

Reconfigurable Computing

Integration

“God does not care about our mathematical difficulties. He integrates empirically.”

– Albert Einstein

Philip Leong (philip.leong@sydney.edu.au)
School of Electrical and Information Engineering

<http://phwl.org/talks>



Permission to use figures have been gained where possible. Please contact me if you believe anything within infringes on copyright.

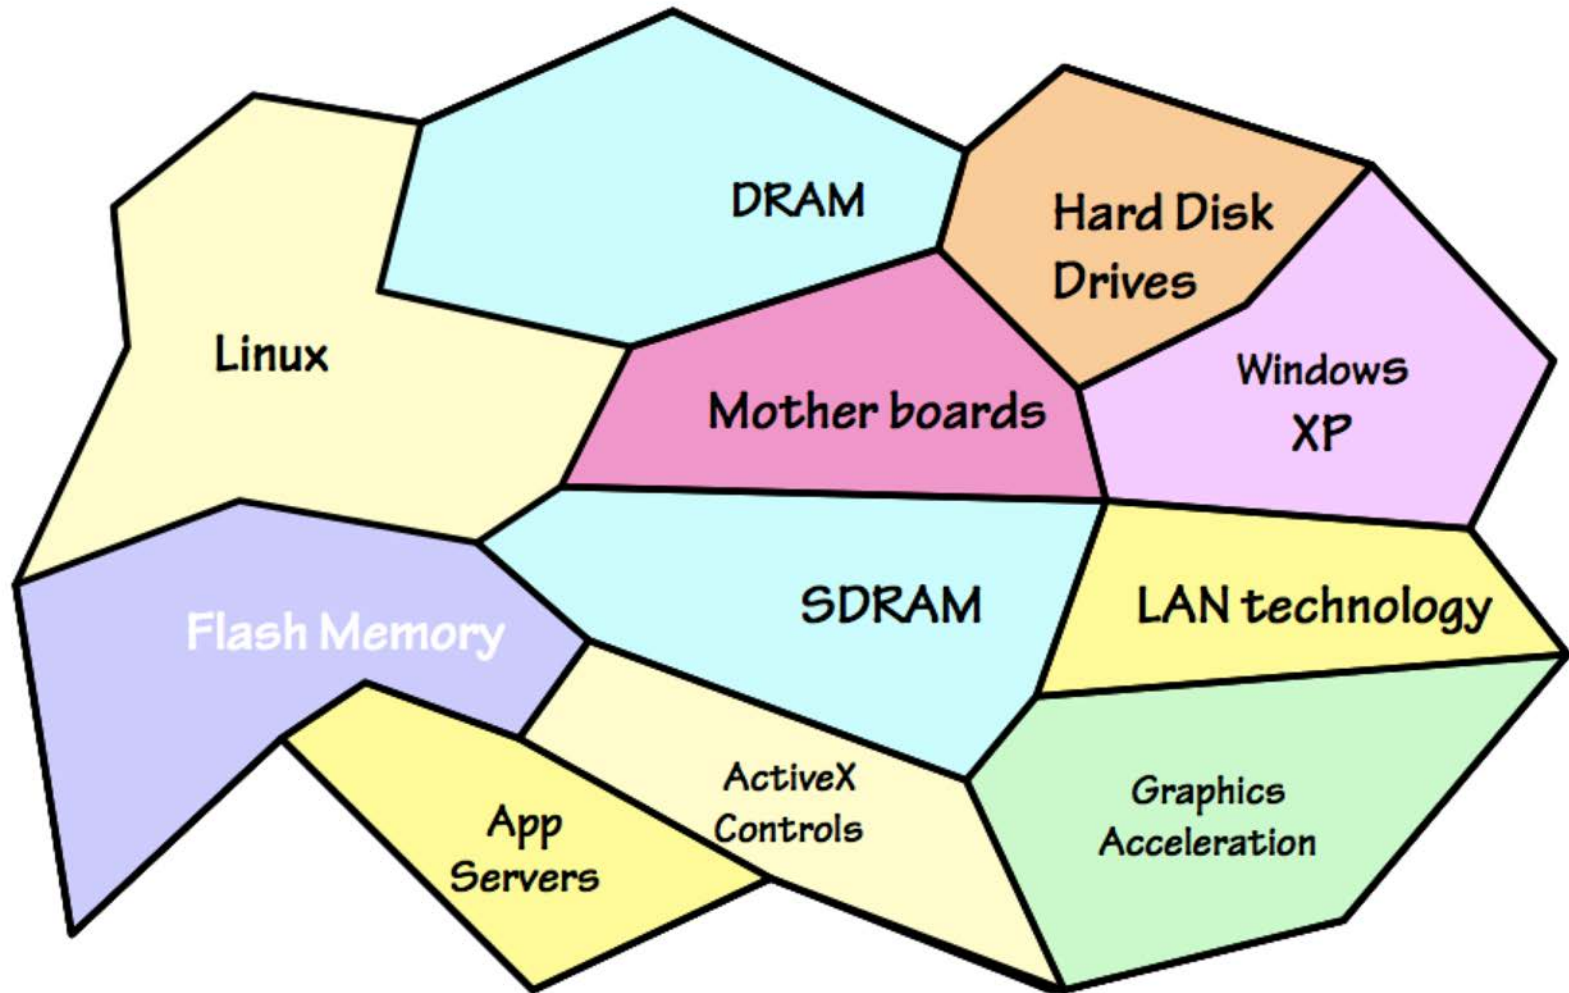
- › Bus and network principles
- › PCIe interface

Busses




Computer System Technologies

What's the most important part of this picture?




- › Interfaces typically deserve more engineering attention than the technologies they interface...
 - Abstraction: should outlast many technology generations
 - Often “virtualized” to extend beyond original function (e.g. memory, I/O, services, machines)
 - Represent more potential value to their proprietors than the technologies they connect.
- › Interface sob stories:
 - Interface “warts”: Windows “aux.c” bug, Big/little Endian wars
 - IBM PC jr
- › ... and many success stories:
 - IBM 360 Instruction set architecture; Postscript; Compact Flash; ...
 - Backplane buses


Ancient Times (Ad hoc connections)



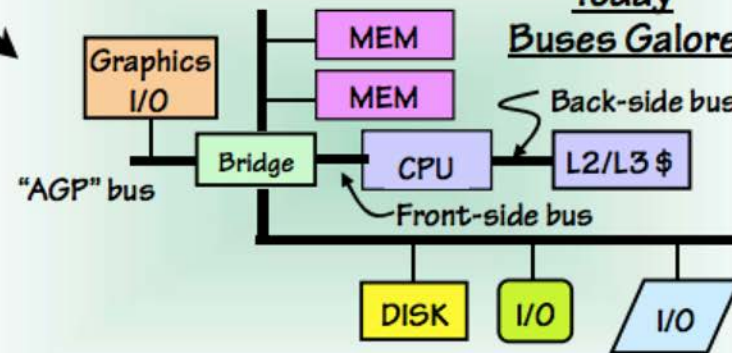
Late 60s (Processor-dependent Bus)



80s (Processor-independent Bus)



Today Buses Galore



?

