♦ Design and Development of a Terrestrial Digital Video Broadcast Demodulation Core: An International Collaborative Effort

Alexei Ashikhmin, Adriaan J. de Lind van Wijngaarden, Zhao Haibo, Bertrand M. Hochwald, Thomas L. Marzetta, Vinay Purohit, Chen Qinghong, Paul A. Wilford, Sheng-Rong Zhou, Michael A. Zuniga, and Edward S. Zuranski

The terrestrial digital video broadcast (DVB-T) system, specified in the ETSI 300-744 standard, is currently the most popular digital television system with more than 120 million receivers deployed globally in integrated television sets, set-top boxes, in-car modules, universal serial bus (USB) sticks, and personal computer (PC) cards. A provider of multimedia application specific integrated circuits (ASICs) in Asia required a DVB-T demodulation core as part of its strategy to broaden its stake in this growing market. To meet this challenge, Bell Labs partnered with Cambridge Industries Group, China, to form an international team to design and implement the underlying multimode orthogonal frequency division multiplexing (OFDM)–based receiver. This paper presents insight into the essential elements of the system design, the design methodology, and project management and describes the underlying signal processing algorithms, the bit accurate model, and register transfer level code. It is demonstrated how close collaboration between Bell Labs and selected global partners can yield added value to Alcatel-Lucent.
© 2007 Alcatel-Lucent.

Introduction

Digital video broadcasting is rapidly emerging and replacing traditional analog video broadcasting systems in an effort to provide higher quality and more flexibility. Standardization efforts are driven by the Digital Video Broadcasting (DVB) Project, an industry-led consortium with more than 250 members, comprising broadcasters, manufacturers, network operators, software developers, and regulatory bodies from more than 35 countries. This worldwide alliance oversees the development of open standards for the global delivery of digital television and data services, which are published by a Joint Technical Committee (JTC) of the European Telecommunications Standards Institute (ETSI), the European Committee for Electrotechnical Standardization (CENELEC), and the European Broadcast Union (EBU). As a result, a suite of standards regulate the distribution of digital video data by satellite (DVB-S), cable (DVB-C), terrestrial transmission (DVB-T), and terrestrial transmission for handheld devices (DVB-H). Services using DVB standards are available on every continent with more than 120 million DVB receivers deployed to date.

Bell Labs Technical Journal 12(2), 97–118 (2007) © 2007 Alcatel-Lucent. Published by Wiley Periodicals, Inc.
Published online in Wiley InterScience (www.interscience.wiley.com). • 10.1002/bltj.20239

WILEY InterScience® DISCOVER SOMETHING GREAT

Panel 1. Abbreviations, Acronyms, and Terms

ACI—Adjacent channel interference
ADC—Analog-to-digital converter
AGC—Automatic gain control
ASIC—Application specific integrated circuit
ATSC—Advanced Television Systems Committee
AWGN—Additive white Gaussian noise
BAM—Bit accurate model
BCH—Bose-Chaudhuri-Hocquenghem
BER—Bit-error rate
CCI—Co-channel interference
CENELEC—European Committee for Electrotechnical Standardization
C/N—Carrier to noise
DVB—Digital video broadcasting
DVB-C—Digital video broadcasting-cable
DVB-H—Digital video broadcasting-handheld
DVB-S—Digital video broadcasting-satellite
DVB-T—Digital video broadcasting-terrestrial
EBU—European Broadcasting Union
EDTV—Enhanced definition television
ETSI—European Telecommunications Standards Institute
FEC—Forward error correction
FFT—Fast Fourier transform
FIFO—First-in, first-out
FPGA—Field programmable gate array
FSM—Finite state machine
HDTV—High definition television
HP—High priority
I^2C—Inter-integrated circuit
ICI—Inter-channel interference
IF—Intermediate frequency
IP—Internet Protocol
IPTV—Internet Protocol television
ISI—Intersymbol interference
I/Q—In-phase/quadrature-phase
JTC—Joint Technical Committee
LDPC—Low-density parity-check
LDTV—Limited definition television
LP—Low priority
LPF—Low-pass filter
LUT—Look-up table
MPEG—Motion Picture Experts Group
NCO—Numerically controlled oscillator
OFDM—Orthogonal frequency division multiplexing
PAL—Phase alternating line
PC—Personal computer
PRBS—Pseudo-random binary sequence
QAM—Quadrature amplitude modulation
QEF—Quasi-error-free
QPSK—Quadrature phase shift keying
RAM—Random access memory
RF—Radio frequency
ROM—Read only memory
RS—Reed-Solomon
RTL—Register transfer level
SDTV—Standard definition television
SFN—Single-frequency network
TPS—Transmission parameter signaling
TS—Transport stream
UHF—Ultra high frequency
USB—Universal serial bus
VHDL—VHSIC Hardware Description Language
VHF—Very high frequency
VHSIC—Very high speed integrated circuit

The DVB-T system is the most popular digital terrestrial television system in the world, adopted in more countries than any other. It has been successfully deployed in the United Kingdom, Germany, Sweden, Finland, Spain, Italy, the Netherlands, Switzerland, Singapore, Taiwan, and Australia. DVB-T trials are ongoing in China, Malaysia, Thailand, Vietnam, Ukraine, Azerbaijan, Croatia, South Africa, and other countries.

The DVB-T system, specified in the ETSI 300-744 standard [2], has been designed for digital terrestrial television services, targeting receivers in homes and offices, and to some extent mobile receivers [5, 9]. In order to use the existing very high frequency (VHF) and ultra high frequency (UHF) bands for analog transmission effectively, the DVB-T system must operate in three nominal bandwidths at 6 MHz, 7 MHz, and 8 MHz and provide sufficient protection against high levels of co-channel interference (CCI) and adjacent channel interference (ACI) emanating from existing broadcast services. The DVB-T system uses two OFDM modes with different levels of quadrature phase shift keying (QPSK) modulation and quadrature amplitude modulation (QAM), and a multirate concatenated forward error correction (FEC). This helps the system to meet the bandwidth and efficiency requirements and to support two-level hierarchical

channel coding and modulation, including multiresolution constellations for the transmission of two independent Motion Picture Experts Group (MPEG) transport streams. A program service can thus be simulcast as a low bit rate, error resilient version or as a higher quality, but more error-prone version. Alternatively, entirely different programs can be transmitted on the separate streams with different levels of ruggedness. Good spectral efficiency in the VHF and UHF bands can be obtained by using a single frequency network (SFN) of DVB-T transmitters that operate synchronously to suppress interference.

Bell Labs has a strong interest in the research and development of DVB because of the growing importance of video to telecommunications service providers, and consistently with its rich technical tradition in television-related research and development, which dates back to the demonstration of the first over-the-air broadcast of television in 1927. Subsequent milestones include developing the terrestrial networks that enabled nationwide transmission of television signals, demonstrating the first television broadcast across the Atlantic with the *Telstar* satellite in 1962, and developing OFDM, which is the basis for DVB and other wireless communications systems. More recent milestones include developing the first-generation high definition television (HDTV) demodulator and decoder chipset that are used by many major television manufacturers, and its pioneering contributions to the Advanced Television Systems Committee (ATSC) HDTV system, for which Bell Labs received an Emmy Award in 1997. Current research focuses on integrating Internet Protocol television (IPTV) with Alcatel-Lucent's IP Multimedia Subsystem (IMS) solution to enable network operators to offer compelling new services that integrate telephony features, Web browsing, and television capabilities in exciting ways.

Early in 2005, Bell Labs responded to an opportunity to develop a DVB-T demodulation core for a provider of multimedia ASICs in Asia as part of its strategy to broaden its stake in the growing DVB-T/DVB-H market. The design and development of a DVB-T demodulation core constitute an exceedingly complex undertaking requiring the effective collaboration of staff with wide ranging expertise. To meet this challenge, a collaborative design and development team was formed that includes members from two Bell Labs core research departments, the Video and Packet Networking Research Department and the Communications and Statistical Sciences Department with staff in Murray Hill, New Jersey; Fairfax, Virginia; and Largo, Florida; as well as members of Cambridge Industries Group in Shanghai, China. Because of an aggressive time-to-market schedule, the support for DVB-H and other features enhancements was to be deferred until completion of the DVB-T application specific integrated circuit (ASIC).

This paper is aimed at providing insight into the international collaborative effort to design and implement a DVB-T demodulation core with the performance requirements specified in the ETSI 300-744 standard [2] and the NorDig unified requirements [7]. We first give a concise overview of the DVB-T system, followed by a detailed description of the demodulator core, the target design specifications, design methodology, and project management plan. Then we focus on the design and implementation details of the demodulator and describe the underlying signal processing algorithms, the bit-accurate modeling, and register transfer level (RTL) code development. We conclude with a summary that highlights the project milestones and noteworthy technical accomplishments to date and provides perspectives on future extensions to the work.

