Active Safety Control for Automobiles

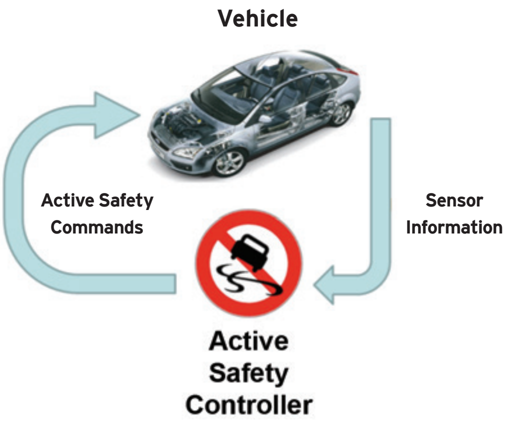

The rapid evolution of technology over the last 20 years has made automobiles much safer than ever before. Active safety is a relatively young branch of the automobile industry whose primary goal is avoiding accidents and at the same time facilitating better vehicle controllability and stability, especially in emergency situations.

The driver + vehicle + environment form a closed-loop system, with the driver providing control actions by manipulating three primary actuators: the steering wheel and the brake and accelerator pedals. In certain cases, as a result of environmental or vehicle conditions, or the driver's actions, the car may end up in an unsafe state, with the driver's ability to control the vehicle curtailed. Active systems correct such situations by automatically applying differential braking and cutting engine torque (and in the near future, correction of wheel turn).

Some Active Safety Control Mechanisms

- Antilock braking systems (ABSs; available today)
- Traction control (TC; available today)
- Electronic stability control (ESC; available today)
- Automatic steering correction (future)

Antilock Braking

The first active electronic safety system was the anti-skid Sure-Brake system proposed by Chrysler and Bendix in 1971; a previous all-mechanical system was introduced by Dunlop in 1950 for aircraft. The first production use was in 1978 when Bosch mounted an ABS on trucks and the Mercedes-Benz S-Class.

The main objectives of ABS are to minimize stopping distance under braking and to avoid wheel locking to maintain the drivability of the vehicle. Since wheel locking occurs when the slip ratio between road and tire (that is, the normalized difference between the peripheral velocity of the tire and the longitudinal velocity of its hub) exceeds a maximum value, the ABS tries to avoid this situation.

As depicted in the figure below, the driver, through the brake pedal, imposes a certain pressure in the hydraulic system. The inlet and outlet valves initially work for normal braking, that is, open and closed, respectively (the opposite of the situation in the figure); in this case, the brake fluid (in the red branch) pushes the caliper into the braking disk. If this braking action determines a slip ratio on the wheel close to the maximal slip ratio, the control strategy changes the state of valves by closing the red branch and opening the blue one so that the pressure on the caliper decreases (and hence the slip ratio). The inversion of fluid flow causes a "feedback" vibration at the pedal. Notice that the opening/closing actions of the hydraulic system are cyclical (a form of high-frequency switching control) such that the slip ratio is kept close to its maximal value. The principal manufacturers are Bosch, Delphi, Continental Teves, and Kelsey-Hayes, which formed a group in 2000 called the ABS Education Alliance and estimated that almost 28% of the accidents on wet roads are caused by uncontrolled braking.

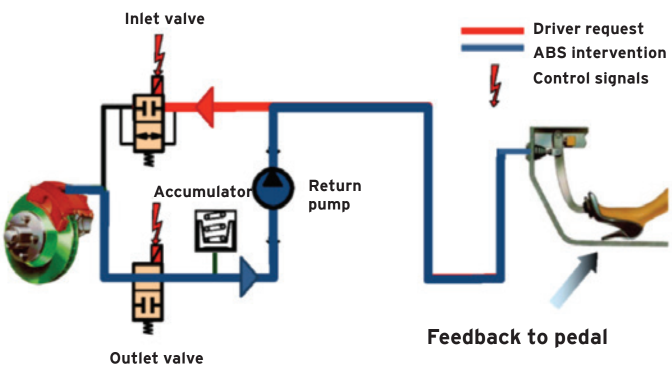

ABS in operation: the automatic release phase when the inlet/outlet valves are in closed and open status, respectively

Contributor: Luigi Glielmo, Università del Sannio, Italy

Cost-benefit analyses of these systems for EU-25 show that in the decade 2010-2020, the use of ESC can return benefits (in terms of accident avoidance) of €2.8-4.4 for each euro spent. This has convinced governments to make the installation of ESC systems on all cars in the European Union and the U.S. mandatory from 2012. Mandates for ABS are also in effect or in progress, including for motorcycles.

Traction Control

Traction control (TC) systems (or anti-slip regulators) have the opposite goal of ABS in that they try to keep the wheels from spinning in acceleration. This is done by maintaining the slip ratio (opposite in sign with respect to the braking situation) within a certain threshold, modulating the traction torque on the wheels. TC is available in two different versions: one, produced by Saab in collaboration with Teves and Hella, uses the braking system and engine torque variation; the other one, produced by Honda and Bosch, uses only the engine torque variation.

Electronic Stability Control (ESC)

ABS only works well during longitudinal panic braking and TC in start-up maneuvers; neither is effective when vehicle stabilization involves lateral dynamics (sideslips). ESC systems fulfill this need. They act on individual brakes and possibly engine torque, based on measurements or estimated errors of two vehicle variables and their respective (computed online) reference signals: the yaw velocity (the angular velocity around the vertical axis) and the sideslip angle (the angle between the longitudinal axis of the vehicle and the direction of the velocity vector). In particular, the yaw velocity must track a reference trajectory computed on the basis of the steering wheel angle and the vehicle velocity, and the sideslip angle must not exceed a certain threshold. The whole control action (estimation + actuator command generation) is performed in a strict sampling time (10-20 ms). Human drivers would not be able to simultaneously coordinate braking of four individual wheels and cutting of engine torque (if longitudinal velocity is too high) so as to correct the vehicle direction.

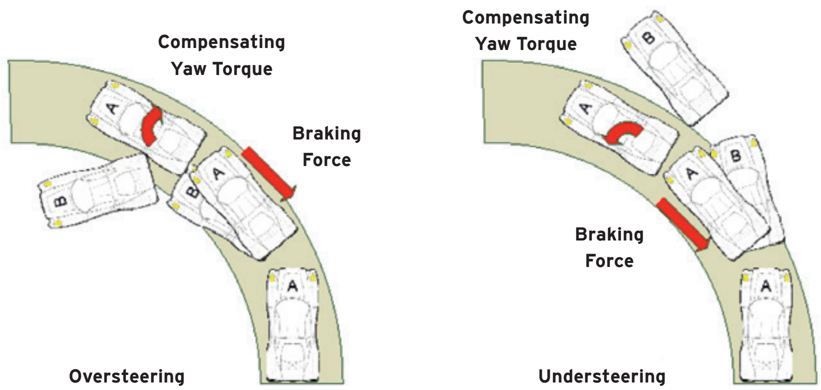

Correcting oversteer and understeer with ESC

The first commercial ESC was developed between 1987 and 1992 by Mercedes-Benz and Robert Bosch GmbH. Today ESCs are available under trade names such as AdvanceTrac, Dynamic Stability Control (DSC), Dynamic Stability and Traction Control (DSTC), Electronic Stability Program (ESP), Vehicle Dynamic Control (VDC), Vehicle Stability Assist (VSA), Vehicle Stability Control (VSC), Vehicle Skid Control (VSC), Vehicle Stability Enhancement (VSE), StabiliTrak, and Porsche Stability Management (PSM). These products differ in the combination of actuators used and the conditions for activating the control strategy.

The Future: Advanced Model-Based Control

Active safety control systems are typically designed using gain-scheduled single-input, single-output controllers whose calibration is obtained after extensive real-time simulations and tests on the track. Furthermore, the coordination among multiple subsystems is done through heuristic rules that determine activation conditions and manage shared resources. Limitations of this approach are that new actuators or sensors are difficult to integrate and it cannot take into account from the beginning the multivariable and constrained nonlinear nature of the global problem. Hence, research is under way to introduce more complex model-based and robust control design methods, exploiting the increased computational power available on board.

Advanced Control Design for Automotive Powertrains

The automotive industry currently spends approximately $1 billion each year in the development and calibration of powertrain control. According to one survey, the number of lines of code in a vehicle is increasing by a factor of 10 every eight years, and the development cost for software will exceed that of hardware before 2020.

Modern automotive control problems are outpacing the design techniques that have been traditionally used to solve them. The legislated and consumer demands of reduced fuel consumption and emissions have driven increased complexity into modern powertrains in terms of sensors, actuators, and new subsystems. Engineering, calibration, and test cell time and cost are all increasing. At the same time, the optimal coordination of subsystems remains elusive.

The OnRAMP Design Suite provides a systematic framework for end-to-end model-based powertrain control design. Controllers can be designed and deployed in weeks instead of months, as was previously the case.

Problem Characteristics

Automotive powertrains comprise many different configurations—single-stage and multistage turbocharging, high- and low-pressure exhaust gas recirculation, variations on existence and arrangement of throttles and valves, and various exhaust aftertreatment options and configurations—depending in part on the intended application and target market for the vehicle.

From a control perspective, any one of these configurations represents a highly nonlinear, multivariable plant with significant model uncertainty due to both manufacturing variability within a production line and in-service aging. At the same time, any designed control must satisfy constraints on both actuator and output (e.g., turbocharger speed or air-fuel ratio) while hosted on an embedded hardware platform with limited memory and processor power (less than 60 MHz and only a few megabytes of flash memory) at fast sample times typically measured in milliseconds.

Current State of the Art

The current process for automotive control design is highly manual and labor intensive. It typically involves many months of experimental work in an engine test cell and vehicle to tune or "calibrate" standard production controllers over all conditions (engine speed, load, ambient temperature, pressure, etc.) they may encounter in practice. The finalization of the control and diagnostics can easily take two to three years as the vehicle passes through more than one season of testing. Systematic model-based control design techniques are still relatively rare in production automotive applications.

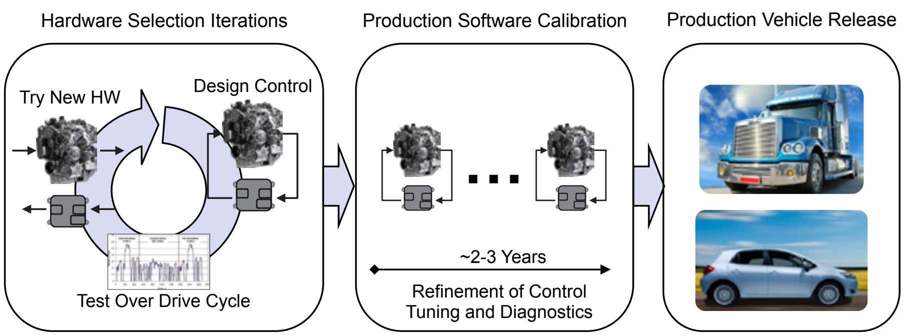

Every step of the powertrain development process involves control design. Today's process typically requires a few years from requirements specification to engine certification.

Contributor: Gregory Stewart, Honeywell, Canada

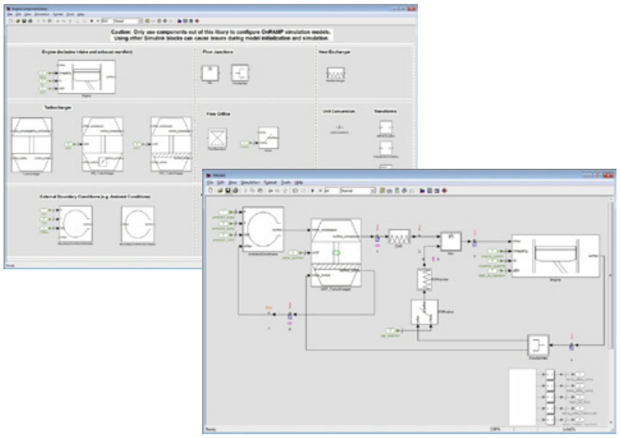

The Innovation: OnRAMP Design Suite for Powertrain Modeling and Control Design

OnRAMP supports end-to-end powertrain control design. The user is guided through three phases—modeling, control design, and controller deployment—each of which is supported by software tools based on a control-theoretical foundation. Multivariable control over a transient drive cycle can typically be achieved in two to three weeks, with the current record being four days for a new engine.

Modeling

A fully identified model of the engine or aftertreatment device may be obtained in less than a week. The modeling uses a physical-component-based library to produce the ordinary differential equations (ODEs) suitable for control design while avoiding the complex wiring challenges in ODE configuration. A second key innovation has been the automatic and robust identification of nonlinear models useful for engine applications. A hierarchical identification strategy first fits individual model components, and then a systemwide nonlinear optimization of all model parameters simultaneously over all of the recorded data—including all inputs, outputs, and operating points—is performed.

Control Design and Tuning

The control approach automatically generates both the feedforward and feedback control required by powertrain applications. As constraints are of key importance, the feedback uses a version of explicit model predictive control (MPC) that fits within the processor and memory limitations of modern ECUs.

Tuning is intuitive with a slider bar for speeding up or slowing down a given output or actuator action. MPC weights are then computed such that the resulting (linearized) closed-loop transfer functions will satisfy a small-gain theorem condition for robust stability. Tuning is thus user friendly without risk of generating unstable control in the face of real-world model uncertainty.

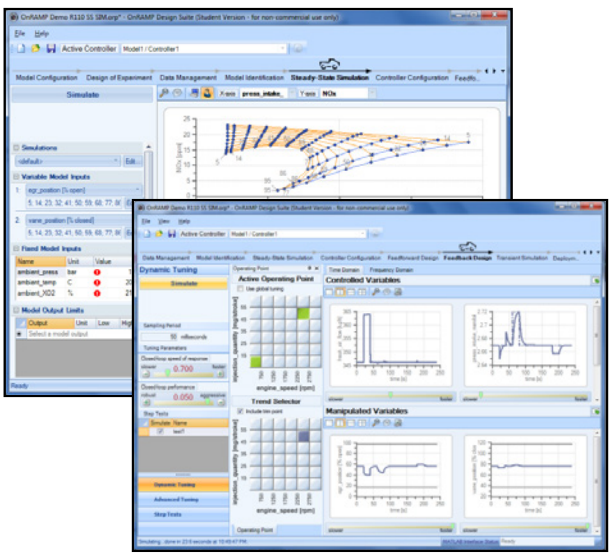

The OnRAMP Design Suite was released as a product in late 2011. Some 35 users have been trained, and the technology has been applied for several applications by engine manufacturers. The clean-sheet development time to achieve transient control is reduced in most cases from several months to a few weeks.

Award

"IEEE Control Systems Technology Award" from the IEEE Control Systems Society *"for the design, implementation and commercialization of the* ***OnRAMP Design Suite for Powertrain Control***," awarded to Francesco Borrelli, David Germann, Dejan Kihas, Jaroslav Pekar, Daniel Pachner, and Greg Stewart in 2012.

For more information: G.E. Stewart et al., Toward a systematic design for turbocharged engine control, in L. Del Re et al. (eds.), Automotive Model Predictive Control, Lecture Notes in Control and Information Sciences, vol. 402, pp. 211-230, Springer, 2010; R. Beňo, D. Pachner, and V. Havlena, Robust numerical approach to mean-value modeling of internal combustion engines, Proc. 7th IFAC Symposium on Advances in Automotive Control, Tokyo, September 4-7, 2013; www.honeywellonramp.com.

Advanced Control for the Cement Industry

The cement industry of the 21st century is confronted with disparate goals that at first glance seem to conflict. For example, the enormous pressure to increase profit and margins is juxtaposed against the considerable public interest in the sustainable and environmentally friendly use of natural resources. In other words, plant operators find themselves in a situation where they need to react fast and optimally to continuously changing conditions while still meeting various, and probably conflicting, objectives. Thus, there is a need for tools that bring the plants to their optimal economic performance allowed by technological, environmental, and contractual constraints. From a technological standpoint, these tools are related to mathematical programming: optimization subject to constraints. The cpmPlus Expert Optimizer (EO) was developed to address these challenges, in particular for cement plants.

Solution Overview

Over many years, a variety of strategies for control and optimization of key industrial processes have been developed and implemented in EO, with particular focus on control and optimization in the cement industry:

- Raw materials blending
- Vertical mills for raw meal grinding
- Calciners and rotary kilns
- Vertical and ball mills for cement grinding

The technology has been deployed in cement plants worldwide. Most installations have been made in blending, kiln, and grinding operations. More than 45 blending systems, 195 rotary kilns, and 90 ball mills have been commissioned by the ABB team in recent years.

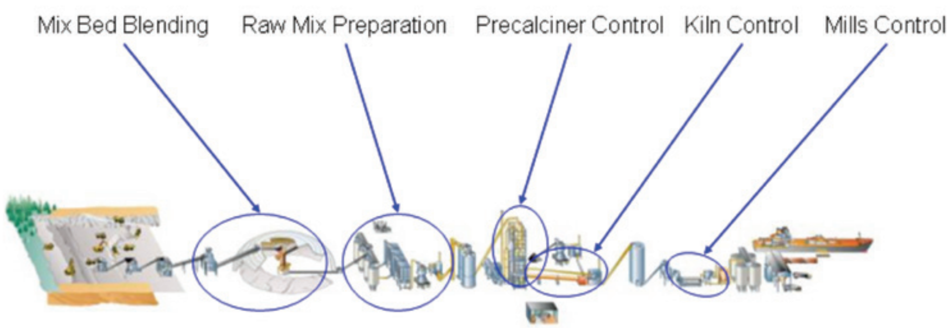

cpmPlus Expert Optimizer applications scope in the cement industry

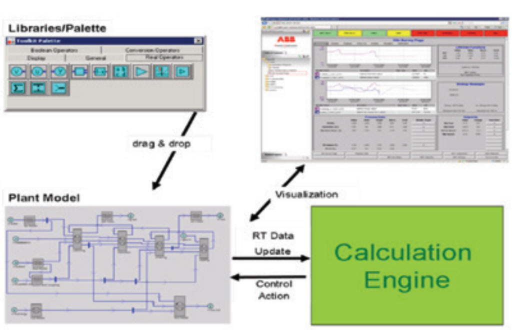

$$x(t+1) = Ax(t) + B_1u(t) + B_2\delta(t) + B_3z(t)$$
$$y(t) = Cx(t) + D_1u(t) + D_2\delta(t) + D_3z(t)$$
$$E_2\delta(t) + E_3z(t) \le E_1u(t) + E_4x(t) + E_5$$

Energy Efficiency and CO_2 Reduction

The cpmPlus Expert Optimizer is a generic platform for development of advanced process control solutions at ABB. It is primarily designed for closed-loop control, optimization, and scheduling of industrial processes, although it can also be used for open-loop decision support applications. When this platform is used, the problems described above can be attacked with techniques such as model predictive control (MPC) in its mixed logical dynamical (MLD) systems formulation, which includes Boolean variables and logical constraints.

For ease of use, the technology has been embedded in a graphical modeling toolkit that allows maximal flexibility during model and cost function design while hiding the mathematical complexity from the user.

Contributors: Eduardo Gallestey and Michael Stalder, ABB, Switzerland

Global Fuels Award, 2008

In 2008, the cpmPlus Expert Optimizer received the "Global Fuels Award for most innovative technology leading to electrical energy savings." The award was granted by the Global Fuels 2008 conference in London.

Selected Success Stories

Switzerland: Material Blending at Untervaz

The Untervaz plant wanted to reduce raw mix quality variability, reduce the associated material costs, and increase the useful lifetime of the quarry. This would also allow the plant to have better process parameters in the kiln, getting closer to clinker quality targets, increasing production, and reducing the risk of process disruptions. In March 2007, ABB extended Untervaz's Expert Optimizer to include ABB's Raw Mix Preparation (RMP) solution. The technologies used are MPC and MLD systems. The benefits achieved by the installation are that raw mix quality variability has been reduced by 20% and kiln process variability has also been reduced. New daily clinker production records have been achieved in the time since RMP has been online.

Germany: Precalciner With Alternative Fuels at Lägerdorf

The Lägerdorf plant wanted to increase alternative fuels utilization, get closer to optimal calcination conditions, and reduce the risk of process disruption. In August 2006, ABB successfully installed Expert Optimizer, encompassing a Precalciner Temperature (PCT) control solution, on the calciner at Lägerdorf. The technologies used are MPC and MLD. The installation achieved a dramatic increase in the use of alternative fuels. Furthermore, it was possible to reduce temperature variability, bring the precalciner average temperature toward optimal values, and reduce the risk of cyclone blockages.

Italy: Cement Grinding at Guidonia

Buzzi Unicem wanted a solution for its Guidonia plant that would increase the productivity of its cement grinding system, consisting of three mills. ABB installed Expert Optimizer on the mills at the Guidonia plant between December 2006 and January 2007. The EO team overcame the challenges at the Guidonia plant by applying the MPC approach together with a tailor-made parameter adaptation and process supervision procedure. The benefits are better grinding process parameters and operation closer to process constraints. The specific energy consumption was reduced by as much as 5%.

Turkey: Full Process Optimization at Adana

Adana Cement Industry Inc. operates four separate clinker production lines, two coal mills, and five cement mills at its Adana plant. Over a period of a few months in 2010 the company deployed Expert Optimizer to several kilns, mills, coolers, and calciners. The strategies were based on model predictive control and achieved runtime factors of over 90% while reaching the targeted improvements in energy efficiency and production.

Advanced Control of Pharmaceutical Crystallization

Nearly all pharmaceutical manufacturing processes use crystallization as the primary means for providing high purity, efficacy, and potency. Process modeling, monitoring, automation, and control systems are now widely used for the design and development of pharmaceutical crystallizers.

Modern control system technologies have reduced the time required to develop robust, scalable, and reliable crystallization processes; enabled the robust isolation of metastable and stable crystal forms of active pharmaceutical ingredients; and enabled the removal or simplification of post-crystallization processing—with associated increases in productivity, product quality, and product consistency.

The technologies have produced substantial technical and economic benefits.

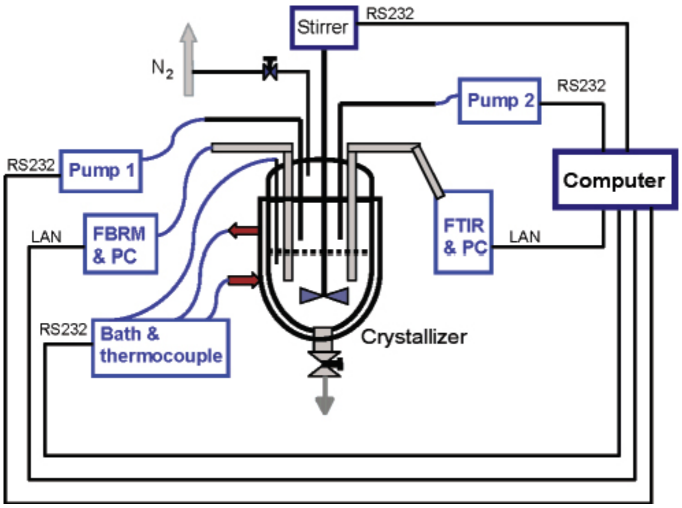

Schematic of an apparatus and instrumentation setup for a pharmaceutical crystallizer

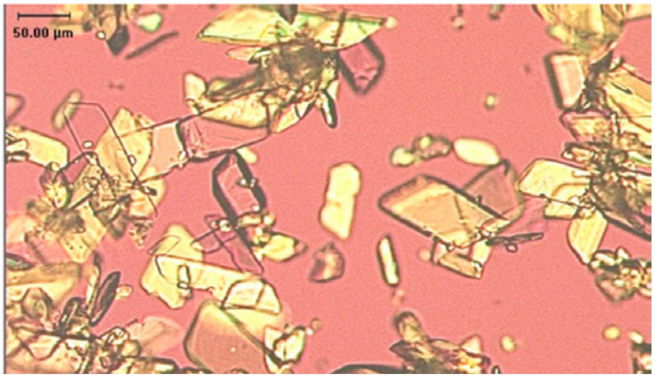

Microscope images of highly pure crystals produced by an automated process monitoring and control system (Source: G. Zhou et al., Evolution and application of an automated platform for the development of crystallization processes, Organic Process Research and Development, vol. 17, pp. 1320–1329, 2013; Copyright 2013 American Chemical Society, reprinted with permission)

Successful Applications Worldwide

Advanced process monitoring and control system technologies have been implemented in pharmaceutical crystallizations in many companies, including:

- AstraZeneca, United Kingdom
- AbbVie, United States
- Bristol Myers-Squibb, United States
- Merck & Co., United States and United Kingdom
- Novartis Pharma AG, Switzerland
- Sanofi-Aventis Deutschland GmbH, Germany
- Syngenta, Münchwillen, Switzerland

Implementations of Control

Control systems technologies that have been implemented in the pharmaceutical industry include:

- An automated procedure that designs nearly optimal batch control policies for crystallization processes. The states of the liquid solution and the crystal size distribution are monitored, and nonlinear state feedback control provides low sensitivity to disturbances.
- Monitoring techniques based on multivariate statistics that are applied to experimental data collected from attenuated total reflection – Fourier transform infrared (ATR-FTIR) spectroscopy to achieve highly accurate *in situ* solution concentration estimates in dense crystal slurries.
- A feedback control system that is provably robust to the large variations in the crystallization kinetics for cooling, solvent addition, and combined operations. Feedback control enables the production of large high-purity crystals, even with varying contaminants in the feed solutions and deviations in the seeding and temperature and solvent addition profiles.

