

Multicore Hardware Abstraction Items Specification (v1.2)

Kiyeon Lee, Charmgil Hong, and Sangyeun Cho
Dept. of Computer Science
University of Pittsburgh
lee, chh50, cho@cs.pitt.edu

June 18, 2008

Contents

1	Introduction	2
2	HABS Items	3
2.1	System level hardware description	3
2.2	Processor chip level hardware description	3
2.3	Memory chip level hardware description	4
2.4	Chip connectivity hardware description	5
2.5	Core cluster level hardware description	6
2.6	Hardware IP description	6
2.7	On-chip memory hardware description	7
2.8	Processor Core hardware description	8
2.9	In-cluster memory hardware description	9
2.10	On-chip connectivity	9
2.11	In-core memory hardware description	10

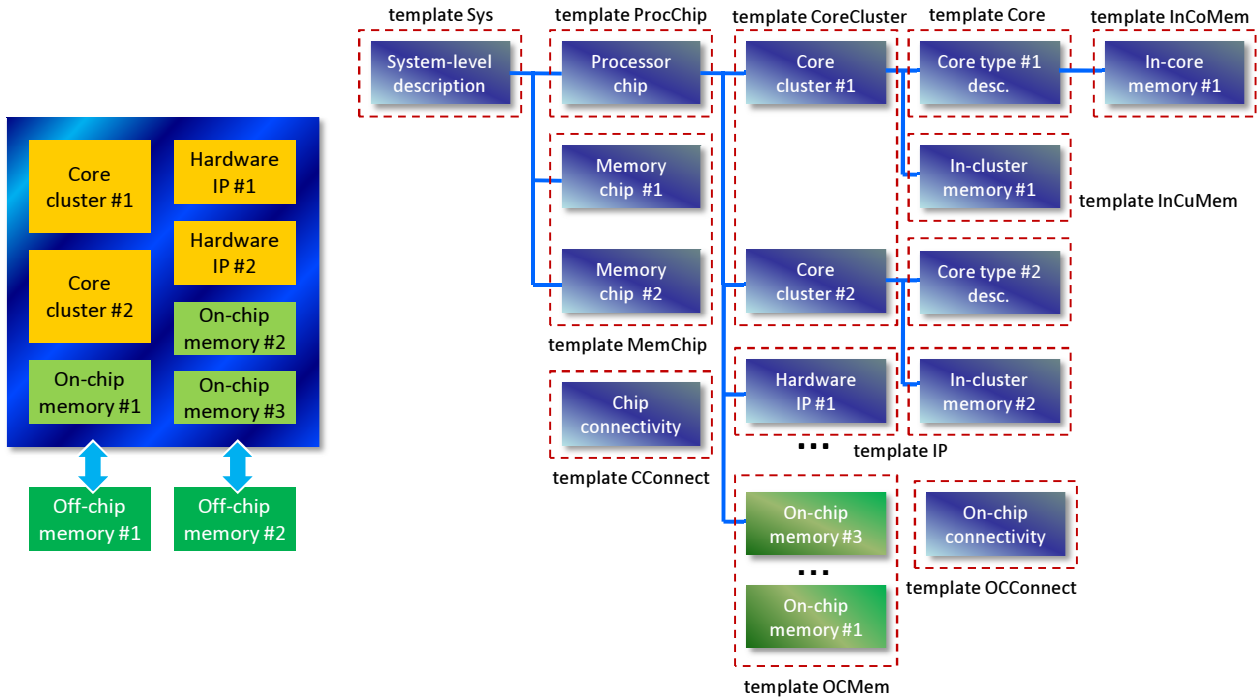


Figure 1: (a) HABS system model. (b) The hierarchy of HABS hardware descriptions. Templates for actual HABS description are highlighted.

1 Introduction

In this report we present a specification of the hardware abstraction items to be included in the proposed multicore hardware abstraction (HABS). Figure 1(a) shows the system model we assume and how the machine information is organized in a hierarchical manner. The system model is at the board level and is composed of processor chips, memory chips, and their connection.

