

The AM08 Associative Memory ASIC Design, Architecture and Evaluation methodology

Alberto Annovi
Sezione di Pisa
Istituto Nazionale di Fisica Nucleare (INFN)
Pisa, Italy
alberto.annovi@pi.infn.it

Alessandro Cerri
Department of Physics
University of Sussex
Falmer, UK
ORCID 0000-0002-1904-6661

Pascal Corona, Francesco Crescioli, David Martin, Eric Pierre
LPNHE
CNRS/IN2P3
Paris, France
{pascal.corona, francesco.crescioli, david.martin, eric.pierre} @lpnhe.in2p3.fr

Sebastian Dittmeier, Gunnar Föhner, Dirk Gottschalk, André Schöning
Physikalisches Institut
Ruprecht-Karls-Universität Heidelberg
Heidelberg, Germany
{dittmeier, foehner, gottschalk, schoning}@physi.uni-heidelberg.de

Luca Frontini, Valentino Liberali, Alberto Stabile
Dipartimento di Fisica, Università degli Studi di Milano
and INFN - Sezione di Milano
Milano, Italy
{luca.frontini, valentino.liberali, alberto.stabile}@unimi.it

Kostas Kordas
School of Physics
Aristotle University of Thessaloniki
Thessaloniki, Greece
kostas.kordas@cern.ch

Tommaso Lari
Sezione di Milano
INFN
Milano, Italy
tommaso.lari@mi.infn.it

Matteo Monti
Dipartimento di Milano
INFN
Milano, Italy
matteo.monti@mi.infn.it

Erdem Motuk, Matt Warren
Department of Physics and Astronomy
University College London
London, United Kingdom
{h.motuk, matthew.warren} @ucl.ac.uk

Andreas Vgenopoulos
School of Physics
Aristotle University of Thessaloniki
Thessaloniki, Greece
andreas.vgenopoulos@cern.ch

Abstract—The Associative Memory (AM) ASIC reached its version 8 in 2020 when it was submitted for fabrication. The AM08 has all the functionalities of the final chip, AM09, which was planned for the ATLAS experiment’s Hardware Track Trigger (HTT) system, at the High-Luminosity Large Hadron Collider (HL-LHC) at CERN. It is made in a 28nm CMOS technology with 10 metal layers and comes in a 15×15 FCBGA package. Being a digital chip with a full-custom CAM cell design that can store 12,000 patterns (16 bit \times 8 words per pattern), and can achieve 6.25×10^{12} comparisons per second. The design and architecture of the chip is presented in this paper. Additionally we discuss the behavioral simulations that run and also the generation of the test vectors purposed for industrial and in-house testing, in VCD and STIL file formats. The laboratory test-benches both for the bare-die chips and the packaged ones are also presented, including the related test-boards. Finally, we discuss the preliminary power measurements and compare these with the post-layout simulations.

Index Terms—associative, content addressable, memory, ASIC, pattern recognition, ATLAS, HEP, high-energy physics

I. INTRODUCTION

This paper outlines the design of the AM08 Application Specific Integrated Circuit (ASIC) and then details the evaluation and verification updates. The chip specifications were set in 2017, so it was under design for 3 years and submitted for production to IMEC and ultimately TSMC in December 2020. The production batch contains 100 chips, of which 15 were delivered without further processing (bare die chips) to the team in July 2021 while the remaining 85 are being packaged with a 12-layer substrate and are expected to be delivered in Spring 2022. The chip design was done by the AM08 team, and has been funded mainly by the ATLAS Collaboration. AM08 has the full functionality of the AM09pre/AM09 chip which has 384k patterns and is the version which was intended to be used in the ATLAS Phase-II Upgrade for the TDAQ system, and specifically in the Hardware-based Tracking for

the Trigger (HTT) system [1]. The Collaboration ultimately decided to adopt a commodity hardware approach based on commercially available CPU-servers, with the option to add accelerators (FPGAs and/or GPUs) to assist the CPUs, for the Phase-II TDAQ Event Filter tracking. In the next section the design and architecture of the chip are described.