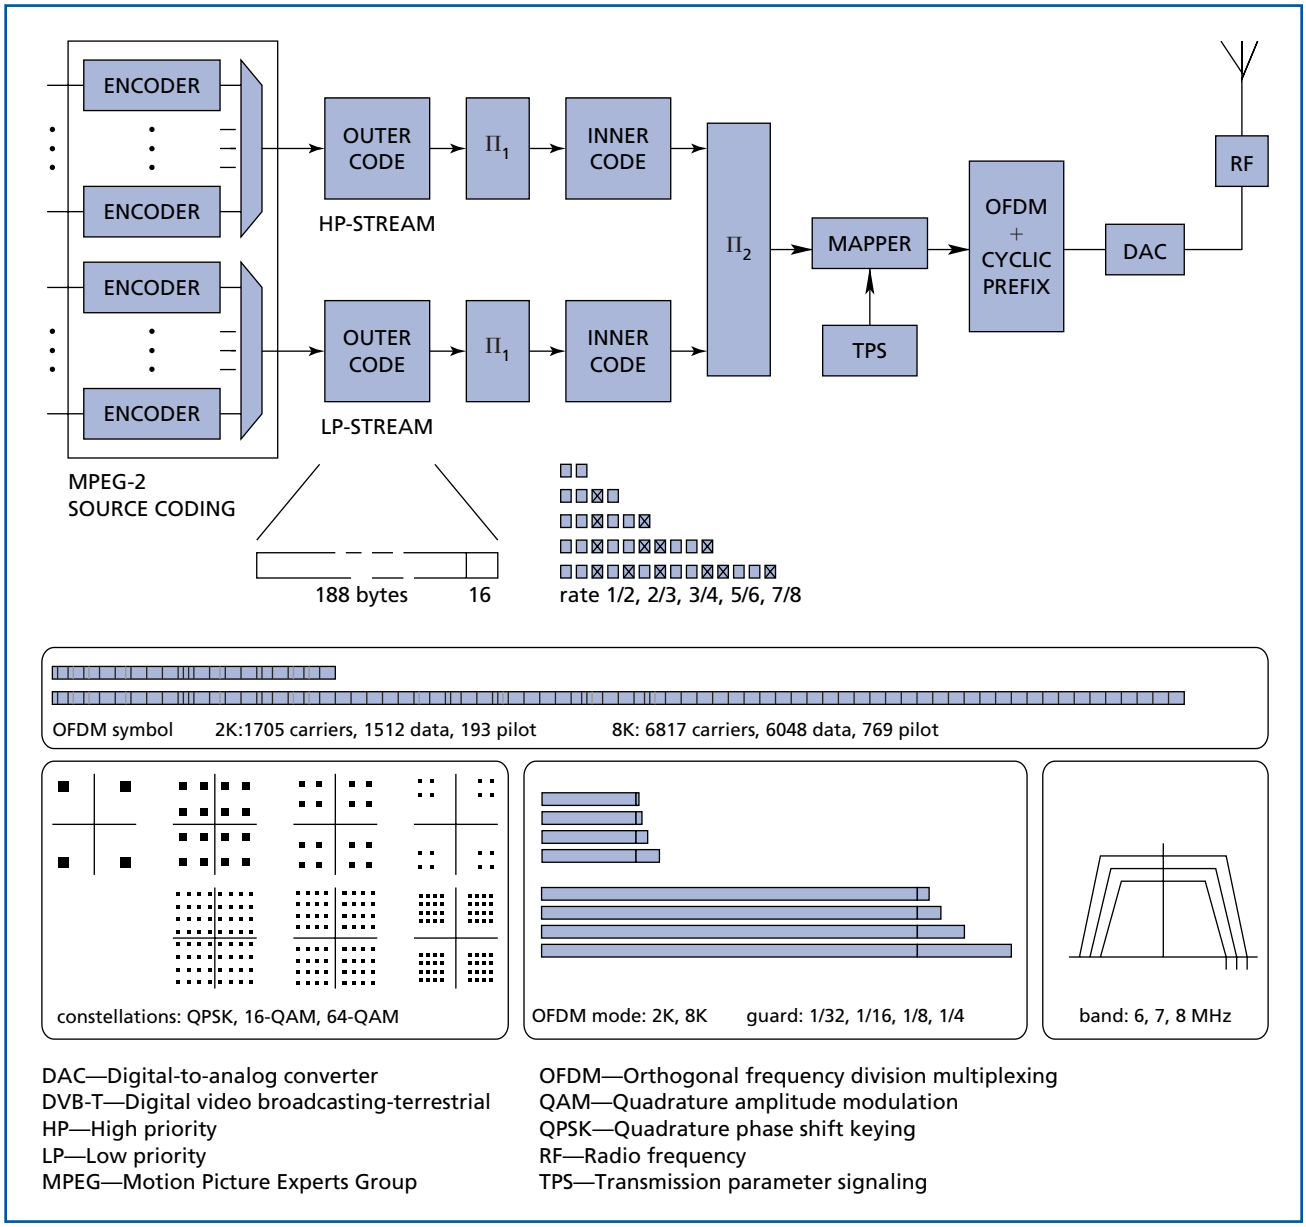
System Overview

The DVB-T system is described in the ETSI 300-744 standard [2], which specifies a baseline transmission system for digital terrestrial television broadcasting that employs multiple framing structures, multirate channel codes, and multiple modulation formats to support digital multiprogram limited definition television (LDTV), standard definition television (SDTV), enhanced definition television (EDTV), and high definition television (HDTV) terrestrial services. The system has been designed to operate within the existing VHF and UHF bands allocated for analog transmission, and to be resilient against high levels of co-channel interference (CCI) and adjacent channel interference (ACI) emanating from existing broadcast services. It supports different network configurations,

from a single DVB-T transmitter system to a large SFN of DVB-T transmitters that operate synchronously to achieve a high spectral efficiency and to suppress interference.

The digitally modulated signal emanating from the DVB-T transmitter is specified in detail to ensure compatibility of DVB-T equipment developed by different vendors. **Figure 1** illustrates the high-level architecture of a typical DVB-T transmitter that supports the transmission of high priority (HP) and low priority (LP) MPEG-2 streams.

The DVB-T standard stipulates a channel spacing of 6 MHz, 7 MHz, and 8 MHz; that is, the nominal bandwidth B_n is 6 MHz, 7 MHz, or 8 MHz. The basic specification is the same for each of the three nominal bandwidths, except for the effective bandwidth,



DAC—Digital-to-analog converter
DVB-T—Digital video broadcasting-terrestrial
HP—High priority
LP—Low priority
MPEG—Motion Picture Experts Group
OFDM—Orthogonal frequency division multiplexing
QAM—Quadrature amplitude modulation
QPSK—Quadrature phase shift keying
RF—Radio frequency
TPS—Transmission parameter signaling

Figure 1.
Schematic block diagram for the architecture of the DVB-T transmitter that supports the transmission of two independent HP and LP MPEG-2 streams.

$B = 7/8\ B_n$, and the corresponding elementary sample period, $T = 1/B$.

The DVB-T standard defines two orthogonal frequency division multiplexing (OFDM) modes of operation. The "2K mode" uses OFDM symbols with $N = 2048$ carriers, of which $K = 1705$ are active. The bandwidth B of the OFDM signal is given by $B = K/(NT)$, because no power is applied to the higher-frequency OFDM bins. It can be easily verified that $B < 7/8\ B_n$. This mode uses $K_U = 1512$ carriers to transport information, and the other 193 active carriers are used as pilots for synchronization, equalization, and conveying of transmission parameter signaling (TPS) information. The 2K mode is suitable for single transmitter operation and for small SFNs with a limited radius. The "8K mode" uses OFDM symbols with $N = 8192$ carriers, of which $K = 6817$ are active. This mode, which uses $K_U = 6048$ carriers to transport information and 769 carriers as pilots, can be employed both for single transmitter operation and for small and large SFNs .

The useful OFDM symbol interval, T_U, is the fundamental period of the harmonically related sine waves that compose the OFDM signal. In order to mitigate inter-channel interference (ICI) caused by multipath delay, the OFDM symbol is cyclically extended in time by an amount known as the *guard interval*. The guard interval is created by a cyclic continuation of the useful part, T_U, and is inserted before the useful part to form a cyclic prefix. The DVB-T system accommodates four guard intervals of duration $T_G \in \{1/4T_U, 1/8\ T_U, 1/16\ T_U, 1/32\ T_U\}$ to enable the system to support different network configurations.

The DVB-T system uses 68 OFDM symbols per OFDM frame, and 4 OFDM frames per superframe. The duration of a superframe equals $T_S = 4 \times 68 \times (T_U + T_G)$: that is, it ranges from 62.8 milliseconds (ms) for an 8 MHz channel where a 2K mode with $T_G = 1/32\ T_U$ is used, to 406 ms for a 6 MHz channel where an 8K mode with $T_G = 1/4\ T_U$ is used.

The DVB-T system supports the selection of the most appropriate level of error correction for a given service or data rate by employing multirate channel coding and modulation, including uniform and multiresolution constellations to be able to transmit two independent MPEG transport streams. A program service can thus be simulcast as a low bit rate, error resilient version or as a higher quality, but more error-prone version. Alternatively, entirely different programs can be transmitted on the separate streams with different ruggedness. To meet bandwidth and efficiency requirements and to provide ruggedness for a wide range of practical applications, the DVB-T system uses QPSK modulation, 16-QAM or 64-QAM signal constellations of rate $R_M \in \{2, 4, 6\}$, respectively. Both QAM constellations can be either uniform or nonuniform. This is to support the simultaneous transmission of two independent MPEG transport streams, referred to as the HP and the LP streams. The parameter $\alpha \in \{1,2,4\}$ specifies the ratio of the minimal distance separating two constellation points carrying different high priority bit values and the minimal distance of the constellation.

The modulation schemes are used in combination with a multirate concatenated FEC. The inner code is a rate 1/2 convolutional code with 64 states. Puncturing this code produces a code of rate $R_C \in \{1/2, 2/3, 3/4, 5/6, 7/8\}$. The outer code is an 8-error correcting [204,188] Reed-Solomon (RS) code. The number of RS codewords in a superframe is given by $1/6 \times K_U\ R_C\ R_M$. The frame structure of the OFDM superframe is chosen such that an integer number of RS codewords fit exactly in an OFDM superframe, thus precluding the need for any stuffing. The number of RS codewords in a superframe ranges from 252 for the 2K mode with a code rate $R_C = 1/2$ and QPSK modulation up to 5292 for the 8K mode with a code rate $R_C = 7/8$ and a 64-QAM constellation. This allows the useful bit rate to be varied from 4.98 Mbps to 31.67 Mbps, depending on the type of service and the channel conditions. It becomes clear from the standard and the preceding description that there are a few hundred distinct system configurations. To convey this information to the receiver, the system uses 17 TPS pilot carriers in 2K mode and 68 TPS pilot carriers in 8K mode to convey the boundaries of the OFDM frames, the frame index, the cell information, and the signaling parameters of the next superframe, which comprise the constellation, hierarchy information, code rate (HP and LP stream), guard interval,

and transmission mode. This information, protected by a Bose-Chaudhuri-Hocquenghem (BCH) code, is represented by a 68-bit sequence, which is differentially coded and transmitted using the TPS pilot carriers.

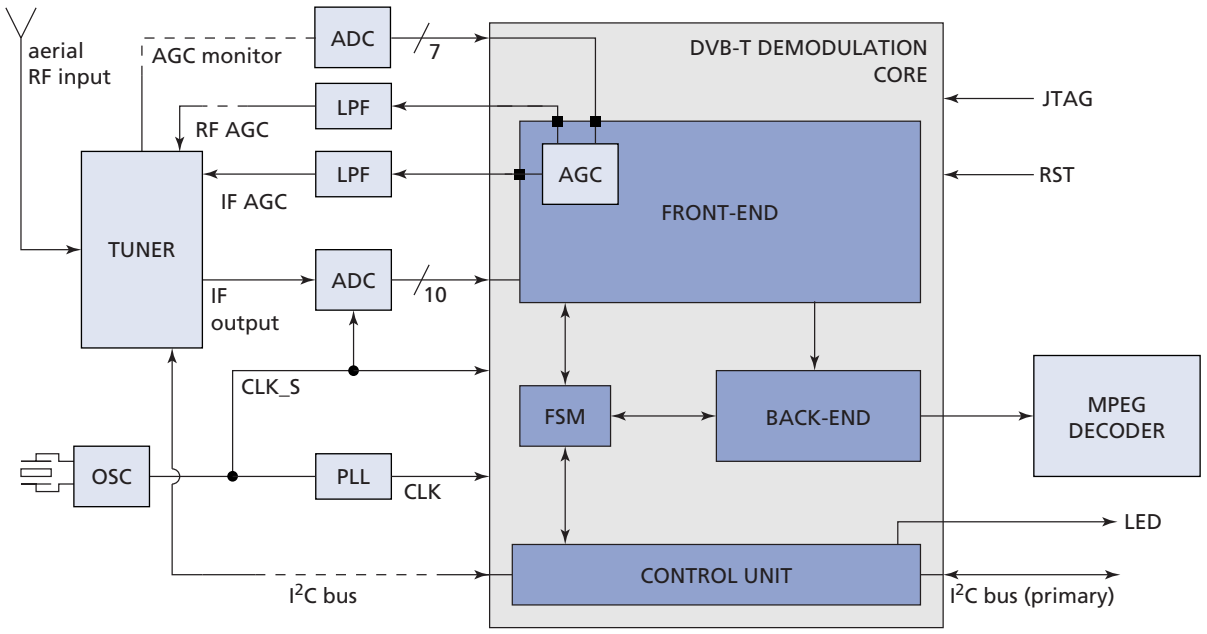
The functionality of the receiver is to a large extent defined by the detailed specification of the signal processing blocks at the transmitter. The processing at the receiver is left open to different implementation solutions as long as its performance complies with the ETSI standard [2] and the NorDig performance specification [7]. In the next sections, we will detail the design and implementation of a DVB-T demodulation core.