Contributor: Richard D. Braatz, Massachusetts Institute of Technology, USA

Innovations

Advanced process monitoring and control systems have enabled many advances in productivity:

- The combination of multivariate statistical methods and ATR-FTIR spectroscopy provide *in situ* solution concentration estimates that are an order of magnitude more accurate than prior technologies.
- An automated software system that implemented advanced process monitoring and control was transferred from academia to Merck & Co., where it reduced process development times by more than an order of magnitude.
- A key challenge in pharmaceutical crystallization is to determine process operations that reliably produce the desired crystal structure, known as the polymorph. An undesired polymorph can have very different bioavailability and effects on the body than the desired polymorph. Automated process monitoring and control enables users to select which polymorph to produce in the crystallization. The specified polymorph, either stable or metastable, is reliably and repeatedly produced for a wide variety of organic compounds and solvents.
- In a typical application at Merck & Co., the technology produced crystallizer operations that ensured the robust isolation of the thermodynamically most stable crystal form of an active pharmaceutical ingredient. The process was robust, scalable, and reliable and enabled the removal of post-crystallization product milling.

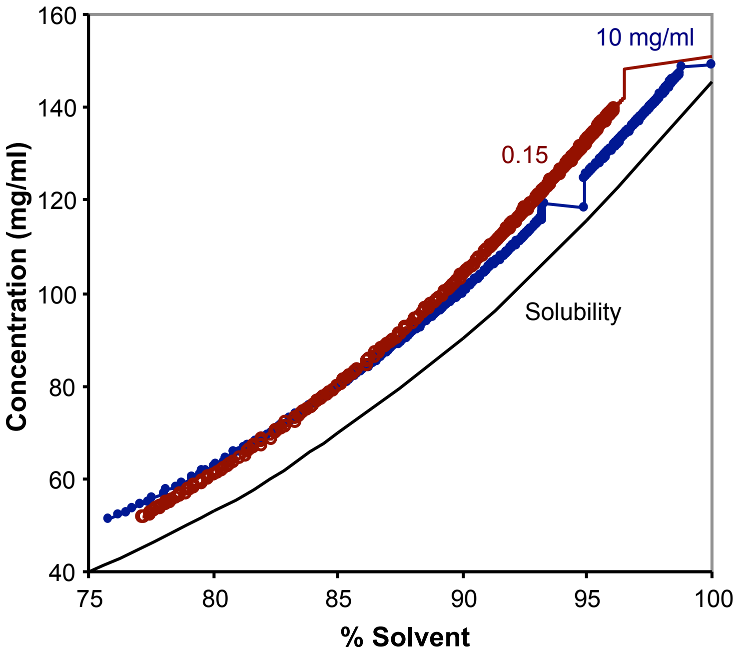

Two paths in the crystallization phase diagram followed by using an advanced process control system during a crystallization at Merck & Co. of a pharmaceutical compound from a mixture of solvents. The sharp deviation in the path with a setpoint of 10 mg/ml around 120 mg/ml was caused by an extremely large external disturbance in the pump flow rate that was introduced while temporarily turning off the control systems. When the pump was returned to automatic control mode, the robust nonlinear feedback control system quickly returned to the desired path in the phase diagram.

Advanced monitoring and control technologies in pharmaceutical crystallization have resulted in order-of-magnitude or greater improvements in chemical concentration estimates and process development times!

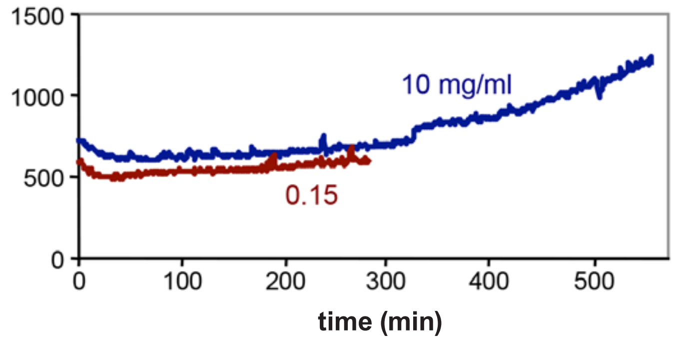

Measurement of number density in a crystallizer shows some crystal nucleation for controlled operations with a setpoint of 10 mg/ml, whereas a setpoint of 0.15 results in negligible nucleation. The minimal nucleation enables the production of large uniform crystals in a much shorter batch time of about 280 min.

For more information: G.X. Zhou et al., Direct design of pharmaceutical antisolvent crystallization through concentration control, Crystal Growth & Design, vol. 6, pp. 892–898, 2006; A. Cote, G. Zhou, and M. Stanik, A novel crystallization methodology to ensure isolation of the most stable crystal form, Org. Process Res. Dev., vol. 13, pp. 1276–1283, 2009.

Advanced Energy Solutions for Power Plants

Fuel costs, energy conversion efficiencies, and environmental impacts of fossil-fueled plants have become priorities in both developed and developing countries. Advanced Energy Solutions (AES), a product of Honeywell Process Solutions, is an advanced process control product that significantly improves power plant efficiency and reduces plant emissions.

AES provides combustion control in boilers; coordinates multiple boilers, turbines, and heat recovery systems for optimal operation of entire power plants; and provides dynamic balancing of power production to demand.

The AES solution is effective for both fossil-fueled power plants and industrial steam plants and has been used for applications covering boilers, steam/gas turbines, and heat recovery steam generators in Europe, Africa, and Asia.

Successful Applications Worldwide

AES and its component technologies have been implemented in plants worldwide, including the following:

- Co-generation plant Otrokovice, Czech Republic
- ECG Kladno, Czech Republic
- Samsung Fine Chemicals, Korea
- Nam JeJu power plant, Korea
- Sinopec JinShan power plant, China
- SASOL steam plant, Secunda, South Africa
- REPSOL steam plant, La Coruna, Spain

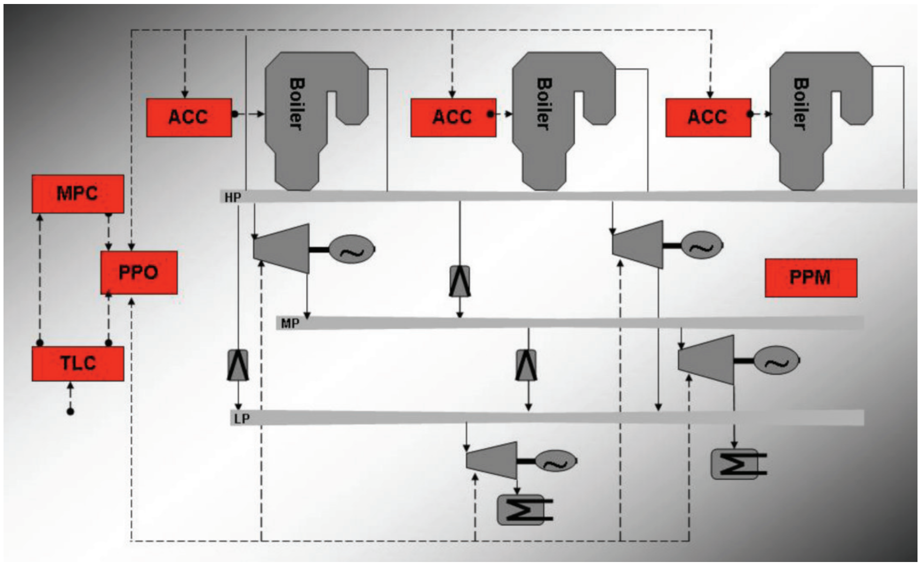

Solution Overview

AES is a software-based product that can be implemented as a hierarchical application layer on baseline distributed control systems (DCSs). Several modules are available:

Advanced Combustion Controller (ACC) optimizes air distribution and tightly coordinates control of fuel and air ratio for advanced control of the combustion process.

Plant Performance Optimizer (PPO) increases the efficiency and reliability of the power plant by optimizing the utilization of steam for electricity generation and process or heating needs. PPO also analyzes key performance indicators (KPIs) so business objectives can be achieved.

Master Pressure Controller (MPC) stabilizes steam pressure and prevents boiler and turbine outages using advanced predictive control algorithms. It continuously balances produced and consumed steam and increases asset life by minimizing wear.

Tie-Line Controller (TLC) is a power quota planning and real-time execution toolkit for management of energy supply and demand.

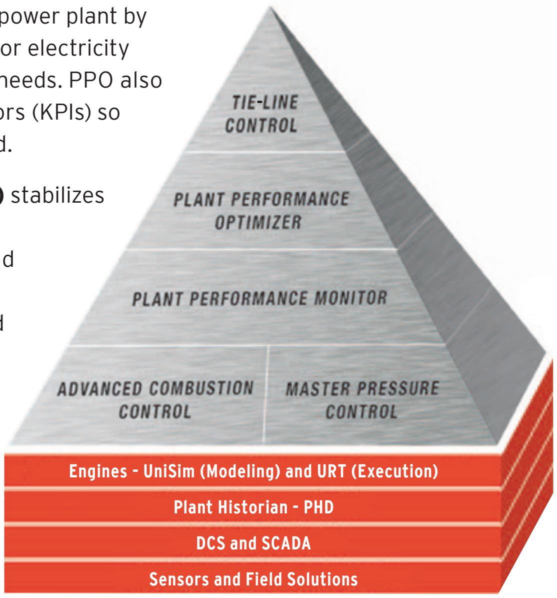

Contributor: Vladimir Havlena, Honeywell, Czech Republic

"As the first company in the world to apply advanced control application technology to CFB units, Sinopec significantly enhanced the effectiveness and control performance of the distributed control system at the CFB boiler level and for the entire plant. Even more impressive, all improvements were achieved by implementing software rather than executing a major hardware refurbishment at the plant. We have also to date achieved an estimated $1 million of savings on the supply of energy to our refinery."

— Zhao Weijie, Chief Engineer, Sinopec Shanghai Petrochemical Company (2008)

Inventions and Innovations

Advanced Energy Solutions incorporates innovative concepts to improve energy efficiency, reduce emissions, and improve the economic operation of industrial and utility fossil-fueled power plants:

- Dynamic coordination of the air-fuel ratio (AFR) in the boiler reduces the AFR variation and enables combustion optimization. An extension of linear model predictive control technology for ratio control was developed.
- Turbulence during combustion results in emissions being highly stochastic. Deterministic optimization methods were unable to provide satisfactory performance. AES's "cautious optimization" strategy takes uncertainty into account.
- One of the key challenges for coal-fired power plants is the variability in the BTU content of the coal. With advanced estimation and inferential sensing technology, leaking air variation and coal quality variation are identified and combustion parameters are optimized online.
- The solution has been extended for circulating fluidized bed (CFB) boilers. CFB boiler dynamics depend significantly on the accumulated char in the bed. An inferential bed fuel inventory (BFI) sensor was developed to estimate the accumulated char level and adapt the model used for predictive control accordingly.
- Another innovation is the plantwide optimization of boilers, turbines, and heat recovery systems to improve the end-to-end efficiency of a power plant.

Most Innovative Power Technology of the Year Award from *Asian Power* magazine, 2008

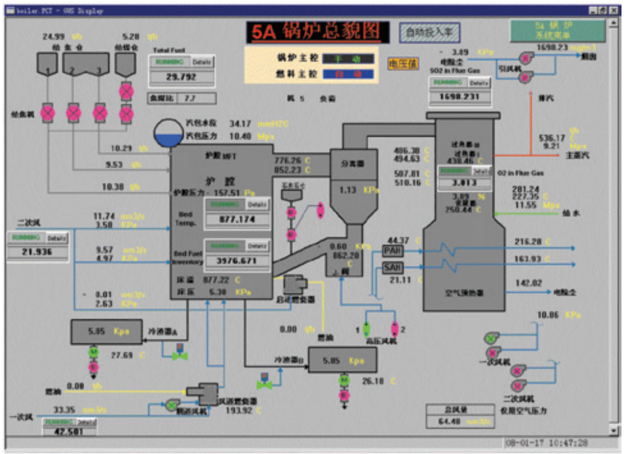

For the application of AES to Sinopec's Shanghai Petrochemical Company Principal Power Plant in Shanghai, Honeywell received the 2008 Most Innovative Power Technology of the Year Award from Asian Power, the leading publication for energy professionals in Asia.

For more information: V. Havlena and J. Findejs, Application of model predictive control to advanced combustion control, Control Engineering Practice, vol. 13, pp. 671-680, 2005.

Advanced Zinc Coating Control in Galvanizing Lines

Hot dip galvanizing lines (HDGLs) are an industrial process where cost savings in terms of increased production, more stringent tolerances on final product quality, and reduced raw materials utilization are especially needed. One aim of an HDGL is to apply a uniform protective zinc coating to the surface of steel coils, which are then used, for instance, in the construction industry for making rustproof parts.

After being dipped in a molten zinc bath, metal strips are blown on with special air knives (devices that produce bladelike jets of compressed air) to remove excess zinc. As zinc prices continue to soar, closed-loop control is needed to keep the coating mass to a constant, minimum value.

Controlling this process presents many challenges, as it is nonlinear and multivariable in nature and some key variables cannot be measured directly. In addition, coating thickness is conventionally measured after a considerable delay, when the strip has cooled down.

The solution proposed by Danieli Automation solves this problem with a multivariable controller built around an adaptive model of the process, leading to more stringent tolerances and reduced zinc utilization.

Inventions and Innovations

- Innovative use of a multivariable feedback controller for coating mass in hot dip galvanizing lines
- Online identification of crucial, unmeasured inputs to the nonlinear model of the air knives that is part of the controller

Galvanized steel coils ready to be shipped

Enabling Technologies

- Key to coating control using air knives is the availability of reliable estimations of variables whose measurements are delayed or are not available.
- The final thickness of the coating is conventionally measured about 100 m downstream, so a Smith predictor is built on top of a model of the air knife process.
- The distance between the air knives and the strip surface cannot be measured and must be estimated.
- The model of the air knife process is online and recursively tuned.

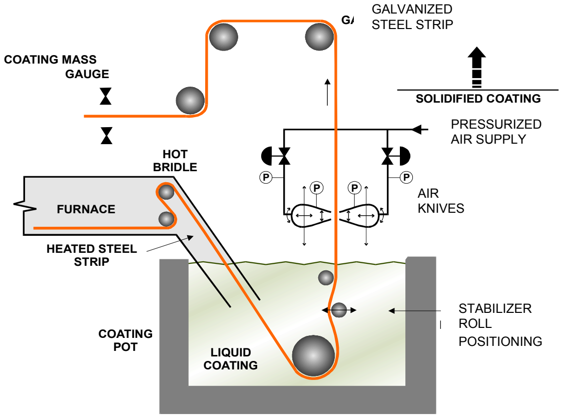

Schematic view of an HDGL coating section with air knife actuators and a coating mass sensor

***Contributors:** Francesco Cuzzola, Danieli Automation, Italy; Riccardo Ferrari, Danieli Automation, Italy; and Thomas Parisini, Imperial College London, U.K.*

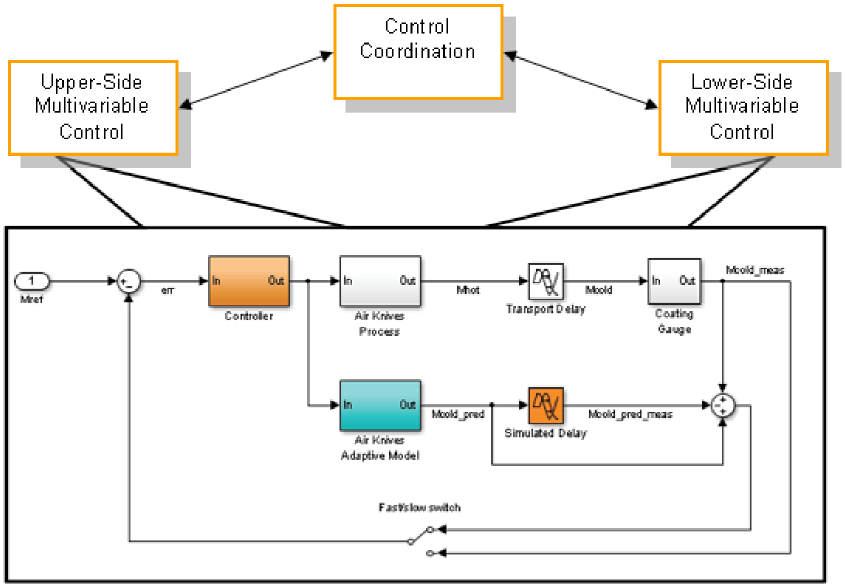

The coating control architecture. The air knife model is subject to online fine-tuning. The orange blocks are implemented in the regulatory control level, the cyan block in the supervisory control level. Only one side of the strip is represented here for simplicity.

Control Architecture

- A multivariable Smith predictor is implemented for each strip side with a controller, an air knife adaptive model, and a simulated delay implemented in the regulatory control level.
- Another synchronizer controller is in charge of coordinating the two multivariable controllers corresponding to each strip side.
- An adaptive, nonlinear model of the air knives is recursively tuned online by the supervisory control level.
- To control the cold coating mass M_{ref}, the multivariable controller acts separately on the air knives' pressure and on their distance from the strip (which is not conventionally measured).

Realized Benefits

- The Advanced Coating Control architecture is regularly commissioned on Danieli mills.
- Use of the section control leads to a dampening of coating mass fluctuations and a
  - 9% increase in average material processing speed
  - 15% reduction in coating weight (about 1.3 kg of zinc per ton of material)
  - 0.45–1.5% reduction in costs
- Considering a plant producing 350 kTon per year and a zinc price of 1700 €/kg, this leads to savings of 760 k€ per year.

Conclusion

The advanced thickness control solution described here leads to

- Reduction in overcoating, which translates to reduced production costs
- No need for extra hardware
- Robustness to uncertainties in air knife actuators

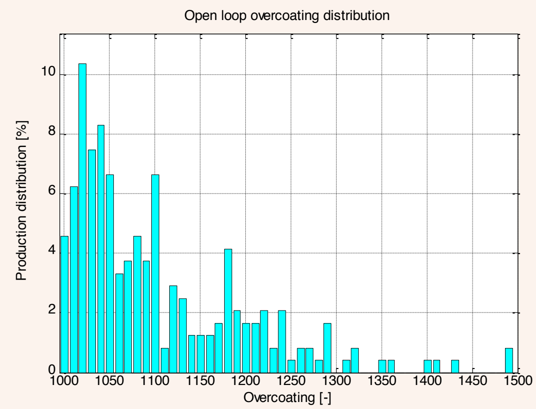

Typical distribution of coils overcoating (target = 1000, meaning no excess coating is present) when no closed-loop coating control is applied

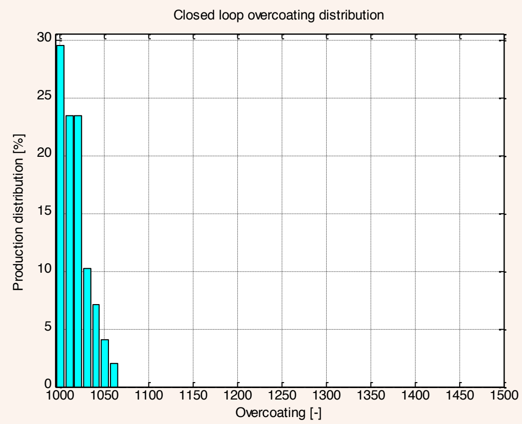

Distribution of coil overcoating when closed-loop coating control is present. As can be seen, most of the coils are very close to the desired target (1000).

For more information, visit http://www.dca.it/Libraries/Papers/Danieli_Automation_Advanced_Multivariable_Control_scheme.sflb.ashx.

Automated Manual Transmissions

Shift buttons on the steering wheel of a FIAT Bravo (Source: www.fiat.it)

The automated manual transmission (AMT) is an intermediate technological solution between the manual transmission used in Europe and Latin America and the automated transmission popular in North America, Australia, and parts of Asia. The driver, instead of using a gear stick and clutch pedal to shift gears, presses an upshift or downshift button and the system automatically disengages the clutch pedal, shifts the gear, and engages the clutch again while modulating the throttle; the driver can also choose a fully automated mode. AMT is an add-on solution on classical manual transmission systems, with control technology helping to guarantee performance and ease of use.

AMT Overview

An AMT is composed of a dry clutch, a gearbox, and an embedded dedicated control system that uses electronic sensors, processors, and actuators to actuate gear shifts on the driver's command. This eliminates the need for a clutch pedal while still allowing the driver to decide when to change gears. The clutch itself is actuated by electronic equipment that can synchronize the timing and the torque required to make gear shifts quick and smooth. The system is designed to provide a better driving experience, especially in cities where congestion frequently causes stop-and-go traffic patterns.

AMTs have been used in racing cars for many years, but only recently have they become feasible for use in everyday vehicles with their more stringent requirements for reliability, cost, and ease of use.

Benefits of AMT

- Changing gears without using a foot to operate the clutch
- No engine or gear modifications
- Less physical or psychological stress
- More comfortable than manual transmissions
- More "fun" factor compared to fully automatic transmissions

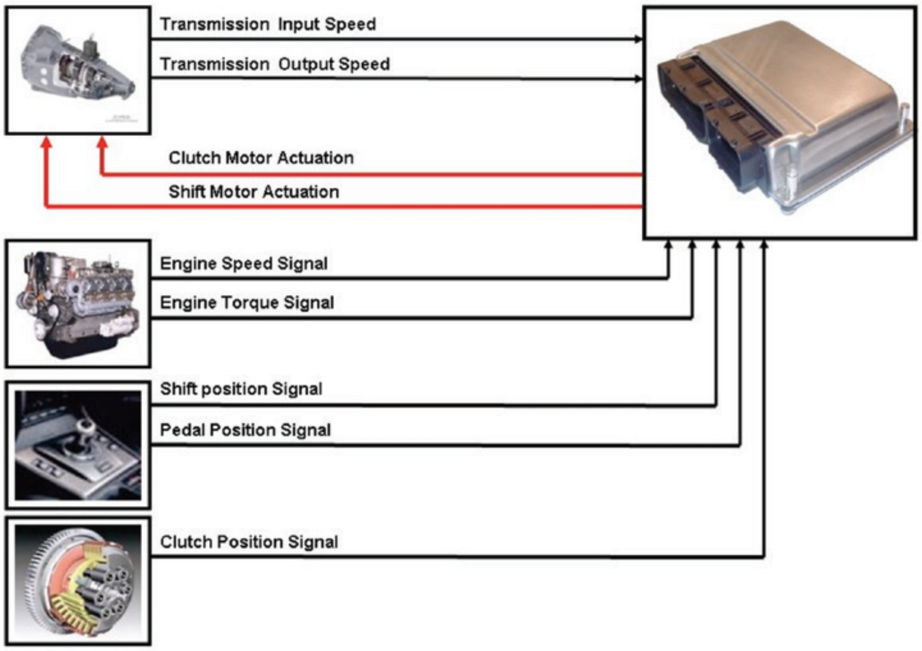

Inputs and outputs for a typical AMT system (Source: www.itri.org.tw)

Contributor: Luigi Iannelli, Università del Sannio, Italy

AMT systems are currently installed by several automakers under different commercial names, such as SeleSpeed by FIAT, Sequential Manual Gearbox by BMW, 2Tronic by Peugeot, SensoDrive by Citroen, and EasyTronic by Opel.

Commercial dual-clutch transmission (DCT) systems include the Direct-Shift Gearbox by Volkswagen Group and the Dual Dry Clutch Transmission by FIAT Group.

Inventions and Innovations

AMT is an interesting example showing the importance and potential of automatic control. The control of the clutch engagement on AMT systems must satisfy different and conflicting objectives:

- It should result in the same or better shifting times as with manual transmissions.
- It should improve performance in terms of emissions and facing wear.

In a typical AMT control scheme, a constant engine speed is requested during the engagement so as to equalize engine and clutch torques as well as possible. In this case, the clutch control provides a clutch torque reference, and through a suitable model (or map), the torque reference is converted into a position reference for the clutch actuator position control (see figure below).

Commercial implementations of AMT today rely on enhancements of PID controllers with feedforward actions and controller gain scheduling.

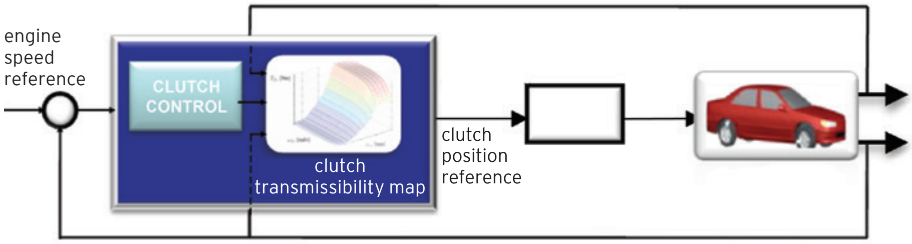

Future View: Toward Model-Based Control of AMTs

Model-based approaches are attracting increasing interest as evidenced by several control strategies that have recently been proposed in the literature. These strategies are based on optimal control, predictive control, decoupling control, and robust control.

Innovative AMT technology uses a dual-clutch transmission consisting of one clutch for odd gears and another for even gears. The goal is to improve the speed and comfort of the gear shift. But effective AMT controllers, particularly for dual-clutch systems, are difficult to design without an accurate model of the clutch torque transmissibility characteristic, or the relationship between the clutch actuator position (or the pressure applied by the clutch actuator) and the torque transmitted through the clutch during the engagement phase.

The clutch transmissibility model, key to advanced control of AMTs, is difficult to attain: it depends on various parameters and phenomena, such as friction pad geometries, cushion spring compression and load, and slip-speed-and-temperature-dependent friction. Accurate clutch transmissibility models will allow the use of advanced model-based control strategies aimed at improving the overall behavior of the system with respect to current commercial solutions.

Automotive Engine-Based Traction Control

The availability of inexpensive embedded controllers and sensors has enabled an ever-increasing number of advanced functions for automotive safety, driver comfort, and convenience. After microprocessor-based engine controls were introduced, anti-lock brakes (ABS) provided the first instance of embedded controls for chassis and vehicle dynamics. Shortly thereafter, traction control (TC), the tractive analog to ABS for vehicle acceleration, appeared. Continuous improvement of traction systems led to Ford Motor Company's in-house development of engine-only traction control. This system provides the majority of the safety-related TC function with improved refinement at a greater value for the vehicle buyer.

The Traction Challenge

Vehicle motion is caused almost entirely by forces generated at the tire/road interfaces, and these forces provide the greatest nonlinearity and uncertainty in vehicle dynamics control problems.

