HIGH THROUGHPUT FULLY PROCESSED PAYLOAD FOR BROADBAND ACCESS NETWORKS

Robert Hughes, Rosalind Warren, Alijah Ahmed and Alan Couchman, Astrium Ltd., Gunnels Wood Road, Stevenage, Hertfordshire SG1 2AS, UK. Tel. +44 (0)1438 313456, e-mail: robert.hughes@astrium.eads.net
Giuseppe Chiassarini and Eugenio Rossini, Space Engineering S.p.A., Via Dei Berio 91, I-00155, Rome, Italy, Tel. +39 06 22595224, e-mail: chiassarini@space.it
Piero Angeletti, ESTEC, Keplerlaan 1, Postbus 299, 2200 AG Noordwijk, The Netherlands, Tel. +31 (0)71 5656565, e-mail: piero.angeletti@esa.int

Abstract

In order to compete with terrestrial services, satellite systems for broadband access networks must maximise the useful capacity provided by each satellite. The High Throughput system has been carefully optimised in terms of capacity as part of an ESA R&D contract. This system provides Ka-band links between approximately one hundred user beams and a network of gateway stations over a regional European coverage from geostationary orbit. On-board digital processing is essential in providing the routing and frequency plan flexibility needed to use the resources efficiently across so many spot beams. Despite the large bandwidths involved, recent advances in ASIC, mixed signal conversion and digital interconnect technologies, make a fully processed payload handling up to 50 GHz of processed bandwidth feasible in the relatively near term. As well as fine channel demultiplexing, fully flexible routing and gain control, the digital processor also supports digital beamforming across a 500 MHz band, allowing for flexible frequency reuse and coverage that can be reconfigured in orbit. All these functions can be demonstrated at speed for a reduced number of processing chains using currently available commercial off-the-shelf FPGA-based processing cards.

Introduction

The High Throughput concept [1-2] was proposed as a future high capacity satellite system designed to maximise the traffic throughput of a single satellite in order to minimise the cost per user and thereby provide broadband access to a European mass market at prices competitive with the terrestrial networks. The broad spectrum available at Ka-band is needed to support the very high throughputs envisaged, which are in excess of 50 Gbps per satellite. To keep user terminal costs down a geostationary orbit is assumed for the satellite with a large number of high gain spot beams to meet the capacity, link budget and European coverage requirements. To maximise the utilisable capacity through the system, on-board digital processing is highly desirable to provide the necessary flexibility of routing and bandwidth allocation to beams.

The system concept, payload architecture and detailed design were developed under an initial ESA contract by Astrium and partners. A solution was proposed in which a variety of antenna types could be supported by a payload architecture with at its heart a transparent digital processor of modular and scalable design. Although the high bandwidths involved are challenging for a digital processor, the study found the implementation to be feasible with technology advances expected within a timescale of about 5 years. The key digital technology items required are for the most part already available for terrestrial use and the development of space-qualified equivalents is going ahead with the support of ESA-sponsored TRP and FP7 programmes. In the near term, Astrium's Next Generation Processor development [4] is based on a similar architecture and uses today's technology to provide a smaller scale product with generic applicability to a wide range of mobile and fixed satellite services. Meanwhile, a second High Throughput contract with ESA is now under way to develop a hardware demonstrator of a slice through the broadband digital processor using commercial off-the-shelf technology. The aim of this paper is to report progress on the development of the full scale High Throughput processor design towards a hardware realisation.

System and Payload Architecture

The High Throughput system architecture and its generalisation in subsequent work [3] is characterised as a distributed broadband access network. It provides high capacity point-to-point and multicast services for ubiquitous internet access. Network access is distributed between multiple network access providers (NAPs) sharing the satellite resources, resulting in a multi-star topology illustrated in the figure below. The space segment consists of a geostationary satellite handling forward and return Ka-band links between a large number of user spot beams (typically >100) and a smaller number of gateway beams (typically >10).