To describe different hardware components in a convenient and efficient manner, we define a number of *hardware description templates*. In the current specification (v1.1), we have defined 11 templates as shown in Figure 1(b). They are: template for system-level description (Sys), template for processor chip (ProcChip), template for memory chip (MemChip), template for chip-level connectivity (CConnect), template for processor core cluster (CoreCluster), template for hardware IP (IP), template for on-chip memory (OCMem), template for processor core description (Core), template for in-cluster memory (InCuMem), template for on-chip connectivity (OCCConnect), and template for in-core memory (InCoMem). Some templates contain structural information while some others not. For instance, the MemChip, OCMem, IP, InCuMem, and InCoMem templates are descriptive of what they model, but do not contain structural information such as links to other objects. On the other hand, the Sys, ProcChip, CConnect, CoreCluster, Core, and OCCConnect templates are predominantly structure-oriented.

The actual descriptions in each template are classified into three categories depending on their objective: (i) functional; (ii) performance-related; and (iii) power-related. “Functional” descriptions are concerned with the structure of hardware (*i.e.*, how processing and memory components in the system are connected via interconnection components) and the behavioral aspects of each component. “Performance” descriptions are for associating

with a hardware component properties that help compute related performance (*e.g.*, data width and clock for a bus interconnect). Similarly, “power” descriptions are for computing power consumption of hardware components.

HABS items provide the lowest-level information for the HABS API (Application Programming Interface). That is, the low-level HABS API functions retrieve the information associated with the corresponding HABS items in a machine description document (MADD). Furthermore, the HABS API defines a set of higher-level functions that combine a few low-level HABS API functions and derive the target system’s performance properties. For instance, one may obtain time needed for synchronizing a set of processor cores by calling one of the high-level API functions.

In the following section, we will list the HABS items. The HABS API will be defined in a separate document. The actual description of a MADD is done using an XML (eXtensible Markup Language) format. The XML description formats are again defined in a separate document.

2 HABS Items

2.1 System level hardware description

2.1.1 Items for functional description

In the following items, the information for the ones with a remark (*) can be either included within the template or can be derived from the lower-level objects that are specified within the template.

1. Name: the name of the system
2. NProcChips: the number of processor chips
3. NMemChips: the number of memory chips
4. NProcCores: the total number of processor cores
5. MemCapacity: the total memory capacity in the system (in bytes)

2.1.2 Items for performance

1. MaxSystemPerf(ClockFreq)*: the total maximum performance (in MIPS) of the system given the clock frequency ClockFreq (in Hz) for the processor cores inside the system

2.1.3 Items for power

1. SystemPower: the system power (Min, Avg, Max) (in W/MHz)

2.2 Processor chip level hardware description

2.2.1 Items for functional description

1. Name: the name of the processor chip
2. NClusters: the number of processor core clusters
3. ClusterName(ClusterNum)*: the name of the cluster ClusterNum
4. CoreType(ClusterNum)*: the description of a core in the cluster ClusterNum

5. NHWIPs: the number of hardware IPs
6. HWIPName(HWIPNum)*: the name of an IP specified by IPNum
7. NOCMem: the number of on-chip memory
8. OCMemName(OCMemNum)*: the name of the memory indexed by OCMemNum
9. OCMemType(OCMemNum)*: the type of the memory indexed by OCMemNum ¹
10. NDMA*: the number of DMA units in the cluster
11. DMASrc(DMANumber)*: the list of source nodes of the DMA unit indexed by DMANumber
12. DMADst(DMANumber)*: the list of destination nodes of the DMA unit indexed by DMANumber

2.2.2 Items for performance

1. MaxSpeed: the maximum clock speed (in Hz)
2. MaxChipPerf(ClockFreq)*: the total maximum chip performance (in MIPS) given the clock frequency ClockFreq (in Hz) for the processor cores inside the processor chip

