ELSEVIER

Journal of Hazardous Materials 130 (2006) 42–47

Journal of Hazardous Materials

www.elsevier.com/locate/jhazmat

Potential use of advanced process control for safety purposes during attack of a process plant

James R. Whiteley *

School of Chemical Engineering, Oklahoma State University, 423 Engineering North, Stillwater, OK 74078-5021, USA

Available online 17 November 2005

Abstract

Many refineries and commodity chemical plants employ advanced process control (APC) systems to improve throughputs and yields. These APC systems utilize empirical process models for control purposes and enable operation closer to constraints than can be achieved with traditional PID regulatory feedback control. Substantial economic benefits are typically realized from the addition of APC systems.

This paper considers leveraging the control capabilities of existing APC systems to minimize the potential impact of a terrorist attack on a process plant (e.g., petroleum refinery). Two potential uses of APC are described. The first is a conventional application of APC and involves automatically moving the process to a reduced operating rate when an attack first begins. The second is a non-conventional application and involves reconfiguring the APC system to optimize safety rather than economics. The underlying intent in both cases is to reduce the demands on the operator to allow focus on situation assessment and optimal response planning.

An overview of APC is provided along with a brief description of the modifications required for the proposed new applications of the technology.

Keywords: Process safety; Advanced process control; Process threat management; Hazard mitigation

1. Introduction

Domestic petroleum refining and petrochemical facilities have been identified by the U.S. Department of Homeland Security as potential terrorist targets [1]. The new process safety implications of a terrorist attack have been described in two previous papers [2,3]. We use the term 'Process Threat Management' to distinguish the new issues outside the realm of traditional process safety management.

The most significant aspects of a terrorist attack are: (1) the intentional nature of the initiating event and (2) the unpredictable consequences over an unknown time frame. Traditional safety systems are designed to deal with one unintentional event followed by a sequence of predicable consequences. As discussed in Ref. [2], the problem is extremely complex. This paper focuses on one specific aspect, operability. The traditional process safety strategy is to shut down or move to a pre-specified fail-safe position when a problem occurs. We have argued that such an approach may not be appropriate during a terrorist attack using the following simple analogy.

1.1. Traditional (unintentional) safety event

An operator is driving a car when one of the front tires experiences a sudden blowout after unintentionally running over a nail. The operator is not sure why the tire failed but has no reason to expect other problems. The operator can anticipate the consequences and makes steering and speed adjustments to bring the car to a stop at the side of the road. This is similar to the response of the passive safety systems used in most plants.

1.2. Terrorist or criminal (deliberate) safety event

The operator is driving a car when a terrorist suddenly jumps from the side of the road and throws a spike strip in

* Tel.: +1 405 744 9117.
E-mail address: rob.whiteley@okstate.edu.

0304-3894/$ – see front matter © 2005 Elsevier B.V. All rights reserved.
doi:10.1016/j.jhazmat.2005.07.038

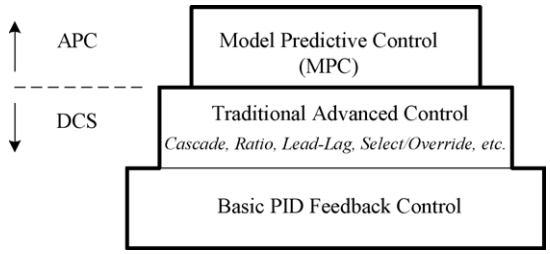

Fig. 1. Automation used for control of continuous commodity processes.

front of the car to cause a deliberate blowout of one or more tires. The result is the same as with the unintentional blowout but the operator has reason to expect additional hostile action. Rather than bring the car to a stop, the operator would slow the car to a controllable speed but continue driving in a direction and manner that maximizes his perceived chances for survival. The responses of the passive safety systems are utilized in this case but continued operation of the process (car) is required.

There are two points to be taken from this analogy. The first is that the concept of the traditional "fail-safe" condition for a plant may not be appropriate during a terrorist attack. The second is that evasive action can produce a response that minimizes the impact of a thinking adversary. In both cases, the operator is responsible for making the decisions that determine the degree to which an attack is mitigated.