Interface Standard: Backplane Bus

Modular cards that plug into
a common backplane:


CPU's

Memories

Bulk storage

I/O devices

S/W?

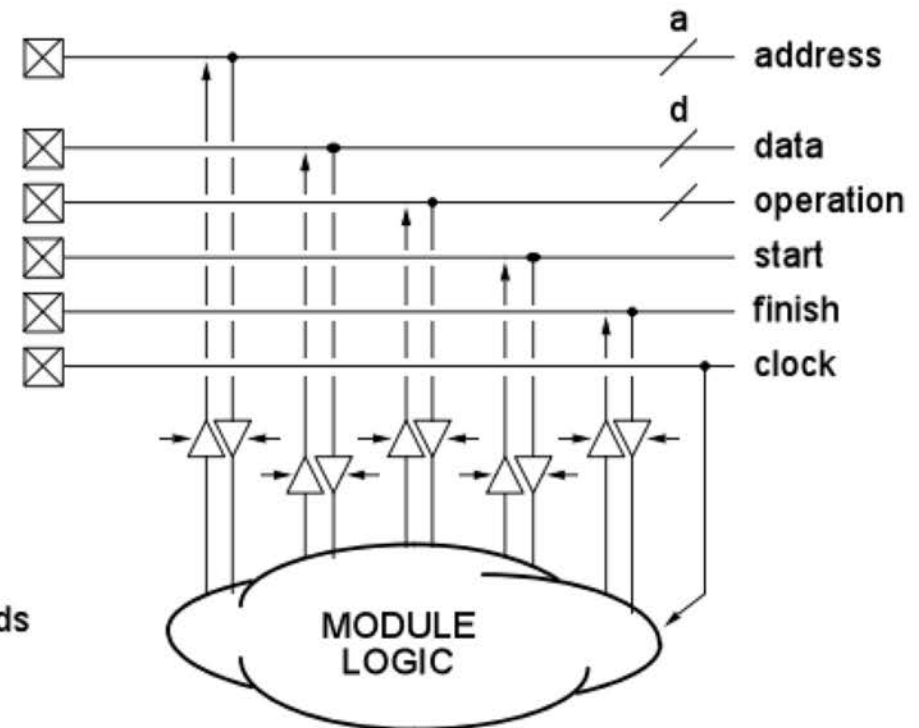


The backplane provides:

Power

Common system clock

Wires for communication



Original primitive approach --

Just take the control signals and data bus from the CPU module, buffer it, and call it a bus.

Ah, you forget,
Unibus, S-100,
SWTP S8-50,
STB, MultiBus,
Apple 2E, ...
/

ISA bus (Original IBM PC bus) -

Pin out and timing is nearly identical to the 8088 spec.

Pin	Signal	Pin	Signal
B1	Ground	A1	I/O Channel Check
B2	Reset Driver	A2	Data 7
B3	+5VDC	A3	Data 6
B4	Interrupt Request 9	A4	Data 5
B5	-VDC	A5	Data 4
B6	DMA Request 2	A6	Data 3
B7	-12 VDC	A7	Data 2
B8	Zero Wait State	A8	Data 1
B9	+12 VDC	A9	Data 0
B10	Ground	A10	I/O Channel Ready
B11	Real Memory Write	A11	Address Enable
B12	Real Memory Read	A12	Address 19
B13	Input/Output Write	A13	Address 18
B14	Input/Output Read	A14	Address 17
B15	DMA Acknowledge 3	A15	Address 16
B16	DMA Request 3	A16	Address 15
B17	DMA Acknowledge 1	A17	Address 14
B19	Refresh	A18	Address 13
B20	Clock	A19	Address 12
B21	Interrupt Request 7	A20	Address 11
B22	Interrupt Request 6	A21	Address 10
B23	Interrupt Request 5	A22	Address 9
B24	Interrupt Request 4	A23	Address 8
B25	Interrupt Request 3	A24	Address 7
B26	DMA Acknowledge 2	A25	Address 6
B27	Terminal Count	A26	Address 5
B28	Address Latch Enable	A27	Address 4
B29	+5 VDC	A28	Address 3
B30	Oscillator	A29	Address 2
B31	Ground	A30	Address 1
		A31	Address 0

<http://www.techfest.com/hardware/bus/pci.htm>

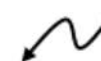
NuBus, PCI...

Isolate basic communication primitives from processor architecture:


- Simple read/write protocols
- Symmetric: any module can become “Master” (smart I/O, multiple processors, etc)
- Support for “plug & play” expansion

Goal: vendor-independent interface standard

TERMINOLOGY –

 PCI: initiator

BUS MASTER – a module that initiates a bus transaction. (CPU, disk controller, etc.)

 PCI: target

BUS SLAVE – a module that responds to a bus request. (Memory, I/O device, etc.)

BUS CYCLE – The period from when a transaction is requested until it is served.

Aren't buses simply logic circuits with long wires?

Wires: circuit theorist's view:

Equipotential "nodes" of a circuit.

Instant propagation of v, i over entire node.

"space" abstracted out of design model.

Time issues dictated by RLC elements; wires are timeless.

Wires: interconnect engineer's view:


Transmission lines.

Finite signal propagation velocity.


Space matters.

Time matters.

Reality matters.



ANALOG ISSUES:



- Propagation times

- Light travels about 1 ft / ns
(about 7"/ns in a wire)

- Skew

- Different points along the bus see the signals at different times

- Reflections & standing waves

- At each interface (places where the propagation medium changes) the signal may reflect if the impedances are not matched.
- Make a transition on a long line – may have to wait many transition times for echos to subside.

We'd like our bus to be *technology independent*...

- **Self-timed** protocols allow bus transactions to accommodate varying response times;
- **Asynchronous** protocols avoid the need to pick a (technology-dependent) clock frequency.

BUT... asynchronous protocols are vulnerable to analog-domain problems, like the infamous


WIRED-OR GLITCH: what happens when a switch is opened???




