

SimpCon – a Simple SoC Interconnect Draft

Martin Schoeberl
Vienna University of Technology, Austria
mschoebe@mail.tuwien.ac.at

November 28, 2005

This document proposes a simple interconnection standard for system-on-chip (SoC) components. It is intended to provide pipelined access to devices such on-chip peripherals and on-chip memory controller with minimum hardware resources.

1 Introduction

The intention of the following SoC interconnect standard is to be simple and efficient with respect to implementation resources and transaction latency.

SimpCon is a fully synchronous standard for on-chip interconnections. It is a point-to-point connection between a master and a slave. The master starts either a read or write transaction. Master commands are single cycle to free the master to continue on internal operations during an outstanding transaction. The slave has to register the address when needed for more than one cycle. The slave also registers the data on a read and provides it to the master for more than a single cycle. This property allows the master to delay the actual read if it is busy with internal operations.

The slave signals the end of the transaction through a novel *ready counter* to provide an early notification. This early notification simplifies the integration of peripherals into pipelined masters.

Slaves can also provide several levels of pipelining. This feature is announced by two static output ports (one for read and one write pipeline levels).

Off-chip connections (e.g. main memory) are device specific and need a slave to perform the translation. Peripheral interrupts are not covered by this specification.




Figure 1: Basic read transaction

1.1 Basic Read Transaction

Figure 1 shows a basic read transaction for a slave with one cycle latency. The acknowledge signals are omitted from the figure. In the first cycle, the address phase, the rd signals the slave to start the read transaction. The address is registered by the slave. During the following cycle, the read phase, the slave performs the read and registers the data. Due to the register in the slave the data is available in the third cycle, the result phase. To simplify the master, the read data stays valid till the next read request response.

1.2 Slave Acknowledge

Flow control between the slave and the master is usually done by a single signal in the form of *wait* or *acknowledge*. The *ack* signal, e.g. in the Wishbone specification, is set when the data is available or the write operation has finished. However, for a pipelined master it can be of interest to know it *earlier* when a transaction will finish.

For a lot of slaves, e.g. a SRAM interface with fixed wait states, this information is available inside the slave. In the SimpCon interface this information is communicated to the master through the two bit signal *rdy_cnt*. *rdy_cnt* signals the number of cycles till the read data will be available or the write transaction will be finished. Value 0 is equivalent to an *ack* signal and 1, 2, and 3 are equivalent to a wait request with the distinction that the master knows how long the wait request will last.

To avoid too many signals at the interconnect *rdy_cnt* has a width of two bits. Therefore, the maximum value of 3 has the special meaning that the transaction will finish in 3 or *more* cycles. As a result the master can only use the values 0, 1, and 2 to release actions in it's pipeline.

Idle slaves will keep the former value of 0 for *rdy_cnt*. Slaves, that don't know in advance how many wait states are need for the transaction can produce sequences that omit any of the numbers 3, 2, and 1. The master has to handle this situations.

Figure 2 shows an example of a slave that needs three cycles for the read to be processed. In cycle 1 the read command and the address are set by the master. The slave registers the

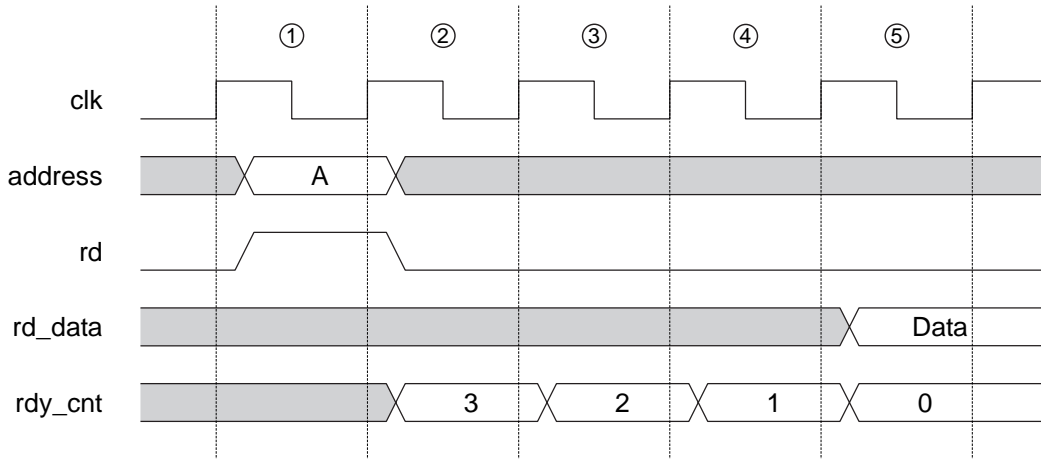


Figure 2: Read transaction with wait states




Figure 3: Write transaction with wait states

address and sets rdy_cnt to 3 in cycle 2. The read takes three cycles (2–4) during which rdy_cnt is decremented. In cycle 4 the data is available inside the slave and gets registered. It is available in cycle 5 for the master and rdy_cnt is finally 0. Both, the data and rdy_cnt will keep their value till a new transaction is requested.