Demodulator Core

A typical DVB-T receiver architecture comprises a tuner, an analog-to-digital converter (ADC), a demodulator, and an MPEG decoder. A high-level block diagram of the receiver is shown in **Figure 2.**

The tuner outputs an analog bandpass signal that is characterized by intermediate frequency (IF) f_I and bandwidth B. This signal is digitized using an analog-to-digital converter with sampling rate f_S. The demodulator processes the digitized passband tuner output to reconstruct the transmitted MPEG transport stream in either serial or parallel mode.

The demodulation core consists of two major signal processing blocks, a front-end and a back-end, which are controlled by a finite state machine (FSM), as depicted in Figure 2. The front-end takes the passband tuner output signal at IF frequency f_I, sampled at frequency f_S, as input. Its tasks are to detect and demodulate the OFDM signals in the incoming stream of samples, and to output the high- and low-priority in-phase/quadrature-phase (I/Q) data streams. The



ADC—Analog-to-digital converter
AGC—Automatic gain control
CLK—Clock
DVB-T—Digital video broadcasting-terrestrial
FSM—Finite state machine
IF—Intermediate frequency
I²C—Inter-integrated circuit
JTAG—Joint Test Action Group
LED—Light emitting diode
LPF—Low-pass filter
MPEG—Motion Picture Experts Group
OSC—Oscillator
PLL—Phase-locked loop
RF—Radio frequency
RST—Reset

Figure 2.
System application diagram with the main interfaces to the DVB-T demodulation core.

back-end accepts the I/Q data from the front-end and outputs the MPEG transport stream in either serial or parallel mode.

To ensure interoperability with a wide variety of DVB-T tuners, an automatic gain control (AGC) is used to control the tuner IF gain, and optionally the radio frequency (RF) gain, so that the signal from the tuner is within the dynamic range of the ADC. An RF AGC monitoring input signal is provisioned to monitor the RF AGC if a tuner provides such an output. Typically, most DVB-T tuners operate with IF frequencies at $f_I = 4.57$ MHz or at $f_I = 36.125$ MHz.

Both the front-end and the back-end blocks are controlled by an FSM, which, in turn, can be controlled and monitored through the on-chip control unit. This control unit is accessible through an inter-integrated circuit (I^2C) bus [8].

The complexity of a DVB-T demodulator is such that the most cost effective implementation, suitable for use in consumer electronics, is an ASIC. Moreover, the demodulation core is designed to be used with a low-cost ADC and an oscillator with a tolerance up to 100 ppm. The clock internal to the ASIC is to be derived from this oscillator.

The design is to comply with the ETSI standard [2] and the performance with the NorDig performance specification [7]. The latter stipulates the required carrier-to-noise ratio *C/N* to achieve quasi error-free (QEF) reception, where *quasi error-free* is defined as a stationary situation where less than one uncorrected error event per hour occurs.

An objective measure to attain QEF reception quality is the bit error rate (BER) at the input of the MPEG-2 demultiplexer, which needs to be of the order of 1×10^{-11}. For channels where the channel characteristics are Gaussian-like, it suffices to determine the BER at the input of the inner (Viterbi) decoder. A BER of 2×10^{-4} needs to be attained at the output of the inner (Viterbi) decoder to obtain a BER of 1×10^{-11} at the input of the MPEG-2 demultiplexer. A different metric is required in situations where impulsive noise is significant or where the channel conditions deteriorate significantly for a short period, for example, in the presence of interfering phase alternating line (PAL) signals. In this case, the NorDig unified requirements [7] prescribe that the quality measure is to be based on a BER of 1×10^{-11} at the input of the MPEG-2 demultiplexer, or, alternatively, on a subjective measurement method, where a specified video test sequence is observed for a certain period to determine whether it corresponds to a good reception quality.

An alternative measure is the *picture failure point*, which is defined as the minimum carrier-to-noise, *C/N*, or carrier-to-interference, C/N_i, value such that on average more than one transport stream (TS) packet error occurs in 10 seconds. This is a much more convenient measure for some of the measurements than the QEF performance criterion, because it requires much fewer simulations to determine the picture failure point. Four scenarios are to be considered:

- The quasi-stationary channel.
- Reception in the presence of impulsive noise.
- Reception in the presence of co-channel interference.
- The single-frequency network.

In summary, the design target performance indices may be qualitatively defined as follows:

- Threshold *C/N* for QEF reception for quasi-stationary channels.
- Threshold impulse noise level for QEF reception.
- Threshold C/N_i for QEF reception for co-channel interference.
- Threshold carrier-to-echo ratio as a function of delay for QEF reception in an SFN.

In addition to meeting the NorDig performance specifications, the DVB-T demodulation core has several mandatory and optional high-level interfaces and inputs, all of which are depicted in Figure 2.

The front-end and back-end of the DVB-T demodulation core each comprise key signal processing modules that form the heart of the demodulator. These modules, and the aforementioned interfaces, are described in more detail in the next sections. The remainder of this section discusses the procedures that are followed to design and implement a field programmable gate array (FPGA) prototype and a demodulator ASIC.

Design-Level Methodology

The implementation of a DVB-T demodulator ASIC is characterized by the following steps common to ASIC designs of this complexity:

- *Step 1*. System design, algorithm selection, and high-fidelity simulation (e.g., in Matlab*, C).
- *Step 2*. Development of a bit accurate model (BAM) for system design, simulation, and verification. The BAM is typically fixed-point C code, which is used to simulate and verify the system design realistically.
- *Step 3*. Implementation and testing of the design in an FPGA to enable rapid prototyping, real-time simulation, and field tests with the ability to modify the code prior to finalizing the design and manufacturing an ASIC. A hardware description language such as Verilog or VHDL is used to implement the system design in RTL. Simulations in RTL are used to refine the logic of the design and to ensure performance is being met by comparing the results with the BAM simulations.
- *Step 4*. Concept verification using a hardware setup in which an FPGA is programmed using the developed RTL, in order to test the designed demodulator with real signals. At this stage, the objective is usually to tune the parameters of the design to optimize performance rather than to make any major modifications to the design, although this will be done if necessary.
- *Step 5*. Development of the ASIC RTL using ASIC synthesis tools. In this step, the ASIC RTL design is developed from the FPGA RTL design, and a gate level netlist, specific to the manufacturer of the ASIC, is synthesized. The gate-level netlist is the resulting collection of standard cells, plus the necessary electrical connections between them.
- *Step 6*. Placement of the standard cells into regions representing the final ASIC.
- *Step 7*. Routing. A routing tool takes the physical placement of the standard cells and uses the netlist to create the electrical connections between them. The output is a set of photo masks enabling a semiconductor manufacturer to produce the ASIC chips.

It is important to note that once the first four steps are completed, the remaining steps to produce the DVB-T demodulator ASIC are straightforward.

Project Planning and Management

The collaborative design and development team spanned two continents and included staff from two Bell Labs departments in the United States and staff from Cambridge Industries Group in China with locations and areas of responsibility as follows:

- The Alcatel-Lucent Video and Packet Networking Research Department, based in New Jersey, Virginia, and Florida, was responsible for project management, target design specifications, development of the BAM, and simulation of the system and the support of customer laboratory and field testing.
- The Alcatel-Lucent Communications and Statistical Sciences Department, based in New Jersey, was responsible for target design specifications, system architecture and design, algorithm design, and simulation and support of customer laboratory and field testing.
- Cambridge Industries Group, based in Shanghai, was responsible for FPGA RTL architecture, design, and simulation.

The customer, although not an integral part of the RTL design and development team, had responsibility for specifying the demodulator performance and interface requirements, providing sample data for use in simulations, developing the FPGA-based test and evaluation hardware, and conducting laboratory and field tests with real signals.

To facilitate project execution and communications, the design and development team held weekly teleconferences to review the status of the design and development work and to discuss issues requiring resolution. A teleconference summary and an evolving list of action items with dates for resolution were distributed to the design and development team each week. A weekly written status report of the project was e-mailed to the customer, who indicated a preference to communicate by e-mail rather than by phone because of language barrier problems.

At the beginning of the project, the team developed an aggressive composite project schedule that encompassed all phases of the project through development of the RTL. Because of the complexity, it was decided that the demodulator core should be designed in three, manageable stages, as follows:

- *Stage 1*. Back-end, including inner/outer deinterleavers, Viterbi decoder, frame alignment, Reed-Solomon decoder, MPEG transport interface.