COMMON COMPROMISE: Synchronous, Self-Timed protocols


- Broadcast bus clock
- Signals sampled at "safe" times
- * DEAL WITH: noise, clock skew (wrt signals)

Synchronous Bus Clock Timing



Allow for several "round-trip" bus delays so that ringing can die down.





MASTER:

- 1) Chooses bus operation
- 2) Asserts an address
- 3) Waits for a slave to answer.

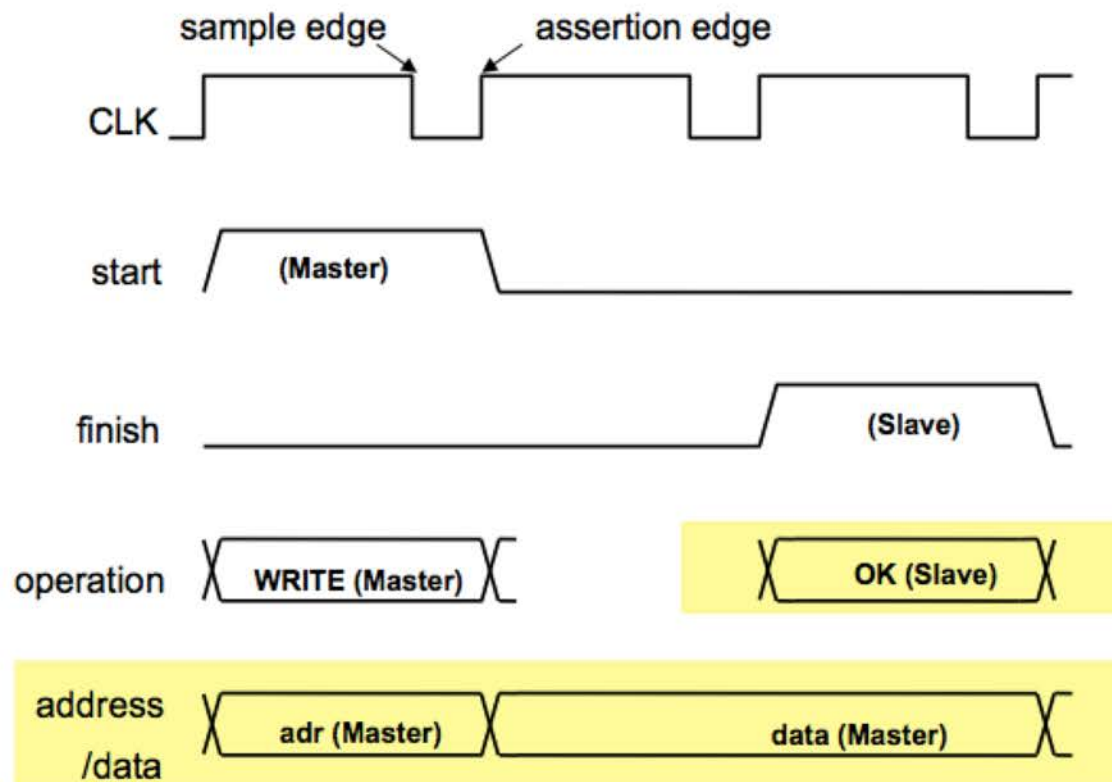
SLAVE:

- 1) Monitors start
- 2) Check address
- 3) If meant for me
 - a) look at bus operation
 - b) do operation
 - c) signal finish of cycle

BUS:

- 1) Monitors start
- 2) Start count down
- 3) If no one answers before counter reaches 0 then "time out"

Multiplexed Bus: Write Transaction: More efficient use of shared wires



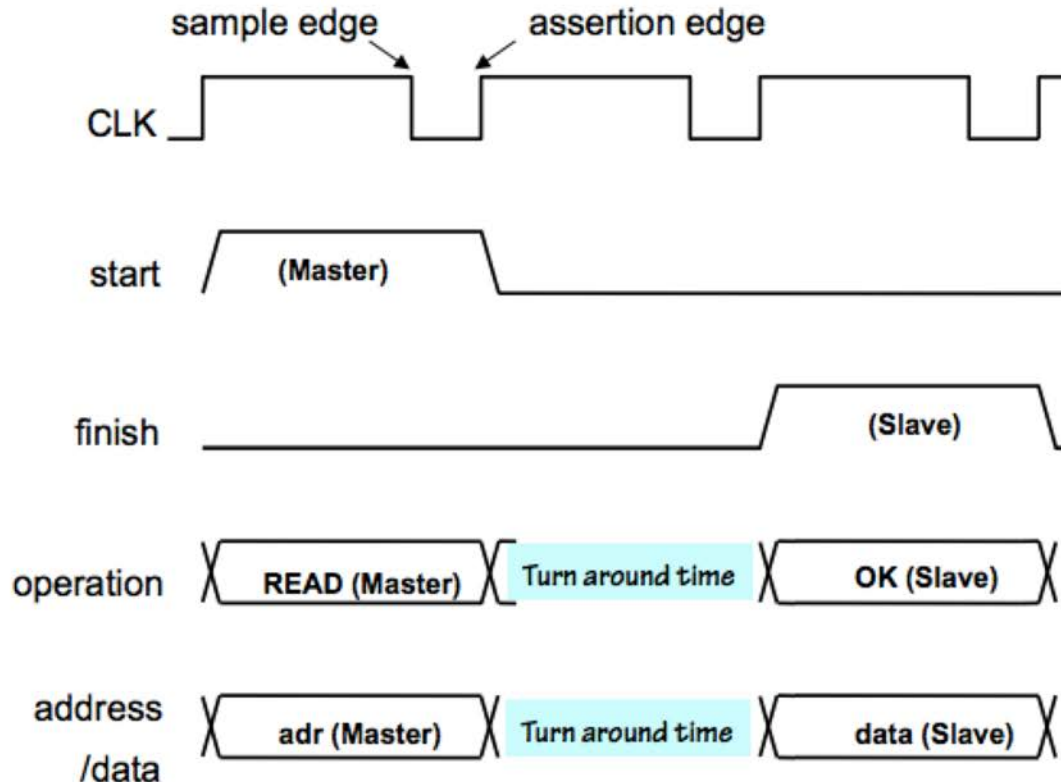
We let the address and data buses share the same wires.

Slave sends a status message by driving the operation control signals when it finishes. Possible indications:

- request succeeded
- request failed
- try again

A slave can stall the write by waiting several cycles before asserting the finish signal.