II. DESIGN & ARCHITECTURE

This AM chip is a massively parallel system able to perform bit-wise comparisons of incoming data with the pre-stored templates (patterns). In total, the AM08 chip can store 12288 patterns organized in 3 cores of 4096 patterns, each core with a slightly different design and feature set. These cores are based on the KOXORAM+ Content Addressable Memory(CAM) technology developed by INFN. The maximum expected operating core clock speed is at 400MHz. Unlike previous versions of the chip, one of the cores of AM08 allows for the possibility of a fast read-back of the CAM contents. Regarding the area it occupies, the Total Area Estimation for AM08 is 8.67 mm² [2].

Very good power efficiency is expected from the AM08, achieving a 1femto-Joule/comparison/bit assuming an average bit-flip rate of 50%. The AM08 (and AM09) total Power Consumption is calculated by the following formula

$$P(AM08) = 1W + \langle \text{input rate} \rangle \times 0.05W/MHz$$

As regards the I/O interface: it operates in DDR mode LVDS18 I/O has been designed to work as baseline at 1.0 Gbit/s.

A. Outline of building blocks

The Functional Block Diagram of the AM chip is important for the overall understanding of the chip and it is shown in Fig. 1. A detailed, online pdf version of it can be found in page 8 of the official specifications document [2].

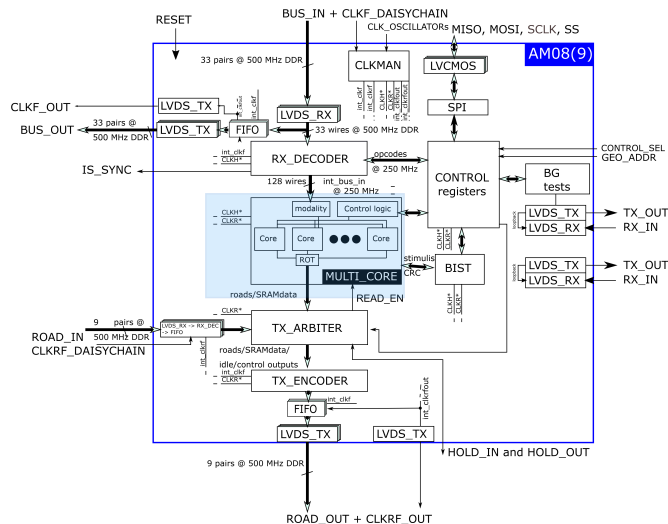


Fig. 1. Functional Block Diagram of the AM08 ASIC

The complete list of functional blocks that make up the AM08 are briefly described below:

The fast serial links of the chip are (LVDS18 protocol): BUS_IN (33 pairs), BUS_OUT(33pairs), ROAD_IN(9 pairs) ROAD_OUT(9 pairs), with I/O speed of 500MHz.

- **Cores:** in AM08 there are 3 cores, the processing-memory heart of the chip. There are 3 Enable-cores, of which only one supports the SRAM read-back mode. The Core, is made of the CAM Block or 64-block, containing 64 CAM cells plus the important Quorum logic near the output and supportive logic like Flip-Flops.
- **RX Decoder:** first stage to decode data, takes as *input* the 33-bit BUS_IN plus 1-bit control. Its output is two 128-bit buses, connected to the Cores and the Controller.
- **Cascade TX:** is used when a neighboring AM08 chip is connected. The BUS_IN data, are transmitted through this block and via the BUS_OUT. Minimizing skew for this procedure is needed and also *Clock Domain Crossing (CDC)* is implemented, in order to re-sync the data.
- **BIST:** Built-in Self Test, checks fast if there are stuck-one or -zero CAM cells.
- **ROAD Decoder:** propagates the ROADs from the ROAD_IN towards the TX Arbiter using a cascaded FIFO to change clock domains from CLKRF to CLKR. It can be disabled if the chip is the last in the daisy chain.
- **TX Arbiter:** This block merges multiple data sources (Controller, Cores, ROAD_IN) into a single stream and then passes it to the TX Encoder. It includes a HOLD signal which can stop data sending from a source, when received.
- **TX Encoder:** the block is responsible for reliably transferring data to the FPGA of the motherboard. It sends 8 data units of 32-bit length in 8 LVDS pairs, plus one control and one source-synchronous clock output (CLKRF_OUT).
- **Controller:** it contains the Control and Status Registers that can be accessed by either the input interface or the SPI slave interface. Access to the registers is done via a paging mechanism. The register page is available on [2].
- **SPI:** implements a version of the Serial Peripheral Interface (SPI) protocol with maximum operating frequency of 25MHz. It uses star topology for multiple slave configuration/operation.
- **CLKMAN:** the Clock Manager generates secondary clock signals from the CLKF reference clock, for the different clock domains which are CLKH, CLKR, CLKRF, CLKRFOUT, CLKFOUT.
- **BG Test:** includes two loopbacks each one using a RX_IN and a TX_OUT port.