Operators are highly skilled with extensive training in how to respond to equipment failure and other emergencies. The operator represents the best opportunity to minimize the consequences of a terrorist attack. Anything we can do to help the operator focus on the big picture during an attack and not worry about details will pay maximum dividends in terms of protecting the public welfare and preserving company assets. The purpose of this paper is to consider how existing advanced process control (APC) systems could be used in a novel way to help the operator during a terrorist attack.

The remainder of this paper is organized as follows. The next section provides background on plant control systems. The emphasis is on how APC works and is integrated with the base plant regulatory control system. Next is a description of two different ways APC could be used as an operator aid during an attack on the plant. The discussion includes concerns and technical issues that would need to be addressed. The final section summarizes the key points raised in the paper.

2. Plant control systems and APC

Control of the process[1] is performed by the distributed control system (DCS). APC is a generic term that refers to automation installed on top of the DCS (Fig. 1). The purpose of APC is to maximize profitability by operating the process closer to constraints than possible with only the base regulatory control system. Before considering how APC could be used to maximize safety rather than economics, it is necessary to first understand how APC fits into the overall plant control system.

2.1. Base regulatory control system

The basic building block of the plant control system is the traditional single input, single output (SISO) feedback control loop. The elements of a SISO feedback control loop are the measurement sensor for the controlled variable (CV), sensor signal lines, controller, and final control element. In most cases, the final control element is an automatic control valve used to throttle the flow of the manipulated variable (MV). The controller compares the measured CV value to the desired or set point (SP) value and adjusts the MV to push the CV back to target. The scheme is predicated on a cause–effect relationship between the MV and CV, respectively. The cruise control in an automobile is a good example of a SISO feedback control loop. The CV is the vehicle speed and the MV is the engine throttle position.

Design of the base regulatory control system involves identifying the key process variables that must be maintained at specified values in order for the plant to produce on-spec product at the desired rate. These are the CVs. Each CV is paired with an MV. An MV can be paired with only one CV. Each CV–MV pair constitutes a single regulatory control loop. As an example, consider a temperature control loop used to regulate the process outlet temperature from a heat exchanger. The CV is the process outlet temperature; the MV is the steam flow rate to the heat exchanger. The temperature control loop adjusts the energy balance around the heat exchanger to achieve the desired temperature. The base regulatory control system for the plant is the collection of SISO loops. There are an equal number of CVs and MVs in the base regulatory control system.

The hardware used to perform the control calculations and interface with the operators is the DCS. The DCS has the capability of augmenting the base regulatory control system with additional functionality. Examples include feedforward, ratio, split-range, valve-position, and other types of control. The DCS also has the ability to impose constraints on allowable CV or MV values. Operator interaction with the process occurs via the DCS. The operator inputs set points, constraints, alarm thresholds, tuning parameters, etc. via the DCS.

The number of SISO control loops required to control a process can vary from tens to hundreds depending on the complexity of the process. Under normal conditions, the operator places all loops in AUTO mode and inputs the set points for the desired steady-state conditions. In this mode, the operator delegates all responsibility for changing MVs to the SISO controllers. If desired, the operator can switch a controller to MANUAL mode to allow direct manipulation of an MV.

[1] This paper focuses on large-scale, continuous commodity processes (e.g., petroleum refining) that utilize APC.

2.2. APC

There are limitations on the quality of control that can be achieved by a collection of SISO controllers (base regulatory control system). The biggest problems occur when the process is multivariable with slow or complex dynamics. APC was developed to handle this situation.

The challenge of controlling a multivariable process with a collection of SISO controllers is easily illustrated by a residential shower with separate hot and cold water adjustments. There are two CVs for the shower: CV_1 = water flow rate and CV_2 = water temperature. Likewise, there are two MVs: MV_1 = hot water flow rate and MV_2 = cold water flow rate. Assume that we construct our base regulatory control system by pairing CV_1 with MV_1 and CV_2 with MV_2. Controller 1 will adjust the hot water flow rate to maintain the desired shower water flow rate and controller 2 will adjust the cold water flow rate to maintain the shower water temperature. Changing one MV produces a change in both CVs. This interaction dictates that corrective action by one controller will always produce a problem for the other controller with long settling times for both CVs. The model-based nature of APC allows anticipation of this interaction problem. APC proactively adjusts both MVs to achieve the desired change in one CV while holding the second CV constant.

