

Real-Time Signal Acquisition, High Speed Processing and Frequency Analysis in Modern Air Data Measurement Instruments

Theodoros Katsibas¹, Theodoros Semertzidis¹, Xavier Lacondemine² and Nikos Grammalidis¹

¹*Centre for Research and Technology Hellas, Informatics and Telematics Institute
Greece*

²*THALES Aerospace Division
France*

1. Introduction

There is no doubt that the signal processing task is a very critical issue in the majority of new technological inventions and challenges in the area of electrical engineering. Despite the fact that during the last thirty years numerous algorithms and techniques have been developed in the field and classical methods of DSP technology were established, it is still an extremely skilled work to properly design, build and implement an effective signal processing unit able to meet the requirements of the increasingly demanding and sophisticated modern applications. This is especially true when it is necessary to deal with real-time applications of huge data rates and computational loads.

It is the objective of this book chapter to present all the successive steps of such a process in the framework of an ambitious plan to develop a functional and reliable lidar-based air data measurement instrument, launched by the FP6 Aeronautics and Space NESLIE project. Designed to perform in real-time (onboard the plane) and to provide vital flight parameters to the cockpit of the aircraft, this remote sensing equipment dictates the implementation of a robust, flexible and very fast signal processing unit.

This chapter starts by introducing the characteristics of the expected signal as well as the overall requirements of the system. Taking into account the very low signal-to-noise ratio (SNR) of a signal resulting from aerosol particles light backscattering and the different flight levels in a non-uniform and unstable environment, at which the measurement unit has to operate, we specify the appropriate algorithms and develop the necessary techniques for an effective signal processing scheme. The high speeds of modern aircrafts imply the use of a correspondingly impressive sampling rate which inevitably necessitates the use of parallelism of the computations and the optimum usage of all available resources. Very powerful FPGA boards have been selected for that purpose, making the fast processing implementation feasible. A "slow signal processing" module, performed in the CPUs and using input from these fast FPGA boards, produce the final measurements at a slower data

rate. This software integrates the entire system by estimating the air data parameters, presenting them to a suitable MMI and synchronizing all the modules of the signal processing unit.

Many simulations and real flight test results are also shown in order to prove the ability of the developed unit to perform real-time measurements of the vital flight parameters in a very effective and reliable way.

2. The Doppler lidar measurement mechanism

The well known Doppler effect is widely used in many radar, sonar and ultrasonic applications providing accurate velocity measurements of various remote objects or flows. Using much higher frequency waves and lasers instead of the radio and acoustic waves of the technologies above, it comes to the relative new technique of the Light Detection And Ranging or LIDAR. This technique has already found great application in many fields of science such as geology, seismology, geography, atmospheric chemistry and physics as well as meteorology. Calculating the frequency shift that moving targets may cause to the backscattered optical signal and estimating the corresponding radial component of their velocity, numerous important parameters can be measured performing advanced surveying and mapping tasks.

More specifically, Doppler lidars illuminate the aerosol particles or molecules in the atmosphere (depending on the target we have IR or UV emission respectively) and have already been proven to be an important tool for remote measurement of atmospheric wind fields and detailed observation of weather phenomena (Bilbro et al., 1984, Bilbro et al., 1986, Post & Cupp, 1990).

Starting from ground-based lidar systems, ongoing advances in optics and laser industry have led to the development of efficient airborne and space-based lidar units (Browell et al., 1998). Coherent CO₂ Doppler lidar, operating in NASA Marshall Space Flight Center at a wavelength of 10.6- μm (Bilbro et al., 1984, Bilbro et al., 1986), has been followed by the introduction of solid-state Doppler lidar, (Kavaya et al., 1989, Henderson et al., 1993, Frehlich et al., 1994), and




Fig. 1. The laser-based single particle detection and velocimetry

demonstrated the capability of the continuous wave system to perform single-particle detection. In addition, the fact that the Doppler frequency shift for optical wavelengths is high enough enables us to estimate the radial velocity for each backscattered pulse (Levin, 1965, Mahapatra & Zrnica, 1983, Frehlich & Yadlowsky, 1994). Furthermore, the advent of the

eye-safe 1.5- μm infrared laser technology that proved of great performance (Korb et al., 1997) allowed the implementation of low cost compact optical systems.

These last technological achievements showed the road to the development of reliable, airborne, lidar measurement instruments for commercial use in civil and military aircrafts. The demonstration of the fact that such an evolution may be realizable nowadays constitutes the aim of the EU funded NESLIE research project in the framework of which a laser-based, air data standby channel has been built. This system has been designed to perform real-time measurements of vital flight parameters such as the true air speed (TAS), the angle of attack (AOA) and the slide slip angle (SSA) to the cockpit of the aircraft. The unit is going to be used complementary to the existing, traditional, pneumatic, measurement systems (Pitot tubes) significantly increasing the safety of the flights.

2.1 Signal characteristics

It is pointed out that the ultimate goal of the signal processing task of such a sophisticated, real-time application is to constantly guarantee clear signal particle detection, accurate Doppler frequency shift estimation and finally the necessary vector velocity calculation. That is to provide the speed of the aircraft relative to the air mass in which it flies (aerosol scatterers). The longitudinal component of this 3-D air speed vector (projection to the laser axis) is given by the following equation

$$V_{long} = \frac{f_D \cdot \lambda}{2} \quad (1)$$

where, V_{long} is the longitudinal velocity component, f_D the Doppler frequency shift due to the relative motion between the aircraft and the air mass and λ the wavelength of the laser. During the time a particle crosses the beam in the small focal volume of the laser it is illuminated with an optical power of very high density and scatters back a significant amount of energy. The principal of the single particle laser-based velocimetry is schematically depicted in Fig. 1. Resulting from such a process, the received signal is a sequence of random bursts at random time instants plus noise. Fig. 2 shows two such modeled sinusoidal carriers with Gaussian envelope in a noisy environment. Of course, this is an idealized representation of the signal to be processed in order to show the randomness of the shape and the time duration of the expected, typical signal bursts as well as of the time instants they occur. In reality, on the other hand, a signal burst cannot be clearly distinct in the time-domain received signal since it has a very low SNR and it is deeply buried in noise.