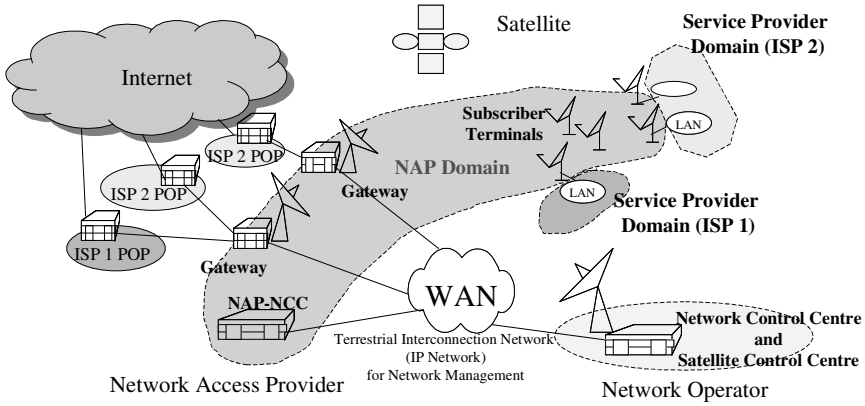

The reference system scenario considers a regional European coverage with user beams carrying up to 500 MHz bandwidth and the gateway beams up to 2 GHz per polarisation. Full connectivity is provided between gateway and user beams and the very large number (>1000) of possible links are supported by a digital processor providing fine granularity channel routing in order to make efficient use of the available bandwidth.

The target market is internet traffic and estimates of the forward:return capacity requirements vary from 4:1 to 1:1. The High Throughput processor is designed to be modular and scalable, both to reduce costs and to suit different mission scenarios. The reference system scenario assumes a DVB-S2 air interface on the forward link and DVB-RCS on the return, but the on-board processing is transparent so as not to restrict the choice of air interface. The adaptive coding and modulation provided by DVB-S2 is important to minimise the link margins required at Ka-band. Market analyses of traffic have given highly variable forecasts in terms of geographical take-up as well as link asymmetry and the expectation is that the traffic distribution will be highly non-uniform. These considerations serve to highlight a second key requirement of the High Throughput space segment that it must be very flexible. In addition to the flexible beam-to-beam routing, flexible allocation of capacity to each user beam is needed to serve traffic “hot spots”. This in turn requires flexible power to beam allocation and reconfigurable frequency reuse schemes. Ideally the coverage should be tunable, e.g. to allow narrow spot beams over high traffic density areas and wider beams over low density regions. These difficult requirements are best met by an active array-fed reflector (AFR) antenna system with on-board digital beamforming and flexible TWTAs providing the high power amplification. However, the additional mass and power required by such a complex payload reduces the maximum supportable capacity and so a variety of architectures including active and single feed per beam (SFPB) antennas is retained to support a range of missions.

The target spacecraft platform is the Alphabus currently in development and which can support payload power dissipations in the range 12-18 kW. Detailed payload modelling has shown that it will be possible to accommodate an advanced payload supporting ~140 user beams and a maximum throughput of 70 Gbps. Key technologies for such a payload will be the development of miniaturised and flexible Ka-band TWTAs and the use of high density packaging schemes for the RF and digital electronics. The calculated throughput is a theoretical figure based on a uniform traffic assumption and in reality the processed capacity will be oversized typically by a factor of two to support high traffic user beams with up to the full 500 MHz spectral allocation. Less oversizing is needed for the AFR architecture with digital beamforming in which the beam pattern, frequency reuse and bandwidth per

beam are all fully configurable. Several antenna designs have been considered including a 155-feed AFR with a 2.7 m diameter reflector and capable of supporting a European coverage with 100-175 user beams and a theoretical throughput of up to 50 Gbps. Dual polarisation schemes have also been considered for the AFR but this doubles the on-board processing required. A suitable compromise between payload complexity and performance can be achieved with a 112-feed dual polarisation array. Modelling of DRA antennas providing European coverage has also been performed using techniques to reduce the number of control points to within the 256 feeds that is the limit of scalability of the High Throughput DSP given the current technology assumptions. A simplified payload diagram representative of the AFR design and supporting up to 30 gateway beams is reproduced from [3] below.