There are many situations in a typical process plant where one MV affects more than one CV. The reflux rate in a distillation column affects the purity of all product streams (distillate, bottoms, and side streams). The flow rate through an individual pass in a multi-pass furnace affects the outlet temperatures of all passes. Use of feedback alone (i.e., deviation from set point) cannot provide fast, high performance control when interactions exist.

There are two costs for use of a collection of SISO controllers for a multivariable process, long settling times and the need to specify CV set points unnecessarily far from constraint values. Long settling times result in excessive off-spec production. Interaction effects produce larger deviations from set point during transient periods. This translates to larger 'safety margins' when establishing set points for CVs operating near a constraint. The result is reduced profitability as the economic operating point frequently lies at the intersection of constraints [4].

Model predictive control (MPC) technology was developed specifically to address the control problems of multivariable processes [5]. APC is implemented using MPC technology. For the purposes of this paper, the terms APC and MPC are synonymous. The key to MPC is the use of explicit process models to predict the future response of the process. The process models, not used or required by conventional SISO single loop controllers, allow consideration of interaction effects when making control decisions.

An excellent survey of industrial MPC technology is provided by Qin and Badgwell [6]. The article includes a brief history of the development of MPC and a list of commercially available MPC products. The technical discussion is aimed at the controls community but the qualitative discussion dispersed throughout the article should be enlightening for anyone interested in understanding how MPC works. Additional references for MPC include [5,7].

All MPC control algorithms employ the same two-step process. The first step is to use the process models to predict the future behavior of the CVs assuming no changes are made to any of the MVs.[2] The second step is to determine the optimal sequence of MV changes to eliminate any gaps between the desired CV trajectories and those calculated in step one. The first element in the sequence of calculated changes for each MV is then input to the process. The rest of the MV changes are discarded as the entire two-step process is repeated at the next control instant.

The input MV and CV values to the MPC controller are provided by the DCS. Likewise, the output from the MPC controller is sent to the DCS. Interaction with the process occurs through the DCS at all times.

The MV changes calculated by the MPC controller are not valve position changes but flow or pressure set point changes for cascade SISO control loops. This arrangement protects against failure of the MPC controller, as the base regulatory control system will maintain the process at the last set of valid set points. Consequently, installation of APC typically requires conversion of the base regulatory control system to a cascade configuration as indicated in Fig. 2.

The CV set points input to the MPC control algorithm are typically generated by a steady-state economic optimizer provided with the APC system. A linear program (LP) is commonly used for this purpose. The operator is responsible for inputting upper and lower constraints for all MVs and CVs.

Because the MPC controller is set up to solve an optimization problem, the number of CVs and MVs do not have to be the same. Additional CVs are typically specified beyond those included in the base regulatory control system. The number of MVs in the MPC controller is the same as the base regulatory control system. As an example, a crude unit at the front end of a petroleum refinery may utilize an MPC controller with 50 CVs and 20 MVs. The following list for a generic two-pass preheat furnace illustrates how the number of CVs exceeds the number of MVs.

MPC configuration for preheat furnace

11 CVs: Total process inlet flow rate, process outlet temperature, excess O_2 in stack gas, furnace draft, delta outlet temperature (between passes), delta flow (between passes), maximum skin temperature for pass A, maximum skin temperature for pass B, A pass valve position, B pass valve position, fuel gas rate

4 MVs: Fuel gas pressure at burners, damper position, flow rate through A pass, flow rate through B pass

A typical MPC controller generates a new control action once a minute. The sequence of steps is:

[2] For purposes of simplification, disturbance variables (DVs) are not included in any of the MPC discussion.

