Selected Data Analysis Techniques for Equipment Monitoring Using Drive's Control Signals

Michal Orkisz[*], Maciej Wnek[†], Pieder Joerg[†], Klaus Ruetten[†], Edgar Jellum[‡]

Abstract -- **Variable speed drives are becoming the method of choice for powering AC motors in many industrial applications. They allow for flexible speed control, substantial energy savings, and reduced stress during motor start-up, all of which lead to lower lifecycle costs for the shaft line. Modern drives, in order to perform their controlling function, are equipped with sensors, hardware and software to measure or compute a number of signals (for example, current, torque, speed, frequency, power, flux, etc.). This data is available and can be used for diagnostic purposes. These signals can be analyzed for their spectral content, used for monitoring the operating point, for examining cyclic processes, etc. The data is freely available, but the format, quality and quantity are determined by the hardware. Furthermore, variable speed drives change the output frequency to follow desired process parameters. This work presents selected methods for handling these difficulties and turning them to our advantage. Some examples of these methods application to industrial data are presented. These results prove that a variable speed drive can be a valuable source of diagnostic data not only for the drive itself, but also for the entire driven process.**

Index Terms -- **Algorithms, Fault diagnosis, Motor drives, Spectral analysis**

I. Introduction

In order to provide quality service, correct diagnosis of equipment state is needed. It can be based on Condition Monitoring data. The benefits of Condition Monitoring are obvious [1]-[2], but it is not used as widely as it might. One reason is the costs associated with sensors, cabling, installation, measuring/processing hardware and software, as well as time invested by the personnel. Lowering the costs would make Condition Monitoring more ubiquitous. What if the data were already available? This is often the case when Variable Speed Drives are used to power rotating machinery. Drives collect and generate large quantity of data for the purpose of control and operation. Often this data can be used to perform diagnostics, not only of the drive itself, but of the whole process chain. This is because the drive is always connected to other equipment. On the input side there are the supply network, transformer, circuit breaker, etc. On the output side – electric motor, that in turn drives pumps, compressors, fans, conveyor belts, gearboxes, steel mill rolls, etc. There are several potential problems, though. We are limited to the signals that the drive provides. Data gathering capability (number of samples, sampling rate, A/D converter resolution) is also fixed. Also, true to their name, variable speed drives may change the driving frequency, challenging the simple-minded application of spectral analysis methods (as explained in [3]-[4]). Drive's control system will attempt to compensate for any speed/torque/ current deviations due to motor defects. On the positive side, data is available in vast quantities, at any time, allowing for selection of "best" items, looking at both short-time and long-time effects. This work focuses on taking advantage of these opportunities.

[*] ABB Corporate Research Center, Starowiślna 13A, 31-038 Kraków, Poland; e-mail: michal.orkisz@pl.abb.com
[†] ABB Schweiz AG, Austrasse, CH-5300 Turgi, Switzerland
[‡] ABB AS, Ole Deviks vei 10, NO-0666, Oslo, Norway

DOI: 10.1109/DEMPED.2009.5292762

II. Equipment

The capacity to measure, store and transmit data will depend on particular drive and its control platform. Here we focus on the ACS family of ABB AC Variable Speed Drives, with particular focus on the Medium Voltage drives (but also applicable to the Low Voltage drives of the same family, as they share the control platform). Internally, the signals (measured and computed values, such as speed, frequency, torque, flux, current, power, temperatures, etc.) and parameters (configurable drive settings) are stored in a regularly updated memory table. Current values can be read from this table as OPC values, or into hardware data loggers, available via OPC in bulk. Data loggers are programmable hardware buffers capable of storing values from several selected variables concurrently with a specified sampling rate. They are triggerable in different ways, such as by a fault or alarm, by selected variable's value crossing a threshold, by software command, etc. The sampling frequency is high enough to make the data useful for spectral analysis (though no anti-aliasing filtering is present). The data is collected using DriveMonitor™ system consisting of an industrial PC and collection/analysis software. The details of this system can be found elsewhere [5]-[7].

