

Architecture and DSP Implementation of a DVB-S2 Baseband Demodulator

P. Savvopoulos, N. Papandreou[†] and Th. Antonakopoulos

Department of Electrical and Computer Engineering

University of Patras, Rio-Patras, 26500, Greece

Email: psavvop@upatras.gr, npapandr@ece.upatras.gr, antonako@upatras.gr

Abstract—This paper presents the design and implementation of a baseband demodulator for DVB-S2 satellite receivers. In order to meet the requirements of different complex and multi-domain signal processing stages of the DVB-S2 baseband signal-flow, the presented architecture is based on efficient fixed-point implementation of the various demodulation algorithms and on the use of a dynamic time-sharing scheduler for the various DSP software tasks. The prototyping of the demodulator and its verification in the design of a complete digital DVB-S2 satellite receiver using a versatile testbed is also presented.

I. INTRODUCTION

DVB-S2 comprises the newest European standard for broadband satellite communications that exploits new achievements in the fields of modulation and coding [1], [2]. DVB-S2 meets the high performance requirements of today's satellite broadcasting and interactive communications, in terms of capacity and power efficiency, while keeping the complexity of the terminal receiver at acceptable levels. Both functional characteristics are based on the versatility of the DVB-S2 physical layer with frame-by-frame adaptability according to the channel conditions [3].

In order to support high data-rate applications, like multimedia applications and point-to-point backbone connectivity, the DVB-S2 standard supports high-order M-ary modulation formats, such as 16APSK and 32APSK, which provide high spectral efficiency and robustness under non-linear distortion introduced by the high-power amplifier (HPA) onboard the satellite [4]. In addition, since powerful channel coding (LDPC/BCH) is adopted [1], the receiver is able to recover the transmitted data under low signal-to-noise ratio (SNR) conditions. Therefore, the application of efficient synchronization techniques are imperative in order to achieve the high performance of the DVB-S2 standard.

In this paper, we present the architecture and implementation of a digital baseband demodulator for DVB-S2 satellite receivers. The demodulator's architecture is based on a powerful DSP processor with on-chip memory and external high-speed data ports, thus providing a stand-alone solution. The demodulator implements the full signal-processing chain at the physical layer including: symbol timing recovery, frame synchronization, carrier frequency and phase recovery, automatic gain control and constellation decoding. The integration of the different multi-domain signal processing stages in a single

processor is realized by the development of efficient fixed-point implementations of the various demodulation algorithms that exploit the parallel-instruction execution capabilities of the target CPU. Efficient time-scheduling of the various software tasks enables the dynamic intercommunication of the different signal processing stages that are executed sequentially or in parallel, thus enabling real-time execution.

The rest of this paper is organized as follows. Section II presents the functional requirements of the DVB-S2 demodulator and describes the circuits at the various baseband signal processing stages. Section III describes the architecture of the demodulator, discusses the implementation details of the various algorithms in DSP software threads and presents their time-scheduling according to the different operational modes. Finally, Section IV describes the prototyping and verification of the presented demodulator using a versatile testbed that utilizes powerful commercial equipment along with custom devices.

II. DVB-S2 DEMODULATOR CIRCUITS

The block diagram of a DVB-S2 terminal is presented in Fig. 1. It consists of the following main components: the outdoor unit (ODU), the analog-front end (AFE) and the digital receiver. The ODU consists of a parabolic dish antenna and the Low Noise Block (LNB) that converts the received RF signal into the respective L-Band (950-2150 MHz). The signal is further down-converted to the IF-Band (70 MHz) by a programmable analog tuner that includes output power control and fixed bandpass filtering. The incoming IF signal is digitized by a high-speed analog-to-digital converter (ADC) and then a second digital down-converter (DDC) provides the baseband complex I/Q signal to the digital signal processing receiver. Adjustment of the incoming IF signal to the dynamic range of the used ADC is achieved by an automatic gain controller (AGC) that is driven by a power measurement circuit implemented in the DDC.

