

SystemC abstractions and design refinement for HW-SW SoC design

© CoWare, Inc. 2001

Dündar Dumlugöl

Vice President of Engineering,
CoWare, Inc.




Overview

© CoWare, Inc. 2001

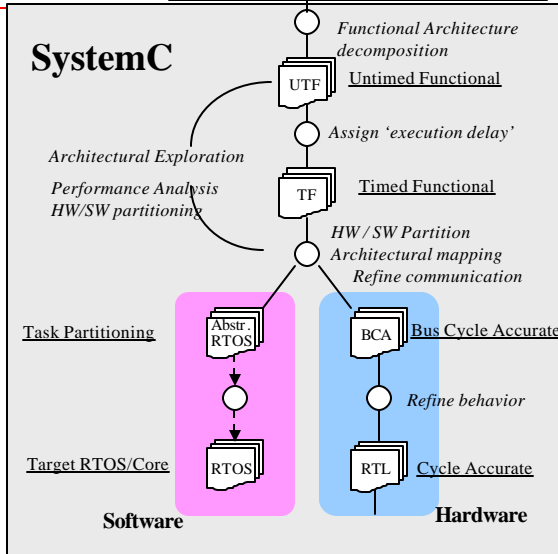
- ❑ SystemC abstraction levels & design flow
- ❑ Interface Synthesis
- ❑ Analyzing system performance
- ❑ Examples
- ❑ Key modeling paradigms & semantics
- ❑ Conclusions



Abstractions



© CoWare, Inc. 2001



3



Simulation speed for different abstractions

© CoWare, Inc. 2001

	Simulation speed	Real time (s) simulated in 10h	Comments
SoC	100 MHz	3,300,000x	WCDMA test platform @ 100 Mhz Real chip
Emulation	300 kHz	10,000x	Estimated
UTF	3 MHz	100,000x	
Timed Trans.	300 kHz	10,000x	Minimum required for SW/firmware developer
CATrans.	30 kHz	1,000x	Minimum required for platform developer
BCA	300 Hz	10x	With ISS for core
RTL	30 Hz	1x	No ISS

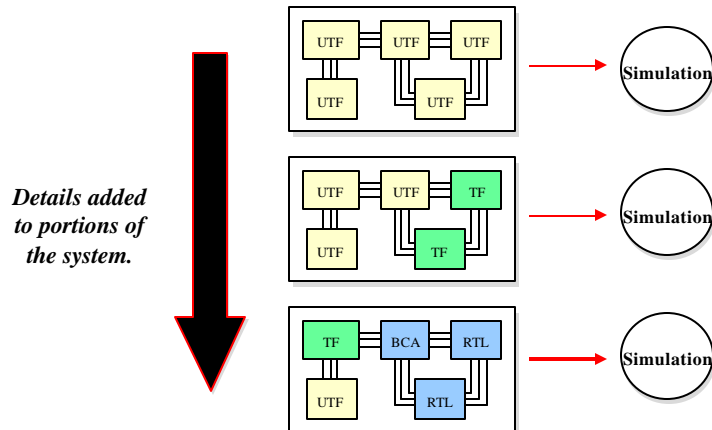
4



Gradual Refinement of the Design

© CoWare, Inc. 2001

Key to the methodology is that a design may be refined in a gradual step-wise fashion, rather than in one giant step... it need not be "all or nothing".



5

CoWare

Algorithmic level

© CoWare, Inc. 2001

- ❑ Algorithm described in a sequential language such as C
- ❑ Functional structure has no bearing to implementation architecture
- ❑ Example: JPEG/ MPEG algorithm development
- ❑ Design solutions: Matlab, SPW, Cossap, etc

6

CoWare

Untimed Functional (UTF)

© CoWare, Inc. 2001

- ❑ An architectural/structural decomposition of the system into functional blocks which are structurally interconnected over abstract communication channels
- ❑ Models control- and data-flow at an abstract functional level with abstract data types
- ❑ Enables easy refinement to HW-SW architecture
- ❑ Re-use of test benches & simulation results as references in subsequent design refinements

7

CoWare

Sequential (in-lined) process execution

© CoWare, Inc. 2001

- ❑ A process triggers execution of a slave process in-lined with itself
- ❑ Emulates function execution as in C programs