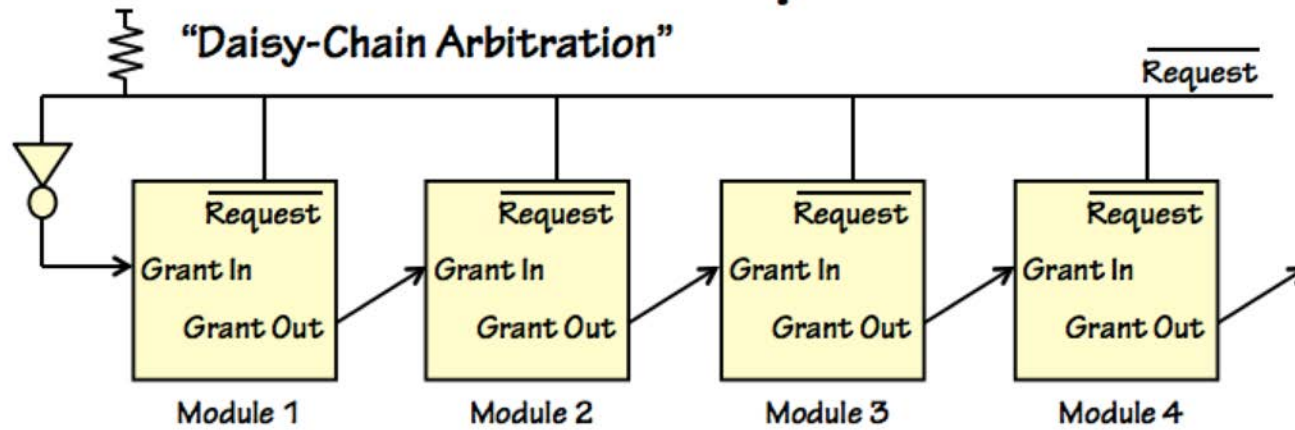
Multiplexed Bus: Read Transaction



On reads, we allot one cycle for the bus to “turn around” (stop driving and begin receiving). It generally takes some time to read data anyway.

A slave can stall the read (for instance if the device is slow compared to the bus clock) by waiting several clocks before asserting the finish signal. These delays are sometimes called “WAIT-STATES”

Throughput: 3+ Clocks/word



ISSUES:

- **Fairness** - Given uniform requests, bus cycles should be divided evenly among modules (to each, according to their needs...)
- **Bounded Wait** – An upper bound on how long a module has to wait between requesting and receiving a grant
- **Utilization** - Arbitration scheme should allow for maximum bus performance
- **Scalability** - Fixed-cost per module (both in terms of arbitration H/W and arbitration time).


STATE OF THE ART ARBITRATION: N masters, $\log N$ time, $\log N$ wires.

Networks



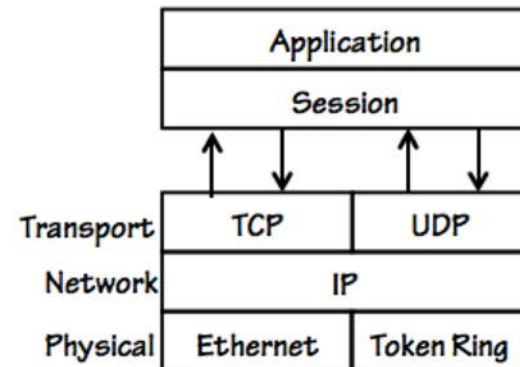
The Network as an interface standard

ETHERNET: In the mid-70's Bob Metcalf (at Xerox PARC, an MIT alum) devised a bus for networking computers together.



- Inspired by Aloha net (radio)
- COAX replaced "ether"
- *Bit-serial* (optimized for long wires)
- Variable-length "packets":
 - self-clocked data (no clock, skew!)
 - header (dest), data bits
- Issues: sharing, contention, arbitration, "backoff"

EMERGING IDEA: Protocol "layers" that isolate application-level interface from low-level physical devices:



Serial, point-to-point communications....

Becoming standard at all levels?

ETHERNET: Broadcast technology

- Sharing (contention) issues
- Multiple-drop-point issues...




Figure by MIT OpenCourseWare.

Serial point-to-point bus replacements

- Multi Gbit/sec serial links!
- PCIe, Infiniband, SATA, ...
- Packets, headers
- Switches, routing
- Trend: localized, superfast, serial networks!




Figure by MIT OpenCourseWare.


Evolution: Point-to-point

- 10BaseT, separate R & T wires
- Each link shared by only 2 hosts
- Network riddled with switches, routers

Busses and Bridges in Modern Computers

FPGA
(usually)
goes here

Serial ATA
Firewire
USB 1.1/2.0
...



PCI Express x16...
(point-to-point)

1-dimensional approaches:

"Low cost networks" – constant cost/node

BUS


ONE step for random message delivery (but
only one message at a time)



RING

$\Theta(n)$ steps for random message delivery






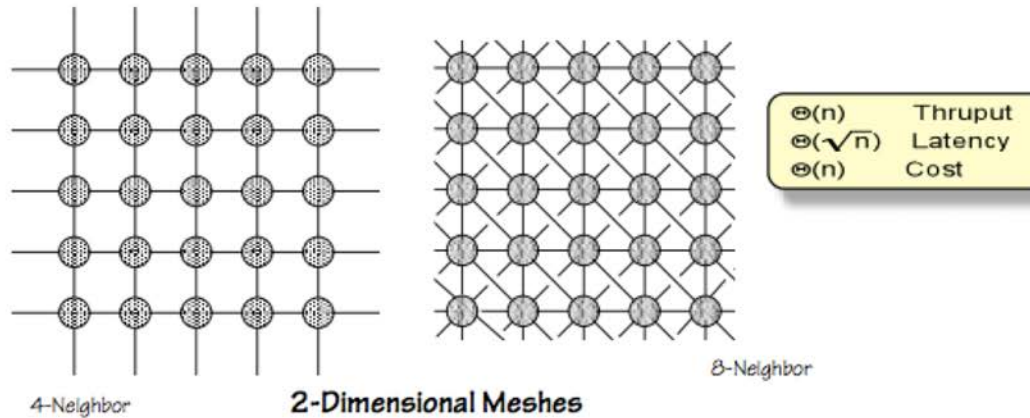
COMPLETE GRAPH:

Dedicated lines connecting each pair of communicating nodes. $\Theta(n)$ simultaneous communications.