The digital DVB-S2 receiver consists of the baseband demodulator, the forward error correction (FEC) decoder, the MPEG or IP data stream framer and the network interface. The demodulator implements the full baseband signal-processing chain including: symbol timing recovery, frame synchronization, carrier frequency and phase error recovery, automatic gain control and M-ary constellation decoding. In particular, QPSK, 8PSK, 16APSK and 32APSK modulation formats are

[†]N. Papandreou is currently with IBM Zurich Research Laboratory




Fig. 1. DVB-S2 receiver block diagram and detailed baseband demodulator signal processing stages.

supported. FEC decoding is carried out by the concatenation of an LDPC inner decoder and a BCH outer decoder [1]. A block de-interleaver is also preceded when high-order modulation formats, such as 8PSK, 16APSK and 32APSK, are used. DVB-S2 supports two major transport modes: the MPEG transport stream and the DVB-S2 generic stream.

This work is focused on the design and implementation of the digital demodulator. Fig. 1 also shows the various signal processing stages of the demodulator. Detailed description of the demodulator circuits follows.

A. Symbol Timing Recovery

The first stage of the demodulator is the symbol timing recovery (STR). The functionality of this module is to track the symbol rate fluctuations introduced during signal transmission. STR is implemented as a second order feedback loop utilizing a Farrow structured cubic interpolator along with the non-data-aided (NDA) timing error detector (TED) proposed by Gardner [5], which provides the loop with the required error signal. Gardner TED is capable of operating under random symbols and unknown carrier frequency offset error without precise carrier synchronization.

At the output of the STR interpolator, the synchronized samples stream is matched filtered and finally downsampled leading to a rate of one sample per symbol that feeds the rest of the signal processing chain. Fig. 3 shows the closed-loop STR using the NDA Gardner TED. The performance of the STR




Fig. 2. Differential PL header detection.

feedback loop under non-linearities, AWG noise and residual carrier frequency offset is studied in [6].

B. Frame Synchronization

After the STR has reached steady state, the next function to be performed is frame synchronization. This is done by searching the physical layer (PL) header using an appropriate correlator that operates on a symbol-by-symbol basis. Using differential detection, accurate frame synchronization is possible even in the presence of strong carrier frequency errors [7]. Fig. 2 shows the shift-register structure used for the detection of the DVB-S2 frame boundaries. Each frame starts with the

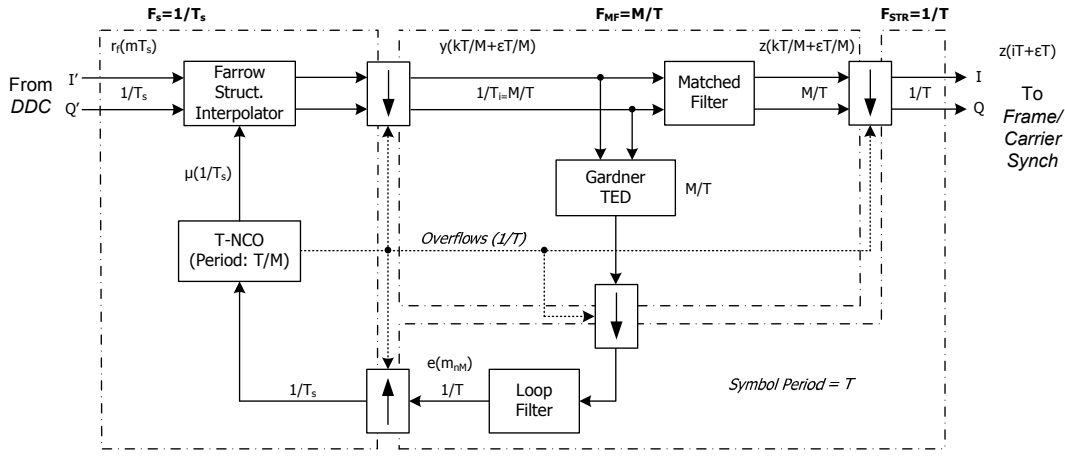


Fig. 3. Closed-Loop STR using the NDA Gardner TED.

