

On the Way to the 2.5 Gbits/s ATM Network

ATM Multiplexer Demultiplexer ASIC

Jacobo Riesco[†], Juan C. Díaz[†], Luis A. Merayo[†], José Luis Conesa[†], Carlos Santos^{*} & Eduardo Juárez[◇]

[†]Telefónica Investigación y Desarrollo.

Emilio Vargas 6. 28043 Madrid (Spain).

^{*}SIDSA

Isaac Newton, 1. 28760 Tres Cantos - Madrid (Spain).

[◇]Universidad Politécnica de Madrid.

ETSI Telecomunicación. Ciudad Universitaria, s/n. 28040 Madrid (Spain)

Abstract

The present paper describes the AMDA integrated circuit (ATM Multiplexer/Demultiplexor ASIC). The circuit has two operation modes: in multiplexer mode an ATM low speed flow (up to 622 Mbits/s) is inserted in the empty slots of a high speed ATM flow (2.5 Gbits/s); in demultiplexer mode, the cells belonging to the low speed channels are extracted from the high speed ATM flow. A specific algorithm of distributed control has been developed, simulated and implemented, in order to guarantee an even bandwidth distribution independently of the network node position. The circuit is able to handle 8K connections, with four different qualities of service; it manages a local queue of up to 16K ATM cells using an external high speed SSRAM. The maximum clock frequency of the circuit is 155,52 MHz and it has been processed with the LSI-LOGIC's LCB500K technology (0,5 μ m CMOS). It contains 34800 equivalent gates, 48 Kbit of single port memory and 8,5 Kbit dual port memory, using an area of 6,7 x 6,7 mm and it is packaged in a 208 pins QFP.

1. Introduction

One of the main advantages of an asynchronous transfer mode (ATM) network is its ability to obtain a gain on efficiency by relaying on the statistical multiplexing effect of variable bit rate (VBR) sources, on the condition that enough sources are multiplexed and they are not correlated [1].

While in a network with fixed bit rate (FBR) coding, the required bandwidth is the sum of all sources peak bit rates, an ATM network will only allocate a certain bandwidth for each source, determined by the acceptable cell loss rate,

given by the probability that a number of uncorrelated VBR or bursty sources multiplexed on a single link surpass its capacity. This probability decreases as the number of sources increases, i.e., the larger the link bit rate, the higher multiplexing gain can be achieved [1, 6].

However to exploit this benefits, the means to multiplex the data from the sources on a high speed link and also to extract the multiplexed channels to the destination terminals are needed.

The main concern with multiplexing is the medium access control (MAC), as all the sources share a common link. The MAC must be distributed, since the sources may be geographically disperse, scalable, and must distribute the total bandwidth evenly. Multiplexers must also provide an input queuing mechanism to store source data, when the link is saturated, until they can be inserted.

Demultiplexers must identify the data destined to them and, as bursts at the link bit rate may appear in the receiver side, data must be queued and presented to the terminal at the proper speed.

2. AMDA general description

The AMDA ASIC development was aimed to cope with the multiplexing and demultiplexing tasks of a high speed ATM network (figure 1). It can be programmed to perform either function, but not both at the same time.

As a multiplexer an ATM high speed flow (2.5 Gbits/s) is filled up with the cells coming from a low speed one (up to 622 Mbits/s). The cells awaiting to be inserted are stored in two different queues, one for high priority traffic (i.e. voice, real time video) and one for low-priority traffic (i.e. data), until an empty cell is detected in the high speed flow. The cells stored in the high-priority queue are inserted in the high speed flow before those of the low-priority queue (i.e.,