8

CoWare

Example: sequential C program

© CoWare, Inc. 2001

```
static int sum = 0; // state variable
```

```
void generate_data() {  
    for (int i = 0; i < 10; i++) {  
        accumulate(i);  
    }  
}
```

```
void accumulate(int in1) {  
    sum += in1;  
    cout << "Sum = " << sum << endl;  
}
```


9

CoWare

Same example with in-lined process execution

© CoWare, Inc. 2001

Sequential behavior of C-code is preserved
Separation in 2 processes allows
structural decomposition and re-use



Structure is
key for re-use

Transaction triggers slave process execution
Equivalent to function call but without function pointer

10

CoWare

Example: producer module

© CoWare, Inc. 2001

```
SC_MODULE(producer) {  
    sc_outmaster<int> out1;  
  
    void generate_data() {  
        for (int i = 0; i < 10; i++) {  
            out1 = i; // this will invoke  
                    //the slave;  
        }  
    }  
  
    SC_CTOR(producer) {  
        SC_METHOD(generate_data);  
    }  
};
```

11

CoWare

Slave process

© CoWare, Inc. 2001

```
SC_MODULE(consumer) {  
    sc_inslave<int> in1;  
  
    int sum; // state variable  
  
    void accumulate() {  
        sum += in1;  
        cout << "Sum = " << sum << endl;  
    }  
  
    SC_CTOR(consumer) {  
        SC_SLAVE(accumulate, in1);  
        sum = 0; // initialize  
    }  
};
```

12

CoWare

Structural module

© CoWare, Inc. 2001

```
SC_MODULE(top) { // structural module
    producer *A1;
    consumer *B1;
    sc_link_mp<int> link1;

    SC_CTOR(top) {
        A1 = new producer("A1");
        A1->out1(link1);
        B1 = new consumer("B1");
        B1->in1(link1);
    }
};
```

13

CoWare

Sequential (in-lined) process : summary

© CoWare, Inc. 2001


- SC_SLAVE process executes in-line with another process
 - Slave process has a single slave port that connects to a sequential link (sc_link_mp)
 - Caller process connects with a master port to a sc_link_mp
 - Invocation occurs by master accessing master port
 - Slave process can call other slaves

14

CoWare

Untimed Functional example

© CoWare, Inc. 2001



System is described in terms of 3 concurrent threads with in-lined slave processes. The concurrent threads (IU, CU, AU shown as loops) are synchronized over events: CR (Command Ready), BF (Buffer Free), AUR (AU Ready), SAU (Start AU)

15

CoWare

What the example does

© CoWare, Inc. 2001

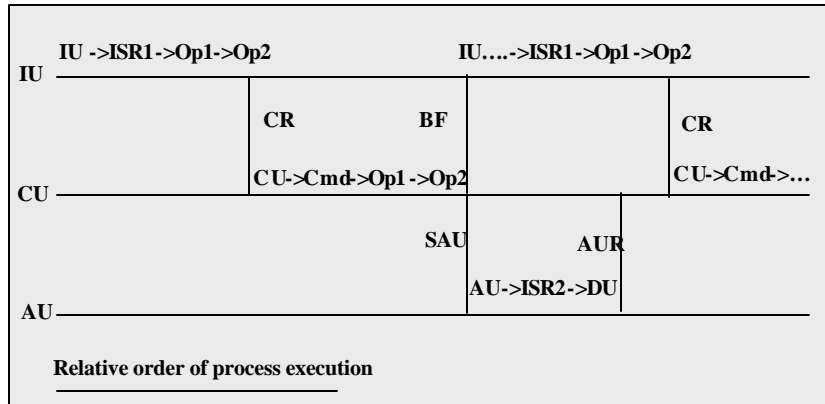
- ❑ A simple system with an Input Unit (IU), a Controller Unit (CU) and an arithmetic co-processor (AU) unit takes a triplet consisting of an arithmetic operator and two operands, calculates the result and displays it.
- ❑ The control flow is as follows: An Input Unit generates triplets (can be a test bench) asynchronously. Sends the triplet to the Control Unit (CU) if the latter's input buffer is free. CU writes the triples into the register of the Arithmetic Unit (AU) if the latter is not busy and generates a buffer free event. This will start the AU. The CU then waits for an AU ready event. The AU calculates the result and sends it back to the CU (ISR2) which then dispatches the result to the Display Unit and generates an AU ready event.