The figure at right depicts typical tire forces generated as a function of longitudinal tire slip for a given steering angle and road surface friction. Longitudinal tire force initially increases with slip (the difference between the driven and nondriven wheel speed) but falls off at high slip. Lateral force capability decreases with slip. In the high slip range, accelerating and steering both diminish. The purpose of a traction control system is to manage vehicle acceleration and steerability by targeting an appropriate driven wheel slip, based on driver control inputs and vehicle state, and meeting this target using closed-loop control through powertrain torque modulation.

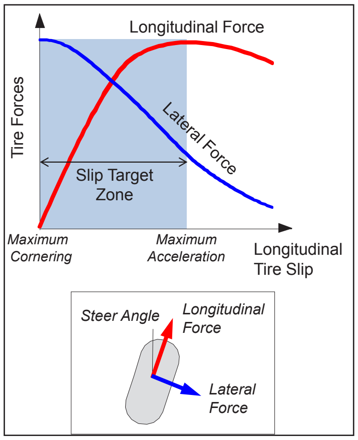

Tire characteristics

The following figure depicts the generalized TC structure with powertrain torque as the control actuation. The wheel speed (slip) target is based on vehicle speed, steering wheel angle, accelerator pedal position, and estimated road surface friction. When the vehicle is cornering, low longitudinal tire slips are targeted to produce the necessary higher lateral tire forces. When large accelerator pedal positions are present and the driver's steering input is low, larger slip levels are targeted to produce the larger longitudinal forces required for better acceleration.

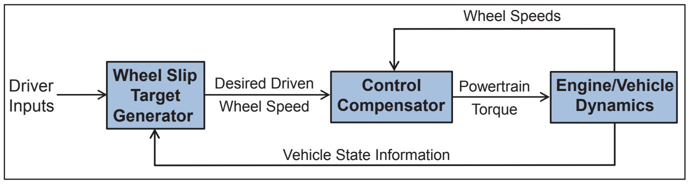

Traction control system structure

Traction control design is challenging due to the torque production dynamics of the powertrain, which exhibits sizable variation in transport lag depending on engine speed and torque level. Further complicating the problem is the variability of the tire force/slip relationship that is strongly dependent on the road surface condition (ice, snow, gravel, etc.). This effect can lead to a locally unstable plant. Powertrain output torque is modulated using electronic throttle control, spark advance, cylinder cutting, cylinder air/fuel ratio, and transmission shifting. Each of these actuation methods operates with its own bandwidth, limited authority, and in some cases transport delay.

Development Process

Control system development in an industrial setting begins with appropriate control-oriented plant modeling and controller design, including stability analysis. For automotive applications, great care is taken to guarantee closed-loop robustness across a wide range of operating conditions. This is followed by work to ensure that the resulting control design is compact, computationally efficient, and fail-safe.

The initial control analysis for the Ford Traction Control system modeled the powertrain as a lumped rotating mass with significant transport delay to account for the intake-to-power delay in the engine. A state feedback controller with optimal target tracking and actuator use was constructed with linear quadratic (LQ) design. The resulting control structure was then realized in the form of a classical proportional-integral structure with a cascaded lead filter (dubbed "PI+") tuned through gain scheduling to address the varying engine transport delay. This form preserves all the elements of the LQ design while fitting into the familiar classical PID-type control structure and

(continued on next page)

Contributors: Davor Hrovat and Michael Fodor, Ford Motor Company, USA

"Traction control on the V-6 test car was just right—perhaps unique in all the industry. It allowed tire spin when starting forcefully on slick roads and gradually eased the spinning without trying to stop it, allowing the car to keep moving forward as traction was gained. It should be unusually effective in winter and whenever some spinning helps forward progress."

—James R. Healey, *USA Today*,
2006 Ford Fusion Review,
"Fusion Charges Off the Blocks,"
October 28, 2005

Development Process (continued)

providing tuning based primarily on the desired closed-loop bandwidth for the system.

In addition to the core control law, the controller includes a mu-estimator for surface friction detection and corresponding PI+ control gain adaptation, a wheel speed target generator, and a state machine for control initiation and error handling. The figure at right above shows a typical TC event initiated by full accelerator pedal application on snow with the accompanying mu-estimator response.

Computer code for implementation was written to minimize memory footprint, stack size, and computation time (chronometrics). Furthermore, the entire system design was scrutinized using the Boundary Diagram, P-Diagram, a full failure mode and effects analysis, and thorough in-field testing.

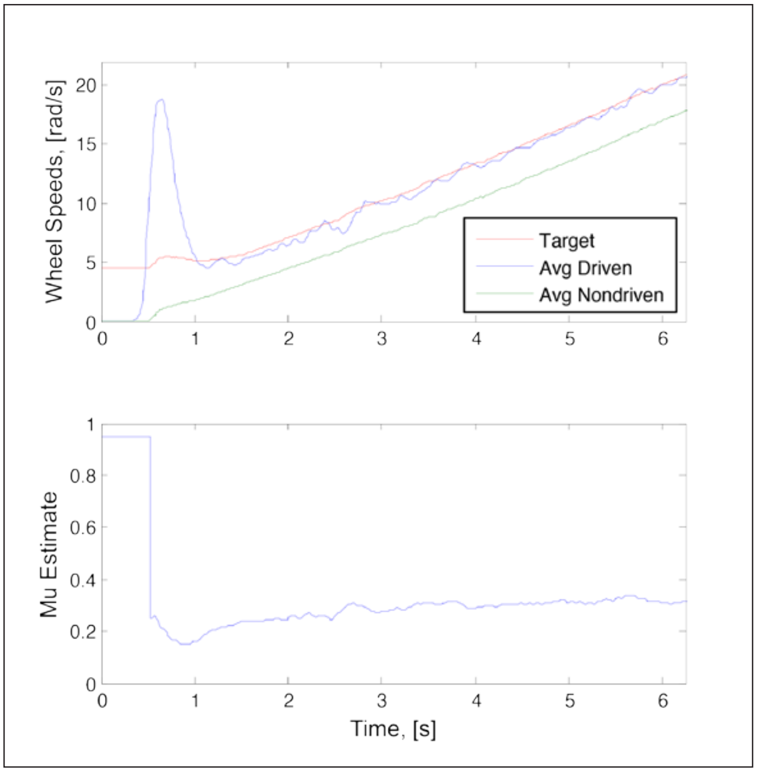

Typical traction control operation on snow—angular velocities of driven and nondriven wheels and mu-estimator (friction detection) response

Benchmarking and Market Reception

To provide assurance of its performance capability, the previously described controller was benchmarked against a hybrid model predictive controller (MPC) designed using the same plant model. The MPC system was implemented using the "explicit" form, which prestores all possible control actions in a searchable collection of piecewise affine control regions. The PI+ controller compared favorably, losing approximately 10% target tracking performance compared to the benchmark MPC system while employing a significantly smaller control structure and tuning that is familiar to present calibrators.

For straight-line acceleration, data collected for several drivers driving with and without TC show that with TC an inexperienced, perhaps less agile, driver will perform at least as well if not better than test drivers experienced at driving on slippery surfaces without TC. Similar improvements in vehicle steerability and stability are also evident with traction control engaged.

The Ford Engine-Only Traction system was introduced in 2006 on the Ford Fusion and the F150 light truck. It has been well received in the market as evidenced by favorable reviews (see the sidebar) and high customer take rates where it is offered as an option. It has since been introduced on several other Ford and Lincoln vehicles, including Ford's large and commercial truck lines.

Awards: Best Paper of Automotive Track, 1998 Digital Avionics Systems Conference, Seattle, WA; 2004 Henry Ford Technology Award, Ford Motor Company, Dearborn, MI.

For more information: D. Hrovat, J. Asgari, and M. Fodor, Automotive mechatronic systems, in Mechatronic Systems, Techniques and Applications: Vol. 2—Transportation and Vehicle Systems, C.T. Leondes, Ed., Gordon and Breach Science Publishers, 2000.

Autopilot for Small Unmanned Aerial Vehicles

Small unmanned aerial vehicles (UAVs) have numerous applications in civilian sectors. These include terrain monitoring for agriculture, assessment of damage caused by natural or manmade hazards, archaeological discovery, and exploration of remote and inaccessible areas.

In all these applications, stable, controllable flight of the vehicle is essential. The UAV must be able to follow a commanded trajectory and maintain its attitude so as to ensure high-quality sensor data. The small size and low weight of the craft make it susceptible even to low-level wind disturbances, plus small UAVs exhibit significant nonlinearities in their dynamics—presenting challenges for flight control and autopilot design.

Autopilots

The controllable flight of aircraft requires the precise manipulation of aerodynamic surfaces such as elevators, ailerons, and rudders. Pilots do not affect these actuators directly. Instead, a flight computer, the "autopilot," translates higher-level commands (e.g., heading and altitude changes) into appropriate commands to the surfaces.

In many cases, and especially for small UAVs, automated and systematic approaches for autopilot design are lacking. Autopilots are based on simple single-variable PID controllers. Extensive manual tuning is required for adequate performance.

New "robust control" techniques have been developed that automate much of the autopilot design process for UAVs and also allow accurate flight in a substantially broader range of environmental conditions.

Robust Design of Multiloop Autopilots

- Robust control theory suggests a rigorous, holistic approach to designing flight controllers. The interactions of different sensors and actuators are directly and elegantly handled by these mathematical techniques.
- The protracted manual and heuristic trial-and-error process of tuning PID controllers is replaced by a model- and tool-based framework. Autopilots can be designed in hours instead of weeks, and UAV performance is significantly improved!
- Several sources of variation related to manufacturing and operation must be addressed. These include wind gusts, changes in payload for different missions, and the lack of repeatability of low-volume manufacturing.
- A probabilistic robust controller has been developed that takes into account model and environmental parametric uncertainties for small UAVs.
- The controller is implemented as an open-source autopilot that can be reprogrammed in flight if required.

The MicroHawk UAV family has been developed at Politecnico di Torino (Italy) to promote innovative scientific techniques for Antarctic exploration, as well as archaeological and other applications of societal interest within project ITHACA (Information Technology for Humanitarian Assistance and Cooperation Actions—in cooperation with the UN World Food Program). Robust multivariable autopilots have been designed and implemented for the MicroHawks, with demonstrated improvements in both controller design time and UAV performance.

Contributors: Elisa Capello, Giorgio Guglieri, Fulvia Quagliotti, and Roberto Tempo, Politecnico di Torino and CNR-IEIIT, Italy

MicroHawk flights have been completed in urban and rural areas. Missions have been flown successfully in various weather conditions.

UAV operators can provide high-level commands such as waypoints to which the vehicle should fly and customized trajectories the vehicle should take.

The adoption of the model-based framework also allows the development of operator training systems and simulators.

Archaeological site monitoring with a MicroHawk 2000. The photos were taken from an onboard camera during UAV flight. In the top photo, the aircraft's shadow can be seen over the Roman Amphitheatre of Bene Vagienna, Italy.

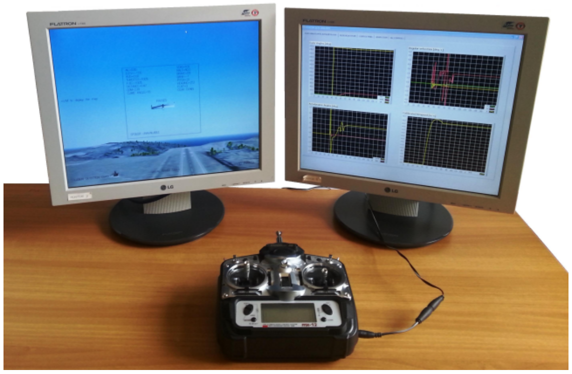

MicroHawk educational and training flight simulator

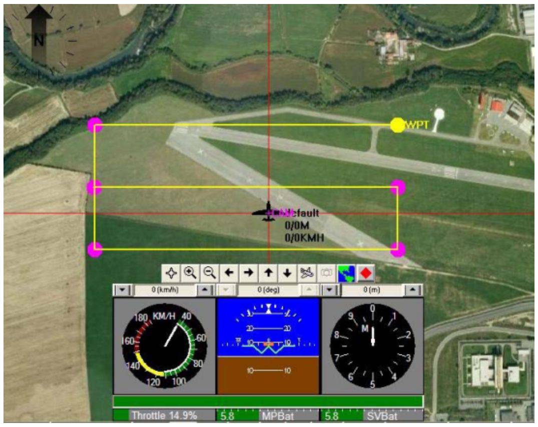

Waypoint assignment via the operator interface during a flight test at the airport Aeritalia, Torino, Italy

Auto-tuners for PID Controllers

Despite all the progress in advanced control, the PID remains the most popular controller. Any stable system can be controlled with an integrating controller; performance can be increased by adding proportional and derivative action. There is ample evidence that many manually tuned PID controllers do not work well. Automatic tuning has improved performance and simplified use.

PID controllers come in different varieties: as stand-alone components, as elements of distributed control systems, or embedded in instruments and systems.

PID control is used everywhere—in cellular phones, vehicles, process control, heating, ventilation, air conditioning, machine tools, and motor drives. Many PID controllers are found in cars, for example, in engine, cruise, and traction control. PID control is also embedded in instruments such as atomic force microscopes and adaptive optics. Because of their widespread use, it is difficult to precisely estimate the number of control loops installed each year, but an educated guess is that it is in the billions.

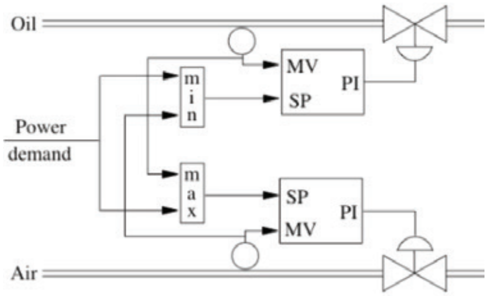

The PID controller is based on very simple ideas. As illustrated in the idealized formula below, the controller output is a combination of three terms:

- The proportional term reacts to current errors.
- Past errors are accounted for by the integral term.
- The derivative term anticipates future errors by linear extrapolation of the error.

$$u_{PID}(t) = k_p e(t) + k_i \int^t e(\tau)d\tau + k_d \frac{de(t)}{dt}$$

A remarkable property of a controller with integral action is that it gives the correct steady state, if a steady state exists, even for nonlinear processes.

Predicting a noisy signal by linear extrapolation is difficult; it is also difficult to find values of derivative gain k_d that give a robust system (tuning the derivative gain is more difficult than tuning the proportional and integral gains). Most PID controllers are in fact used as PI controllers.

A Real PID Controller

PID control is much more than what is captured by the simple idealized formula. To get a functioning controller, one must consider set-point weighting, filtering of the measured signal, protection for integral windup, as well as bumpless mode and parameter changes.

Complex System

The PID controller is a simple system. Well-developed architectures exist for building complex systems from the bottom up by combining PID controllers with linear and nonlinear elements such as cascade, mid-range, selector control, and gain scheduling. The figure on the left shows a system with PI controllers and selectors for controlling a burner that guarantees there will always be excess air.

Automatic Tuning

Traditionally, PID controllers were tuned manually using simple rules that date back to Ziegler and Nichols in the 1940s. The rules were based on process experiments. The step response method is based on measurement of the open-loop step response. The frequency response method is based on a closed-loop experiment where the system is brought to the stability boundary under proportional control. Unfortunately, the traditional rules resulted in systems with poor performance.

Automatic tuning has increased the use of derivative action. It has even been said: "This controller must have automatic tuning because it uses derivative action."

Automatic tuning can be done in many ways. In rule-based methods that mimic an experienced instrument engineer, features of the closed-loop response are calculated and controller parameters are adjusted based on empirical rules. Other methods are based on estimation of low-order process models, typically first-order dynamics with time delays. The controller parameters are then determined by a variety of control design methods.

Relay auto-tuning is another widely used approach that has proven to be robust and that brings attractive theoretical properties as well.

Contributors: Karl Johan Åström and Tore Hägglund, Lund University, Sweden

PID auto-tuners are in widespread use, especially in the process and manufacturing industries. All major instrumentation and control suppliers offer auto-tuning as a feature in their products. Auto-tuning software is also commercially available for PC, SCADA, and DCS platforms and in the simulation programs Simulink and LabView.

PID auto-tuners

Relay Auto-tuning

In relay auto-tuning, the process is first brought to oscillation by replacing the PID controller with a relay function (see figure below). The controller parameters are then determined from the period and the amplitude of the oscillation. An interesting feature of relay auto-tuning is that it automatically generates signals that are customized for modeling critical aspects of the process. The relay can also be applied to a closed-loop system.

For typical process control applications, the relay auto-tuners can be designed so that tuning can be executed simply by pushing a button; there is no need to set any parameters. The auto-tuner can also be used to generate gain schedules automatically.

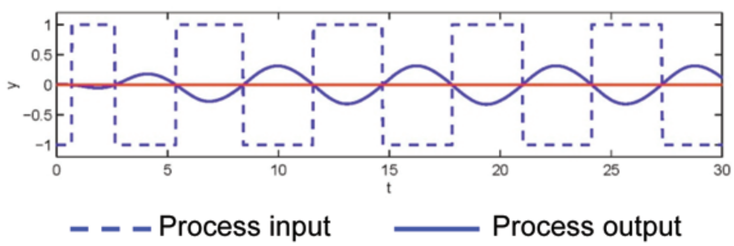

The data are from a recorder where time runs from right to left. A PI controller produced oscillations as seen in the top plot. The PI controller was switched to manual at time 11:15. The oscillation stops but the process drifts. An auto-tuner was installed and tuning was initiated at time 14:00 by pushing the tuning button; no further manual interaction was involved. Tuning is completed at time 20:00 and the controller switches to automatic with good control performance. The auto-tuner reduced the proportional gain, increased the integral time, and introduced derivative action with prediction time 1100 sec.

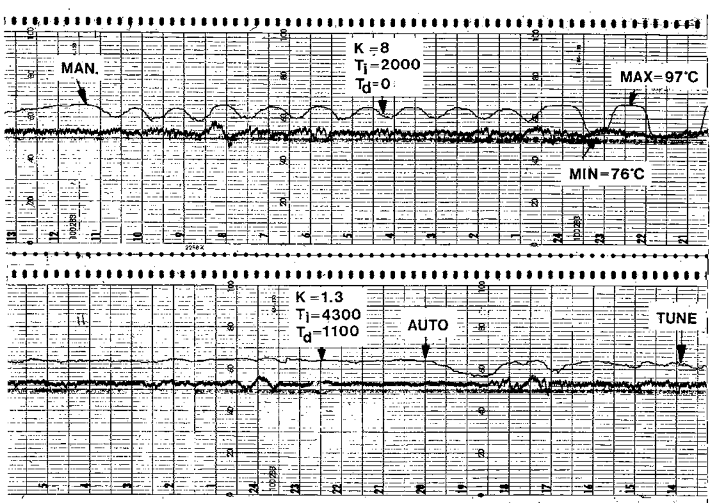

Relay auto-tuning of a temperature control loop on a distillation column

For more information: K.J. Åström and T. Hägglund, Advanced PID Control, ISA, Research Triangle Park, NC, 2004; T.L. Blevins et al., Advanced Control Unleashed: Plant Performance Management for Optimum Benefit, ISA, Research Triangle Park, NC, 2003.

Control for Formula One!

In August 2008, the deployment of a novel mechanical control device in Formula One racing was announced. Developed at the University of Cambridge by Malcolm Smith and colleagues, the device, called an "inerter," was deployed by the McLaren team in 2005 in Barcelona.

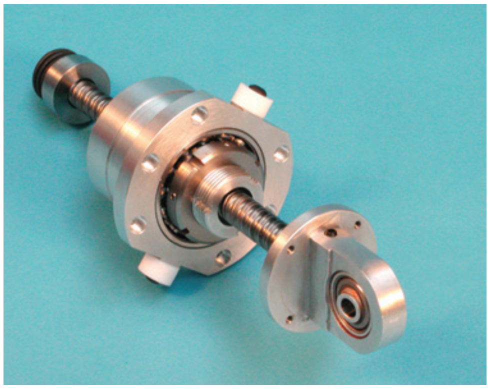

A ballscrew inerter (flywheel removed) made at Cambridge University, Department of Engineering, in 2003, designed by N.E. Houghton

What Is an Inerter?

The standard analogy between mechanical and electrical networks relates force to current and velocity to voltage. The following correspondences exist between standard modeling elements:

spring ↔ inductor　　damper ↔ resistor　　mass ↔ capacitor

The correspondence is perfect for the spring and damper, but the mass element is analogous to a grounded electrical capacitor and not to a general two-terminal capacitor. Without a two-terminal capacitor equivalent, mechanical systems are unable to provide the same flexibility in dynamic response that electrical systems can. The two-terminal electrical model suite above can be used to produce any "passive impedance" device.

The inerter overcomes this limitation of mechanical systems—this two-terminal element has the property that the applied force at the terminals is proportional to the relative acceleration between them.

Schematic of an inerter with a flywheel driven by a rack, pinion, and gears

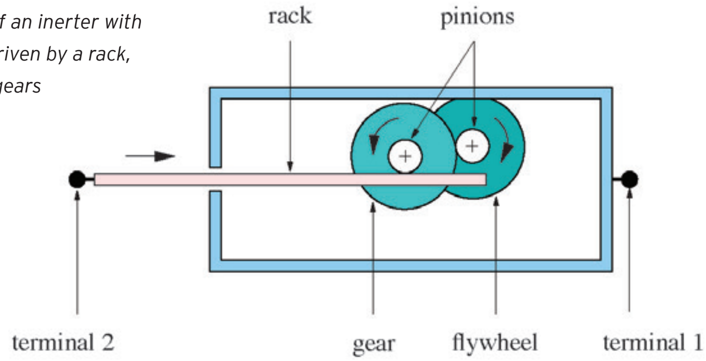

Kimi Raikkonen crosses the finish line to take victory for McLaren in the first car to race the inerter. (Photo courtesy of LAT Photographic)

The First Application: Vehicle Suspensions

Malcolm Smith's group at Cambridge University, in attempting to build high-performance mechanical impedances for car suspensions, realized that the lack of a true capacitor equivalent was a fundamental limitation.

After several fruitless efforts to prove that such a device could not exist, they realized it could be built, and in a relatively simple manner. They ultimately developed several prototypes of the device they called the inerter.

Contributor: Malcolm Smith, Cambridge University, UK

From the Laboratory to the Racetrack

Analyses of inerter-based suspensions indicated a potential performance advantage for vehicle suspensions that might be large enough to interest a Formula One team. Cambridge University filed a patent on the device and then approached McLaren in confidence. McLaren signed an agreement with the university for exclusive rights in Formula One for a limited period.

After a rapid development process, the inerter was raced for the first time at the 2005 Spanish Grand Prix by Kimi Raikkonen, who achieved a victory for McLaren.

Stolen Secrets . . . and the Truth Ultimately Comes Out

During development, McLaren invented a decoy name for the inerter (the "J-damper") to keep the technology secret from its competitors for as long as possible. The "J" has no meaning and was just a ruse, and of course the device is not a damper. The idea behind the decoy name was to make it difficult for personnel who might leave McLaren to join another Formula One team to transfer information about the device and in particular to make a connection with the technical literature on the inerter, which Malcolm Smith and his group were continuing to publish.

This strategy succeeded in spectacular fashion during the 2007 Formula One "spy scandal," when a drawing of the McLaren J-damper came into the hands of the Renault engineering team. The FIA World Motor Sport Council considered this matter at a hearing in December 2007. According to the council finding, "[a drawing of McLaren's so-called J-damper] was used by Renault to try to have the system that they thought McLaren was using declared illegal. This failed because Renault had certain fundamental misunderstandings about the operation of the J-damper system." A full transcript of the decision is available on the FIA website: *http://www.fia.com/mediacentre/Press_Releases/FIA_Sport/2007/December/071207-01.html*.

Neither the World Motor Sport Council nor McLaren made public what the J-damper was. Thereafter, speculation increased on Internet sites and blogs about the function and purpose of the device. Finally, the truth was discovered by Craig Scarborough, a motor sport correspondent from *Autosport* magazine. *Autosport* ran an article on May 29, 2008, which revealed the Cambridge connection and that the J-damper was an inerter.

Further Applications

With the truth out, and McLaren's exclusivity expired, Cambridge University entered a license agreement with Penske Racing Shocks USA, enabling Penske to supply inerters to any team in Formula One as well as in other domains of motor sport and elsewhere. The use of inerters in vehicle suspensions has continued to spread. In 2012, inerters were allowed in IndyCar racing for the first time. The Cambridge University research group is working with partners to develop other applications of the inerter. One particular focus is their use in railway vehicle suspensions, where improvements have been found in theory and simulation for ride quality and track wear.

Kimi Raikkonen leading the field in the McLaren-Mercedes MP4-20 at the Spanish Grand Prix, May 8, 2005, Circuit de Catalunya, Barcelona, Spain (photo courtesy of LAT Photographic)

For more information: M.C. Smith, Synthesis of mechanical networks: The inerter, IEEE Transactions on Automatic Control, vol. 47, no. 10, October 2002; http://www.admin.cam.ac.uk/news/dp/2008081906; http://www.eng.cam.ac.uk/news/stories/2008/McLaren.

Control in Mobile Phones

Mobile phones have made a huge impact on the world in a short time period. They are now affordable for those with daily incomes as low as a dollar, and they have brought communication infrastructure to new areas. In addition to enabling convenient and low-cost telephone services, mobile phones have also made information available at subscribers' fingertips. For many, their first contact with the Internet is with a mobile phone, not a computer.

Mobile phones as affordable and attractive consumer products would not be possible without control. Each phone has at least a half dozen function-critical control loops. Control is used to reduce cost, size, and power consumption to levels where mass-produced, battery-operated products are feasible.

Control has been embedded in mobile telephones since the first large, bulky, barely portable handsets and continues to be a key technology for today's smartphones (images not to scale).

