Proceedings of the 16th European VLBI Network Symposium
Ros, E., Benke, P., Dzib, S., Rottmann, I., & Zensus, J.A. (eds.)
September 2nd-6th 2024, Bonn, Germany

DBBC4 - A 256 GHz bandwidth flexible VLBI environment

G. Tuccari[1,2], W. Alef[2], S Buttaccio[1], S. Dornbusch[2], A. Felke[2], A. Gupta[2], H. Rottmann[2], A. Roy[2], M. Wunderlich[2]

[1] INAF-Istituto di Radioastronomia,via Gobetti 101, 40129 Bologna, Italy
[2] Max-Planck-Institut für Radioastronomie, Auf dem Hügel 69, 53121 Bonn, Germany

Abstract. The development of the newer version of the VLBI digital front- and back-end system belonging to the DBBC systems' family is under way as a work-package of the EU-funded RadioBlocks Project. This instrument dedicated to increasing the VLBI observation capabilities in terms of bandwidth and output data rate involves a number of relevant novelties ranging from the full 32 GHz digitized input band in a number of up to 8 RF/IFs per system, to a fast recording system, and includes for the first time in a VLBI system a first implementation of a hardware AI processor.
The DBBC4 can be implemented as a distributed system of elements to be positioned in different parts of the radio-telescope site, in contrast to the more traditional way of in a single box. Distributed elements which can operate also stand-alone are: DiFrEnd28 with full 28 GHz bandwidth, DiFrEnd4T with 6 GHz bandwidth in the range $0 - 33$ GHz, DiFrEndVGOS covering the entire VGOS range $2 - 14$ GHz. This last one can also be used in conjunction with the DBBC3 to digitize the full VGOS band and replace any analogue frequency conversion. All the distributed elements can have backend functionality being able to produce channelized VDIF output packets to be recorded or to be sent through fibre to the correlator.
The progress of the DBBC4 development is reported together with a comprehensive description.

1. Introduction

The DBBC4 will be a key technical component in enabling new scientific applications in the rapidly evolving field of wide-band, multi-frequency astronomical and geodetic VLBI and will set a new standard in the area of VLBI backends. Technically, the DBBC4 will incorporate the latest state-of-the-art sampling technology enabling an increase in the processed bandwidth by a factor of eight compared to the DBBC3, the predecessor system and current de-facto standard for astronomical and geodetic VLBI. The use of artificial intelligence algorithms will allow a number of new possibilities ranging from the radio feature detection up to the mitigation of radio frequency interference (RFI) in near real-time. This in particular is one of the most severe issues to be addressed when increasing the observing bandwidth. The DBBC4 is the latest in the successful family of DBBC backends (DBBC, DBBC2, DBBC3) developed in a long-lasting collaboration between INAF (Istituto Nazionale di Astrofisica, Italy) and the MPIfR . The DBBC4 key technologies are based on and extend the developments of the BRAND (BRoad-bAND) digital receiver project covering the $1.5 - 15.5$ GHz band, which fully includes the $2 - 14$ GHz IVS VGOS band.

The DBBC4 backend is intended to offer the following capabilities and features:

- **Input bandwidth:** 274.4 GHz maximum full aggregate bandwidth realized by 8×28.8 GHz in digital front- or backend plus 8×5.5 GHz in ancillary digital front-end.
- **Output data rate:** up to: 1 Tbps @ 2-bit, 2 Tbps @ 4-bit, 4 Tbps @ 8-bit
- **Processing modes:** DSC (full band for data transfer), OCT (wide bands defined in the input band), DDC (narrow band tunable down-conversion)
- **New features:** Burst-mode, AI-mode for RFI-mitigation and transient detection, net-to-memory/disk capability

2. DBBC4 Architecture

The new system provides vastly greater bandwidth and agile signal processing capabilities, while simultaneously offering a feature-compatible upgrade path from earlier DBBC installations. The main architectural difference in the DBBC4 is that it is a distributed system. Traditionally, in radio-astronomical data acquisition systems, the digital back-end is well separated from the analogue part of the receiver, they being in widely-different physical locations. Typically the interface between analogue and digital domains, the sampler, is located with the digital system since these systems must run synchronously. This architectural choice was revisited in the BRAND-EVN project, where we split the traditional monolithic back-end into a network-attached ultra wideband digital front-end (called 'DI-FR-END') and a remote digital back-end; the front-end carries out simple digitization right at the receiver, while the bulk of signal processing is done by the digital backed at a different location where, e.g. cooling and RFI suppression is easier.