16

CoWare

Process execution order for the example

© CoWare, Inc. 2001



Vertical arrows depict synchronization events between the 3 concurrent threads
Horizontal lines show in-lined process executions within each concurrent thread (IU, CU, AU)

17

CoWare

UTF summary

© CoWare, Inc. 2001

- ❑ Combines sequential and concurrent processes
- ❑ Transaction & computation order is modeled, time is not
- ❑ Allows abstraction of (initially unwanted or unknown) implementation details : data/ control flow, clocking, concurrency, pipelining, time
- ❑ Assumptions: infinite resources which are infinitely fast

18

CoWare

Benefits of UTF

© CoWare, Inc. 2001

- ❑ Allows control/data flow description before timing is available
- ❑ Can be easily partitioned into SW-HW
- ❑ Very compact descriptions of complex data- and control-flows (10x-100x less code)
 - No FSM's required for behavior or communication
 - Abstract data types
- ❑ 5 orders of magnitude faster simulation than at RTL level

19

CoWare

Timed Functional Abstractions

© CoWare, Inc. 2001

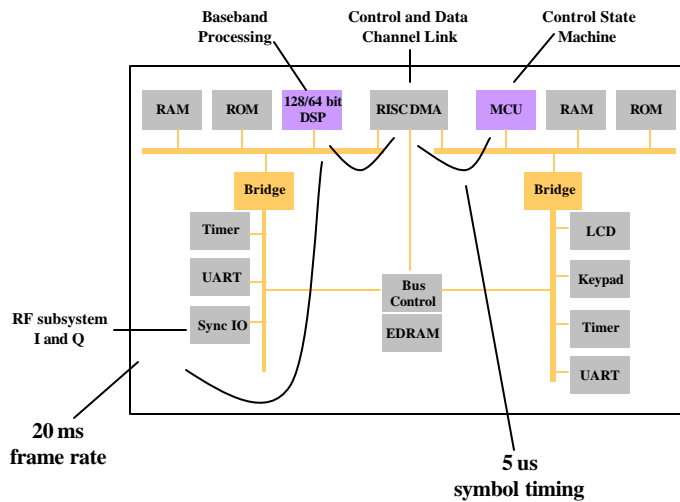
- ❑ Pre-partitioned (HW-SW)
 - Add timing to the UTF system in the form of timing estimates, timing constraints or time budgets
 - Time delays are in absolute time units
 - System is timed but not clocked
- ❑ Post-partitioned (HW-SW)
 - Add timing to the UTF system for HW-SW partitioning, architecture exploration and optimization (bus hierarchy, memory architecture, cache size, etc)
 - Map SW blocks to cycle accurate models of processor core & bus/memory architecture
 - HW is either untimed (assuming fast enough to not affect the critical path) or timed with estimated cycle delays

20

CoWare

Optimizing Wireless Handset architecture

© CoWare, Inc. 2001



21

Optimizing SOC bus architecture

© CoWare, Inc. 2001

- ❑ Bus architecture - choosing the correct bus architecture is one of the biggest problems in SOC design.
- ❑ Parameters
 - Standard (e.g. AMBA) versus custom
 - Arbitration (scheduling algorithms)
 - Bus width
 - Bus pipelining (wait state versus split transaction)
 - DMA
 - Hierarchy (e.g. system bus versus peripheral bus)

