A Concept of Resilience in Production Systems

Felicjan Rydzak
Centre for Advanced Manufacturing Technologies
Wroclaw University of Technology
ul. Lukasiewicza 5, 50-371 Wroclaw, Poland
Phone: +48 71 3204184 / Fax: +48 71 3280670
Email: felicjan.rydzak@pwr.wroc.pl

Piotr Magnuszewski
Institute of Physics
Wroclaw University of Technology
Wybrzeze Wyspianskiego 27, 50-370 Wroclaw, Poland
Phone: +48713202918
Email: piotr.magnuszewski@pwr.wroc.pl

Jan Sendzimir
International Institute for Applied Systems Analysis
Schlossplatz 1, A-2361 Laxenburg, Austria
Phone: +43 2236 807 471 / Fax: Fax: +43 2236 71 313
Email: sendzim@iiasa.ac.at

Edward Chlebus
Centre for Advanced Manufacturing Technologies
Wroclaw University of Technology
ul. Lukasiewicza 5, 50-371 Wroclaw, Poland
Phone: +48 71 3202046 / Fax: +48 71 3280670
Email: edward.chlebus@pwr.wroc.pl

Abstract: *The article presents the results of ongoing research on resilience in production systems. It refers to the term resilience as used in socio-ecological systems and applies it to assess the long-term functionality and effectiveness of industrial organizations in an uncertain world. It concentrates on refineries and chemical plants which, due to complex and sophisticated production processes and products themselves, are especially vulnerable to various kinds of disturbances. To illustrate the idea of resilience we examine machine reliability improvement programmes using System Dynamics models to investigate situations when internal or external stress can lead to regime shift in production systems and make them move to an undesirable configuration e.g. from proactive to reactive maintenance mode. Resilience proved to be a useful concept for production managers interested in organizational transition to a more desired operational domain as well as remaining in the desired domain. System Dynamics methodology offers powerful tools for examining the impacts of various policy options on the dynamics of resilience in production systems.*

Key words: resilience, production system, system configuration, System Dynamics

1. Introduction

Modernization of production plants has so fully exploited miniaturization, speed, automation and robotics through advanced computer applications that the notion of reliability, availability, maintainability and safety are becoming of great importance for industrial organizations. Precise synchronization of production processes itself is not enough to shorten manufacturing lead time, time to market and to serve clients with customized products. Efficient production is possible only if processes are not interrupted by e.g. machine failures and breakdowns. Thus it seems appropriate to extend the idea that quality depends on processes (Deming 1986) and agree with Nakajima (1988), an advocate of Total Productive Maintenance, that production output – as well as quality – all depends on using those processes for expert coordination of equipment.

Total Productive Maintenance (TPM) is a concept carried out by all employees through small group activities (Nakajima 1988). TPM originates in an American approach - 'preventive maintenance', which was introduced in Japan in 1951.'Preventive maintenance' focuses on daily maintenance and periodic inspections to lower the costs of long-term maintenance and repair of periodic failures. This concept went through several stages of development – machine operators were involved in routine maintenance, the equipment started to be modified in order to improve reliability – and finally that experience founded the TPM programme (Venkat 2003). The dual goal of TPM is zero breakdowns and zero defects, which result in increased productivity, costs reduction and inventory minimization.

The idea of reliable equipment and uninterrupted production is even more important for refineries and chemical plants, due to characteristics of the production processes and products themselves. Petro-chemical manufacture is extreme both in scale and danger of operation. Normally huge complexes with highly automated infrastructure are dedicated to realize continuous operation of refining and synthetic processes. Most of the processes carried out in refineries involve high (often lethal) levels of temperature and pressure in combination with highly flammable materials or hazardous chemical components. This means that any equipment failure or malfunction can lead not only to financial dissatisfaction but also to catastrophic results.

Reliability improvement programmes have the potential to eliminate breakdowns and defects in the production systems. However, they are very often perceived as expensive and often do not survive management's policy revisions that seek to cut costs for short-term profitability. As recognized by Ledet and his team in DuPont Chemical Company (1999), only by engaging the entire production organization for sufficiently long periods is it possible to achieve substantial results in reliability improvement. For that purpose, based on benchmarking study and modelling effort, an interactive simulation called The Manufacturing Game was developed to shorten the transition process by engaging the very people involved in implementing such changes. The BP Lima Refinery is one of many actual applications where the usefulness of The Manufacturing Game workshop was demonstrated. In playing and discussing the game, the real-life actors experienced how long-term reliability can emerge from sustained application of proactive maintenance. In the aftermath of that success the leaders of programme in Lima refinery announced 'A New American TPM' (Griffith, Kuenzli, and Monus 1999).

The benchmarking study in DuPont (Ledet et al. 2005) identified the three most common approaches of conducting maintenance - '*reactive*' (equipment is repaired after it breaks), '*planned*' (equipment is repaired before it breaks) and '*precision*' (equipment is not only repaired but improved). Among these three practices the best performance is achieved when the 'precision' operating domain is realized by adjusting certain relationships between agents and processes in the resource allocation and production systems. The fact that an organization remains in the 'precision' domain does not mean that it will not move to a lower performance domain, often for the very reasons why short-term cost-cutting often causes proactive maintenance to be abandoned before its full potential was realized. Even progressive companies that vigorously try to transform production system performance may fail. The dynamics embedded in the structure – relations and processes – of the production system as well as various disturbances affecting the system might cause the organization to shift to a domain of inferior operations (a regime of sub-par performance).

Enterprises achieve long-term profitability and competitive advantage by developing management regimes that sustain their production system's functionality in the face of disturbances. The desirability of a domain or regime can be evaluated from a business perspective (profitability) or a socio-ecological perspective (sustainability). Transitions between desirable and undesirable domains have been analysed for more than thirty years as part of the development of Resilience Theory (Gunderson, Holling, and Light 1995; Olsson et al. 2006; Van de Brugge 2007). This theory has promoted the study of surprising and non-linear dynamics of ecosystems and society (Carpenter et al. 2001; Gunderson 2000). In this article we used the experience of studying resilience in social-ecosystems and apply it to assess the long-term competitiveness of businesses in an uncertain world. As Stolz (2004) argues '*Resilience is the only sustainable, portable strategic plan. Resilient individuals, teams, and organizations consistently outlast, outmaneuver and outperform their less resilient competitors'*.

This paper presents one perspective of *resilience* as a way to deal with uncertainty and disturbances in production systems. Applied at different levels of an organization, resilience helps to balance functionality and efficiency in the face of different forms of disturbances. The paper focuses on refineries and chemical plants, but a concept of resilience can also be recognized in other production systems. The paper is organized as follows. In the next section a concept of resilience is introduced. Section 3 discusses dynamics of system configurations for two case studies – DuPont Chemical Company and Lima refinery. Further investigation of resilience using System Dynamics models is carried out in section 4 and 5. The final section makes some closing remarks.

2. Resilience

Resilience is an increasingly popular term in today's business organizations, which try to apply that concept in various areas and at different levels of operation. Some consider resilience as an ability to deal with personal stress and misfortunes. (Coutu 2002) Others define it as a strategic ability of enterprises to continuously adjust their operations to

requirements of a turbulent age (Horne 1997; Hamel and Välikangas 2003). The resilience concept also crosses the boundaries between individual enterprises. Supply networks are becoming more complex and tightly coupled as companies start to compete through their entire supply chains, and are more vulnerable to such events as terrorists attacks and natural disasters (Rice and Caniato 2003; Christopher and Peck 2004; Sheffi 2005).
This paper refers to the term *resilience* as originally proposed by C. S. Holling (1973) and further developed in the study of the dynamics of socio-ecological systems. *Resilience* addresses the dynamics of dealing with disturbance - how a system absorbs the impacts of stress or shock and how it re-organizes afterward. In that context it is defined as the capacity of the system to undergo disturbance and still maintain its functions, structures and controls. Three system characteristics broadly describe how to assess its resilience: the amount of change the system can undergo and still remain in the same configuration (retain the same controls on structure and function), the degree to which the system is capable of self-organization, and the degree to which the system can build the capacity to learn and adapt (Carpenter et al. 2001; Walker et al. 2002). Our world is a hierarchical set of systems nested within each other, from the micro- to the macro-scopic, and they often are at different stages of development. What is a single element to some is seen as a vast system to others. So care must be taken to define the system in question in terms of its bounds and its state of development (nascent to mature) relative to its surroundings. Then one can precisely answer the two key questions, as Carpenter et al. (2001) argue: what system state is being considered ('resilience *of what'*) and what disturbances are of interest ('resilience *to what'*).

Applying this idea to socio-technical systems, it can be assumed that the more resilient configuration of the production system is, the larger the extent and number of certain disruptions the system can undergo without serious losses. Extent relates to the intensity, duration and spatial footprint of a stress or shock. Being resilient is highly attractive to any enterprise provided that the production system is in a favourable *configuration*. A management regime is defined by the web of relationships linking its agents and processes in the allocation of strategic resources to sustain system functionality and fulfil company goals. A *domain* or *configuration* represents a diversity of system states that are possible under a particular *regime* of actors, processes and functions. For example, most healthy humans (where *healthy* is one configuration) can assume many different states, from running a marathon to standing on their head, and their heart will still function well (stay in a healthy stability domain). However, it is naive to use such robust examples to cast resilience as a positive property. Those who suggest 'resilience' as a cure to current problems in organizations ignore a frequent need of many systems: the need to change configurations. Resilience has to do with the capacity to persist in a stability domain, but that domain may be undesirable. For example, some enterprises survive in an unprofitable state, staving off fundamental reform through clever use of bankruptcy law. Companies that perceive their configuration as undesired may have to lower their resilience in the short-term in order to escape their current domain. They can do this by initiating improvement programs, but they may suffer failures and losses in the process. To effect real change to a more desirable domain one must develop and use the system characteristics that can build or undo resilience – self-organization, learning and adaptive capacity, which are embedded in a production system, to influence the dynamics within the organization. Such characteristics can help management navigate the rough waters of transition, where profound and meaningful change may require many kinds of sacrifice: regime change (shifts in personnel and policy), sustained investment in learning for the future even as profits stay low during the transition. To illustrate the idea of resilience in production systems we will examine examples of machine reliability improvement in the DuPont Chemical Company (Ledet 1999) and in the BP Lima Refinery (currently Valero - Lima) (Kuenzli, Griffith, and Monus 1998).

3. Resilience in practice

Resilience and system stability is widely described in literature in form of ‘ball and cup’ heuristic (Carpenter, Ludwig, and Brock 1999; Gunderson 2000). The term heuristic emphasizes that this graphic is a tool for discussion, learning and proposing experiments, not for prediction. Accuracy is not a question here. This is a simplification of complex dynamic situations that is useful in dealing with high uncertainty. Still, it is a very convenient way of presenting the idea of resilience in production systems and can be used as a conceptual tool with managerial applications. An example of the heuristic is presented in Figure 1.