III. Adjusting the Data

A saying has it that "there ain't no free lunch". This is also true here. The data is free, but, coming from a device that has not been designed for condition monitoring, it may exhibit problems and limitations that require extra work to circumvent or correct. This is a general problem encountered whenever we deal with hardware limitations, whether the data comes from drives, motor control centers, protection relays, or other devices. Therefore we will devote some space to discuss it here.

Ideally one would like to have a sufficiently high sampling rate (≥40kHz), long acquisition times (≥4s), and constant operating point conditions when acquiring the data. In practice often the buffer length is limited, so a tradeoff between sampling rate and acquisition time is needed. The main problems are

- inverter-introduced noise
- varying output frequency
- lack of anti-alias filtering
- limited buffer length
- limited A/D converter resolution

Some of these can be circumvented, others we have to live with. For example, A/D converter resolution may be treated as a source of noise. For an N-bit converter the RMS value of the quantization error is $2^{-N}/\sqrt{12}$ times the full scale amplitude and, for M-point FFT, the average value of the noise contained in each frequency bin is 10 $\log_2(M/2)$ dB below that RMS value [8].

A. Inverter-introduced noise

This is a noise that is introduced by the inverter's output being only approximately sinusoidal. It is important to understand its nature and magnitude. In simple Pulse Width Modulation (PWM) schemes, the switching frequency is held constant (independently of the load and speed), and only the duty cycle is modified. The switching frequency harmonics are easy to identify and eliminate from the output current spectrum. Direct Torque Control (DTC) scheme employed in the investigated drives produces a non-deterministic switching pattern, where the switching frequency is constant only in an average sense. DTC spectrum is non-stationary: it varies even when the speed and load are constant. Fig. 1 shows an example of two sets of spectral data (inverter output current) under constant conditions. It is obvious that high variability and presence of "random" peaks would make it difficult to interpret an individual spectrum.

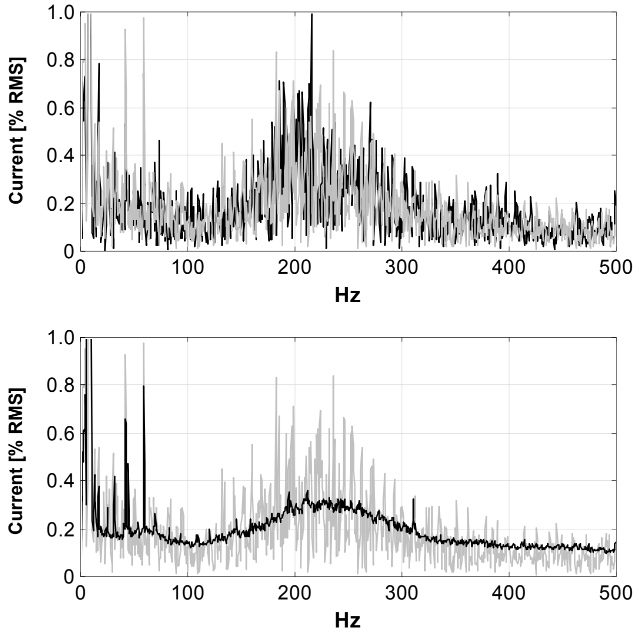

Fig. 1. Individual output current spectra for identical operating points (above), and an averaged spectrum (below).

B. Non-constant output frequency

True to their name, variable speed drives can change the output frequency to modify motor's speed or torque. This can happen due to an external request (e.g., command to pump more water), or due to process changes (e.g., more load on a conveyor belt increases the slip of an asynchronous motor), or a combination of both. When collecting data we have several options. One would be to collect only when certain speed/load conditions are met. As the measurements are on-line, it is not difficult to implement such scenario. Another alternative would be to collect all the data and to bin them according to the current operating point (e.g., speed, torque and current). Finally, we can average the spectra obtained from the data, as described below.