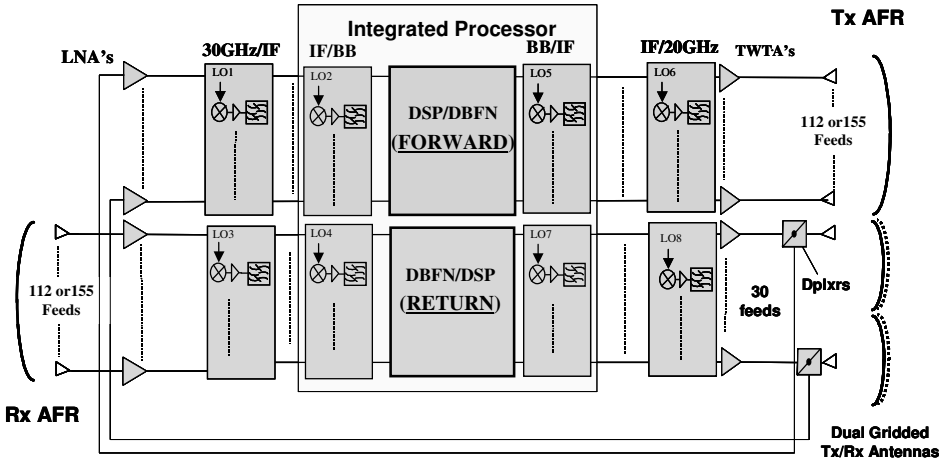

Full Scale Flight Processor Architecture

The digital signal processors (DSPs) sit at the heart of the payload performing fine channelisation, routing and digital beamforming functions on separate forward and return links. They also provide gain control and automatic level control on a per channel basis. The High Throughput design assumes processing of 500 MHz useful bandwidth per DSP interface port, based on the latest A/D and D/A converters which provide sampling clock rates well in excess of 1 GHz. This is obviously well matched to the 500 MHz Ka-band allocation for satellite user terminals. On the gateway side, additional analogue pre- and post- processing functions are required to separate up to 2 GHz of spectrum into 500 MHz sub-bands for digital processing.

The crucial challenge facing a high capacity DSP is minimisation of its power dissipation since this is in general proportional to the processed capacity. A target processed capacity of 50 GHz corresponds to a DSP with 100 active input and 100 active output ports each of 500 MHz useful bandwidth. This capacity is more than two orders of magnitude greater than that handled by the Inmarsat 4 processor, which represents the state of the art for currently flying DSPs, and which performs similar transparent processing functions within a mobile L-band system. The Inmarsat 4 DSP dissipates 1.8 kW of power, which is about the same power budget that is available for the DSP in the High Throughput system. Therefore the High Throughput DSP design must improve on the current state of the art by more than two orders of magnitude in power efficiency.

The required efficiency gains can be made by improvements in DSP algorithms, architecture and technology. Channel demultiplexing is performed using an efficient fast convolution algorithm which offers both a fine channel granularity of just 1.25 MHz and a low oversampling factor on the channelised outputs for a modest decrease in performance. Oversampling is needed to preserve the information in the transition bands of the channelisation filter to permit perfect reconstruction of the signals at the boundaries between frequency slots by the remultiplexing algorithm. In this way the processor supports wider channels with width that can be any multiple of 1.25 MHz. Minimisation of the oversampling factor reduces the rate at which the routing and beamforming functions must be performed and therefore reduces the power. Architectural improvements can be made by recognising