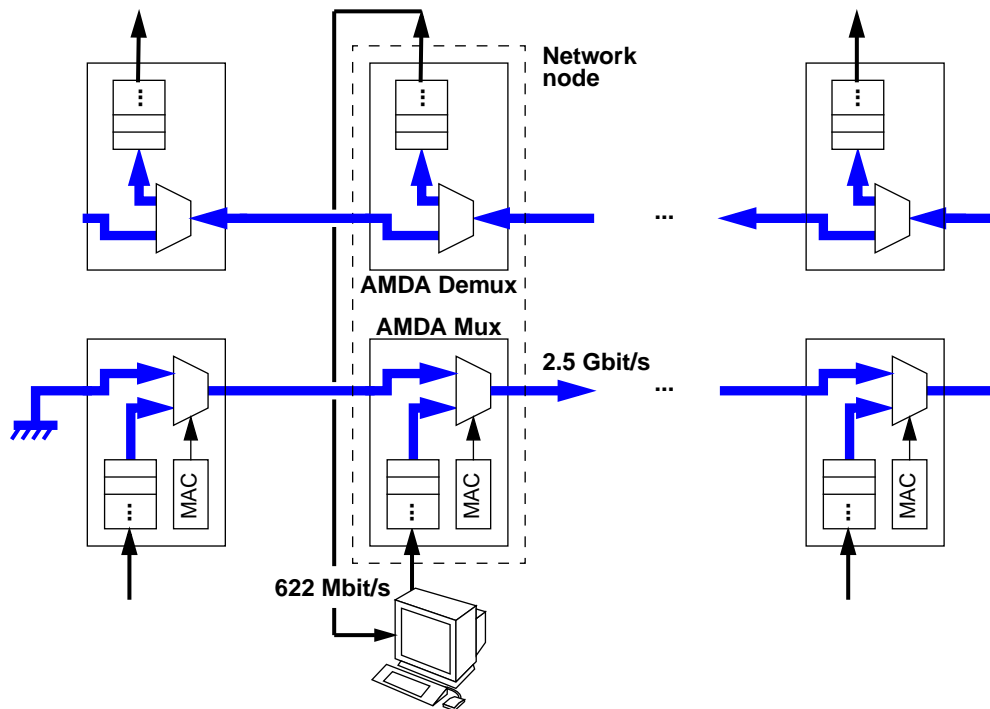


Figure 1: Network configuration

while there are cells in the high-priority queue those of the low-priority queue are not inserted). To ensure to all the sources a similar probability of inserting its traffic [2-5] it is necessary a distributed mechanism for the medium access control (MAC).

In demultiplexer mode, the ATM cells are identified and extracted depending upon their headers. This function is done in a way that allows for point to multipoint and broadcast communications. The different traffic rates between the high speed link and the receiver are decoupled by two queues: the high-priority cells are extracted towards the receiver before the low-priority ones. Additionally, when the low-priority cells queue of any demultiplexer gets saturated, (the number of stored cells exceeds a certain programmable threshold) it informs (through a *backpressure* signal) all the upstream multiplexers, so that they stop sending low-priority cells, and thus avoiding the queue overflow.

The AMDA has been designed so that the communications between them can be transported on high speed optical fibre, to allow the construction of a distributed local or metropolitan area network (LAN or MAN).

3. Architectural and functional description.

The AMDA internal architecture (see figure 2) is made-up by a high speed path, transporting a 2.5 Gbit/s ATM

cell flow, a low speed path transporting cells at a binary rate of up to 622 Mbit/s, and the control and auxiliary blocks.

3.1. High speed path.

The high speed bus is processed through a pipeline path composed by an input interface, an input FIFO, two blocks for inserting or extracting the low speed cells, and an output interface.

The input and output blocks have functionally compatible interfaces with the ATM Forum's UTOPIA Level 1 standard (16 bit with octet protocol, ATM RX/TX layer [7]) extending the maximum clock frequency up to 155,52 MHz. Their functions are the conversion between the external 16 bits data format and the internal 64 bits width bus, the extraction/insertion of the *request* and *backpressure* signals in programmable positions within the cells header, the empty cells identification [9] and the checking/generation of a 8 bits parity field. Besides that, the input block is in charge of cell integrity control and maintenance functions, and counting the invalid cells to measure the quality of the link.

The input FIFO, a DPRAM memory of 56 words 64 bits wide, acts as data *buffer* when physical layer circuits (SDH/PDH) are used for transmission [10] and allows pleiosynchronism between receiver and transmitter clocks.