The choice to locate the sampler with the receiver offers us superior performance in terms of bandwidth, phase

stability, higher dynamic range, and offers the greater ease of transporting digital rather than analogue signals to the backend area. Such digital data transport is robust and simple and can carry pure sampled or preprocessed digital data. This architecture comes with its challenges though, in terms of RFI shielding between the digital sampler and the nearby and very sensitive front end.

Although we envisage that most DBBC4 installations would use the distributed configuration between digital front- and back-ends, it is still possible to combine the entire functionality in a single unit in the back-end area. This can be useful in particular when existing receivers are already routed to this area, or when very high frequency (sub-millimeter) receivers are used, and they include frequency conversions in the antenna focal area. To accommodate such solutions, dedicated analogue conditioning and sampler modules are provided.

The distributed system is considered not only for the digital front-end, but includes other possible 'dislocated' elements in support of the more advanced functions offered by the DBBC4. These elements are sensors which collaborate with the DBBC4 main unit to provide information in support of new functionalities. Some of those will be described later in this document, while still a larger number will be defined during the period of development and even at a later stage when the DBBC4 will be operational in the field.

The DBBC4 contains several functional entities already present in previous backend models, but introduces a number of new elements:

- 100GCoMo Module, analogue conditioning
- ADCore4 Module, A/D converter and digital processor
- FILA100G, data storage and network interface
- A-EYE, AI deep neural network controller
- DiFrEnd28, digital 28 GHz front-end,
- DiFrEndVGOS, implementation of the DiFrEnd28 dedicated to VGOS observations to be used even in conjunction with a DBBC3
- DiFrEnd4T, digital 40 GHz front-end
- CONE, a number of different elements with dedicated functionalities to operate with the A-EYE Controller
- ROD, a number of different elements with dedicated functionalities to operate with the A-EYE Controller

The signal coming from the analogue front-end requires conditioning to be applied before being converted to digital format. For this purpose, the 100GCoMo module performs the functions of optimizing the amplitude, measuring the total power in pre-determined frequency ranges inside the input band, and applying ad-hoc filters. The output signal from the 100GCoMo is connected to the analogue input of the ADCore4.

Alternatively to the analogue input at the DBBC4, the analogue signal of the receiver can be digitized by the digital front-end with 28 GHz bandwidth DiFrEnd28 at or close to the receiver and be transported and inserted into the system through the digital input. Similarly a 5.5 GHz bandwidth in the range up to 40 GHz input can feed the DiFrEnd4T part.

The ADCore4 is the central element of the 'control room' system and is able to perform the double functionality of analogue to digital conversion and digital data processing. After conversion, the functions as required by the particular observation are applied. The modes available are DSC, OCT, and DDC, as already well established in the previous versions of the DBBC. Improvements to these modes are being applied, but still maintaining compatibility with the existing modes. More details are provided in section 4.

The data with the final bandwidth and data rate, ready to be transferred to the correlator or to be recorded, are sent via the FILA100G to prepare the final aggregate format in single- or multi-stream, depending on the output data rate. Before the composition of the final format it is possible to store an amount of data to be used for the burst mode function. An additional possibility is offered by performing the data storage on external NVMe SSD units. Here, the direct connection net to PCI-e offers the possibility to skip any intermediate data transfer with great advantage to the writing data rate.

Notice particularly the newest addition in the DBBC family offered by the DBBC4, namely the Artificial Intelligence controller, called A-Eye. This has great potential in a number of modes it provides for both single-dish and VLBI observations, for example in RFI mitigation. To operate in real time, the controller can make use of a number of additional elements, named CONE and ROD. The first type supports the more complex functionalities to preprocess the signal than does the second, which simply forwards the required information to the mixed hardware-software deep neural network that performs the planned functionality. The A-Eye controller can then interact in both directions with the elements mentioned above to perform the required functionality. More details are described in the dedicated section of this document.

3. Broadband Analogue Conditioning Module - 100GCoMo

The 100GCoMo is the analogue conditioning module responsible for coupling the analogue input signal (0 – 40 GHz) to the digital conversion step. A DBBC4 system can contain up to four 100GCoMo modules (each processing two analog input signals). The core functionalities of 100GCoMo are automatic gain control (AGC) or manual power level control for optimal conversion of the signal with the 8-bit converter, and total power measurement in defined frequency ranges. Optionally, the component can contain a section for the ad-hoc band definition used by the DiFrEnd4T sampling unit (see Sec. 8). The communication with the general DBBC4 controller is realized through the traditional PCI method already adopted in previous DBBC systems. A 100GCoMo prototype unit covering the frequency range up to 33 GHz has been built and successfully tested.