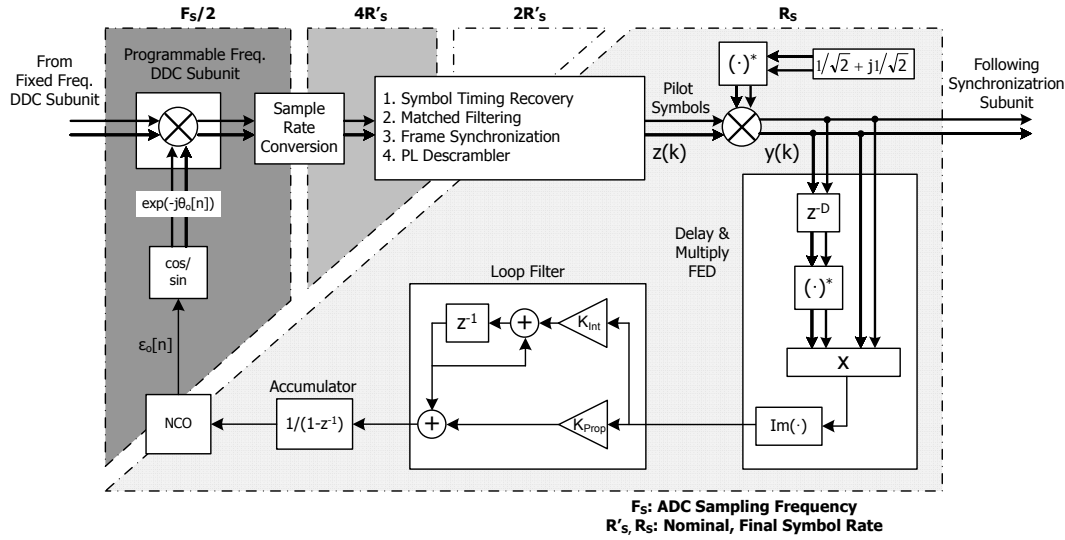


Fig. 4. Coarse carrier frequency recovery.

PL header that consists of two parts, the Start of Frame (SOF) and the Physical Layer Signalling code (PLSC). SOF is a known 26-symbol pattern, while PLSC is a 64-bit linear binary code [1]. The shift register in Fig 2 is partitioned into two sections. The first is associated with the SOF, the second with the PLSC. The output of the correlator drives a peak detector and has its maximum value when the whole PL header appears in the shift register.

A finite-state machine (FSM) controls the frame synchronization unit during acquisition and normal operation. After frame synchronization has locked, the data symbols (I/Q) of each frame are descrambled using a special Gold sequence in the complex domain [1]. The scrambling process at the transmitter provides symbol randomization and is used for energy dispersal. In our implementation, the descrambling sequence is precalculated and stored locally in the DSP memory. Based on the correct framing alignment, it is possible to demultiplex (DEMUX) the pilot symbols from the incoming

frame in order to drive the various pilot-aided carrier and phase synchronization circuits.

C. Carrier Frequency Recovery

Carrier frequency recovery (CFR) is invoked directly after the frame synchronization unit has detected the start of the transmitted frames and thus the locations of the respective pilot fields. CFR is performed in two sequential steps, which both use the known pilot symbols that are regularly repeated in the form of fields during the payload transmission of the DVB-S2 frames.

The first step of carrier synchronization consists of a coarse carrier recovery mechanism, which compensates large frequency offset errors up to several MHz and is implemented as a second order feedback loop based on a delay-and-multiply (DM) frequency error detector [8]–[10]. The compensation of this loop is applied through the numerical control oscillator (NCO) embedded inside the DDC logic. Fig. 4 shows the

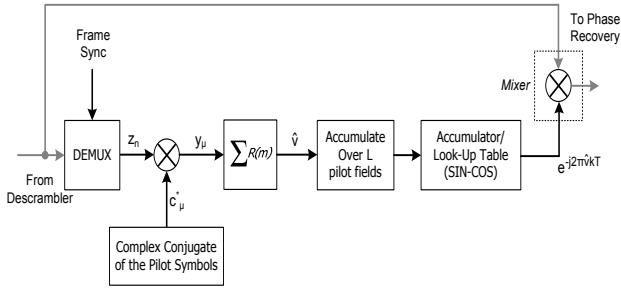


Fig. 5. Fine carrier frequency recovery.

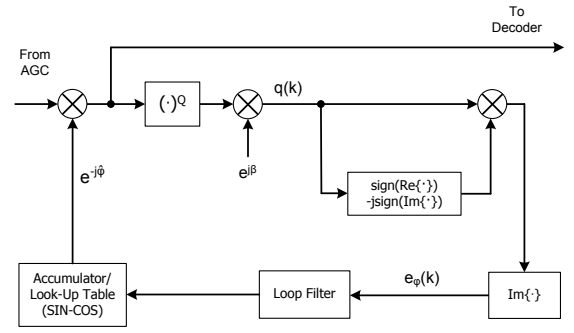


Fig. 7. Fine phase recovery.




Fig. 6. Coarse phase recovery.