C. Averaging of spectra

The technique of spectra averaging is normally used to decrease the variance of the computed spectrum [9]. Each point of the final spectrum is computed by taking the mean of the corresponding points of the K spectra being averaged. In the process any transient features are averaged out, while stationary features common to all spectra remain. This is illustrated in the bottom part of Fig. 1, where the averaged transient features form a smooth background (including a "hump" around the average switching frequency), while the stationary features are identifiable as sharp peaks.

The procedure described above averages individual points based on their absolute frequency (e.g., amplitudes at 100Hz for each spectrum are averaged together). Other choices of x-axis scaling are possible. For example, individual spectra can be scaled to the orders of the output frequency. In such case, even if the spectra being averaged had been collected at different values of that frequency, the output frequency harmonics would fall in the same place (e.g., amplitudes at 6 times the output frequency for each spectrum are averaged together). Each scaling choice "preserves" certain families of peaks, while "averaging out" others, when spectra from a sufficient number of operating points are considered. Absolute frequency scaling keeps the frequencies that do not change between data sets. These include the line frequency (supplying the drive), or structural resonance peaks (e.g., torsional oscillations). Scaling to the orders of the output frequency preserves the output frequency harmonics. Scaling to the rotating frequency keeps the features related to the mechanical speed.

It remains to be discussed how to average the spectra when the x-values fall in different places (e.g., when scaled to orders of the output frequency). We proceed by defining equispaced frequency "bins", and assigning each frequency in any spectrum to its nearest "bin". Then all the values in a given "bin" are averaged. Alternatively, instead of averaging, median can be taken (it is more "robust" than the average [9]). Even though more sophisticated schemes can be devised (e.g., weighting each frequency's contributions to the neighboring bins), the scheme just described gives adequate results and is very straightforward.

D. Spectral aliasing

This section describes both the problem and the solution to the phenomenon called aliasing. The sampling theorem [10] tells us that a continuous signal is uniquely determined by its discrete values probed at rate R, but only if it contains no spectral content at frequencies higher than the Nyquist frequency R/2. The usual practice is to low-pass filter the signal prior to digitizing it – to ensure this condition. In our case there is no pre-filtering – the signals were not designed for condition monitoring. They are probed/generated at rates up to 40kHz, but are available at lower frequencies (e.g., multiples of 1kHz or 10kHz), so spectral content above the Nyquist frequency is usually present. This is a problem, but may be a blessing in disguise, as discussed below.

Spectral features above the Nyquist frequency are mapped onto the computed spectrum range according to:

$$F_{app} = R/2 - |(F_{true}\ MOD\ R) - R/2| \quad (1)$$

where F_{true} is the true frequency of a spectral feature (peak),

F_{app} is its apparent position in an aliased spectrum, and R is the sampling rate ($R/2$ is the Nyquist frequency). A good analogy is a piece of transparent foil with the true (i.e., sampled infinitely fast) spectrum drawn. Folding the foil accordion-like along the multiples of $R/2$ produces the apparent spectrum.

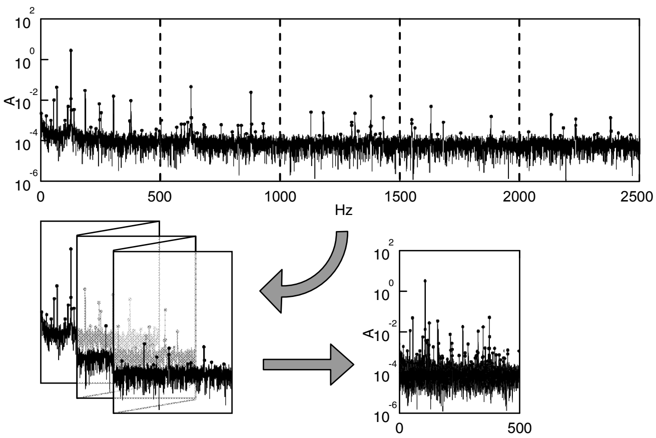

Fig. 2. Folding foil analogy of aliased spectral content. True spectrum (above) is "folded" (below left), so all present frequencies overlap in the aliased low-frequency spectrogram (below right).