2.2.3 Items for power

1. ChipPower: the chip power (Min, Avg, Max) (in W/MHz)
2. ChipPowerManagement: the description of the power management mechanism in the chip: {None, Chip-Level DVFS, Cluster-Level DVFS, Core-Level DVFS}
3. LeakageCurrent: the leakage current (in A)
4. PowerDownMode: is there a power-down mode?
5. PowerDownCurrent: the current in the power-down mode (in A)
6. Voltage: voltage of the chip (in V)

2.3 Memory chip level hardware description

2.3.1 Items for functional description

1. Name: the name of the memory chip
2. Type: {SRAM, SDRAM, DDR SDRAM, One Nand, NOR flash, NAND flash, Mask ROM}
3. Capacity: the number of bytes in the memory chip
4. NBanks: the number of banks
5. NRows: the number of memory rows in a bank
6. NColumns: the number of columns in a bank
7. WordWidth: the number of bits in a memory word

¹Many attributes of an on-chip memory can be derived from the on-chip memory description given OCMemNum. Refer to Section 2.7 for these attributes.

8. AddrWidth: the address width (in bits)
9. DataWidth: the data width (in bits)

2.3.2 Items for performance

1. MinCycle: the minimum cycle time (in s)
2. MinLatency: the minimum latency (cycle count) in terms of MinCycle^2

2.3.3 Items for power

1. EnergyRead: the energy (in J) per read access
2. EnergyWrite: the energy (in J) per write access
3. LeakageCurrent: the leakage current (in A)
4. PowerDownMode: is there a power-down mode?
5. PowerDownCurrent: the current in the power-down mode (in A)
6. Voltage: voltage of the chip (in V)

2.4 Chip connectivity hardware description

2.4.1 Items for functional description

1. Name: the name of the chip connectivity structure
2. Type: the chip connection type: {Flat Bus, Hierarchical Bus, Crossbar}
3. NChipConnections: the number of chip-to-chip connection pairs
4. AddrWidth(ChipConnectionNum): the chip connection address width (in bits)
5. DataWidth(ChipConnectionNum): the chip connection data width (in bits)
6. Shared(ChipConnectionNum): is the connection ChipConnectionNum shared?
7. DegreeSharing(ChipConnectionNum): the degree of sharing in the connection ChipConnectionNum

2.4.2 Items for performance

1. Speed(ChipConnectionNum): the speed (in Hz) of the connection ChipConnectionNum

2.4.3 Items for power

1. LoadCap(ChipConnectionNum): the loading capacitance of the connection ChipConnectionNum (in F)
2. Voltage(ChipConnectionNum): the operating voltage of the connection ChipConnectionNum (in V)

²To compute actual latency correctly, we may need to expand the parameter set.

2.5 Core cluster level hardware description

2.5.1 Items for functional description

1. Name: the name of the core cluster
2. Type*: the processor core type
3. MPSSupport: support and type for multiprocessing: {Not Supported, Shared Address, Non-Shared Address}
4. NProcCores: the number of processor cores in the cluster
5. OSRunnable: if an OS may run on a processor in the cluster
6. NetworkTopology: the network type in the cluster: {Bus, Mesh, Ring, Crossbar}
7. CacheCoherent: is the network cache coherent?
8. NDMA: the number of DMA units in the cluster
9. DMASrc(DMANumber): the list of source nodes of the DMA unit indexed by DMANumber
10. DMADst(DMANumber): the list of destination nodes of the DMA unit indexed by DMANumber

2.5.2 Items for performance

1. MaxBW: the maximum bandwidth requirements of the cluster (in B/s)
2. MaxSpeed: the maximum speed of the cluster (in Hz)
3. ClusterPerf(ClockFreq)*: the aggregate performance of the cluster (in MIPS) given the clock frequency ClockFreq (in Hz) for the processor core in the cluster

2.5.3 Items for power

1. PowerManagement: the description of the power management mechanism in the cluster: {None, Cluster-Level DVFS, Core-Level DVFS}
2. Voltage: voltage of the chip (in V)
3. PowerDownMode: is there a power-down mode?