Fig. 8. Digital automatic gain control (DAGC).

coarse CFR mechanism, where F_s is the sampling frequency, R_s is the nominal symbol frequency and R'_s is the recovered symbol frequency. After the convergence of the coarse carrier recovery loop, fine carrier recovery is initiated where the frequency offset error is in the order of a few hundreds of kHz. Fine frequency recovery deploys a feed-forward estimation algorithm, derived from an alteration of the L&R technique [11], and an integrator look-up table for the respective frequency offset removal. Fig. 5 depicts the fine frequency recovery scheme, which is also based on a closed loop structure.

D. Phase Recovery

Phase recovery needs to cope with residual frequency offset resulting from both carrier recovery procedures. In case of low order modulation transmissions (QPSK or 8PSK), a pilot-assisted maximum-likelihood (ML) feed-forward estimator is used for computing the average phase of each pilot field [8]. Phase compensation is based on the interpolation between the estimations of two consecutive pilot fields. When high-order modulations of 16APSK or 32APSK are used, an additional phase synchronizer is needed, which consists of a closed-loop based on the NDA phase error detector of Q-th power (Q=3 for 16APSK and Q=4 for 32APSK) [12]. This NDA feedback loop is located after the digital amplitude control unit, which is described in the next subsection, and is initiated as soon as the ML feed-forward coarse estimator has reached its steady state. Figures 6 and 7 show the coarse and fine phase recovery units respectively.

E. Digital Automatic Gain Control

Digital automatic gain control (DAGC) takes place after the first level of phase recovery and is based on a data-aided (pilot-assisted) vector tracker mechanism (DA-VT) [13], which utilizes the known pilot symbols in order to determine the amplitude multiplication factor. Therefore, the DAGC is activated during the transmission of the pilot fields and it is frozen during data symbol transmissions, where proper amplitude adjustments are applied. Fig. 8 shows the DAGC unit.

F. Constellation Decoder

After the various synchronization parameters have been compensated, the complex I/Q symbols of the incoming frame are demapped into the data bit sequence based on the used modulation efficiency. The DVB-S2 standard supports the following modulation formats: QPSK, 8PSK, 16APSK and 32APSK. The output bit sequence is grouped into coded data frames (blocks) that feed the FEC decoding unit.

Fig. 9 describes the format of the physical layer frame (PLFRAME) [1]. The FEC encoded data, i.e. the LDPC codeword, form the FECFRAME, the length of which is 64800 or 16200 bit depending on the type of the DVB-S2 frame, namely *normal* or *short* frame respectively. Based on the modulation efficiency, the FECFRAME is divided into S slots of complex symbols, which form the XFECFRAME. The size of each slot is 90 symbols. The number of slots is given by $S = \eta_{LDPC}/(90 \cdot \eta_{MOD})$, where η_{LDPC} is the size of the LDPC

TABLE I
PARAMETERS OF THE *normal* AND *short* PLFRAME

	normal frame: $\eta_{LDPC} = 64800$ bits				short frame: $\eta_{LDPC} = 16200$ bits			
η_{MOD}	S	α_{PIL}	K	η (%)	S	α_{PIL}	K	η (%)
QPSK: 2	360	22	33282	97.35	90	5	8370	96.77
8PSK: 3	240	14	22194	97.32	60	3	5598	96.46
16APSK: 4	180	11	16686	97.09	45	2	4212	96.15
32APSK: 5	144	8	13338	97.17	36	2	3402	95.24




Fig. 9. Format of the DVB-S2 PLFRAME.

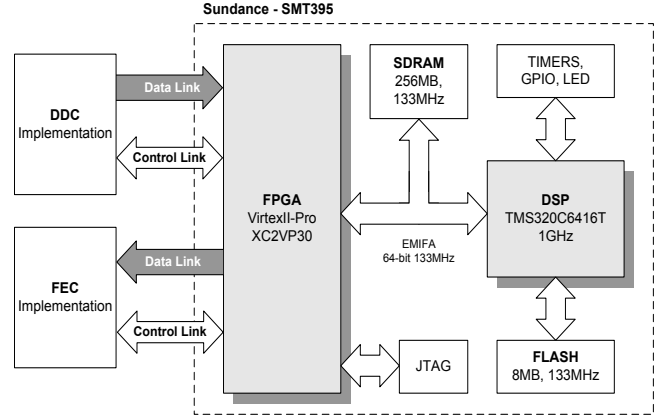


Fig. 10. Architecture of the demodulator's TIM development module.

codeword in bits and η_{MOD} is the modulation efficiency in bits-per-symbol.

