

DESIGN OF ROBUST SERVICE OPERATIONS USING CYBERNETIC PRINCIPLES AND SIMULATION

D. C. Wynn, S. Cassidy and P. J. Clarkson

Keywords: service design, cybernetics, uncertainty, robustness, simulation, Cambridge Advanced Modeller (CAM)

1. Introduction

Information flows in a service organisation allow business units to co-ordinate their response to changes in the operating environment. Processes and interactions can be designed so that the right information flows to the right people, at the right time to make effective decisions regarding job priorities and allocation of limited resource. This paper develops an analysis framework and simulation approach to identify the internal information flows an organisation needs to “tune itself” for changing conditions, thus making itself more robust to uncertainties. The ideas are developed and illustrated through a case study with a major telecoms company.

The telecoms company, like other organisations that develop and deliver complex services, does so through many interlocking processes. Certain processes develop new products. Others handle customer requests. Others balance the organisation’s limited resources across projects that request them. Different processes are triggered by different events or observations, and serve different purposes to the organisation. Each process enables a different aspect of the overall business function. They are co-ordinated by information flowing between them.

If such a “process system” operated in a well-known and well-understood environment, it could in theory be tuned to give optimal performance for that particular pattern of demand. For instance, all the business units might be given the particular resources they would need to operate most effectively and efficiently. However, most service organisations are not so static. The products they offer have lifecycles throughout which volumes change. A new product usually involves new technology, such that new processes must be put in place quickly and refined later. Unforeseen or unexperienced events, such as a sudden increase in customer demand, cause resource shortages and workflow congestion. Decreases in demand lead to over-resourcing and inefficiency.

Mature organisations incorporate metrics and balances to handle such issues. If congestion in a certain process increases, this will be noted and more resource sought to ease the problem. It is usually necessary to balance such actions across a product/service portfolio, considering how re-allocating scarce resource to relieve one area of a business can cause knock-on problems elsewhere. Allocation decisions also need to take into account the company’s plan for the future of its portfolio.

To take relevant issues into account while making these decisions, information must flow across business units in the organisation. In the example given above, information flow occurs from the business units responsible for portfolio and strategy into that responsible for resource planning. Appropriate flows allow decision-makers to take balanced decisions, by helping them understand the impact of their actions on their colleagues and the goals of the organisation. However, all information cannot be sent to everybody, because information flows in an organisation incur costs. Too much information is difficult to assimilate—especially if it appears to be contradictory, as usually happens

when a decision involves trade-offs between multiple objectives. The communication must provide clear benefit to both the provider and consumer of information, or it will not be accepted.

2. Case study

Our analysis of information flows in service operations draws on a case study of a data connection service offered by BT to other businesses. In overview, each request to provide this service is handled by a process called “Lead to Cash” (L2C). A customer places an order for a new service or for modification of an existing service. The order is handled by a sequence of processes in the operations business unit. For example, the customer details and credit-worthiness must be validated, the installation must be designed, customer premises equipment must be procured and installed, and the service initiated and tested. Some of these jobs may be very quick. In other cases further checks and approval might be sought, according to the characteristics and complexity of the order. Meanwhile, the physical infrastructure and network capacity on which the data service runs must be maintained and expanded as the number of installed services increases. When a service is uninstalled, the capacity and equipment must be recovered for use elsewhere. Other processes deal with fixing problems encountered by customers. Human resource for running all these processes must be maintained at an appropriate level, and balanced against the needs of other business areas that draw on the same people. Many challenges underlying data service provision relate to managing variability of orders. For instance, some installations are essentially standard, while in other cases the installation is personalised to the customer’s requirements. This variability in order types creates potential for flow problems in the data service provision processes, as certain orders may take longer to process than others. These can cause congestion in flowing jobs through the L2C process. Other issues can affect the capacity of the process system. For instance, major projects may draw resources away from normal service provision. Many of these issues are common across the industry—as reported, for instance, in the work of [Akkermans and Vos 2000], [Le et al. 2010] and [Delgado et al. 2011].

BT is keen to avoid the delays and knock-on consequences that can arise from variability and uncertainty. It also needs to run the processes as ‘hot’ as possible, to keep resourcing costs under control. Our case study aimed to help balance these competing concerns by identifying how an appropriate structure of information flows and a suitable set of operating policies can be designed.

3. Cybernetic principles applied to service operations

Information flows between business units naturally involve communication between the people who perform given functions. In considering such communication flows, this paper does not focus on the form of information encoding or the mechanisms of communication, which could include custom information systems, email, or informal networks between people. Rather, we concentrate on identifying which channels should be provided, and on understanding whether the information thus transmitted allows process participants to make decisions that are ‘better’ for the organisation as a whole. We approach this analysis of information flows in context through the lens of cybernetic systems theory.

The cybernetic view of an organisation presents it as a system of people and processes that interacts with its environment through information flows across the system boundary [Beer 1972]. The organisation-as-a-system may be decomposed hierarchically into divisions and sub-divisions (business units). Some such divisions interact only with other parts of the organisation, while others have direct interfaces to customers, suppliers, competitors etc. The cybernetic perspective supposes that an organisation regulates its own processes through feedback and feed-forward of information between different business units. For instance, one process might observe another. If bottlenecks in processing are seen to occur, the first process might respond by assigning more resource to the second.