Fig. 2. Two successive model signal bursts buried in noise

However, the frequency of such a burst corresponds to the value of the Doppler frequency shift and depends on the longitudinal velocity component as already written in Eq. (1). The time duration of each burst is a function of the transversal velocity component and the laser's beam diameter whereas the amplitude a function of the particle's cross-section and its characteristic trajectory crossing the laser beam. Of course, we mostly expect to receive only noise, especially at higher flight levels where low concentration of aerosol scatterers occurs.

2.2 Lidar system description

The developed lidar measurement unit follows the general principal of the coherent single particle Laser Doppler Anemometry (LDA) and it is concisely described in the block diagram form of Fig. 3. A continuous wave optical signal is being produced by the laser source and emitted in the atmosphere via the beam splitter and the separator. The received backscattering signal, which comes from a single, every time, aerosol particle and has been shifted in frequency according to the Doppler effect, is finally driven to the interferometer for comparison with the local oscillator reference beam (an identical copy of the emitted signal) coming from the beam splitter. The two beams are superimposed and feed the photo-detector which converts the incoming optical power into a practical electrical HF signal analogous to the frequency difference that these input waves present each other. That is the Doppler frequency shift information we are seeking for and can be effectively derived, in real-time, after the proper signal processing procedure we have developed and present in the following paragraphs.

Since we aim to deliver at a constant rate the three-dimensional vector of the true air-speed, at least three identical such laser units are needed to be installed on board the plane. However, one additional unit is going to be used for better precision and control as well as to secure continuing measurements in case of failure of one of the laser modules. All laser beams are focalized at the same measurement volume and at a properly chosen distance close to the fuselage, since we are interested in single particle detection and thus the focal volume should be relatively small. Fig. 4 illustrates the onboard set up of the lidar measurement equipment.

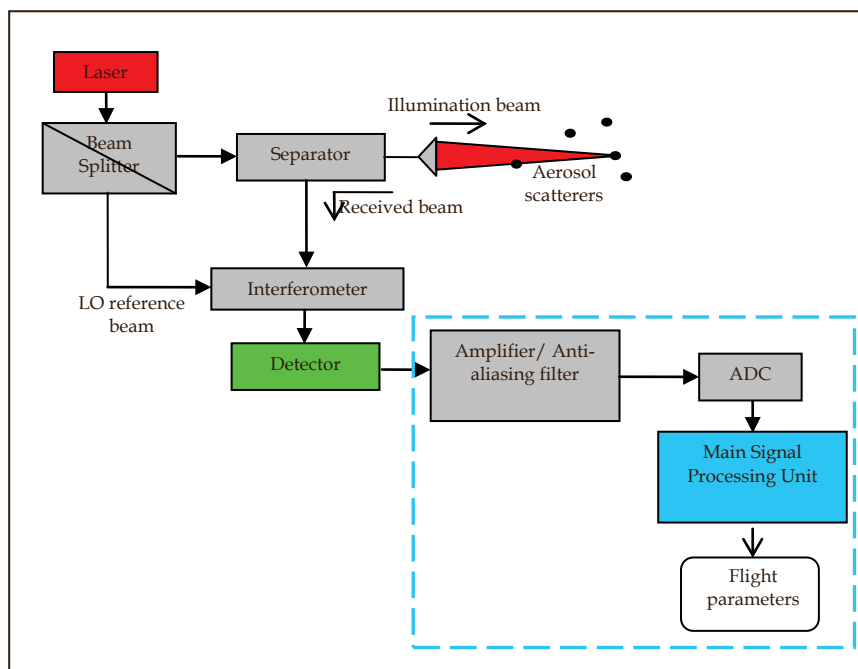


Fig. 3. Schematic diagram of the airborne lidar measurement unit

3. The signal processing unit implementation

The NESLIE signal processing task has been developed for the real-time acquisition and processing of the HF signals received from the detectors, in order to extract the aircraft velocity information and the other useful flight parameters. It is a composition of a fast signal processing firmware and a host software which both work in close association each other. The fast signal processing firmware is coming first to perform the signal burst detection and the Doppler frequency shift information. These operations have been implemented in properly selected powerful FPGA boards since such a real-time application dictates the processing of a data sequence that comes at an extremely high rate. Then, the slow processing software, executed in the host CPU, is responsible for the collection and processing of the produced data in the FPGA boards. Finally, the 3D air-speed and the other flight parameters are computed, recorded and sent to a suitable MMI for presentation.




Fig. 4. Onboard installation of the lidar equipment

3.1 Signal preprocessing

Just before proceeding to the main signal processing functions it is necessary to perform some basic preprocessing steps of the acquired signals. Thus, the amplification and the low-pass filtering (actually band-pass filtering since we are not interested in very low frequencies) of the continuous, time-domain signals, aim to eliminate the aliasing effect. That is the input signal at or above half the sampling frequency must be severely attenuated at a level below the dynamic range of the analog-to-digital (ADC) converter. The signal is then digitized in the ADC with m bits resolution, i.e. 2^m quantization levels are used to cover the full amplitude range of the received signal. It is noted that there is a limitation on the selected resolution since it is strongly related with the SNR of the signal in question. A relative high level of noise restricts the possibility for accurate resolution up to a certain number of bits. Any effort for increasing the quantization levels in order to perform further analysis does not add anything more than just measuring noise. The sampling frequency according to the Shannon-Nyquist sampling theorem should be higher than twice the highest frequency of interest. This is the maximum Doppler frequency shift that is expected to be caused to the received signal and it is determined from the corresponding maximum aircraft speed we wish to measure every time. As a consequence, the data samples have to be obtained at an extremely high rate.

3.2 The fast signal processing module

The constitutive tasks of the firmware implemented in the FPGA for the fast signal processing of the raw data stream, coming from the ADC, are presented in this paragraph. A simplified block diagram of the main FPGA firmware tasks are shown in Fig. 5.