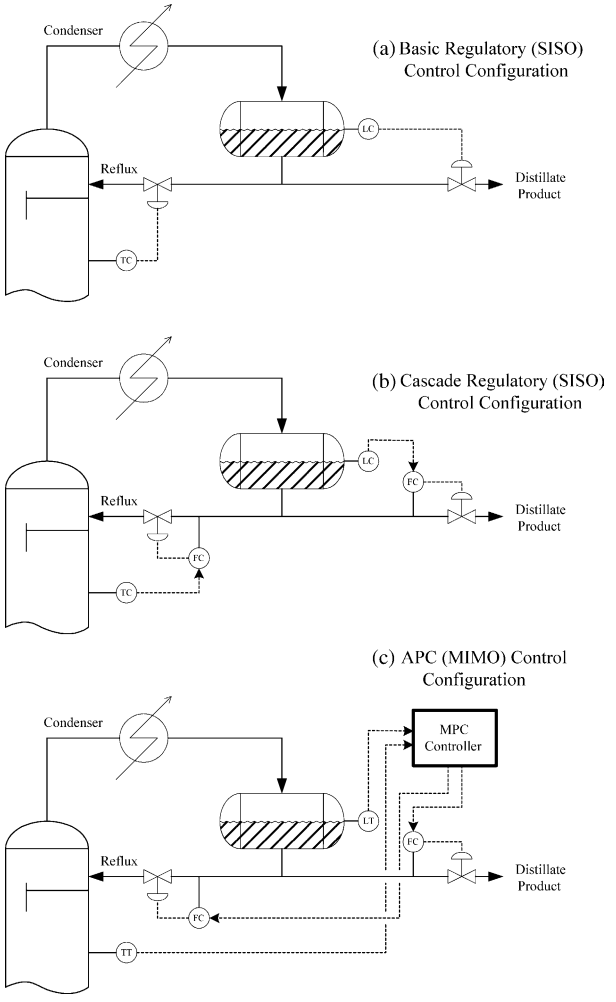

Fig. 2. Cascade control configuration employed by MPC controllers.

1. obtain current values of process variables (CVs, MVs, etc.);
2. predict future CV trends without any change to MVs;
3. determine optimum CV values (steady-state optimization);
4. determine MV changes that optimize CV trajectories along desired path (dynamic optimization);
5. send MV changes to DCS for input to process.

The key to the practical implementation of MPC is step 4, the solution of a dynamic optimization problem. Most MPC controllers minimize an objective function of the following general form:

$$\min J(u) = e^{\mathrm{T}}Qe + \Delta u^{\mathrm{T}}R\Delta u + u^{\mathrm{T}}Su$$
$$\text{subject to : CV, MV, and rate of change constraints} \quad (1)$$

where u is the vector of MV values and the three terms on the right-hand side quantify the deviation from desired CV trajectory, magnitude of control effort, and deviation from optimum MV resting values, respectively. Q, R, and S are adjustable tuning matrices that allow the user to modify the weighting for individual terms.

2.3. Summary of plant control with APC

Only a portion of the total control loops in a plant are typically configured for APC control. The rest operate as conventional SISO feedback loops.

Under normal operating conditions, the following summarizes the interactions between the operator, DCS, APC, and the process:

Operator:	The operator retains ultimate supervisory control of the plant. The operator provides the following information: Non-APC loops: CV set points APC: CV and MV constraints, economic coefficients for CVs and MVs (for steady-state economic optimization)
DCS:	Real-time (second to second) control of the process is performed by the DCS for all loops (APC and non-APC). CV and MV constraints configured in the DCS are rigorously enforced. The DCS constraints are input separately and typically differ from the APC constraints.
APC:	The APC uses the information input by the operator to determine the optimum values for the MVs (as cascade set points) once a minute. This information is passed to the DCS for implementation by the base regulatory control system. APC optimizes economic performance of the plant by pushing the process to operate as close to constraints as possible. APC is totally separate from safety instrumented systems (SIS) used for emergency shutdown of runaway reactions and other potentially unsafe conditions identified by HAZOP.