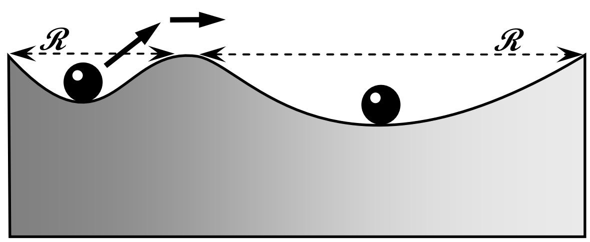

Figure 1 Ball and cup heuristic of system stability

Valleys represent stability domains (desirable or undesirable attractors of the system), balls represent the system state (current state of the production system), and solid arrows represent changes in a system state, often caused by disturbances (peoples’ actions focused on change or perturbations affecting system – unexpected machine breakdowns, shortage of supplies, etc.). However, change between system states can also come without disturbance, as links between actors can shift stochastically. How the ball moves about this valley depends on the energy of movement and the valley’s shape. The energy of movement comes from the play between the excitation of outside drivers, be they shock and disturbance or energy sources, and the pull of gravity as the system ages and energy dissipates. The shape of the domain, which changes while system experiences disturbances, reflects the configuration of infrastructure and relationships that define the system’s self-organizing, learning and adaptive capacity – the characteristics of resilience. Resilience (ℛ) itself is most simply approximated by the width of domain as depicted by dashed arrows (domains that are wider can be considered as more resilient), though the valley shape (steepness and volume) that signifies the likelihood the ball will escape the valley would be a more precise measure.

3.1 DuPont Chemical Company story

The reliability improvement programme conducted by Ledet (1999) and his team in DuPont was a result of a benchmarking study investigating the best maintenance practices world-wide. The study revealed that maintenance operations in DuPont had a reactive nature – machines were repaired only *after* they broke down. The main purpose of the improvement programme was to change the focus from defect correction to defect prevention and elimination (Sterman 2000).

From the perspective of resilience and the 'ball and cup' heuristic, while starting the improvement programme, DuPont Chemicals remained in a 'Reactive Maintenance Domain'. Theoretically any alternative to the present domain does not have to exist at all, but as far as DuPont is concerned, it seems that an alternative domain existed – a 'Planned Maintenance Domain' (Figure 2, landscape#1).

> 'We already know that planned maintenance is a good idea', 'We tried those policies and they didn't work'. (Sterman 2000, p.74)

DuPont tried to move the system – solid arrow in Figure 2, landscape#1 – from the wide (resilient), undesirable domain but their motivational pressure was not sufficient.

> Ledet realized that the client group for the project – the group of people whose behaviour had to change for any results to be realized – was far broader than the management team responsible for maintenance. (Sterman 2000, p.74)

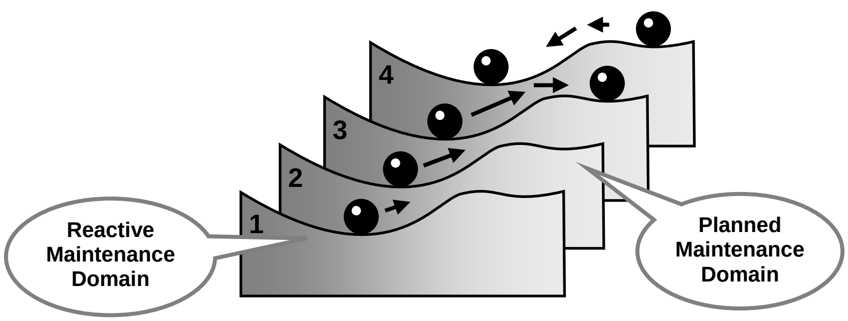

Figure 2 Ball and cup heuristic on DuPont Chemicals story

In order to engage more employees The Manufacturing Game was built at the request of the Maintenance department's management. People could attend a 2 day workshop and right afterwards do planned maintenance. Playing The Manufacturing Game can be perceived as introducing 'disturbance' to a system stuck in the 'Reactive Maintenance Domain' and is presented as the solid arrow in Figure 2, landscape#2. However, the magnitude of disturbance caused by one workshop for several participants was 'absorbed' by a still resilient, undesirable domain.

> The game and learning Laboratory proved popular throughout the company. But playing it once with a small group of managers wasn't enough. The team found that they had to run several workshops for a given plant before a critical mass of people emerged to lead action team that put proactive maintenance policies into practice. (Sterman 2000, p. 75)

Only involvement of a significant majority of workers in planned maintenance created enough trust and understanding in the long-term potential of proactive approach to maintenance that the vicious cycle of the reactive mode was broken and led to change in system configuration as presented in Figure 2, landscape#3.

Any transformation within an enterprise requires 'technical' changes, corresponding to know-how, production processes, standards and procedures, and internal regulations, and also changes in peoples' behaviour, the way they perform their tasks (Heifetz and Linsky 2002). These two kinds of internal changes influence the dynamics of domains' shapes. Decreasing the resilience of the undesirable domain and building the resilience of the desired one would be of great help during the DuPont improvement programme. Though people engaged through The Manufacturing Game workshop into proactive maintenance moved the system to the Planned Maintenance Domain, that domain was still very narrow – not resilient.

Unfortunately DuPont Chemicals did not sustain the results of the improvement programme – the resilience of the new, planned maintenance domain was not sufficient to catalyze and secure the forthcoming changes as long-term changes in policy.

> 'As soon as you get the problem down, people will be taken away from the effort and the problems will go back up'. In fact, cost-cutting programs mandated by corporate headquarters did cause significant downsizing throughout the entire company and limited their ability to expand the program. (Sterman 2000, p.77)

The understanding gained by middle management in game-playing and experimentation in the production system had not been communicated to upper management. Interventions from company headquarters, completely out of synch with local improvements, can be also interpreted as disturbances. However, unlike some systems that thrive on opportunities for renewal that periodic disturbances can offer, such disturbances are not the basis for company evolution, quite the opposite. They affected the system state and made it return to the undesired 'Reactive Maintenance Domain' as presented in Figure 2, landscape#4.

3.2 BP Lima Refinery story

In the 1990's BP Lima Refinery experienced some serious organizational problems. Poor performance and unsuccessful improvement programs made BP think first about selling and later about closing the plant. This threat was a driver for workers to save the refinery through actions focused on reliability improvement using the DuPont experience with The Manufacturing Game (Griffith, Kuenzli, and Monus 1999). Prior to this latest reliability improvement programme, maintenance practices concentrated on reactive operations, keeping the system in 'Reactive Maintenance Domain' (see Figure 3, landscape#1). Once the programme started the 'Precision Maintenance Domain' appeared.

> "What are you doing so that this won't happen again?" [Midwest State System Manager] was the first person I heard use the expression, "Don't Just Fix It, Improve It." (Houshower 1999, p. 26)

Again The Manufacturing Game workshops can be perceived as a disturbance introduced to probe the system configuration for new possible links and alignments (solid arrow in Figure 3, landscape#2). For example, internal regulations were realigned to promote the idea of improvement, and workers were allowed to follow this message. Encouraged by an emerging new regime (new attitudes of their leaders and new rules in the organization) employees started to develop new norms for actions. This directly reshaped the stability domains and enhanced resilience of the precision domain as suggested by the widening right valley in Figure 3, landscape#2.

> We were not just talking about pumps, but our other organizational and management decisions. At the time, we did not realize it, but it became our policy. And what an easy to remember policy it was!
>
> If we were trying to do a rush job, the mechanics quickly pointed it out to us: "Don't just fix it, improve it, huh?"
>
> Management was willing to say, "Nobody likes to spend money, but OK, you're convinced this is a problem, so spend. We did and we solved the problem and saved a lot of money.
>
> The action teams felt they had the power to fix things and they were given money and authority to make changes. (Houshower 1999, p. 26-27)

Spending time and money on improvements, especially those that did not immediately offer benefits, can lower resilience to stress or shock. But that may be a necessary price of transition toward higher resilience. Lowering the resilience of the undesirable domain and increasing the resilience of the desirable one, facilitated through The Manufacturing Game workshops, was followed by the financial success of the refinery (Figure 3, landscape#3). Some results of the successful organizational change in Lima Refinery from 1991 to 1998 include such statistics as increase of pump Mean Time Between Failures (MTBF) from 12 to 58 months and improvement in planned work fulfilment up to 95%. The safety of employees and contractors improved by factor of 4. The indicator of hydrocarbon loss to the atmosphere decreased from 1.75 to 0.35. The estimated value added from the project to the refinery was about $43 million per year. (Griffith, Kuenzli, and Monus 1999)
The Lima refinery avoided closure and was sold to Clark. It was not the end of the effort of change, however. People experienced that improving machines rather than just fixing them is the appropriate way of doing maintenance. They continued to develop habits, which combined with policies, procedures, reward systems and social values increased the resilience of Precision Maintenance Domain further, as presented in Figure 3, landscape#4.

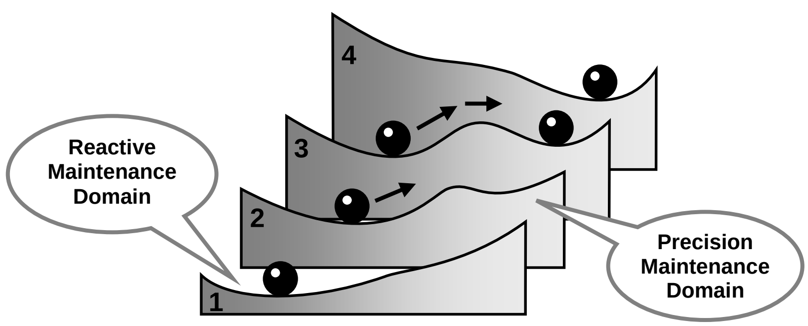

Figure 3 Ball and cup heuristic on BP Lima Refinery story

These examples highlight two key aspects of organizational change in a context of resilience – transition to a more desired operational domain and remaining in the desired domain. As illustrated through 'ball and cup' heuristics for both DuPont and Lima Refinery the most effective way to conduct a project of change is to decrease resilience of existing, undesired domain and force the system to move. However, the change will not be sustained unless the alternative, desired operational domain is resilient enough. Lessons learned have to be secured as new policies and operating procedures. Any attempt to change the existing operations is usually expensive, involves people, time and the risk of experimentation. The threat of losing valuable lessons mandates thorough investigation of which structures and processes should be improved and which influences should be reduced in order to make the transition more effective and successful. Regarding the sustainable changes the conducted analysis should also investigate what elements should be preserved to enable the system to renew and reorganize itself following a disturbance. These latter elements constitute the system's memory (Berkes and Folke 1998; Berkes 2002; Berkes, Colding, and Folke 2003) and form the nuclei around which the system self-organizes post-disturbance. As a tool for this kind of analysis, System Dynamics methodology is proposed (Rydzak and Chlebus 2005).

