte	8 th Byte	...	4229 th Byte
Page Read	DA	0Xh	RA	RA	RA	EDC				
Page Read for Copy	DA	1Xh	RA	RA	RA	EDC				
Burst Data Read	DA	2Xh	CA	CA	EDC					
Burst Data Load Start	DA	4Xh	CA	CA	EDC	DATA	DATA	DATA	...	DATA
Burst Data Load	DA	5Xh	CA	CA	EDC	DATA	DATA	DATA	...	DATA
Page Program	DA	6Xh	RA	RA	RA	EDC				
Block Erase Address Input	DA	8Xh	RA	RA	RA	EDC				
Erase	DA	AXh	EDC							
Operation Abort	DA	CXh	EDC							
Read Status Register	DA	F0h	EDC							
Read Device Information Register	DA	F4h	EDC							
Read Link Configuration Register	DA	F7h	EDC							
Write Link Configuration Register	FF	FFh	DATA	EDC						

DA = Device Address, CA = Column Address, RA = Row Address, X = Bank Number, EDC = Error Detection Code

NAND Interface Comparison

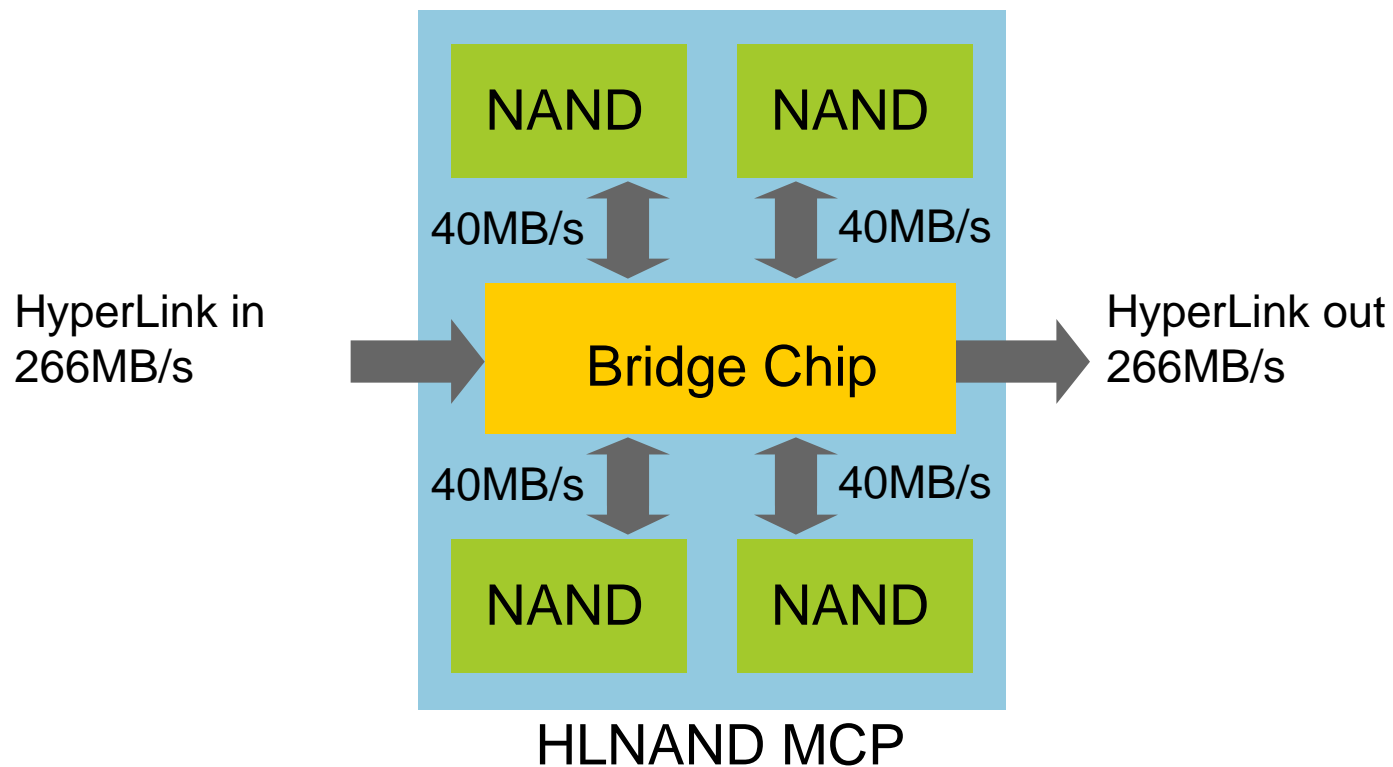
	Traditional NAND	ONFI 2.0, toggle mode	HLNAND HL1	HLNAND HL2
Bandwidth	40MB/s	166MB/s	266MB/s	800MB/s
# devices	4	4	255	255
I/O supply	3.3V	1.8V	1.8V	1.8V
I/O driver	50Ω	18Ω	50Ω	50Ω
# signal pins	15	16	24	24
Read/write	simplex	simplex	duplex	duplex
PLL	no	no	no	yes

HyperLink Advantages

- **Bandwidth**
 - Point-to-point signaling extends to much higher data-rates than multi-drop bus
- **Power**
 - No termination resistors required
 - Point-to-point I/O drivers are much smaller than multi-drop bus drivers and have significantly lower capacitance
 - Packet truncation at destination device also reduces power
- **Scalability**
 - Up to 255 devices in a single ring

HLNAND Multi-Chip Package

- Use legacy NAND die to achieve high performance
- External HyperLink interface - 8 bits
- Internal NAND interfaces – 4 x 8 bits