2.6 Hardware IP description

2.6.1 Items for functional description

1. Name: the name of the hardware IP
2. DataInput: the hardware component providing data
3. DataOutput: the hardware component to send data

2.6.2 Items for performance

1. DataInputBW: the input bandwidth (in bits)
2. DataOutputBW: the output bandwidth (in bits)
3. DataProcessingLat: the data processing latency

2.6.3 Items for power

1. MinClockFreq: minimum clock frequency (in Hz)
2. MinVoltage: the voltage (in V) when minimum clock frequency is used
3. MaxClockFreq: maximum clock frequency (in Hz)
4. MaxVoltage: the voltage (in V) when maximum clock frequency is used
5. Voltage(ClockFreq): voltage (in V) at the clock frequency ClockFreq (in Hz)
6. Capacitance: the capacitance value (in F)
7. LeakageCurrent: the leakage current (in A)
8. PowerDownMode: is there a power-down mode?
9. PowerDownCurrent: the current in the power-down mode (in A)
10. Power(ClockFreq): power consumption (in W) at the clock frequency ClockFreq

2.7 On-chip memory hardware description

2.7.1 Items for functional description

1. Name: the name of the on-chip memory
2. Type: the memory type: {SRAM, NOR, NAND, E-DRAM, Mask ROM, Shared Cache}
3. Capacity: the memory capacity (in bytes)
4. NBanks: the bank count
5. NRows: the number of memory rows in a bank
6. NColumns: the number of columns in a bank
7. WordWidth: the number of bits in a memory word
8. NSets: the number of sets if the on-chip memory is a cache
9. N Ways: the number of ways (associativity) if the on-chip memory is a cache
10. BlockSize: the block size (in bytes) if the on-chip memory is a cache
11. ReplacementPolicy: the replacement policy if the on-chip memory is a cache: {LRU, FIFO, Random, Round-Robin}

2.7.2 Items for performance

1. AccessLatency: the memory access latency (the number of cycles based on on-chip connectivity speed)

2.7.3 Items for power

1. EnergyRead: the energy per read access (in J)
2. EnergyWrite: the energy per write access (in J)
3. LeakageCurrent: the leakage current (in μA)

2.8 Processor Core hardware description

2.8.1 Items for functional description

1. Name: the name of the processor core
2. Type: the basic type of the processor core: {In-Order, Out-Order, VLIW, RP}
3. NThreads: the number of simultaneous thread contexts
4. NIssueWidth: the number of instructions issuable in a cycle
5. NFUTypes: the number of functional units
6. FUTypeName(FUTypeNum): the name of a functional unit type
7. NInCoMem: the number of in-core memory
8. ICache: is there an i-cache?
9. DCache: is there an d-cache?

2.8.2 Items for performance

1. NominalIPC: the nominal IPC value
2. NMemPort: the number of used memory ports: {1: Von Neumann Architecture, 2: Harvard Architecture}
3. MIPS(ClockFreq): the number of MIPS (Million Instructions per second) given the clock frequency ClockFreq (in Hz)

2.8.3 Items for power

1. MinClockFreq: minimum clock frequency (in Hz)
2. MinVoltage: the voltage (in V) when minimum clock frequency is used
3. MaxClockFreq: maximum clock frequency (in Hz)
4. MaxVoltage: the voltage (in V) when maximum clock frequency is used
5. Voltage(ClockFreq): voltage at the clock frequency ClockFreq (in Hz)
6. Capacitance: the capacitance value (in F)

7. LeakageCurrent: the leakage current (in A)
8. PowerDownMode: is there a power-down mode?
9. PowerDownCurrent: the current in the power-down mode (in A)
10. Power(ClockFreq): power consumption (in W) at the clock frequency ClockFreq (in Hz)