22

CoWare

Optimizing SOC memory architecture

© CoWare, Inc. 2001

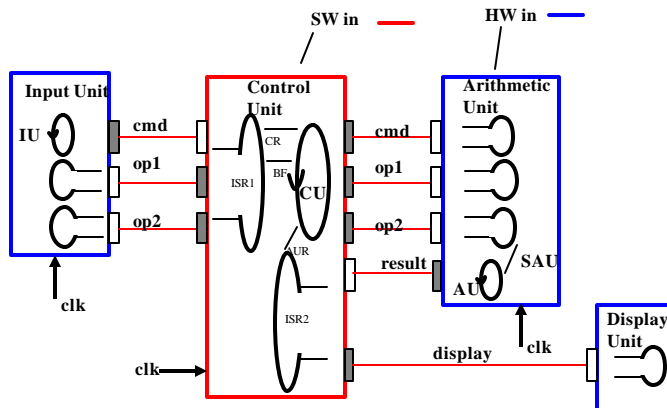
- ❑ Memories - choosing memory components and hierarchy can have large affect on system performance
 - cache, SRAM, DRAM, flash, etc.
 - size and hierarchy
 - power consumption
 - wait states and system throughput

23

CoWare

AU example after HW-SW partitioning

© CoWare, Inc. 2001



Map the processes in the Control Unit block to SW running on a MCU core. Use interrupt scenarios on cmd (from IU) and result channels with ISR1 and ISR2 as respective Interrupt Service Routines.

24

CoWare

Interface Synthesis (IFS)

© CoWare, Inc. 2001

- ❑ What you get from IFS
 - Ability to configure your platform with multi-master, multi-layer, multi-bus and hierarchical memory architectures
 - Cycle Accurate processor model, e.g. ARM926-EJS CCM model
 - Fast Cycle Accurate model of the above platform
 - Automatic mapping of your UTF system to this platform including generation of boot code, device drivers, address maps, address decoders in the bus and bus bridges
 - IFS is an enormous productivity tool since you don't have to know details about the core, its bus protocol, pin timing, etc. because the IFS tool has knowledge about the core and its platform.
 - You can now very easily (in a matter of hours) change your HW-SW partition at the UTF level or choose a different core or modify your platform, regenerate the cycle accurate model of the system and analyze the impact of your changes on system performance

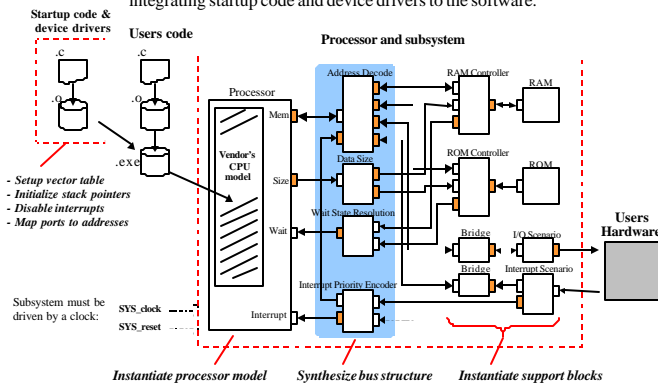
25

CoWare

Interface Synthesis

© CoWare, Inc. 2001

Interface Synthesis divides the system into distinct hardware and software portions, connecting a processor and support subsystem to the hardware, and integrating startup code and device drivers to the software.