Figure 3 shows an example of a slave that needs three cycles for the write to be processed. The address, the data to be written and the write command are valid during cycle 1. The slave registers the address and write data during cycle 1 and performs the write operation during cycles 2–4. The rdy_cnt is decremented and a non-pipelined slave can accept a new command after cycle 4.

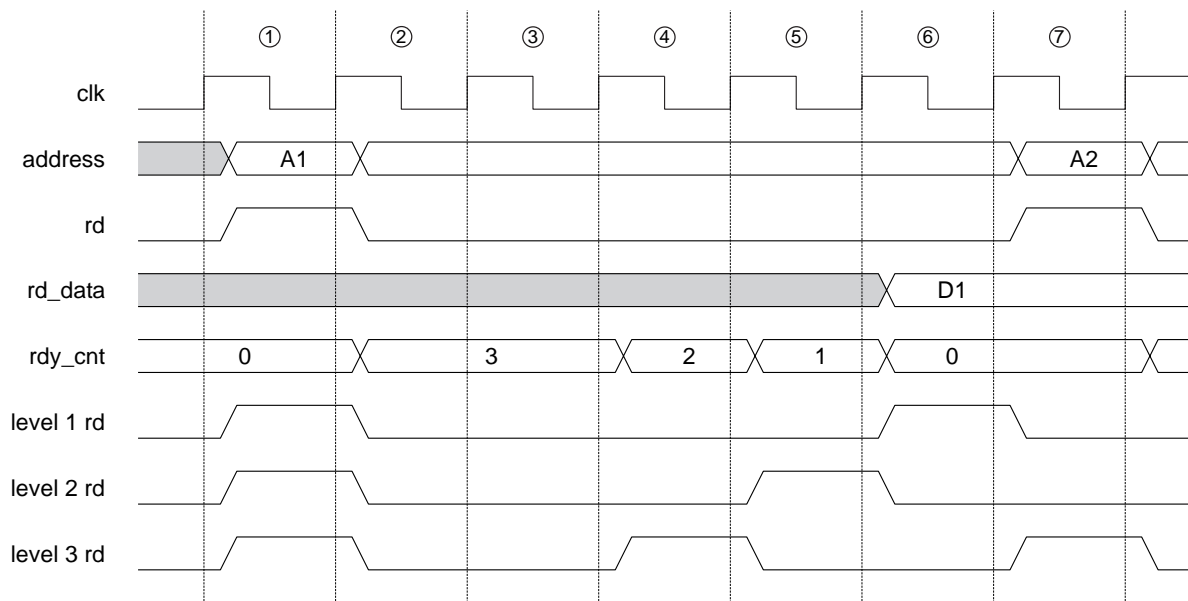


Figure 4: Different pipeline levels for a read transaction

2 Pipelining

Figure 4 shows a read transaction for a slave with four cycles latency. Without any pipelining the next read transaction will start in cycle 7 after the data from the former read transaction is read by the master. The three bottom lines show when new read transactions will be started for different pipeline levels. With pipeline level 1 a new transaction can start in the same cycle when the former read data is available (in this example in cycle 6). Higher levels mean that the next read will start earlier as shown for level 2 and 3.

Implementation of level 1 in the slave is trivial (just two more transitions in the state machine). It is recommended to provide level 1 at least for read transactions. Level 2 is a little bit more complex but usually no additional address or data registers are needed.

To implement level 3 pipelining in the slave at least an additional address register is needed. However, to use level 3 the master has to issue the request in the same cycle as `rdy_cnt` goes to 2. That means this transition is combinatorial. We see in Figure 4 that `rdy_cnt` value of 3 means three or more cycles till the data is available and can therefore not be used to trigger a new transaction.

3 Multiple Master

SimpCon defines no signals for the communication between a master and an arbiter. However, it is possible to build a multi master system with SimpCon. The SimpCon interface can be used as interconnect between the masters and the arbiter and the arbiter and the slaves. In this case the arbiter acts as slave for the master and as master for the peripheral devices.