that an increasingly dominant contributor to the total dissipation is the power consumed in driving interconnect signals between the ASIC (application specific integrated circuit) components. Significant improvements are possible by optimising the DSP architecture to ensure that the digital beamforming network routes only those signals that are absolutely required to each feed element on the user link. This is achieved by combining the beamforming with the routing in the central rank of switch ASICs. An even larger improvement in power efficiency can be made by using the latest ASIC technology to implement more functions on each chip, thereby reducing the number of interconnect stages in the processing chain between input and output. The continued rapid advance of semiconductor technology since the Inmarsat 4 processor has made it possible to reduce the length of the chain from approximately ten to three ASICs, which is the minimum needed to retain a fully flexible switch network. The resulting architecture, shown below, is built up of chains of three ASICs, sandwiched between the A/D and D/A converters. Each chain consists of a Demux ASIC, a Switch/Beamformer ASIC and a Mux ASIC in series with cross-connections between chains to give full routing flexibility.

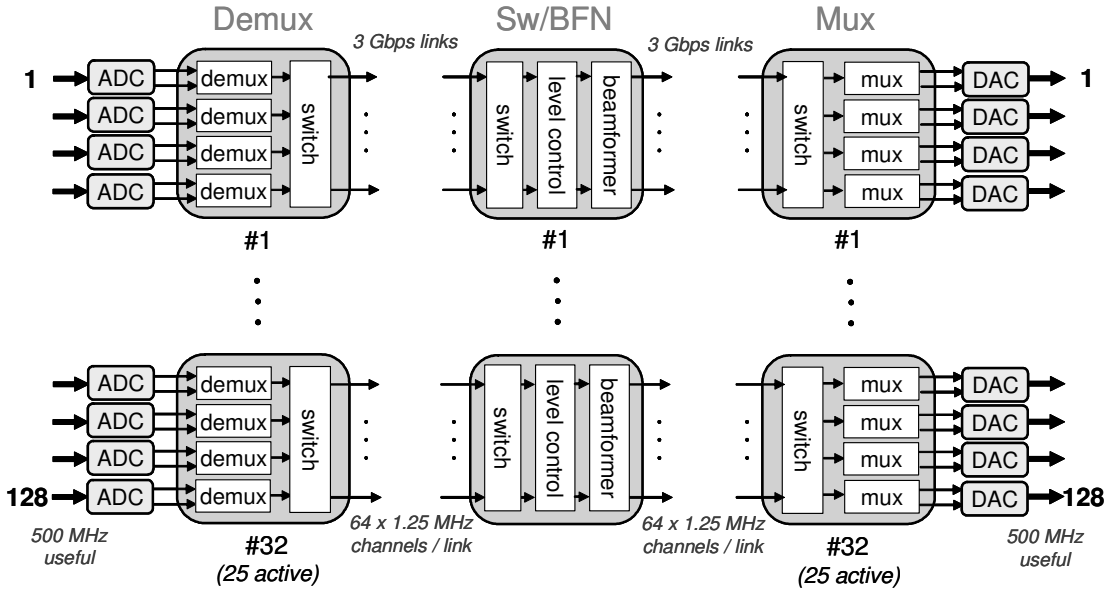

For SFPB payloads, the central rank of Switch/Beamformer ASICs operate in switch only mode with the beamforming disabled. Where an active antenna is employed, the beamforming function is activated, equipped with sufficient flexibility in the choice of beamforming parameters to encompass all the antenna types. Thus the same archietcture is appropriate, with or without an active antenna. The maximum number of input or output feeds is limited to 128 in the figure above based on the assumed maximum interconnect capacity of the central rank of ASICs. This can be extended to 256 feeds where required by splitting each beamforming function between a pair of ASICs and adding an additional rank of 2-way splitter/combiner ASICs, though this obviously increases the mass and power consumption. Any size processor up to these limits can be constructed by removing rows of ASICs. Internal redundancy is provided by providing spare ASIC rows and using the routing functions inherent in the design to perform redundancy switching.