B. The CAM Cell

The KOXORAM+ (and previously KOXORAM for the AM07 [3]) is a new type, full-custom CAM cell designed in HPC 28nm technology, implementing a Kill Signal which can disable individual cells and propagate from one to the next. Thus the switching activity of internal nodes is reduced,

minimizing power consumption.

The three cores of the AM08 are based on KOXORAM+, but arranged differently: CoreA has 64-pattern full custom blocks, CoreB has 256-pattern full custom blocks and CoreC has the same arrangement as CoreB with also the SRAM readout capability.

The KOXORAM+ block has been simulated and consumes 0.42 fJ/comp/bit.

C. Modes of Operation and the Comparison Mode

In total, the AM08 supports four modes of operation. The 2-bit STATE signal is used to inform all blocks about the common operations of the chip. The modes are identified by the following states

- STATE = 00 - Stand-by/Idle mode
- STATE = 01 - Write CAM
- STATE = 10 - Compare CAM
- STATE = 11 - SRAM-Read CAM

III. THE TEST BENCH SETUP

Due to global delays in the semiconductor industry, the package of the chip had a significantly longer lead time with respect to the die production. To obtain early results it was decided to plan a two step strategy: first, perform tests on the bare die at low data rate (limited by the technology used to connect to the bare die), and later perform high rate tests using an FPGA-based test board (see Section III-B).

A. For the bare die chips

All the components and the test-setup have been designed/developed by University of Heidelberg/PI. The PCBs developed are:

- 1) ADA78 V.10 Board which is the main control, readout and interconnects board and the
- 2) AM08 VPC x1 Socket Board on top of it, which embodies the bare-die chips under test.

The system is connected, via a breadboard interconnection, to an Arduino UNO micro-controller board, which in turn is attached via USB to a PC, for programming and for reading-out the measurements. The ADA78 is receiving an external clock, in our case from a SiliconLabs Clock Board. A photograph with these components is shown in Fig. 2

B. For the packaged chips

Full high speed tests will be performed with the use of a different test-setup based on a motherboard called **TB AM08 V1 Board**, which is designed and assembled in LPNHE. This board has a socket to mount the ADA75 or the AM08 VPC mezzanine, programmable DC-DC converters to provide and measure power to the chip, a programmable clock generator and it is equipped with an Enclustra Mercury+ AA1, which is an FPGA System-on-Chip (SoC). It incorporates the Intel® Arria® 10 SoC Module, which is made of an Intel FPGA and an ARM dual-core Cortex™-A9 processor [4]. The system will run the custom Firmware plus a very lightweight Linux distribution. This setup will be capable of handling both the




Fig. 2. Initial test-setup, to probe the bare-dies. Top photo: Overview, Bottom: The chip socket (cap opens)

packaged and the bare-die chips. A functional block diagram of the V1 Board is shown in Figure 3.

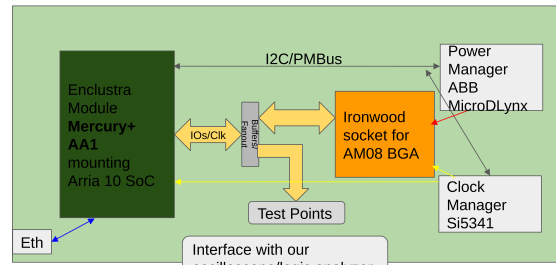


Fig. 3. Functional block diagram of the V1 motherboard, that will host the high speed and overall AM08 tests

IV. PRODUCTION OF TEST VECTORS

A. Industrial Testing

The initial plan was to produce a large number of AM09pre/AM09 chips for ATLAS, in the order of 10,000 chips; each one was planned to undergo exhaustive physical and functionality tests. Such tests are normally performed on Automatic Test Equipment (ATE), a type of specialized semiconductor testing equipment, and the chip which is being tested is called Device Under Test (DUT). For this reason, the team evaluated methods to produce the set of Test Vectors (or Patterns) required by this specialized equipment in order to use them and compare the chip behavior/outputs to these reference




Fig. 4. ADA75 - Board with socket to host the packaged AM08 chips

Vectors. Two established commercial software solutions were evaluated; VectorPro by Test Spectrum [5], and Solstice-TDS by TSSI [6].