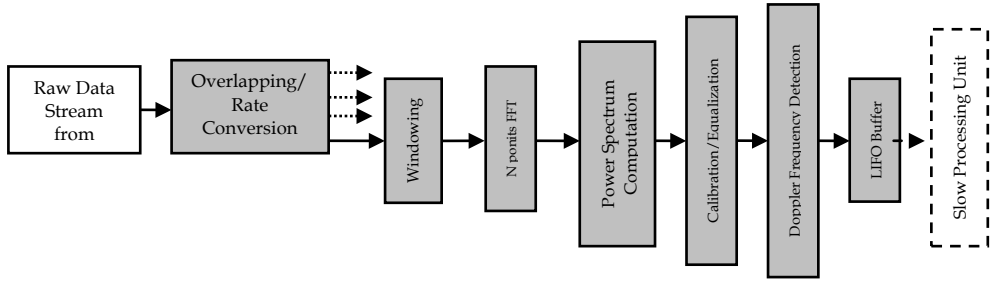


Fig. 5. The FPGA firmware schematic diagram

FFT-based signal analysis

Since the whole concept relies on the exact calculation of the frequency shift caused in the received signal due to the Doppler effect, the Fast Fourier Transform (FFT) and the power spectrum computation and analysis constitute powerful tools for effectively treating and processing the acquired data sequence (Proakis & Manolakis, 2006, Oppenheim & Schaffer, 1999).

In order to achieve an accurate Doppler frequency estimation the frequency range and resolution at the horizontal axis of the spectrum plot must be carefully examined and determined. These parameters depend on the defined sampling rate as well as the number of points of the acquired signal time-window (N points, the FFT size). The proper selection of these parameters is a very important issue since they are closely connected to the final accuracy of the whole set of measurements.

The number of frequency points or lines or FFT bins in the single-sided power spectrum is given from the equation:

$$FFT_{bins} = \frac{N}{2} \tag{2}$$

The start frequency is at 0 Hz (the DC term) whereas the last frequency bin corresponds to the value:

$$f = \frac{f_s}{2} - \Delta f \tag{3}$$

where Δf is the interval between two successive frequency bins or else the frequency resolution. This parameter is computed by the following equation:

$$\Delta f = \frac{f_s}{N} = \frac{1}{N \cdot \Delta t} \tag{4}$$

Δt is the sampling period and thus the product $N \cdot \Delta t$ is the length of the time-window record we process every time. Consequently, it is pointed out that the sampling frequency determines the obtained frequency range or else the bandwidth of the spectrum and that for a given sampling frequency we can determine the desired frequency resolution by choosing, in the time domain, the appropriate length of the signal record to be processed.

As a conclusion, the signal data are properly time-windowed in frames of N samples and driven to the FFT block for the frequency content calculation.

Windowing and spectral leakage

Processing, every time, a specific finite interval of the continuous signal data sequence is equal to applying the FFT to the product of the sampled signal and a rectangular window function (convolution of the corresponding spectra in the frequency domain).

The assumption in the FFT algorithm that the finite, in time, record is exactly repeated throughout all time and thus being a periodic signal may cause spectral leakage. This is the effect that the Fourier transform of a windowed, single frequency, signal has non-zero values at the neighbouring frequencies (side lobes). Fig. 6 shows the Fourier transform of a 100 MHz cosine windowed by a rectangular (a) and a Gauss (b) window function. It can be observed in both cases that the power spectrum does not contain only the single frequency of 100 MHz, as it would be expected, but shows some power content at the neighbouring frequencies as well. This is an effect of windowing the signal sequence in the time-domain and should be taken under serious consideration in this signal processing task.

The rectangular window is characterised as low-dynamic-range window and presents excellent resolution characteristics since the width of the main lobe is extremely narrow. The Gauss window, on the other hand, has a lower sensitivity but is of much higher dynamic range than the first. Since we are searching for a single frequency signal (i.e. the Doppler frequency) in a very noisy environment (narrowband application) and despite the fact that the generated bursts are expected to be Gaussian shaped by nature it is preferred the data frames be properly weighted by a Gaussian function and not just using the rectangular window. This is because the Gauss weighting window substantially improves the obtained SNR due to the high dynamic range it presents and at the same time it also reaches the excellent resolution characteristics of the rectangular window. Thus, it constitutes a good choice among other moderate window functions in an effort to reach the best trade-off between a satisfying SNR and an increased frequency resolution.

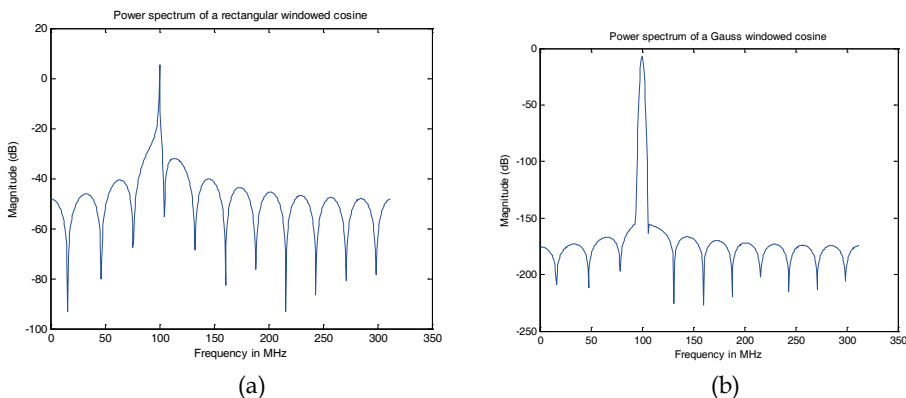


Fig. 6. Power spectrum of a 100 MHz rectangular (a) and Gauss (b) signal window

Window adjustment to the exact signal burst

It has already been mentioned that the signal bursts are not going to be of the same length. The obtained SNR can thus be further improved by fitting the time-window to the exact length of the shortest expected burst in an effort to ease the detection process. Fig. 7(a) depicts a relatively long signal time-window containing a backscattering burst, the length of which is marked by a rectangular frame, as well as the respective power spectrum. On the other hand, Fig. 7(b) shows the same burst in a time-window that almost fits its length and the corresponding power spectrum presenting a clearly improved SNR. This happens because it is just the bursts that contain all the useful information of the signal. The other parts contain pure noise and when they are subtracted from the data window the SNR is notably raised. This is, of course, at the cost of the obtained frequency resolution due to the fewer data samples used (FFT points).