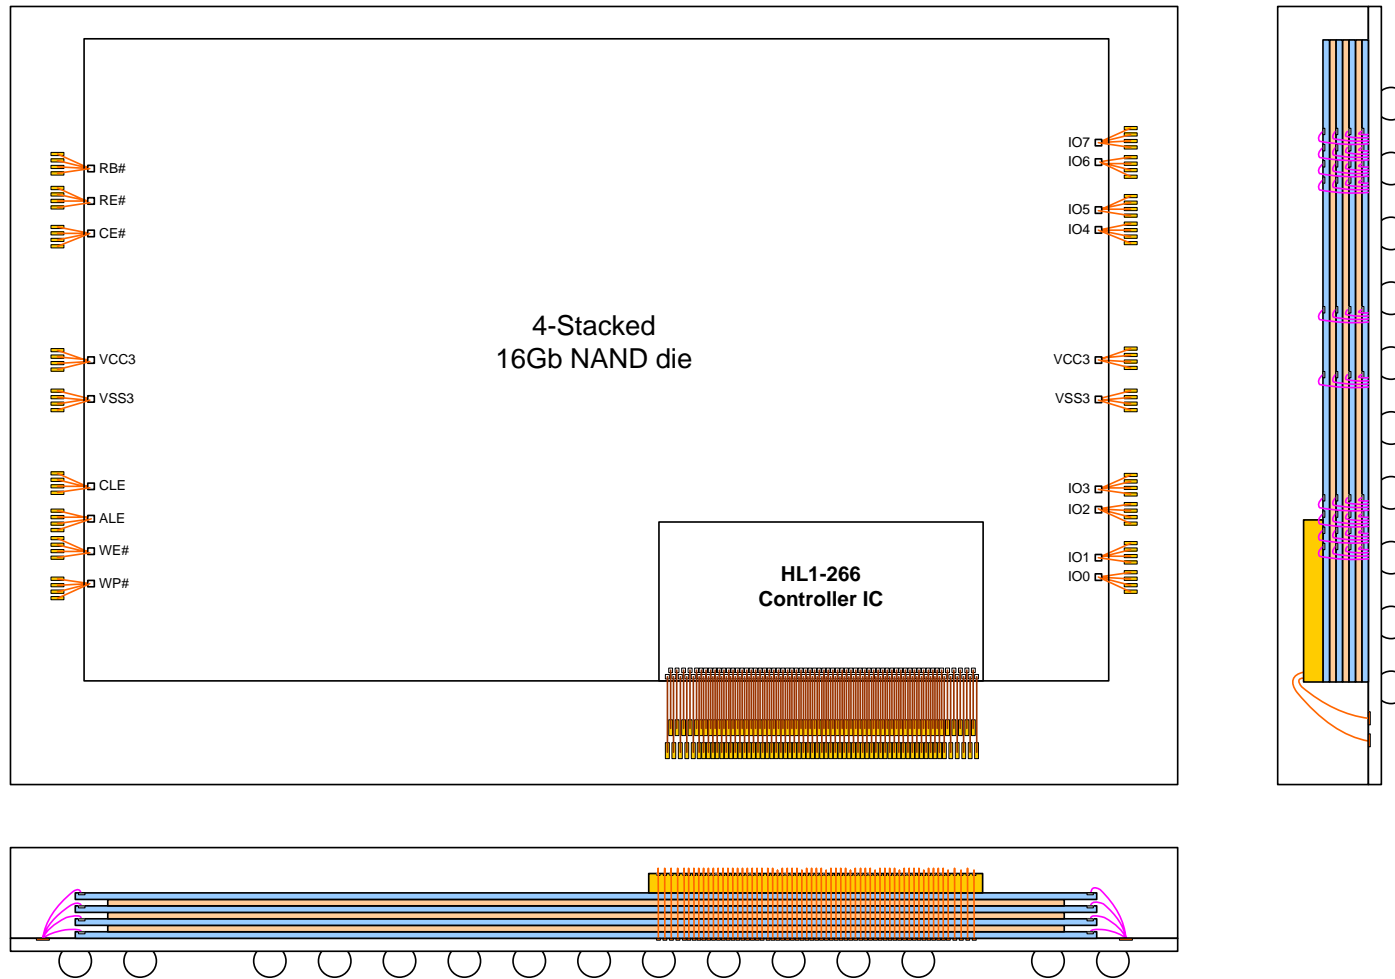


HLNAND MCP Features

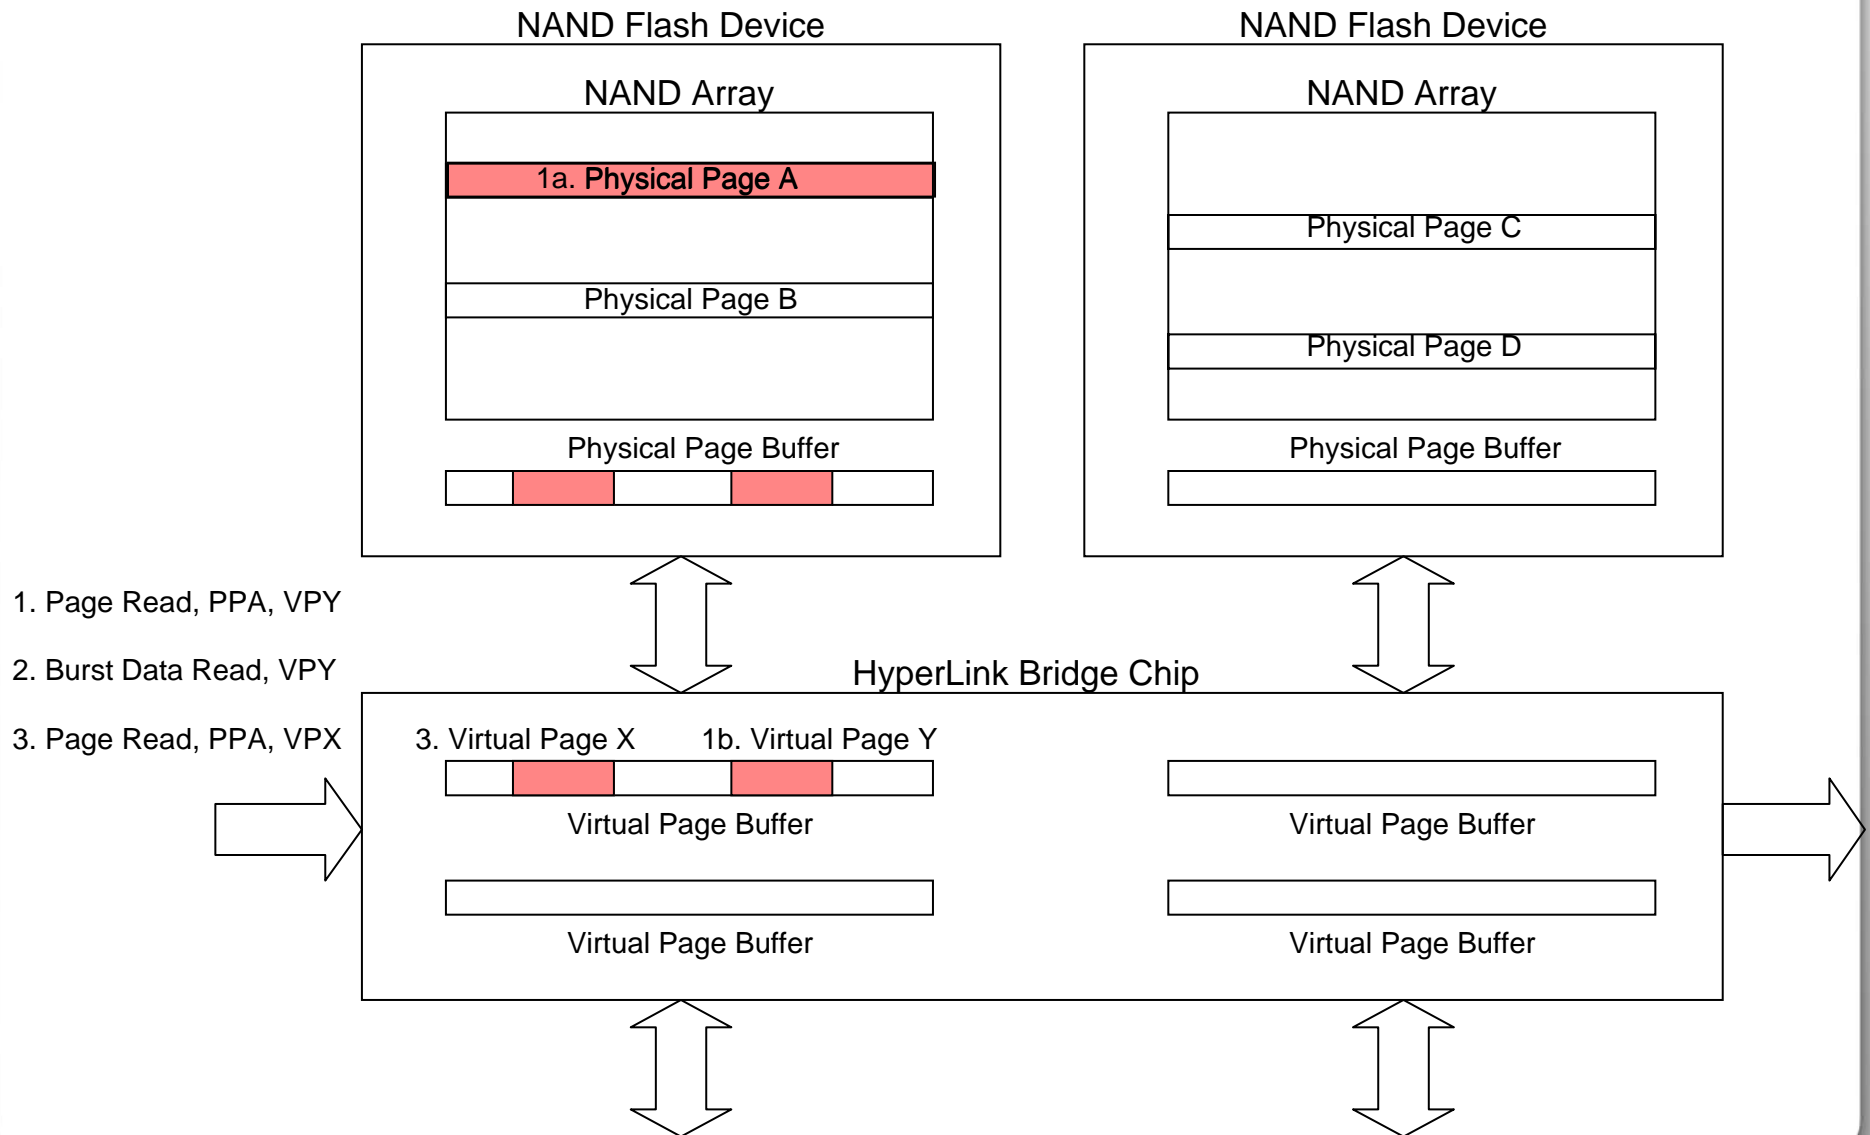
- **Concurrent Operation of NAND Die**
 - Simultaneous read or write data transfers to all four NAND die providing up to 160MB/s in a single package
- **Power Savings**
 - Hierarchical structure isolates inactive die from data transfers to save power
 - External bus operates at 1.8v while legacy internal bus uses 3.3V
- **Programmable Virtual Page Buffer**
 - Bridge chip includes 4 x 4KB virtual page buffers to mirror physical page buffer on NAND die
 - Virtual Page Buffer can be programmed from 128B to 4KB to transfer only the data required to/from the NAND die
 - No special commands required – controller issues the same commands to monolithic HLNAND and HLNAND MCP devices