CROSSBAR SWITCH:

- Switch dedicated between each pair of nodes
- Each A_i can be connected to one B_j at any time
- Special cases:
 - A = processors, B = memories
 - A, B same type of node
 - A, B same nodes (complete graph)






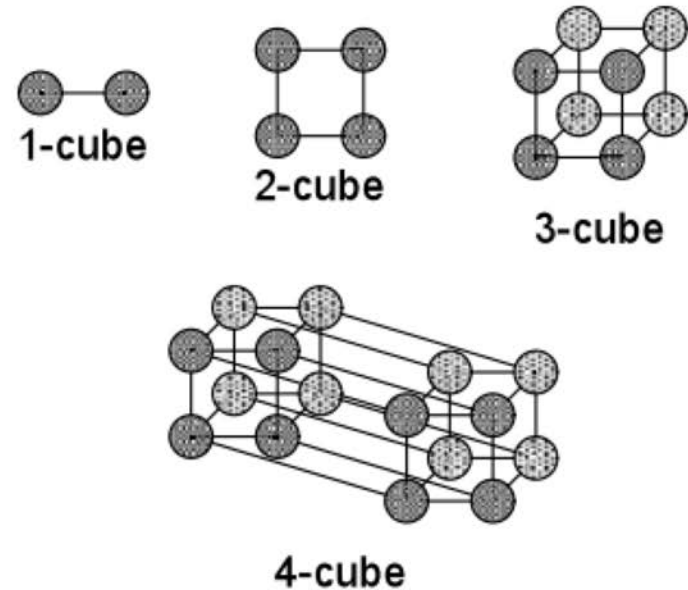
Nearest-neighbor connectivity:
Point-to-point interconnect
 - minimizes delays
 - minimizes "analog" effects
Store-and-forward
 (some overhead associated with communication routing)



HYPERCUBE (n-cube):

Cost = $\Theta(n \log n)$

Worst-case path length = $\Theta(\log n)$



BINARY TREE:

Maximum path length is $\Theta(\log n)$ steps;

Cost/node constant.

Communication Topologies: Latency

Theorist's view:

- Each point-to-point link requires one hardware unit.
- Each point-to-point communication requires one time unit.

Topology	\$	Theoretical Latency	Actual Latency
Complete Graph	$\Theta(n^2)$	$\Theta(1)$	$\geq \Theta(\sqrt[3]{n})$
Crossbar	$\Theta(n^2)$	$\Theta(1)$	$\Theta(n)$
1D Bus	$\Theta(n)$	$\Theta(1)$	$\Theta(n)$
2D Mesh	$\Theta(n)$	$\Theta(\sqrt{n})$	
3D Mesh	$\Theta(n)$	$\Theta(\sqrt[3]{n})$	
Tree	$\Theta(n)$	$\Theta(\log n)$	$\geq \Theta(\sqrt[3]{n})$
N-cube	$\Theta(n \log n)$	$\Theta(\log n)$	$\geq \Theta(\sqrt[3]{n})$

IS IT REAL?

- Speed of Light: ~ 1 ns/foot (typical bus propagation: 5 ns/foot)
- Density limits: can a node shrink forever? How about Power, Heat, etc ... ?

OBSERVATION: Links on Tree, N-cube must grow with n ; hence time/link must grow.

The Old Standbys:

- In box: Backplane buses: parallel, shared data paths
 - Arbitration, skew problems
- Local area: shared, single “ether” cable
 - Contention, collisions

New “switched fabric” tech (in & out of box):


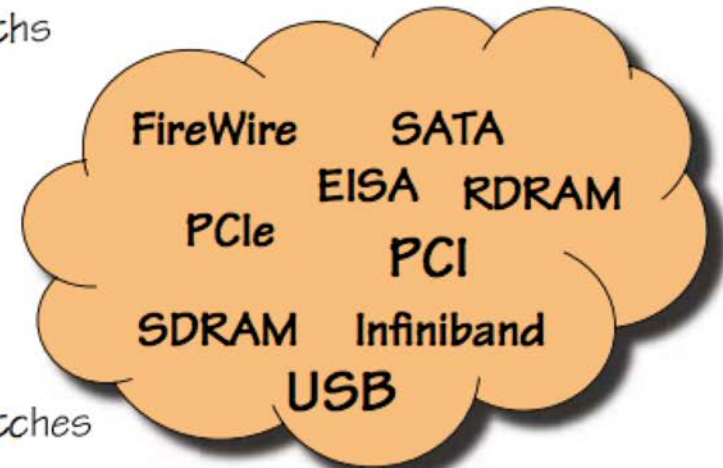
- Shared wires replaced by point-to-point serial
- Parallel data paths replaced by serial “packets”
- Communication network extended via active switches

Topological Invariants:

- Asymptotic performance/cost tradeoffs...
- Log-latency topologies: a useful fiction
- Best-case scaling with 3D mesh

Watch this space!

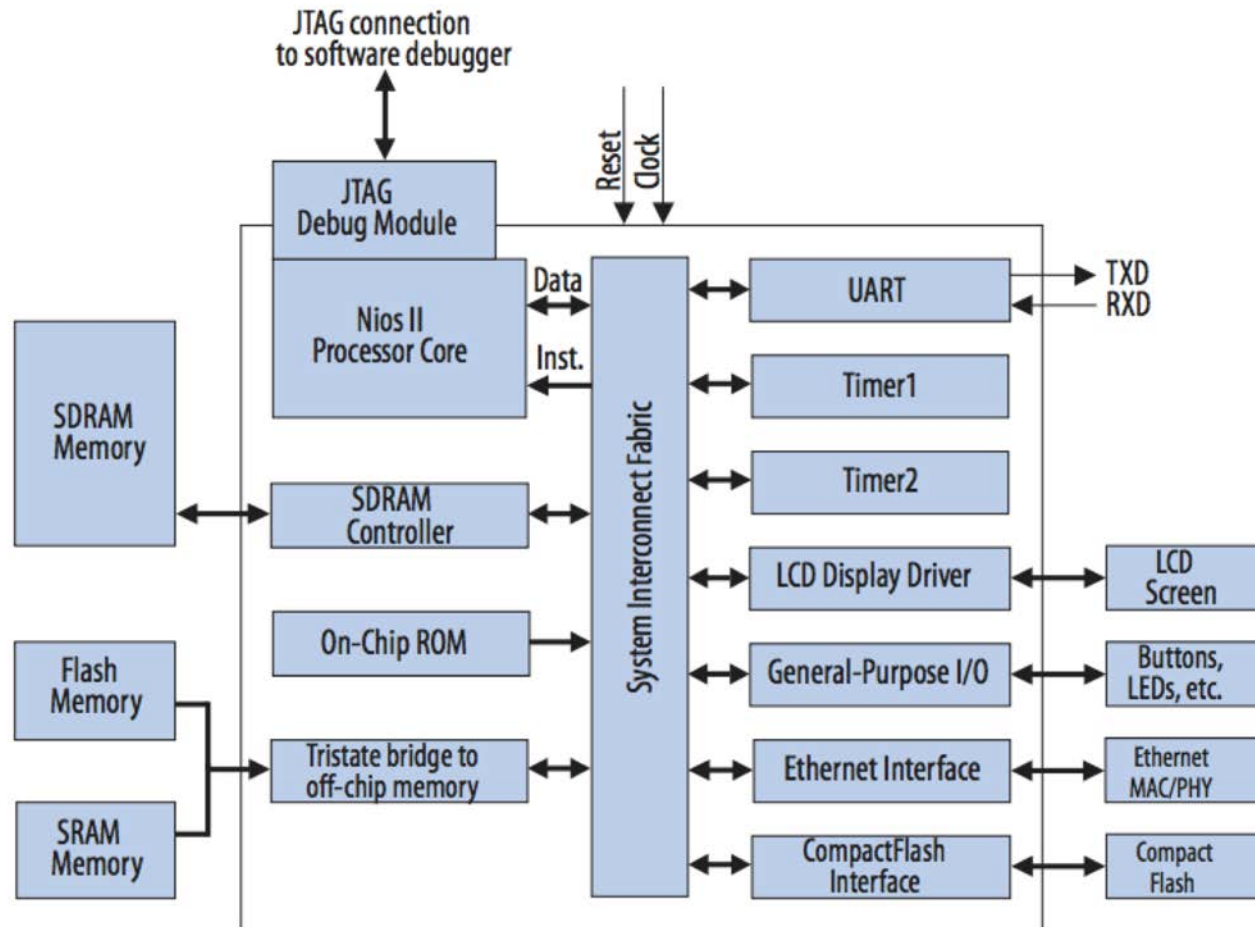
- Technologies: optical, proximity,
- 3D packaging, interconnect
- ???