The zero-padding technique may be employed to overcome that drawback by adding zeros and increasing the length of the data time-window to the desired magnitude. In this case, the detection is slightly improved due to the scalloping loss reduction and at the same time the frequency estimation is better enough, without the need for post processing.

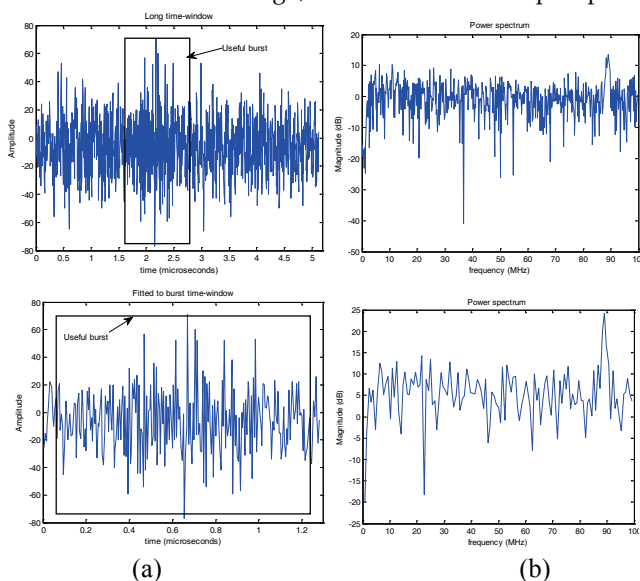


Fig. 7. Spectral analysis of a relatively long time-window containing a burst (a) and a time-window almost fitting the burst (b)

Overlapping

Since we acquire and process successive time windows of the signal sequence and we do not know the exact time a burst may occur, it is possible to lose the desired adjustment of the window to the full length of the burst if, for example, the latter happens between two successive data acquisitions. The SNR in such a case dramatically deteriorates making the detection doubtful. In order to tackle with the problem above the overlap processing is taking place. That means we can calculate every time the spectrum of a data sequence that contains part of the samples from the old time-window and part of the samples from the

new time-window. Fig. 8(a) illustrates the successive data windows for the 50% overlapping procedure.

Conducting overlapping in signal processing, however, has the result that the amount of the data to be processed be dramatically increased. In the 50% case, particularly, the computational cost is doubled. A serial-to-parallel conversion with adequate overlapping can fully secure the effective processing of the data sequence. In that way, a number of parallel data streams are generated and processed in blocks of N samples using all the available computational resources of the FPGA boards.

Calibration

Before detection, it is also necessary to calibrate the power spectrum for the bias removal and the equalization reasons. In order to calculate the calibration coefficients, spectrums are regularly acquired with the laser power set to off.




Fig. 8. Overlap processing of the data stream and power spectrum calculation, frequency analysis and detection

Doppler frequency detection

The burst detection and the Doppler frequency estimation are then performed by spectral analysis comparing the values of the obtained power spectrum with a properly defined threshold. The power values exceeding the threshold level constitute the detected signal whereas the FFT bin of the peak value determines the corresponding Doppler frequency. Fig. 8(b) shows the spectrum of the powerful burst depicted in Fig. 8(a) that can provide a strong signal and an easy detection.


Finally, the detected frequencies are registered in a FIFO buffer, from where they are read in the slow processing unit for the aircraft velocity calculation.

3.3 The real-time processing


Aiming to provide the real-time processing of the NESLIE signal data we had to study and define the highest frequency at which we can calculate a spectrum and the other necessary parameters without missing any data. That is we have to give careful consideration to the data collection time as well as to the processing speed of our FPGA processor or else to the

time it takes to perform the FFT spectrum analysis and all the needed additional computations and tasks described above.


There are three possible cases regarding the relation between the data acquisition time and the processing time as Fig. 9 shows. If the signal time-windows come with a rate higher than that at which the FPGA processor can perform the corresponding calculations, certainly there will be gaps in the acquired data (Fig. 9(a)). That results in the inevitable loss of an amount of data which may enclose useful information (i.e. the Doppler frequency shift of a burst). In such a case the whole procedure fails and the signal cannot be effectively analysed. On the other hand, if the processing speed is higher or equal to the speed that the signal windows are being collected we can really speak about a real-time application (Fig. 9(b) and Fig. 9(c) respectively).



(a)



(b)



(c)

Fig. 9. The three possibilities regarding the relation of the data acquisition time and the processing time.

3.4 Hardware selection and development tools

The Fast Signal Processing module was implemented on a special-purpose FPGA development board for programming a Xilinx Virtex 5-SX95T FPGA with more than 46,000 logic cells. This board, making optimum use of the inherent parallelism of the logic resources in the FPGA, can meet the application's advanced requirements, satisfy the strict timing constraints and effectively manage the huge computational throughput at a considerably high clock rate. The hardware platform is complemented with an ADC board and a PCI 64 bit carrier card.

The FPGA firmware was developed using the 3L's Diamond FPGA tool in association with the Xilinx ISE Foundation 9.2. 3L Diamond FPGA is a software development tool that shortens the development cycle by defining self-contained blocks of code or tasks that communicate each other. A task can be built for a DSP or an FPGA board and consists a simple but efficient abstraction from the actual hardware, hiding many of the low-level details and adding flexibility in the firmware design.