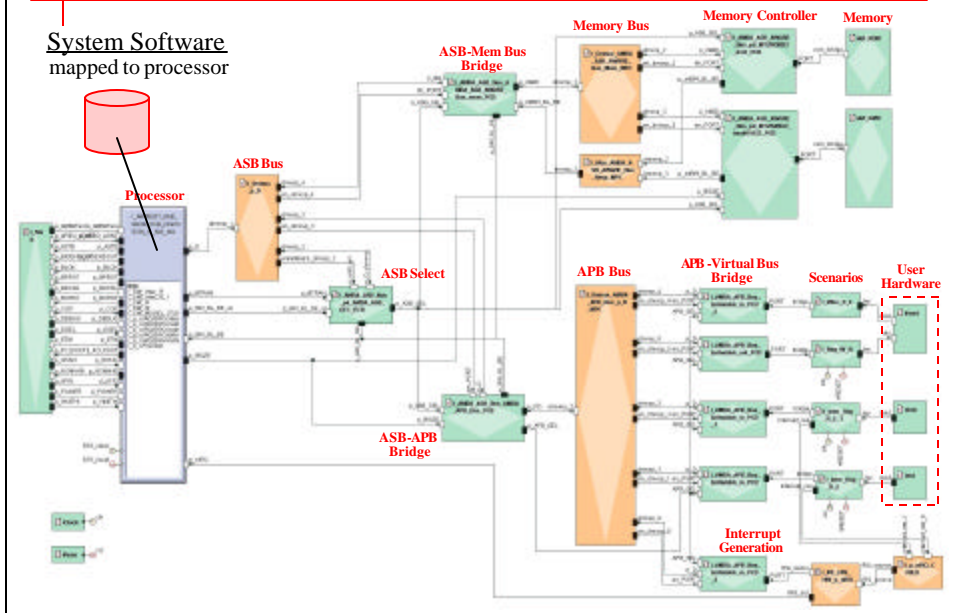


26

CoWare

Sample IFS-generated System

© CoWare, Inc. 2001



Analyzing System Performance

© CoWare, Inc. 2001

- ❑ Architectural analysis
 - Analyze token-based performance model
 - Study throughput and bottlenecks
 - Look at bus switching and cache usage to reduce power
 - Optimize bus & memory architecture
- ❑ Functional analysis
 - Look at system response and task scheduling
 - Analyze complexity to drive partitioning
 - Profile software for optimization

Analysis views (HW/SW)

© CoWare, Inc. 2001

HW Views

- Gantt chart
- Call Graph
- Operation Count
- Variable Info / Operator Count.
- Variable Tracing
- Function/Thread Call Stack

SW Views

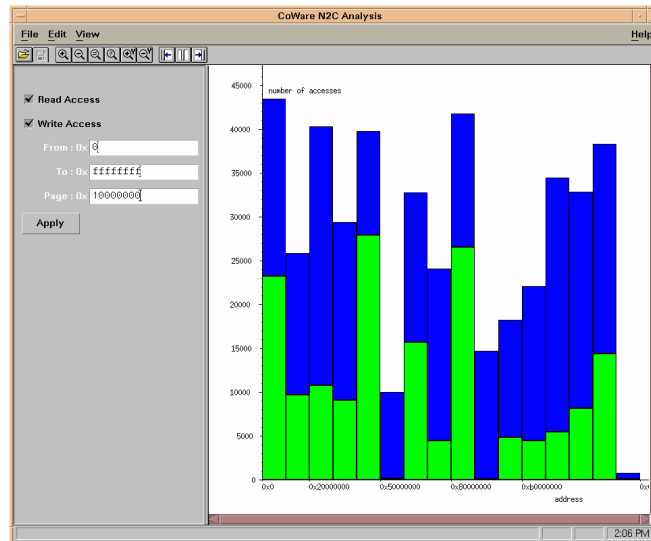
- Gantt chart
- Function/task call stack
- Memory access / time
- Memory page access
- CPU load
- Memory access + SW trace
- Memory map counts

29



Memory access patterns

© CoWare, Inc. 2001

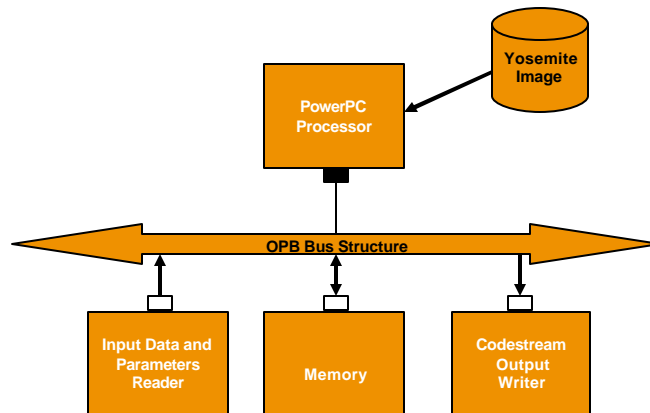


30



JPEG2000 VirtexII Pro (Xilinx) SW Platform

© CoWare, Inc. 2001



33

CoWare

JPEG2000 platform

© CoWare, Inc. 2001

- ❑ Architectural exploration focused on two alternatives
 - A) Implement JPEG algorithm with Wavelet transform in SW
 - B) Implement Wavelet algorithm in HW
- ❑ Architectural analysis reveals that bottleneck is not in the Wavelet algorithm as initially expected but instead in the code stream output writer due to memory accesses
- ❑ Without this analysis, RTL designers would have selected the wrong HW-SW partitioning
- ❑ Greatest gain in performance and power is achieved in architectural optimizations

