The UniBoard: a multi-purpose scalable high-performance computing platform for radio-astronomical applications

Arpad Szomoru[1]

[1]Joint Institute for VLBI in Europe
PO Box 2
7990 AA Dwingeloo
The Netherlands
szomoru@jive.nl

1. Introduction

The concept underlying the UniBoard was originally proposed by Sergei Pogrebenko, system scientist at the Joint Institute for VLBI in Europe (JIVE, where VLBI stands for Very Long Baseline Interferometry), in the late 90s. As one of the architects of the MarkIV correlator [1], he knew like no other the difficulties and technological pitfalls involved in the construction of such a complex instrument, consisting of many separate custom-made electronic components. A single-board, all-station correlator would do away with the need to transfer large, highly synchronized data streams between all these different components, with their associated complicated messaging systems and timing issues. Such a board should have all the CPU power that could be fitted on, in the form of Field Programmable Gate Arrays (FPGA), and as much I/O capacity as possible. However, in those days FPGA technology was not far enough advanced to make this feasible. When the preparation for the RadioNet FP7 proposal got underway in 2007, the situation had changed completely, with new generations of FPGAs combining massive computing power with ease of programming and fast development. The concept now became the basis for the UniBoard, one of the Joint Research Activities in the RadioNet FP7 programme (EC Contract no. 227290).

2. The Project

The UniBoard project has as its aim the creation of a generic high-performance FPGA-based computing platform for radio astronomy, along with the implementation of several applications (correlator, digital receiver, pulsar binning machine). It is a 3-year project that kicked off on January 1, 2009. Now past its half-way point, the first prototype board has been delivered and is undergoing tests, and design documents and a large amount of firmware have been produced. The board has attracted attention in the radio-astronomical community because of its high computing and I/O capacity, its potentially excellent computing/power consumption ratio and its use of generic interfaces. At this time concrete plans exist to use it as the basis for the next-generation Euopean VLBI Network (EVN) correlator, the Apertif correlator and beam former system and at least one all-dipole Low Frequency Array (LOFAR) correlator.

3. Project Structure and Participants

Originally the collaboration consisted of 7 participants, and in the course of the project two more partners joined. The first partners and their original roles in the project were:

- Joint Institute for VLBI in Europe (JIVE): project lead, VLBI correlator, control code
- Netherlands Institute for Radio Astronomy (ASTRON): hardware and test firmware development
- University of Manchester: pulsar binning machine
- Instituto Nazionale di Astrofisica (INAF): digital receiver
- Unversity of Bordeaux: digital receiver
- University of Orléans: Radio Frequency Interference (RFI) mitigation in pulsar binning application
- Korean Astronomy and Space Institute (KASI): VLBI correlator

978-1-4244-6051-9/11/$26.00 ©2011 IEEE

later joined by:

- Shanghai Observatory: VLBI correlator
- University of Oxford: all-dipole LOFAR correlator

At this time the VLBI correlator, digital receiver and pulsar binning machine are all under development, while the RFI mitigation project has expanded to include both pulsar binning and digital receiver applications. In addition to the original applications, work has started or is expected to start soon on an Apertif correlator and beam former (ASTRON), all-dipole LOFAR correlators (ASTRON + University of Amsterdam, University of Oxford). Several other applications are being considered.

4. The Hardware

At the start of the project, an inventory was made of the hardware requirements posed by the different applications. Considerations of price, availability and pin lay-out led to the selection of the Altera Stratix IV EP4SGX230KF40C2 chip (40 nm, 1288 18x18 multipliers, 14.3 Mb internal block RAM, 24 + 12 transceivers). With 1288 multipliers at 400 MHz each of these chips could yield a maximum of about 0.5 TMAC/s.