Diamond 3L tasks can be created with a wide range of FPGA tools, including Matlab Simulink, Celoxica DK suite, Impulse-C, Xilinx System Generator. It can build individual tasks into bitstreams that are finally incorporated into the complete application. Tasks are a good mechanism for sharing and exchanging ideas and technology.

Xilinx ISE® Foundation™ software integrates everything we need in a complete logic design environment for all leading Xilinx FPGA and CPLD products. The build-in tools and wizards make I/O assignment, power analysis, timing-driven design closure and HDL simulation. The Xilinx ISE tool is Xilinx's solution in compiling the design, targeting their hardware.

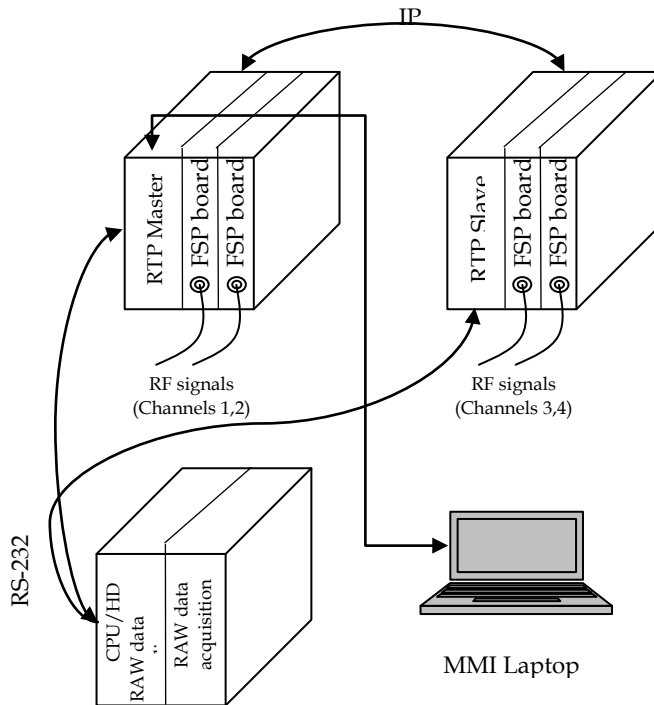


Fig. 10. The general architecture of the SP unit.

3.5 The architecture of the signal processing unit

Since there are four independent signal inputs, corresponding to the four identical laser measurement units installed on board the plane, the system has to perform the processing of these inputs in parallel so that the real-time estimation of the 3D air-speed vector becomes possible. The selected architecture uses two almost identical PC units, each using a dedicated FSP board to perform signal acquisition and processing of one channel. These two PCs, referred as the Real Time Processing (RTP) unit, exchange data between them via TCP/IP (Ethernet). This dual PC setup was selected in order to cope with the computational requirements and secure enough PCI bandwidth by splitting the data in the both PCI buses. Furthermore, it allows the system to function (reduced performance) even if only one PC board (two signal channels) is available.

An auxiliary PC board, referred as the Raw Data Recorder (RDR), is responsible for the control of the lasers as well as for raw data (digitized signal) recording for validation and post processing purposes.

One of the RTP PCs, referred as the RTP Master, can transmit information through to the RDR using RS-232 connection. However, both RTP PCs are able to receive information from RDR, using a simple T shaped cable connection. This approach was seen to reduce significantly the communication overhead required to achieve synchronization between the 3 PCs and makes the setup less vulnerable to Ethernet delays and/or communication errors.

A laptop is used as an MMI to display results of the system in a Web browser, as well as to allow the modification of an RTP configuration file containing all necessary system parameters. The entire system architecture is illustrated in Fig.10.

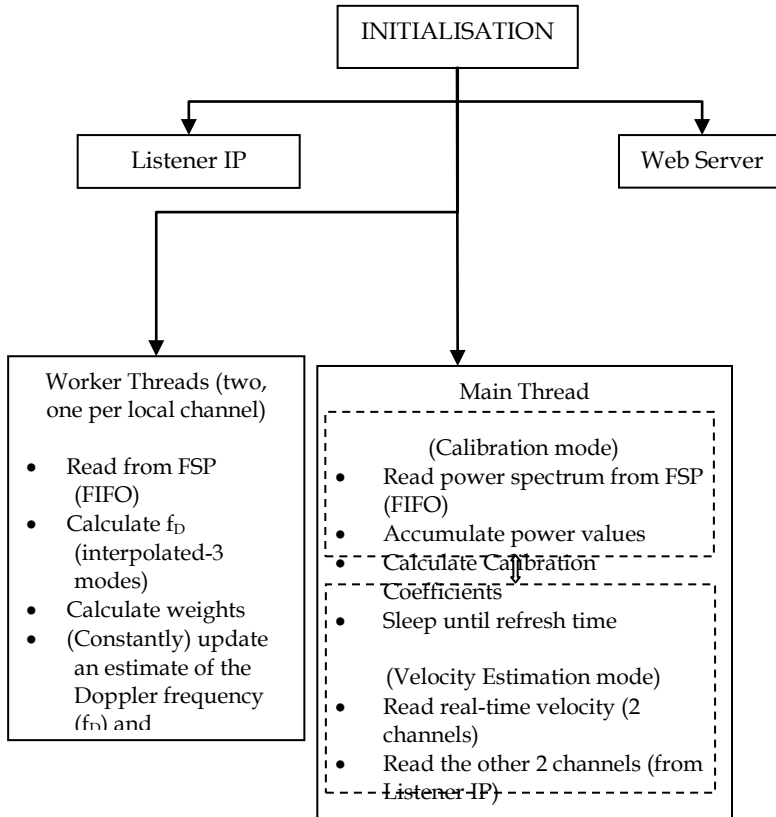


Fig. 11. Concise diagram of the SSP software.

3.6 The slow signal processing module

An efficient C++ software module was developed for the last stage of processing as well as the data storage and display. It is a multi-threaded application that ensures the necessary interconnection between the two independent units, the RDR and the MMI, the accurate computation of the 3D air-speed vector and finally the adequate output presentation at the User Interface (MMI). Fig. 11 depicts the threads that were defined as well as the main functions performed, which are almost the same for both RTP CPUs.