The bad news is evident from Fig. 2 – all spectral features are crammed into the computed range, so it is difficult to interpret the peaks. The good news is that the spectral content is still present in our signal (not removed by anti-alias filtering). With some work it can be recovered, at least partially.

The procedure is based on spectra averaging described above. The requirement is that the output frequency varies slightly between individual data sets. Begin by "unfolding" each individual spectrum: appending alternatively reversed and straight copies of the computed spectrum (see Fig. 3). We repeat this procedure till the highest frequency of interest is included (e.g., 3-5 times). In this manner each peak is repeated once in each "fold". One of these copies will occur at the actual frequency. For example, Table I considers a spectrum sampled at 1000Hz and the 12th harmonic of the 100Hz output frequency (1200Hz). The Nyquist frequency is 500Hz and, according to (1) the 12th harmonic will appear in the spectrum at 200Hz. Repeating the "unfolding" three times will produce additional copies at 800Hz (reversed copy), 1200Hz (straight copy) and 1800Hz (another reversed copy). We managed to reconstruct the true frequency (order 12) in the 3rd fold.

TABLE I

FREQUENCIES (HZ AND ORDERS) OF VARIOUS APPARENT PEAKS DUE TO ALIASING OF THE 12TH HARMONIC FOR DIFFERENT OUTPUT FREQUENCIES.

Freq. &	12th harm	base		2nd fold reversed		3rd fold straight		4th fold reversed	
Hz	*Hz*	*Hz*	*Ord.*	*Hz*	*Ord.*	*Hz*	*Ord.*	*Hz*	*Ord.*
99	1188	188	1.9	812	8.2	1188	**12**	1812	18.3
100	1200	200	2.0	800	8.0	1200	**12**	1800	18.0
101	1212	212	2.1	788	7.8	1212	**12**	1788	17.7

Generally, during the operation of the drive, the output frequency will shift somehow (or a lot), due to changing process requirements. When the spectra are scaled to the orders of the output frequency, the spectral peaks related to the output frequency will all match; while any other peaks will shift their relative position (see other rows of Table I).

Continuing our example, consider the output frequency shifting to 99Hz. The 12th harmonic is now 1188Hz, aliased into the spectrum at 188Hz=order 1.90. Looking at Table I we notice that the copy of the peak corresponding to the true frequency stays at the correct order (12th) independently of the value of the output frequency. This is not true for the copies in other "folds". Thus, when averaging multiple spectra, the "true" peaks will overlap, producing a high average, while other copies will get "averaged out" (Fig. 3).