The feasibility of the DSP design relies on technological advances in a number of key areas. Although the pace of development in microelectronics is such that terrestrial technology has already overtaken these requirements, qualification of the technologies for use in space still lags some way behind. The key areas for investment in technology qualification are:

- Large, deep sub-micron ASICs with a characteristic feature size of 130 nm or less. Radiation-hardened devices currently exist at 150-180 nm while commercial devices are available off-the-shelf at 65 or 90 nm.
- Fast, low power 10-bit data converters (ADCs and DACs) with sampling rates in the region of 1.3 GHz or more. The latest commercial devices are improving the performance and bringing down the power requirements of >1 GHz converters into the 1-2 W range needed to support the High Throughput mission.

- High speed serial interconnect operating at rates of 3 Gbps to carry the huge volume of signal data between ASICs. These are currently available in consumer electronics (hard disk drives) but little progress has been made on qualifying them for use in a radiation environment.
- A highly integrated, low mass packaging solution with improved thermal properties sufficient to remove 1-2 kW of dissipated power with typically 10 W dissipation per ASIC.

Given these key technologies, the estimated DSP power consumption is 1.3 kW for the 100-beam SFPB architecture and 1.8 kW for the 155-feed AFR design. Although there are challenges still to be met, the fact that such a large processor is becoming feasible only a few years after the launch of the first fully processed payloads demonstrates the rapid pace of advances in digital technology and the impact that this could have on future satellite systems.

Testing Digitally Processed Payloads

Current efforts to develop the High Throughput processor are focused around the design and implementation of a demonstrator using commercial FPGAs (field programmable gate arrays) in place of the flight ASICs. One clear goal of this exercise is to demonstrate the feasibility of the processor design at speed using representative technologies, including the data converters and the high speed serial links. However, a second principal aim is to validate the DSP design by means of a full characterisation of the equipment performance, including the digital beamforming function operating over a 500 MHz bandwidth. This is in many ways more challenging than for a classical analogue payload due to the effects of aliasing and subtle non-linearities due to quantisation effects.

One example that illustrates the difference between analogue and digital systems is in their response to sinusoidal signals. Single or two tone testing is a common measurement technique to measure the linearity of analogue systems. Low order harmonics or intermodulation products can be observed as spurious peaks in the output signal spectrum and compared in magnitude to the test tone. Whereas in this case the non-linearity typically only has low order (principally third order) terms, in a digital system the non-linearity is due to the staircase characteristic of the quantiser in the analogue-to-digital converter and subsequent arithmetic functions. The power series expansion of this characteristic has many thousands of small but significant terms. Furthermore, because of the discrete time sampling of the digital system, the harmonics of a sinusoidal input signal do not extend in a neat series beyond the Nyquist frequency but instead repeatedly alias back at baseband to form a quasi-uniform quantisation noise floor. However, the noise floor will not be uniform if the input signal tone frequency is coherently related to the sampling clock frequency by a rational factor since in this case sets of aliased harmonics will coincide to generate spurious peaks in the output spectrum. The spurious levels are a sensitive function of the test tone frequency and unrepresentative of the behaviour of the real multi-carrier system. These characteristics render sinusoidal testing of digital systems inappropriate.

Instead, the NPR or noise power ratio is commonly used as a figure of merit for digital signal processors. In this measurement the test signal consists of a flatband noise signal in which a narrow notch has been filtered out of the input spectrum. At the output, the noise introduced into the notch by the device under test is measured and compared to the power in the occupied part of the spectrum to give the NPR. This technique gives far more representative results for the multi-carrier signal scenario on which the DSP operates and is a standard measurement for high power amplifiers. In the case of the DSP, however, the NPR also includes a contribution due to adjacent channel interference since aliasing in the digital channelisation process superimposes incompletely rejected copies of many adjacent channels at a low level on the test channel. These individual contributions can be deduced from measurements in which the relative signal levels and positions of adjacent carriers are varied.