Immediately after its initialization, the procedure initiates four auxiliary threads responsible for:

- receiving and processing of information from the other SP unit and the RDR
- presentation of the final results (Web server)

- receiving and processing data from the FSP board (two “worker threads”, one for each local channel)

Two modes of operation are defined, i.e. the “real-time velocity estimation” and update (calculation of TAS components at the corresponding laser axes) mode, performed in a constantly repeated loop as well as the calculation mode, which is periodically used to update the calibration coefficients.

While the “Worker threads” perform their tasks as fast as possible, the main thread uses a fixed time cycle (defined by output (measurement) rate requirements). Within this time cycle, communication with the Worker threads and the remote RTP PC is used to obtain the four samples required to estimate the 3D velocity vector as well as the associated flight parameters, such as the True Air Speed (TAS), the Angle of Attack(AOA) and the Side-Slip Angle (SSA) and various error statistics. Finally, within the same data cycle, the results are being recorded into disk logs and are provided to the Web server thread, so that they can be displayed in the MMI.

The real-time velocity estimation function

In this function a set of seven values is being repeatedly read from the FIFO buffer in the FSP board. The first is a timestamp corresponding to the detection and the next is the index x of the maximum (FFT bin) which corresponds to the estimated Doppler frequency. Finally, five values $y[i]$, $i=1,\dots,5$ correspond to the useful power spectrum values of the detected burst (Fig. 12). More specifically, $y[3]$ is the maximum value of the power spectrum while its four neighboring values (two on each side) are $y[1]$, $y[2]$ and $y[4]$, $y[5]$ respectively.

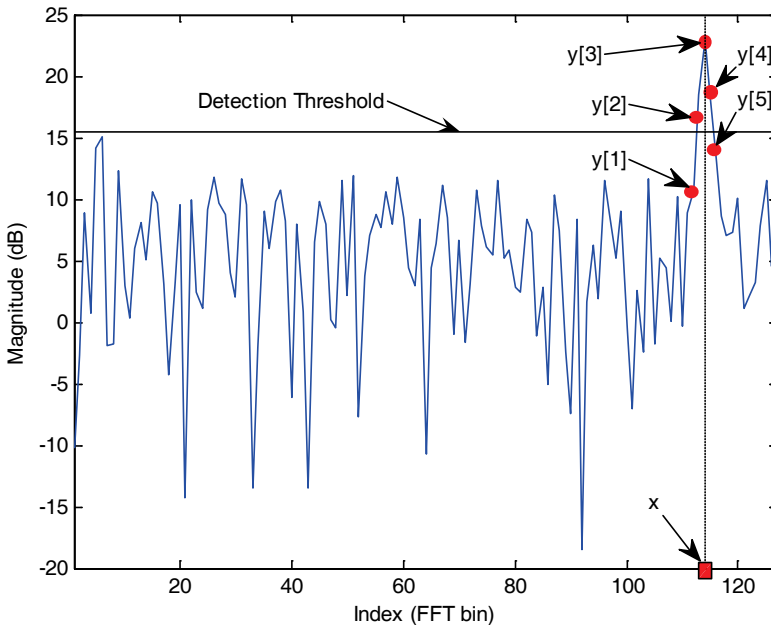


Fig. 12. The power spectrum of a detected burst and the six values read in the CPU from the FIFO buffer for the SSP.

In order to improve the resolution for the estimated Doppler frequency, an interpolation method can optionally be applied to the acquired data. Two interpolation techniques with low computational cost are parabolic interpolation using either 3 or all 5 available spectrum samples:

Three point parabolic interpolation

With that mode three values of the useful power spectrum are being used to form the interpolating parabola and calculate the interpolated frequency from:

$$x_{\text{int}} = x + \frac{1}{2} \cdot \frac{y[4] - y[2]}{2y[3] - y[4] - y[2]} \quad (5)$$

Five point parabolic interpolation

All five data points are being used to calculate the Least Squares parabola:

$$y = f(x) = \alpha_2 x^2 + \alpha_1 x + \alpha_0 \quad (6)$$

Fig. 13 depicts the useful spectrum (five data values) and the interpolated frequency value which is being computed from:

$$x_{\text{int}} = -\frac{\alpha_1}{2\alpha_2} \quad (7)$$

After every new detection and estimation of the corresponding Doppler frequency x_{int} , a weight coefficient is calculated and associated to that value in an effort to rate the reliability of the measurement. Two criteria can be used for the calculation of the weight w that is assigned to the estimated frequency value:

a) The power of the useful spectrum, in each signal detection, with respect to the detection threshold. The mean value of the five recorded spectrum magnitudes is being computed for that purpose:

$$\mu = \frac{\sum_{i=1}^5 y[i]}{5} \quad (8)$$

If that value is higher than the detection threshold ($\mu > \text{th}$) the weight w is set to: $w = \mu / \text{th}$, while otherwise it is set to zero.

b) The maximum frequency index of the detected burst has to lie within a periodically updated frequency window. The expected frequency window depends on the previous velocity estimations.

The next step is to calculate an estimate of the Doppler frequency f_D from all detections (valid frequencies estimates) obtained for the specific channel, within each cycle of the slow processing unit. For this calculation, three possible options are being supported:

a) A weighted mean f_D of the estimated Doppler frequency values using their corresponding weights:

$$f_D = \frac{\sum_{i=1}^n x_{int}(i)w(i)}{\sum_{i=1}^n w(i)} \tag{9}$$

The weighted variance is used as a reliability estimate in this case.

b) An estimate of the pdf of f_D is generated by estimating a histogram (or weighted histogram) of its values within each cycle. If the underlying function is known or can be assumed (e.g. Gaussian), then this function can be fit on the estimated pdf. The mean and the variance of the estimated distribution yield the estimated Doppler frequency and its reliability, respectively.

c) Robust estimate: This technique is similar to (b), however instead a robust estimate (e.g. the median) of the values is being used. This method is computationally efficient, avoids outliers, and has satisfactory resolution. The reliability of the estimate in this case can be defined based on the desired percentage of detections (percentage of samples within the window versus the total number of samples).