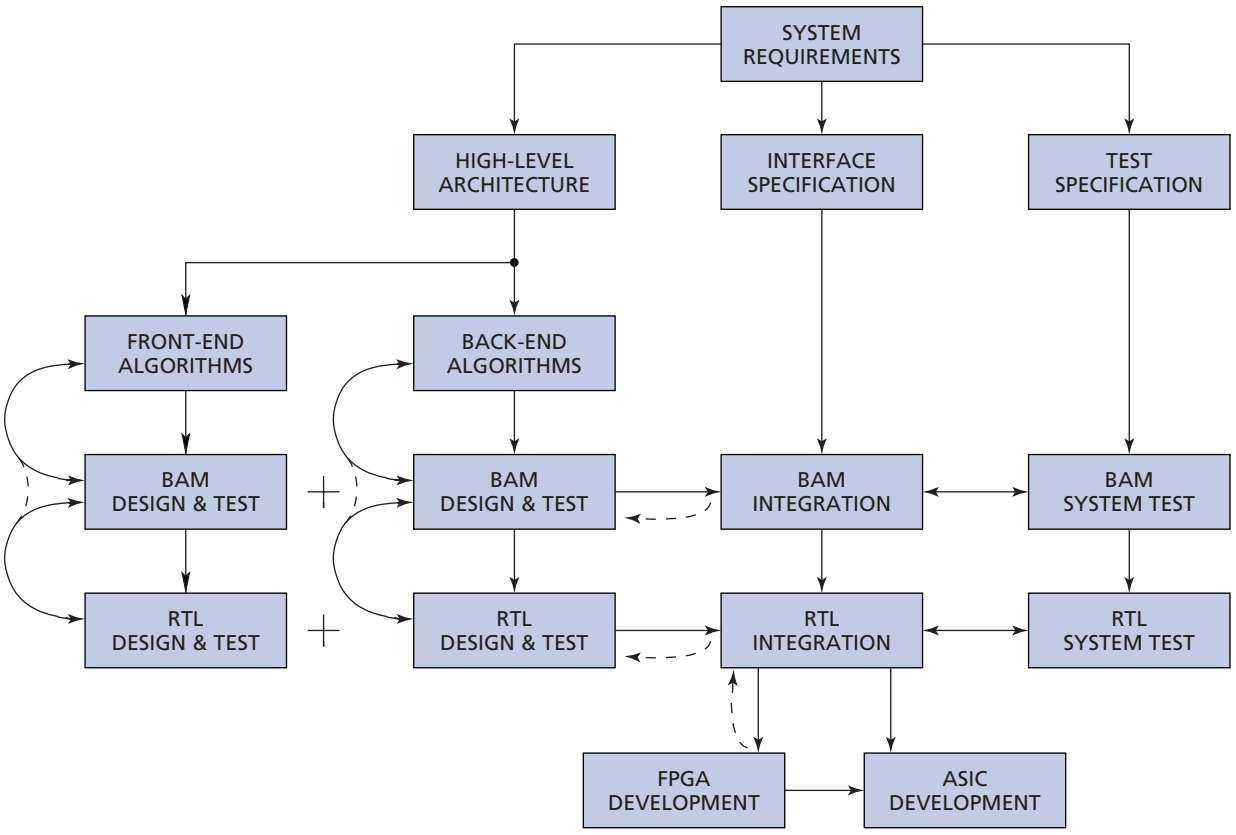


Figure 3.
Design and development flow for the DVB-T demodulator, from system requirements to FPGA and ASIC implementation.

- *Stage 2*. Front-end, including baseband converter, fast Fourier transform (FFT), channel estimator, cyclic prefix removal, demapper.
- *Stage 3*. Front-end, including pre/post FFT synchronization, AGC, equalizer, impulse noise suppressor and TPS decoder, and integration of the front-end and back-end.

Although the design methodology described in the preceding section, and depicted in **Figure 3**, was applied to each of the aforementioned design stages, there was some degree of concurrency between the stages. For example, as the algorithm staff completed the design and simulation of stage *k*, it would commence the design of the modules of the next stage while the bit accurate model (BAM) and RTL design teams were completing their tasks related to stage *k*. The ability of the design and development team to meet the aggressive project schedule was in large part attributable to their expertise and versatility; their ability to handle multiple tasks, sometimes concurrently; and most importantly their ability to work as a team.

Front-End Design

The front-end of the DVB-T demodulator takes the analog passband tuner output signal at IF frequency f_I, sampled at frequency f_S, as input and processes this stream to obtain I/Q (channel encoded) data stream(s) as shown in **Figure 4**. The sampled signal is down-converted and filtered to form a

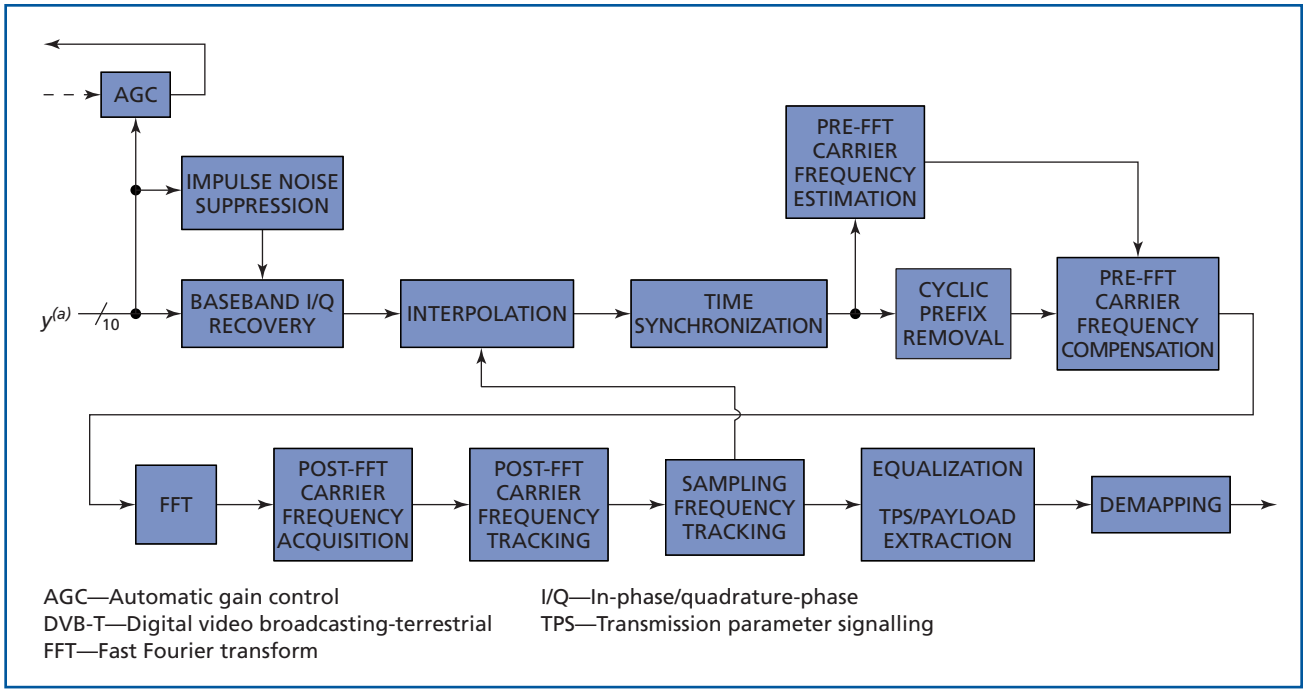


Figure 4.
Schematic block diagram of the front-end architecture of the DVB-T demodulation core.

corresponding stream of baseband I/Q samples. The nominal bandwidth B_n of the oversampled signal $y^{(a)}$ is known a priori, but both the transmission mode and the guard interval must be determined during acquisition. Upon synchronization, an interpolator is used to resample $y^{(a)}$ in order to be able to perform an FFT. Post-FFT blocks use the pilots embedded in the OFDM symbols to extract the frequency error and to equalize the FFT output. The modulation format and the code rate are extracted from the TPS information provided by the TPS pilot carriers embedded in the 68 consecutive OFDM symbols of an OFDM frame. The payload is extracted and demapped, after which it is passed to the back-end for further processing. A block diagram of the resulting DVB-T receiver front-end is depicted in Figure 4.

An important design parameter is the choice of the sampling rate f_S and the resolution of the ADC. The signal input to the ADC is a bandpass analog signal that is characterized by bandwidth B and IF frequency f_I. The signal-output of the ADC is a real-valued digital sequence that corresponds to the sampling rate f_S. The effect of sampling at rate f_S is to replicate the spectrum at intervals of $l \times f_S$ for all integer values of l. Aliasing occurs when a displaced left sideband interferes with the right sideband, or vice versa. We determine ranges of sampling rates that prevent aliasing and facilitate the design of a digital filter that separates the right sideband and the displaced versions of the left sideband. The objective is to maximize the minimal frequency separation between the right sideband and the displaced versions of the left sideband. For $f_I = 36.125$ MHz frequency, it follows that a sampling rate of $f_S = 28.9$ MHz emerges as a viable choice, since the next lower candidate, $f_S = 20.64$ MHz, would possibly require more complicated fractional-sample delay filters. The same sampling rate of 28.9 MHz is completely satisfactory for the 4.57 MHz IF frequency as well. The resolution for the ADC should be more than 6 bits in order to handle 64-QAM constellations. In order to minimize the cost of the design, a 10-bit ADC was selected. In the next sections we will revisit the impact of this choice on the performance.

Front-End Modules

The development of the front-end modules basically follows the signal flow, except for the FFT and the TPS decoder, which have a well-defined interface.

Automatic gain control. This module observes the ADC output over time by using an adjustable low-pass filter (LPF) that measures the average amplitude of the input. This is then compared with a predefined set point. The difference is determined periodically and used to produce feedback control information to the tuner in the form of RF AGC and IF AGC control signals that are generated by using a digital LPF followed by a sigma-delta converter.

Impulsive noise suppression. Impulsive noise, typically caused by appliances such as switches and electric motors, can be characterized as random occurrences of pulses having random amplitude and spectral content. It was shown in [6] that impulsive noise can be adequately modeled as a burst of n_p pulses of duration T_p with a specified minimal and maximal spacing, $\tau_{\min}$ and $\tau_{\max}$ between adjacent pulses in a burst. The duration of a pulse, T_p, is typically about 250 nanoseconds (ns). This model is known as the double-gated Gaussian noise model [6]. It implies that the burst length, T_b, ranges from $n_pT_p + (n_p - 1)\tau_{\min}$ to $n_pT_p + (n_p - 1)\tau_{\max}$. The effective burst length, T_e, is equal to $n_p\, T_p$. The spacing between bursts ranges from about 0.14 ms to about 15 ms, depending on the source of the impulse noise. According to [6], impulsive noise typically falls into one of six noise impulse profiles representative of traffic, heating systems, fluorescent lights, and appliances. An impulsive noise suppression module is used to scan the ADC output and to detect the six types of impulsive noise by using an FSM combined with threshold mechanisms. Once an impulse is detected, the affected samples are attenuated to reduce the effective impulse noise energy. The attenuation levels are adjustable.

Baseband conversion. Baseband conversion is the first step of the demodulation process. Its task is to convert the digitized IF signal $y^{(a)}$ into a complex I/Q signal $y^{(b)}$. The baseband I/Q recovery is implemented by multiplying the samples by $\exp(-i2\pi f_{\mathrm{NCO}})$, generated by a numerically controlled oscillator (NCO), and then passing the result through a low-pass filter. The NCO uses a reference frequency that is preset to 4.570 MHz and 7.225 MHz to support IF frequencies of 4.57 MHz and 36.125 MHz, respectively. The baseband converter uses two 2048×11 bit read-only memory (ROM) to store partial sine and cosine waveforms in conjunction with control logic that exploits symmetries to reproduce the entire waveforms.