HLNAND MCP Package

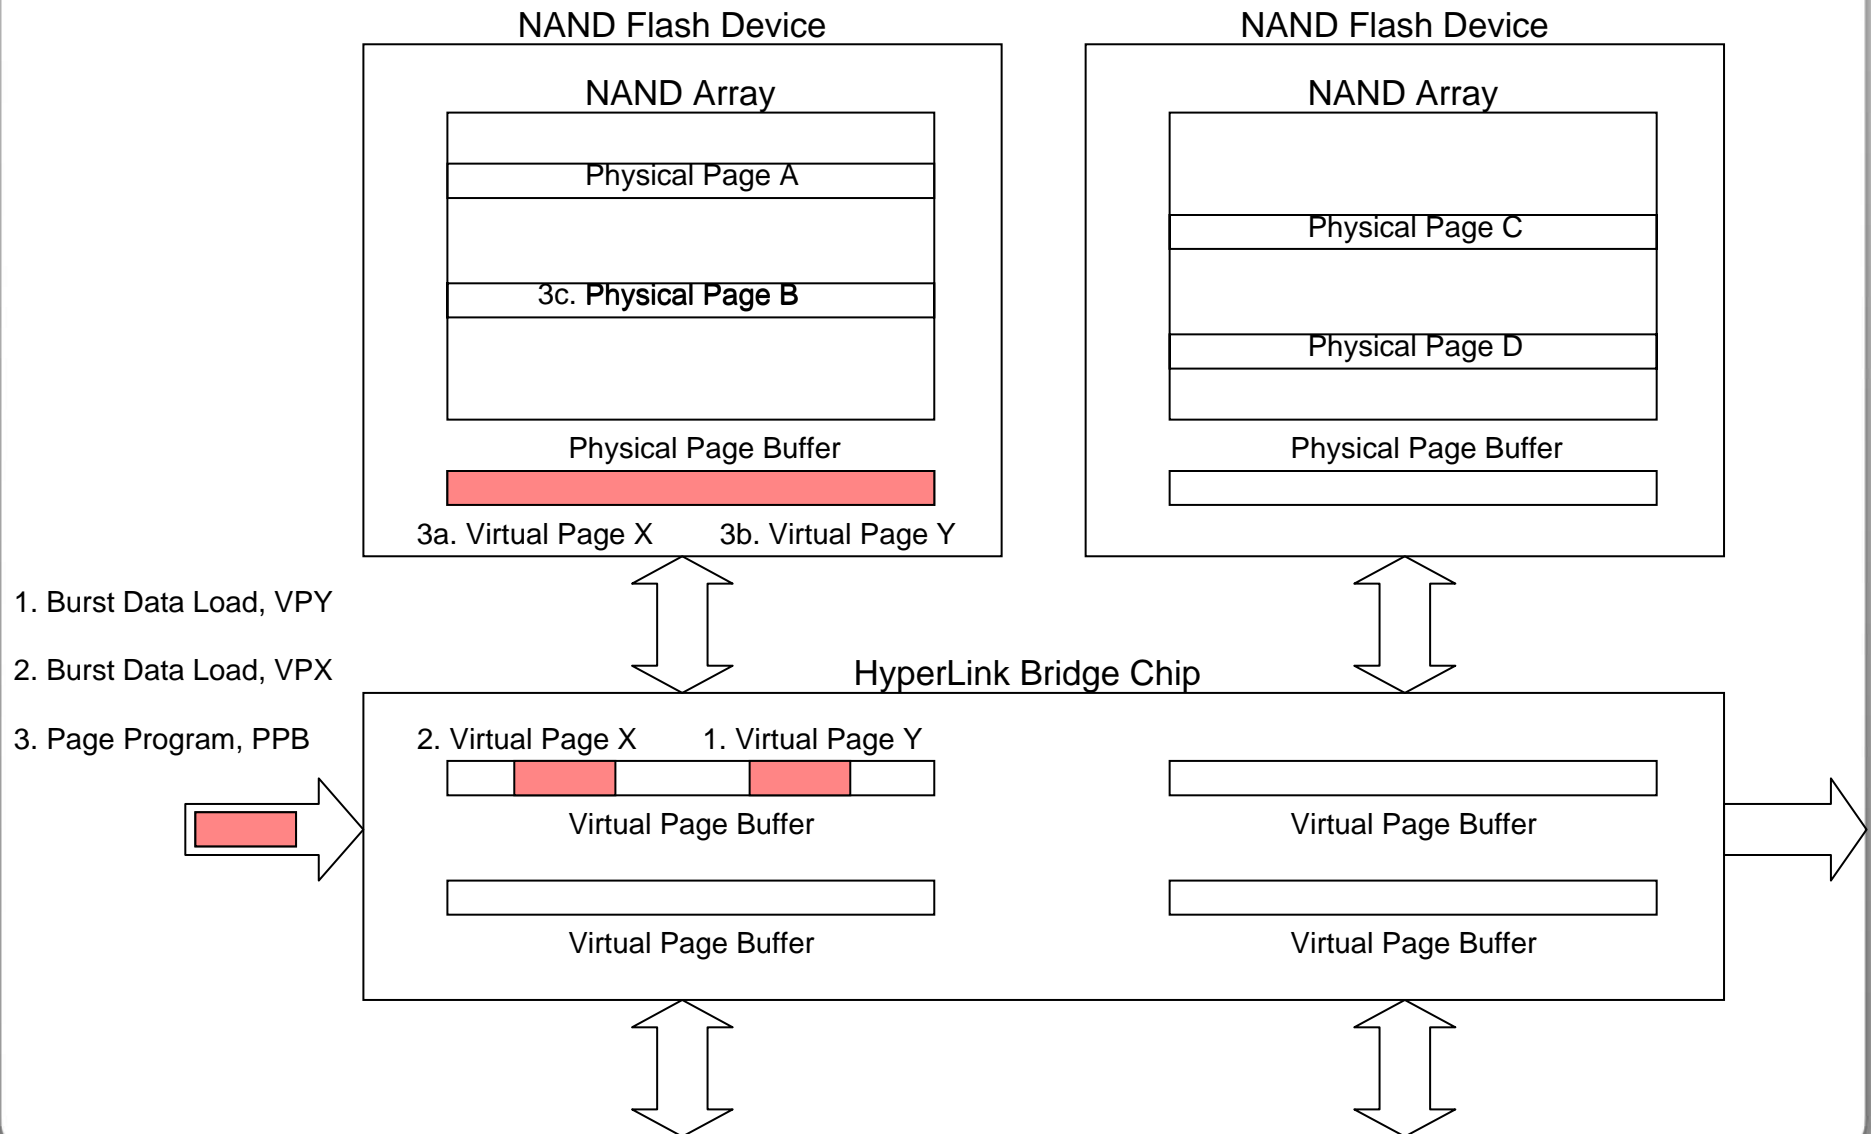
- 12mm x 18mm, 100 ball BGA



HLNAND MCP Read Operations



HLNAND MCP Program Operations




64GB HLNAND Module

- 200-pin DDR2 SO-DIMM form factor
- 8 x 64Gb HLNAND MCP (4 on each side)
- High speed signals fully shielded for signal integrity