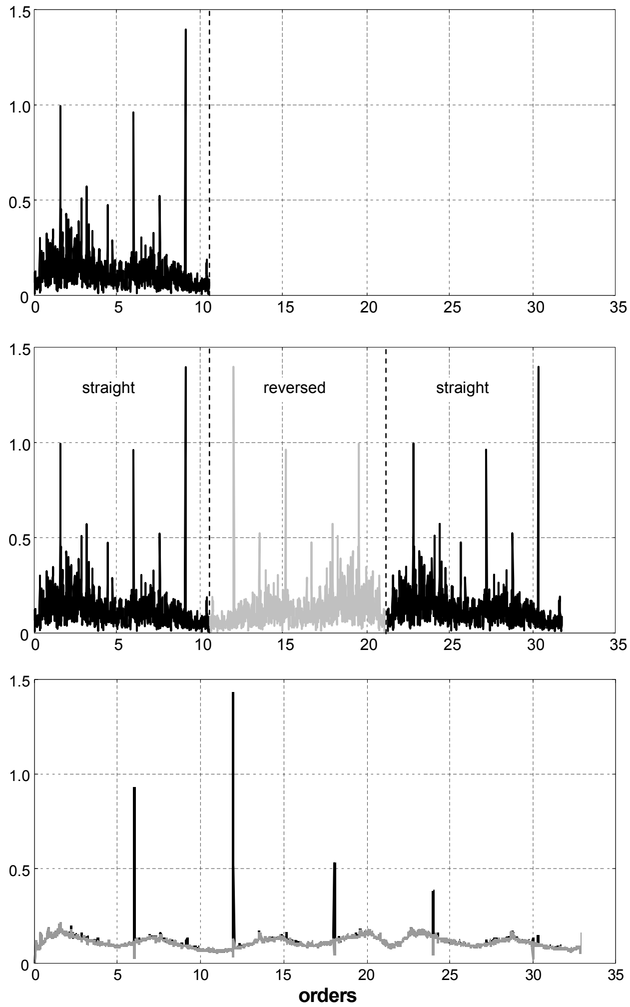

Fig. 3. Individual electric torque spectrum with aliased peaks (top), "unfolded" spectrum (middle) and averaged, "unfolded" spectrum (bottom).

This method allows us to recover spectral content even beyond the Nyquist frequency.

IV. DIAGNOSTIC OPPORTUNITIES

This section discusses examples of diagnostic information that can be obtained from properly processed drive data. A general approach is to provide a set of Key Performance Indicators that can be tracked over time. In some cases their absolute value can have a well-defined meaning, in others it would be their change over time that provides diagnostic information.

One of the simplest applications is (for drives equipped with active rectifier unit) a power quality meter. Total harmonic distortion of drive's input voltage and current signals can be indicative of improper power supply quality.

Another obvious example is looking at the total harmonic distortion of the motor stator current (drive's output current). Measurement of this current can also be the basis of Motor Current Signature Analysis, taking into account the closed-loop control of the drive [3]-[4].

Other examples, such as operating point tracking, transient oscillation detection and cyclic load variability are described elsewhere [7].

One advantage of using the signals from the drive is that we have access to more than just the currents and voltages. As the drive controls the speed and/or torque, it has to keep track of these and other quantities. It is also aware of the flux, RMS current, electric power, instantaneous *cos* φ, and many other intermediate signals.

Below are two further examples of information obtainable from drive-supplied data.

A. Determining the output frequency

Generally, the output frequency can be read from the drive. However, sometimes it is impractical to do it, e.g., all channels are used to collect other signals. In such case we can determine the output frequency directly from the data. Many signals show harmonics of 6 times the output frequency. We can use these peaks (even if aliased) to enhance the estimate of this frequency significantly. To see why, let's return to the example in Table I. Examining the 12^{th} harmonic peak, we see that its apparent position's change is amplified 12 times: shifting the output frequency from 100 to 101Hz changes the peak's position from 200 to 212Hz. This means that establishing the peak's location with uncertainty ε, the uncertainty of the fundamental (computed based on the 12^{th} harmonic) is reduced to ε/12.

A crude way of determining the peak's frequency is by taking the location of its maximum amplitude. The uncertainty ε is then ± half the Fast Fourier Transform's resolution. A simple improvement is fitting a parabola to this maximum and the two points around it and using the calculated location of this parabola's maximum. More sophisticated schemes along this general idea give even better results.