Finally, the component of the 3D air-speed vector at the corresponding laser axis is given by the equation:

$$v = \frac{f_D \cdot \lambda}{2} (V_{axis}) \tag{10}$$

where λ is the wavelength of the emitted signal. The value of this parameter is very important, so it is constantly measured by the RDR and is provided to the RTP in a regular basis. All the parameters of the estimation accumulators and histograms are reset to zero at the start of each cycle of the slow processing unit.




Fig. 13. The parabolic interpolation of the five spectrum values and the estimated index of the maximum.

Calculation of calibration coefficients

Since the lower frequencies in the pure noise power spectrum show the tendency to be stronger than higher frequencies it is necessary to calibrate the data after the FFT computation in order to remove the bias and ease the detection process. The calibration coefficients used for that purpose in the fast signal processing module (FSP) are calculated in a separate mode of the slow processing module (SSP).

During this mode, the laser power is set to 'OFF' and RTP accumulates successive noise spectrums. Fig. 14 shows the existing noise bias after accumulating 1000 noise spectrum frames. These values constitute the coefficients with which the corresponding data of the signal spectrum have to be divided (within the FSP) before each detection in order to eliminate the inherent inclination and equalise the noise level.

Another alternative to is interpolate the accumulated spectrum and divide the data with the new coefficients coming from liner, quadratic or polynomial interpolation (Fig. 14).

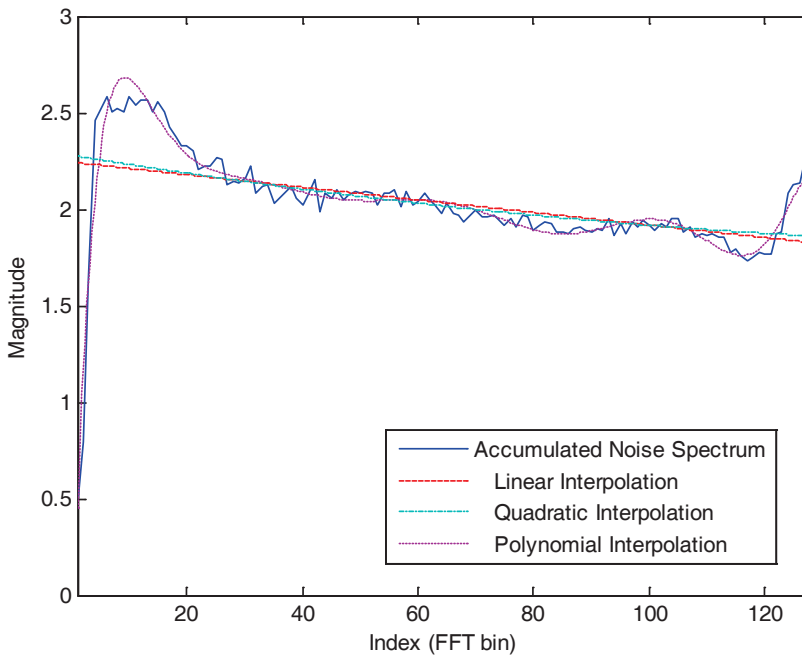


Fig. 14. Accumulation of 1000 noise spectrums and interpolation of the data.

3-D Velocity estimation at the main thread

As already described, the main thread reads velocity estimates for its two local channels, as well as for the two remote laser channels from the remote RTP PC via Ethernet. If all four channel estimates are available, we form a system of four equations with three unknowns (V_x , V_y , and V_z), which is solved by standard least squares:

$$\begin{bmatrix} \alpha_{1x} & \alpha_{1y} & \alpha_{1z} \\ \alpha_{2x} & \alpha_{2y} & \alpha_{2z} \\ \alpha_{3x} & \alpha_{3y} & \alpha_{3z} \\ \alpha_{4x} & \alpha_{4y} & \alpha_{4z} \end{bmatrix} \begin{bmatrix} V_x \\ V_y \\ V_z \end{bmatrix} = \begin{bmatrix} V_{axis1} \\ V_{axis2} \\ V_{axis3} \\ V_{axis4} \end{bmatrix} \tag{11}$$

where α_{ik} are the x, y, z coordinates of the unit vectors at the four laser axis ($i=1,2,3,4$ and $k=x,y,z$). A solution can also be found even if only three channels are available.

4. Flight results

A number of flight tests have been planned and performed in order to test the operation of the developed unit as well as to prove its effectiveness for on board real-time air-data measurements.

During the flights, the calculated parameters as well as all the other useful system’s data are being recorded, constantly updated and properly displayed in a functional Man Machine Interface (MMI). The Doppler frequency histogram as well as the power spectrum (SNR) histogram corresponding to the two local channels are also presented.

Fig. 15 shows the MMI at such a state during one of the flights.