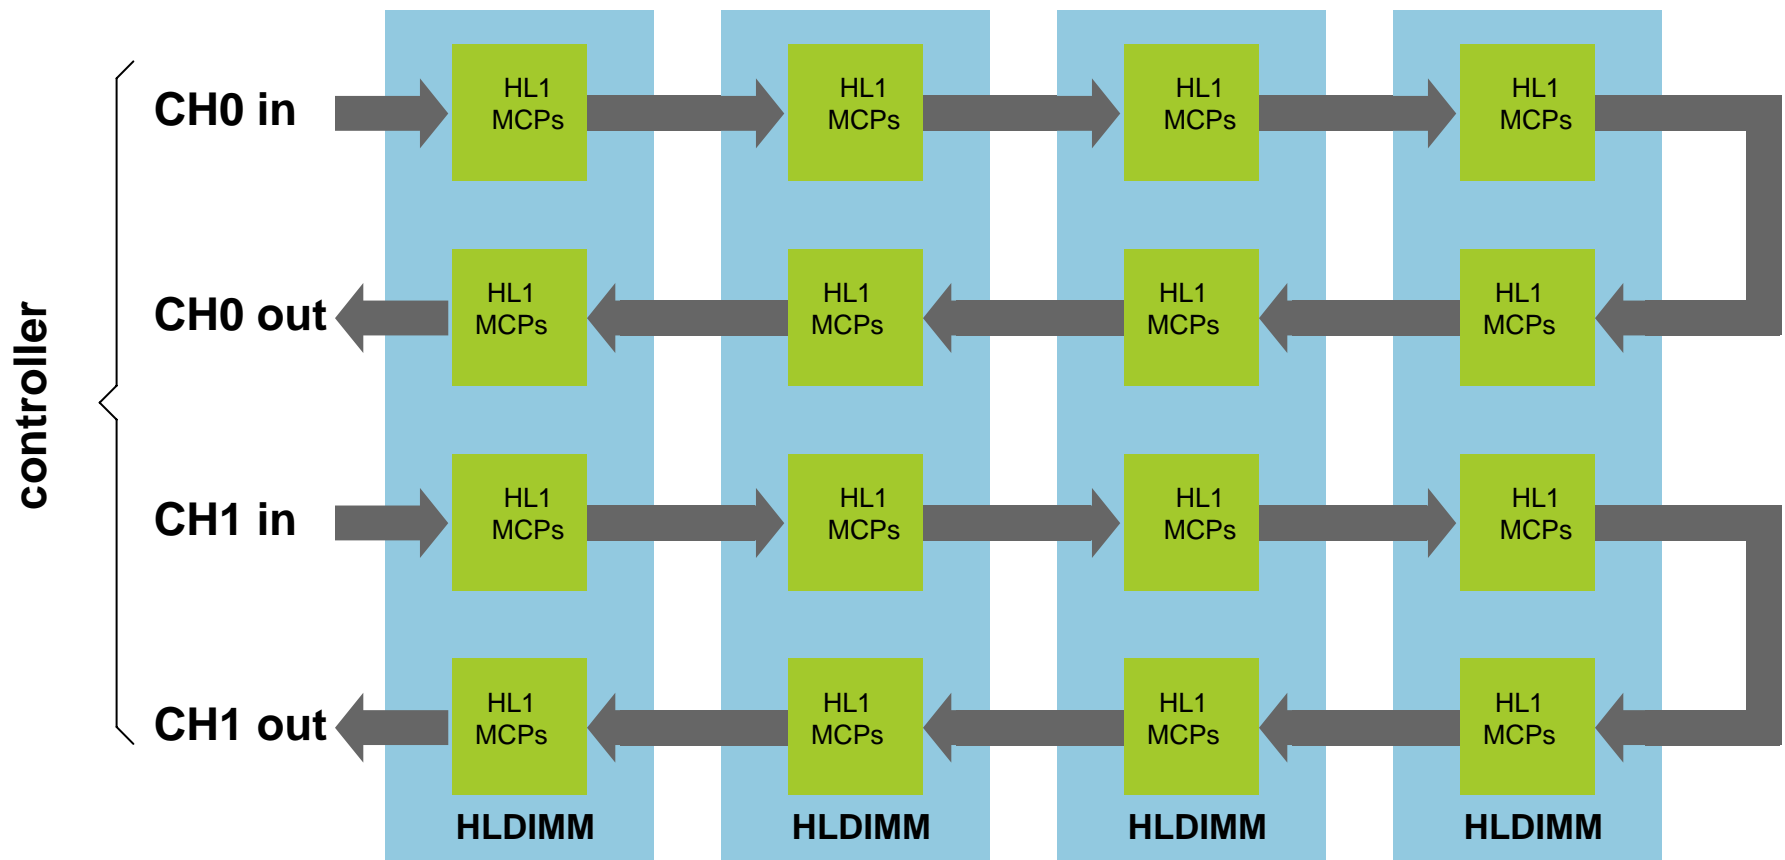
HLDIMM Port Configuration

- 2 x HyperLink HL1 interfaces with 533MB/s read + 533MB/s write = 1066MB/s aggregate throughput