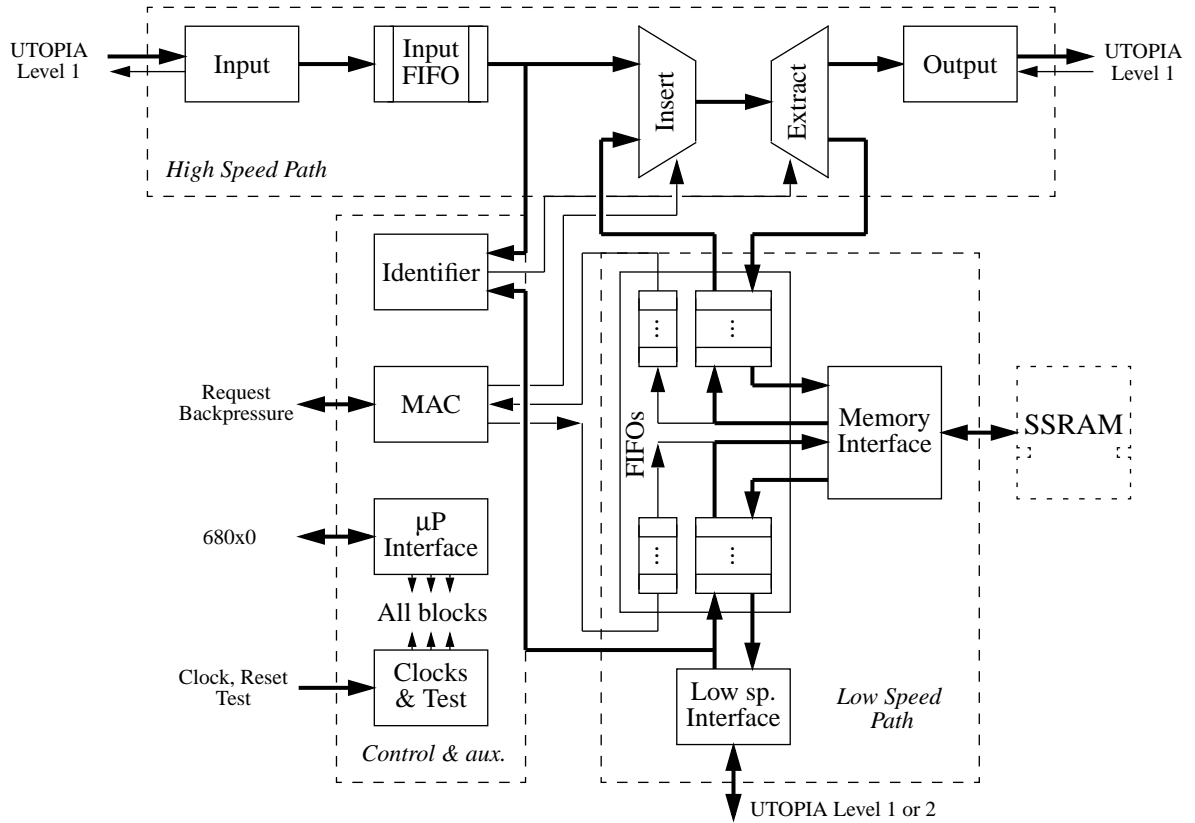


Figure 2: AMDA block diagram

In multiplexer mode the insertion block replaces, under the MAC indication, the empty input cells with cells coming from the low speed interface; it also signals a congestion status, asserting the Explicit Forward Congestion Indication (EFCI) [9], when the low-priority queue reaches a certain programmable threshold.

In demultiplexer mode the extraction block writes the input cells in the low speed FIFOs and may replace them with empty cells, depending on the indications of the identification block; besides, it has a group of counters to maintain statistics measures on the high speed channel cells flow (both in mux and demux modes).

3.2. Low speed path

It is composed by the FIFOs block, external memory interface and low speed interface.

The FIFOs block controls the processes to store the cells in two different queues depending on their priority. Each queue is physically implemented with several FIFOs acting as internal buffers connected to the external mass storage memory interface. This block contains four DPRAMS of sizes (words x bits) 42x64, 28x64, 8x16 and 8x17 respectively.

The external memory interface arbitrates and controls the writing and reading processes in an external SSRAM [11] and generates the threshold and overflow flags. The AMDA is able to address up to 256 Kwords of 32 bits, which means a storage capacity of up to 16 K ATM cells. From the logical point of view, the memory behaves as two FIFOs (low and high-priority) with programmable sizes and thresholds. The memory is accessed in burst mode to increase the bandwidth; address and data buses are multiplexed to reduce the interface pin count. The memory interface can operate with a clock of up to 110 MHz to provide a total bandwidth of 3 Gbit/s.