In order to form the PLFRAME, a header is inserted in the beginning of each frame. If pilots are used, then a 36 symbols pilot field is inserted every 16 slots of data symbols. Given that a PLFRAME cannot terminate to a pilot field, the total number of pilot fields is $\alpha_{PIL} = \lfloor (S - 1)/16 \rfloor$. The total length of the PLFRAME in symbols is

$$K = \begin{cases} 90 \times (S + 1) & \text{without pilots} \\ 90 \times (S + 1) + 36 \times \alpha_{PIL} & \text{with pilots} \end{cases} \quad (1)$$

while the efficiency of the frame format is $\eta = (90 \times S)/K$. Table I displays the different parameters of the *normal* and *short* DVB-S2 frame for all constellation formats.

III. SOFTWARE ARCHITECTURE OF DEMODULATOR ALGORITHMS

The demodulator presented in this work comprises the basic component of a software-defined radio (SDR) implementation of a DVB-S2 satellite receiver developed on a powerful and reprogrammable hardware platform that utilizes four Texas Instrument Modules (TIMs) [14]. The demodulator is implemented in a single TIM that combines a C6416T DSP running at 1GHz along with a Virtex-II Pro (XC2VP30) FPGA [15]. As shown in Fig. 10, the demodulator's TIM is interconnected via dedicated high-speed data and control links with the digital

down-converter (DDC) unit, which feeds the input to the demodulator, and with the forward error correction (FEC) unit, which receives the output bit stream. The interconnection links are implemented in the FPGA. The development of the demodulator's subunits is based on a multi-thread DSP implementation, where the various signal processing algorithms are organized and executed into different threads.

A. Mapping of Demodulator Subunits into Threads

Fig. 11 shows the structure of the demodulator's thread-based implementation. The various signal processing units are divided into two threads. The first thread is fed with the complex I/Q samples derived from the DDC unit, which is located at a different TIM, and implements part of the demodulator's signal processing chain. The second thread completes the demodulation process and drives the output bitstream to another TIM that implements the FEC operations of the DVB-S2 receiver. The two threads are controlled by an appropriate interrupt handler associated with a DSP timer whose frequency is proportional to the sampling rate of the input signal.

As shown in Fig. 11, the DDC input samples are transferred to the demodulator's TIM via a dedicated interface logic. The interrupt handler holds a logical FIFO that stores the incoming samples. Using a frequency of twice the nominal symbol rate, in order to account for the STR mechanism (the Gardner




Fig. 11. Software thread architecture of the DVB-S2 demodulator.

detector needs 2 samples/symbol), the interrupt handler reads a block of N samples and then invokes the first thread by signalling a specific semaphore.

The first thread consists of the synchronization mechanisms that recover the actual symbol rate (STR), identify the start of a each new frame (PL framer) and compensate the carrier frequency offset errors (coarse and a fine CFR). These procedures are executed on a block of N samples that has been transferred from the logical FIFO to a dedicated input buffer, which is directly accessible by the various threads. The output of the first thread is enabled as long as the underlying synchronization stages have converged. A two output buffer strategy is used in order to preserve real-time operation. These buffers are named A and B in Fig. 11. The size of each buffer equals to the size of a PLFRAME. When a complete PLFRAME has been written into buffer A, the second thread is initiated via a dedicated semaphore for processing the stored samples. At the same time, the first thread stores the compensated samples of the next PLFRAME into buffer B. The interrupt handler provides efficient synchronization between the two threads in order to preserve real-time functionality.

The second thread, whose activation is frame-based, implements carrier phase recovery, automatic gain control and demapping. Carrier phase is recovered using two separate steps: a coarse step that is adequate for single ring constellations (QPSK and 8PSK) and a fine step, which is necessary for two and three ring constellations (16APSK and 32APSK). The automatic gain control mechanism is embedded between the two steps of the phase recovery, while the demapper comprises the last processing stage of the demodulator. The output of the second thread is the decoded bitstream that feeds the FEC unit of the DVB-S2 receiver.

In order to support the different modes of the DVB-S2 signal (see Table I), it is important that the first thread is able to complete the processing of each block of N samples between two successive interrupt signals, even when most of its processing stages are active. This requirement imposes that the first thread is called immediately after the interrupt, while the second thread is paused periodically by the interrupt

handler and resumed after the completion of the first thread.