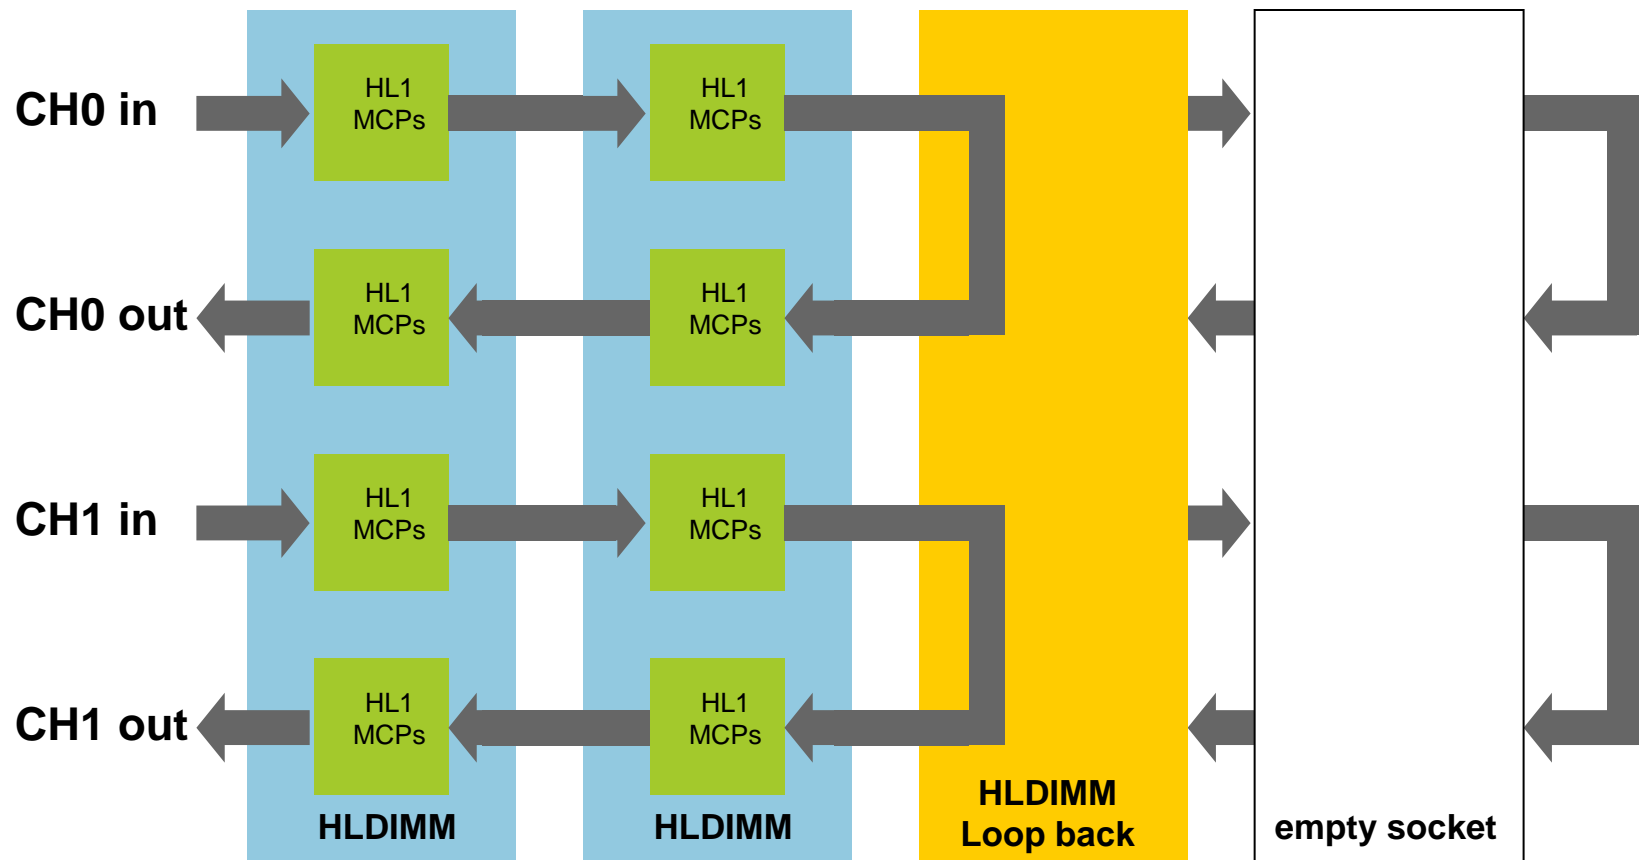
HLDIMM Module Interconnection

- Short, point-to-point connections for high speed
- Each channel loops back at last socket