B. Torsional resonance

An example of diagnostic information obtained from drive-supplied signals relates to torsional oscillations on a compressor. The system consists of a medium voltage drive, induction motor, gearbox and a compressor. It has a resonant frequency at 16.7Hz, which sometimes gets excited. Before stopping production to fix the problem, the plant operator was interested in discovering whether it is possible to tract it. It turns out that many signals contain a trace of it. One example is the electric torque spectrum. Fig. 4 shows two examples of such spectra – one containing the resonant peak, the other not. By collecting the data every few minutes and computing the amplitude of the peak near 16.7Hz it is possible to continuously track the extent of the torsional oscillations, thus providing a measure of the problem's severity. Fig. 5 illustrates a histogram of the peak's amplitudes. For 2/3 of the cases the peak was discernable. So far we failed to correlate the amplitude with any other quantities. However, the very presence of the resonance is undesirable, as it puts extra stress on the motor, gearbox, bearings, etc.

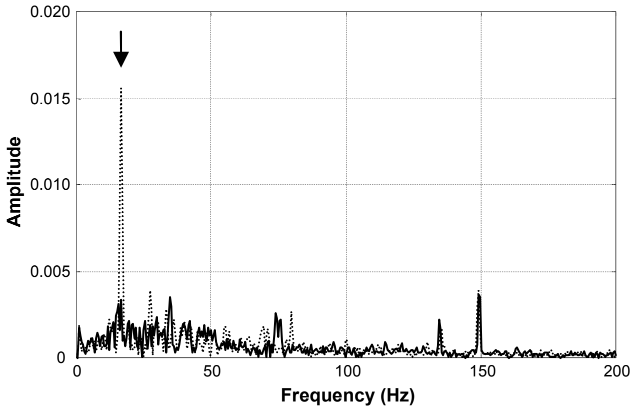

Fig. 4. Two spectra – one shows the evidence of the resonance peak at 16.7 Hz (dotted line), in the other the peak is absent (solid line).

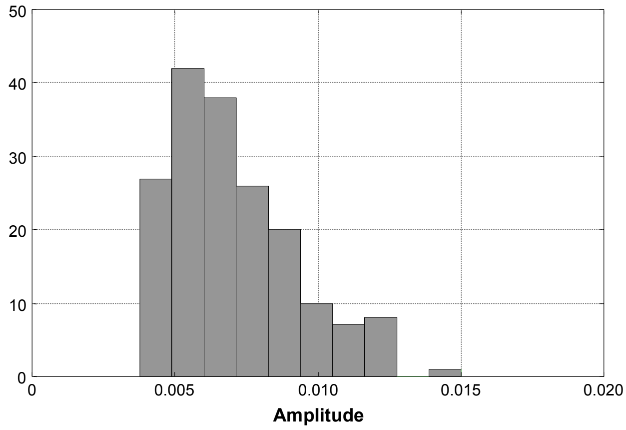

Fig. 5. Histogram of the resonance peak amplitudes in computed spectra. This is for the 66% of the cases when the resonance is manifested.

V. CONCLUSIONS

Drives are often used to power equipment important to the plant process. They have access to and themselves generate large quantities of data related not only to their own operation, but also to the connected equipment. Though this data is normally used to support drive's controlling function, it can be used for diagnostic purposes. This information is available at no extra cost, without additional hardware investment. It can be obtained 24 hours a day, seven days a week. "Bad" data points may be detected and discarded. However, both the format and the resolution of this data are pre-determined by the existing hardware. Data that is sufficient for control purposes may need to be specially treated to be useful for condition monitoring and diagnostics. Problems, such as varying speed, drive-introduced noise and spectral aliasing require special processing techniques to ensure reliable results. In many cases the quality is rescued by the quantity, which allows for averaging of many data sets to improve the reliability.

This work has focused on some techniques for processing this data as well as on examples of useful information that can be thus obtained. These results validate the approach of using drive-supplied data as a supplemental source of diagnostic information about the drive itself, as well as about various aspects of the driven process.