As part of the design validation it is necessary to elucidate all of these subtle effects and estimate their magnitude in order to determine the worst case performance, usually involving large numbers of measurements with varying signal scenarios. The validation activity often involves a large simulation effort to identify the various contributions to the overall signal degradation. However, this is limited by the long execution times and large number of possible multi-carrier signal scenarios involved for a complex transparent DSP. A key aim of the current demonstrator activity is therefore to supplement the simulation results with real-time measurements using a bit-precise FPGA emulation of the DSP

hardware. To this end, half of the FPGA resources are dedicated to providing a sophisticated signal generation and analysis testbed which allows considerable flexibility in the test signal scenarios, allowing the complex DSP behaviour to be fully investigated.

A detailed test plan has been proposed for the FPGA demonstrator. This includes both single input single output (SISO) measurements to characterise a single chain and multiple input or output (MISO and SIMO) measurements to characterise the performance of the digital beamforming network in receive and transmit modes of operation. SISO tests include overall SNR performance at feed level, in-band amplitude and phase responses and distortion and spurious characterisation. MISO and SIMO tests include amplitude and phase tracking measurements between feeds and overall SNR performance at beam level, which requires some modelling of the antenna to interpret the results. The tracking between feeds for digital beamforming is of crucial importance since, even at baseband, maintaining phase coherence between feeds over 500 MHz to within a few degrees translates into a timing uncertainty of only 25 ps. This is beyond the tolerances of the electronic components and requires an automatic calibration system to maintain coherence. The testbed will include facilities for calibration signal generation and fine-tunable digital delays (implemented as a Farrow interpolator) to demonstrate the calibration phase measurement performance under controlled circumstances.

FPGA Demonstrator

The FPGA demonstrator will be implemented using two separate FPGA racks, one emulating four chains from the full scale digital processor and representing the unit under test (UUT), the other constituting the testbed signal generation and analysis functions. Implementation of the testbed in FPGAs provides the desired flexibility and is a cheaper alternative to banks of specialist test equipment. Each rack includes its own A/D and D/A converters so that the signal links between them are analogue. This partitioning allows the UUT and its testbed to be developed independently, improving the quality of validation, while the analogue links allow the testbed to be substituted by or supplemented with standard test equipment. Independent clocks can be supplied to the two racks to avoid measurement artefacts due to coherence between signal and clock.

The rack representing the digital processor provides four inputs and four outputs, each of 500 MHz useful bandwidth and implemented using data converters clocked at the same rate as in the full scale processor. Each input is demultiplexed into 1.25 MHz frequency slots by a full channeliser function before routing and beamforming. The routing function is made as representative as possible of the full scale processor by implementing a subset of the switch network using full scale switch components and exercising the unused inputs by means of dummy traffic generation logic added to the processor module. The beamforming function is, however, necessarily cut down from the full scale version because only four elements of the array are implemented in the demonstrator. As well as the gain control, an RF sensing function is included within the processor. This measures the phase relationships between beacon signals incident on the uplink array elements in order to provide information as to the satellite pointing. After beamforming and routing of the signals to the four output chains, the frequency slots are remultiplexed into 500 MHz bands for digital-to-analogue conversion. The signal connections between FPGAs are implemented using high speed serial links running at approximately 3 GHz, as in the full scale processor architecture. The demonstrator implementation must therefore solve the same protocol and synchronisation issues that will be faced in the flight processor.

The signal generation and analysis functions are colocated on the second FPGA rack shown at the bottom of the figure. This includes the capability to generate four multi-carrier outputs of 500 MHz bandwidth to exercise the processor module. The test signals comprise modulated carriers, tones and high quality Gaussian noise which are generated from a pool of signal generation resources and routed as required to the four outputs. This routing includes a weighting function and a high resolution programmable delay in order to test the digital beamforming functions of the processor. For signal analysis purposes all four outputs from the processor module are sampled and a single test channel is extracted and passed to the measurement functions. These include spectral analysis, from which NPR measurements can be derived, cross-correlation for analysis of the beamforming network and a demodulation function which provides an E_b/N_0 measurement and allows analysis of carrier distortion.