4. Regime shifts in production systems – a System Dynamics model

To further investigate how resilience changes in production systems a System Dynamics model was built. Though this is a model of a fictional refinery it addresses the problem of machine reliability improvement in DuPont Chemicals and Lima Refinery. The model incorporates the idea of equipment flow and machine availability, due to the number of defects in the production system, presented by Ledet (2002) and also described in Sterman (2000). Allocation of resources for production, maintenance and repair reflects 'working harder' and 'working smarter' strategies investigated by Repenning and Sterman (2001). The general structure and key variables of this new model, trying to investigate dynamics of transition between alternative operating domains, are presented in Figure 4. The model variables and equation formulations are presented in Appendix A.

Total production depends on availability of operators and machines. Since organizations are often unwilling to invest in infrastructure development, any increase in desired production level creates pressure on operators to work harder (*Pressure to Do Work*). The production processes can be disturbed by machine breakdowns. Broken machines decrease the level (available number) of working machines, which impairs productivity. If there are broken machines pressure is put on workers to repair them.

Should a production gap occur an alternative strategy would suggest taking machines down and putting pressure on maintenance to limit the number of unplanned interruptions (*Pressure to Do Maintenance*). This policy requires shutting down and inspecting some working machines since inspection takes less time than repairing actions.

The machine breakdown rate depends on machine reliability. The higher the reliability of machines the less machine failures occurs. Total machine reliability is increased mainly through maintenance efforts but also through repairs. Since reliability is decreased by day-to-day machine utilization continuous involvement is needed to sustain this characteristic on a desired level. On the other hand, be it maintenance, repair or production, one operation excludes others. In a case of breakdown or installation of a new unit, operators very often help maintenance people. Even if operators are not directly involved in machine service they cannot execute their work, since the equipment they are supposed to use is occupied by maintenance people. Similarly if there is a great pressure on production the maintenance will be not carried out until machine will break down.

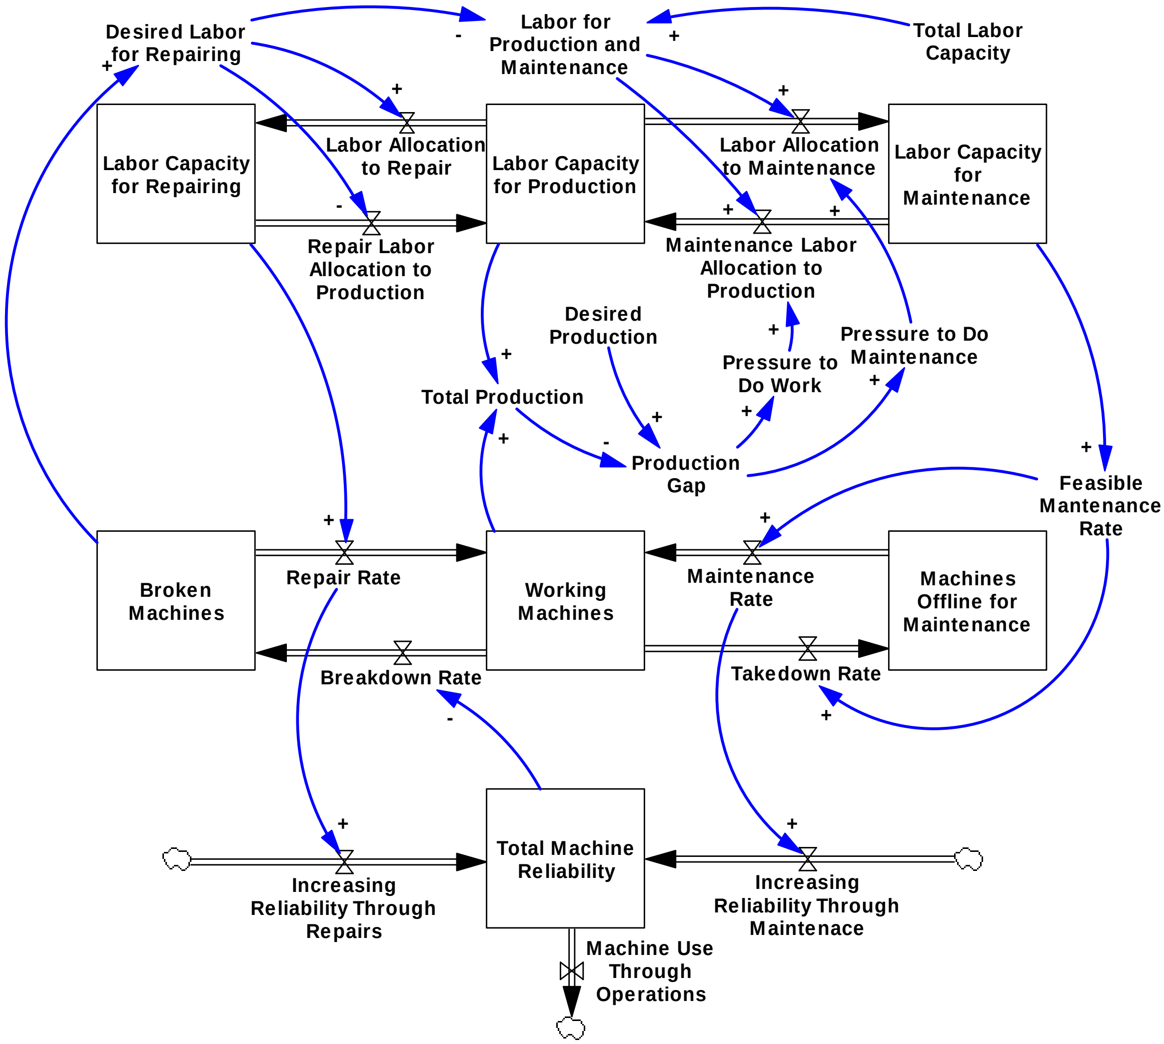

Figure 4 Structure and key variables of the model investigating machine reliability improvement programme with focus on production system resilience

Various internal and external disturbances can lead the organization to change the system configuration and move from one to another domain. In order to illustrate the behaviour of the system to an external stress in this section we will concentrate on a temporary increase in customer orders. The next section will discuss internal disturbances in the system. We may assume the production system undergoes a 20% increase in the desired production level (from 7500 up to 9000 products/month), which occurs in month 5 and lasts for one year. Since the difference between actual and desired production puts a great pressure on production efficiency, more labour capacity is allocated to production operations, simultaneously decreasing labour capacity for machine inspections (Figure 5, Graph#1). Around month 8 more labour capacity must be assigned for repairs to address a sudden increase in a number of broken machines. Eventually only 25 machines are available for production, 75 are broken and none is taken down for maintenance (Figure 5, Graph#2).

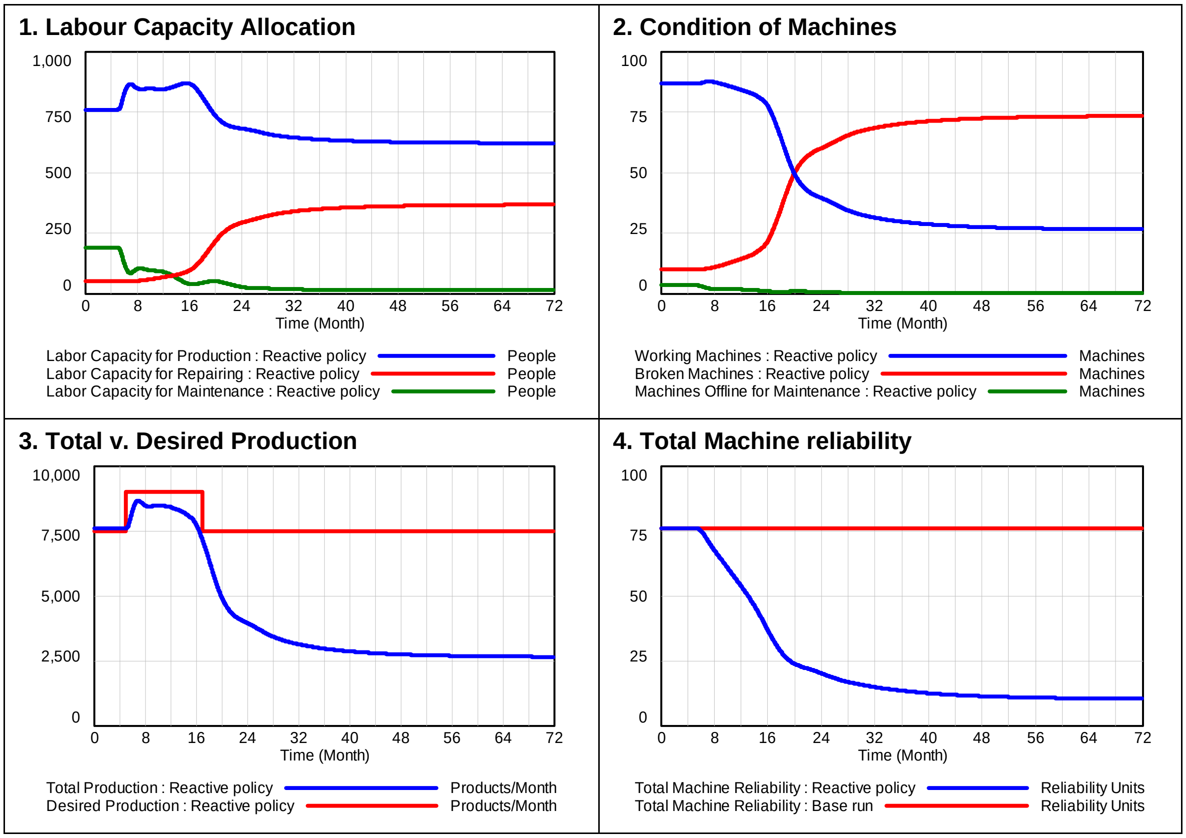

Figure 5 Graphs presenting the reactive response of the production system to disturbance in desired production level

Total production demonstrates 'better before worse' behaviour (Repenning and Sterman 2001). Although the desired production level drops back to its initial value in month 17, the strained system condition makes the feasible production plummet to a highly undesirable level of about 2500 products/month (Figure 5, Graph#3). The reactive policy (fix machines after they break) applied as the answer to increased orders, caused total machine reliability to fall down to a level of 10 (Figure 5, Graph#4). Machine reliability decrease was the reason of more frequent machine breakdowns and led the system to lose its desired functionality.