The low speed interface block consist of two unidirectional programmable interfaces, bidirectional pins are used to reduce the pin count. In multiplexer mode the low speed interface is an UTOPIA standard input interface (ATM RX layer); it can be programmed as a level 1 octet protocol interface [5] or a level 2 (multi-PHY) interface [6], 8 or 16 bits wide, and with 54 or 53 bytes cells in the 8 bit mode. In demultiplexer mode works as an UTOPIA standard output interface (ATM TX layer), with the same programming possibilities. In both cases the maximum traffic rate is 622 Mbit/s, using a 50 MHz clock.

3.3. Control and auxiliary blocks

These blocks are: MAC block, cells identification block, microprocessor interface, clocks generation block and test control block.

The MAC (medium access control) block, in multiplexer mode, arbitrates in a distributed way, the access to the 2.5 Gbit/s channel between the different AMDAs sharing it, controlling the insertion block, and generating and propagating the *request* signals. The algorithm implemented, ADAM, minimizes the queues lengths and the cell delay variations (CDV) independently of the AMDA position in the network. Thus all the nodes have the same bandwidth available, and the network behaviour is equivalent to a distributed queue with a *First-Come-First-Served* (FCFS) service discipline [3,5].

The identification block assigns a set of connection parameters to every cell, and performs selective cell discarding and statistics measures on the cells flow. The cells are identified with 13 bits of the VPI and VCI cell header fields [9], which allows a maximum of 8K different channels. Internally it has a 2K words of 24 bits memory, so there are 6 control bits per connection: propagation control bit (demux mode), extraction control bit (demux mode), priority control bit, loss priority control bit, selective discarding enable bit and selective discarding status bit.

When the selective cell discarding mechanism is enabled for a given connection, if a cell within a frame is rejected, the remaining cells until the end of the frame indication (flagged by a given value on the PTI field [9]) will also be rejected, thus avoiding further FIFOs overflows. The priority and loss priority qualifiers allow four quality of service (QoS) levels in communications (see table 1).

The microprocessor interface controls the communications with 680x0 family micros. It allows the programming of the different operation modes and the statistical measures acquisition, through the access to the internal AMDA registers. It also handles the wrong conditions detection through the interruptions system (8 interrupt reasons). The AMDA has a total of 31 registers (of 8 bits).

Table 1: QoS Levels in AMDA

		Delay sensitivity	
		High	Low
Loss Sensitivity	High	Coded real time video	Data
	Low	Voice	WWW

The clocks block generates the different internal operation clocks of the AMDA from the main system 155 MHz clock and synchronizes the asynchronous system *reset* and software *reset* signals with these clocks.

The test control block generates the signals for the control of the foundry tests (Iddq, High-Z, V_{IH}/V_{IL}) structural tests (ad-hoc) and BIST of the memories.

4. Design methodology

The AMDA circuit has been designed following a top-down methodology approach with the use of the Verilog HDL and logical synthesis.

The first step was to develop a very high level functional Verilog representation of the whole circuit and simulate a simple network with 16 nodes in order to determine the best MAC algorithm and to dimension the mux and demux queues [3].

After that, a functional description of each circuit block was done, simulated and integrated with the other blocks. Automatic traffic generators and analysers were used to automate functional verification. Also, a simulation was run with two AMDAs connected, to ensure the correct interoperability.

The most critical parts of the circuit, in terms of speed, were described in a nearly structural (but technology independent) way, while slower blocks had higher level RTL descriptions. The other main concern in design, besides speed, was power consumption, and so, low power (clock enabling, multiphase clocks) design methodologies were used for the high speed (155 MHz) input and output block descriptions.

Technology selection key point was I/O high speed buffers availability and consumption, as one of the design main targets was to be able to package the ASIC in a cheap 208 plastic quad flat pack (PQFP).