The missing arbitration protocol in SimpCon results in the need to queue $n - 1$ requests in

an arbiter for n masters. However, for this additional HW we get zero overheads for the bus request. The master, which gets the bus will start the slave transaction in the same cycle.

TODO: add a timing diagram to explain this concept.

4 Examples

This section provides some examples for the application of the SimpCon definition.

4.1 IO Port

TODO: Show how simple an IO port can be with SimpCon. We need no addresses and can tie `bsy_cnt` to 0. We only need the `rd` or `wr` signal to enable the port.

4.2 SRAM interface

The following example is taken from an implementation of SimpCon for a Java processor. The processor is clocked with 100MHz and the main memory consists of 15ns static RAMs. Therefore the minimum access time for the RAM is two cycles. The slack time of 5ns forces us to use output registers for the RAM address and write data and input registers for the read data in the IO cells of the FPGA. These registers fit nice with the intention of SimpCon to use registers inside the slave.

Figure 5 shows the interface for a non-pipelined read access followed by a write access. Four signals are driven by the master and two signal by the slave. The lower half of the figure shows the signals at the FPGA pins where the RAM is connected.

In cycle 1 the read transaction is started by the master and the slave registers the address. The slave also sets the registered control signals `ncs` and `noe` during cycle1. Due to the IO cell registers, the address and control signals are valid at the FPGA pins very early in cycle 2. At the end of cycle 3 (15ns after address, `ncs` and `noe` are stable) the data from the RAM is available and can be sampled with the rising edge for cycle 4.

The master reads the data in cycle 4 and starts a write transaction in cycle 5. Address and data are again registered from the slave and are available for the RAM at the beginning of cycle 6. To perform a write in two cycles the `nwr` signal is registered by a negative triggered flip-flop.

In figure 6 we see a pipelined read from the RAM with pipeline level 2. With this pipeline level and the two cycles read access time of the RAM we get the maximum bandwidth possible.

We can see the start of the second read transaction in cycle 3 during the read of the first data from the RAM. The new address is registered in the same cycle and available for the RAM in the following cycle 4. Although we have a pipeline level of 2 we need no additional address or data register. The read data is available for two cycles (`rdy_cnt` 2 or 1 for the next read) and the master is free to select one of the two cycles to read the data.