The effort to change and apply an alternative, proactive strategy can make the system escape the vicious cycle and shift the regime to a desired operational domain. Assume the organization in month 20 shifts the pressure away from the 'working hard' approach to focus more on preventive maintenance. As illustrated in Figure 6, Graph#1 labour capacity for production drops and capacity for maintenance returns to its initial level. After month 20 labour capacity for repairing is still required because the number of broken machines, which increases for the next two months, reaches a peak and falls. This is a result of continuous planned maintenance and significant improvement in total machine reliability (Figure 6, Graph#4). More machines are available for operations (Figure 6, Graph#2) and total production recovers to its desired level (Figure 6, Graph#3).

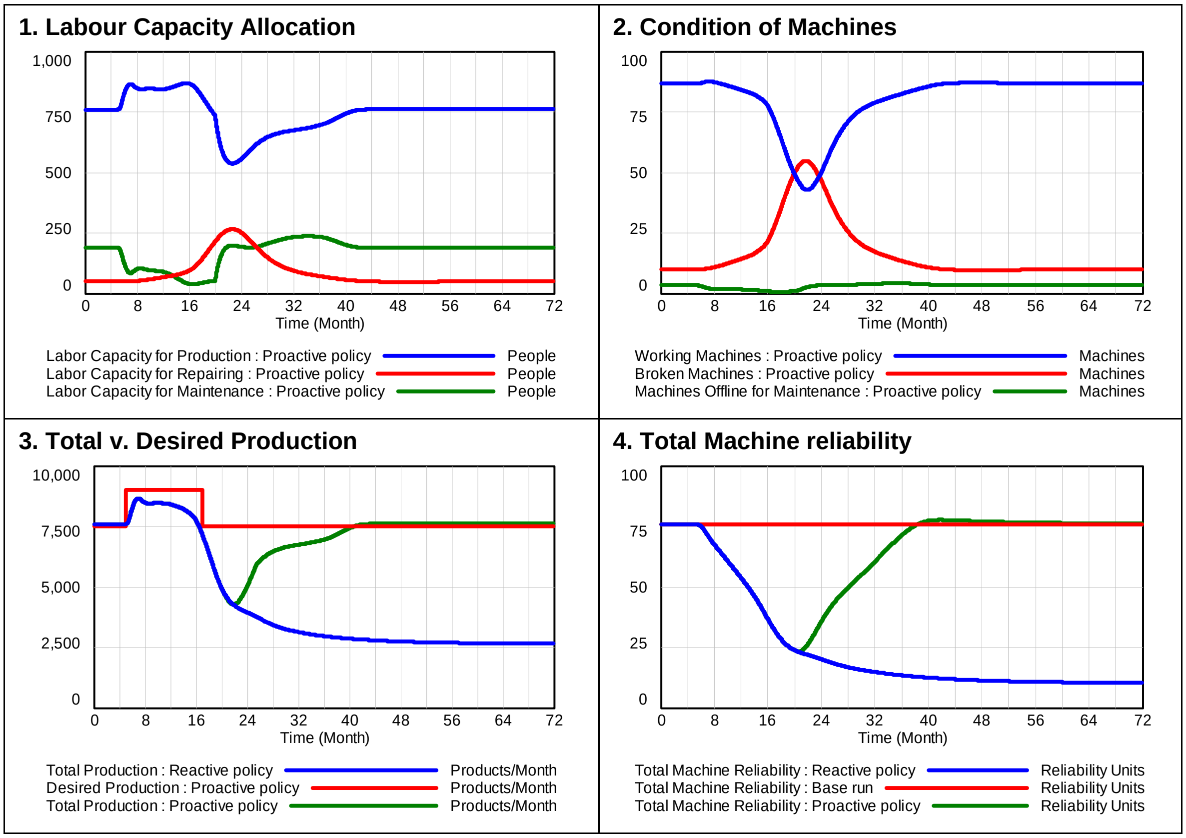

Figure 6 Graphs presenting the proactive strategy to recover after disturbance to the desired domain of operations

Comparing results for both strategies we note that system state variables (e.g. labour capacity, machines, reliability) approach two stable states. Either the organization remains in an undesired domain of attraction, or it achieves a desired one. Which state will the organization reach as the effect of stress and variations, depends on the system configuration – relationships between agents and processes in the system (e.g. 'work hard' policy) as well as availability and allocation of strategic resources (e.g. machines, labour capacity). If the desired configuration is resilient enough it can absorb the impacts of the disturbance and still preserve the functionality of the system.

For both resilience analysis as well as the operational point of view, it is critical to determine which system configuration is desired. Only then can the appropriate measures and steps be taken to increase the resilience of that configuration. The model's example shows how the wrong choice can lead to catastrophic results. However, one has to be aware that this kind of decision is not as easy as it may seem while looking at 'ball and cup' heuristics presented in section 3. Graphs in Figure 5 and Figure 6 provide a deeper sense of understanding of what difficulties decision makers face.

As an example of the model behaviour we analyse the effects of the reactive strategy (Figure 5). Within the first two months this strategy seemed to work considerably well. The total production did not completely cover the desired level but there was a noticeable increase in efficiency. Managers would claim that everything is under control. While people need to get used to new circumstances some more pressure on operators and additional allocation of resources would probably do the work. There might arise some hesitation among decision makers around month 8, when a slight decrease in productivity is noticed, but it would quickly turn into blaming maintenance people, who dare to mention inspections even though all machines are required online for production. About month 15 managers start to worry. The situation is not satisfactory; total production is falling and repair teams seem to 'mess up' with their job. The next month the desired production is coming back to its initial level. On one hand managers try to analyse whether it is the answer of customers to the late orders fulfilment or just market variation. On the other hand they are happy this difficult year has finished, and everything will be in the past. Nothing could be more wrong! The performance over the next four months would lead to changes in management board or to the decision to sell the refinery (as happened in Lima, section 3.2). The results over the next year would certainly lead to plant closure. The reactive strategy, which initially seemed to be satisfactory, happened to be catastrophic.

The dilemma the organization faces concerns day-to-day efficiency and long-term performance: a choice between 'doing things right' and 'doing right things', which at first sight seem to be contradictory concepts. For example, one can take today's machines offline for maintenance, which will impair performance capability, in order to secure functionality and even better performance in the next months. This problem is even more profound in refineries and chemical plants where a high pressure exists to keep production running. They are exposed to sudden and unexpected shifts in the stability of their configurations. Since reliability decreases, the whole system self-organizes and adapts to new conditions. And once the transition happens, it is a real challenge to recover. The System Dynamics model presented here does not encompass all aspects which influence dynamics of the organization and the resilience of both desired and undesired system configurations, e.g. human factor or safety and security matters. However, it outlines the main idea of resilience in production systems and demonstrates its importance for sustained operations.

5. The stylized model of resilience in production systems

Further consideration of resilience and machine reliability improvement will be based on an abstract, stylized model of resilience in production systems. The stylized model is a simplification and focuses on the main characteristics of the model presented in the previous section. It is used to calculate and analyse the potential of the system, the idea of which will be discussed in the following part of the paper. The stylized model explores a concept of quality improvement developed by Repenning and Sterman (1997). Here this idea will be applied in the context of breakdowns and two alternative strategies of machine maintenance: preventive and reactive. The structure of the model is presented in Figure 7.

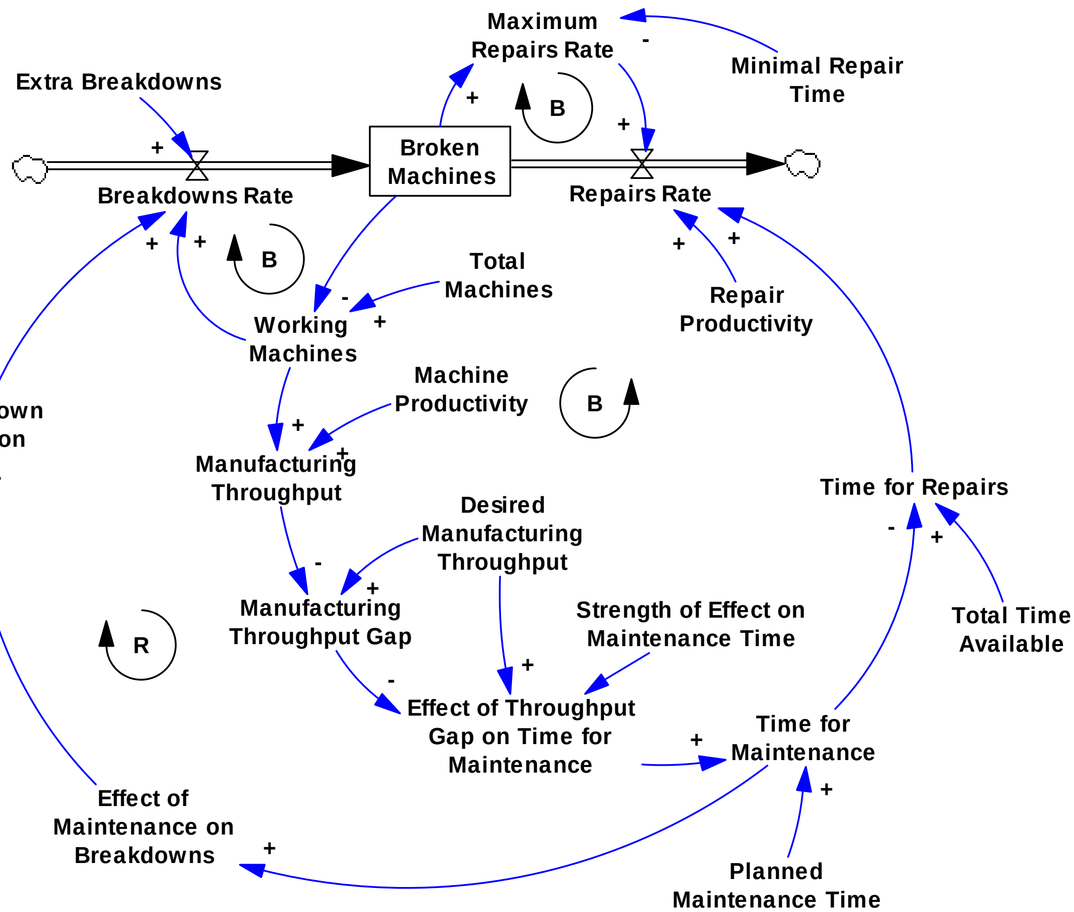

Figure 7 The stylized model of resilience in production systems incorporating the idea of planned and reactive maintenance.

