Brace yourself

Chemical process control simulation with stochastic events

by Carl Sandrock

The problem seems deceptively simple: You are at a cocktail party. Try to stay upright and walk from the door to the dessert table while people are pushing you from all sides. The human brain is such an advanced computational tool that few people have given thought to the calculations required to complete such a task. In chemical processes, variables have limits (staying upright), target values (the table) and unexpected disturbances (people jostling for position) in addition to free variables (the placement of your arms or direction of your shoulders). Process control can be summarised as the science of reaching the target while staying within limits in the face of disturbances.

Process control has come a long way since the first controllers were used to regulate water clocks more than 2000 years ago, but really boomed with the advent of digital computers. Suddenly it was no longer required that the control laws be implemented in analogue hardware, opening up a vast new field by incorporating a model of the process in the controller. Before widespread use of computers, control engineers would design controllers 'once off' using mathematical manipulation, and the structure of the controller was fixed as it was installed. Nowadays, all aspects of industrial controllers can be changed with minimal impact to the system. Even so, modern control systems are usually based on so-called 'linear' models, basically a simplification of the complexity of realistic models that reduces the computational effort involved in solving the control equations. In addition, control systems are largely 'reactive' in that they react to measured events only after they have happened.

Picture yourself back at the party. Once you have reached the table, you may find yourself leaning back slightly to compensate for the effect of people pushing up behind you. This positioning of your free variables makes you more resistant to disturbances. What is required to allow a control system to emulate this bracing behaviour and how do we evaluate the effects of such a strategy?

Firstly we require a method of simulating realistic inputs into the system – you need to know that a push is more likely to come from behind you than from the table. On a chemical process plant, problems are usually due to identifiable events. The identification of events in past input signals can give us clues about how they are distributed in time and we can use Markov processes to generate a similar series of events. This in itself is not a simple problem, as process data usually contains much noise, and the highly interactive nature of modern plants makes it difficult to identify which input led to the failure. Let us imagine for a moment, however, that we have a single input with already idenfified events. The Markov chain is simply represented as a directed graph containing states at the nodes as shown in Figure 1. One starts at some initial state and at each evaluation step one moves to the next state with a certain probability which can be represented as a matrix with rows and columns for every state.

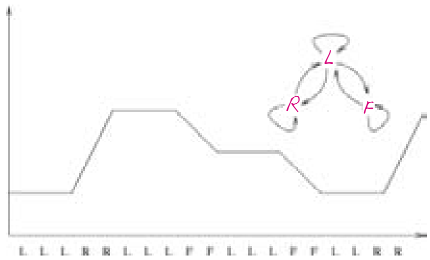

→ 1. Describing variable events using Markov Chains and level (L), rising (R) and falling (F) elements.

Next we need a good model of the system in question and a method of simulating the system's response to inputs over time. Much research has gone into such systems, but no freely available dynamic chemical process simulator exists that can accept stochastic inputs and model their effects on both inside and outside the controller. By combining aspects of the current state of the art hybrid simulators, and introducing novel interaction between the controller and the model, systems can be simulated that are not only stochastic, but "aware" of their nature. We can use the simulator to do many runs of the same process with random changes obeying the original distributions in a Monte Carlo simulation. These runs can then be analysed to find the statistical properties (like the average or most likely value) of any variable. This only makes sense when the results are presented in an intelligible way, so we need methods of visualising these distributions. The results from the simulation are then used in a control algorithm that uses the statistical properties as inputs. In this way, the controller can position the system to be resistant to a likely disturbance or more amenable to a likely command.

A framework combining all these elements (event identification, stochastic hybrid simulation and result visualisation) is currently being developed. While still a work in progress, some interesting insights can already be gleaned from the results obtained modelling simple systems. As a simple example – simpler than our earlier cocktail party analogy – consider the cart shown in Figure 2. A pendulum is mounted on the cart, which can move from left to right. The goal is to move from one place to another without dropping the pendulum.

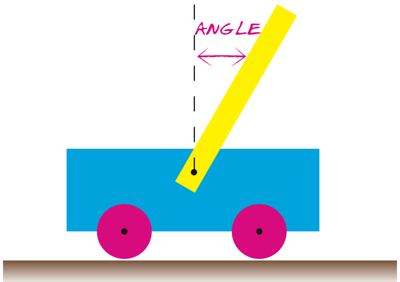

→ *2. The cart problem is a classic control problem*

Let us consider the response of the system to a gusty wind that changes direction every so often. In Figures 3 and 4, the angle of the pendulum is shown when it is started at a zero angle (pointing straight up) and then moves due to the gusty wind. For Figure 3, the wind is more likely to blow from the left, while in Figure 4 it is more likely to blow from the right.