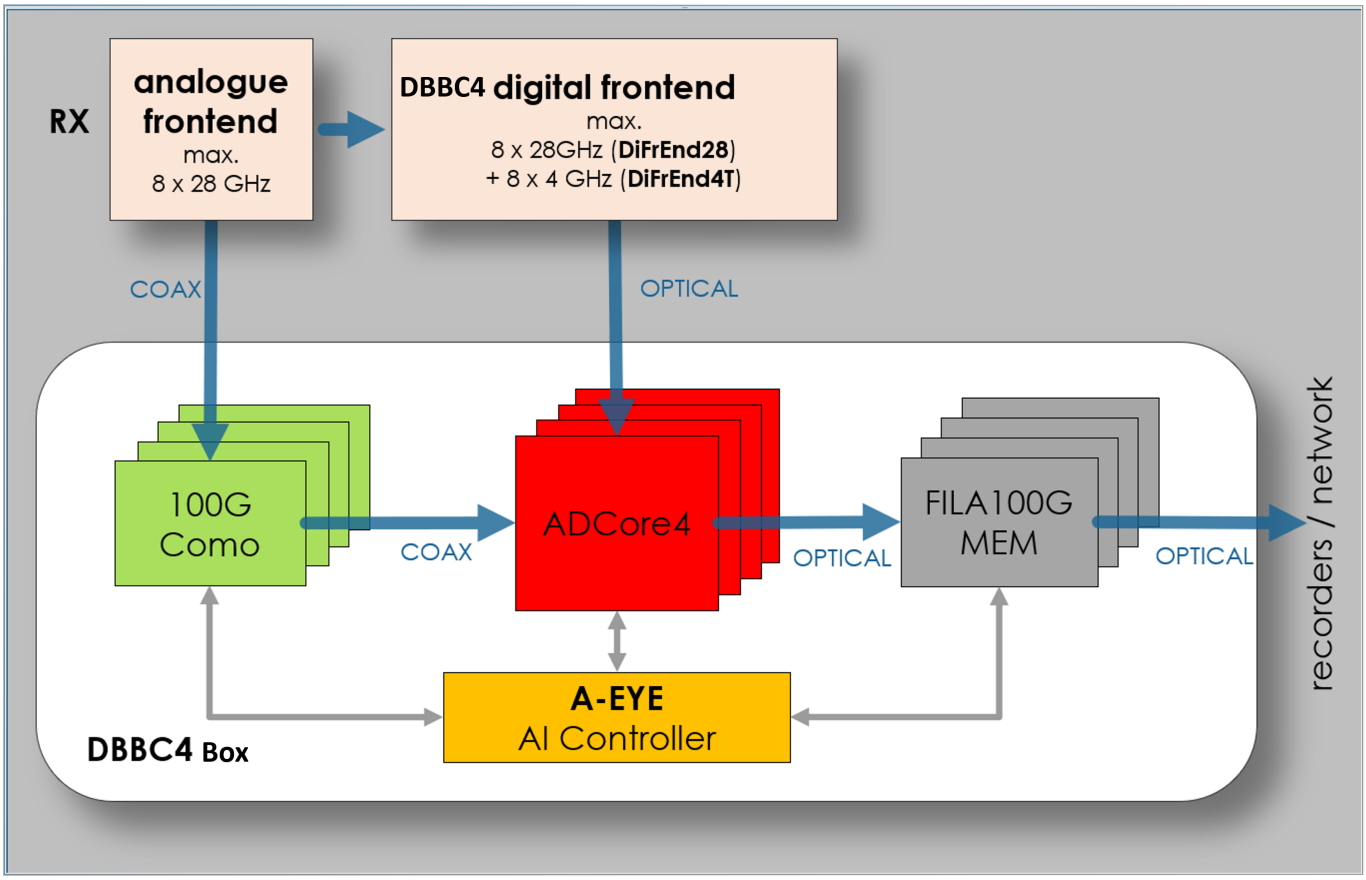

Fig. 1. DBBC4 architecture

4. AD Converter and processing unit - ADCore4

The module ADCore4 is a key element in the DBBC4 and again, similarly to the previous versions in the DBBC family, includes the functionality of the analogue to digital conversion and data processing, but with state of the art components and then capabilities. The flexibility in defining the functionality is like in the DBBC4 an important element. A number of different options can be selected in order to optimize performance and requirements. The functionality is straightforward: the ADB4 unit samples the signal conditioned by the 100GCoMo and then the digital format of it is transferred to the CORE4 FPGA using an aggregate channel of serial lines.The required functionality between DSC, OCT, DDC is then processed and available to the output channels in order to be sent out to the correlator/recorder or to the FILA100GMEM. The ADCore4 component can indeed be operated in two modes: it can either accept 2×28.8 GHz of *digitized* input bandwidth, or alternatively the ADCore4 can be equipped with an ADC stage (ADB4) performing direct sampling of two $0 - 28.8$ GHz *analogue* input signals. The optional sampler component will make use of a high-end specialized ASIC device capable of sampling at 2×57.6 Gsps. In both scenarios the data will be passed to the CORE4 FPGA element for the digital processing in the desired mode (DSC, OCT, DDC). The processed data will be sent out to the data recorders and/or the FILA100Mem component (see Sec. 5). An ADCore4 prototype has been realized and is now under testing.