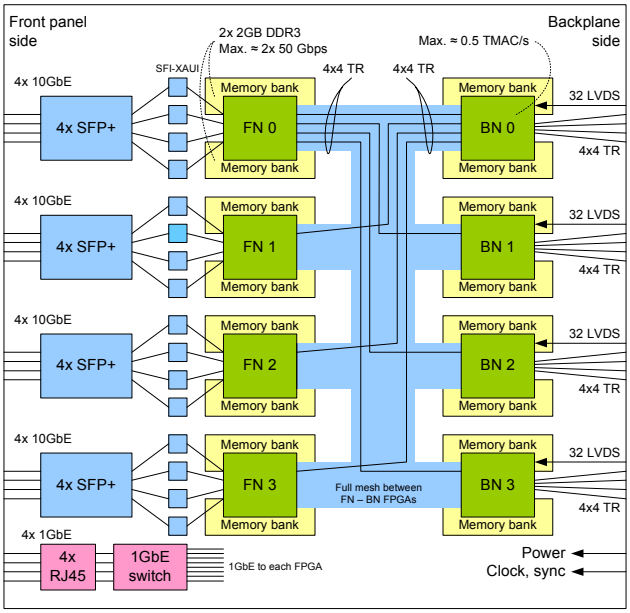

Figure 1: high level UniBoard design

The actual design was in the hands of Gijs Schoonderbeek and Sjouke Zwier, hardware engineers at ASTRON. A configuration of eight FPGAs per board was found to be optimal in terms of computing power, power consumption, density and complexity of the board (Figure 1). Each FPGA is connected to two Double Data Rate 3 (DDR3) memory banks, mounted on the back side of the board. Four times four 10GE links connect to the front nodes (FN) via four Small Form-factor Pluggable Plus (SFP+) cages. A high speed mesh connects each FN to all back nodes (BN). The BNs in their turn connect via four times four 8-bits Low Voltage Differential Signaling (LVDS) connections to a backplane connector. To make the board completely symmetrical, a 10GE break-out board (the XGB) has been designed in the form of a mini-backplane, with a total of 16 CX4 connectors. For system management, each FPGA also has a 1Gb/s Ethernet connection to an onboard ethernet switch, which offers 4 x 1Gb/s connectivity on RJ-45 connectors. The central power supply of -48V is distributed on the board via DC/DC convertors and regulators, the Printed Circuit Board (PCB) itself has 14 layers. Control and configuration are done via the embedded NIOS processor. The actual PCB production and board assembly were

outsourced. The prototype was delivered May 17, 2010 (Figure 2). No major design flaws were identified, although power consumption at full load turned out somewhat higher than the original estimate (~400W versus 280W).

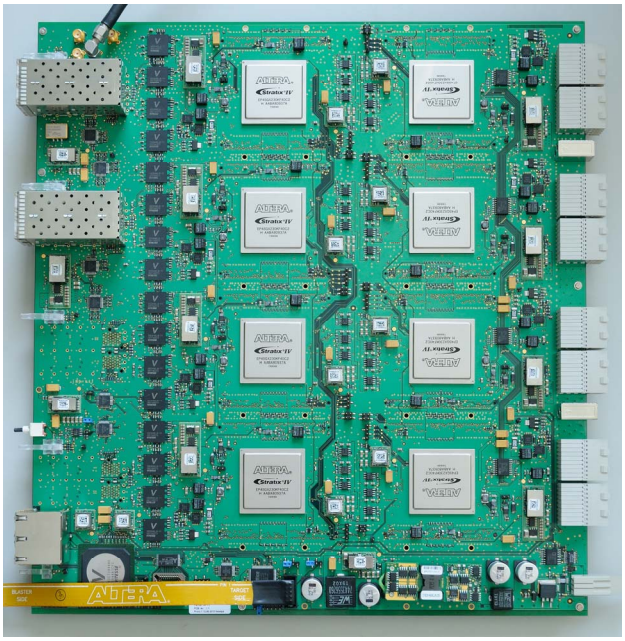

Figure 2: prototype UniBoard, deliverd May 17, 2010

5 The Applications

Throughout the board development phase, work on the various applications progressed; design documents were produced and refined, simulations were done and firmware was written. All documentation is posted on a project wiki, and all code is shared through a common repository. For board control Erlang was selected, a high-level programming language that provides robustness and completeness and enables a very short code development cycle. At the same time, a general correlator control system is being designed at JIVE. Several configurations for different applications are being considered, some of which are illustrated in Figures 3 and 4.

