

Multi-layer AHB

V2.0

Technical Overview



Multi-layer AHB

Technical Overview

Copyright © 2001, 2004 ARM Limited. All rights reserved.

Release Information

The following changes have been made to this book.

Change History

Date	Issue	Confidentiality	Change
16 March 2001	A	NonConfidential	First issue
10 May 2004	B	NonConfidential	Second issue

Proprietary Notice

Words and logos marked with ® or ™ are registered trademarks or trademarks of ARM Limited in the EU and other countries, except as otherwise stated below in this proprietary notice. Other brands and names mentioned herein may be the trademarks of their respective owners.

Neither the whole nor any part of the information contained in, or the product described in, this document may be adapted or reproduced in any material form except with the prior written permission of the copyright holder.

The product described in this document is subject to continuous developments and improvements. All particulars of the product and its use contained in this document are given by ARM in good faith. However, all warranties implied or expressed, including but not limited to implied warranties of merchantability, or fitness for purpose, are excluded.

This document is intended only to assist the reader in the use of the product. ARM Limited shall not be liable for any loss or damage arising from the use of any information in this document, or any error or omission in such information, or any incorrect use of the product.

Where the term ARM is used it means “ARM or any of its subsidiaries as appropriate”.

Confidentiality Status

This document is Non-Confidential. The right to use, copy and disclose this document may be subject to license restrictions in accordance with the terms of the agreement entered into by ARM and the party that ARM delivered this document to.

Product Status

The information in this document is final, that is for a developed product.

Web Address

<http://www.arm.com>

Contents

Multi-layer AHB Technical Overview

Chapter 1	Multi-layer AHB Technical Overview	
1.1	Introduction	1-2
1.2	Implementation	1-3
1.3	Advanced options	1-4
1.4	Example implementation	1-7

Chapter 1

Multi-layer AHB Technical Overview

This technical overview describes the functionality of the multi-layer AHB in the following sections:

- *Introduction* on page 1-2
- *Implementation* on page 1-3
- *Advanced options* on page 1-4
- *Example implementation* on page 1-7.

1.1 Introduction

Multi-layer AHB is an interconnection scheme, based on the AHB protocol, that enables parallel access paths between multiple masters and slaves in a system. This is achieved by using a more complex interconnection matrix. Key advantages are:

- You can develop multi-master systems with an increased available bus bandwidth.
- You can construct complex multi-master systems that have a flexible architecture. This removes the requirement to fix design decisions about the allocation of system resources to particular masters at the hardware design stage.
- You can use standard AHB master and slave modules without requiring modification.
- Each AHB layer can be very simple because it only has one master, so no arbitration or master-to-slave muxing is required. These layers can use the AHB-Lite protocol, meaning that they do not have to support request and grant, or retry and split transactions.
- The arbitration effectively becomes point arbitration at each peripheral and is only necessary when more than one master wants to access the same slave simultaneously.
- The only hardware you have to add to the standard AHB transport infrastructure is the multiplexor block to connect the multiple masters to the peripherals.
- Because the multi-layer architecture is based on the existing AHB protocol, you can reuse previously-designed masters and slaves without modification.

Figure 1-1 shows a block diagram of the basic multi-layer concept.

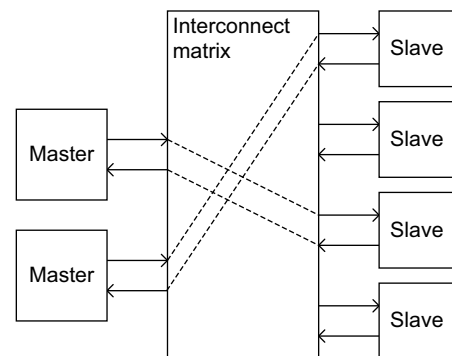


Figure 1-1 Basic multi-layer concept

1.2 Implementation

In the simplest implementation of a multi-layer system, each master has its own AHB layer and is connected to the slave devices by an interconnect matrix, as shown in Figure 1-2.

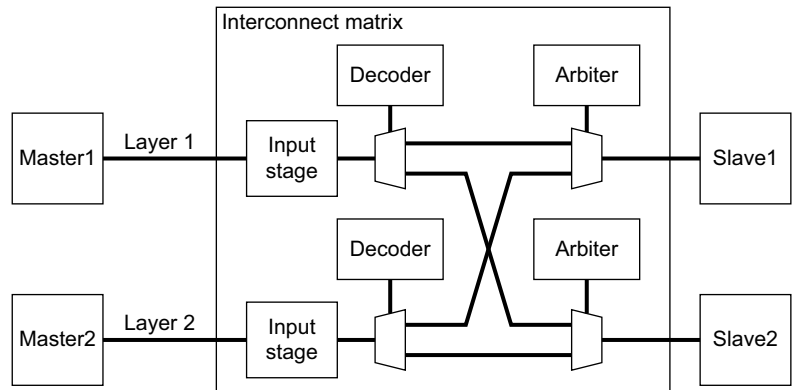


Figure 1-2 Multi-layer interconnect topology

Within the interconnect matrix:

1. Every layer has a Decode stage that determines which slave is required for a transfer.
2. A mux routes the transfer from the appropriate layer to the required slave.