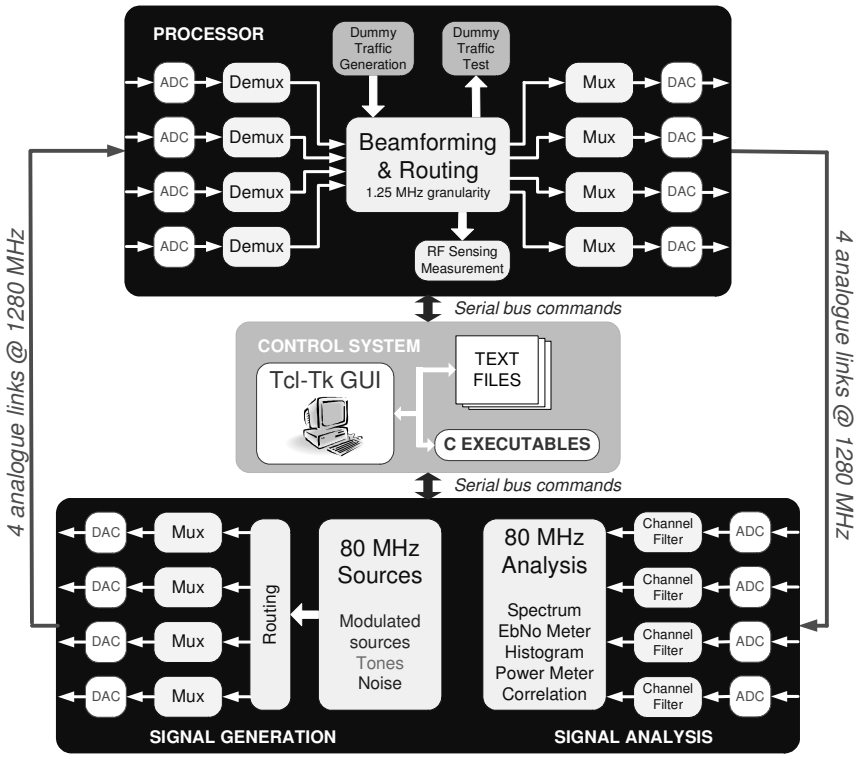

Both modules are configured via a control system comprising software running on a PC with fast serial interfaces to both FPGA racks. The control system is used to configure the signal generation scenario and the processor beamforming and routing functions in a coordinated fashion and to process and display measurement results from the signal analysis functions. A graphical user interface is provided to allow intuitive and interactive investigation of the processor functionality and performance while the system can also be script-driven to support long, repeatable test series. The control PC may also be connected to a network for remote access to the testbed.

To reduce costs, the processor unit and the signal generation and analysis functions are implemented on identical FPGA racks. Each rack comprises a central FPGA board and four I/O boards, each incorporating an ADC, a DAC and a powerful FPGA, connected in a star topology over a backplane using high speed serial links. Because of the high bandwidth, flexibility and signal quality required, the signal generation function is more demanding in terms of processing load than the processor unit under test. A major effort has gone into the design of the signal generation and analysis functions to ensure that the performance of the test system is very well characterised and all potential systematic errors are understood. As part of the test plan the testbed will be calibrated by connecting the signal generation and analysis functions back to back before integration with the processor unit.

As an example of the flexibility provided by the testbed, the architecture of the modulator component is illustrated in the figure below. Each modulated carrier is highly configurable, with programmable centre frequency, power, bandwidth, filtering and modulation scheme up to 64-QAM. Pseudo-random data to simulate the information stream is grouped into symbols and mapped to any of the DVB-S2 constellations. These symbols are passed through a pulse-shaping filter with programmable coefficients and rate changing to provide the carrier bandwidths of up to 80 MHz with much higher signal quality than can be achieved using look-up tables. Finally the carrier is mixed up to the desired centre frequency by a digital mixer using the CORDIC algorithm. Numerous modulators are provided to allow a wide range of frequency plans to be simulated, together with Gaussian noise generators to ensure that the whole 500 MHz band can be suitably loaded. Within the signal analysis an equivalent demodulation function is provided using the inverse algorithm. This outputs the recovered symbol centres for signal-to-noise and distortion analysis. The demonstrator architecture allows the demodulator to operate in the same clock domain as the modulator and be supplied with the known symbol stream so that implementation losses due to carrier and symbol timing recovery are avoided in the measurements.