With a world penetration of more than 4 billion users, the number of control loops in mobile phones is in the range of 10^{10} to 10^{11}. If you choose any control loop in the world at random, it is likely located in a mobile phone, making the application area one of the major success stories of control in recent times. The area is heavily patented, with thousands of new patents granted each year, a large share of them describing control inventions.

Access Control

Each phone contains a transceiver unit that makes radio access possible with one or several base stations. Designing a low-cost transceiver that is easy to mass produce and has sufficient power efficiency, receiver sensitivity, and linearity is a major technical challenge. Some of the control loops that have enabled transceiver design with the technology components available today are automatic gain control (AGC), automatic frequency control (AFC), transmission power control, timing control, and feedback control of coding and modulation.

Radio Unit

Application CPU
Access CPU

Radio Unit Clock Rate: ~2 GHz
Application and Access CPU: ~500 MHz
Memory: 512 MB RAM + 1 GB Flash
(data for high-end phones)

Circuit Design Level Control

Control loops are also heavily used on the electronic circuit design level, for example, in the design of low-noise amplifiers (LNAs), voltage conversion units, operational amplifiers, and power-efficient sigma-delta analog-to-digital and digital-to-analog converters. Feedback control on the circuit level is typically used to compensate for component variations due to temperature, voltage, and aging.

Application Control

In mobile telephones, application control refers to the control of on-device resources. Boundaries between mobile phones and computers are disappearing. A major challenge is to facilitate distributed application development on scalable architectures, where the amount of available computational resources, memory, and power is unknown until runtime. Thus, feedback control loops are also becoming important for controlling computational resources in mobile phones. Reliable temperature control is also important for products that lack the ability to survive critical situations by starting a cooling mechanism such as a fan.

Contributor: Bo Bernhardsson, Lund University, Sweden

Power Control

In the most-used version (WCDMA FDD) of the 3G radio standard introduced at the beginning of the millennium, all mobile phones in a radio cell transmit simultaneously on the same frequency. A clever design of the coding scheme makes it possible to filter out and amplify the wanted part of the received signal. All other transmissions will act as noise. Thus, controlling the power of all transmitted signals is critical; failed power control in one mobile can destroy the operation of an entire cell.

The base station (BS) and mobile phones cooperate to control both the downlink signal power (BS to mobile) and uplink signal power (mobile to BS) using two control loops. An interesting coupling between these loops arises because failure in downlink control will have an impact on the communication of control commands for the uplink power, and vice versa.

Because the controller includes integral action, anti-windup control must be used when the control loop is broken. The 3G standard includes tests for proper anti-windup. An interesting control situation also occurs during the so-called "soft handover," when several base stations simultaneously try to control the output power of the mobile.

Automatic Frequency Control (AFC)

For correct reception of radio signals, the local oscillator in the mobile phone must have the same frequency as the signal to be received. The relative frequency accuracy targeted for good reception is on the order of 0.01 to 0.1 parts per million (ppm).

Without feedback control, achieving this specification would require crystal oscillators with high power consumption. The crystal oscillators would also be large and expensive. The accuracy achievable with open-loop control and at reasonable cost is on the order of 10 ppm today. Thus, feedback extends the technological frontier by a factor of 1000. The main disturbances for which feedback is essential are due to temperature variations, the Doppler effect for moving users, variations over battery voltage, and oscillator frequency and aging.

The AFC control loop locks the oscillator phase to the phase of the received radio signal using known transmitted signals, digital "pilot" symbols. The controller can be of proportional-integral (PI) type, and the main design tradeoff is to achieve good noise rejection and fast tracking of frequency variations simultaneously. Gain scheduling is typically used, with faster control for rapidly moving phones.

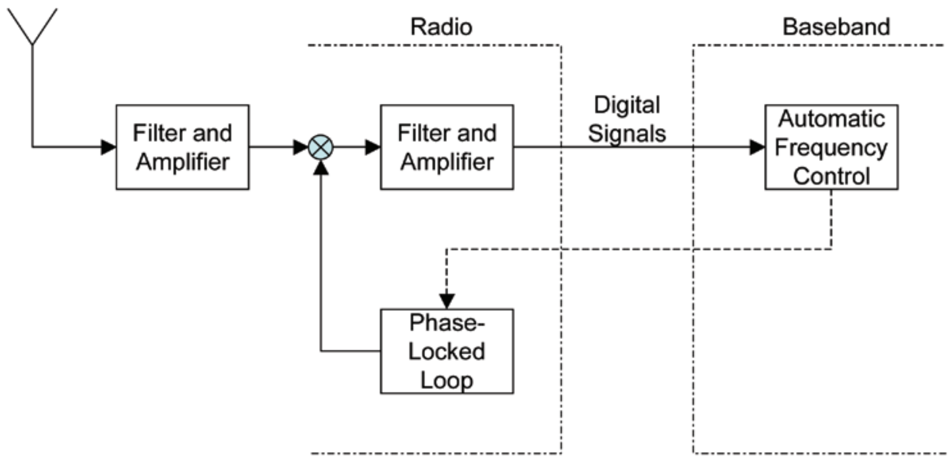

Mobile phones include a radio unit, which works with analog signals at high frequencies using analog circuits, and a baseband unit, which works with digital signals using digital hardware blocks and special-purpose digital signal processors. Automatic frequency control thus controls the analog phase-locked loop using digital symbols.

Automatic Gain Control (AGC)

The strength of the received signal shows large variations depending on the distance from the transmitter to the receiver. The receiver must show linear behavior for an operating range between -25 and -115 dBm (where 0 dBm equals 1 mW); that is, a 10^9 power change on the input (comparable to the power ratio between a lamp and a nuclear power plant).

Low-cost electronic components with such dynamic range are not feasible today. Feedback is used in several stages to control the gain of each block in the receiver chain so that the output signal fits the dynamic range of the succeeding block. The AGC loops must be sufficiently fast to track channel propagation variations. A PI controller with gain scheduling is often used.

Control in Stroke Rehabilitation

Stroke is the foremost cause of disability in developed countries. Less than 15 percent of patients with upper-limb impairment following stroke regain full function, which restricts their ability to perform everyday reaching and grasping tasks. Functional electrical stimulation (FES) used to assist stroke patients in moving their impaired limbs has been shown to increase upper-limb function; however, the benefits of FES are greatest when combined with maximal voluntary effort from the patient to perform the movement. This presents a control problem: to provide the right amount of FES to assist with movement while also encouraging maximal voluntary effort.

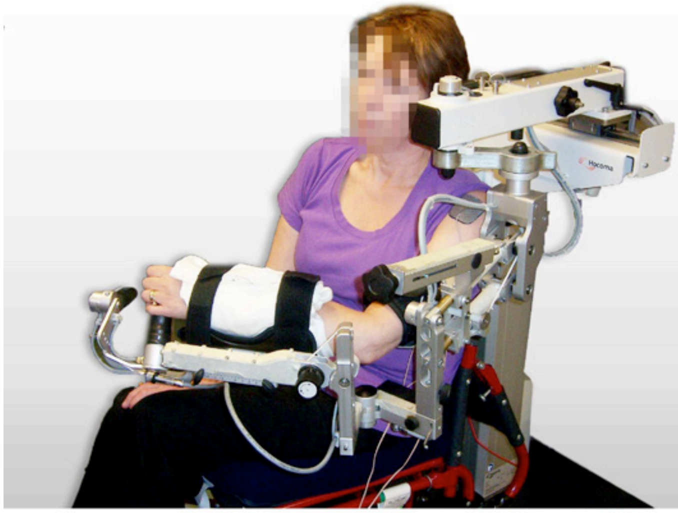

An upper-limb rehabilitation system has been developed at the University of Southampton that uses iterative learning control (ILC), FES, and robotic support. Significant improvement in patient arm movement has been realized.

Development of Upper-Limb Rehabilitation Systems

This development was the outcome of three main research programs:

- ***Planar reaching:*** Initial proof-of-concept experiments incorporated movement in one plane and stimulated one muscle group (triceps) to control movement around the elbow joint. Patients tracked a moving trajectory with their hand while FES was applied to assist with the movement. Following each trial, ILC updated the FES signal for the subsequent trial. Results showed improvements in tracking accuracy during the sessions.

- ***3-D virtual reality:*** Following the successful proof of concept, the system was extended to movements in 3-D space using a virtual reality tracking task. Patients' arms were supported by an Armeo robotic support (Hocoma, Switzerland), with FES applied to the triceps and anterior deltoid muscle groups to control movement around the elbow and shoulder joints. An experimental trial demonstrated the system's effectiveness, with improvements shown in tracking accuracy and in Fugl-Meyer clinical assessment scores. See the figure at top right.

- ***Functional reach and grasp:*** The most recent system advances the work to include control of the hand and wrist during functional tasks. ILC-controlled FES is now also applied to the extensors of the wrist and hand to assist with picking up and manipulating real-world objects. Minimal robotic support is provided by a spring system (SaeboMAS, USA), and patient tracking is achieved using Microsoft Kinect. A recent study reported improvements in patients' performances of functional tasks and clinical scores (Fugl-Meyer and ARAT). See figure at right.

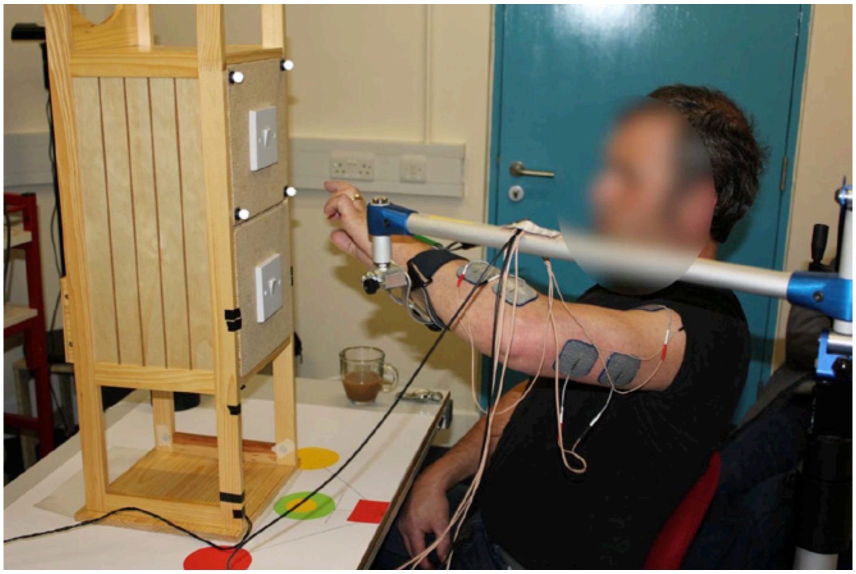

***Contributors:** Timothy A. Exell, Chris T. Freeman, Katie L. Meadmore, Ann-Marie Hughes, Emma Hallewell, Eric Rogers, and Jane H. Burridge, University of Southampton, U.K.*

Relearning of Functional Tasks

Rehabilitation Tasks

The current system incorporates common tasks of daily living, utilizing instrumented objects associated with daily life. The tasks currently performed include closing a drawer, pressing light switches, picking up and repositioning a drink, and stabilizing objects to assist the unaffected arm. Spasticity in stroke patients often restricts flexion of the shoulder, extension of the elbow, and extension of the wrist and fingers. Therefore, the anterior deltoid, triceps, and wrist and hand extensor muscles were selected for stimulation.

Control Approach

A simplified dynamic model of the arm-support system incorporates a biomechanical description of the human arm and a representation of the spring support (see below right). A proportional-integral-derivative (PID) controller is currently used in this system in parallel with phase-lead ILC, based on joint angle reference signals from unimpaired movement. The repetitive performance of the rehabilitation tasks used in this system make it an ideal application for ILC to control the FES signals for each muscle group. Performance error from each trial is used by the ILC to update the FES control parameters of the subsequent trial in an attempt to reduce error. This approach reduces the stimulation following successful performance, increasing the effectiveness of rehabilitation by requiring maximal patient effort.

Clinical Improvements

Reaching the clinical trial stage is often difficult in rehabilitation engineering. The FES-ILC technology has reached this stage, and initial clinical trials involving 15 chronic stroke patients over six-week periods have shown significant improvements in clinical assessments of arm movement across all patients.

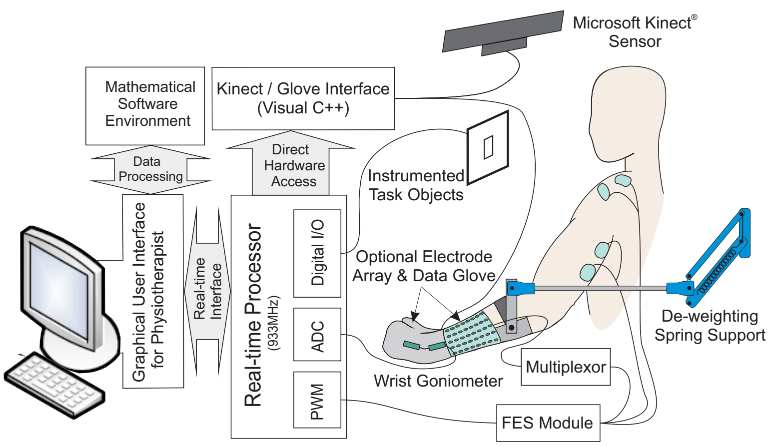

Left: Architecture of the current system, incorporating ILC and real-time controlled FES applied to each muscle group

Below: (a) Kinematic model of the spring support, and (b) anthropomorphic arm used in the control model

Future Directions

- Incorporate patient-customized models of movement within control design
- Reduce size and cost of system hardware for use within patients' homes
- Develop greater selection of functional tasks that can be incorporated, utilizing instrumented "real-world" objects
- Undertake a large-scale clinical trial through national stroke units

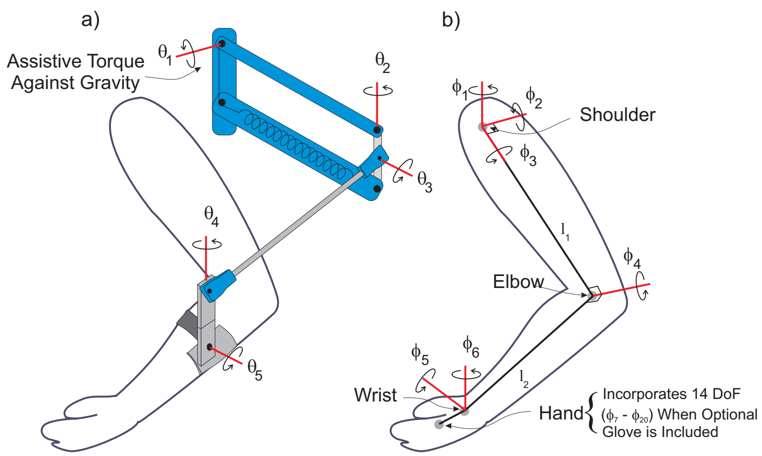

For more information: C.T. Freeman et al., Iterative learning control of FES applied to the upper extremity for rehabilitation, Control Engineering Practice, vol. 17, no. 3, pp. 368-381, 2009; K.L. Meadmore et al., Functional electrical stimulation mediated by iterative learning control and 3D robotics reduces motor impairment in chronic stroke, J NeuroEng Rehabil, vol. 9, no. 32, 2012; T.A. Exell et al., Goal orientated stroke rehabilitation utilizing electrical stimulation, iterative learning and Microsoft Kinect, in IEEE International Conference on Rehabilitation Robotics, Seattle, WA, 2013.

Control of NASA's Space Launch System

The flight control system for the NASA Space Launch System (SLS) employs a control architecture that evolved from Saturn, Space Shuttle, and Ares I-X while also incorporating modern enhancements. This control system, baselined for the first unmanned launch, has been verified and successfully flight-tested on the Ares I-X rocket and an F/A-18 aircraft. The development of the launch vehicle itself came on the heels of the Space Shuttle retirement in 2011 and will deliver more payload to orbit and produce more thrust than any other vehicle, past or present, opening the way to new frontiers of space exploration as it carries the Orion crew vehicle, equipment, and experiments into new territories. The initial 70-metric-ton vehicle consists of four RS-25 core stage engines from the Space Shuttle inventory, two five-segment solid rocket boosters that are advanced versions of the Space Shuttle boosters, and a core stage that resembles the External Tank and carries the liquid propellant while also serving as the vehicle's structural backbone. Just above SLS's core stage is the Interim Cryogenic Propulsion Stage (ICPS), based on the payload motor used by the Delta IV Evolved Expendable Launch Vehicle (EELV).

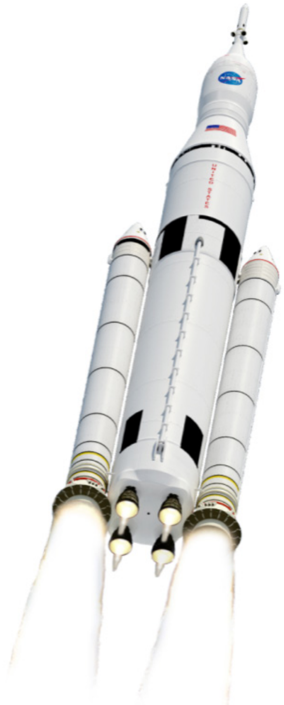

NASA's Space Launch System Block I Vehicle

Challenges of Controlling the Space Launch System

As with all large launch vehicles, the SLS control system must balance the competing needs of maximizing performance while maintaining robustness to limitations in preflight models. With its high thrust, large size, and multistructural load paths, SLS exhibits a high degree of structural flexion with nonplanar characteristics. Control commands are allocated to each of its six engines that are gimbaled along orthogonal axes by hydraulic thrust vector control actuators, leading to a control resource management challenge. The massive core propellant tanks and long-moment-arm upper-stage tanks have lightly damped lateral sloshing modes. The SLS trajectories are highly optimized to maximize payload-to-orbit performance. These challenges restrict the available design space for the flight control system and leave very little margin to share among the disciplines and subsystems.

SLS Flight Control System Architecture

Gyro blending: Optimizes attenuation of low-frequency structural modes by blending the sensed rotational rate from multiple sensors.

PID + bending filters: Classical control architecture with bending filters specifically tailored to attenuate or phase-stabilize vehicle bending and slosh.

Disturbance Compensation Algorithm (DCA): Generalization of Ares I-X anti-drift/load-relief algorithm designed to cancel external moments on the vehicle, minimize lateral drift, and alleviate gust loading.

Optimal Control Allocation (OCA): Allocates the desired control command to the engine thrust vector control actuators and includes logic to accommodate engine-out, closed-loop throttling and staging events.

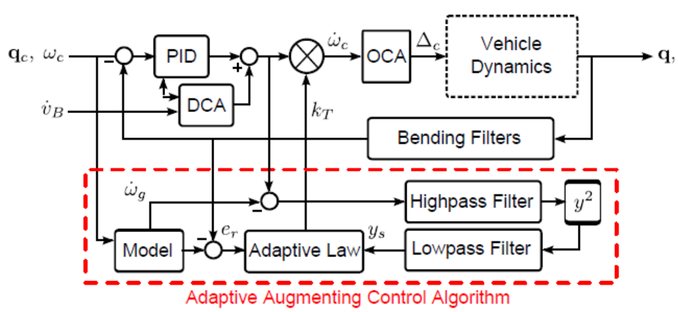

Simplified block diagram of the SLS flight control system

Adaptive Augmenting Control (AAC): Most novel addition to the SLS control architecture. Modifies the total attitude control system response to provide the fixed-gain control architecture with additional robustness. Increases performance when excessive tracking error is present and decreases the responsiveness when undesirable frequency content is observed in the control command as a result of internal dynamics (flexibility, fuel slosh, and actuators). Expands the envelope under which the Flight Control System (FCS) is capable of safely flying the vehicle, thereby maximizing vehicle survivability and crew safety.

Contributor: Tannen S. VanZwieten, NASA Marshall Space Flight Center, USA

Infusion of Advanced Technology: The Adaptive Augmenting Control Algorithm

In the absence of vehicle or environmental uncertainty, a fixed-gain controller could be optimized prior to flight such that there would be no need for adaptation; however, a review of historical reusable launch vehicle data from 1990 to 2002 revealed that 41% of failures resulting from other malfunctions might have been mitigated by advanced guidance, navigation, and control technologies. Traditional barriers to capitalizing on the benefits of advanced control techniques that are particularly relevant for human-rated systems include algorithm and code complexity, predictability of the response, ability to reconcile the stability analysis in the context of classical gain and phase margin, and flight certification. Thus, an algorithmically simple, predictable adaptive augmenting control design with direct ties to classical stability margins was implemented for SLS. It was initially formulated and tested during the Constellation Program, then refined as part of the baseline autopilot design and flight software prototype for SLS and flight tested on an F/A-18.

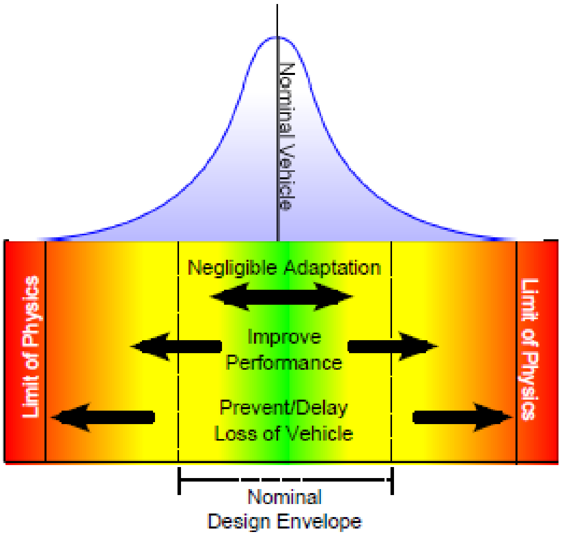

AAC design philosophy

Ares I-X flight test launch (October 28, 2009)

Flight Testing

A majority of the components of the SLS FCS for vehicle ascent were flown on Ares I-X: sensor blending, PID, bending filters, DCA, and parameter identification maneuvers. A series of F/A-18 research flights was used to test SLS prototype software, including the previously untested AAC and OCA components. Except for disabling the DCA, the control parameters for this test were identical to the SLS design set. The fighter aircraft completed a series of trajectories during multiple sorties with the SLS FCS enabled while matching the aircraft's pitch rate to that of the SLS and matching attitude errors for various nominal and extreme SLS scenarios through the use of a nonlinear dynamic inversion controller. The emphasis of the 100+ SLS-like trajectories was on fully verifying and developing confidence in the adaptive augmenting control algorithm in preparation for the first unmanned launch of SLS.

Photographs depicting the trajectory flown repeatedly by NASA's F/A-18 while testing the SLS AAC (2013, through a partnership among the NASA organizations MSFC, AFRC, NESC, and STMD-GCT)

For more information: J. Orr et al., Space launch system ascent flight control design, Proc. AAS Guidance and Control Conf., 2014; J. Wall, J. Orr, and T. VanZwieten, Space launch system implementation of adaptive augmenting control, Proc. AAS Guidance and Control Conf., 2014; T. VanZwieten et al., Adaptive augmenting control flight characterization experiment on an F/A-18, Proc. AAS Guidance and Control Conf., 2014.

Controller Performance Monitoring

A typical industrial process, as in a petroleum refinery or petrochemical complex, includes thousands of control loops. Instrumentation technicians and engineers maintain and service these loops, but rather infrequently. Routine maintenance of such loops can result in significant savings. Controller performance monitoring (CPM) can identify and diagnose incipient problems. CPM implementations have been successfully deployed in large sites and have substantially improved the performance of control loops.

Source: BASF SE

Identifying and Fixing Control Loop Problems

Studies indicate that on average only 40% of industrial control loops are delivering satisfactory or optimal performance. As many as 60% of control loops may have poor tuning or configuration or actuator problems and thus may be responsible for suboptimal process performance. As a result, monitoring of such control strategies to detect and diagnose the cause(s) of unsatisfactory performance has received increasing attention from industrial engineers. Specifically, the methodology of data-based controller performance monitoring is able to answer questions such as the following:

"Is the controller doing its job satisfactorily, and if not, what is the cause of the poor performance?"

In many of today's plants, performance of the process control assets is monitored on a daily basis and compared with industry benchmarks. The monitoring system also provides diagnostic guidance for poorly performing control assets. Many industrial sites have established reporting and remediation workflows to ensure that improvement activities are carried out in an expedient manner. Plantwide performance metrics can provide insight into companywide process control performance. Closed-loop tuning and modeling tools can also be deployed to aid with improvement activities.

Industrial Implementations

CPM software is now readily available from most distributed control system (DCS) vendors and has already been implemented successfully at several large-scale industrial sites. Large-scale industrial implementations of CPM technology provide clear evidence of the impact of this control technology and its adoption by industry.

Operational Applications of Controller Performance Monitoring

- As part of its OPAL 21 (Optimization of Production Antwerp and Ludwigshafen) excellence initiative, BASF has implemented the CPM strategy on more than 30,000 control loops at its Ludwigshafen site in Germany and on more than 10,000 loops at its Antwerp production facility in Belgium.
- As part of its process control improvement initiative, Saudi Aramco has deployed CPM on approximately 15,000 proportional-integral-derivative (PID) loops, 50 model predictive control (MPC) applications, and 500 smart positioners across multiple operating facilities.

***Contributors:** Sirish L. Shah, University of Alberta, Canada; Marcus Nohr, BASF SE, Germany; and Rohit Patwardhan, Saudi Aramco, Saudi Arabia*

Plantwide Performance Assessment

The key to using this technology effectively is to combine process knowledge, basic chemical engineering, and control expertise to develop solutions for the indicated control problems that are diagnosed in the CPM software. The operational philosophy of the CPM engine is incorporated in the continuous improvement process at BASF and Aramco, where all loops are monitored in real time and a holistic performance picture is obtained for the entire plant. Unit-wide performance metrics can be displayed in effective color-coded graphic forms as shown at right. Detailed reports can be accessed for every loop in units that require attention followed by diagnosis of poor performance, as shown below.