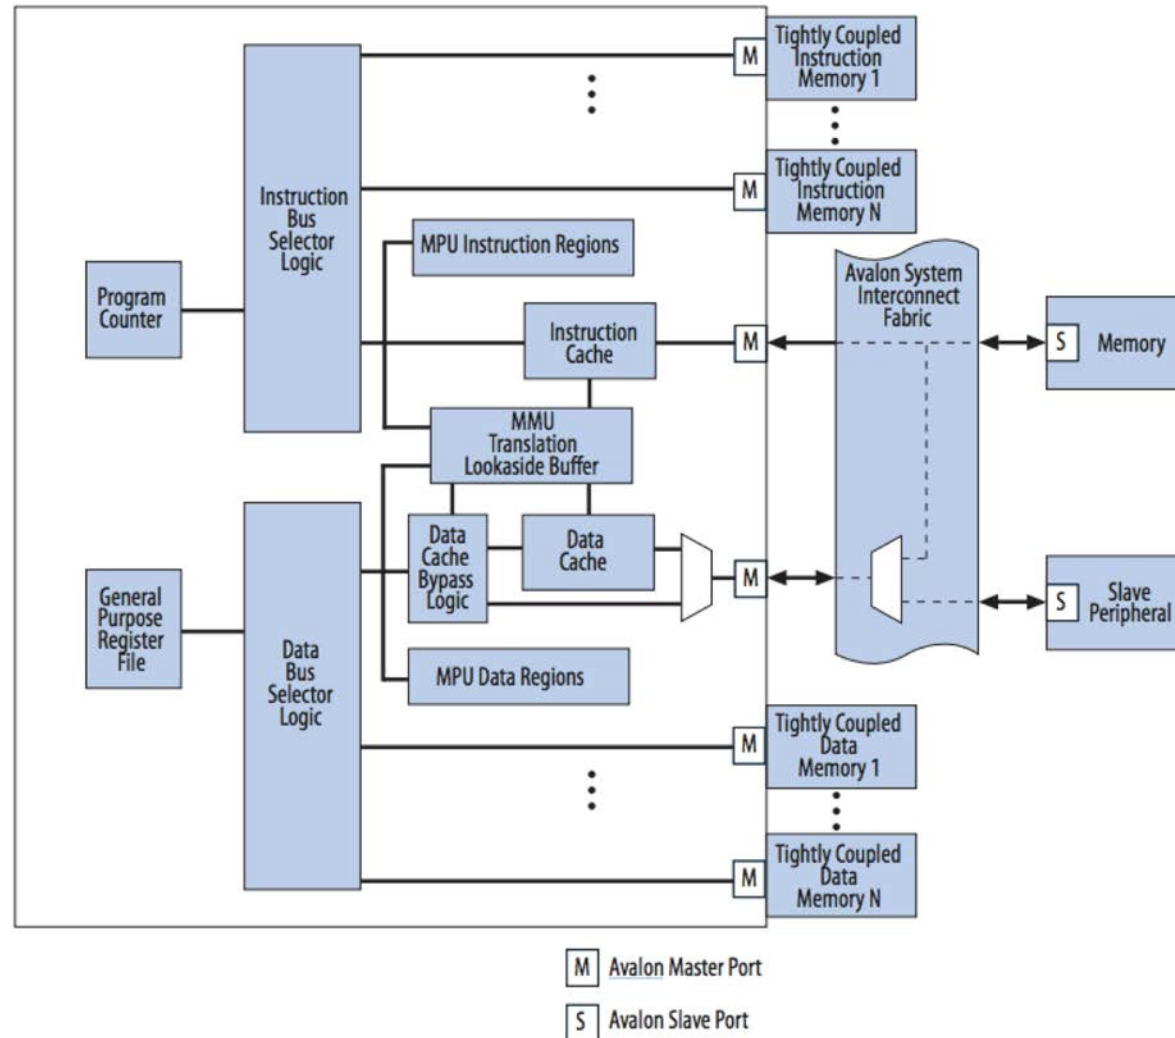
Altera NIOS Processor



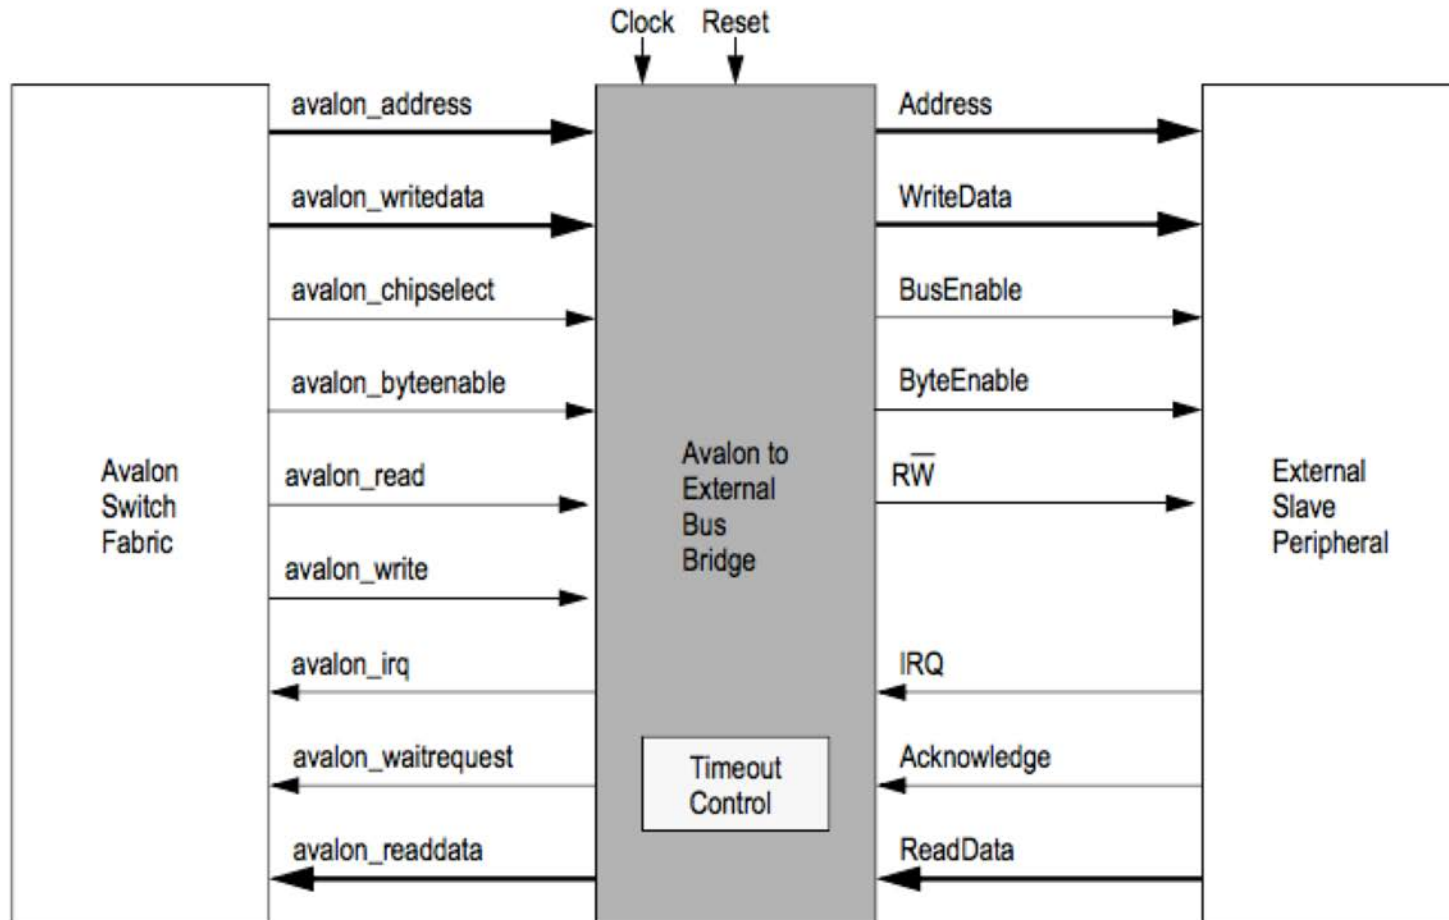
- › NIOS II Classic is a soft processor from Altera
 - Has floating point, MMU, caching,
 - Uses the Avalon bus
 - Excellent supporting tools
 - SOPC Builder generates designs from parameters such as data width, address range
 - › Describe Avalon to External Bus Bridge
-