Manufacturing Throughput is determined by the number of *Working Machines* and a constant called *Machine Productivity* (the number of products possible to produce per machine per month). *Working Machines* are defined as *Total Machines* less *Broken Machines*. Any increase in *Desired Manufacturing Throughput* or decrease in actual *Manufacturing Throughput* enlarges *Manufacturing Throughput Gap*. One should perceive *The Effect of Throughput Gap on Time for Maintenance* as a managerial decision leading to preserving or decreasing time for maintenance. The effect was modelled as a function and parameterised. The shape of the effect for different values of *Strength* is presented in Figure 8, Graph#1. *The Strength of Effect on Maintenance Time* represents the strength of the relation between *Manufacturing Throughput Gap* and *Time for Maintenance*. Greater values of the *Strength* lead to bigger *Time for Maintenance* reductions while *Manufacturing Throughput Gap* is growing. When the *Strength* is equal zero the *Effect the Manufacturing Throughput Gap* does not influence *Time for Maintenance* – it is equal to *Planned Maintenance Time*.

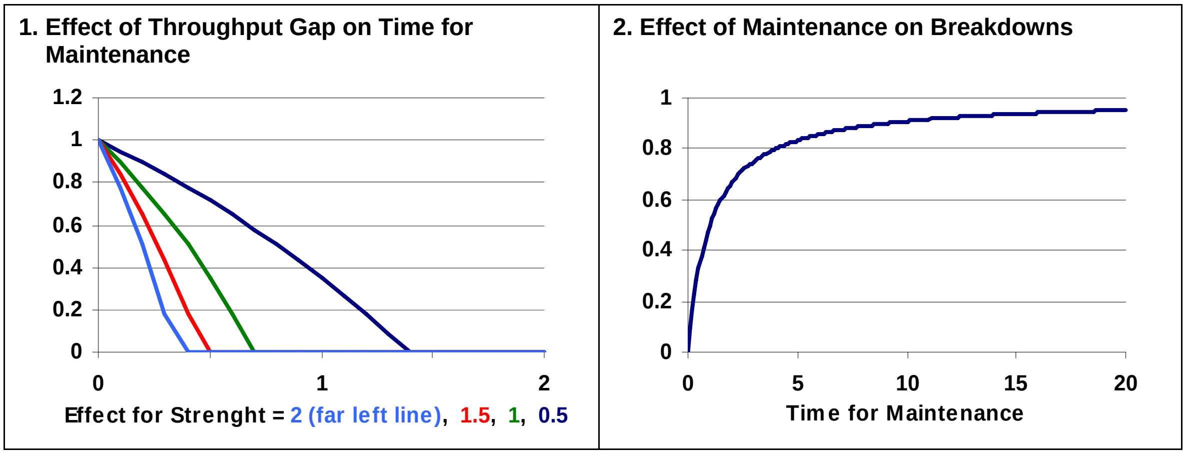

Figure 8 Graphs presenting nonlinear functions used in the general System Dynamics model of planned and reactive maintenance

Time for Maintenance reduces time available for repairs. *Repair Rate* decreases the number of *Broken Machines* and is controlled by *Maximum Repairs Rate* and *Time for Repairs*. Cutting *Time for Maintenance* and increasing *Repairs Rate* closes the balancing loop representing the reactive strategy of machine maintenance.

Time for Maintenance also influences *Breakdown Fraction*. The *Effect of Maintenance on Breakdowns* is presented in Figure 8, Graph#2. Initially, when *Time for Maintenance* increases the *Effect* grows significantly and decreases *Breakdown Fraction*. For greater *Time for Maintenance* the *Effect* increase slows down and asymptotically approaches 1, which means that there is not a significant reduction of *Breakdown Fraction* when *Time for Maintenance* rises from 15 to 20 hours per month. Smaller *Breakdown Fraction* reduces *Breakdown Rate*, which leads to lower level of *Broken Machines*. The policy, when *Time for Maintenance* is preserved in order to decrease machine breakdowns, is considered as planned machine maintenance. The model variables and equation formulations are presented in Appendix B.

Based on the stylized model it is possible to calculate the potential of the system, which is the original idea that inspired the 'ball and cup' heuristic presented in section 3. Figure 9 presents potentials of the discussed system for various values of the parameter *Strength of Effect on Maintenance Time*. On each graph there was placed a 'ball' in order to illustrate the likely system state. Arrows indicate possible disturbances affecting the current state. As the value of the parameter *Strength of Effect on Maintenance Time* increases, the shape of the system potential changes and the alternative, undesirable domain (reactive maintenance) becomes more resilient (the alternative domain becomes wider). At the same time the desirable domain (planned maintenance) is losing its resilience. Eventually, for the last example (Figure 9, Graph#6) even a relatively small disturbance can cause the system transition from planned to reactive maintenance domain. For information on potential calculation see Appendix C.

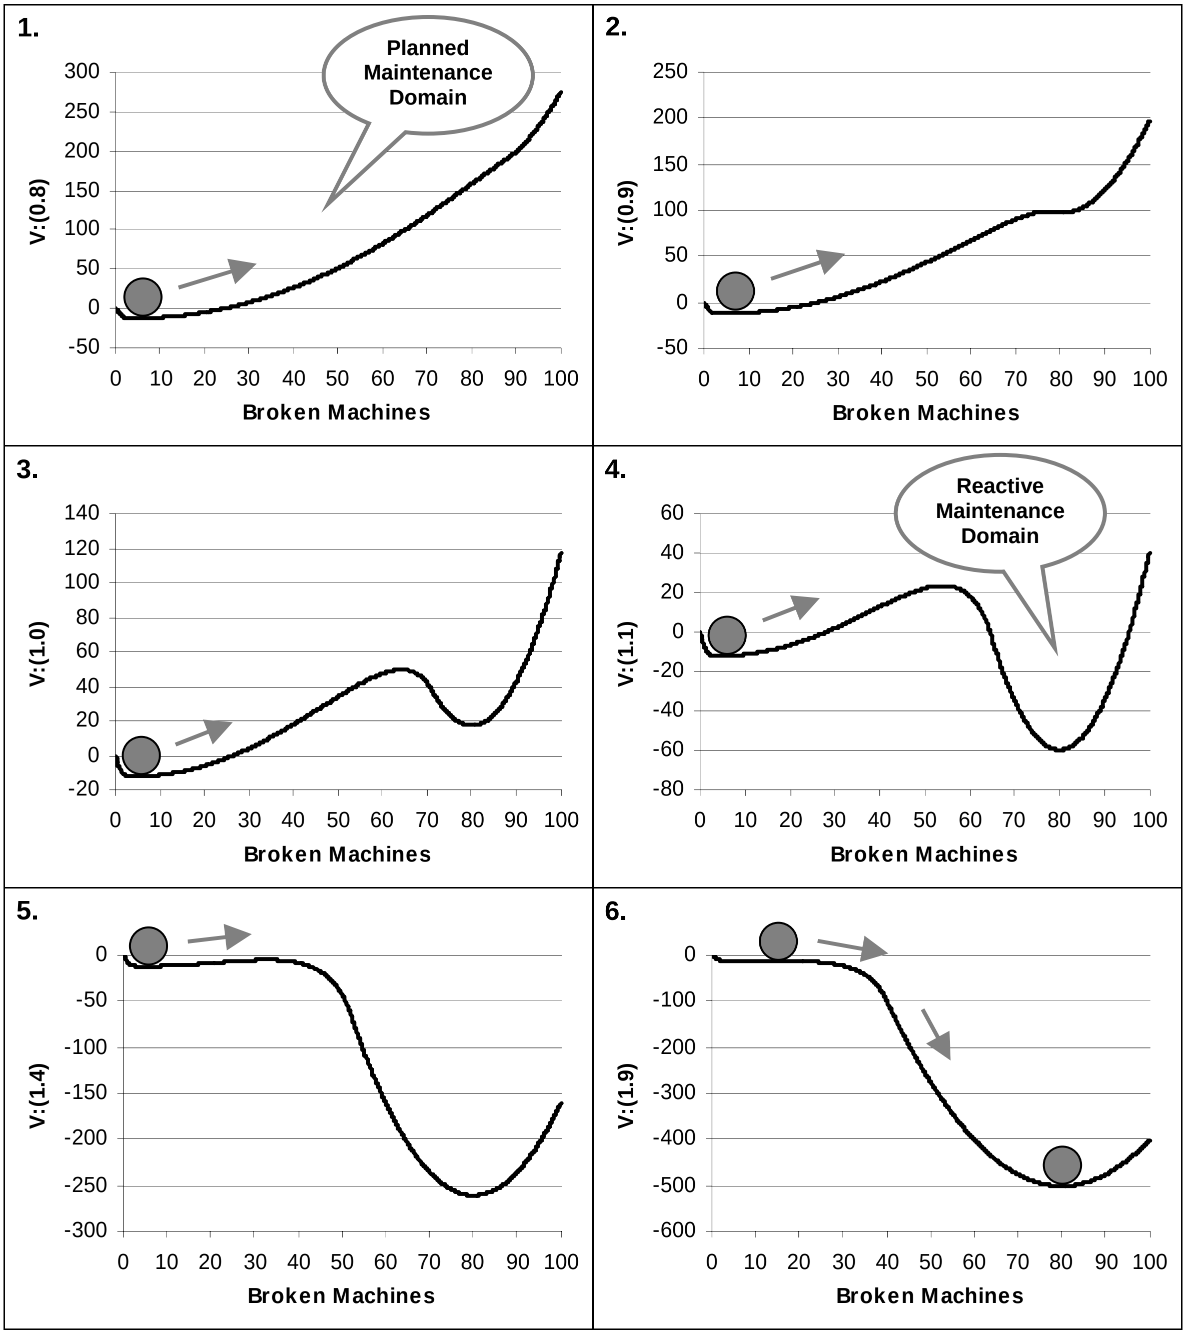

Figure 9 Potential of the system for increasing value of the parameter *Strength of Effect on Maintenance Time*. Added 'balls' represent current state of the system and arrows indicate disturbances.

The dynamics of the transition between desirable and undesirable domains due to two different types of disturbances will be investigated through simulation analysis of the stylized model.

5.1 Scenario 1 – extra breakdowns

The first scenario focuses on sudden, temporary increase in machine breakdown. The simulation analysis of the model starts in equilibrium. There are 100 machines and 5 of them are broken. Machine productivity equals 10 products per month, so feasible throughput is equal to 950 products per month. Since desired throughput is 950 products per month there is no spare production capacity. Such high goals fixed by managers for production personnel and manufacturing infrastructure are currently a common situation in most companies. Highly optimized production processes with reduced slack between operations and minimized buffers are nowadays models of advanced production organization. On one hand extra inventory and spare assets are expensive and, as demonstrated by the lean manufacturing approach, can reduce the quality (Womack and Jones 2003). On the other hand many companies, also refineries and petrochemical plants, seem to ignore the uncertainty inherent in production systems, the reality that production technologies fail unexpectedly and in unforeseeable ways. (Wolf 2001) Following this idea, let us assume the breakdown rate spikes by two machines per month during month 5. The results of this disturbance are presented in Figure 10.