Color-coded graphics for plantwide CPM. The tiled rectangles refer to plant subsystems (e.g., units); the size of each rectangle can be proportional to energy consumption, alarm counts, or a user-defined property. Colors indicate the current assessment of each subsystem. (Source: BASF SE)

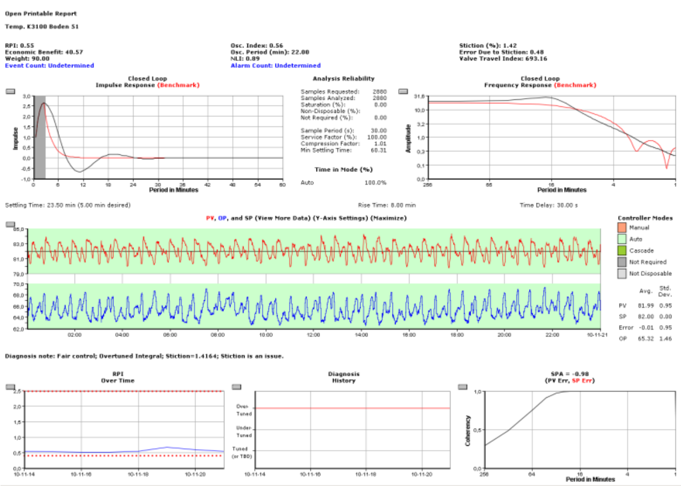

Detailed diagnosis of individual loops. Several parameters are calculated and tracked to identify incipient performance or safety issues. (Source: BASF SE)

CPM for Controller Optimization

The main objective in implementing CPM is to facilitate controller optimization. CPM monitors performance and aids in the diagnosis and remediation of poorly performing loops. The figure at lower left shows results of controller reconfiguration attained through controller performance monitoring and subsequent diagnostic analysis. The reconfigured control loop is able to reduce variability, resulting in smoother process operation close to optimum constraints with increased throughput. Other ways to use the additional degrees of freedom from controller optimization are to allow for reduced energy consumption or improved product quality. The key benefit of this technology is improved performance from the regulatory and advanced control layers, resulting in:

- Improved plant stability,
- Reduced operator load,
- Reduced process variability and as a result operation closer to economic constraints, and
- Improved economic margins for the process.

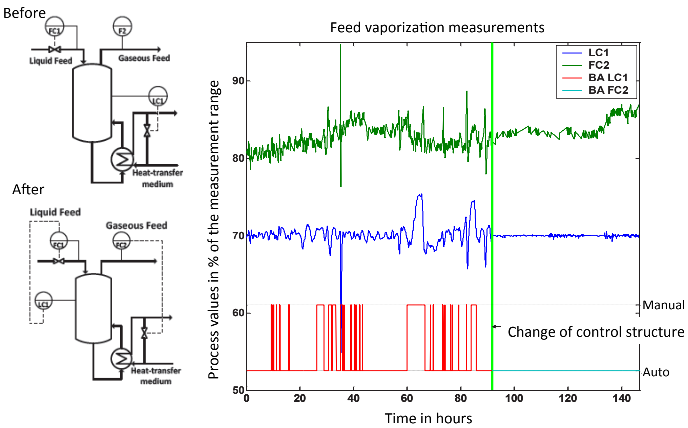

Remediation of poor performance of a feed evaporator in a petrochemical plant after CPM diagnosis (LC1/FC2: level/flow measurements under level/flow control; BA: manual or automatic controller modes) (Source: BASF SE)

For more information: B. Huang and S.L. Shah, Performance Assessment of Control Loops: Theory and Applications, Springer-Verlag, October 1999; F. Wolff, M. Roth, A. Nohr, and O. Kahrs, Software based control-optimization for the chemical industry, VDI, Tagungsband, Automation 2012, Baden-Baden, Germany; L.D. Desborough and T.J. Harris, Performance assessment measures for univariate feedback control, Can. J. Chem. Eng., vol. 70, no. 6, pp. 1186-1197, 1992.

Controlling Energy Capture from Wind

Wind energy is currently the fastest growing power-generation technology worldwide, reaching a 30% annual growth rate and an installed capacity of 300 GW. To realize these achievements, wind turbine designs have overcome multiple technical challenges to be competitive with predominant energy sources. Control technology has played a crucial role in this quest. The control system dynamically adapts to a wide range of wind conditions and maintains structural integrity while maximizing energy production. In addition, the controller must manage weather conditions, abnormal wind disturbances, and fault scenarios that may occur unexpectedly during the life span of the turbine.

Leveraging the experience of more than 22,000 units installed worldwide, General Electric has developed a comprehensive model-based control system for wind energy that includes algorithms to adapt to variable wind conditions, fault-tolerant strategies to accommodate failures and extreme disturbances, and a supervisory system that manages both the turbine and farm-level fully automated operations.

Challenges of Controlling Wind Turbines

Wind turbines are rather complex systems, with multiple flexible structures coupled through highly nonlinear dynamics and subject to varied wind disturbances. Their designs need to satisfy stringent requirements to ensure safe and reliable operation for a 20-year or longer life span. Indeed, international standards such as IEC 61400 require extensive numerical simulations to ensure survival in conditions ranging from normal operation under turbulent wind to a multiplicity of fault events combined with abnormal wind or weather scenarios. An example requirement is to consider the most likely sudden failures of the blade pitch system or storm conditions and show that the turbine can shut down without overstressing the blades, the tower, or the drive train. Maximizing energy production under these requirements is achieved with only a limited set of actuators that command blade pitch, rotor yaw, generator torque, and brakes to stop the rotor. In addition, to avoid the cost of on-site manual maintenance, the turbine controller includes extensive logic to perform automatic calibration procedures, self-diagnostics, and supervisory functions.

Contributor: Fernando D'Amato, GE Global Research, USA

Solution Approach

The control system for GE wind turbines includes:

- Turbine control algorithms that adapt operation to each wind condition
- Supervisory controls to handle failures and braking procedures
- Farm-level control to maximize energy capture and manage grid integration

Wind turbine controls rely heavily on model-based technology at every layer of implementation. Namely, the basic dynamic equations of the drivetrain, tower, and rotor are included as an integral part of controller algorithms. Model-based estimation is used to learn the state of the wind and structural components. The main turbine controller uses this model to achieve consistent disturbance rejection response across the operating range. In addition, first-principle concepts of dissipative systems theory are implemented to ensure safe braking and shutdown procedures. Finally, optimization algorithms based on wind models are used to generate the turbine control setpoints to maximize energy capture and minimize noise at the farm level.

Results and Benefits

The limiting loads of wind turbine components are highly dependent on the ability of the control system to cope with all design requirements. Decreasing the maximum load of components directly translates into cost reduction or increased performance. For example, improvements in the control system for 1.6-MW turbines enabled use of larger blades, resulting in a 20% increase in annual energy capture. For turbines in the range of 2.5 MW, development of fault-tolerant controls to handle extreme wind events reduced component loads by up to 30% in several turbine components.

In the foreseeable future, control design improvements will continue to increase operating efficiency and reliability and reduce turbine costs. Advances in controls are crucial for more competitive wind power generation!

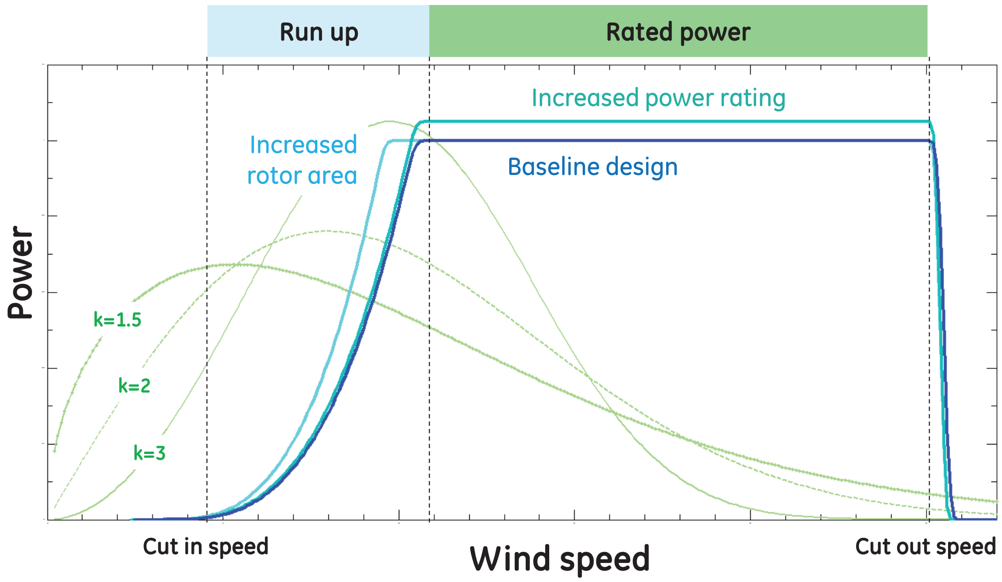

By effectively managing structural loads, control technology enables significantly higher power to be captured from wind resources, as shown by these power curves. Lower loads allow using generators with increased power rating or increased rotor area (i.e., longer blades) than the baseline design. The three green lines show typical wind site characteristics and represent the proportion of time expected for prevailing wind speeds. Each characteristic is approximated by a Weibull distribution with different shape parameter k.

For more information: visit http://www.ge-energy.com/products_and_services/products/wind_turbines/ges_2.5_120_wind_turbine.jsp.

Coordinated Ramp Metering for Freeways

Freeways were originally conceived to provide virtually unlimited mobility to road users, but the continuous increase in car ownership and demand has led to a steady increase (in space and time) of recurrent and nonrecurrent freeway congestion, particularly in and around metropolitan areas. Freeway congestion causes excessive delays, increases fuel consumption and environmental pollution, and deteriorates traffic safety.

Ramp metering, the most direct and efficient way to control freeway networks, aims at improving traffic conditions by appropriately regulating inflow from the on-ramps to the freeway mainstream. Coordinated ramp-metering strategies make use of measurements from a freeway network to control all metered ramps included therein. A new traffic-responsive feedback control strategy that coordinates local ramp-metering actions for freeway networks has been developed at the Dynamic Systems and Simulation Laboratory of the Technical University of Crete, Greece. The proposed coordination scheme is named HERO (heuristic ramp metering coordination) and has been extensively tested via simulation as well as in field implementations. The developers are the recipients of the 2010 IEEE CSS Transition to Practice Award, a prize awarded by the IEEE Control Systems Society to recognize outstanding university–industry collaboration that enables the transition of control and systems theory to practical industrial or commercial systems.

Solution Overview

HERO is simple and utterly reactive, that is, based on readily available real-time measurements, without the need for real-time model calculations or external disturbance prediction. HERO is modular in structure and includes many interacting and cooperating feedback control loops (such as mainstream occupancy control, ramp queue-length control, waiting time control), as well as two Kalman filters for estimation of ramp queue length and mainstream critical occupancy. Generic software has been developed that implements the HERO coordination scheme for any freeway network via suitable input configuration. Several extensions and improvements have been incorporated into the HERO system based on the experience gained from implementation, making it more efficient and generally applicable.

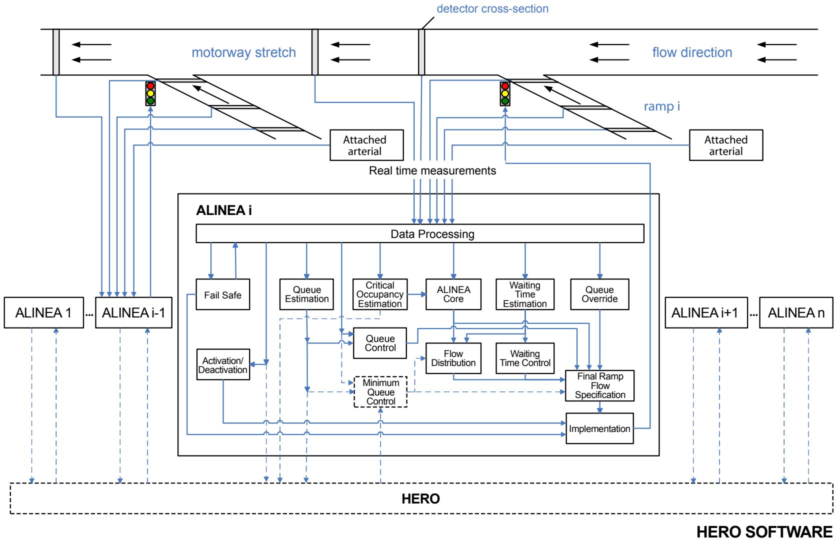

Modular structure of the HERO coordination software, including the ALINEA regulator for local ramp metering

Contributors: Markos Papageorgiou and Ioannis Papamichail, Technical University of Crete, Greece

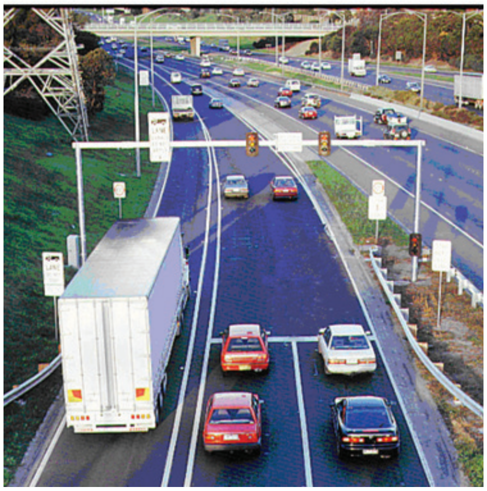

Field Application at the Monash Freeway, Australia

Since 2000, Melbourne's freeways had become increasingly congested with extended periods of flow breakdown. The Monash Freeway, a six-lane dual carriageway that carries in excess of 160,000 vehicles per day, of which up to 20% are commercial vehicles, was experiencing long periods of congestion lasting between three and eight hours a day.

In early 2008, to address this problem, the responsible road authority, VicRoads, implemented HERO at six consecutive inbound on-ramps of the Monash Freeway. This $1M (Australian) pilot project was part of the Monash-CityLink-West Gate (MCW) upgrade and received two 2009 Victorian Engineering Excellence Awards, one for Technology and one for Engineering Innovation. Significant benefits were demonstrated over the previous metering policy. The control logic has proven to be robust and transparent to traffic engineers. Transition to HERO has been seamless to motorists and has provided significant flexibility and capability to operate the freeway close to optimal conditions. The pilot project economic payback period was estimated at just 11 days. The successful implementation and evaluation of HERO led to its rollout during 2009–10 at 63 sites across the entire 75-km route of the MCW upgrade project; furthermore, 12 new sites of M80 and two new sites of M8 were added in 2012–13.

An evaluation of HERO's field performance was undertaken by VicRoads. HERO reduced the space-time extent of freeway traffic flow breakdown and provided significant improvements in throughput and travel speed. The a.m. peak evaluation revealed a 4.7% increase in average flow (over the previous system) and a 24.5% increase in average speed, whereas the p.m. peak evaluation showed an 8.4% increase in average flow and a 58.6% increase in average speed.

Other Field Applications

In 2006, HERO was also implemented in a 20-km stretch of the inbound A6 freeway in the south of Paris, France, albeit in a preliminary and simplified form due to lack of real-time on-ramp data in the control center. Nevertheless, results indicated a clear improvement over the existing system. In 2011, HERO was deployed at six on-ramps along a section of the Pacific Motorway and South East Freeway (M1/M3) in Queensland, Australia. The successful implementation and evaluation of HERO led to its extension during 2012 to two more sites. Finally, HERO is currently being implemented in a North American freeway, and further implementations are in a planning or discussion phase.

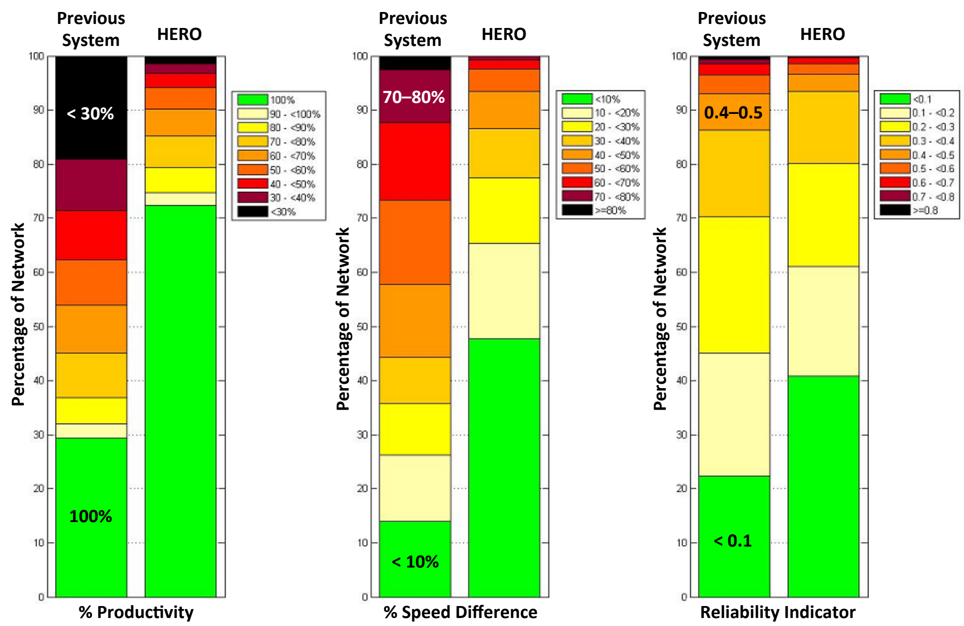

Results from the HERO implementation on the Monash Freeway using Austroads National Performance Indicators (ANPI). Three indicators are shown, with side-by-side before-and-after comparisons for each. Left to right, the indicators are: productivity (a combination of high speed and high volume on the freeway), mean speed deviation from the posted speed limit, and reliability (reflecting travel time differences from day to day). The green-to-black color codes reflect best-to-worst performance.

For more information: I. Papamichail and M. Papageorgiou, Traffic-responsive linked ramp-metering control, IEEE Transactions on Intelligent Transportation Systems, vol. 9, pp. 111–121, 2008; I. Papamichail et al., Heuristic ramp-metering coordination strategy implemented at Monash Freeway, Australia, Transportation Research Record, Journal of the Transportation Research Board, no. 2178, pp. 10–20, 2010.

Digital Fly-by-Wire Technology

Digital fly-by-wire (DFBW) is one of many success stories where technology developed under the U.S. space program has proven beneficial in other areas. Based in part on a recommendation from Neil Armstrong, who was directly familiar with the Apollo Guidance Computer through his historic lunar landing, NASA's Dryden Flight Research Center chose to work with Draper Laboratory to adapt the concept for aircraft, beginning with experimentation on a U.S. Navy F-8 Crusader in 1972.

From top to bottom: Space Shuttle, Airbus A320, B-2 Stealth Bomber, Boeing 777, Dassault Falcon 7X, and Joint Strike Fighter X35 (Sources: NASA, Airbus, Boeing, and Dassault)

Draper developed DFBW as an extension of its work on the Apollo Guidance Computer. The concept uses a highly reliable computer and electronic flight control system, rather than mechanical or hydraulic-based systems, to stabilize and maneuver a vehicle. The computer is able to execute far more frequent adjustments than a human pilot, thus helping maintain stability while offering increased maneuverability.

The 15-year DFBW technology research program also demonstrated adaptive control laws, sensor analytical redundancy techniques, and new methods for flight testing digital systems remotely.

Real-World Applications

The F-8 digital fly-by-wire program served as the springboard for DFBW technology to be used in both military and civilian aircraft. Today, commercial launch service providers and satellite manufacturers also routinely use the technology in their vehicles and spacecraft. Below is a partial list of aircraft and spacecraft with DFBW technology:

- Space Shuttle
- Launchers: Ariane, Vega, Titan, Delta, Proton
- Airbus A320 (first airliner with DFBW controls)
- Boeing 777 and 787
- Jet fighters: F-18/22, Dassault Rafale, Eurofighter, Joint Strike Fighter X35
- Stealth Fighter/Bomber: F-117, B-2
- Dassault Falcon 7X (first business jet with DFBW controls)
- Rotorcraft: V-22 Osprey, RAH-66 Comanche, AH-64 Apache, NH-90, Sikorsky S-92
- Several unmanned aerial vehicles (UAVs)

Apollo computer interface box used in the F-8C digital fly-by-wire program (Source: NASA)

NASA used an F-8C for its digital fly-by-wire program, the first DFBW aircraft to operate without a mechanical backup system. This photo shows the Apollo hardware jammed into the F-8C. The computer is partially visible in the avionics bay. (Source: NASA)

Contributor: Christian Philippe, European Space Agency, Netherlands

Space Technology Hall of Fame 2010

NASA's Dryden Flight Research Center, Draper Laboratory, The Boeing Company, and Airbus were inducted into the Space Technology Hall of Fame in 2010 for the development of digital fly-by-wire technology that makes modern aircraft easier and safer to operate.

Left: The RAH-66 was the first helicopter that was capable of backflips. Right: X-31 demonstrating high angle of attack maneuver.
(Sources: Boeing/Sikorsky and NASA)

Major benefits of DFBW flight controls include:

- Overall cost reduction
- Overall airframe weight reduction
- Increased safety and reliability
- Fuel efficiency
- Reduced CO_2/NO_x emissions
- Improved flying (or handling) qualities
- Improved passenger comfort
- Reduced pilot workload
- Ease of assembly and maintenance
- Improved survivability
- Improved mission performance

The following features, enabled by DFBW, are currently implemented onboard fighter aircraft:

- Reconfigurable flight control system allowing mission continuation or safe recovery following system failures or battle damage
- Flight envelope protection such as bank angle protection, turn compensation, stall and overspeed protection, pitch control and stability augmentation, and thrust asymmetry compensation
- Online system identification for verification of the aerodynamic effects on aircraft flexible modes

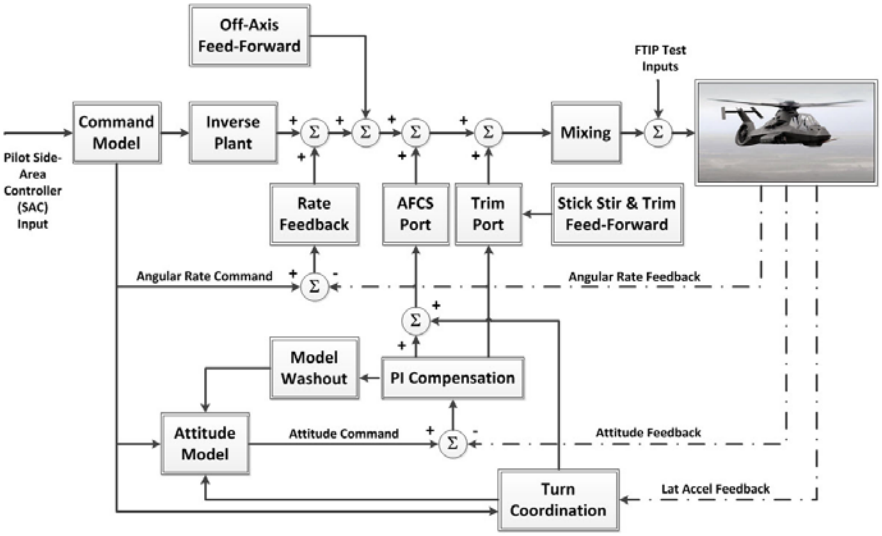

RAH-66 Comanche multimode control law architecture
(Source: Boeing/Sikorsky)

For more information: J.E. Tomayko, Computers Take Flight: A History of NASA's Pioneering Digital Fly-By-Wire Project, The NASA History Series, NASA SP-4224, National Aeronautics and Space Administration, 2000; NASA Dryden Technology Facts – Digital Fly By Wire, http://www.nasa.gov/centers/dryden/about/Organizations/Technology/Facts/TF-2001-02-DFRC.html.

Digital Printing Control: Print Shop in a Box

DocuColor® 8002

iGen4®

iGen4® 220 Perfecting Press

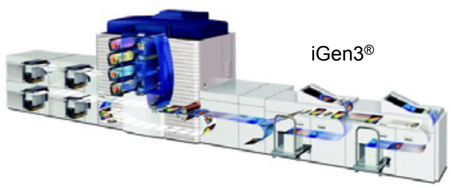

iGen3®

Xerox® Color 800/1000

ColorQube™

Xerox® Color 560/570

The digital print process is remarkably challenging because it involves many process and digital actuators for applying a range of advanced control techniques. Many of the new challenges listed here opened the door for the insertion of new control theory. Numerous Xerox® printing systems (for example, iGen3®, iGen4®, DocuColor® 7002, DocuColor® 8002, DocuColor® 5000, DocuColor® 8000, iGen4® 220 Perfecting Press, Xerox® Color 800/1000 Presses, ColorQube™) produce high-quality prints using these control innovations to generate several billion dollars in annual revenue.

Control Challenges for Digital Printing

- Optimize job workflows via feedback from the press (streamline workflow and free operators to focus on running the print jobs)
- Increase productivity with automated color management tools
- Provide consistent color image quality (first page, between pages, job to job, operator to operator, machine to machine)
- Provide offset look and feel with the best image quality, no nonuniformity, and no defects
- Provide automated calibration (completely hands-free), spot color (Pantone® matching), and color management profiling for more stable color
- Allow mix-and-match of press configurations (any application to any printer, finisher, and feeder on any marking engine)
- Manage time-sensitive activities of various machine modules
- Adjust color dynamically using internal process control feedback loops
- Provide active control of registration of all color separations
- Compensate for sheet-to-sheet differences in the paper as well as drive system wear, temperature variations, and the like

Contributors: Lalit K. Mestha and Alvaro E. Gil, Xerox Corporation, USA

From: *The Impact of Control Technology*, 2nd ed., T. Samad and A.M. Annaswamy (eds.), 2014. Available at www.ieeecss.org.

Digital printing today is a complex, high-technology process requiring advanced sensors and actuators and state-of-the-art control algorithms. Processing of images occurs at multiple levels within and outside a hierarchical printing and publishing system. Time-based separation is adopted at each level, with higher-level functions (spot colors, ICC profiles, gray balance) occurring at a slower time scale in the digital front end (DFE) and faster real-time control functions (Levels 1, 2, and 3) typically occurring in the print engine (PE).