NIOS II Memory and I/O Organisation

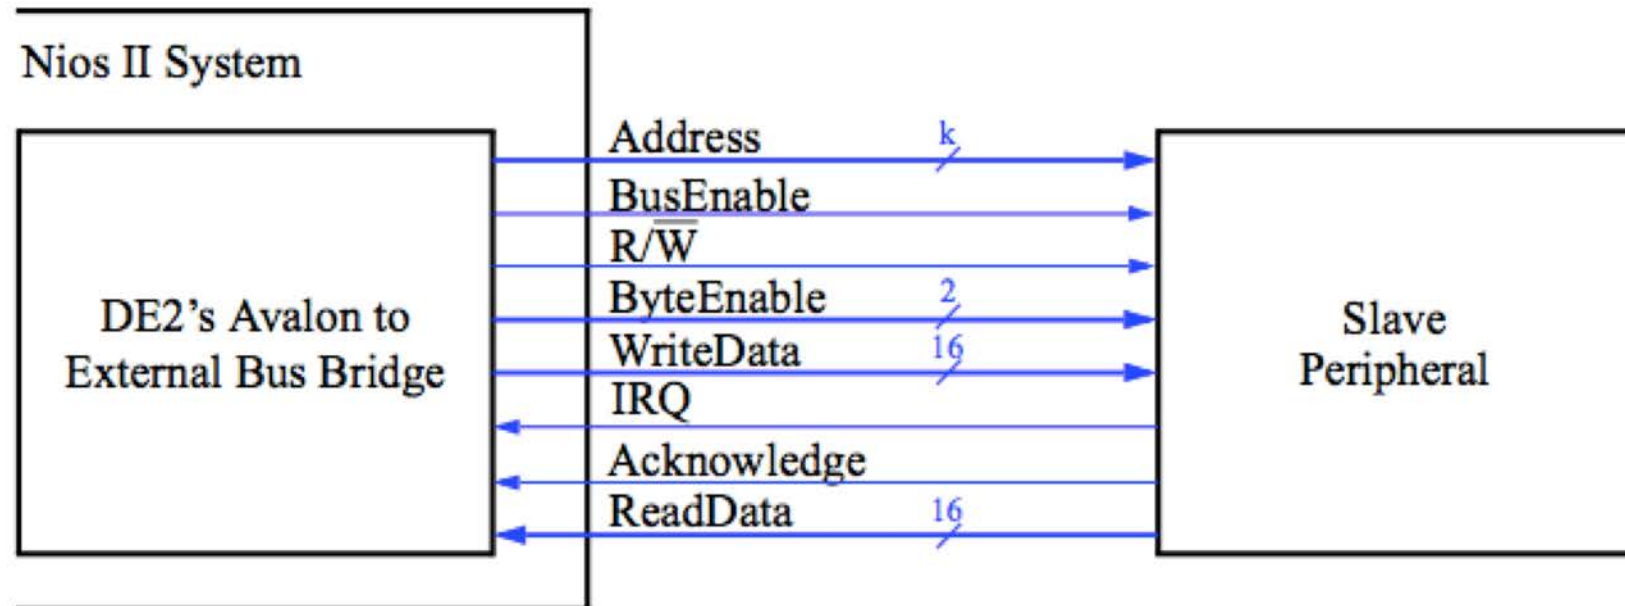


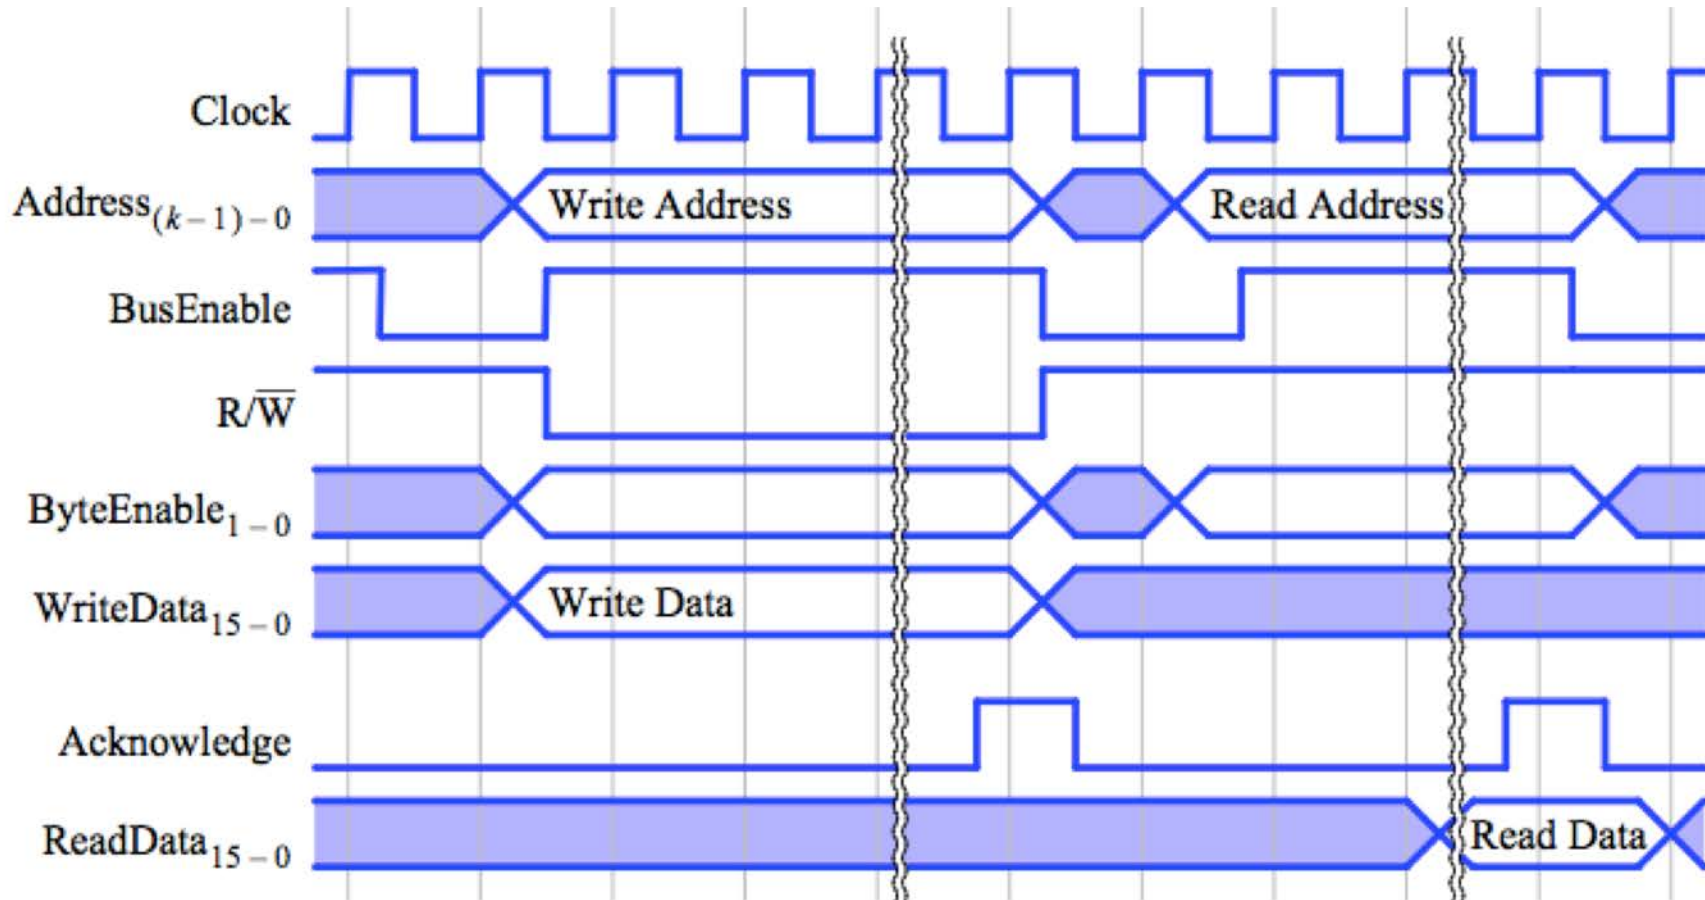
Avalon to External Bus Bridge



- › Address—kbits (up to 32).
- › BusEnable—1bit. Indicates that all other signals are valid, and a data transfer should occur.
- › RW—1bit. Read(1), Write(0)
- › ByteEnable — 16, 8, 4, 2 or 1 bits. Each bit indicates whether or not the corresponding byte should be read or written. These signals are active high.
- › WriteData — 128, 64, 32, 16 or 8 bits. The data to be written to the peripheral device during a Write transfer.
- › Acknowledge — 1 bit. Used by the peripheral device to indicate that it has completed the data transfer.
- › ReadData—128, 64, 32, 16. Data read from peripheral during a Read transfer.
- › IRQ—1bit.Used by peripheral device to interrupt the processor.

- › Synchronous bus – all transfers occur on rising edge of clock
- › Bus has time out





- › <http://wl.altera.com/literature/lit-nio2.jsp>
- › https://www.altera.com/content/dam/altera-www/global/en_US/pdfs/literature/manual/mnl_avalon_spec.pdf