If two layers require access to the same slave at the same time, the arbitration within the interconnect matrix must determine which layer has highest priority. The layer that is not given access is waited using **HREADY** until it is given access to the required slave. When a layer is waited an Input Stage is used to store a copy of the pipelined address and control information until the access to the shared slave is given.

Each slave port has its own arbitration and a number of different schemes can be used. For example:

- input layers can be serviced in round-robin manner, changing every transfer or every burst
- the arbitration can use a fixed priority scheme where certain high priority layers are always given access in preference to lower priority layers.

The number of input/output ports on the interconnect matrix is completely flexible and can be adapted to suit the system requirements.

1.3 Advanced options

This section describes:

- *Bus configurations*
- *Multi-port slaves* on page 1-6.

1.3.1 Bus configurations

As the number of masters and slaves in a system increases, the complexity of the interconnect matrix can become significant. You can use the following techniques to optimize the system architecture:

- Figure 3 shows how you can make slaves local, or private, to a particular layer. This reduces the complexity of the interconnect matrix. You can use this when it is acceptable that a slave can only be accessed by masters on the same layer.




Figure 1-3 Local slaves

- Figure 1-4 on page 1-5 shows how you can make multiple slaves appear as a single slave to the interconnect matrix. This is useful to combine a number of low-bandwidth slaves. You can also use it where a set of slaves are normally accessed by just one master, such as a DMA controller, and the interconnect matrix is used only to give access to other masters under special circumstances, such as debugging the system.




Figure 1-4 Multiple slaves on one slave port

- In the simplest implementation of a multi-layer system, each master has its own layer. Figure 1-5 shows how you can also build a system where multiple masters share a layer. This is well suited to combine masters that have low-bandwidth requirements or masters, such as the *Test Interface Controller (TIC)*, that have particular characteristics.

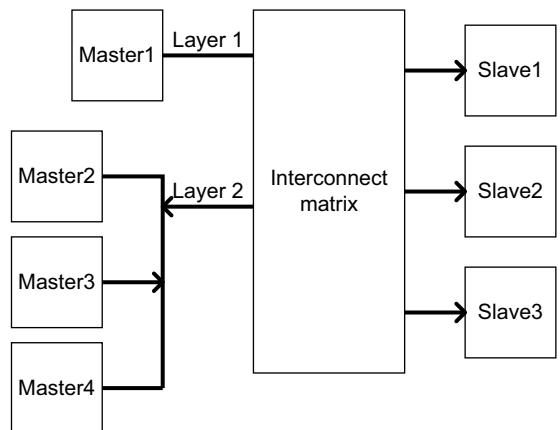


Figure 1-5 Multiple masters on one layer

- Each layer can be a complete AHB subsystem, with the interconnect matrix being used to enable communication between the two subsystems. The example in Figure 1-6 shows only a single slave is shared. Typically this is an on-chip memory slave that is used as a buffer area between the two subsystems.

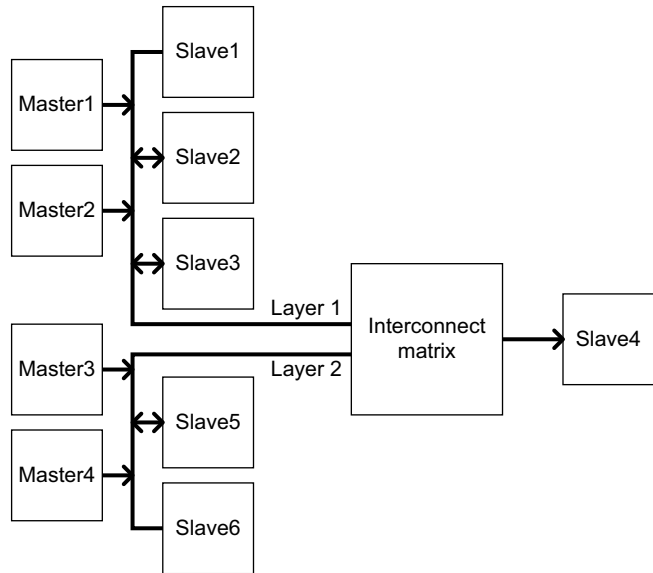


Figure 1-6 Separate AHB subsystems

1.3.2 Multi-port slaves

In a multi-layer AHB system certain slaves, such as an SDRAM controller, are able to operate more efficiently by processing transfers from different layers in parallel. Figure 1-7 shows how you can do this by designing the slave with multiple AHB slave ports.




Figure 1-7 Multi-port slaves

1.4 Example implementation

The techniques described in *Advanced options* on page 1-4 can be combined into a single system. The example in Figure 1-8 shows:

- CPU1 has a dedicated AHB layer, layer 1.
- CPU2 and the DMA engine share AHB layer 2.
- There is a dedicated SRAM block connected to layer 1, only accessible from CPU1.
- The LCD controller is only connected to layer 2, because this is the layer on which the DMA engine operates.
- The interconnect matrix has two slave ports, each accessible from either layer. One of the devices on these slave ports can be an AHB to APB bridge, enabling the connection of many, low-bandwidth peripherals.
- The external SDRAM interface has its own two-port connection to both layers.




Figure 1-8 Example multi-layer AHB system