34

CoWare

Modeling HW in the TF Abstraction

© CoWare, Inc. 2001

- ❑ Some blocks may remain untimed if their timing is not critical to the system function or performance (they're assumed to be fast enough)
- ❑ Other blocks are in UTF annotated with cycle delays
 - Concurrent HW threads must be attached to a clock
 - Slave processes get their clock tick from the concurrent thread that calls them

35

CoWare

SW design flow

© CoWare, Inc. 2001

- ❑ Abstract RTOS level
 - Functional threads are partitioned into SW processes per core processor in a multi-core system
 - SW processes may be created dynamically
 - SW processes are scheduled using an abstract (generic) RTOS; scheduling policy and process priorities are defined
 - Inter-process communication is implemented (shared memory, semaphores, sockets, mailboxes, etc)
- ❑ Target RTOS level
 - Abstract RTOS calls are now mapped to a target RTOS with further refinement of communication and scheduling

36

CoWare

The RTL Implementation level

© CoWare, Inc. 2001

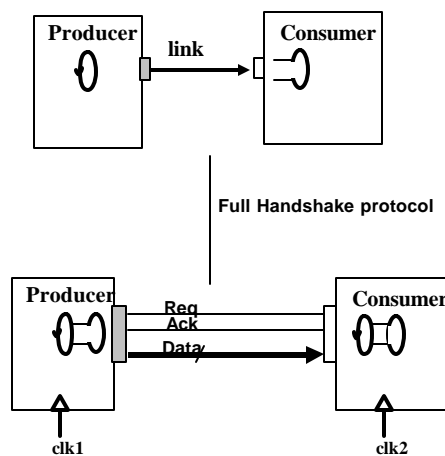
- ❑ IFS generates automatically RTL implementation of the platform from the specifications of the TF platform
- ❑ HW blocks are implemented at the RTL level as follows:
 - All HW (concurrent and slave) processes are implemented as clocked processes
 - Functional communication channels are refined to handshake bus protocols

37

CoWare

BCA example: full-handshake protocol

© CoWare, Inc. 2001



After in-lining of communication channel processes into the master and slave blocks

38

CoWare

Conclusions

© CoWare, Inc. 2001

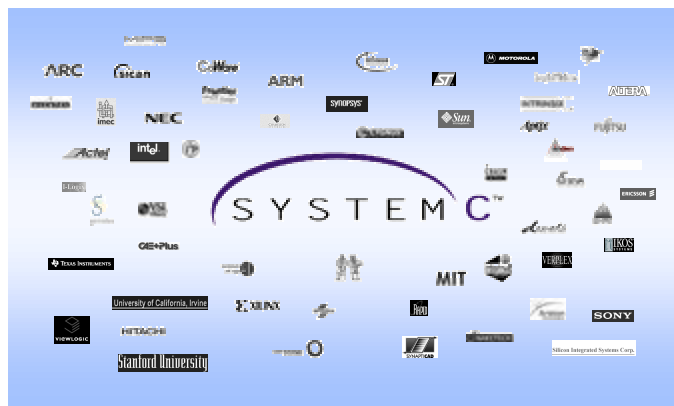
- ❑ SystemC supports modeling abstractions for HW/SW co-design from Functional to RTL
- ❑ Abstraction levels can be mixed for gradual refinement
- ❑ Benefits
 - Fast design exploration & HW/SW partitioning
 - Enables synthesis of communication
 - Fast IP embedding & retargeting
 - Orders of magnitude faster than RTL verification

39

CoWare

SystemC™ Community

© CoWare, Inc. 2001



40

CoWare

Backup slides

© CoWare, Inc. 2001

- ❑ TF communication details
- ❑ Key modeling paradigms

41

CoWare

TF communication revisited

© CoWare, Inc. 2001

- ❑ Timed Transactional (TT)
 - Mostly for SW modeling
 - Delay is in absolute time units
- ❑ Cycle Accurate Transactional (CAT)
 - For HW modeling including Bus/ Memory architecture
 - Delay in clock cycles