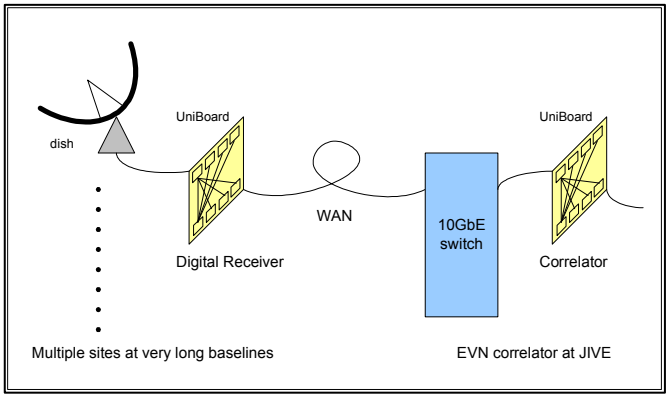

Figure 3: UniBoard as digital receiver and VLBI correlator, connected via internet

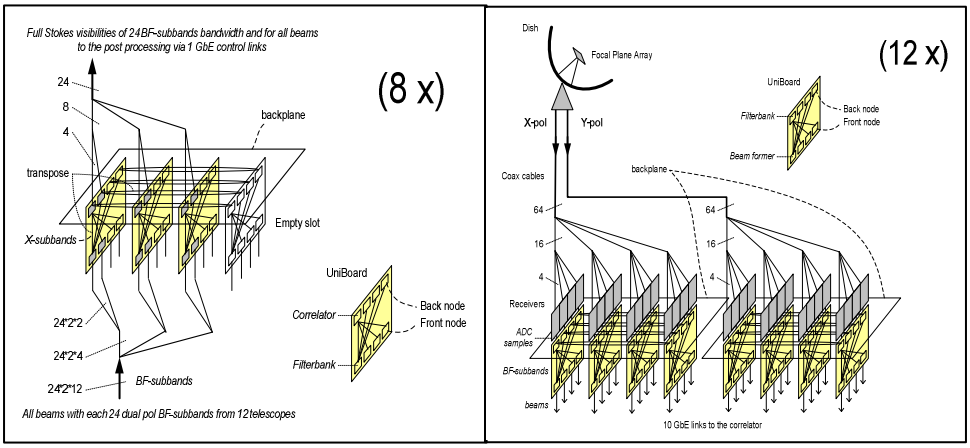

Figure 4: UniBoard as Apertif correlator (left) and beam former (right), interconnected via a custom-made backplane, with Analog to Digital Converters (ADC) connected to the opposite side of the backplane

6 Conclusion

The UniBoard project is well on its way. The performance of the prototype board has met the specifications without revealing any major design flaws. A production run is taking place now (February 2011) after which the hardware will be distributed among the partners. At the end of 2011 a number of working applications will be in place, demonstrating in practice the capabilities of the platform.

A Square Kilometer Array (SKA) Phase 1 correlator/beamformer could presumably be constructed using the UniBoard as a building block. However, it is obvious the SKA will not be built using 2009 technology. The current RadioNet project ends in 2012. If RadioNet3 receives EC funding, the UniBoard effort will continue in UniBoard[2], which has as its aim to create a completely re-designed platform with several innovative features, that will be production-ready by the time the construction of SKA gets underway. This new project will among others address power efficiency, first of all by using the newest available hardware (at this time this would mean replacing the current 40nm by 28nm FPGAs), but also by investigating techniques offered by FPGA manufacturers under names such as HardCopy and EasyPath. This enables one to develop on standard FPGAs and then to freeze the design into ASICs with the same footprint, allegedly cutting power consumption by as much as 50%. Other "green" measures will include the use of non-leaded components, the careful balancing of system parameters and performance and the optimisation of firmware designs and algorithms.

7 References

1. A. Whitney et al., "Mark 4 VLBI correlator: Architecture and algorithms", *Radio Science*, Vol. 39, RS1007, 2004, pp. 24.