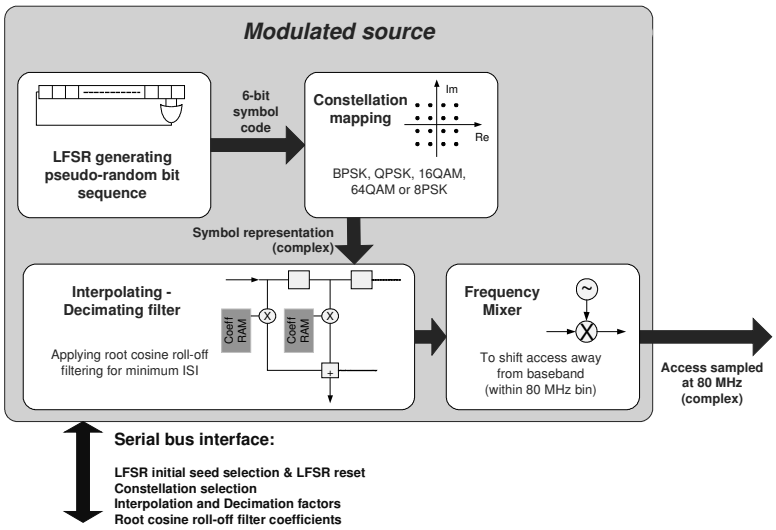

The integration of the whole UUT and testbed into a programmable system allows a fundamental limitation of the demonstrator, that it only supports four feeds, to be overcome. By building hardware triggers into the testbed, the demonstrator can coherently cycle through short signal sequences four elements at a time and combine them, according to an antenna model stored in the control system, to emulate the full digital beamforming network. This is in effect a hardware-accelerated simulation of the full processor and allows the beamforming performance to be mapped out over the coverage region taking into account the exact performance degradations of the digital processor, a task that would be impractically slow by software simulation.

Summary

The High Throughput system proposes an ambitious fully processed space segment with links at Ka-band to provide broadband internet access over a wide coverage region to mass markets at costs that are competitive with terrestrial services. Thanks to rapid improvements in digital technology and advances in design, high capacity on-board processing is becoming feasible, offering significant operational advantages to a Ka-band satellite communications system. Efforts are currently underway to implement a significant slice of the processor using commercial off-the-shelf FPGA technology. An integrated testbed with flexible signal generation is proposed to allow thorough design validation and full characterisation of the processor performance. The processor emulator is sufficiently representative to provide a powerful demonstration of on-board digital processing capabilities to be expected in the medium term.

References

[1] “High Throughput Processor For Future Broadband Satellite Access Networks”, X. Maufroid, F. Coromina, B. Folio, H. Göckler, H. Kopmann, 8th International Workshop Signal Processing for Space Communications (SPSC’03), Catania, Italy, Sept. 2003.

[2] “Next generation of transparent processors for broadband satellite access networks”, X. Maufroid, F. Coromina, B. Folio, R. Hughes, A. Couchman, S. Stirland, F. Joly, 22nd AIAA International Communications Satellite Systems Conference (ICSSC), Monterey, California, May 2004.

[3] “Next Generation Ka-Band Payloads”, S. Taylor, D. Jones, A. Couchman, R. Hughes, 10th Ka and Broadband Communications Conference, Vicenza, Italy, October 2004.

[4] “Next Generation Processed Payload”, D. Jones, P. Cornfield, 13th Ka and Broadband Communications Conference, Turin, Italy, September 2007.