Interpolation. The interpolator resamples the incoming stream at $f_S = 28.9$ MHz to the OFDM sampling rate $1/T$. A compact design is obtained with a 5-tap Farrow filter [3], which computes the output sample at offset τ on the basis of the five nearest input samples using a 5×5 matrix with 19 nonzero entries. This would imply the use of up to 19 multiplications per sample output. The need for multipliers can be eliminated by using a $2^7 \times 80$ bit ROM. A 7-bit quantization of the fractional difference δ is considered to be sufficient.

Time synchronization. The transmission mode and guard interval are typically not known to the receiver at initial acquisition. The tasks of the time synchronization module are to detect the transmission mode (2K or 8K) and the guard interval automatically, and to identify the boundaries of the OFDM symbol. This module is implemented by using a sliding window correlator and an algorithm that stacks a configurable number of correlation vectors to detect the size and location of the cyclic prefix. The module adjusts the sliding window so as to check all possible combinations of modes and guard intervals in a predefined order. The stacking algorithm amplifies the correlation peak and decreases the noise by averaging. Once the peak becomes sufficiently distinct, symbol synchronization is achieved, the transmission mode and guard interval are deemed to be detected, and the resulting information is passed to the guard removal module for further processing. Otherwise the search for the correct mode is continued by switching to the next candidate mode and guard interval. By adjusting the number of stacked correlation vectors, this procedure proves to be very robust under adverse channel conditions such as inter-symbol interference (ISI) and multipath distortion. In a multipath environment, the receiver detects a signal that is the weighted sum of several delayed copies of each OFDM symbol, thus

widening the correlation peak. In this case, the developed algorithm measures the width of the correlation peak in order to locate a position for which the corresponding sample is the weighted sum of sample points of the cyclic prefix.

Pre-FFT carrier frequency estimation. The carrier frequency offset can be expressed as a multiple of $1/T_U$. The pre-FFT carrier frequency estimation module determines the fractional part of this multiplication factor. The module is based on the correlation of the cyclic prefix of an OFDM symbol with the data part of the OFDM symbol. The size of the cyclic prefix is passed to the module from the time synchronization module. The fractional offset is estimated separately for each OFDM symbol and is passed through a smoothing filter, which averages the fractional offset estimates of several OFDM symbols. This makes the resulting estimate very accurate and robust.

Cyclic prefix removal. This module uses information from the time synchronization module to remove the cyclic prefix. It recovers the useful part of the OFDM symbol, being either a 2048-point or a 8192-point complex vector.

Pre-FFT carrier frequency compensation. This module compensates for the fractional carrier frequency offset by phase-shifting the samples of the OFDM symbol by an amount that is determined by accumulated and instantaneous offset estimates provided by the pre-FFT carrier frequency estimation module.

FFT module. The FFT is a critical component in an OFDM receiver because of its size, complexity, power consumption, and timing requirements, dictated by the minimal OFDM symbol duration in the 2K and the 8K mode. It follows that the FFT module has to compute a 2048-point complex FFT within 224 microseconds (μs) and an 8192-point complex FFT within 896 μs. The objective is to design a single FFT module with a desirable low power consumption and the smallest silicon area that can operate in the 2K mode as well as the 8K mode. This is accomplished with a mixed radix-2/radix-4 architecture. In the 8K mode, the FFT is performed by six serially concatenated segments that perform radix-4 butterfly operations, followed by one segment that performs radix-2 butterfly operations. The seven segments are surrounded by two buffers, implemented as random access memory (RAM). A block-floating scaling algorithm is used to achieve a higher precision. The 2K mode is supported by bypassing the last segment that performs radix-4 butterfly operations.

Post-FFT carrier frequency acquisition. This module compensates the integral carrier frequency offset. The estimation is based on computing the correlation vectors between several shifts of the received OFDM symbol and the predefined values of its continual pilots. The module searches for the most reliable shift in the range from −15 to +15 samples. The model outputs the frequency offset, once it has been determined, to shift the FFT output in order to align the OFDM symbol. The operation of the post-FFT carrier frequency acquisition module is controlled by the FSM. Generally, this module performs the search only at acquisition, and at reacquisition if needed. Once the offset has been determined, it switches to a flywheel mode where it monitors the offset.

Post-FFT carrier frequency tracking. This module detects and removes any residual fractional frequency offset that is left at this point. The algorithm uses the observation that the residual offset is typically almost constant for all samples of an OFDM symbol. The offset estimate is then obtained by averaging the estimates of the offset for all continual pilots.

Post-FFT sampling frequency tracking. This module estimates the sampling frequency offset, based on the observation that the magnitude of the offset grows approximately linearly with the index of a data sample in an OFDM symbol. The continual pilots in the negative and positive frequency portions of an OFDM symbol are used to estimate the linear growth of the offset. The rate of the linear growth is used as an estimate of the offset, which is fed back to the interpolator. This offset increases very slowly but eventually leads to significant errors, and therefore the interpolator occasionally adjusts the sampling rate.

TPS decoder. The TPS decoder observes the TPS pilot carriers in each OFDM symbol and uses differential detection to retrieve the coded transmission parameters. The TPS decoder consists of three parts. The first part extracts and combines the

TPS pilots to obtain a soft decision value for the TPS symbol that is contained in an OFDM symbol. Each OFDM symbol contains a single, differentially encoded TPS symbol. In acquisition mode, the objective is to find the starting position of a 68-symbol OFDM frame. A frame synchronization module that comprises a correlator is used to locate the sync word that is embedded in the 68-bit TPS sequence. Once synchronization is obtained, the boundaries of the OFDM frames and their relative position in a superframe are known. A third module is used to decode the TPS information, which is protected by a double-error correcting shortened BCH code. A hardware-efficient direct solution BCH decoder has been developed to decode the TPS information. After decoding, the sequence is parsed to extract the frame index and the cell information, as well as the transmission signaling parameters for the next superframe.

Equalization. This module estimates the channel frequency response using the embedded pilots and corrects the data of each carrier on the basis of the estimated frequency response of the channel. A two-dimensional linear equalization algorithm was developed to interpolate and equalize effectively over time and frequency. The equalization in the frequency domain after transmission over a multipath channel is considered first. As noted previously, the time synchronization module locates a sample for which the component sample of each path is a part of the cyclic prefix. This sample is used as the starting sample of the 2K or 8K sample vector that is fed into the FFT. The obtained vector is thus a sum of shifts of the received OFDM symbol. After performing the FFT, this results in the multiplication of each sample of an OFDM symbol by a factor, which for m paths can be represented as a sum of m exponentials, the magnitudes of which are proportional to the sample index and the magnitudes of the cyclic shifts of the received OFDM symbol. The equalization module is used to correct these multiplicative errors. The design of the equalization module is a nontrivial problem since the scattered and continual pilots do not cover all OFDM samples. In particular, a significant number of samples are not covered by any pilot. Thus, the estimates of the multiplicative errors for those samples are to be computed on the basis of adjacent samples that are covered by either scattered or continual pilots. This problem is solved with the help of asymmetric low-pass filter interpolation. A similar approach is followed for the equalization in the time domain.

Demapping module. The demapping module takes the constellation points at the output of the equalizer and attempts to extract soft decision data samples from the Gray-mapped constellation. The demapping algorithm depends on the constellation (QPSK, 16-QAM, or 64-QAM) and the hierarchy (α = 1, 2, or 4).

Front-End Validation

The validation of the front-end system design is performed at a component level. The basic components, such as the baseband converter, filters, and the interpolator, are validated using a randomized input. The major signal processing blocks, in particular, the synchronization algorithms and the equalizer, require a full-blown simulation with a functional DVB-T transmitter and a transmission channel. This requires a test bench that comprises a random signal generator, a transmitter, a channel, and a chain of receiver modules up to at least the module under test. For this purpose, modules were implemented in Matlab and C that together provide the functional behavior of the transmitter and its underlying components. Various channel modules were implemented, including the channels specified in the NorDig performance specification [7]. System validation is discussed as a part of system integration.

Back-End Design

The back-end of the DVB-T demodulator decodes the received I/Q data samples and aims to reconstruct the MPEG-2 transport system (TS). The back-end architecture of the DVB-T receiver consists of modules that are the counterparts of the channel coding, interleaving, and modulation and signal mapping elements of the DVB-T transmitter, as specified in [2]. The resulting block diagram of the DVB-T receiver back-end is depicted in **Figure 5**.

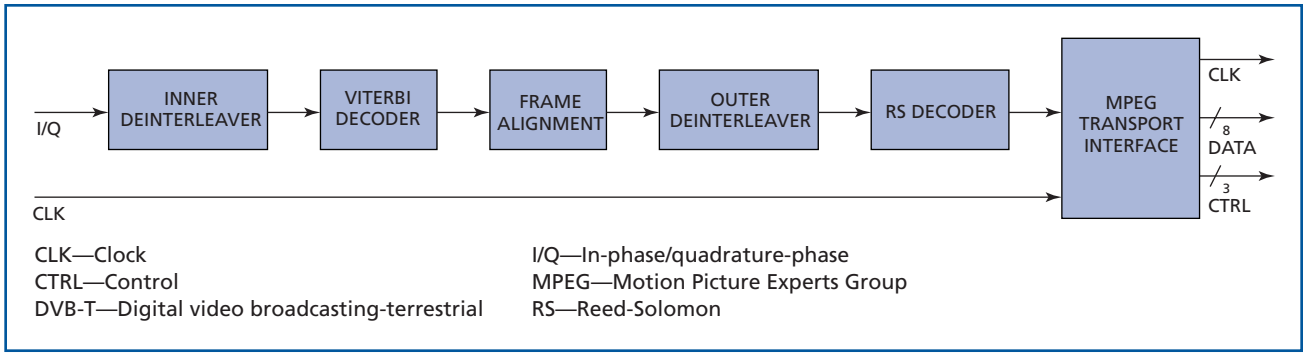


Figure 5.
Schematic block diagram of the back-end architecture of the DVB-T demodulation core.

The first two blocks, the symbol deinterleaver and the bit deinterleaver, together invert the operation of the inner interleaver at the transmitter side. The next two blocks constitute the two decoding stages of the convolutional inner code. The transmitter uses a rate 1/2 convolutional code with 64 states. This so-called mother code is punctured to obtain codes with rates of 2/3, 3/4, 5/6, or 7/8. The depuncture module reconstructs the nonpunctured rate 1/2 stream by inserting 0s at the punctured positions. The resulting stream is processed by the Viterbi decoder. The frame alignment module checks the frame alignment at the output of the Viterbi decoder. The outer deinterleaver then inverts the outer interleaving operation and produces frames of length 204 bytes, which are in turn decoded by the RS decoder. Upon data extraction and correction, the 188-byte output is passed through the descrambler and forwarded to the MPEG-2 transport stream generator.