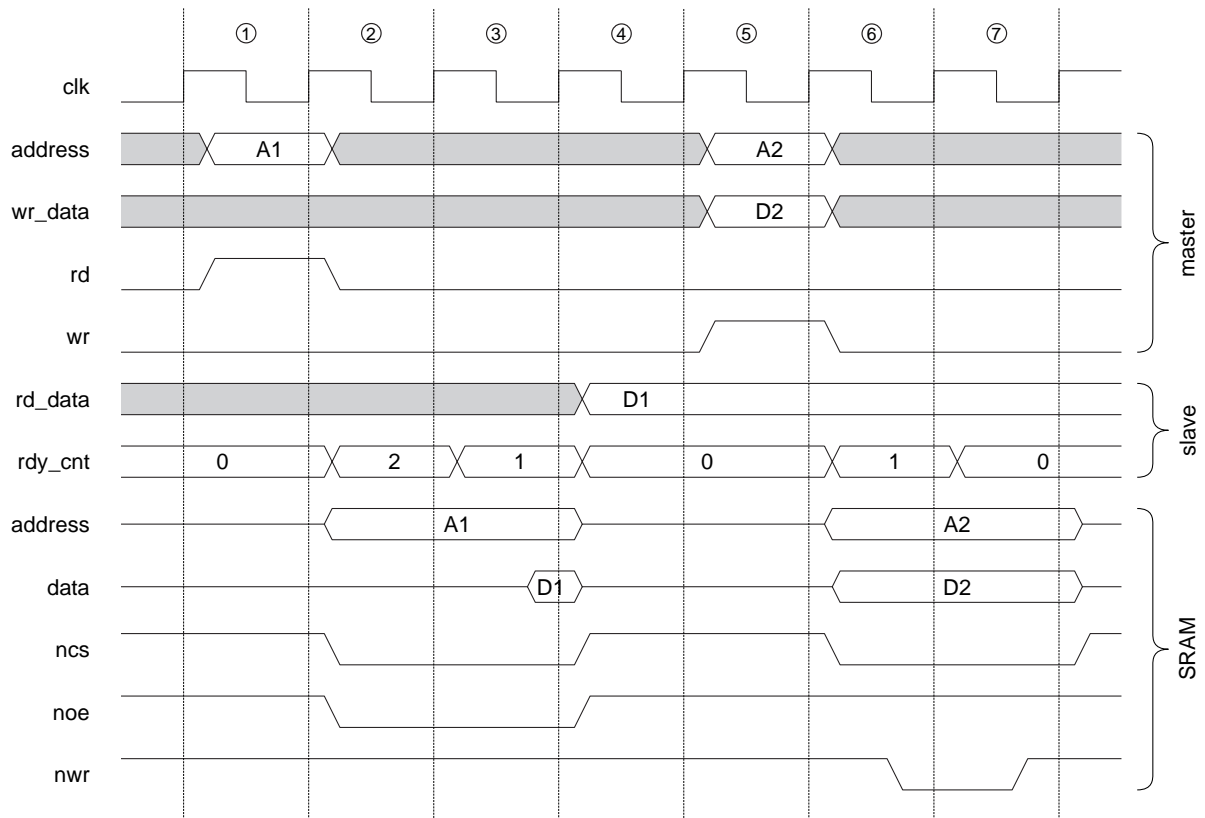


Figure 5: Static RAM interface without pipelining

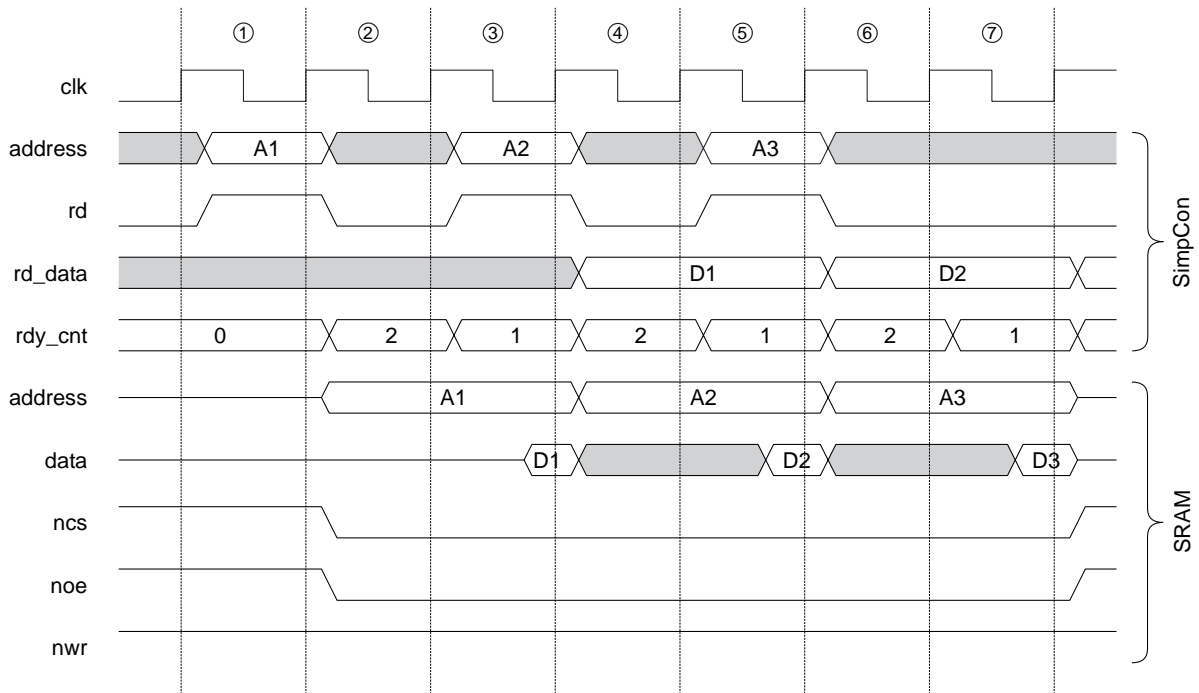


Figure 6: Pipelined read from a static RAM

4.3 Master Multiplexing

To add several slaves to a single master the rd_data and bsy_cnt have to be multiplexed. Due to the fact that all rd_data signals are registered by the slaves a single pipeline stage will be enough for a large multiplexer. The selection of the multiplexer is also known at the transaction start but needed at most in the next cycle. Therefore it can be registered to further speed up the multiplexer.

TODO: add a schematic for the master rd_data multiplexer.

5 Status

- First timing diagrams drawn
- SimpCon SRAM interface for JOP on Cyclone and Spartan-3 is available
- Project at opencores.org is requested

Next steps:

- Continue this document
- Provide more SimpCon examples (e.g. a UART)

- Change JOPs IO interface to SimpCon
- Provide Wishbone bridges