In the European Journal of Control (1998) 4:294-297.

Discussion on: 'Advanced Motion Control: An Industrial Perspective' by M. Steinbuch and M. L. Norg

Daniel Y. Abramovitch
Hewlett-Packard Laboratories
Storage Technologies Department
1501 Page Mill Road, M/S 4U-12
Palo Alto, CA 94304 USA
Phone: (650) 857-3806
E-mail: danny@hpl.hp.com

This paper gives an excellent quasi-survey of area of control for storage systems and how advanced control techniques can be applied. It is worth elaborating on a few of the points made in the paper from the perspective of this author's own limited experience.

The first point is that a complete design using the only modern analysis and design techniques is a rare case in the storage industry. This is partly because of the difficulty (mentioned in the paper) of posing the industrial problem into the modern frameworks. However, a compounding factor is that of time constraints on design time. Any design method which requires the controls engineer to use different techniques in every aspect of the drive design is bound to take far more time to implement than making marginal improvements on a current design method.[1] Thus, controls engineers – sensing that their performance evaluations are on the line – are very cautious about deviating from practices that have worked before.

It has been this author's experience that the researcher stands a much better chance of having a new idea implemented into a storage system if it can be dropped as a piece into an existing design methodology. A perfect example of this is in the area of adaptive feedforward cancellation schemes [5, 11] which were introduced as add on schemes [12, 7] for eliminating repetitive disturbances from rotating machinery (such as a disk drive). Because an adaptive feedforward controller could be *added in* to an existing design method, it was relatively easy for such a scheme to be adopted into storage systems. Today, adaptive feedforward cancellation of the repetitive runout is a fairly standard feature of disk drive designs. The success of creating a drop in addition to current control designs has motivated this author's research to follow those lines (see [1]).

Another major issue was stated quite eloquently by Dr. Babatunde Ogunnaike, during his plenary talk at the 1998 American Controls Conference in Philadelphia, PA. Dr. Ogunnaike was speaking about the chemical process control industry, but much of his talk could apply to any industrial control setting. He said that building effective control systems in the chemical process industry was 90% process understanding and 10% control design. These ratios are certainly true in the data storage industry. While modern control methods put much effort into the creation of nice algorithms for control design, industrial control problems are dominated by an understanding of the particular physical problem. Achieving a good understanding of the particular control problem is justifiably where most the effort is spent. Thus, a new design methodology which requires substantially more time to understand and implement stands little chance of being used when time is tight. Furthermore, the improvements that can be made by fundamentally understanding the problem and then applying simple solutions typically swamp out the improvements achieved by an improved control design.

An example from this author's work [4, 6, 2, 3] and that McAllister [10, 9, 8] is the generation of air flow by the spinning magnetic media. As the disks spin in a disk drive, they carry along air. This in turn generates the *air bearing* that the magnetic heads ride upon. The air flow also generates a process disturbance as it impacts upon the magnetic read/write head and the actuator arm. This effect gets worse as the spindle speed increases and, in fact, the air flow can become turbulent. This is analogous to an airplane design problem. A design that is appropriate for an airplane with a top speed of 300 km/hour would fail dramatically for an airplane

---

[1] The rule of thumb is that for each drive design, the engineer gets to make at most one major change. If a new method requires many changes in order to get any results, it is unlikely to be tried.

with a top speed of Mach 2. Likewise, drive designs that were acceptable for lower spindle speeds run into problems for the higher speed spindles.[2]

The method of improved control design would identify the characteristics of the process disturbance and attempt to design a controller which rejected disturbances in this frequency range. However, by understanding the physical problem it becomes apparent that an aerodynamic redesign of the interior of the disk drive and of the actuator arm achieves such a large reduction in the turbulent air flow that the conventional control designs can be used. Thus, this is where the effort is spent. Returning to the aircraft design example, there is precious little that one can do with the autopilot on a Boing 747 to make the plane travel safely at Mach 2.

It is a rather humbling experience but an important lesson to learn. Unfortunately, it is not the only one. As mentioned in the paper, there are severe limitations on the number of calculations that can be done by a control algorithm between sample points. First of all, the sample rate for most hard disk drives is between 4 and 12 kHz. For optical disk drives the sample rates can range between 20 and 50 kHz. At such high speeds, and given the cost constraints of storage systems, the DSPs that one can afford to put into a disk drive have only a few hundred clock cycles between sample points. One might believe that this still leaves enough operations to calculate some fairly sophisticated control algorithms, but a further restriction is present in that the DSP spends most of its time doing something other than the control law calculation. A number quoted to this author from an engineer at Hewlett-Packard's former Disk Memory Division was that the DSP spends 95% of its time doing something other than the control law calculation. Such non-glamorous tasks as calibration, normalization of inputs, presaturation of outputs, etc. consume most of the code space and clock cycles. For a new, sophisticated algorithm to be used in a disk drive requires that it either: be compact enough to replace the simple control law *or* provide such a large improvement in performance to justify a more expensive DSP. When stated this way, the problem of using modern control design methodologies becomes more than a technical issue.