We can see the difference between these cases clearly – the pendulum is more likely to fall in the direction that the wind is blowing. Doing a simulation like this gives us additional information about what the system will be doing in the future, giving the controller more information. Simulations with controllers on this system shows an improvement in performance over traditional systems due to the placement of internal states.

Due to the incredible computational effort required to do such simulations, the program is currently being run on the University of Pretoria's 24 node Velocity cluster. Even on a system like this one, the effort required is vast. To simulate the five seconds shown in the figures, required 10 seconds of computer time. That doesn't sound so bad on its own, but taking into account that control action has to take place faster than once a second, it becomes clear that the multiple computer route is the only way to go. The optimisation of the program and simulation code is clearly also an important element of the project. However, if computer capacity keeps doubling every 18 months, we can expect simulations like this one to become commonplace within the next five years.

This system is not limited to simple simulations like this one – it has been extended to accommodate chemical systems including many interacting nonlinear components. Even though it is still in development, there has been much interest from industry in the results of the simulations, as several industries have inputs that are not well characterised by a constant input at a given time, but could be any of a range of values. The combination of event detection, stochastic simulation and control based on statistic properties of future values makes this project unique, but this also means that there are interesting times ahead. We are waiting to see what will emerge in the future.

Carl Sandrock *is a lecturer in the Department of Chemical Engineering at the University of Pretoria.*
carl.sandrock@up.ac.za

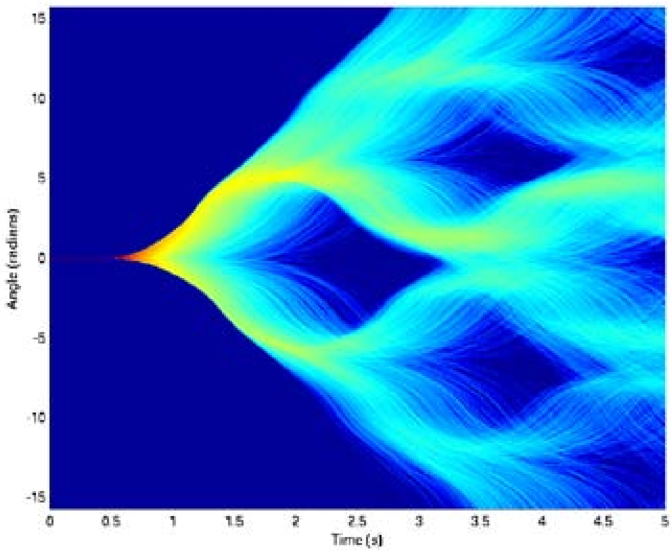

→ *3. System response to biased force to the left – more red colour means more likely*

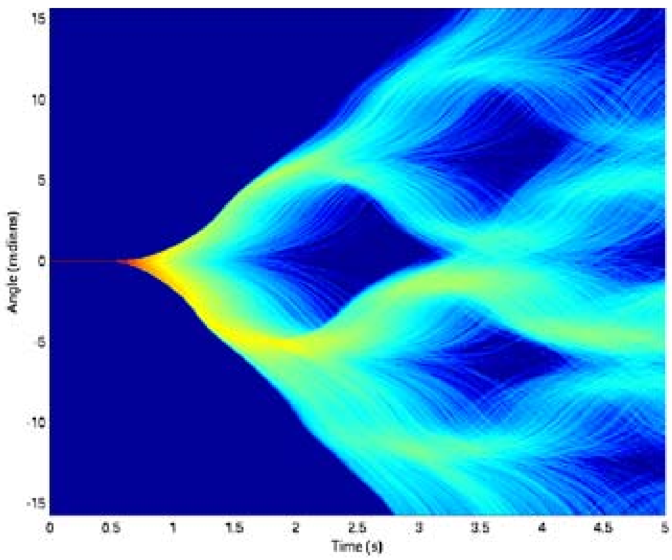

→ *4. System response to biased force to the right – more red colour means more likely*