VI. REFERENCES

[1] J. S. Mitchell, *Physical Asset Management Handbook*, Clarion Technical Publishers, USA, 2002, p.185

[2] M. Wnek, J. Nowak, M. Orkisz, M. Budyn, S. Legnani, "Efficient use of Process and Diagnostic Data for the Lifecycle Management", Proceedings of Euromaintenance and 3rd World Congress on Maintenance, Basel, Switzerland, 2006, pp. 73-78

[3] C.C. Martins Cunha, B.J. Cardoso Filho "Detection of Rotor Faults in Squirrel-Cage Induction Motors using Adjustable Speed Drives", Proceedings of 41st IAS Annual Meeting, 2006, pp. 2354-2359

[4] C. Concari, G. Franceschini, E. Lorenzani, C. Tassoni, A. Toscani, "Severity Assessment of Rotor Faults in Closed Loop Induction Drives by Different Approaches", in Proceedings of SDEMPED 2007, Cracow, pp 309-315

[5] M. Wnek, J. Nowak, M. Orkisz, B. Kosiba, and S. Legnani, "Practical Approach to Condition Monitoring of MV Drives", in Proceedings of SDEMPED 2007, Cracow, pp 486-490

[6] M. Wnek, M. Orkisz, C. Pinto "Utilizing the Variable Frequency Drive as a Condition Monitoring Tool in Rotating Machinery Diagnostics", Presented at ELROMA 2008, Mumbai, India, Jan. 2008

[7] M. Orkisz, M. Wnek, K. Kryczka, and P. Joerg, "Variable frequency drive as a source of condition monitoring data", in Proceedings of SPEEDAM 2008, Ischia, pp 179-183

[8] K.H. Lundberg, "Analog to digital converter testing", 2002, Available: http://web.mit.edu/klund/www/papers/UNP_A2Dtest.pdf

[9] W.H. Press, B.P. Flannery, S.A. Teukolsky, W.T. Vetterling, *Numerical Recipes: The Art of Scientific Computing*, Cambridge University Press, 1986

[10] C.E. Shannon, "Communication in the presence of noise", Proc. Institute of Radio Engineers, vol. 37, no.1, pp. 10-21, Jan. 1949. Reprint as classic paper in: Proc. IEEE, Vol. 86, No. 2, (Feb 1998)

VII. BIOGRAPHIES

Michał Orkisz was born in Kraków, Poland. He received B.S. degrees in Computer Science, Physics and Mathematics at M.I.T. in 1988, and Ph.D. in Condensed Matter Physics in 1994 at M.I.T. He has worked in Gel Sciences, Inc., Genome Therapeutics, and, since 1998, in ABB's Corporate Research, where he has dealt with condition monitoring, risk analysis and computer vision.

Maciej Wnek since 2009 is responsible for service products at ABB Medium Voltage Drives and sales of the drives upgrades. He received PhD in Solid State Physics from Jagiellonian University in Krakow in 1993. After working at Tohoku University, Japan and Universita La Sapienza di Roma in the area of LCDs he joined ABB Corporate Research center in Krakow in 1997. There for 10 years he was leading the research group active in the area of condition monitoring, diagnostics, and service.

Pieder Jörg is responsible for the drive systems business of ABB Medium Voltage Drives. He joined ABB in 1995, starting at Corporate Research in the area of power electronics. In 2002 he joined the business unit Medium Voltage Drives as head of product development. Since 2008 he is responsible for the drive systems business and its products within Medium Voltage Drives. Pieder Jörg received his M.Sc. degree in electrical engineering from the Swiss Federal Institute of Technology.

Klaus Ruetten since 2008 is responsible for the support of diagnosis functions inside the ACS control software. Since he joined ABB in 1991 he dealt with embedded control software for traction converters and variable speed drives and was responsible for the ACS 1000 control software since 2004. He received his graduate of electrical engineering from the 'Bergische' University in Wuppertal, Germany.

Edgar Jellum is working for ABB Oil & Gas in Norway as a principal scientist. He joined ABB in 1994 starting at Corporate Research in the area of process control and optimization. Edgar Jellum received his M. SC. degree in 1979 in electrical engineering from the Norwegian Institute of Technology in Trondheim.