Back-End Modules

The development of the back-end modules was done in parallel, because all modules have a well-defined interface.

Symbol deinterleaver. The symbol interleaver takes the output I/Q data from the front-end modules and requires two control signals to determine the start of a superframe and the transmission mode (2K or 8K). The symbol interleaver consists of a 6048 × 18 bit RAM to store the (soft) input symbols and a control unit that comprises a 12-bit shift register and some combinatorial logic to generate the indices for the even and odd symbols [2]. Since the even and odd symbols are using different address generation schemes, it is possible to implement the deinterleaver using a single-port memory.

Bitwise deinterleaver. The bitwise deinterleaver uses shift registers to invert the operation of the bit-interleaver present in the transmitter. The symbol sequence emanating from the symbol deinterleaver is split into six substreams, which are permuted and delayed using a serial concatenation of two shift registers with a total delay of 105. Some or all substreams are multiplexed according to the modulation scheme and to uniform or hierarchical mapping. The processing is based on bit vectors instead of bits to support soft decision Viterbi decoding.

Depuncture module. The depuncture module reverses the puncturing performed at the transmitter to expand the received codewords of rate 2/3, 3/4, 5/6, or 7/8 by inserting "0s" at the punctured positions and by marking these positions by erasure indicators. Next, the serial input is demultiplexed to obtain the two input streams $x^{(v)}$ and $y^{(v)}$ to the Viterbi decoder. The module supports an automatic probing mode that uses feedback from the Viterbi decoder to establish synchronization of $x^{(v)}$ and $y^{(v)}$.

Viterbi decoder. The Viterbi decoder is designed to decode the rate 1/2 convolutional "mother code" at a line speed of about 36 Mbps (useful data rate), and to support both hard-decision and soft-decision Viterbi decoding. The Viterbi decoder was implemented using

a sliding window design. The trace-back module finds the optimal path through the Viterbi trellis by tracing through a fixed number of add-compare-select results. The length of the trace-back buffer is typically 10 times the constraint length for punctured codes. For this application, a trace-back buffer of length 128 was used. A best-state option is available to select the starting location for the trace-back from the state with minimal cost in all 64 states. A special trace-back survivor memory management algorithm [4] was selected to cope effectively with the high data rate, and three single-port RAMs of depth 128 are used to hold the 64-bit state message for rotation. It was found that 3-bit soft decoding formed the right trade-off between performance and complexity, because its error correction performance is within 0.2 dB of floating-point performance.

Outer deinterleaver. The outer deinterleaver is a convolutional deinterleaver consisting of 12 branches, numbered from 0 to 11, with the property that branch J has a delay of $11 \times J$ bytes. This delay is implemented using a first-in, first-out (FIFO) buffer. The input bytes are cyclically fed to the 12 branches and the output is formed by cyclically retrieving single bytes from the 12 branches. It should be noted that branch 0 has zero delay. This deinterleaver is implemented by stacking the 11 FIFO buffers to obtain a single 1122×8 bit RAM. A control unit is used to generate the memory addresses to transmit write operations to the 12 branches in consecutive memory segments of the RAM, and to retrieve the output. A single-port memory thus suffices to realize the required functionality.

Frame alignment. The frame alignment module recovers the MPEG TS packet alignment by searching for the sync bytes that mark the start of each length-204 packet. The first byte of the first packet of an 8-packet segment is equal to SYNC byte 01000111, while the other seven packets in the segment start with the inverted SYNC byte 10111000. To simplify detection, we search for the differential sequence 1100100. An FSM with three states, "search," "presync," and "lock-sync," is used to control this module. Feedback from the RS decoder is used to determine whether a loss of synchronization has occurred.

Reed-Solomon decoder. The RS decoder is used to correct the (spurious) errors that the Viterbi decoder could not correct. These errors are typically burstlike. An RS code, usually referred to as an $[n,k]$ RS code over Galois field $GF(2^m)$, is a code of n m-bit symbols, out of which k are systematic information symbols [1]. This code can correct up to $t = (n - k)/2$ symbol errors and is therefore particularly effective in correcting burst errors when used in combination with symbol interleavers. If a block is received with more than t errors, the decoder will fail. The RS decoder generally detects that more than t errors have occurred and indicates a failure, but there is a small fraction of error patterns that form a codeword with at most t errors. In this case the decoder adds the "corrected" error pattern to the codeword and as such introduces up to t additional errors.

The 8-error correcting [204,188] RS code, used for the DVB-T system, reduces a BER of 2×10^{-4} at the output of the Viterbi decoder to a BER of 1×10^{-11}. The RS decoder consists of consecutive modules that perform syndrome computation, the key equation solver, error value computation, and error correction. In this implementation, Horner's rule [1] is used to compute the syndrome values efficiently, requiring 16 8-bit registers and some combinatorial logic. The next module, the key equation solver, takes the 16 syndrome values as input and determines the coefficients of the error locator polynomial and the error evaluator polynomial. The key equation solver is implemented using the Massey-Berlekamp algorithm [1], which provides a hardware-efficient implementation based on shift registers. The next module, the error value computation module, evaluates the error locator polynomial and the error evaluator polynomial at all positions using Chien's search and Forney's equation [1] in order to determine the error locations and the error values. This is realized by using an architecture that is similar to syndrome computation. If there are more than 8 errors, there are likely to be inconsistencies between the degree of the error locator polynomial and the number and position of located errors. This can be detected by comparing the number of error positions with the degree of the error locator polynomial and by using

extra RAM to buffer the received vector. The errors are then only corrected in the buffered word if there are no inconsistencies.

Descrambler. The function of the descrambler is to derandomize the decoded data by performing a bitwise modulo-2 addition of the data sequence with a pseudo-random binary sequence (PRBS) sequence. The PRBS sequence, defined by polynomial $x^{15} + x^{14} + 1$, is generated on the fly using a 15-bit shift register. The descrambler is initialized with the register contents 100101010000000 and is reinitialized at the start of every eight transport packets.

MPEG transport interface. The function of the MPEG transport interface is to prepare the MPEG-2 transport multiplex packets for transmission in parallel or serial mode. This module uses a FIFO buffer to accommodate the different time domains.

Back-End Validation

To validate the back-end effectively, in terms of both functionality and performance, we implemented the corresponding components in the transmitter. Components that take a discrete input and produce a discrete output, that is, the deinterleavers, the depuncture module, and the MPEG transport interface, are the easiest to validate. The Viterbi decoder and the RS decoder were tested by simulating error-free situations as well as for various additive white Gaussian noise (AWGN) channels, and the results were compared to known performance results. The frame alignment module was tested during the back-end system test. Test vectors generated in the BAM were used to test the RTL design, and end-to-end tests of the BAM and RTL code confirmed the correct operation of the back-end.

System Integration

The final step in the design and development of the DVB-T demodulator is integration of the individually validated front-end and back-end blocks with each other and with two additional modules, the FSM and the register map. The tasks of the FSM are to determine the state of the demodulator and control the process from initial acquisition to steady-state operation. The register map is used for both monitoring and control. It is accessed through the I^2C bus. These two modules interact with both the front-end and the back-end and therefore are dealt with separately in the system integration phase.

Synthesis Results

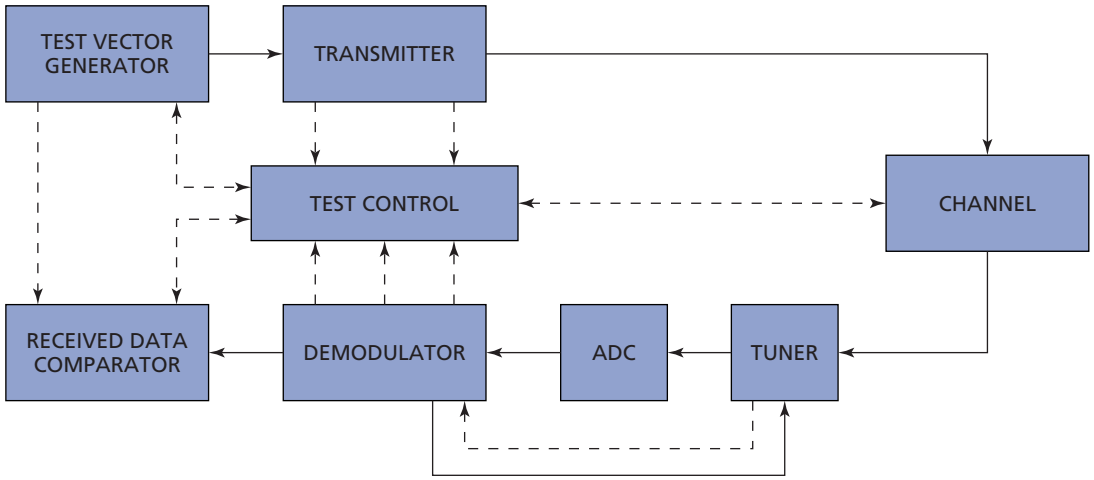
The back-end RTL implementation in a Xilinx* FPGA requires 2752 registers, 6155 look-up tables (LUT) and 160 kilobits of memory. The FFT implementation that is part of the front-end requires 1885 registers, 2822 LUT, and 416 kilobits of memory.

System Validation

System validation, in terms of both function and performance, is essential in the development process. As discussed in the two previous sections, component tests and partial systems tests were an integrated part of the design process. The high-level interactions between the modules and the overall system performance need to be determined upon system integration. This is accomplished by reusing the functional modules that simulate the transmitter and by enhancing the channel models and simulation test benches. To validate the DVB-T demodulator design a BAM testbench of the end-to-end DVB-T system was constructed. The system-level test flow for the DVB-T demodulator is illustrated in **Figure 6**.

The simulation parameters provided in the ETSI 300-744 standard [2] are given for 47 different channel code rates and modulation types, both non-hierarchical and hierarchical. The required *C/N* values are given for a Gaussian channel, a Ricean, and a Rayleigh fading channel as specified in [2]. The simulation results, provided for 8 MHz channels only, assume that channel estimation is perfect and that there is no phase noise. The results are of an informative nature and, according to the standard, subject to confirmation by testing. Consequently, these results are to be treated as indicative of the performance instead of stringent performance requirements. The NorDig unified requirements [7], on the other hand, specify the channel conditions in more detail. In particular, three profiles are specified, under which a DVB-T receiver is to be tested:

- *Profile 1*. Gaussian noise *N* is applied together with the carrier *C* in a signal bandwidth of 7.61 MHz (for 8 MHz channels) or 6.66 MHz (for 7 MHz channels). No echo is applied.



Solid lines represent signal flows, dashed lines represent optional control and/or monitoring signals.

ADC—Analog-to-digital converter
DVB-T—Digital video broadcasting-terrestrial

Figure 6.
System-level test flow for the DVB-T demodulator.