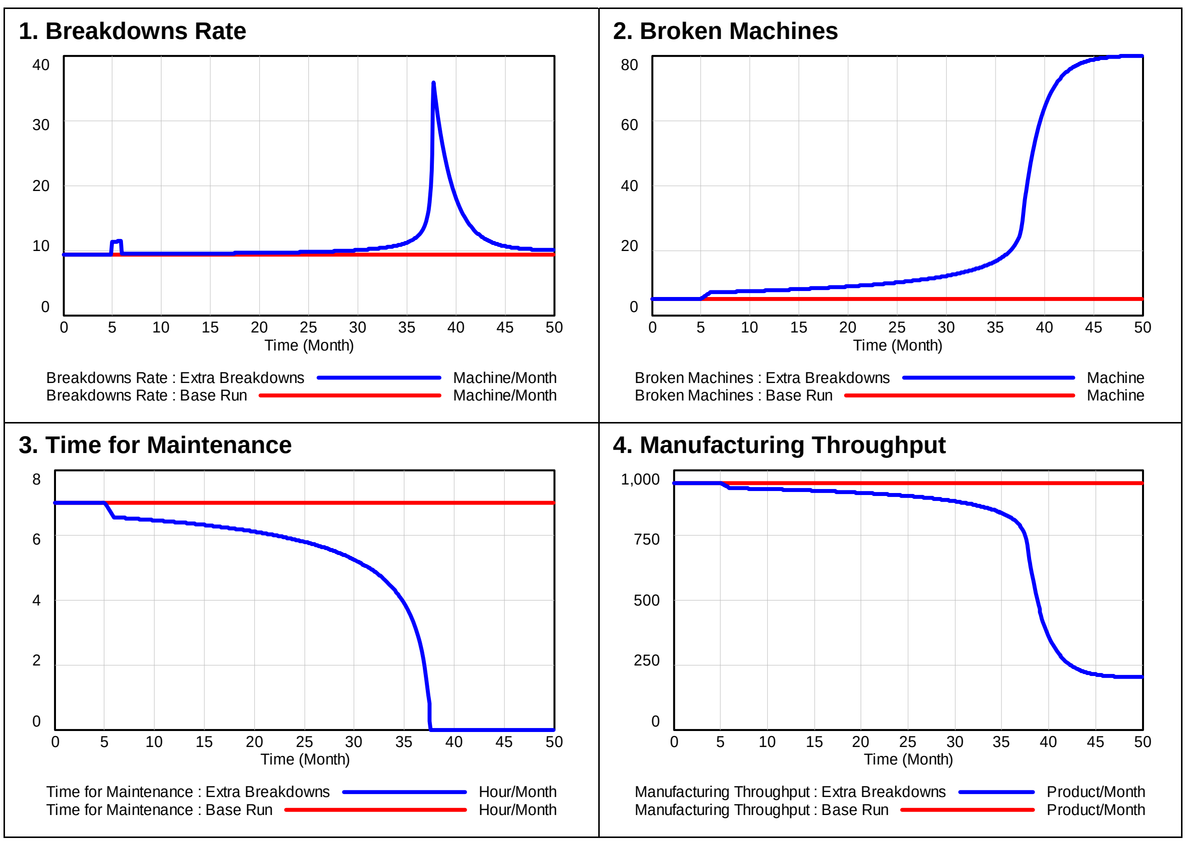

Figure 10 Graphs presenting the response of the system to extra breakdowns (initial system conditions: Desired Manufacturing Throughput = 950 products/month; Strength of Effect on Maintenance Time = 3; Planned Maintenance Time = 7 hours/month)

In month 6 the breakdowns rate drops back to its initial level. This disturbance spike destabilises the production system, however. In the following months the number of broken machines starts to increase (Figure 10, Graph#2) and since there is a strong relationship between manufacturing throughput gap and maintenance time (*Strength of Effect on Maintenance Time* = 3), the time devoted to planned maintenance starts to shrink and around month 35 maintenance is neglected (Figure 10, Graph#3). As a consequence in month 37 the *Breakdowns Rate* reaches a peak of 35 machines per month (Figure 10, Graph#1) and the production system shifts to a highly undesirable configuration. The number of *Broken Machines* reaches a level of 80, bringing the *Manufacturing Throughput* down to 200 products per month.

The potential of the system for the initial model conditions used in the first scenario is presented in Figure 11. The extra machine breakdowns can be perceived as pushing / pulling the 'ball' (current state of the production system), trying to move it to the second domain. The desirable domain (planned maintenance), the system is initially remaining in, is stable but very narrow - not very resilient. Once the system state is destabilised by extra breakdowns it moves to the alternative, undesirable domain of attraction (reactive maintenance) and eventually reaches the minimum for 80 Broken Machines. Since the undesirable domain is very wide (very resilient) it is difficult to restore the system back to its initial state.

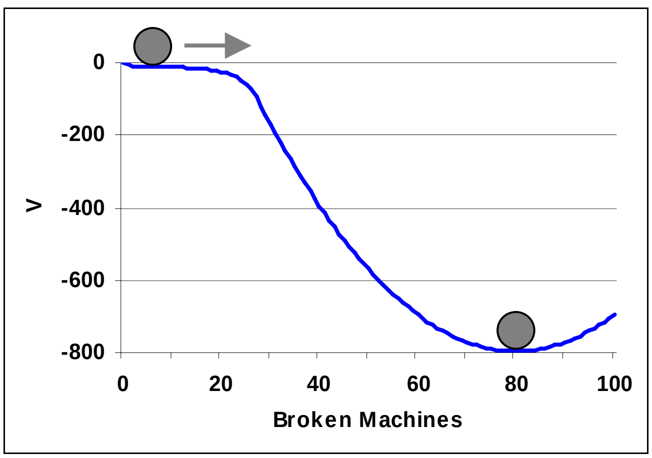

Figure 11 Potential of the system for the initial model conditions used in 'extra breakdowns' scenario. The arrow represents the change in system state in response to disturbance – extra breakdowns; balls relate to initial and final system state.

5.2 Scenario 2 – cutting time for maintenance

Unexpected events like extra machine breakdowns are not the only triggers causing an avalanche of events leading to an undesired system condition. Even in a case of a less demanding production-organization model, when managers set lower goals for manufacturing throughput, production system dynamics can be experienced which are similar to the one presented in Figure 10. The second scenario considers a slow change of time devoted to planned maintenance. Figure 12 shows results of *Planned Maintenance Time* degradation over a period of 50 months.

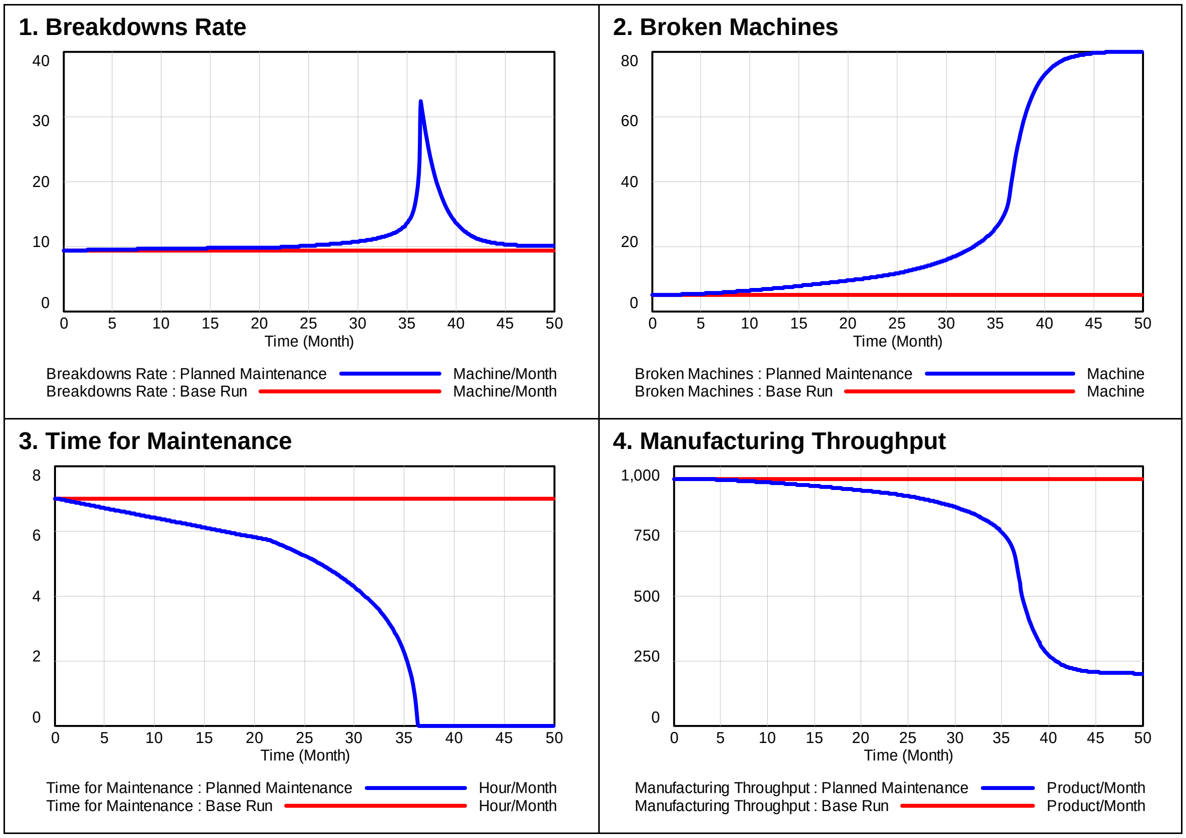

Figure 12 Graphs presenting the response of the system to degradation of planned time for maintenance (initial system conditions: Desired Manufacturing Throughput = 900 products/month; Strength of Effect on Maintenance Time = 2.5; Planned Maintenance Time = RAMP from 7 to 4 hours/month over 50 month)

As the *Planned Maintenance Time* decreases, the actual time spent on maintenance shrinks, causing the machine *Breakdowns Rate* to increase. This in turn leads to a higher level of *Broken Machines*. Emerging and rising differences between feasible and desired manufacturing throughput causes further reduction of the time for maintenance (Figure 12, Graph#3). Around week 37 the *Breakdowns Rate* reaches a peak of about 32 machines per month (Figure 12, Graph#1), and eventually the system moves to the undesired domain.

In the case of the second scenario, the planned maintenance time eroded, causing changes to shape of the system potential. The initial and final potential of the system is presented in Figure 13. Initially the system remains in a stable, resilient domain. At a certain point, following the change in time devoted to planned maintenance, the initial domain suddenly lost its stability, and the system self-organized and moved to the undesired domain (the ‘ball’ rolled down to the bottom of the ‘cup’). Balls in Figure 13 indicate the states of the system at the beginning and the end of the scenario.