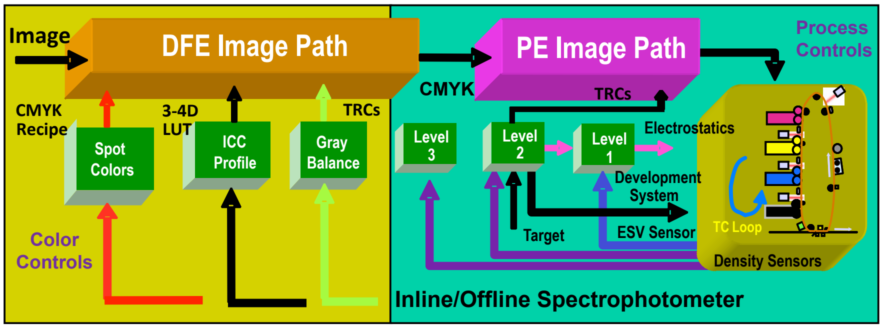

Control system architecture for digital printing (LUT: look-up table; ICC: International Color Consortium; TRC: tone reproduction curve; ESV: electrostatic voltmeter; TC: toner concentration; CMYK, cyan, magenta, yellow, black colors)

Inventions and Innovations

Xerox's new printing systems are tours de force for control technology! Highlights include:

- A hierarchical automation architecture distinguishes between different color control horizon levels.
- The control design includes classical single-input, single-output (SISO) PID-type controllers, with delay and anti-windup compensation, for several subsystems.
- Toner concentration control, although a SISO system, is especially challenging. The solution integrates a Kalman filter to handle noisy and unreliable measurements, a Smith predictor to handle the delay, and an anti-windup compensator for constraint management.
- Several multivariable controllers are also part of the design. These include state feedback, pole placement, model predictive control, and linear quadratic regulator designs. Systems under multivariable control include electrostatic control, spot color calibration, and color management profiling (ICC profile).
- A learning algorithm is incorporated for paper registration control.
- Singular value decomposition is used for dimensionality reduction to reduce gray-level samples while constructing spatial toner reproduction curves or functions on photoconductors.
- Ideas from cooperative control theory and simultaneous perturbation stochastic approximation have been used for gray-component replacement in the color management profiling (ICC profile) system.
- K-means clustering for spectral reconstruction in real time allows the use of a low-cost LED-based spectrophotometer.
- For complex printing jobs, scheduling of paper sheets is a difficult operation; constraint-based scheduling algorithms are used to solve it.
- Motion and unevenness in motion can induce disturbances in the printing process. Repetitive control and adaptive feedforward control algorithms help mitigate the effects of these disturbances.
- The printers integrate diagnostic and self-monitoring features using statistical process control, among other methods.
- A set of objectives customers may be interested in is taken into account (e.g., smoothness, spectral matching, and color difference) and the concept of Pareto front is used to provide optimal solutions to accommodate a wide variety of user preferences for spot color rendering.
- Constrained multi-objective optimization is used to optimize image quality criteria, noise mottle frequency, and smoothness during multidimensional image profile construction. Constraints are imposed in terms of color accuracy and spectral response to achieve robust color match under various illuminants.
- The process is fully automated to optimize Black point compensation (BPC) parameters (inputs), including a virtual sensor that estimates the effect of the parameters on three image metrics (outputs): shadow details, highlight details, and color attributes present in images. A model is developed using input/output data, and parameters for BPC are estimated with multi-objective optimization while constructing ICC profiles.

For more information: L.K. Mestha and S.A. Dianat, Control of Color Imaging Systems, CRC Press, May 2009; L.K. Mestha and A.E. Gil, Color controls — An advanced feedback system, in The Control Handbook, Second Edition, CRC Press, 2010.

Dynamics and Control for Deep-Sea Marine Risers

A marine riser is a pipe that connects a floating platform such as an oil rig or drillship on the ocean surface to the sea floor. Used as a fluid conveyer, it transports undersea energy resources from the seabed to the platform on the surface. Marine risers are also used in relief operations for transporting mud, cement, and other materials to the seabed.

Due to winds, waves, and water currents, the floating platform on the sea surface responds in six-degree-of-freedom motions. The surge, sway, and yaw motions of the floating platform are controlled by means of a dynamic positioning system (DPS) or a passive mooring system. Forces on the marine riser can be mitigated with a riser tensioner control system.

Advanced control of marine risers and floating platforms is a crucial enabler for deepwater drilling and safety operations.

Development Driller III drilling relief well at Deep Water Horizon site (Courtesy of Paddy Ryan, www.ryanphotographic.com)

Deepwater Horizon Relief Well Application

In September 2009, Deepwater Horizon, an ultra-deepwater, dynamically positioned offshore oil drilling rig, drilled the deepest oil well in history at a vertical depth of more than 10,600 m in the Gulf of Mexico. On April 20, 2010, while drilling at the Macondo Prospect, an explosion on the rig caused by a blowout killed 11 crewmen and ignited a fireball visible from 35 miles (56 km) away. On April 22, Deepwater Horizon sank, leaving the well gushing at the seabed and causing the largest offshore oil spill in U.S. history. Over the next several months, relief wells were drilled to cap the well, which was finally declared sealed on September 19, 2010.

Advanced control was crucial to the relief effort. A dynamic positioning system enabled the rig to be quickly positioned automatically in surge, sway, and yaw directions on site without the use of mooring lines. The riser tensioner control enabled the drilling operation to take place despite the heave response of the vessel due to wind, wave, and ocean current forces from the marine environment.

Riser Tensioner System

As the bottom of the riser is fixed on the seabed, forces arising from the motion of the surface vessel must be managed. For vertical motions, if the floating platform heaves downward due to waves or swell, the riser would buckle under its own weight. On the other hand, if the platform heaves upward, the pulling will induce the transmission of high forces through the riser.

This heaving motion is mitigated by the riser tensioner system, which controls the upward pulling force. As shown in the schematic on the left, the upward force is transmitted by a number of wires in contact with sheaves and coupled to the tensioner assembly. Up to eight tensioner assemblies can be used to minimize the force communicated to the riser.

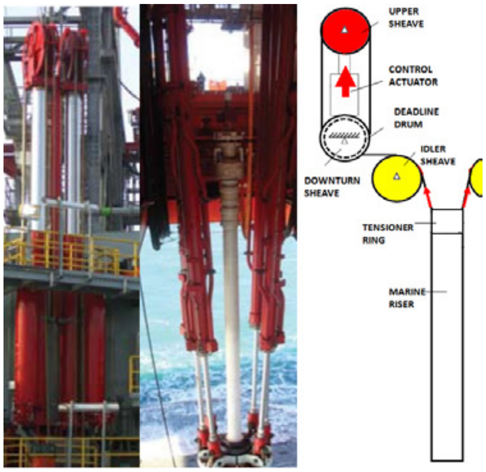

Riser tensioner system

Contributors: Shuzhi Sam Ge, National University of Singapore, Singapore, and University of Electronic Science and Technology of China, China; Yoo Sang Choo and Bernard Voon Ee How, National University of Singapore, Singapore; Wei He, University of Electronic Science and Technology of China, China

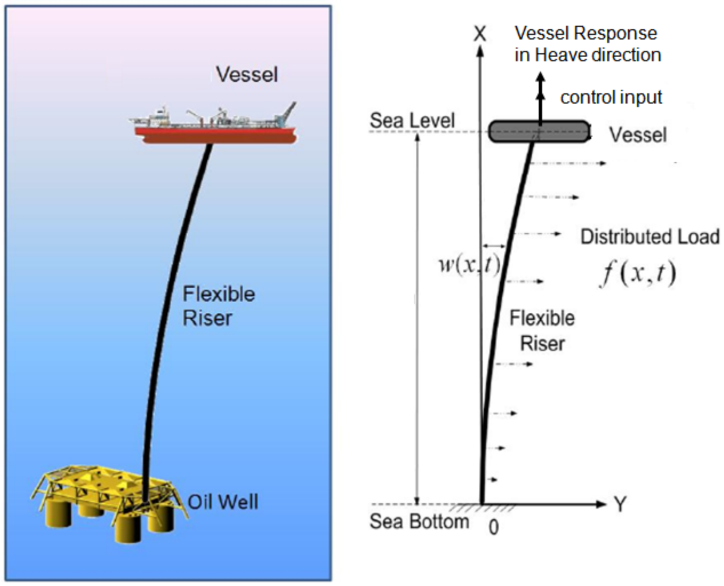

Illustration of the marine riser bottom fixed on the seabed and connected to a floating surface vessel

Vibration Suppression of Marine Risers

The lifting forces in the wire result from the actuation of a tensioner piston using high pressure from hydraulic cylinders. The cylinders are pressurized by a hydropneumatic accumulator that is connected to air pressure vessels via a control valve to maintain the specified pressures. To prevent extreme movement of the tensioner piston, a fluid flow shutoff valve is inserted between the hydropneumatic accumulator and the piston bore chamber of the tensioner. An air shutoff valve is inserted between the accumulator and control valve unit. In the event of a wire break, the control system restricts the flow of the volume via the flow shutoff valves. Thus the top tension of the riser is controlled, allowing the system to operate despite the heave response of the vessel.

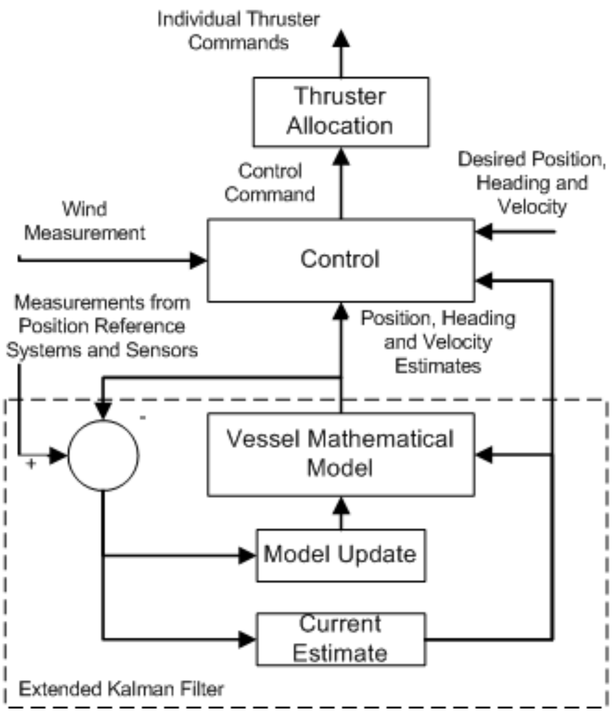

Simplified block diagram showing the dynamic positioning system with control, thruster allocation, and the extended Kalman filter

Dynamic Positioning System (DPS)

The DPS actuates the thrusters to keep the vessel within a fixed envelope determined by the operation being carried out. A high-precision control mode can be used for high-accuracy positioning in any weather conditions, at the expense of power consumption and wear and tear of machinery and thrusters. Another mode of control uses the thrusters more smoothly at the expense of a larger positioning envelope for less position-critical operations or during benign weather.

DPS control typically relies on a mathematical model of the vessel that includes hydrodynamic characteristics such as current drag coefficients, added mass, coriolis, and damping. As the mathematical model itself is not a perfect representation of the actual vessel, the model is continuously updated using an extended Kalman filter. The vessel heading and position are measured using sensors such as gyrocompassess, a position reference system, a differential global positioning system, inertial measurement units, and vertical reference sensors.

Ethylene Plantwide Control and Optimization

Ethylene is the largest-volume industrial bulk commodity in the world. The majority of ethylene is used in the production of ethylene oxide, ethylene dichloride, ethylbenzene, and a variety of homo- and copolymers (plastics ranging from plastic food wrap to impact-absorbing dashboards inside cars).

Ethylene plants are complex, large-scale, flexible factories that can process a wide variety of feedstocks, ranging from gases (such as ethane, propane, and LPG), to naphthas, to distillates and gas oils. Main products are polymer-grade ethylene and propylene. Operational objectives include yield improvement, production maximization, and energy intensity reduction.

Honeywell's advanced control and optimization technology has been applied to ethylene plants worldwide with substantial economic benefits, including millions of dollars from increased production annually and additional benefits from energy savings.

Process and Operating Characteristics

Universal:
- No product blending
- Stringent product quality requirements
- Slow dynamics from gate to gate
- Gradual furnace and converter coking
- Frequent furnace decoking and switching
- Converter decoking

Site-Specific:
- Feed quality variations
- Product demand changes
- Sensitivity to ambient conditions
- Periodic switching (for example, dryers)

Operating Flexibilities and Solution Goals

The main operating degrees of freedom for ethylene plantwide control and optimization include feed selection, furnace feed rates, cracking severity, dilution steam, cracked gas compressor and refrigeration compressor suction pressures, typical column variables (reflux, reboiler, and pressure), and converter temperature and H_2 ratio. Advanced control and optimization goals include:

- Stabilizing operation
- Minimizing product quality giveaway
- Maximizing selectivity and yield
- Minimizing converter over-hydrogenation
- Minimizing ethylene loss to methane and ethane recycle

Combined Control-Optimization Solution

Unlike optimization approaches based on steady-state models, the solution featured here relies on dynamic models of model predictive controllers (MPCs). There is no need to wait for the plant to reach steady state, and economic optimization is augmented into a standard MPC control formulation known as range control. Nonlinearity of the plant is accounted for with successive linear dynamic models. The use of nonlinear dynamic models is in development and has been demonstrated experimentally.

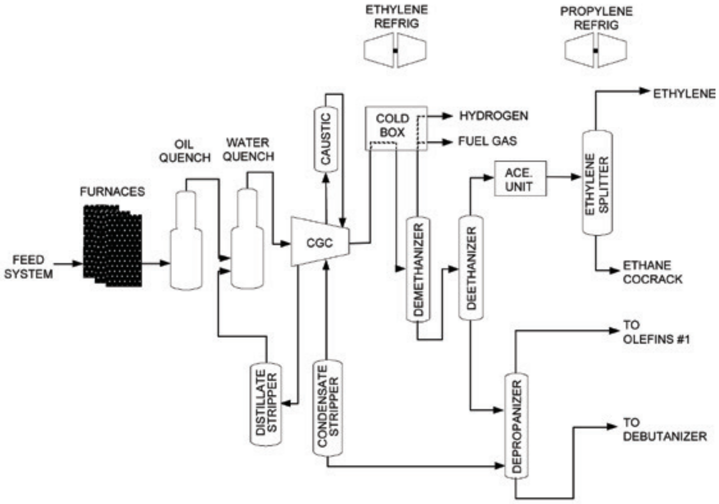

Ethylene plant schematic

Contributors: Joseph Lu and Ravi Nath, Honeywell, USA

Overview of an Ethylene Plantwide Control and Optimization Project

A typical advanced control and optimization solution for ethylene plants comprises a global optimizer (Profit Optimizer) that coordinates 15 to 30 model predictive controllers (Profit Controllers) for separation and quench towers, converters, and a fuel-gas system. MPC controllers execute every 30 to 60 seconds and the global optimizer every minute. Technip's SPYRO nonlinear model is used to update the furnace yield gains every 3 to 5 minutes.

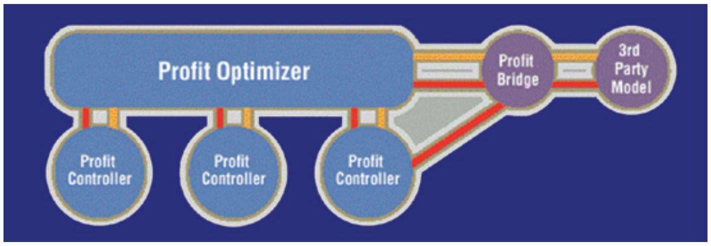

Honeywell's Profit Suite optimization and control products also include Profit Bridge for integrating third-party, nonlinear steady-state models; Profit Stepper for model identification (including closed-loop identification that allows models to be developed while the plant operates); and an advanced single-loop controller, Profit Loop.

The technology is based heavily on dynamic models. Steady-state nonlinear models are used selectively for calculating critical gains. All Profit Controllers operate off linear dynamical models (usually developed with the Profit Stepper application). The base Profit Optimizer model is automatically aggregated from the Profit Controller models; Bridge Models and source/clone structures are added to define interactions among the controllers.

Once a plant has been commissioned, validated, and put into operation (a process that takes 9 to 12 months), little maintenance is typically required. Clients either dedicate a half-time control engineer to monitor and perform minor services or depend on quarterly visits by Honeywell staff. (In contrast, traditional real-time optimization solutions typically require a full-time modeling person and a half-time control engineer to maintain the solution.)

Normalized olefin production rate before and after advanced control and optimization (ACO):
Production is increased and product variability reduced.

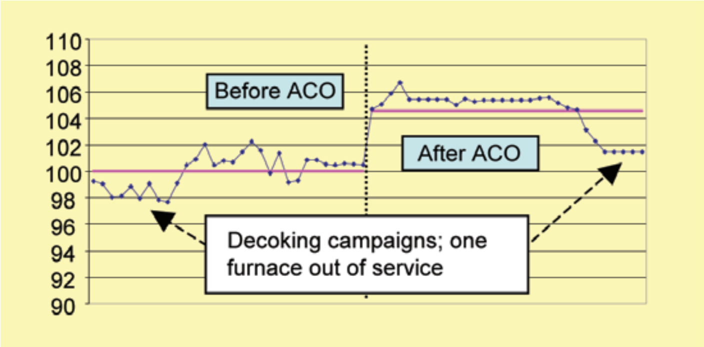

Ethane content (ppm) in product ethylene:
The plantwide optimization and control solution gives tighter control and also reduces quality giveaway; the product purity specification is met without incurring the expense of further reducing ethane content.

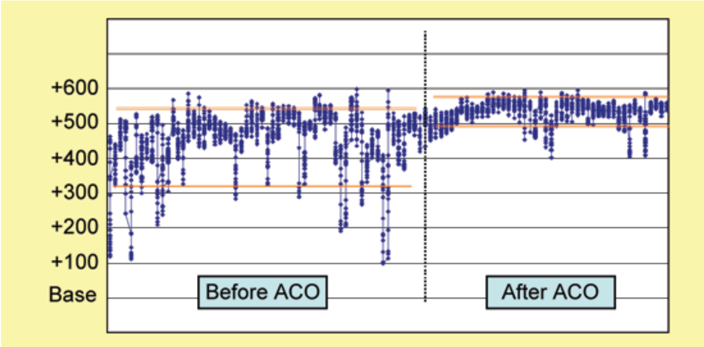

Awards

American Automatic Control Council, 2010, Control Engineering Practice Award for "Innovation in advanced control and optimization with sustained impact on the process industries."

***Control Engineering Magazine*, 1999, Editor's Choice Award for RMPCT (now Profit Controller).**

Broad Process Industry Impact

The sequentially linear, dynamical model predictive control and optimization solution showcased here has been applied since 1995 to more than 10 industries, such as refining, petrochemical, oil and gas, coal gasification, LNG and LPG, pulp and paper, polymer, and aluminum.

Production increases valued at $1.5–$3 million annually are typical for ethylene plants. Energy savings are an additional and significant benefit.

For more information, visit www.honeywellprocess.com/software and search for "Advanced Control and Optimization."

H-infinity Control for European Telecommunication Satellites

Artist's rendering of Eutelsat W2A satellite in orbit, based on the Spacebus 4000 C4 platform, with deployed solar arrays and 12-m-diameter antenna (Source: Thales Alenia Space)

In the early 1990s, European space industries initiated research in robust control, specifically H-infinity (H_∞), through close collaboration with external control laboratories. The research was motivated not only by the desire to gain experience in this new method, but also to evaluate its potential benefits, performance improvements, and development costs when compared to traditional (proportional-integral-derivative and linear-quadratic-Gaussian) controllers commonly used in the 1980s. The increasing performance and dynamic complexity of future space applications were also a source of motivation for the research program on robust control techniques.

The transfer of robust control techniques from research laboratories to industrial space applications covered not only the technique itself, but also the process-oriented engineering tools and methodologies required for modeling, design, and analysis of robust H_∞ controllers.

Telecommunication Satellite Control System Design: Challenges and Needs

Geostationary telecommunication satellite platforms typically consist of a central body and large (deployable) antennae together with low-damping flexible solar arrays that are rotating with respect to the Earth-pointing central body at a rate of one rotation per day. During orbit inclination correction maneuvers, the satellite is submitted to thruster-induced disturbance torques that require some few tens of nanometers control authority to limit the attitude depointing below 0.1 deg. Because of the low damping (typically 10^{-3}) and shifting frequency modes with high resonant peaks of the large rotating solar arrays (see figure below), a stiff filtering controller is required. Using classical control design techniques, the design problem is solved in an ad hoc fashion requiring skilled engineers to initiate the lengthy iterative design procedures, tune the convergence control parameters, and balance the multi-objective performance index.

The limited capability of the classical design procedures to adapt to other space control problems prompted the need to develop automated control design techniques, including systematic procedures to rapidly adapt to changes in dynamic models, to rapidly optimize performance under constraints of parameter uncertainties, and to address "flexible structure control" formulations in the frequency domain. From an industrial perspective, there was also a need to improve European system integrators' competitiveness within the global space market by reducing the overall telecommunication satellite development time. H_∞ techniques enabled fulfillment of these requirements.

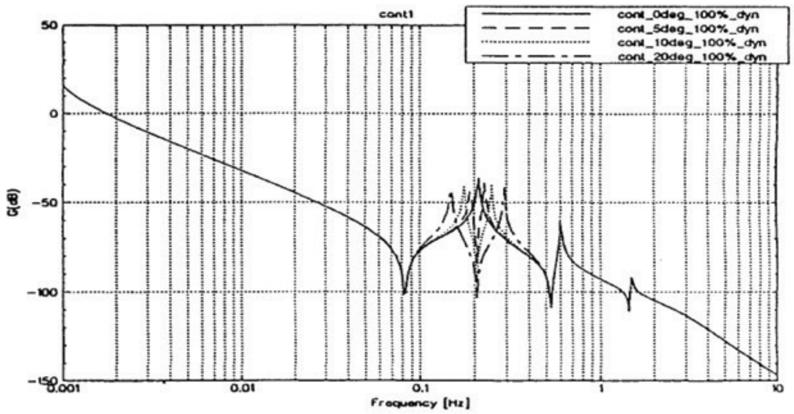

Structural flexible modes (above) of the Astrium communication satellite platform Eurostar 2000+ (left). The shifting of frequency modes corresponds to different angular positions of the solar arrays. (Source: EADS)

Contributor: Christian Philippe, European Space Agency, Netherlands

Standard Four-Block H_∞ Scheme

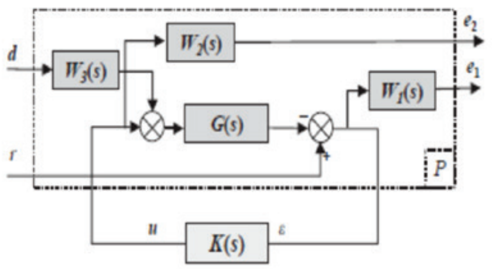

Standard four-block H_∞ scheme (Source: Thales Alenia Space)

The standard four-block H_∞ problem corresponds to the scheme shown on the left, where e_1 and e_2 are signals to be controlled with respect to reference input r and disturbance input d. The closed-loop control objectives are attained through an appropriate tuning of the weighting filters $W_i(s)$ chosen to shape the four transfer functions from r and d inputs to ε (control error) and u (command) outputs. The weighting filters are tuned to ensure good reference tracking and disturbance rejection as well as to meet the desired control bandwidth and rolloff attenuation of the flexible structural modes. The H_∞ problem consists of finding the controller K(s) that fulfills the four main control design objectives:

1. Guaranteeing stability margins
2. Filtering flexible structural modes from solar arrays and/or antennas and propellant sloshing modes
3. Minimizing the attitude degradation in the presence of external and internal disturbances
4. Minimizing sensor noise transmission and fuel consumption

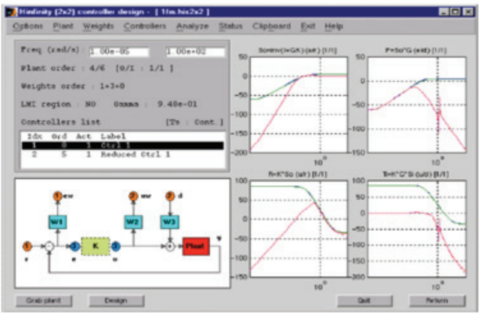

Ready-to-use engineering software tool based on MATLAB from MathWorks (Source: Thales Alenia Space)

The standard four-block H_∞ scheme benefits from attractive numerical and physical properties and analytically guarantees the stability margin and robustness. Although the scheme offers an all-in-one design procedure, it must be used in association with order-reduction techniques to obtain the final controller. The all-in-one approach prevents the designer from having to do repeated iterations between preprocessing, optimal design, and post-processing, as experienced in the classical control design procedure. The all-in-one control design procedure has been implemented in a ready-to-use engineering software tool based on MATLAB from MathWorks and in dedicated control toolboxes (see figure at left).

Benefits

The development of H_∞ controllers for European telecommunication satellite platforms such as Eurostar 3000 and Spacebus 4000 has helped reduce the duration of the orbit inclination correction maneuver by 50%, equivalent to a propellant mass savings of about 20% when compared to the classical control design technique based on proportional-integral-derivative control combined with specific filters in the flexible-modes frequency range.