- *Profile 2*. The carrier *C* includes the direct path signal and an echo. The echo has the same power (0 dB echo) as the direct path signal, is delayed from 1.95 μs to 0.95 times the guard interval length, and has 0 degree phase at the channel center.
- *Profile 3*. Carrier *C* includes the direct path signal and an echo, Doppler shifted by $\pm$ 10 Hz. The echo has the same power (0 dB echo) as the direct path signal, is delayed from 1.95 μs to 0.95 times the guard interval length, and has 0 degree phase at the channel center.

The first two profiles apply to stationary channels and the third profile applies to time-varying channels, typical of fixed rooftop reception (e.g., mast sway, antenna sway) and in-house portable reception (e.g., people walking around the receiving antenna). It should be noted that a Doppler shift of $\pm$ 10 Hz is applied to the direct path as well as the echo rather than to the echo only. It follows that the frequency separation of these two paths can be up to 20 Hz.

The threshold *C/N* for QEF reception under NorDig profiles 1 and 2 is given in [7]. It can be seen that for Gaussian channels the performance requirements are at least 2 dB higher than specified in [2]. This indicates a 2-dB penalty due to imperfect channel estimation and phase noise.

End-to-end simulations are used to demonstrate the correct functional behavior of the implementation and to give an initial indication of the performance. As an example, we present the results of an end-to-end simulation of a 2K, QPSK-modulated DVB-T system with a guard interval $T_G = 1/4\ T_U$ for NorDig Profiles 1 and 2. **Figure 7** illustrates the performance of the synchronization module. Figure 7(a) presents the results of the iterative correlation stacking algorithm for channels corresponding to NorDig 1; it shows that that an echo-free channel produces a well-defined correlation peak that sharpens with the number of iterations. In contrast, Figure 7(b), which presents results for NorDig 2, illustrates the previously discussed broadening effect on the correlation peak caused by the presence of the channel echo. As with the echo-free case, the correlation output in this case also improves with the number of iterations.

Figure 8 depicts the various scatter diagrams for the 2K, QPSK-modulated DVB-T system with a guard

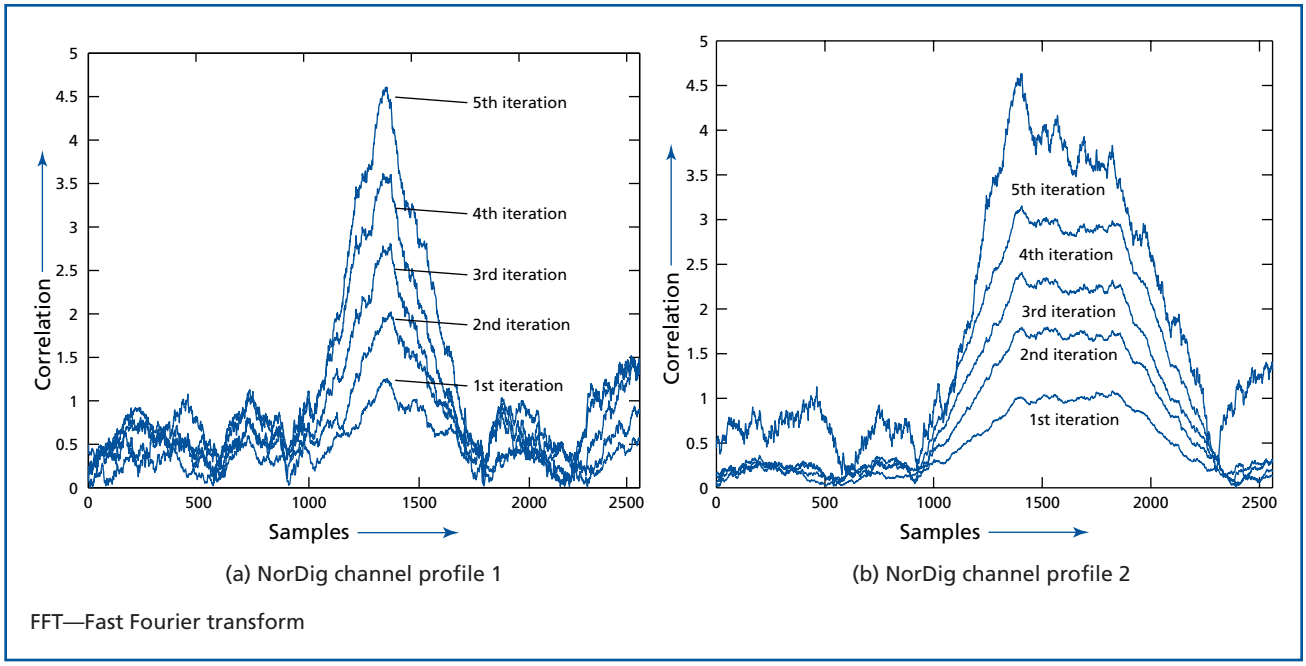


Figure 7.
Pre-FFT synchronization acquisition showing five iterations of the correlation output for NorDig channel profiles with a 0 dB echo at 0.95 times the guard interval.

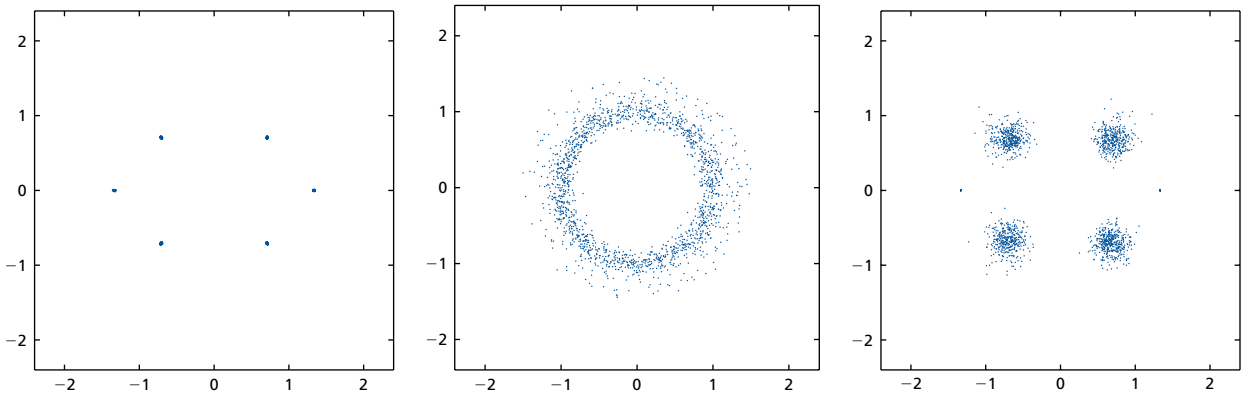
interval $T_G = 1/4\ T_U$. Figure 8(a) shows the scatter diagram of the input to the equalizer for a noiseless channel. One clearly sees a quasi-undistorted reconstruction of the transmitted 4-point QPSK constellation plus the pilot carriers. Figure 8(b) and Figure 8(c) show the scatter diagrams at the input and output of the equalizer for the NorDig Profile 1 channel where $C/N \approx 8$ dB. Similarly, Figure 8(d) and Figure 8(e) show the scatter diagram at the input and output of the equalizer for a NorDig Profile 2 channel for $C/N \approx 8$ dB and a 0 dB echo that occurs at a delay equal to 0.95 times the guard interval. Clearly, the noise and the echo cause a severe blurring of the constellation points that must be corrected by the equalizer. Figure 8(c) and Figure 8(e) show that the constellation points are now clearly identifiable. The resulting I/Q data streams were processed by the back-end, which produced QEF performance at the output.

The simulations are in general extremely time-consuming and for this reason the software simulator is intended to give an initial indication of performance, to generate test vectors, and to identify and resolve implementation errors. Once the developed RTL is tested on an FPGA test board, it will become feasible to determine the QEF performance for different configurations and thoroughly test the design in the field.

Summary

This paper provides insight into the status of an international collaborative effort, spanning two continents and led by Bell Labs, to design and develop a DVB-T RTL demodulation core for a provider of multimedia ASICs in Asia. The complex undertaking to design and develop the underlying OFDM-based system required a diverse and versatile team with experience in signal processing algorithms, simulation, bit accurate modeling, RTL code development, and project management and planning.

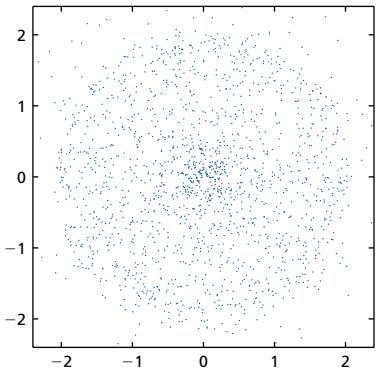
The demodulation core consists of two major signal processing blocks, a front-end and a back-end, which are controlled by an FSM. The front-end block takes the RF tuner IF output signal and outputs the I/Q data streams. The back-end block accepts the I/Q data from the front-end and delivers the MPEG transport stream, in either serial or parallel mode. The processing used



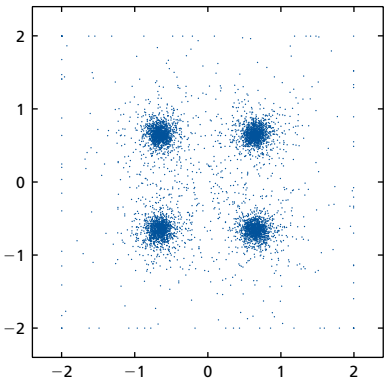
(a) Input to the equalizer for an ideal channel

(b) Input of the equalizer for Nordig channel profile 1 with C/N ≈ 8 dB

(c) Equalizer output for the input depicted in (b)



(d) Output of the equalizer for Nordig channel profile 2 with a 0 dB echo at 0.95 times the guard interval and C/N ≈ 8 dB



(e) Equalizer output for the input depicted in (d).

C/N—Carrier to noise
DVB-T—Digital video broadcasting-terrestrial
QPSK—Quadrature phase-shift keying

Figure 8.
Scatter diagrams for a 2K, QPSK-modulated DVB-T system with a guard interval T_G = 1/4 T_U

in these blocks is left open to different implementation solutions as long as its performance complies with the ETSI standard [2] and the NorDig performance specification [7]. The implementation process developed by the project team has been discussed and is characterized by the following steps:

- System design, algorithm selection and high-fidelity simulation
- Bit accurate modeling for system design, simulation, and validation
- FPGA RTL architecture, simulation, specification of registers
- Implementation and real time testing on an FPGA evaluation board
- ASIC RTL synthesis

Project milestones to date include completion of the FPGA RTL architecture and simulation of the back-end block, and of the system design, algorithm selection, and bit accurate modeling of the front-end block. Noteworthy technical accomplishments to date include a compact overall design, with particularly robust synchronization, and an efficient implementation of the FFT, Viterbi decoder, and RS decoder. In addition, a versatile simulation environment has been established.