5. FILA100GMEM

The FILA100G is an optional component that provides fast buffered memory for burst-mode operations and direct writing of the received packets to SSD NVMe (PCIe mode) disk modules. In addition the component will support functionalities present in previous versions of the DBBC systems (FILA10G) like channel reordering as well as channel extraction allowing e.g. streaming of sub-bands to a correlator for real-time fringe verification. Because packets will be recorded on-the-fly onto the fast NVMe disks without any CPU data handling, we expect to achieve burst speeds of 56 Gbps/disk over 80 s and sustained output data rates of 13 Gbps/disk. The module can allocate a variable number of SSD NVMe units as required by the burst mode duty cycle/number of channels/data rate. The module can allocate a variable number of SSD NVMe units, expandable in memory capability as required by the burst mode duty cycle/number of channels/data rate. More FILA100GMEM modules can be used in parallel in order to improve the memory capacity, or number of channels to be used in burst mode or to be recorded. It is important that disks do not require any CPU data handling for the digital data recorded ‘on the fly’ on the

NVMe disks. This solution allows a burst writing rate (80 s) up to 56 Gbps/disk and a sustained write rate up to 13 Gbps/disk. A dedicated board was defined for this mode which can be adopted in multiple samples for the required maximum data rate and recording time.

6. A-EYE Controller

Artificial Intelligence functionality meets VLBI technology. The A-EYE module is a controller making use of artificial intelligence methods to perform a number of functions useful for both single dish and interferometric observations. It will make use mainly of pre-trained networks, ready to be used for a number of modes which can include: RFI recognition and mitigation, extraction of non-statistical-noise signals, recognition of human-like extraterrestrial emissions, and other similar or different types of application.

The AI controller adopted for such functionality is a multi-CPU FPGA device optimized for this type of applications. The general development working flow consists in a session dedicated to select and pre-train a suitable DNN (deep neural network) configuration with additional training dedicated to the specific purpose to be satisfied. This configuration is then synthesized in a hardware DNN with mixed software dedicated implementation. The entire synthesized solution is running in the above mentioned programmable device which is able to interact with the other units in the DBBC4 in order to drive specific functionalities in the different components of the system. When a larger network is required a direct link is possible with a neural network operating on a cloud.

7. DiFrEnd28 and DiFrEndVGOS

The DiFrEnd28 is a standalone module expected to be allocated in the receiver area in order to minimize the connections between the analogue part and this section where the sampling with a 28.8 GHz bandwidth is performed and the output bands are generated. This unit can be fully independent or connected to the ADCore4 when additional functionalities available in the DBBC4 are required. Indeed the internal capability to generate OCT and DDC filters permits to use its functionality to directly connect a VLBI recorder or correlator through a number of optical fibers.

For VGOS observations, a dedicated unit has been developed making use of the same hardware, but running a specialised firmware version. The so-called DiFrEndVGOS component can be connected to a DBBC3 backend in order to perform the full VGOS sampling in dual-polarization, close to the receiver thus greatly reducing the required analogue connections to the sampling point. The unit can also act as a standalone front- and back-end offering the possibility to provide a large number of tunable DDC channels. The hardware design has been finalised and a prototype unit was realized, which is currently under test. Phase cal not could not be needed any more.

8. DiFrEnd4T

Complementary to the DiFrEnd28, the DiFrEnd4T is a standalone sampling unit which provides 6 GHz of sampled bandwidth in the range of 0-40 GHz. The desired portion of the band is to be selected by an appropriate filter. In the case of the DBBC4 the DiFrEnd4T is planned to cover the frequency range 27.5-33.0 GHz. Similar to the DiFrEnd28, the internal FPGAs can produce filtered output streams that can be directly recorded or processed by a correlator. The development of such unit is under way.

9. Conclusions

The DBBC4 will provide a VLBI front/backend system that is offering solutions for the technical challenges in the era of multi-frequency, large-bandwidth VLBI observations. The development of the DBBC4 VLBI backend is progressing as planned, with the first prototype system components already available and undergoing testing. Additional units are either under construction or at the design stage. Integration of the various components into a first prototype system to be used for end-to-end testing is estimated to be realized by 2025. A dedicated unit is planned to operate with the DBBC3 in in order to allow a direct connection with the receiver without any analogue filter with the possibility then to flexibly digitally select the input band for each DBBC3 IF.