The target file format is Standard Test Interface Language (STIL) for Digital Test Vector Data - a standard known as IEEE 1450-1999 [7]. It provides an interface between digital test generation tools and test equipment. Initially the WGL format was also considered. The creation of the primary test vectors (VCD) was done by running a set of Behavioral Simulations (of 36 Test Benches in total) and then by exporting the results in text files. This was done using Cadence Design Systems Xcelium software. So, this first stage format of the vectors is the Value Charge Dump (VCD), defined with Verilog in IEEE 1364-2001. Examples of testbenches written in SystemVerilog and simulate the various modes of operation are: `tb_SRAM` (simulates `STATE=3`), `tb_BIST` (for the Built-in Self Test), `tb_Regular_CoreB_10event6` where CoreB only is active, with input of 10 test events, the Quorum logic threshold is set to 6. Among others, `tb_SPI` is a generic testbench to send configurations to the chip, also extensively used in the simulations of Sect. V.

Production of vectors for the AM08 chip follows these steps in Fig.5. As for the evaluation of the software to use, both solstice-TDS and VectorPro, were given as input a subset of the 36 testbenches described above. As expected the resulting STIL files were of equal quality, and the user experience seemed to favour the final use of VectorPro, whose capabilities covered more than adequately the needs for STIL files production.

B. In-house Testing

Since the AM08 batch consists of 100 chips, the Industrial Testing is not necessary. However, the exported VCD files

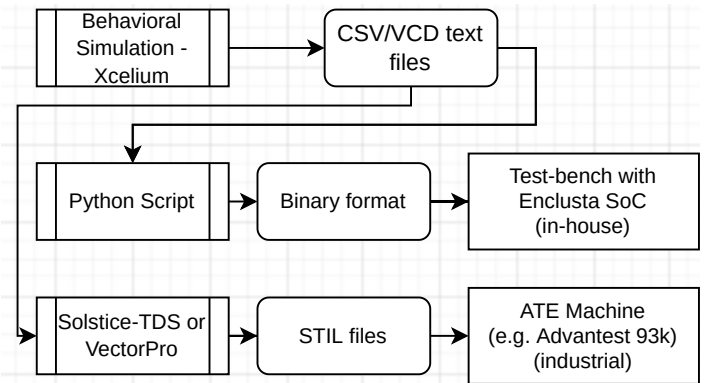


Fig. 5. Production of Test Vectors for in-house and industrial tests.

produced by the 36 Test Benches (or more if required) are used for both the Test Bench setups, described in Section III.

V. POWER MEASUREMENTS AND SIMULATIONS

A. Power Measurements with the test bench setup

As mentioned in Section III, the first deployed test setup is based on the Arduino, so some key results are shown in Table I. Column 1 shows the Reset status or the state of the chip as described in II-C, column 2 the core enable status, e.g. the value 111 means that all three cores are enabled in this measurement. Column 5 shows the total current measured for the 1.0 Volt rails, and the next, column 6, has the result of the same current but for the Gate-level simulation of the chip, using the same configuration as the Arduino test-bench.

TABLE I
POWER MEASUREMENTS ON BARE-DIES AND RELATED SIMULATIONS

STATE	Enable ABC	Clock Speed [MHz]	I(1.8V) [mA]	I(1.0V) [mA]	Sim 1.0V [mA]
reset	000	50	51.8	6	11.45
0	000	0	44.4	7.4	6.6
0	111	0	45.2	6.8	6.6
0	000	50	85.2	14	11.51
0	100	50	89.2	14.8	12.95
0	111	50	89.2	19	15.92
1	000	50	85.2	14	11.51
1	100	50	89.2	14.6	12.95
1	111	50	89.2	30.4	15.92
2	000	50	85.2	14	11.51
2	100	50	89.2	15.4	13.49
2	111	50	91.2	20	17.4
3	000	50	93.2	53.2	11.51
3	100	50	89.2	53.2	12.95
3	111	50	91.2	58.8	15.93