Other Real-World Applications

Ariane 5 launcher (top) and SILEX (bottom) (Source: EADS Astrium)

H_∞ controllers have also been developed, implemented, and successfully flown on the Ariane 5 Evolution launcher, the Automated Transfer Vehicle (ATV), and earth observation, scientific, and exploration satellites, as well as pointing, acquisition, and tracking (PAT) systems. The benefits of applying H_∞ control techniques for Ariane 5 and the first European optical communication terminal in orbit (SILEX) are summarized below:

Ariane 5 Evolution Launcher

- Issue: structural bending and sloshing modes
- H_∞ controller synthesis (atmospheric phase)
- Benefit: thrust vector control actuation effort minimized

Optical Laser Terminal (PAT)

- Issue: performance limitations of traditional controller
- H_∞ controller synthesis (tracking mode at 20 Hz)
- Benefit: pointing stability performance of 0.1 μrad

For more information: B. Frapard and C. Champetier, H_∞ techniques: From research to industrial applications, Proc. 3rd ESA International Conference, Noordwijk, Netherlands, November 26-29, 1996; G. Pignie, Ariane 5 and Ariane 5 evolution GN&C overview, 34th COSPAR Scientific Assembly, The Second World Space Congress, Houston, TX, October 10-19, 2002; C. Charbonnel, H_∞ and LMI attitude control design: Towards performances and robustness enhancements, Acta Astronautica, vol. 54, pp. 307–314, 2004.

Improved Audio Reproduction with Control Theory

To the discerning listener, sound quality in CDs and other standard digital formats leaves something to be desired. It turns out that this "something" is related to high-frequency signal elements. A new advance based on sampled-data control theory, the YY filter, has overcome this problem with audible advantages!

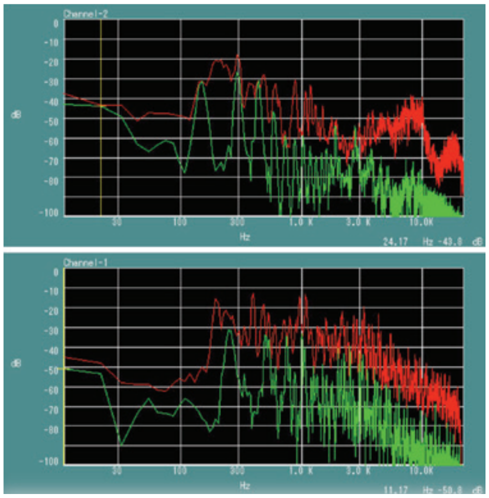

Fast Fourier transforms (0–20 kHz) of an analog record (top) and CD reproduction (bottom). The green trace is the FFT; the red trace is the peak FFT value over the past 10 seconds. The record exhibits a range that extends well beyond 20 kHz; the CD has a sharp cutoff at 20 kHz. (The traces are not from the same sound, but the lower figure shows that the frequency components are sharply cut at 20 kHz, in contrast to the analog source above.)

Conventional Sound Encoding

The audible range is widely accepted to be limited to 0-20 kHz, and anything beyond is sharply cut (filtered out) by a low-pass filter. This is based on the well-known Whittaker-Shannon sampling theorem; all frequencies beyond the Nyquist cutoff are regarded as noise. However, the Shannon formula is noncausal and hence not directly applicable to sound reconstruction/recovery.

In addition, the high-end frequency (so-called Nyquist frequency) of 22.05 kHz (half of the sampling frequency used in digital audio) may not provide sufficient margin against the audible range. Digital filters used today usually cut the frequency components beyond 20 kHz very sharply. But this has the side effect of inducing (1) a large amount of phase distortion (phase error is not considered in the conventional Shannon paradigm); and (2) ringing around 20 kHz due to the sharp-cut characteristic of the filter (Gibbs phenomenon). The latter induces a very "aggressive," sharp, and metallic sound that is likely the main reason for the audiophile's complaints about CD recordings. Undesirable distortions intrude below the Nyquist frequency too.

An Application for Sampled-Data Control Theory

High-frequency components are intersample signals. This observation suggests that modern sampled-data control theory can offer solutions to the problems of sound processing today. Based on recent theoretical results, filters can be designed that optimally interpolate the intersample content (that is, the lost high-frequency components). This is the intuition behind the YY filter, which is able to recover the optimal continuous-time (analog) performance.

Commercial Example

The figure below is an example from a mini-disc (MD, format similar to MP3) player. The horizontal axis is the audio frequency on a linear scale of 0–22 kHz. The top graph shows the frequency response with the MDLP4 standard at 66 kbps. The bottom graph shows the response at the same bit rate with the YY filter implemented. The improved high-frequency response is evident.

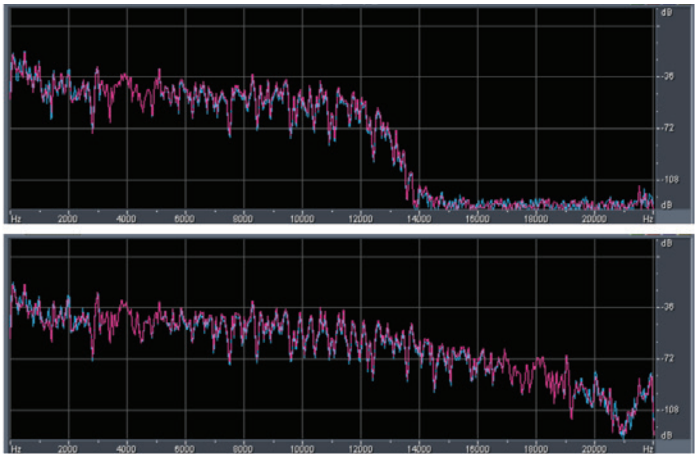

Source: Sanyo Corporation

Contributor: Yutaka Yamamoto, Kyoto University, Japan

From: *The Impact of Control Technology,* 2nd ed., T. Samad and A.M. Annaswamy (eds.), 2014. Available at www.ieeecss.org.

By the Millions!

The YY filter has been implemented in integrated circuits produced by ON Semiconductor for expanding the effective range in such devices as CDs, MP3s, mobile phones, digital voice recorders, and car audio systems. The sound quality has proven superior to the original according to the PEAQ (Perceptual Evaluation of Audio Quality) index; the filter enhances the quality by almost 30% for MP3 128 kbps and by over 30% for advanced audio coding (AAC) on average. Cumulative production has reached 40 million chips during the period 2005–2012.

The YY filter can be employed as an iPhone/iPod app. A new music app is being marketed under the name "AiSound." Unlike the default music player supplied by Apple, this player extends the MP3 and AAC frequency range of 16 kHz up to 20 kHz via the YY filter and produces crisper, sharper, and more natural sounds.

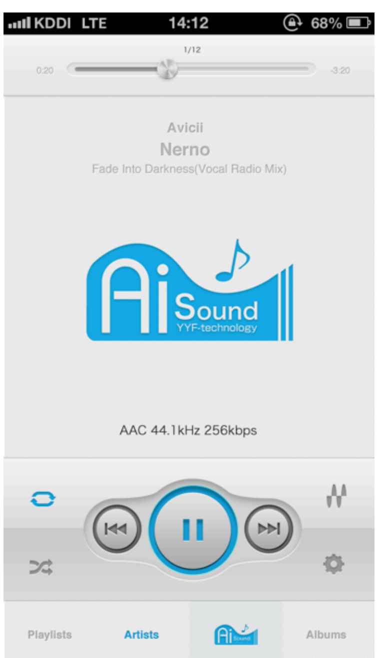

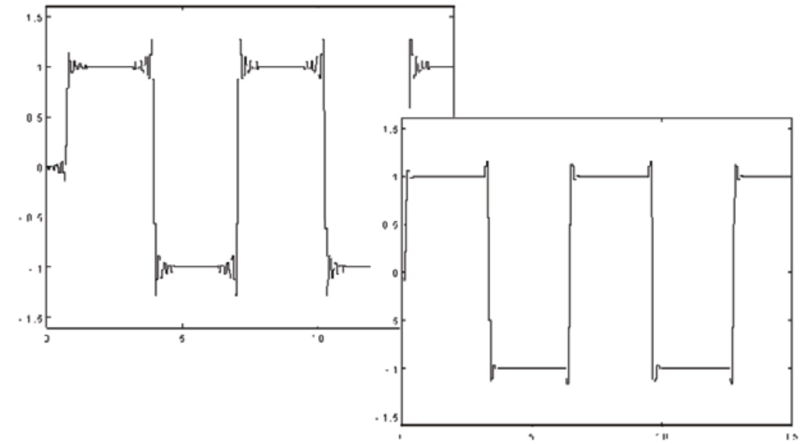

The effect of the YY filter is particularly evident in these reconstructions of a sampled square wave. Conventional reconstruction results in significant high-frequency distortion (the "ringing" observed at the corners of the signal). The YY filter substantially reduces the distortion.

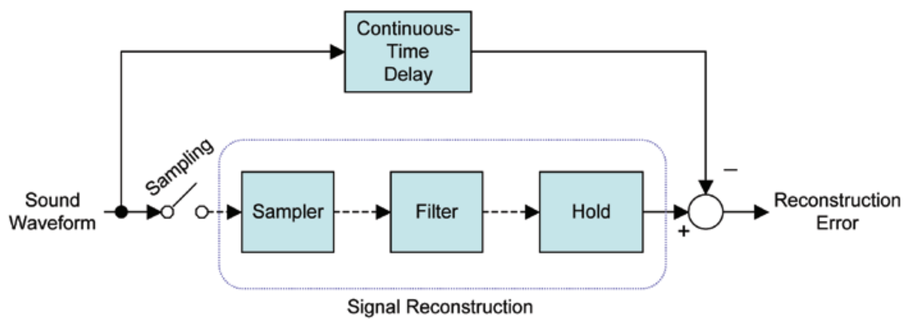

The YY filter design process mathematically optimizes the filter to ensure that the reconstruction error across a desired frequency range—not determined solely by the Nyquist frequency—is less than a design parameter. This is a sampled-data H_∞ control problem.

Applications to Image Processing

The YY filter can be applied to still and moving images as well, as illustrated here. Left image: interpolation via the bicubic filter; right image: interpolation with the YY filter. Visit http://www-ics.acs.i.kyoto-u.ac.jp/~yy/image.html for high-resolution images and http://www-ics.acs.i.kyoto-u.ac.jp/~yy/mimage.html for videos.

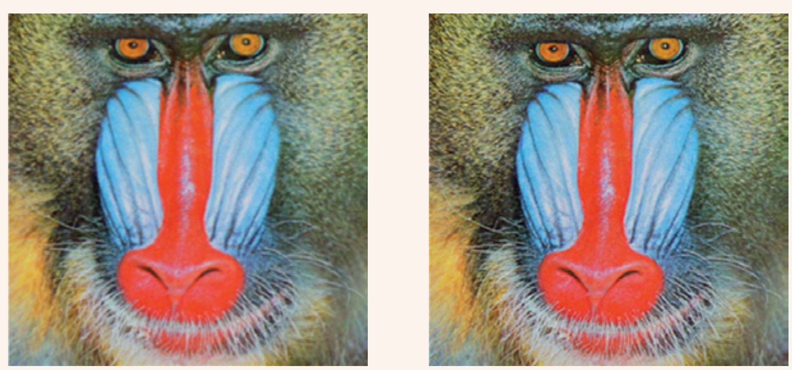

For more information: Can you hear the difference? Visit http://www-ics.acs.i.kyoto-u.ac.jp/~yy/sound.html.

Mobile-Robot-Enabled Smart Warehouses

Order fulfillment is a multibillion-dollar business. Existing solutions range from the highly automated—whose cost-effectiveness is inversely related to their flexibility—to people pushing carts around warehouses manually filling orders, which is very flexible but not very cost-effective. Kiva Systems uses a new approach to order fulfillment in which operators stand still while the products come to them. Pallets, cases, and orders are stored on inventory pods that are picked up and moved by hundreds of mobile robotic drive units. As a result, any product can go to any operator.

Successful Installations Worldwide

Kiva Systems has deployed dozens of installations worldwide, including a 1,000-mobile robot system for a retail company in the United States. Customers include:

- Crate and Barrel
- Diapers.com
- Dillard's
- Gap
- Gilt Groupe
- Office Depot
- Saks Fifth Avenue
- Staples
- Timberland
- Toys "R" Us
- Von Maur
- Walgreens

Contributor: Raffaello D'Andrea, ETH Zurich, Switzerland

Founders Mick Mountz, Peter Wurman, and Raffaello D'Andrea received the 2008 IEEE/IFR Invention and Entrepreneurship Award in Robotics and Automation, whose aim is to foster innovation paired with entrepreneurial spirit, and to promote the best possible use of synergies between science and industry in the fields of robotics and automation.

Kiva Systems was purchased by Amazon.com in 2012 for $775 million.

System Description

Kiva uses hundreds of mobile robots and powerful control software to provide a complete fulfillment solution: storing, moving, and sorting inventory. Instead of being stored in static shelving, flow racks, or carousels, products are stored in inventory pods in the center of the warehouse while operators stand at inventory stations around the perimeter.

- When an order is received, robotic drive units retrieve the appropriate pods and bring them to the worker, who selects the correct items and places them in the carton. Completed orders are stored on separate pods, ready to lift up and move to the loading dock when the truck arrives.
- The Kiva drive units are differential-drive, two-wheeled robots with a patent-pending mechanism for lifting pods off the ground. This mechanism is essentially a large actuated screw—by rotating a drive unit underneath a pod and simultaneously counterrotating the screw, a pod may be lifted off the ground.
- A suite of sensors on the drive units and custom control software and algorithms allow the vehicles to safely navigate the warehouse. Coordination is aided by a hierarchical layer similar to that used in air traffic control systems.
- The drive units share information about their environment and use that knowledge to adapt. As a result, the performance of the vehicles, and hence the system, improves over time. In addition, adaptation and learning ensure that the system is robust to changes in the environment.

Select Customer Quotes

"Our customers expect to get great value and service from Crate and Barrel, but they also care about our carbon footprint. This played a role in our selection of Kiva Systems," said John Ling, vice president of supply chain management and logistics at Crate and Barrel. "Kiva's mobile robotic approach is not only the most cost-effective way to automate pick, pack, and ship operations, but also the greenest. The robots themselves are energy efficient, plus the entire robot zone can be operated with almost no lighting."

"Using a flexible, automated order fulfillment system helped our Piperlime operations scale to increased capacity over the critical holiday season," said Chris Black, vice president of operations at Gap Inc. Direct. "The system freed up our employees' time, allowing them to focus on processing a higher volume of customer orders faster and to ensure more accuracy. We're looking forward to leveraging Kiva's system when we expand our online business internationally."

"We have been implementing warehousing technology for over twenty years, and Kiva has the only practical automated solution that supports the diversity of pick-pack-ship operations we run," said Bruce Welty, founder and CEO of Quiet Logistics. "Kiva is the first automation approach we'd be willing to bet our business on. Kiva's innovation immediately cuts hard costs, increases throughput, and improves service levels to give us and, more importantly, our customers a tremendous competitive advantage."

For more information, visit www.kivasystems.com.

Nonlinear Multivariable Flight Control

F-35

Under contracts from NASA and the Air Force Research Laboratory, Honeywell and Lockheed Martin developed and documented a novel approach to flight control design. This multivariable flight control methodology based on nonlinear dynamic inversion was applied to the experimental X-35 military aircraft and the X-38 Crew Return Vehicle program and is now in production for the Lockheed Martin F-35 aircraft.

This work contributed to the advancement of advanced control technology and resulted in the control law being recognized as a design option for flight control development.

Key Innovations

Nonlinear dynamic inversion control is a systematic generalized approach for flight control. Using general aircraft nonlinear equations of motion and onboard aerodynamic, mass properties, and engine models specific to the vehicle, a relationship between control effectors and desired aircraft motion is formulated. A control combination is designed that provides a predictable response to a commanded trajectory. Control loops shape the response as desired and provide robustness to modeling errors. Once the control law is designed, it can be used on a similar class of vehicle with only an update to the vehicle-specific onboard models.

Specific innovations include:

- The dynamic inversion control law
- A control allocation procedure
- An onboard aircraft model (OBAC)

Previous Practice

Nonlinear dynamic inversion is the first systematic approach to nonlinear flight control. Prior to this development, the control law was typically designed from a set of linear plant models and implemented with a gain-scheduled linear controller. The performance capabilities of the aircraft were not fully realized, and the manually intensive development process was time consuming.

Onboard Aircraft Model (OBAC)

The NASA/Honeywell/Lockheed Martin flight control approach includes the first use of an aircraft model in the control law. This model is used to derive coefficients for dynamic inversion and control allocation computations. Changes in vehicle structure during design often only require changing the OBAC model for the controller.

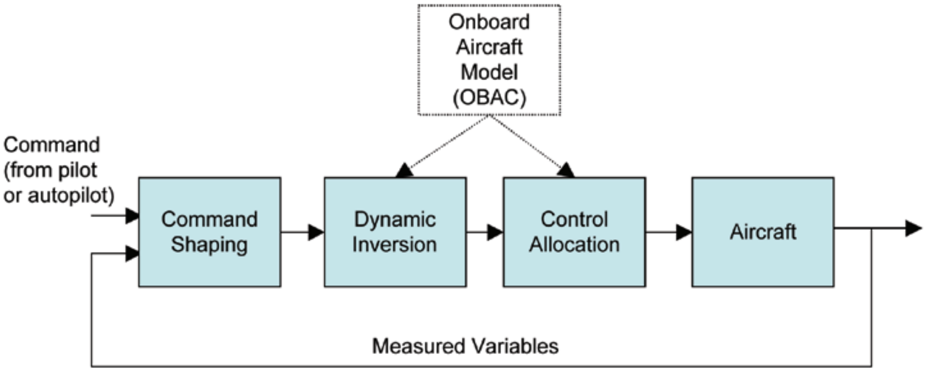

Contributors: John Bosworth, NASA, USA; Dale Enns, Honeywell, USA

F-18

X-38

Nonlinear Dynamic Inversion

Dynamic inversion "inverts" the aircraft model to identify what roll, pitch, and yaw moments will give the desired aircraft trajectory. Sophisticated numerical algorithms are used to ensure rapid computation and to deal with actuation rate and deflection limits. The inversion is calculated onboard at each control loop iteration during flight.

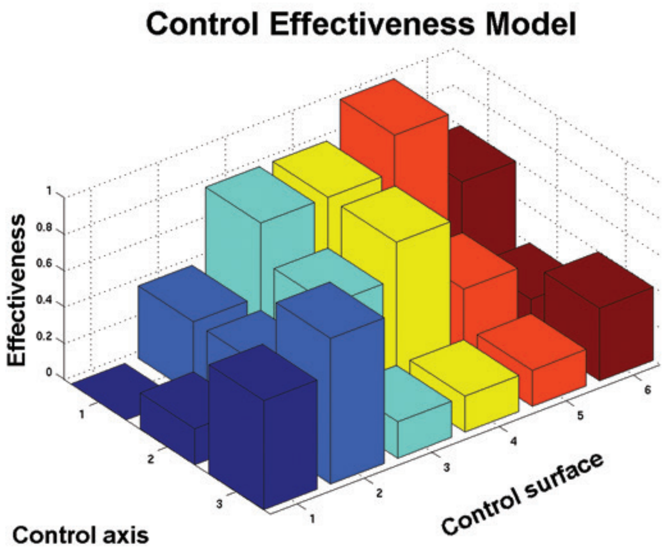

Control Allocation

Modern aircraft have redundant actuation capabilities—the same airframe response can be achieved with different combinations of actuators. The problem of determining which actuators to use, and to what extent, at a given instant is referred to as control allocation. The nonlinear dynamic inversion methodology computes an optimal control allocation, taking into account saturation constraints on actuators. The actuators of interest for the applications developed are the control surfaces (ailerons, elevators, rudder) and thrust vectoring (directing the engine thrust). A control effectiveness model is used in the computation (see figure above).

Program History

The work described in this success story began in the mid-1980s as a theoretical development for an airplane designed for high-angle-of-attack maneuverability. A Honeywell nonlinear dynamic inversion design was selected as the controller for the F-18 research vehicle and implemented in a full hardware-in-the-loop piloted simulation.

In the late 1990s, NASA began the X-38 Space Station Crew Return Vehicle program. The nonlinear dynamic inversion controller was proposed for this program and allowed control updates as the vehicle's structural design changed by simply updating the OBAC model.

These and other foundational projects led to the collaboration between Honeywell and Lockheed Martin, as well as the implementation of nonlinear dynamic inversion on the X-35 prototype and eventually the production F-35 vehicle, the latest state-of-the-art military aircraft. The controller has provided consistent, predictable control through the transition from conventional flight to hover and has also enabled a 4X to 8X reduction in nonrecurring engineering development cost.

For more information: Honeywell and Lockheed Martin, Multivariable Control Design Guidelines, Final Report, WL-TR-96-3099, Wright Patterson AFB, Ohio, USA, 1996.

Optimal Ship-Unloading Solutions

Large quantities of bulk grain are transported in ships. These ships can hold different kinds of grain in up to a dozen hatches and could be damaged in the event of excessive ship tilt or wall stress between hatches.

The port receiving the grain employs its available configuration of unloaders, bulldozers, and conveyors to unload the ship as quickly as possible. Up to four different unloaders are available that move laterally but cannot cross each other. At low grain levels, bulldozers hoisted into the hatch shovel the grain toward the take-up point of the unloader. Up to four conveyors transfer grain from the unloaders to silos.

All this complexity must be managed for unloading grain-laden ships at minimum time and cost without exceeding safety and other constraints. Bühler AG has developed a decision support tool that is used by both sales and technology staff to optimize unloading.

A five-hatch ship with a single unloader digging in hatch 4

An unloader using its take-up point to dig into the grain in a hatch

Two bulldozers feeding residual grain to the take-up point of an unloader in a hatch

Contributors: Robert Haymoz and Mukul Agarwal, Bühler AG, Switzerland

Decisions: Investment and Unloading Sequence

Port personnel must use their experience and know-how to decide the unloading sequence for an incoming ship. They also must decide which changes in (or which new) unloader-bulldozer-conveyor configuration they should invest in. For this purpose, they need to know the expected performance of a projected configuration for a reference ship.

Challenges for a Decision Support Tool

A tool to support the above decisions must address several challenges:

- Find a sequence to minimize unloading time.
- Compute ship tilt and wall stresses dynamically and respect their limitations.
- Account for the three unloading phases in a hatch, with decreasing unloading rate (see figure on next page) and dynamic transition criteria: free digging, maneuvering the take-up point, and residual unloading using bulldozers to shovel grain to feed the take-up point.
- Respect limitations on receiving conveyors.
- Ensure that unloaders do not cross each other.
- Allow two unloaders to service the same hatch together.
- Include unloader travel time between hatches.
- Use different unloading rates for different grains and hatch geometries.
- Enable a variety of ships, grains, unloaders, bulldozers, and conveyors.
- Minimize computation time for interactive use.
- Animate results for verification and credibility.

The tool should enable users to determine the unloading time and sequence for an incoming or reference ship. It should also be able to optimize the configuration of unloaders (number, types, and positions), bulldozers, and conveyors and allow tradeoff between unloading time and investment cost.

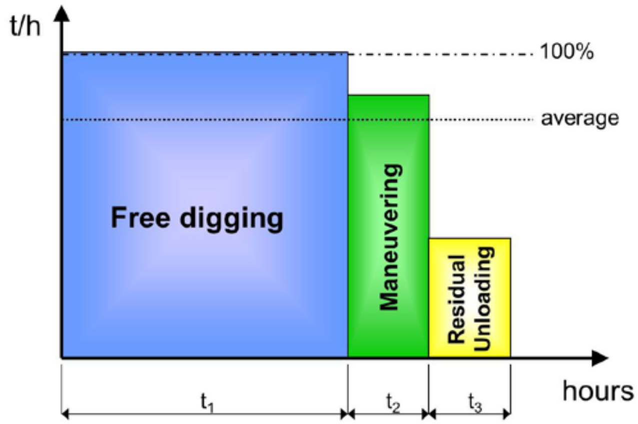

A typical pattern of unloading rate (tons per hour) versus active unloading time for a hatch

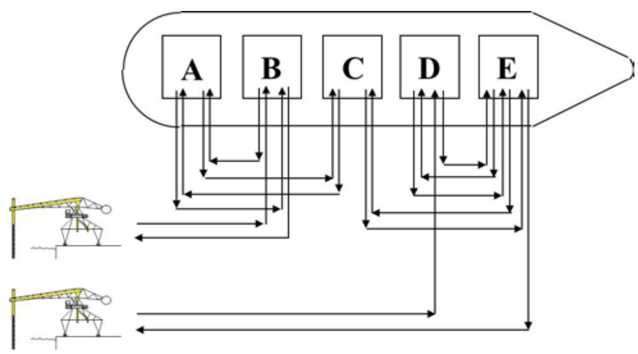

A possible unloading sequence using two unloaders for five grain hatches

Solution Strategy

For a given ship and unloader-bulldozer-conveyor configuration, the tool simulates the entire unloading while making event-triggered decisions based on evaluations of multiple heuristic criteria. It also iteratively optimizes the relative weights of the criteria.

In addition to this "inner-loop" process, an outer genetic-algorithm-based loop optimizes the unloader-bulldozer-conveyor configuration. Semiautomated operation is also possible—for example, if an investment-performance tradeoff is not given, the tool can optimize just the unloader positions, allowing the user to manually, and iteratively, explore options for the rest of the configuration.

Application of the Ship-Unloading Tool

Since 2008, sales and technology personnel at Bühler AG have used the tool routinely for two purposes:

1. About 50 times a year, sales personnel use the tool to determine the performances of manually proposed unloader-bulldozer-conveyor configurations so as to select and convincingly present the best options to port customers worldwide. The tool enables sales personnel to do each job in half to one day instead of a week and to guarantee performance with lower risk.

2. Technologists use the tool to advise customers how to modify their manual unloading strategy. The optimization features of the tool enable up to 30 percent faster unloading of ships.