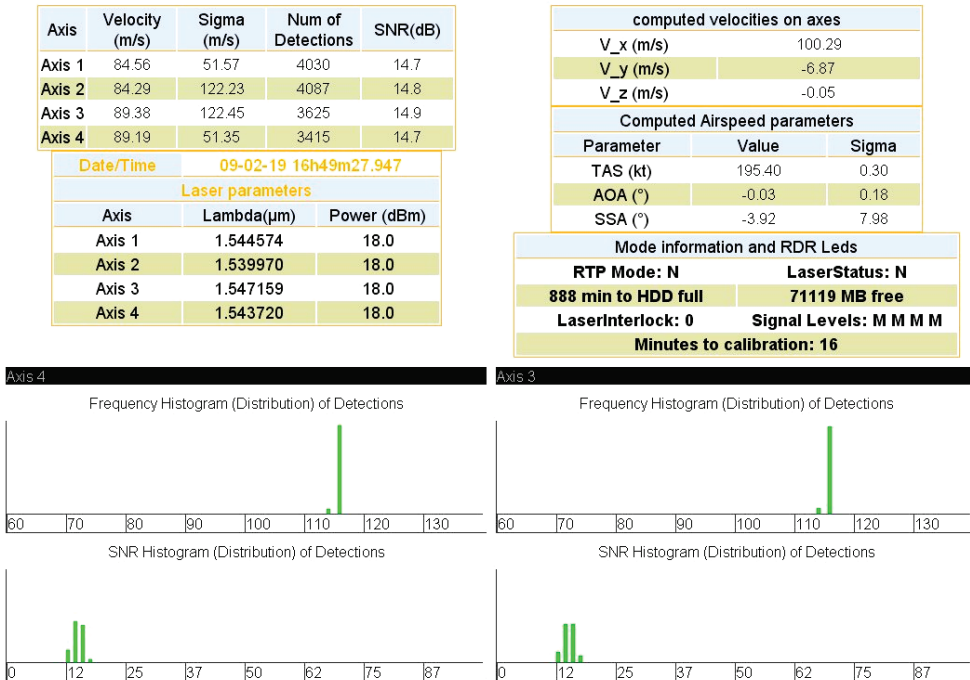


Fig. 15. The MMI of the NESLIE measurement unit at a specific time instant, presenting the computed flight parameters and other useful data




Fig. 16. The flight altitude as a function of time during one of the tests




Fig. 17. Comparison of NESLIE measurements of True Air-Speed (a) and Angle of Attack (b) during a flight with the reference data recorded by the aircraft

The testing flights covered a wide range of altitudes (up to 39.000 feet) as well as different weather conditions in order to better evaluate and validate the performance of the system. Fig. 16 depicts the flight level during such a test according to the time.

The NESLIE measurements were also recorded in order to compare them with the corresponding data from the aircraft’s standard measurement channels (Pitot tubes). The obtained results show a significant agreement between these two completely different modes of air-speed measurements as it can be clearly seen in Fig. 17(a) for the true air-speed (TAS) of the aircraft during the flight and in Fig. 17(b) for the angle of attack (AOA). However, a more thorough analysis of the test results is required and is already being conducted by the NESLIE consortium.

5. Summary

The methodology and the consecutive steps of an efficient signal processing scheme, developed to provide real-time measurements of the true air-speed and other useful flight parameters in the framework of the EU funded NESLIE research project, were analytically presented.

The NESLIE signal processing unit consists of a distributed system that uses two embedded PCs and four FPGA signal processing boards, corresponding to each of the laser axis. The fast signal processing module performs the analog-to-digital conversion, the frequency analysis (FFT) and the burst detection from the adequate spectrum values. These operations are implemented in powerful FPGA boards that finally provide the detected Doppler frequency as well as the useful power spectrum. Using the data above, the slow signal processing module calculates the vector velocity of the aircraft and the other parameters necessary for the safety of the flight. It is the software for disk logging and also responsible for the MMI and communication tasks with the other real-time processing PC and the RDR server.

The whole unit was set on board for flight testing showing more than satisfactory results. That encourages us to proceed to further optimization of the modules demonstrating that such an instrument can really work as a reliable air-data standby measurement channel.

6. Acknowledgements

This work was supported by the European Community (EC) under the FP6 Aeronautics and Space "NESLIE - NEW Standy Lidar InstrumEnt" project (Contract no.: 30721).

7. References

- Bilbro, J.; Fichtl, G.; Fitzjarrald, G. & Krause, M. (1984). Airborne Doppler lidar wind field measurements. *Bull. Am. Meteorol. Soc.*, 65, pp. 348-359.
- Bilbro, J. W.; DiMarzio, C.; Fitzjarrald, D.; Johnson, S. & Jones, W. (1986). Airborne Doppler lidar measurements. *Appl. Opt.*, 25, pp. 3952-3960.
- Browell, E. V.; Ismail, S. & Grant, W. B. (1998). Differential absorption lidar (DIAL) measurements from air and space. *Appl. Phys.*, B 67, pp. 399-410.
- Frehlich, R.; Hannon S. M. & Henderson, S. W. (1994). Performance of a 2 μm coherent Doppler lidar for wind measurements. *Journal Atm. Oc. Tech.*, vol.11, pp. 1517-1528.
- Frehlich R. G. & Yadlowsky, M. J. (1994). Performance of mean-frequency estimators for Doppler radar and lidar. *Journal Atm. Oc. Tech.*, vol.11, pp. 1217-1229.
- Henderson, S. W.; Suni, P. J. M.; Hale, C. P.; Hannon, S. M.; Magee, J. R.; Bruns D. L. & Yuen, E. H. (1993). Coherent laser radar at 2 μm using solid state lasers. *IEEE Trans. Geosci. Remote Sens.*, 31, pp. 4-15.
- Kavaya, M. J.; Magee, J. R.; Hale, C. P. & Huffaker, R. M. (1989). Remote wind profiling with a solid-state Nd:YAG coherent lidar system. *Opt. Lett.*, 14, pp. 776-778.
- Korb, C. L.; Gentry, B. M. and Xingfu Li, S. (1997). Edge technique Doppler lidar wind measurements with high vertical resolution. *Applied Optics*. Vol. 36, No. 24, pp. 5976-5983.

- Levin, M. J. (1965). Power spectrum parameter estimation. *IEEE Trans. Inform. Theory*, vol. 11, pp. 100-107.
- Mahapatra, P. R. & Zrnic, D. S. (1983). Practical algorithms for mean velocity estimation in pulse Doppler weather radars using a small number of samples. *IEEE Trans. Geosci. Remote Sens.*, vol. GE-21, No 4, pp. 491-501.
- Oppenheim A. V. & Schafer, R. W. (1999). *Discrete-Time Signal Processing*. Prentice Hall, 1999.
- Post, M. J. & Cupp, R. E. (1990). Optimizing a pulsed Doppler lidar. *Appl. Opt.*, 29, pp. 4145-4158.
- Proakis J. G. & Manolakis, D. K. (2006). *Digital Signal Processing*. Prentice Hall (forth edition).