B. Post-Layout Power Simulations

In order to perform a power simulation it is necessary to produce first a VCD file with all the internal signals of the chip in the state in which estimate the power consumption. All the simulations that are presented are Gate-level simulations, distinct from the Behavioral simulations mentioned in IV-A, obtained running the simulation on the final netlist with backannotation. Once the VCD has been prepared it is

passed to the Cadence Voltus tool, a standalone power signoff tool. The tool estimates from the VCD the activity of every internal net, loads the power description of each cell and net and then computes the power consumption. We ran static power simulation, but also various kind dynamic simulations are possible. Runtimes for each simulation run, are in the order of one hour using parallel processing in a dedicated 16-core server.

Each Power simulation results in detailed report files containing power estimation for all different rails on the chip in the 1.8V and 1.0V domain. The results for the current are summarized in the last column of table I for direct comparison with the measurements on the bare die. Since the limits of the Arduino setup in terms of frequency are lower than the full high speed testbench, there is a set of measurements at 50MHz, and this frequency was also set as a simulation parameter.

C. Comparison

The simulation and the measurements are comparable and shows almost always similar trend in function of the chip configuration and as expected knowing the features of the chip (e.g. the power consumption increase with the number of cores enabled). In most of the cases the power prediction from the simulation is underestimating the measurement by 15-20%, possibly due to lack of detail in the power model of the full custom part.

Another possible source of difference is that in the measurements done on the bare die most of the inputs and outputs of the chip are left floating, while in the power simulation this is not taken into account. The measurements with the final test board will be in a more similar condition with respect to the simulation.

The most striking difference is observed with read back mode (state = 3): the power simulation predicts slightly more power consumption with respect to the other modes when CoreC is also enabled, while the measurements are a factor 3 more. It is not readily explained by the details of the full-custom models because read mode should affect only Core C, as predicted by the power simulation.

Another significant discrepancy not accountable as a generic offset in the power model of the full custom blocks is the fact that the write mode of CoreC is measured to consume significantly more than the other cores while this is not predicted by the power simulation.

These behaviours need to be investigated more thoroughly in the next campaign of tests with the final board (V1).

ACKNOWLEDGMENT

The development, testing and evaluation of the AM08 ASIC, has been part of the ATLAS Experiment Phase-II Upgrade program and mainly funded by the ATLAS Collaboration. Also funding was received from BMBF - Germany, Istituto Nazionale di Fisica Nucleare (INFN) - Italy.

REFERENCES

- [1] "Technical Design Report for the Phase-II Upgrade of the ATLAS TDAQ System," CERN, Geneva, Tech. Rep., Sep 2017. [Online]. Available: <https://cds.cern.ch/record/2285584>
- [2] A. Stabile, A. Annovi, M. Warren, B. Green, N. Konstantinidis, H. E. Motuk, L. Frontini, V. Liberali, F. Crescioli, G. Fedi, C.-I. Sotiropoulou, F. De Canio, G. Traversi, S. R. Shojaii, T. Kubota, G. Calderini, F. Palla, B. Ceccucci, L. A. Spiller, and P. C. Mcnamara, "Phase-II Associative Memory ASIC Specifications," INFN, CERN; LPNHE; University, Milano, Paris, Melbourne, London, Pisa, Perugia, Bergamo, Tech. Rep., May 2018. [Online]. Available: <https://cds.cern.ch/record/2320701>
- [3] A. Annovi, L. Frontini, V. Liberali, and A. Stabile, "Design and characterization of new content addressable memory cells," in *2018 IEEE International Symposium on Circuits and Systems (ISCAS)*, May 2018, pp. 1–5.
- [4] "Enclustra mercury+ aa1, official webpage," accessed: 08-03-2022. [Online]. Available: <https://www.enclustra.com/en/products/system-on-chip-modules/mercury-aa1>
- [5] "Test spectrum vectorpro, official webpage," accessed: 20-02-2022. [Online]. Available: <http://testspectrum.com/home/products/vectorpro>
- [6] "Tssi solstice-tds, official webpage," accessed: 20-02-2022. [Online]. Available: <https://www.tessi.com/solstice-tds>
- [7] "Ieee standard test interface language (stil) for digital test vector data," *IEEE Std 1450-1999*, pp. 1–140, Sep. 1999.