B. Demodulator Operational Modes and Threads Execution

The presented SDR demodulator has three modes of operation. The first mode is called *initial acquisition* and is the initial mode after power-up. In this mode, the demodulator attempts to recover the symbol rate, as well as to determine the start of the PLFRAME. Therefore, during this mode only the STR and PL framer stages are executed in the first thread, while the second thread remains inactive. The operations during this mode are triggered by the timer-based interrupt handler, which signals the availability of a new block of N samples.

After the STR and PL framer subunits converge, the operational mode of the demodulator is switched to the second mode, which is called *frequency acquisition*. During this mode, the descrambling and carrier frequency recovery stages are also executed in the first thread, while the second thread remains inactive. In addition, the PL framer functions during this mode are reduced. The differential PL header detector is enabled only in a window of M symbols before and after the expected reception of the next frame header. This results in significant savings on CPU resources, which can be devoted to other functions. During this mode, the descrambler is enabled in order to generate the unscrambled symbol sequence. The descrambler operates on a complete PLFRAME excluding the frame header. The coarse frequency recovery mechanism is also enabled and drives the frequency control input of the DDC's NCO. The NCO is only updated whenever a pilot field is present. When the coarse recovery subunit converges, its processing is paused and the fine compensation algorithm is enabled. The latter is also pilot-based and is executed during the pilot fields. The operations during this mode are triggered by the timer-based interrupt handler, as soon as a new block of N input samples is available. The flowchart for the *initial acquisition* and *frequency acquisition* modes is shown in Fig. 12-(a).

When the fine frequency recovery subunit converges, the demodulator's mode changes to *normal operation*. During this




Fig. 12. Execution of threads according to the demodulator's mode of operation: (a) *initial acquisition and frequency acquisition*, (b) *normal operation*.

mode, we consider that all synchronization circuits of the first thread have converged and most of the errors have been compensated. Therefore, the first thread stores the recovered frames at the output buffers A or B and enables the second thread in order to process the data. During this mode the STR, the PL framer, the carrier frequency recovery and the output buffer writing subunits of the first thread are enabled along with all circuits of the second thread. The flowchart of the *normal operation* mode is depicted in Fig. 12-(b).

A critical restriction that affects the performance of the presented multi-thread DSP implementation is related with the total processing time of all threads that should not exceed the total duration of the transmission of a complete PLFRAME. The mapping of the demodulator's subunits as well as the dynamic scheduling of the various DSP circuits according to the operational modes described in this Section ensure that the processing of all subsequent frames is performed in real-time. Table II gives the duration of the *normal* DVB-S2 frame, as well as the duration of one data slot (90 symbols) and one pilot field (36 symbols) for 1Mbaud symbol rate and for all constellation formats.

IV. PROTOTYPING TESTBED

Fig. 13 shows the prototyping testbed, used for the development, testing and demonstration of the DVB-S2 demodulator. It consists of a versatile laboratory setup based on commercially available equipment and custom devices that enable the measurement and analysis of the DVB-S2 receiver functionality and performance under different signal conditions and operating modes. In particular, the setup consists of a set of DVB-S2 compliant signal generators, programmable up and down frequency conversion units, a noise injection and multi-channel fading conditions emulator, the hardware

TABLE II
DURATION OF *normal* PLFRAME WITH PILOT FIELDS FOR 1MBAUD SYMBOL RATE

η_{MOD}	Total PLFRAME (msec)	Data Slot (μ sec)	Pilot Field (μ sec)
QPSK: 2	33.282	90	36
8PSK: 3	22.194	90	36
16APSK: 4	16.686	90	36
32APSK: 5	13.338	90	36

prototyping platform and a host computing environment for visualization of measurements, diagnostics and signal statistics using a custom MATLAB-based application.

The DVB-S2 signal generator consists of a NEWTEC NTC2277 IF modulator at 70 MHz and an IP host traffic generator. System parameters such as baud rate, modulation and coding modes, roll-off factor and transmit power are user selectable. DVB-S2 signal generation can also be performed by Vector Signal Generators (VSG). The programmable frequency conversion unit (FCU) supports multiple signal paths, with independent L-band to IF and IF to L-band conversion modules, thus allowing the addition of various signal impairments at different signal processing stages. The FCU along with the VSG provides various signal distortion options, including white and colored Gaussian noise and RF interference. Additionally, the FCU can also be interfaced with an outdoor unit with a LNB that provides L-band signal under realistic conditions.