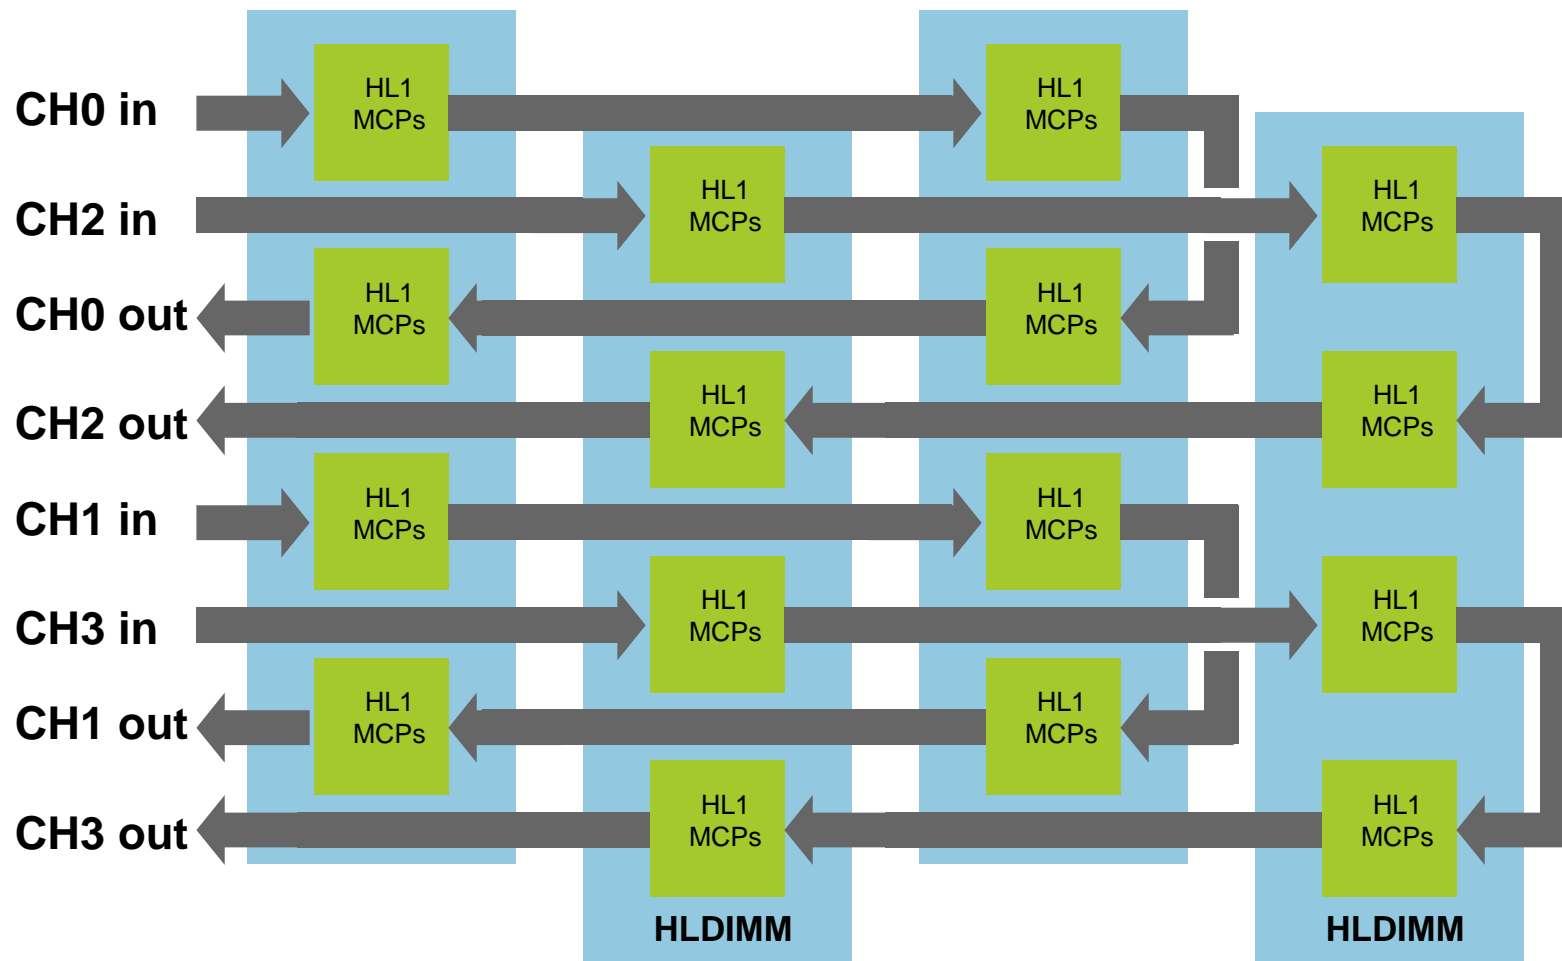
HLDIMM Loop Back Module

- Loop back channel in partially populated system



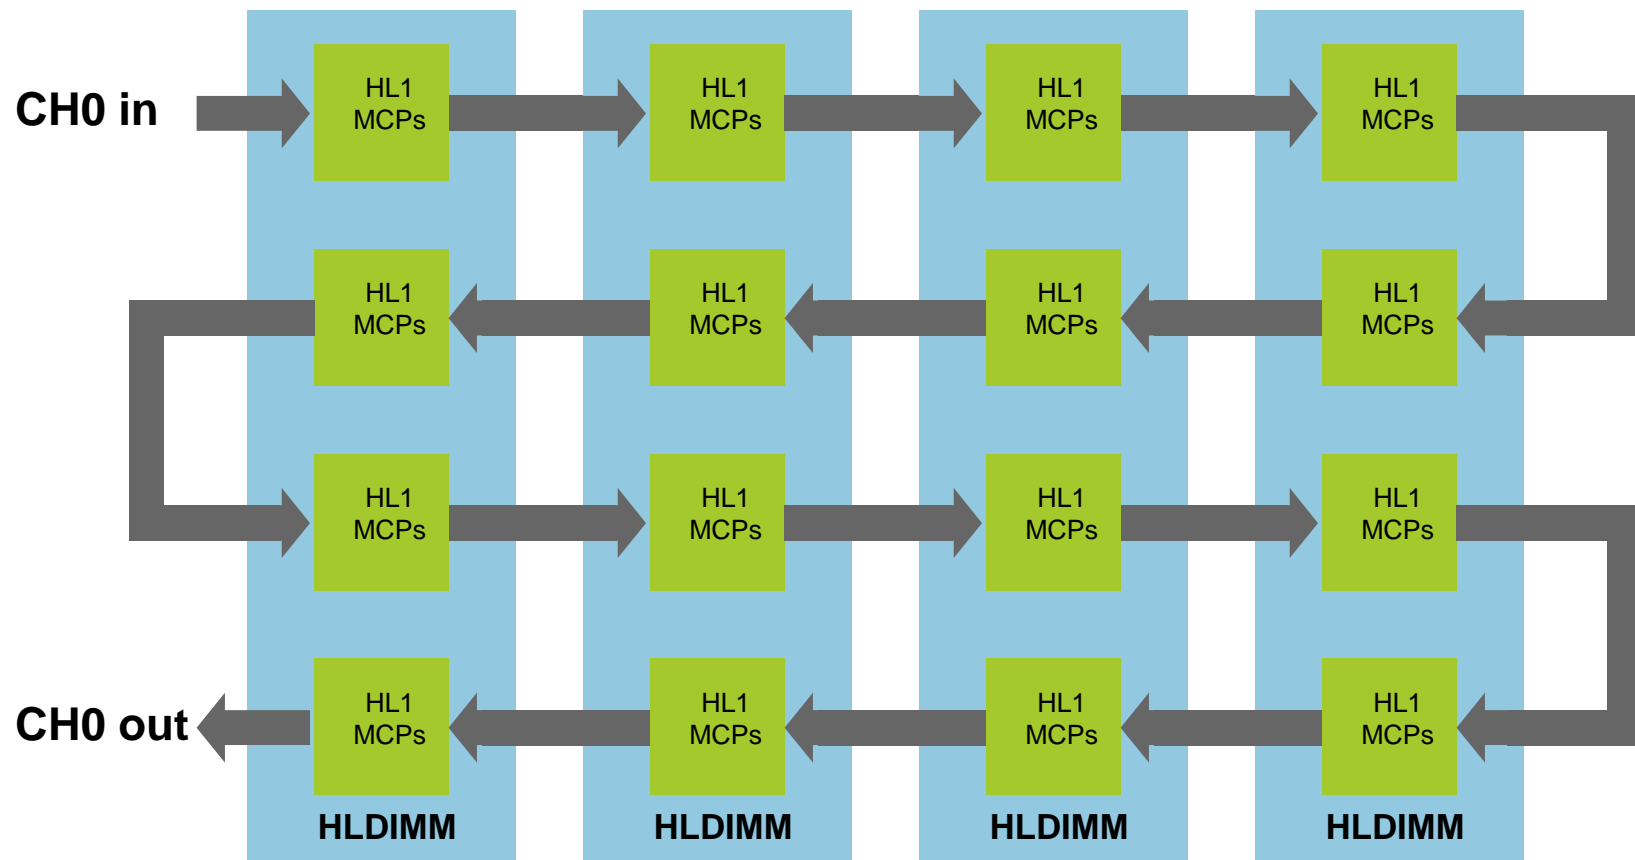
Interleaved HLDIMM Modules

- 4 x HyperLink HL1 interfaces with 1066MB/s read + 1066MB/s write = 2133MB/s aggregate throughput




HLDIMM Single Channel System

- Connect channels at first socket



Motherboard Layout

- 8 DIMMs in vertical mount sockets = 0.5TB in 60cm²



HLNAND - Enabling High Performance, Highly Scalable NAND Flash Systems

- Single HL1 channel provides simultaneous 266MB/s read and 266MB/s write throughput
- Point-to-point signaling allows up to 255 devices per channel
- Roadmap to fully backward compatible 800MB/s HL2
- Support for mission critical applications with command packet error detection
- Hierarchical HLNAND MCP using low cost commodity NAND die
- HLDIMM modules for dense, flexible and configurable NAND Flash systems