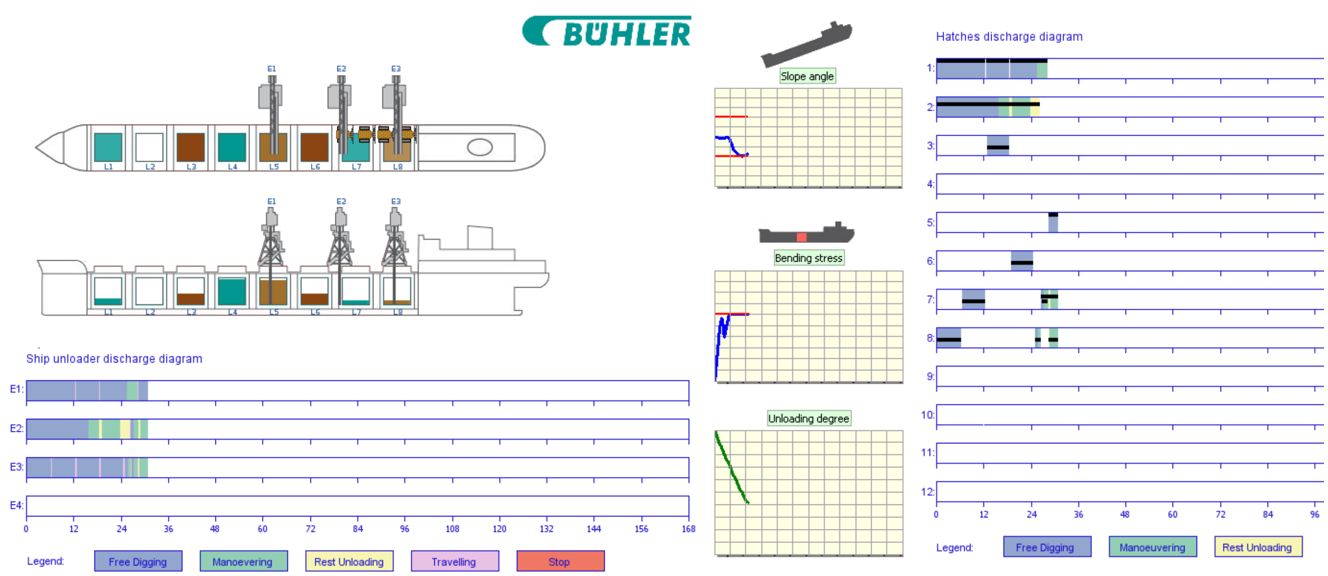

For a real or projected unloader-bulldozer configuration, the tool uses event-triggered heuristic-criteria-based decisions to unload a simulated ship in the shortest time while observing limits on ship tilt and wall stress. The tool outputs the best unloading procedure and the attainable performance (total unloading time).

For more information, visit http://www.buhlergroup.com.

Road Grade Estimation for Advanced Driver Assistance Systems

Modern vehicles are equipped with many assistance control systems that aid the driver in operating the vehicle safely and economically. Several of these systems can be programmed to use stored information from a map, enabling actions to be taken before the driver can even see the road in question. Knowledge of the current and future road grade can be used in engine and gearbox control systems to help meet instantaneous power demand while keeping fuel consumption and environmental impact as low as possible. Heavy-duty vehicles are especially affected by the road grade, and by using information about the future gradient, the energy efficiency of the cruise controller can be greatly increased. On downhill road segments, the vehicle's brakes often need to be used to maintain a safe speed. By automatically coasting over the top of the hill, braking can be minimized and fuel conserved. When going uphill, the number of gear shifts required can often be reduced if the vehicle speeds up before the hill and the gearbox has information about the coming road grade. A map with road grade information can thus reduce the fuel consumption and environmental impact of heavy-duty vehicles. These ideas have been implemented in production systems for Scania trucks and are deployed all over Europe.

Vehicle Fleets as Sensor Networks

As long-distance trucks go about their daily business of delivering cargo, onboard algorithms can develop estimates of the road grades they encounter on their routes. This is much more efficient than the conventional approach of creating digital road maps with specialized probe vehicles that drive around specifically for the purpose of collecting data. In this new approach, only sensors already available in current heavy-duty vehicles are used to estimate road grade information. This effectively turns every truck on the road into a mobile sensor node, and the entire fleet of vehicles forms a constantly moving sensor network. The road grade information is stored in a map and updated each time a vehicle passes over a road segment. The estimation is performed using a Kalman filter that combines information from onboard vehicle sensors and the current driving situation (see figure on the next page). An estimate of the road grade accuracy is also stored in the map. When a vehicle returns to an already mapped road segment, the accuracy is used to determine how much the new data should be trusted relative to the data already stored in the map. If each vehicle is equipped with a communication device, the information gathered can be shared with other vehicles through a server. This way a vehicle gains access to data about roads it has not yet driven. With more measurements of the same road segment, the statistical uncertainty decreases, giving better road grade information and higher performance in the control systems using the estimate.

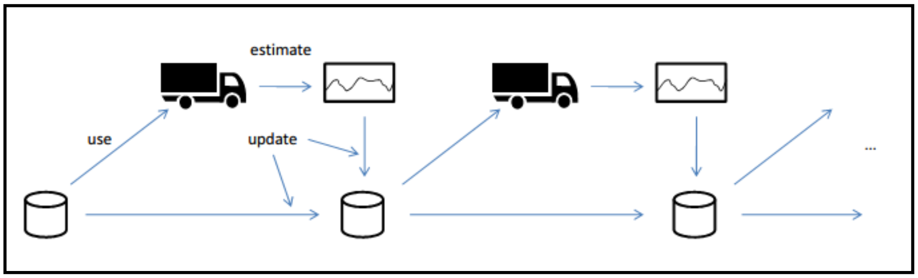

Each vehicle uses the current map, creates a new estimate, and updates the stored map (either on board or on a remote server).

***Contributors:** Per Sahlholm, Scania CV AB, Sweden, and Karl Henrik Johansson, KTH Royal Institute of Technology, Sweden*

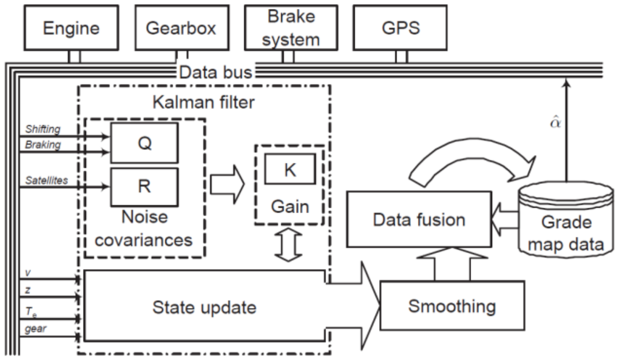

A Kalman filter is used to estimate the road grade based on vehicle sensors and the current driving situation. The estimates are combined with previous estimates stored in a grade map database.

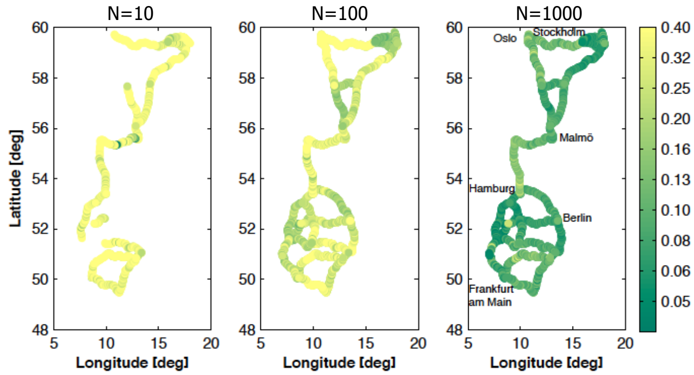

Estimated road network and estimation error (in percent road grade) after N = 10, 100, and 1000 vehicle trips in a region

Vehicle Sensor Network Implementation

Advanced driver assistance systems using mapped road data are already in production and being used by customers on highways. Examples include the Scania Active Prediction and Eco-roll systems. Using road grade information collected by a network of vehicles in daily operation instead of the current manually collected information for these applications would ensure that all relevant roads are covered and that changes in the road network are quickly incorporated into the map. Data collected with traditional probe vehicle methods soon become outdated, and coverage is lacking in many countries.

Scania Active Prediction in use. The cruise controller adjusts the vehicle speed within a range around the desired average speed in order to traverse hills more efficiently.

Future Applications

Many systems yet to be developed may benefit from road grade estimates. Examples include heavy-duty vehicle navigation systems that choose the most energy-efficient route with respect to topography; vehicle platooning systems that automate driving, maintaining short distances between vehicles in hilly terrain; and emergency braking systems that take the road grade into account when calculating the stopping distance. The described approach may be expanded into additional types of data. Examples include road curvature, road friction, and vehicle speed at various times of the day. All these variables are useful in developing the next generation of advanced driver assistance systems for heavy-duty vehicles, and they may all be gathered using a mobile sensor network made up of trucks moving about our highways every day.

For more information: P. Sahlholm, Distributed Road Grade Estimation for Heavy Duty Vehicles, doctoral dissertation, KTH Royal Institute of Technology, Stockholm, 2011; P. Sahlholm and K.H. Johansson, Road grade estimation for look-ahead vehicle control using multiple measurement runs, Control Engineering Practice, vol. 18, no. 11, pp. 1328-1341, 2010.

Robust Adaptive Control for the Joint Direct Attack Munition

Control theory has been the enabling technology in achieving man's dominance over flight. Early experimental aircraft were difficult to control, had limited flight envelopes and flight times, and the pilots had to exercise control over the aircraft's trajectory using mechanical systems. In these early aircraft, the pilots had to adapt to changing environmental conditions and/or failures of any aircraft components. Over several decades, propulsion systems matured, our understanding of flight dynamics and aerodynamics grew, and computers and digital fly-by-wire systems were developed, all of which have helped bring automation to flight control.

With recent advances in control theory, particularly in the area of robust and adaptive control, fully automatic flight is now possible even for high-performance air systems. Among the first application successes of this new technology has been its technical transition to guided munitions, in particular, the Joint Direct Attack Munition (JDAM) system.

Robust Adaptive Control

Two techniques have been developed by control engineers and scientists to accommodate uncertainty in our knowledge of the system we are trying to control. The two techniques are complementary and have been combined to create robust adaptive controllers.

Robust control: Based on a mathematical model of the uncertainty, a formal design procedure is used to develop closed-loop controllers that will provide optimized performance and ensure stability over the range of uncertainty.

Adaptive control: Instead of developing a fixed controller over a space of model uncertainty, adaptive control adjusts the controller online based on detections of plant deviations from a reference model. Adaptive control augments and further extends the performance and robustness of the flight control system.

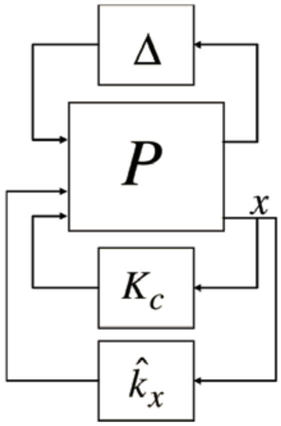

Shown at right is a control engineer's block-diagram representation of robust adaptive control. The nominal plant model P of the system under control (such as a missile) is subject to uncertainties Δ. The baseline flight controller K_c, designed using robust control techniques, is augmented with an adaptive controller. The state vector x is the input to both the baseline and adaptive controllers. The combination provides robust stability and performance over a substantially enhanced space of modeling uncertainties and can accommodate changes in the system under control.

JDAM

The Joint Direct Attack Munition is a guidance kit that converts unguided bombs into all-weather "smart" munitions. JDAM-equipped bombs are guided by an integrated inertial guidance system coupled to a global positioning system (GPS) receiver, giving them a published range of up to 15 nautical miles (28 km). The guidance system was developed jointly by the United States Air Force and the United States Navy. The JDAM was meant to improve upon laser-guided bomb and imaging infrared technology, which can be hindered by potentially bad ground and weather conditions.

Contributors: Kevin A. Wise and Eugene Lavretsky, Boeing, USA

Improving Weapon Control and Effectiveness

The U.S. Air Force Office of Scientific Research (AFOSR) sponsored researchers at The Boeing Company to develop and transition new robust adaptive control algorithms for application in the Joint Direct Attack Munition. The first transition has been to the 500-lb MK-82 JDAM.

Affordability and weapon accuracy (including collateral damage minimization) are among the primary objectives for the JDAM. The new robust adaptive control algorithms provide accurate control of the weapon, accommodating warhead aerodynamic uncertainties and off-nominal mass properties. Without control modifications, these uncertainties can significantly degrade weapon accuracy.

Photo sequence showing JDAM test against mobile ground targets

Other Applications

Other air systems are also prime candidates for robust adaptive control technology. Recently, another variation of the JDAM system has been developed and transitioned into production: a new dual-mode laser-guided JDAM system (L-JDAM) for detecting and prosecuting laser-designated targets (moving or fixed). For the L-JDAM development, the adaptive controller augments the baseline flight control system and only engages if the weapon begins to deviate from nominal behavior. This augmentation approach allowed The Boeing Company to develop and test the new laser variant without expensive wind tunnel testing, reducing development costs and schedule. The hardware modifications to create the L-JDAM weapon included the addition of a sensor nose kit (the sensor fit into the existing fuse well), wire harnesses, straps with barrel nuts, and symmetric tail kit dog ears where the sensor wire harness enters into the tail kit.

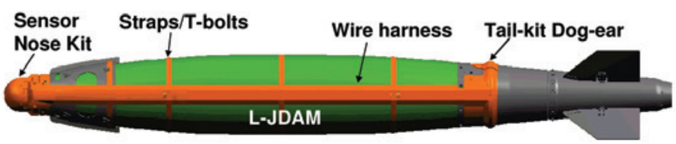

Laser-guided MK-82 scores direct hit against a moving target during tests at Eglin AFB.

Affordable hit-to-kill accuracy minimizes collateral damage; the photograph shows the hole made in the target by a (nonexplosive) weapon.

Trip Optimizer for Railroads

On-time arrival with the least fuel expenditure is a key priority for freight and passenger railroads worldwide. North American railroads consumed 4 billion gallons of fuel in 2008, 26% of operating costs.

Trip Optimizer is an easy-to-use control system that allows the crew or dispatcher to achieve on-time arrival with the least possible fuel use.

Optimal driving solutions are computed onboard and executed in a closed loop using GPS-based navigation. Train and track parameters are adapted online to reduce model errors. Fuel savings of 4%-17% are realized.

Trip Optimizer Modules

- *Trip Planner* finds the driving strategy (speed and throttle) that minimizes fuel consumption for the target arrival time and satisfies speed limit and other train and locomotive operating constraints.
- *Speed Regulator* closes the loop around the plan to correct for modeling errors and external disturbances and provides compensation for slack action in the distributed dynamics of typical mile-long, heavy trains; both hands-off closed-loop and driver-in-the-loop "coaching" solutions are available.
- *Location and Model Estimator* provides precise location of the train, compensation for GPS dropouts, and adaptively tracks train parameters such as weight, length, and drag.

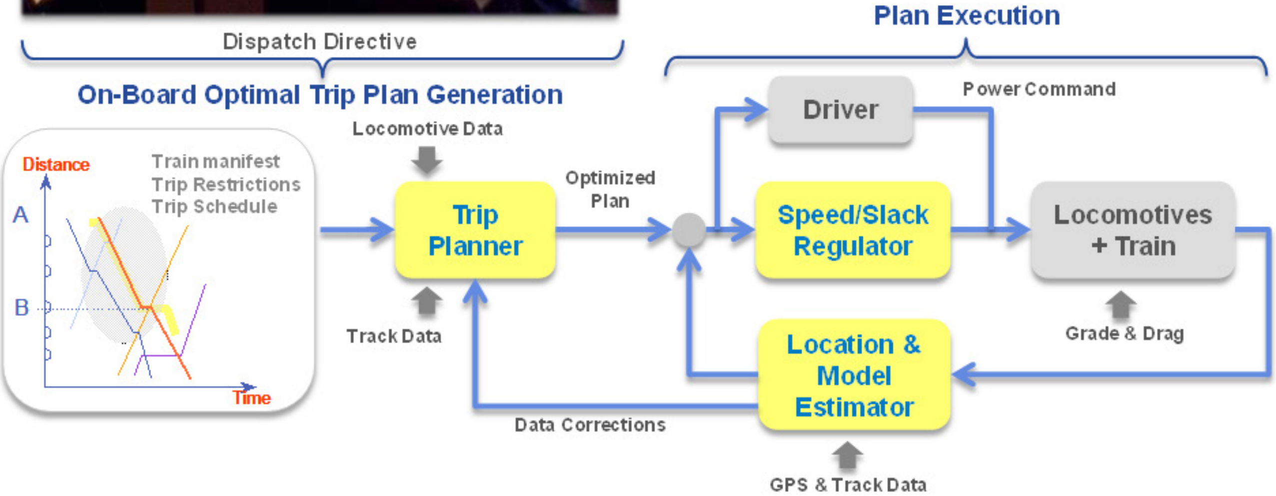

Contributors: David Eldredge, GE Transporation, USA, and Paul Houpt, GE Global Research (retired), USA

Inventions and Innovations

Trip Optimizer provides innovative solutions for optimization, estimation, control, and operator interface requirements to achieve fuel savings and emissions reductions for freight railroads.

- Computation of the driving plan requires solving a math program with thousands of constraints and decision variables in seconds, with time- and fuel-based objectives.
- A robust speed regulator design relies on a loop-shaping algorithm to maintain stable operation and to deal with variations in intercar separation and the resulting forces.
- A location estimator provides precise coordinate tracking via Kalman-filter-based compensation for GPS dropouts. Model-based methods adaptively track key train parameters using GPS and other locomotive data. Tools for extensive offline analysis were also developed to produce high-integrity database sources for use in control and estimation.
- Innovative displays bring intuitive mode awareness and ease of use to the underlying optimal control strategy. Experienced drivers can learn the system in minutes.
- Robust satellite communication from the locomotive provides rapid access to train data (and updates) directly from railroad mainframes with backup from a dedicated 24 x 7 GE facility.

For each Evolution locomotive on which it is used, Trip Optimizer can reduce fuel consumption by 32,000 gallons, cut CO_2 emissions by more than 365 tons, and cut NO_x emissions by 3.7 tons—per locomotive per year. If deployed on the ~10,000 similar locomotives in service in North America, this is equivalent to taking a million passenger cars off the road for a year.

Trip Optimizer is a product of GE Transportation, Erie, Pennsylvania, USA. It has recorded more than 150 million miles in successful revenue service worldwide.

- Adopted by railroads in North America, Australia, and Brazil.
- Total fuel savings to date of over 35 million gallons!

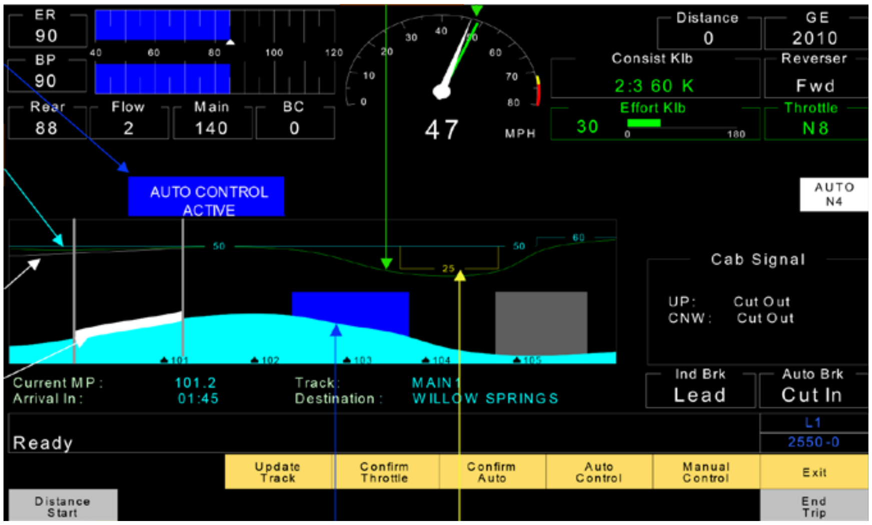

For more information: visit http://www.getransportation.com/.

Verification of Control System Software

Control systems are typically prototyped with graphical design tools such as Simulink; the actual implementation is then obtained by either compilation from these tools or via other high-level languages such as Scade. All of these steps, including early design, may result in bugs slipping into the end product. These bugs may lead to costly product recalls or, in the case of safety-critical systems (aircraft fly-by-wire, medical infusion pumps, safety-critical industrial processes, etc.), to loss of life and limb.

The traditional way to detect bugs in a computer system is through testing: run the programs or components thereof on sample inputs and check for violations of expected system behaviors—not just program crashes, but also functional properties such as actuator constraint violations and inconsistent mode settings. Although coverage criteria for testing typically guarantee that all instructions of the program have been exercised, testing cannot exercise all possible configuration executions on all possible inputs.

The limitations of testing could be overcome if we could *prove* that a program behaves correctly on all possible inputs. This is the goal of formal validation and verification research, which has resulted in practical tools and successes in this area!

The Airbus A380 has advanced fly-by-wire controls implemented in software. A380 software development benefited from static program analysis tools.

From Testing to Proving

Correctness proofs for programs were proposed in the late 1960s by Floyd and Hoare, but the limited technologies available for automating such proofs long confined them to academic examples and idealized versions of crucial algorithms.

A crucial limitation of automated program analysis is that no analysis algorithm can be guaranteed to never give false negatives (failing to point to bugs) or false positives (bugs that cannot occur in reality). This is a basic mathematical result of computability theory. Thus, all automated analysis methods effect a balance between these two kinds of errors/imprecisions.

Within these theoretical limits, though, practical tools can and have been developed and deployed.

From Academia to Industry

The aerospace community's interest in formal methods was renewed in 1996 by the explosion of the maiden flight of the Ariane 5 rocket due to a software bug (an arithmetic overflow). A team of researchers from INRIA was commissioned to design a static analysis system that could detect such kinds of bugs in future. The academic IABC static analyzer was later turned into the PolySpace verifier. (PolySpace, a startup, was later bought by The MathWorks.)

The Airbus A380 was the next major application of static analysis in aerospace. Researchers from École Normale Supérieure of Paris and CNRS developed new analysis techniques for avionics software (e.g., analyzing floating-point computations such as digital filters). The Astrée tool is now marketed through AbsInt GmbH, which also develops the aiT tool for proving bounds on worst-case execution time on modern embedded processors (with pipelined execution units, caches, etc.).

In the United States, the U.S. Food and Drug Administration (FDA) began an initiative in 2010 enforcing the use of static analyzers for programs running infusion pumps; the misbehavior of such programs may result in the death of patients.

Contributors: David Monniaux, VERIMAG, CNRS, France, and Antoine Miné, LIENS, CNRS, France

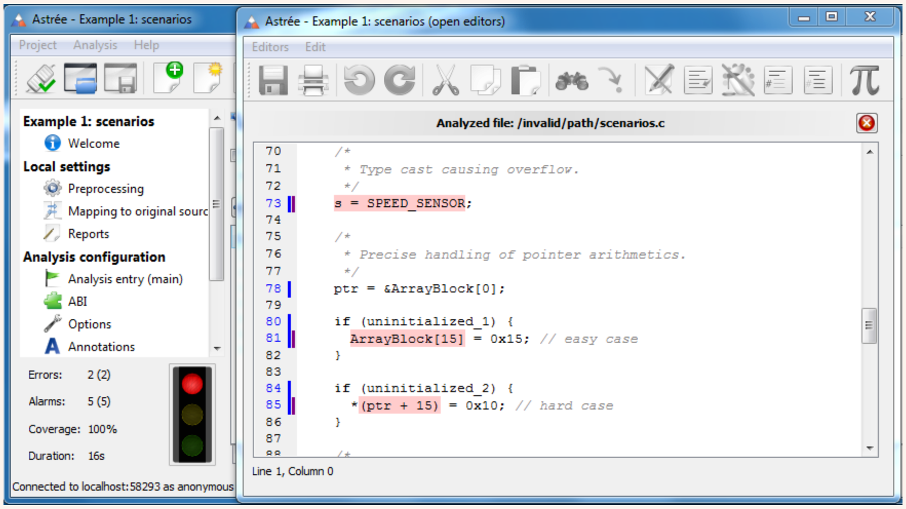

The graphical user interface of the Astrée tool displays program lines that could cause runtime errors and also outputs useful information on the program, such as the range and usage of variables.

The Ariane 5 rocket had to be destroyed on its maiden flight because of a software bug (an arithmetic overflow). This incident renewed interest in formal verification and ultimately resulted in a success story for the technology.

The Astrée Static Analyzer

The Astrée static analyzer takes as input C source code and optional annotations (e.g., range of inputs). After a fully automated analysis, it provides easy-to-understand "traffic-light" indications: a green light for program instructions for which it can prove that no unsafe behavior may occur, a red light for those that it can prove will necessarily result in unsafe behavior if executed, and an orange light for those for which it cannot provide proofs of either safe or unsafe performance.

Some static analysis tools may exhibit false negatives: they may fail to flag possible runtime errors or specification violations. In contrast, Astrée is sound. It performs an exhaustive scan of the control and data space of the program, according to the user-specified inputs and the semantics of the C language (including fine points such as floating-point computations, modular integers, pointer manipulations, and memory layouts). It thus discovers all runtime errors. Such *soundness* of results is often considered to necessarily lead to many false positives (warnings about nonexistent problems), but this is not the case with Astrée when applied to its intended target: safety-critical reactive control code with neither dynamic memory allocation, recursion, nor concurrency. By concentrating on the discovery of runtime errors in such programs, Astrée solves a simpler problem than general-purpose analysis tools, but solves it well.

Astrée is specialized. It is parametrized by a set of abstractions that have been specially tuned for use on embedded control-command software, with a preference for avionic and space software. It includes very specific, mathematically sound analyses for constructs commonly found in such applications (e.g., infinite-impulse-response digital filters or quaternion computations) but not in general-purpose software. Designed to be efficient and precise (few or no false alarms) on these codes, it has also been shown to perform well in other application domains of embedded C software.

Astrée has been successfully applied to the analysis of large industrial codes. In just a few hours, it was able to prove automatically the total absence of runtime error in codes of over 1 million lines. For instance, it analyzed Airbus A380 fly-by-wire control code in 14 hours with no false positives.

For more information on the Astrée tool, visit http://www.astree.ens.fr and http://www.absint.com/astree/.
Also see the companion flyer on "Toward Verifiably Correct Control Implementations" in the Research Challenges section of this volume.