Fig. 13. The DVB-S2 receiver testbed.

V. CONCLUSIONS

In this work we presented the design and implementation of a digital baseband demodulator for DVB-S2 satellite receivers. The architecture of the digital demodulator is based on a multi-thread DSP implementation. Integration of the complex multi-domain signal processing stages of the DVB-S2 signal demodulation in a single processor is realized by efficient fixed-point implementations of the various demodulation algorithms along with dynamic time-scheduling of the different software tasks according to the operational modes of the receiver.

ACKNOWLEDGMENT

This work was supported by the European Regional Development Fund (ERDF) and the Greek Ministry of Development - General Secretariat for Research and Technology (GSRT) in the framework of project 'AKMON'.

REFERENCES

- [1] "Digital Video Broadcasting (DVB); Second generation framing structure, channel coding and modulation systems for Broadcasting, Interactive Services, News Gathering and other broadband satellite applications," ETSI EN 302 307 v1.1.1, 2005.
- [2] H. Zhijie, Y. Zhiqiang, Z. Ming, and W. Kuang, "8PSK demodulation for new generation DVB-S2," *In Proc. International Conference on Communications, Circuits and Systems - ICCAS 2004*, vol. 2, pp. 1447 – 1450, Jun. 2004.
- [3] A. Morello and V. Mignone, "DVB-S2: The second generation standard for satellite broad-band services," *Proc. IEEE*, vol. 94, no. 1, pp. 210–227, Jan. 2006.
- [4] S. Benedetto *et al.*, "MHOMS: High-speed ACM modem for satellite applications," *IEEE Wireless Communications*, vol. 12, pp. 66–77, Apr. 2005.
- [5] F. M. Gardner, "A BPSK/QPSK timing-error detector for sampled receivers," *IEEE Trans. Commun.*, vol. 34, no. 5, pp. 423–429, May 1986.
- [6] P. Savvopoulos and T. Antonakopoulos, "Comparative performance analysis of symbol timing recovery for DVB-S2 receivers," *In Proc. 13th Ka and Broadband Communications Conference*, Turin, Italy, Sep. 2007.
- [7] F.-W. Sun, Y. Jiang, and L.-N. Lee, "Frame synchronization and pilot structure for second generation DVB via satellites," *International Journal of Satellite Communications and Networking 2004*, vol. 22(3), pp. 319–339, 2004.
- [8] U. Mengali and A. N. D'Andrea, *Synchronization Techniques for Digital Receivers*. New York: Plenum, 1997.
- [9] J. W. Park, H. J. Yun, M. H. Sunwoo, P. Kim and D. Chang, "Efficient coarse frequency synchronizer using serial correlator for DVB-S2," *In Proc. IEEE International Symposium on Circuits and Systems - ISCAS 2008*, pp. 1520 – 1523, Seattle, WA, May 2008.
- [10] J. W. Park, M. H. Sunwoo, P. Kim, and D. Chang, "An efficient data-aided initial frequency synchronizer for DVB-S2," *In Proc. IEEE Workshop on Signal Processing Systems - SiPS 2007*, pp. 645 – 650, Shanghai, China, Oct. 2007.
- [11] M. Luise and R. Reggiannini, "Carrier frequency recovery in all-digital modems for burst mode transmissions," *IEEE Trans. Commun.*, vol. 43, no. 2/3/4, pp. 1169–1178, Feb/Mar/Apr 1995.
- [12] E. Casini, R. De Gaudenzi, and A. Ginesi, "DVB-S2 modem algorithms design and performance over typical satellite channels," *International Journal of Satellite Communications and Networking 2004*, vol. 22(3), pp. 281–318, 2004.
- [13] R. De Gaudenzi and M. Luise, "Analysis and design of an all-digital demodulator for Trellis coded 16-QAM transmission over a non-linear satellite channel," *IEEE Trans. Commun.*, vol. 43, no. 2/3/4, pp. 659–668, Feb/Mar/Apr 1995.
- [14] P. Savvopoulos, N. Papanandreou, and T. Antonakopoulos, "The architecture of a software radio DVB-S2 receiver," *In Proc. 13th Ka and Broadband Communications Conference*, Turin, Italy, Sep. 2007.
- [15] Sundance Multiprocessor Technology Ltd., "SMT 395 User Manual," ver. 3.0, Feb. 2001.