3. Use of APC during a terrorist attack

One of the first actions by an operator during a plant emergency is to turn off the APC system (converts control configuration from ‘c’ to ‘b’ in Fig. 2). This makes sense as the purpose of the APC is maximized economic performance. During an emergency, the focus shifts from economics to safety. The reason the APC system is disengaged is not because it cannot be used, but because it is configured for a different purpose than called for by the situation.

The features that make APC attractive for economic optimization are equally attractive for safety optimization. The model-based or predictive nature of APC is potentially of greater value in crisis management during a terrorist attack than during the pseudo-steady-state of routine operations.

3.1. Conventional application for transition management

Discussions with operations personnel reveal that the likely first action by an operator during a terrorist attack would be to ramp the process down to a reduced production level. The intent is to quickly back off to an intermediate point while providing time for the operator to assess the situation and determine the best course of action. This gets the process moving in the direction of a shutdown and is consistent with the tire failure analogy presented earlier.

Ideally, the operator could initiate this transition with a 'one-button' console selection in the same manner as provided by Emergency Shutdown and other SIS systems. A full discussion of the issues of a 'one-button' option is outside the scope of this article. Of current interest is the role of APC in quickly moving from maximum or near-maximum rates (typical operating point) to an intermediate observation point.

Without APC, an operator must modify set points or manually operate a large number of control valves to rapidly ramp down production. The operator must anticipate process interactions while making changes under duress. The effort required to switch between DCS displays and monitor the effect of changes is substantial and occurs at a time when the ability to focus on the big picture rather than details is most beneficial. There is no question that automation of the procedure would be preferred. APC was developed to determine the optimal path between operating points. Use of APC to manage the transition to an intermediate operating point would free the operator to concentrate on situation assessment and how best to proceed.

Conceptually, it should be possible to utilize an existing APC system to manage the transition to an intermediate point. The intermediate point would need to be defined in advance and should be reachable from any operating state. The steady-state economic optimization (step 3 in Section 2.2) would be replaced with the CV targets associated with the intermediate operating point. Replacement of the CV and MV constraints would also be necessary. As before, the constraints need to be defined in advance. The target and constraint replacement process must be automated, as manual entry by the operator is impractical. The configuration of the existing objective function used for dynamic optimization, Eq. (1), should be suitable although modification of the tuning matrices may be desirable. In particular, reduction or elimination of the penalties on large MV changes and deviation from ideal resting values is probably desired.

The potential use of APC for initial transition management during a terrorist attack is conceptually straightforward. In addition to the requirements mentioned in the previous paragraph, the following concerns would also need to be addressed. 'Bumpless transfer' between APC configurations must be provided. Also, the models used by most APC systems are developed at typical production rates. The desired intermediate operating point may correspond to a 50–80% reduction in rates. The existing APC models, linear or nonlinear, may be inadequate at these reduced rates. New models for the transition mode would then be required. Model generation has historically been time consuming and expensive. However, APC vendors are introducing new tools to simplify model creation and generation of transition mode models is now more practical. Current APC systems are intentionally designed with limited robustness to loss of signals from CVs and MVs. Modification of the conditions under which the APC system automatically disengages may be appropriate during transition mode operation.

This application can be thought of as an 'on-the shelf' operator aid. That is, the transition mode APC configuration is prepared in advance. Implementation simply requires substituting the transition mode configuration if/when required. While extremely valuable in terms of freeing up operator time, this application provides benefits over a finite time period. The APC benefits are exhausted once the process reaches the intermediate operating point (neglecting any credit for continuing to hold the process at the intermediate point).

3.2. Non-conventional application for safety optimization

Another potential application of APC during a terrorist attack involves modifying the steady-state and dynamic objective functions to optimize safety rather than economics. This represents a completely new use of APC. In this new case, the APC system would drive the process in an inherently safer direction that minimizes the impact of lost containment or vital services. The key to this application appears to be real-time specification of the appropriate CVs and CV values for the dynamic objective function.