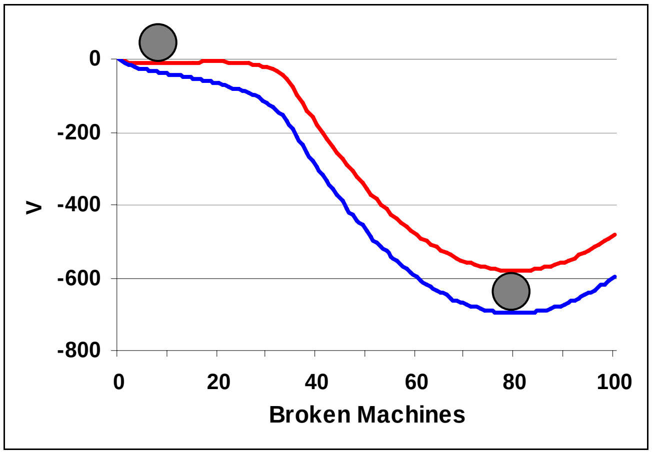

Figure 13 Potential of the system in the 'cutting time for maintenance' scenario. The red line represents the potential of the system for the initial model conditions (*Planned Maintenance Time* = 7); the blue line represents the final potential of the system (*Planned Maintenance Time* = 4). Balls illustrate initial and final state of the system.

The above scenarios illustrated the influence of disturbances and changes on transition of the system between desirable and undesirable domains. In both cases the desirable domain (planned maintenance) was not resilient enough to absorb the impacts of disturbances or changes. As a result the considered production system lost its functionality. Further analysis of the system potential may identify certain conditions which improve resilience of planned maintenance in face of disturbances, e.g. extra breakdowns. At this point it is worth remembering the importance of the statement 'resilience *of* what *to* what' (Carpenter et al. 2001). One has to be aware and clearly specify what system configuration is being considered and what disturbances are of interest because they will impact that particular configuration.

6. Conclusion

The above example illustrates that the main problem is not really the appearance of poor performance but its causes and consequences across the entire production system, resulting from the interdependences between the system's elements. Unfortunately most managers miss this in their excessive focus on performance indicators. Some production systems, such as refineries and chemical plants, due to characteristics of their production processes and products themselves, are more vulnerable and to a larger degree experience effects of any disturbances. Resilience may erode due to production pressure or any other factors that slowly shift a system's configuration, be it at strategic or operational levels of the organization. Failures in managing such complexes might lead to catastrophic results – collapses to domains that are irreversible within an enterprise's horizon of time or money.

This paper presented the concept of resilience as the ability of production systems to experience disturbances and still maintain their functionality. However, resilience should not always be considered as a positive or desirable property. The 'quest for resilience' has to be preceded by the analysis of the current system's overall configuration. Business management needs to classify which structures and processes should be improved or reduced, and what resources preserved, in order to manage resilience, e.g. lower it to facilitate transition or increase it to secure a successful shift to a desirable domain. Managing resilience is not a one-time challenge but a continuous process of dealing with uncertainty.

We found that System Dynamics methodology enhanced with non-linear dynamics tools is a proper tool for conducting analysis focused on management of resilience in production systems. It helps to identify various domains, since there can be more than one alternative domain that the system can shift to. System Dynamics also enables investigation of the functionality of certain system configurations over short and long periods of time. In addition, explicit representation in stylized models of intangible (hard to measure) variables (e.g. pressure to work hard) can contribute significantly to resilience analysis and to making practical improvements in the institutions, policies, tools and methods of production. Modelling for learning and understanding may make very useful contributions to policy even if strict prediction of results is not yet possible.

Acknowledgments

The work presented in this paper was carried out as a part of the doctoral research project of Felicjan Rydzak. It was supported by the grant from The Research Council of Norway within confines of Cultural Agreement Between Norway and Poland 2005/2006, and grant 4T07D01929 from Polish Ministry of Education and Science. This work greatly benefited from discussions with Winston Ledet (Ledet Enterprise), Paul Monus (BP Solar), Doug Parish (Valero - Lima), Warren Burgess (BP), Murray McMillan (The Odyssey Partnership Ltd) and Sean Desmond (INOES Group).

Appendix A

Variables and equation formulations of System Dynamics model presented in section 4:

Average Breakdown Rate = Min Breakdown Rate+(Max Breakdown Rate-Min Breakdown Rate)*Effect of Machine Reliability on Breakdown Rate <*(Machines/Month)/Machines*>

Average Machine Reliability = Total Machine Reliability/Total Machines <*Reliability Units/Machine*>

Average Machine Use Rate = 0.09 <*Reliability Units/(Machines*Month)*>

Breakdown Rate = MIN(Working Machines*Average Breakdown Rate, Working Machines/Working Machines Min Residence Time) <*Machines/Month*>

Broken Machines = INTEG (Breakdown Rate-Repair Rate, Initial Broken Machines) <*Machines*>

Change in Perceived Production Gap =(Production Gap-Perceived Production Gap)/Time to Perceive Production Gap <*Widgets/Month/Month*>

Change in policy = 0.25 <*Dimensionless*>

Change in policy time = 20 <*Month*>

Critical Production Gap = 500 <*Widgets/Month*>

Desired Labour for Repairing = Broken Machines*Labour per Repaired Machine <*People*>

Desired Percent of Labour for Production =Min Percent of Labour for Production+(Max Percent of Labour for Production-Min Percent of Labour for Production)*Effect of Production Gap on Labour Allocation(Perceived Production Gap/Critical Production Gap) <*Dimensionless*>

Desired Production = 7500+STEP(Desired Production Change,5) <*Widgets/Month*>

Desired Production Change = 0 + Desired Production Switch*(STEP(1500,5)-STEP(1500,17)) <*Widgets/Month*>

Desired Production Switch = 0 <*Dimensionless*>

Effect of Machine Reliability on Breakdown Rate = WITH LOOKUP (Machine Reliability Ratio, ([(0,0)-(1,1)],(0,1),(0.1,0.7),(0.2,0.4),(0.3,0.2),(0.4,0.08),(0.5,0.05),(0.6,0.04),(0.7,0.03),(0.8,0.02),(0.9,0.01),(1,0))) <*Dimensionless*>

Effect of Production Gap on Labour Allocation = ([(0,0)-(2,1)],(0,0),(0.226244,0.0280702),(0.538462,0.147368),(0.819005,0.315789),(1,0.5),(1.14932,0.680702),(1.37104,0.849123),(1.63348,0.961403),(2,1)) <*Dimensionless*>

Feasible Maintenance Rate = Labour Capacity for Maintenance*Maintenance Productivity <*Machines/Month*>

Goal for Machines in Maintenance = Feasible Maintenance Rate*Min Machines in Maintenance Residence Time <*Machines*>

Increasing Reliability Through Maintenance = Maintenance Rate*Reliability Increase per Inspection*Reliability Increase Multiplier <*Reliability Units/Month*>

Increasing Reliability Through Repairs = Repair Rate*Reliability Increase per Repair*Reliability Increase Multiplier <*Reliability Units/Month*>

Initial Broken Machines = 10 <*Machines*>

Initial Total Machine Reliability = 76 <*Reliability Units*>

Labour Capacity for Maintenance = INTEG (-Production Labour Allocation, Total Labour Capacity-Labour Capacity for Repairing-Labour Capacity for Production) <*People*>

Labour Capacity for Production = INTEG (Production Labour Allocation-Repair Labour Allocation, Min Percent of Labour for Production*Labour for Production and Maintenance) <*People*>

Labour Capacity for Repairing = INTEG (+Repair Labour Allocation, Desired Labour for Repairing) <People>

Labour for Production and Maintenance = MAX(Total Labour Capacity-Desired Labour for Repairing,0) <*People*>

Labour per Repaired Machine = 5 <*People/Machine*>

Labour Productivity = 10 <*Widgets/Month*Person*>

Machine Productivity = 100 <*Widgets/Month*Machine*>

Machine Reliability Ratio = Average Machine Reliability/Max Machine Reliability <*Dimensionless*>

Machine Use Through Operations = MIN(Average Machine Use Rate*Working Machines, Total Machine Reliability/Min Time of Total Machine Reliability Erosion) <*Reliability Units/Month*>

Machines Offline for Maintenance = INTEG (+Takedown Rate-Maintenance Rate, Goal for Machines in Maintenance) <*Machines*>

Maintenance Productivity = 0.14 <*Machines/(Month*People)*>

Maintenance Rate = MIN(Feasible Maintenance Rate, Maximum Maintenance Rate) <*Machines/Month*>

Max Breakdown Rate = 1 <*(Machines/Month)/Machines*>

Max Machine Reliability = 1 <*Reliability Units/Machine*>

Max Percent of Labour for Production = 0.98+Policy Switch*(-STEP(Change in policy, Change in policy time)) <*Dimensionless*>

Maximum Maintenance Rate = Machines Offline for Maintenance/Min Machines in Maintenance Residence Time <*Machines/Month*>

Maximum Repair Rate = MAX(Broken Machines/Minimum Broken Machines Residence Time,0) <*Machines/Month*>

Min Breakdown Rate = 0.005 <*(Machines/Month)/Machines*>

Min Machines in Maintenance Residence Time = 0.125 <*Months*>

Min Percent of Labour for Production = 0.8 <*Dimensionless*>

Min Time of Total Machine Reliability Erosion = 0.1 <*Months*>

Minimum Broken Machines Residence Time = 1 <*Months*>

Perceived Production Gap = INTEG (Change in Perceived Production Gap, 0) <*Widgets/Month*>

Policy Switch = 0 <*Dimensionless*>

Production Gap = MAX(Desired Production-Total Production,0) <*Widgets/Month*>

Production Labour Allocation = MAX(MIN((Desired Percent of Labour for Production*Labour for Production and Maintenance-Labour Capacity for Production)/Time to Change Labour Allocation, Labour Capacity for Maintenance/Time to Change Labour Allocation),-Labour Capacity for Production/Time to Change Labour Allocation) <*People/Month*>

Reliability Increase Multiplier = WITH LOOKUP (Machine Reliability Ratio, ([(0,0)-(1,1)],(0,1),(0.5,1),(0.6,0.99),(0.7,0.97),(0.8,0.93),(0.85,0.9),(0.9,0.82),(0.95,0.67),(0.975,0.45),(1,0))) <*Dimensionless*>

Reliability Increase per Inspection = 0.3 <*Reliability Units/Machines*>

Reliability Increase per Repair = 0.1 <*Reliability Units/Machines*>

Repair Labour Allocation = MAX(MIN((Desired Labour for Repairing-Labour Capacity for Repairing)/Time to Change Labour Allocation, Labour Capacity for Production/Time to Change Labour Allocation),-Labour Capacity for Repairing/Time to Change Labour Allocation) <*People/Month*>