42

CoWare

Cycle Accurate Transactional Modeling

© CoWare, Inc. 2001

- ❑ Models transactions on a bus between bus masters & slaves
 - Cycle accurate
 - Bus protocol specific
 - Arbitration, address, data cycles
- ❑ Protocol independent generic API allows re-use
 - Mapping to a bus protocol through protocol specific libraries (re-use)
 - Specify attributes of a transaction: data width, data size, burst, status, etc

43

CoWare

Benefits of CA Transactional modeling

© CoWare, Inc. 2001

- ❑ Platform architecture exploration and optimization
 - Bus hierarchy: number of buses, assignment of resources to buses
 - Memory hierarchy / size
 - Cache size
 - HW-SW partitioning
 - Analysis for bus throughput, bottlenecks
- ❑ Much less user code than in RTL

44

CoWare

Key modeling paradigms

© CoWare, Inc. 2001

- ❑ Model of hierarchy and re-use
- ❑ Model of
 - Process execution
 - Process communication
 - Process synchronization
- ❑ Model of Time

45

CoWare

Model of hierarchy and re-use

© CoWare, Inc. 2001

- ❑ All behavior is described in processes; statements in a process execute sequentially
- ❑ Processes are “encapsulated” in modules which have ports
- ❑ All external communication takes place exclusively over module ports (to enable module re-use)
- ❑ Inside a module processes can communicate over shared variables and internal channels
- ❑ Modules are interconnected through their ports to communication channels

46

CoWare

Model of communication

© CoWare, Inc. 2001

- ❑ Abstract functional
- ❑ Transactional
 - Timed
 - Cycle accurate/ bus protocol accurate
- ❑ Bus Cycle Accurate communication
 - Cycle & Pin accurate

47

CoWare

Model of time

© CoWare, Inc. 2001

- ❑ Untimed
- ❑ Timed (but not clocked)
- ❑ Clocked (cycle accurate)

48

CoWare

Model of process execution

© CoWare, Inc. 2001

- ❑ Concurrent process execution
 - Communication & synchronization
- ❑ Sequential (in-lined) process execution
 - Abstract point-to-point & multi-point sequential communication

49

CoWare

Concurrent processes

© CoWare, Inc. 2001

- ❑ SC_THREAD process
 - Thread type (can be suspended)
 - Process execution is triggered from its sensitivity list
 - Process can suspend at a wait() statement & resume
 - Has a stack to store its state
 - Can have static and/ or dynamic sensitivity
- ❑ SC_METHOD process
 - Method process (functional call, can not suspend)
 - Process execution is triggered from its sensitivity list
 - Has static sensitivity only

50

CoWare

Concurrent process communication

© CoWare, Inc. 2001

- ❑ Concurrent processes communicate over a concurrent channel (=user defined module)
 - Shared variables for intra-module communication
 - Signals for intra- & inter-module communication
 - Signals are a special form of concurrent communication channel (for HW, supports evaluate-update paradigm)

51

CoWare

Model of process synchronization

© CoWare, Inc. 2001

- ❑ Concurrent processes synchronize over events
sc_event my_event; // fundamental sync object
- ❑ Two types of process sensitivity:
 - Static : wait() without arguments; sensitivity in constructor
 - Dynamic: changes during simulation
 - ❑ wait(ev1 | ev2 | ...) // on any one or more events
 - ❑ wait(ev1 & ev2 & ...) // on all events
 - ❑ event.notify(0) // next delta cycle
 - ❑ event.notify(delay); // timed notification
 - ❑ event.notify(); // immediate notify, not in v1.2
- ❑ Mutex & semaphores built as libraries on core sync objects

52

CoWare

Sequential (in-lined) process execution

© CoWare, Inc. 2001


- ❑ A process triggers execution of a slave process in-lined with itself
- ❑ Emulates function execution as in C programs

53

CoWare

Multi-point sequential communication

© CoWare, Inc. 2001



Abstract model for a bus communication:

A transaction by a master causes all slaves on the bus to be called;
Based on index value the slaves decide whether to respond or not
Multi-point will be later refined into a real bus communication
with address decoding & arbitration, concurrent behavior

54

CoWare

Sequential & concurrent communication combined

© CoWare, Inc. 2001

BW: blocking write with time-out

BR: blocking read with time-out