Looking ahead, a number of possible enhancements to the DVB-T demodulator core are under consideration, including extensions toward DVB-H and diversity reception/combining. Because of the increasing popularity of portable DVB-T receivers and DVB-T enabled laptops, improving performance in a time-varying, low-velocity channel typical of portable reception indoors and outdoors is of particular importance. To achieve good portable reception performance, several algorithmic modifications will be investigated; these include diversity reception/combining to mitigate channel fading; better estimation and equalization of the time-varying, portable channel; and, possibly, compensation of subcarrier non-orthogonality caused by the time-varying channel. It should be noted that diversity reception/combining significantly increases the complexity of the ASIC design by requiring a diversity combining block and an additional ADC and front-end signal processing block for each additional antenna. Other enhancements of interest to be investigated are improved interference cancellation and optimization of the algorithms to achieve faster synchronization and lower power consumption.

Acknowledgments

The authors would like to thank Jacques Dungee, Hong Jiang, Zulfiquar Sayeed, and Les Wu for technical discussions and support.

*Trademarks

Matlab is a registered trademark of The Mathworks, Inc.

Xilinx is a trademark of Xilinx Inc.

References

[1] R. E. Blahut, Algebraic Codes for Data Transmission, Cambridge Univ. Press, Cambridge, 2003.

[2] European Telecommunications Standards Institute, "Digital Video Broadcasting (DVB): Framing Structure, Channel Coding and Modulation for Digital Terrestrial Television," ETSI EN 300 744, v1.5.1, Nov. 2004.

[3] C. W. Farrow, "A Continuously Variable Digital Delay Element," Proc. IEEE Internat. Symposium on Circuits and Syst. (Espoo, Fin., 1988), vol. 3, pp. 2641–2645.

[4] M. Horwitz and R. Braun, "A Generalized Design Technique for Traceback Survivor Memory Management in Viterbi Decoders," Proc. Symposium on Commun. and Signal Processing (COMSIG '97) (Grahamstown, South Africa, 1997), pp. 63–68.

[5] U. Ladebusch and C. A. Liss, "Terrestrial DVB (DVB-T): A Broadcast Technology for Stationary Portable and Mobile Use," Proc. IEEE, 94:1 (2006), 183–193.

[6] J. Lago-Fernández and J. Salter, "Modelling Impulsive Interference in DVB-T: Statistical Analysis, Test Waveforms and Receiver Performance," BBC R&D White Paper WHP080, Apr. 2004.

[7] NorDig, "NorDig Unified Requirements for Integrated Receiver Decoders: For Use in Cable, Satellite, Terrestrial and IP-Based Networks," NorDig-Unified v1.0.2, Apr. 2005.

[8] Philips Semiconductors, "The I^2C-Bus Specification," v2.1, Jan. 2000.

[9] U. H. Reimers, "DVB—The Family of International Standards for Digital Video Broadcasting," Proc. IEEE, 94:1 (2006), 173–182.

(Manuscript approved June 2007)

ALEXEI ASHIKHMIN is a member of technical staff



in Bell Labs' Communications and Statistical Sciences Department in Murray Hill, New Jersey. He received his Ph.D. degree in electrical engineering from the Institute of Information Transmission Problems, Russian Academy of Sciences, Moscow, Russia. Prior to joining Bell Labs, Dr. Ashikhmin was with the Electrical Engineering, Mathematics and Computer Science Department, Delft University of Technology, Delft, the Netherlands, and served as a Postdoctoral Fellow in the Modeling, Algorithms, and Informatics Group at Los Alamos National Laboratory, Los Alamos, New Mexico. His research interests include classical and quantum information theory, and communication theory. He currently serves as an Associate Editor for Coding Theory for IEEE Transactions on Information Theory. In 2002, Dr. Ashikhmin received the Bell Laboratories President's Gold Award for breakthrough research resulting in the ability to deliver unprecedented wireless bit rates. In 2005, he was honored by the IEEE Communications S. O. Rice Prize Paper Award for work on LDPC codes for information transmission with multiple antennas.

ADRIAAN J. DE LIND VAN WIJNGAARDEN is a member of technical staff in Bell Labs' Communications and Statistical Sciences Department in Murray Hill, New Jersey. He received his M.S. in electrical engineering from Eindhoven University of Technology, the Netherlands, and his doctorate in engineering from the University of Essen, Germany. He has been deeply engaged in both theoretical and application-driven research in communications, information theory, coding, and combinatorial and algorithmic optimization. He provided key contributions to Alcatel-Lucent's 10G and 40G optical systems, in-building wireless and broadband access. He has authored more than 60 technical papers, has received four patents, and has several patents pending. Dr. De Lind van Wijngaarden is a Senior Member of the IEEE and serves as a Publication Editor of the IEEE Transactions on Information Theory. He has co-organized Shannon Day at Bell Labs in 1998, Shannon's Statue Unveiling at Bell Labs in 2001, DIMACS conferences, and several international workshops.



ZHAO HAIBO is the chief technology officer and research and development lead at Cambridge Industries Group, Shanghai, China. He received his M.S. degree in communication engineering from Shanghai Jiaotong University, China. His research interests include access networking, communication VLSI architecture, and ASIC verification.



BERTRAND M. HOCHWALD was a distinguished member of technical staff in Bell Labs' Communications and Statistical Sciences Department in Murray Hill, New Jersey, until November 2005. He is now the chief scientist of Beceem Communications, Inc., in Santa Clara, California. He received an M.S. degree in electrical engineering from Duke University, Durham, North Carolina, and an M.A. degree in statistics and Ph.D. degree in electrical engineering from Yale University, New Haven, Connecticut. Prior to his tenure at Bell Labs, he worked for the U.S. Department of Defense, Fort Meade, Maryland. Dr. Hochwald was the recipient of several achievement awards while employed at the Department of Defense, as well as the Prize Teaching Fellowship at Yale University, and the IEEE Communications S. O. Rice Prize Paper Award in 2006. He has served as an editor for several IEEE journals and has given plenary and invited talks on various aspects of signal processing and communications.



THOMAS L. MARZETTA is the director of Bell Labs' Communications and Statistical Sciences Department in Murray Hill, New Jersey. He received his Ph.D. degree in electrical engineering from the Massachusetts Institute of Technology in Cambridge. Prior to his tenure at Bell Labs, he worked to modernize geophysical signal processing for petroleum exploration at Schlumberger-Doll Research and performed research and development at Nichols Research Corporation under contracts from the U.S. Department of Defense, NASA, and Schlumberger, where he headed a group that improved automatic target recognition, radar signal processing, and video motion detection. He specializes in multiple antenna wireless applications with particular emphasis on techniques for realizing extremely high throughputs with large numbers of antennas. Dr. Marzetta was a member of the Sensor Array and Multichannel Technical Committee of the IEEE Signal Processing Society (2000–2003) and has served as associate editor for two IEEE journals, and as guest associate editor for the IEEE Transactions on Signal Processing Special Issue on Signal Processing Techniques for Space-Time Coded Transmissions (Oct. 2002), and for the IEEE Transactions on Information Theory Special Issue on Space-Time Transmission, Reception, Coding and Signal Design (Oct. 2003). He was elected a Fellow of the IEEE in 2003.



VINAY PUROHIT is a consultant in Bell Labs' Video and Packet Networking Research Department in Murray Hill, New Jersey. He received his B.S. in computer science from the Indian Institute of Technology, Kanpur, and his M.S. in computer science from Texas A&M in College Station. In his prior work at Bell Labs, he served as the principal architect and lead designer of many mission-critical transmission elements used by Sirius Satellite Radio, Inc., including their main studio encoder and mux (SPACE), overlay FEC encoder, and TDM QPSK and OFDM modulators. He was also the principal hardware architect for the GEMS 160 Gbps Ethernet/MPLS router project. His research interests include advanced signal processing and high-speed packet networking.



CHEN QINGHONG is a senior ASIC design engineer at Cambridge Industries Group, Shanghai, China. He received his B.S. degree in material physics from Fudan University, China. His research interests include architectures, algorithms, and digital



integrated circuits design for communications, signal processing, and networking.

PAUL A. WILFORD is a Bell Labs Fellow and director of



the Video and Packet Networking Research Department in Murray Hill, New Jersey. He received the B.Sc. and M.Sc. degrees in electrical engineering from Cornell University, Ithaca, New York. Early on, his research focused on communication theory and predictive coding. He has made extensive contributions to the development of digital video processing and multimedia transport technology and was a key leader in the development of the first HDTV broadcast encoder and decoder. Under his leadership, Bell Labs then developed the world's first MPEG-2 encoder. Currently, he is working on the development and deployment of advanced video services in the network and leading several projects involving packet processing and advanced QoS and buffering algorithms to enhance packet services and reliability. He is interested in applying breakthrough technologies to deliver broadband and multimedia services over a variety of access networks more cost-effectively. He has authored a wide variety of papers and given many presentations and seminars on digital video and data networking.

SHENG-RONG ZHOU is a verification engineer at



Cambridge Industries Group, Shanghai, China. He received his B.S. and M.S. degrees in microelectronics from Tongji University, Shanghai, China. His interests include error control algorithms and digital integrated circuits design for communications, signal processing, and networking.

MICHAEL A. ZUNIGA was a senior technical manager in



Bell Labs' Video and Packet Networking Research Department in Fairfax, Virginia. He is now Chief Scientist of the Acoustics and Marine Systems Operations of SAIC in Arlington, Virginia. He received a Ph.D. degree in physics from the Massachusetts Institute of Technology in Cambridge. Early in his Bell Labs career, he conducted research in the adaptive control of radiative and scattering properties of media, participated in select JASON defense advisory group studies, and managed a major DARPA project involving Bell Labs, Sandia Labs, and the Naval Research Labs. Subsequently, his team pioneered the application of active noise cancellation to cellular terminals, developed sub-band acoustic echo cancellation for the General Motors' OnStar system, and conducted the initial systems engineering design for the TDM-OFDM based Sirius Satellite Radio System. Dr. Zuniga is a Senior Member of the IEEE and a member of the Electromagnetics Academy. He has authored more than 30 technical papers, has received 11 patents, and has 4 patents pending. His current research activities include DVB, intelligent content delivery over 3G networks, adaptive signal processing, electromagnetic wave propagation, and automotive telematics.

EDWARD S. ZURANSKI was a consultant with Bell Labs'



Video and Packet Networking Research Department in Largo, Florida. He is currently a consultant with Custom Manufacturing and Engineering in St. Petersburg, Florida. He received his B.S. degree (magna cum laude) in computer engineering from Syracuse University, and his M. Eng. degree in electrical engineering from Cornell University. He has been involved in the design, modeling, and implementation of a number of wireline and wireless communication products, and he holds 26 patents. ◆