Repair Productivity = 0.05 <*Machines/Month*People*>

Repair Rate = MIN(Labour Capacity for Repairing*Repair Productivity, Maximum Repair Rate) <*Machines/Month*>

Takedown Rate = MIN(Maintenance Rate + (Goal for Machines in Maintenance-Machines Offline for Maintenance)/Takedown Time, Working Machines/Working Machines Min Residence Time) <*Machines/Month*>

Takedown Time = 1 <*Months*>

Time to Change Labour Allocation = 1 <*Months*>

Time to Perceive Production Gap = 1 <*Months*>

Total Labour Capacity = 1000 <*People*>

Total Machine Reliability = INTEG (Increasing Reliability Through Maintenance +Increasing Reliability Through Repairs-Machine Use Through Operations, Initial Total Machine Reliability) <*Reliability Units*>

Total Machines = 100 <*Machines*>

Total Production = MIN(Machine Productivity*Working Machines, Labour Productivity*Labour Capacity for Production) <*Widgets/Month*>

Working Machines = INTEG (+Maintenance Rate+ Repair Rate-Breakdown Rate-Takedown Rate, Total Machines-Broken Machines-Machines Offline for Maintenance) <*Machines*>
Working Machines Min Residence Time = 0.1 <*Months*>
INITIAL TIME = 0 <*Month*>
FINAL TIME = 72 <*Month*>
TIME STEP = 0.01 <*Month*>

Appendix B

Variables and equation formulations of System Dynamics model presented in section 5:

Breakdown Fraction= Maximum Breakdown Fraction-(Maximum Breakdown Fraction-Minimum Breakdown Fraction)*Effect of Maintenance on Breakdowns <*1/Month*>
Breakdowns Rate= Working Machines*Breakdown Fraction+Extra Breakdowns*PULSE(5,1) <*Machine/Month*>
Broken Machines= INTEG (Breakdowns Rate-Repairs Rate, 5) <*Machine*>
Desired Manufacturing Throughput= 950 <*Widget/Month*>
Effect of Maintenance on Breakdowns= Time for Maintenance/(1+Time for Maintenance) <1>
Effect of Throughput Gap on Time for Maintenance= MAX(2-EXP(Strength of Effect on Maintenance Time*Manufacturing Throughput Gap/Desired Manufacturing Throughput), 0) <*1*>
Extra Breakdowns= 0 <*Machine/Month*>
Machine Productivity= 10 <*Widget/(Month*Machine)*>
Manufacturing Throughput Gap= MAX(Desired Manufacturing Throughput-Manufacturing Throughput,0) <*Widget/Month*>
Manufacturing Throughput= Working Machines*Machine Productivity <*Widget/Month*>
Maximum Breakdown Fraction= 0.5 <*1/Month*>
Maximum Repairs Rate= Broken Machines/Minimal Repair Time <*Machine/Month*>
Minimal Repair Time= 0.25 <*Month*>
Minimum Breakdown Fraction= 0.04045 <*1/Month*>
Planned Maintenance Time= 7 <*Hours/Month*>
Repair Productivity= 0.1 <*Machine/Hour*>
Repairs Rate= MIN(Time for Repairs*Repair Productivity, Maximum Repairs Rate) <*Machine/Month*>
Strength of Effect on Maintenance Time= 1.5 <*1*>
Time for Maintenance= Planned Maintenance Time*Effect of Throughput Gap on Time for Maintenance <*Hour/Month*>
Time for Repairs= Total Time Available-Time for Maintenance <*Hours/Month*>
Total Machines= 100 <*Machine*>
Total Time Available= 100 <*Hours/Month*>
Working Machines= Total Machines-Broken Machines <*Machine*>
FINAL TIME= 50 <*Month*>
INITIAL TIME= 0 <*Month*>
TIME STEP= 0.1 <*Month*>

Appendix C

The model shown in Figure 7 is a first-order system (one differential equation):

$$\frac{dBM}{dt} = BR(BM) - RR(BM)$$

where:
BM - Broken Machines
BR - Breakdown Rate
RR - Repairs Rate

In order to visualize dynamics of the first-order system $\frac{dx}{dt} = f(x)$ the potential ***V*** can be defined as (see Strogatz 2000):

$$f(x) = -\frac{dV}{dx}$$

In our model the potential of the system (***V***) is given by:

$$\frac{dV}{dBM} = -(BR(BM) - RR(BM))$$

The direct form of the potential can be calculated by integrating the right-hand side of the above equation (see Strogatz 2000).

References:

Berkes, Fikret. 2002. Cross-scale institutional linkages: perspectives from the bottom up. In *The Drama Of The Commons*, edited by T. D. E. Ostrom, N. Dolsak, P. Stern, S. Stonich, and E. U. Weber. Washington, D.C., USA: National Academy Press.

Berkes, Fikret, Johan Colding, and Carl Folke, eds. 2003. *Navigating Social-Ecological Systems: Building resilience For Complexity And Change*. Cambridge, UK: Cambridge University Press.

Berkes, Fikret, and Carl Folke, eds. 1998. *Linking Social And Ecological Systems: Management Practices And Social Mechanisms For Building Resilience.* Cambridge, UK: Cambridge University Press.

Carpenter, Stephen R., Donald Ludwig, and William A. Brock. 1999. Management of eutrophication for lakes subject to potentially irreversible change. *Ecological Applications* 9 (3):751–771.

Carpenter, Stephen R., Brian Walker, J. Marty Anderies, and Nick Abel. 2001. From Metaphor to Measurement: Resilience of What to What? *Ecosystems* 4:765–781.

Christopher, Martin, and Helen Peck. 2004. Building the Resilient Supply Chain. *The International Journal of Logistics Management* 15 (2):1-13.

Coutu, Diane L. 2002. How Resilience Works. *Harvard Business Review* May:46-55.

Deming, Edwards W. 1986. *Out of the crisis: quality, productivity and competitive position.* Cambridge: Cambridge University Press.

Griffith, James D., Donovan J. Kuenzli, and Paul A. Monus. 1999. A New American TPM: Leadership requirements for breakthrough change. Paper read at NPRA Refinery and Petrochemical Plant Maintenance Conference, May 25-28, at New Orleans, Louisiana, USA.

Gunderson, Lance H. 2000. Ecological resilience – In theory and application. *Annual Review of Ecological Systems* 31:425-439.

Gunderson, Lance H., Crawford S. Holling, and Stephen S. Light. 1995. *Barriers and bridges to renewal of ecosystems and institutions*. New York: Columbia University Press.

Hamel, Gary, and Liisa Välikangas. 2003. The Quest for Resilience. *Harvard Business Review* September:52-63.

Heifetz, Ronald A., and Marty Linsky. 2002. Managing Yourself: A Survival Guide for Leaders. *Harvard Business Review* 80 (6):65-74.
Holling, Crawford S. 1973. Resilience and Stability of Ecological Systems. *Annual Review of Ecology and Systematics* 4:1-23.
Horne, John F. 1997. The Coming Age of Organizational Resilience. *Business Forum* Spring-Fall 22 (2-3):24-28.
Houshower, Hans. 1999. A Voyage Beyond the Horizon and Back. The Heartland Refinery's Continuous Improvement Story. Source: Paul Monus, personal communication.
Kuenzli, Donovan J., James D. Griffith, and Paul A. Monus. 1998. Proactive Manufacturing - accelerating step change breakthroughs in performance. Paper read at NPRA Refinery and Petrochemical Plant Maintenance Conference, May 19-22, at San Antonio, Texas, USA.
Ledet, Winston J. 1999. Engaging the entire organization key to improving reliability. *Oil & Gas Journal* May (24):54-57.
Ledet, Winston J., and Winston P. Ledet. 2002. Dynamic Benchmarking: Experiencing the Best Practices of Others in Your Plant. Source: Winston P. Ledet, personal communication.
Ledet, Winston, Paul Monus, Tony Cardella, and Warren Burgess. 2005. Modeling Sustainable Organizational Change – Why Did Change at BP Lima Sustain While the Change at DuPont Faded Away? Paper read at 23rd International Conference of the System Dynamics Society, July 17 - 21, at Cambridge, MA, USA.
Nakajima, Seiichi. 1988. *Introduction to TPM: total productive maintenance*. Cambridge: Productivity Press.
Olsson, Per, Lance H. Gunderson, Stephen R. Carpenter, Paul Ryan, Louis Lebel, Carl Folke, and Crawford S. Holling. 2006. Shooting the rapids: navigating transitions to adaptive governance of social-ecological systems. *Ecology and Society* 11 (1):18.
Repenning, Nelson P., and John D. Sterman. 1997. Getting Quality the Old-Fashioned Way: Self-Confirming Attributions in the Dynamics of Process Improvement.
———. 2001. Nobody ever gets credit for fixing problems that never happened: Creating and sustaining process improvement. *California Management Review* 43 (4):44-63.
Rice, James B., and Frederico Caniato. 2003. Building a Secure and Resilient Supply Network. *Supply Chain Management Review* September/October:22-30.
Rydzak, Felicjan, and Edward Chlebus. 2005. Analysis and Management of Resilience in Production Systems. Paper read at PhD Colloquium, 23rd International Conference of The System Dynamics Society, July 17 - 21, at Cambridge, MA, USA.
Sheffi, Yossi. 2005. *The Resilient Enterprise: Overcoming Vulnerability for Competitive Advantage*. Cambridge, MA: The MIT Press.
Sterman, John D. 2000. *Business Dynamics: Systems Thinking and Modeling for a Complex World*. Boston: Irwin/McGraw-Hill.
Stoltz, Paul. 2004. Building resilience for uncertain times. *Leader to Leader* 31 (4):16 -17.
Strogatz, Steven H. 2000. *Nonlinear dynamics and chaos: with application to physics, biology, chemistry, and engineering*. Cambridge, Mass.: Perseus Books.
Van de Brugge, Rutger. 2007. Towards a Resilience Approach for Social Transitions. *Ecology and Society*. *in review*.
Venkat, J. 2006. *An Introduction to Total Productive Maintenance*. Plant Maintenance Resource Center 2003 [cited 15 January 2006]. Available from http://www.plant-maintenance.com/articles/tpm_intro.shtml.
Walker, Brian, Stephen Carpente, John Anderies, Nick Abel, Graeme Cumming, Marco Janssen, Louis Lebel, Jon Norberg, Garry D. Peterson, and Rusty Pritchard. 2002. Resilience management in social-ecological systems: a working hypothesis for a participatory approach. Review of Reviewed Item. *Conservation Ecology* 6 (1).
Wolf, Frederick G. 2001. Operationalizing and Testing Normal Accident Theory in Petrochemical Plants and Refineries. *Production and Operations Management* 10 (3):292-305.
Womack, James P., and Daniel T. Jones. 2003. *Lean thinking: banish waste and create wealth in your corporation*. New York: Free Press.