As a rule, reducing the pressure and inventory in the process directionally reduce the potential hazards. This implies use of pressures and liquid levels as CVs in the new APC configuration. It may also be more desirable to shift inventories from process equipment to storage areas where the effects of equipment failure are more confined and less likely to produce domino effects. This would imply including material balance streams as APC CVs. There may be opportunities to modify the hazardous characteristics of a material by blending or reacting with other materials. Again, the APC CVs should include the appropriate material balance streams and liquid levels.

Identification of the appropriate CVs for an APC safety application is an area where more work is needed. It is not clear whether all the CVs needed for the safety application will be available in the list of CVs used by the original profit optimizing application. New CVs would require generation of additional models. The MVs will remain unchanged. Possible sources for the CV target values include:

Source 1: Some CV target values may be fixed under any circumstance. An example is the valve position on a reactor coolant line where full open is desired at all times in an emergency situation. Traditional process safety management is based on the concept of 'fixed' CV safety targets.

Source 2: A second source for CV target values is an analog to the steady-state optimizer used by the original profit optimizing application (step 3 in Section 2.2). This implies the ability to formulate an LP-type safety objective function. The CV target values would not necessarily be fixed although it is possible that the LP solution could be constant. Formulation of a safety objective function is non-trivial.

Source 3: The third source is a higher level optimizer that provides CV target values for multiple APC systems. This third source employs a high level view of the entire plant to determine the optimal dispositions for inter-unit transfers and could automatically compensate to divert streams away from damaged units or units under attack.

The use of an optimizer (Source 3) to establish CV target values for the safety APC application could produce unusual evasive procedures. An example would be intentional overpressure of a vessel to relieve the contents to a flare. Another possibility is maintaining a high firing rate in a furnace to push material up a column and out a different product line to relocate inventory.

The frequency of APC control execution may become an issue during safety optimization. Traditional profit optimizing APC systems typically generate control updates once a minute. This frequency is appropriate for most processes based on the time constants for the economic drivers. The time-scale for safety events is much shorter. Once-a-minute execution by the APC system may be too slow.

All of the implementation issues described previously for the transition mode APC application apply to the safety optimizing application as well.

4. Concluding remarks

This paper described the potential application of APC in two new ways to minimize the impact of a terrorist attack. However, the proposed applications could also provide benefits during non-terrorist plant emergencies. The proposed capability would be available regardless of the initiating event. The need exists to minimize the demands on the operator during any type of abnormal situation. By utilizing APC for safety rather than economic optimization purposes, the operator would be freed to focus on higher level response planning with increased potential for loss reduction in any emergency. The structured control trajectory provided by APC would also provide organizational benefits in terms of speed and consistency of the operational response. Conceptually, the model-based capability of APC should be able to move the plant towards shutdown faster and safer than an operator facing a console or board of active alarms.

As discussed, there are many issues requiring further work to apply APC for the proposed new applications. Hopefully, there will never be an attack or the attack will be thwarted by the additional security measures adopted by the process industries. However, public welfare demands consideration of the potential for a successful attack and identification of mitigating procedures and technology.

The underlying themes of this paper are: (1) the knowledge and creativity of the operator represent the best bet for minimizing damage; the most leverage can be achieved by empowering the operator and (2) APC and other existing plant technologies have capabilities that can be exploited in new ways to assist operators during plant emergencies.

Acknowledgments

The author gratefully acknowledges the contributions of his colleagues with industrial operations and advanced control responsibilities.

References

[1] National Infrastructure Protection Center, Homeland Security Information Update, Information Bulletin 03-003, 2003.
[2] J.R. Whiteley, M.S. Mannan, J. Hazard. Mater. 115 (2004) 163.
[3] J.R. Whiteley, J. Wagner, Process Saf. Prog. 23 (2004) 279.
[4] W. Lee, V.W. Weekman, AIChE J. 22 (1976) 27.
[5] B.A. Ogunnaike, W.H. Ray, Model Predictive Control, Oxford University Press, New York, 1994, Chapter 27.
[6] S.J. Qin, T.A. Badgwell, Control Eng. Practice 11 (2003) 733.
[7] J.A. Rossiter, Model-Based Predictive Control: A Practical Approach, CRC Press, Boca Raton, FL, 2003.