2.9 In-cluster memory hardware description

2.9.1 Items for functional description

1. Name: the name of the in-cluster memory
2. Type: the memory type: {NAND, cache}
3. Capacity: the memory capacity (in bytes)
4. NBanks: the bank count
5. NRows: the number of memory rows in a bank
6. NColumns: the number of memory columns in a bank
7. WordWidth: the number of bits in a memory word
8. NSets: the number of sets if the in-cluster memory is a cache
9. N Ways: the number of ways (associativity) if the in-cluster memory is a cache
10. BlockSize: the block size (in bytes) if the in-cluster memory is a cache
11. ReplacementPolicy: the replacement policy if the in-cluster memory is a cache: {LRU, FIFO, Random, Round-Robin}

2.9.2 Items for performance

1. AccessLatency: the memory access latency (the number of cycles based on the speed of core cluster)

2.9.3 Items for power

1. EnergyRead: the energy per read access (in J)
2. EnergyWrite: the energy per write access (in J)
3. LeakageCurrent: the leakage current (in A)

2.10 On-chip connectivity

One might want to define on-chip interconnection structure between processors within a cluster or among the clusters on a processor chip. When complete descriptions are available, the property of a connection between on-chip objects, such as processor cores, on-chip memory, and hardware IPs, can be derived. The template for on-chip connectivity is used to describe both intra-cluster connectivity and inter-cluster connectivity.

2.10.1 Items for functional description

1. Name: the name of the on-chip connectivity structure
2. Topology: the connectivity topology: {Bus, Mesh, Crossbar, Ring}
3. AddrWidth: the chip connection address width (in bits)
4. DataWidth: the chip connection data width (in bits)
5. CacheCoherent: is the network cache coherent?
6. NConnections: the number of one-to-one connections
7. AddrWidth(ConnectionNum): the address width (in bits) of an one-to-one connection
8. DataWidth(ConnectionNum): the data width (in bits) of an one-to-one connection
9. Speed(ConnectionNum): the speed of an one-to-one connection (in Hz)

2.10.2 Items for performance

1. Speed: the relative speed (clock) of the interconnect to the cluster
2. SrcOverhead: the overhead (# of cycles) at the network source
3. DestOverhead: the overhead (# of cycles) at the network destination
4. SwitchDelay: the delay (# of cycles) at intermediate switches

2.10.3 Items for power

1. MinClockFreq: minimum clock frequency (in Hz)
2. MinVoltage: the voltage (in V) when minimum clock frequency is used
3. MaxClockFreq: maximum clock frequency (in Hz)
4. MaxVoltage: the voltage (in V) when maximum clock frequency is used
5. Voltage(ClockFreq): voltage at the clock frequency ClockFreq (in Hz)
6. Capacitance: the capacitance value (in F)
7. LeakageCurrent: the leakage current (in A)
8. PowerDownMode: is there a power-down mode?
9. PowerDownCurrent: the current in the power-down mode (in A)
10. Power(ClockFreq): power consumption (in W) at the clock frequency ClockFreq (in Hz)

2.11 In-core memory hardware description

2.11.1 Items for functional description

1. Name: the name of the in-core memory
2. Type: the memory type: {NAND, Cache}

3. Capacity: the memory capacity (in bytes)
4. NBanks: the bank count
5. NRows: the number of memory rows in a bank
6. NColumns: the number of memory columns in a bank
7. WordWidth: the number of bits in a memory word
8. NSets: the number of sets if the in-core memory is a cache
9. N Ways: the number of ways (associativity) if the in-core memory is a cache
10. BlockSize: the block size (in bytes) if the in-core memory is a cache
11. ReplacementPolicy: the replacement policy if the in-clustrer memory is a cache: {LRU, FIFO, Random, Round-Robin}

2.11.2 Items for performance

1. AccessLatency: the memory access latency (the number of cycles based on the speed of processor core)

2.11.3 Items for power

1. EnergyRead: the energy per read access (in J)
2. EnergyWrite: the energy per write access (in J)
3. LeakageCurrent: the leakage current (in A)