After technology selection (LSI-Logic LCB500K), the circuit was synthesized. This was not a straightforward process due to the high number (10) of unrelated clocks in the circuit; scripts were written to set the time and design constraints for each block. Extremely careful design was done in all the interfaces between clocks to avoid metastability problems.

Regression tests were then performed on the netlist representation of each block and the whole circuit, first without timing information, and then with estimated delays from LSI-Toolkit. A static timing analysis was also performed in most critical blocks. The high speed interfaces between AMDAs were manually optimised for the target technology.

The test bench was also synthesised, so that a complete functional verification could be downloaded into an ASIC emulator.

The floor-plan, done in cooperation with LSI, followed the data flow structure. After floor-planning and clock buffering, a new regression test was performed.

Some signal buffers and *ad-hoc* test structures (as registers presets) were added to the circuit before final layout. The clock complexity and timing restrictions prevented the use of scan techniques. Test vectors were developed and analysed with fault simulations until a coverage over 90% was achieved.

Once the layout was finished regression tests on functional vectors were done again, as well as the simulation of test vectors.

It is important to notice that all simulation processes were done using the same Verilog test benches developed in the first stages of the design. This has proved to be a great advantage, as simplifies verification processes..

5. Physical implementation

The circuit has been developed with the LSI-Logic LCB500K technology (0,5 μm CMOS). It contains 34800 equivalent gates, 48 Kbit single port memory and 8,5 Kbit of dual port memory, uses an area of 6,7 x 6,7 mm and it is packaged in a plastic 208 pins QFP, of those, 42 are inputs, 35 outputs, 60 bidirectionals, 31 for power supply (3.3 V) and 37 for ground (137 functional pins and 68 supply pins), and 3 are left unconnected. The input and output pins of the high speed interfaces are 3.3V PECL, the remainder pins are (LV)TTL.

6. Conclusions

The AMDA circuit represents one step towards the future high speed broadband ATM network (2.5 Gbit/s). The main advantage of this high speed link rate is the exploitation of statistical gain [6], where the sum of the nodes peak rates can largely exceed the total link capacity as far as the sum of their mean rates do not exceed it. Furthermore, AMDA's media access control algorithm, guarantees an even bandwidth distribution between all network

nodes, and its distributed and scalable nature makes network expansion straightforward.

This circuit can be easily integrated into a fully functional high speed ATM network node implementation.

Acknowledgements

This work has been partially funded by the GAME within the AMDA project.

References

- [1] M. de Prycker. "Asynchronous Transfer Mode. Solution for Broadband ISDN". Ed. Ellis Horwood, 1991.
- [2] C. Reillo, J. C. Díaz, J. Riesco, L. Merayo and P. Lizcano. "A 2.5 Gb/s Mux/Dmux Implementation for ATM Architectures". IEEE 1995 Global Telecoms. Conference (GLOBECOM 95). November 13-17, 1995. Singapore.
- [3] J. Riesco, J. L. Conesa, C. Reillo, J. C. Díaz, L. Merayo. "Performance Evaluation of an ATM Network Using Verilog". EURO-DAC'96. September 16-20, 1996. Geneva (Switzerland).
- [4] "Distributed queue dual bus (DQDB) subnetwork of a metropolitan area network (MAN)", IEEE Std. 802.6, Jul. 1990.
- [5] J. Riesco, L. Merayo, A. Alonso, J. L. Conesa. "ADAM, Algoritmo Distribuido de Acceso a un Medio compartido". (Spanish). Patent request.
- [6] K. Lindberger. "Analytical Methods for the Traffical Problems with Statistical Multiplexing in ATM Networks". ITC 13. Copenhagen 1991.
- [7] "UTOPIA, An ATM-PHY Interface Specification. Level 1", Version 2.01. Mar 21, 1994. The ATM Forum.
- [8] "UTOPIA, An ATM-PHY Interface Specification. Level 2", Version 0.8. April 10, 1995. The ATM Forum/95-0114R1.
- [9] "B-ISDN ATM Layer Specification". ITU-T Rec. I.361 (11/95).
- [10] "B-ISDN user-network interface - Physical layer specification". ITU-T Rec. I.432 (03/93).
- [11] Micron MT58LC32K32/36C5. 32K x 32/36 SYNCBURST SRAM. Data sheet. Rev 1/96.