Another issue in the use of modern tools is that of modelling. Most design tools require explicit models of the system – often in state space form, while most drive models are based on instrument measurements of time and frequency responses. The issue then becomes one of transferring the instrument measurement into the modelling package in a form that is usable by the modern tools. This translation is not trivial and requires a different set of low level programming functions for each new instrument. This again requires extra effort, especially when compared to using a classical control design directly from a Bode plot measured off of a drive.

One final issue not mentioned in the paper was the role that *lack of communication* plays in limiting the use of modern control methodologies. As mentioned earlier, a tremendous performance improvement can be achieved by understanding the underlying system and how it affects the control system. The converse is also true. Understanding of control systems can point to the key areas where the system design must be improved. Thus, the actual control design problem should be an iterative process involving not only the control engineers, but the mechanical engineers, the material scientists, and the device physicists as well. In other words, good industrial control design cannot be separated from good system design and good system design requires a lot of interaction between a lot of different disciplines. When this interaction happens, performance improvements happen largely because a problem which may be hard to fix in one domain is easily eliminated in another domain. Once again this is the issue of process understanding dominating the control design. Unfortunately, the reality of the many industrial situations is that this cross disciplinary communication is far too rare. Thus, engineers in one area finish their part of the problem and "throw it over the wall" to the next team. This leaves precious little time for a new control scheme.

All this is not to say that there is no place for modern control techniques in storage system control problems. Over time, many modern control techniques are slowly creeping into the designs. However, they come in one piece at a time and only when conventional methods are deemed inadequate because that is the only way to justify the effort to a budget and time conscious manager. This author and many others in the field would laud anyone who could show us how to *easily* "drop in" modern control design into all of our control systems for storage devices.

References

[1] Daniel Abramovitch. Customizable coherent servo demodulation for disk drives. *IEEE/ASME Transactions on Mechatronics*, 3(3):184–193, September 1998.

[2] Daniel Abramovitch, Terril Hurst, and Dick Henze. Decomposition of baseline noise sources in hard disk position error signals using the PES Pareto Method. In *Proceedings of the 1997 American Control Conference*, pages 2901–2905, Albuquerque, NM, June 1997. AACC, IEEE.

[3] Daniel Abramovitch, Terril Hurst, and Dick Henze. An overview of the PES Pareto Method for decom-

---

[2] The increase in spindle speed is driven by customer desires for faster access times to support such new applications as streaming video from a disk drive. Thus, there is no avoiding the issue.

posing baseline noise sources in hard disk position error signals. In *Digests of The Magnetic Recording Conference 1997*, Minneapolis, MN, September 1997. IEEE Magnetics Society, IEEE.

[4] Daniel Abramovitch, Terril Hurst, and Dick Henze. The PES Pareto Method: Uncovering the strata of position error signals in disk drives. In *Proceedings of the 1997 American Control Conference*, pages 2888–2895, Albuquerque, NM, June 1997. AACC, IEEE.

[5] Marc Bodson, Alexi Sacks, and P. Khosla. Harmonic generation in adaptive feedforward cancellation schemes. *IEEE Transactions on Automatic Control*, 39(9):1939–1944, September 1994.

[6] Terril Hurst, Daniel Abramovitch, and Dick Henze. Measurements for the PES Pareto Method of identifying contributors to disk drive servo system errors. In *Proceedings of the 1997 American Control Conference*, pages 2896–2900, Albuquerque, NM, June 1997. AACC, IEEE.

[7] Carl Kempf, William Messner, Masayoshi Tomizuka, and Roberto Horowitz. Comparison of four discrete-time repetitive control algorithms. *IEEE Control Systems Magazine*, 13(6):48–54, December 1993.

[8] Jeffrey S. McAllister. Characterization of disk vibrations on aluminum and alternate substrates. *IEEE Transactions on Magnetics*, 33(1):968, May 1996.

[9] Jeffrey S. McAllister. The effect of disk platter resonances on track misregistration in 3.5 inch disk drives. *IEEE Transactions on Magnetics*, 32(3):1762–1766, May 1996.

[10] Jeffrey S. McAllister. Disk flutter: Causes and potential cures. *Data Storage*, 4(6):29–34, May/June 1997.

[11] Alexi Sacks, Marc Bodson, and William Messner. Advanced methods for repeatable runout compensation (disc drives). In *Digests of The Magnetic Recording Conference 1994*, Pittsburg, PA, August 1994. IEEE.

[12] E. Tung, G. Anwar, and M. Tomizuka. Low velocity friction compensation and feedforward solution based on repetitive control. In *Proceedings of the 1991 American Control Conference*, pages 2615–2620, Boston, MA, June 1991. AACC, IEEE.