Thus, from a systemic point of view, the organisation can be considered as a set of intertwined process cycles that interact to regulate each other, as suggested by Figure 1. A given business unit as indicated in Figure 1 can be considered to comprise a structure of *processes*, *information flows* and *models*. Some of these processes may directly trigger others. For instance, when an order is placed by the customer it flows into the in-tray of an L2C process and will be handled as soon as time allows. Other interactions are asynchronous, because some processes interact through the access and modification of

shared information stored in *models*. These models represent stores of information, material or resource within the organisation. The stores can be concrete, such as an information system, or can represent more abstract forms of persistent information, such as the mental models of process participants. Over time, different processes in an organisation draw upon and modify the models it contains. Models thus allow asynchronous processes that operate on different cycle times to communicate. For instance, a resource allocation process might occasionally draw upon ‘mean utilisation’ statistics stored in an information system, using this information to guide a decision. Those same statistics might be updated several times a day by the processes that depend upon the resources.

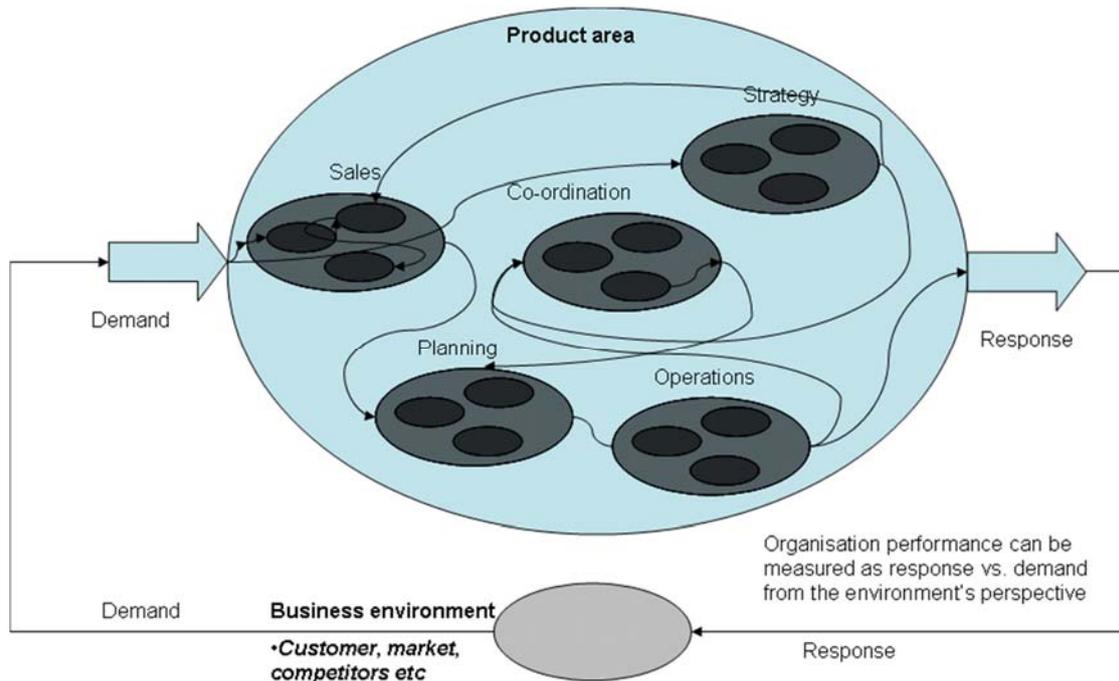


Figure 1. View of an organisation as a cybernetic system of interconnected business units

According to this perspective, each process associated with a business unit can be viewed as a set of functions (tasks) and the information flows between them. A given process is initiated by an event, which could be generated by a business function, by an external stimulus, or by observation of a particular state in the organisation. Once generated, the event flows through the different functions in the process, perhaps being routed by logic gates, and perhaps being processed several times by the same function if the process is ‘iterative’. Eventually, the message will reach the end of the process and a model will be updated. Many such processes cross the boundaries of business units.

So-called cybernetic principles [Wynn et al. 2010] can help to model this situation more concretely, and to predict the qualitative implications of given configurations. Key principles used in our analysis are described in the subsections below.

3.1 Mechanisms of feed-forward and feedback

Perhaps the most basic principle of cybernetics states that a system (here, an organisation) delivers desired behaviour in an uncertain and changing environment through feedback and feed-forward mechanisms. The system has a set of inputs which are converted into a set of outputs, and some control parameters that can be used to adjust its behaviour. In feedback, the difference between the system’s outputs (in this case, aspects of the process performance) and their target values is measured, thereby creating information about the interaction between system and environment. This information is processed and used to change the control parameters to bring the output closer to its desired state. In feed-forward, the process is similar but measurements are made at the system input. This offers the possibility of mid-stream course correction, in advance of problems manifesting at the system output.

3.2 Importance of control laws

In feed-forward and feedback, information must be processed to determine the adjustments which are made in response to a given input. Processing can be based on apparently simple and essentially generic “control laws” (for instance, respond to an increase in process congestion by increasing the amount of available resources in proportion). It can also embody a more complex “model” of the system. In an organisation, as with a physical control system, the appropriate laws depend on the dynamics of the system to be controlled as well as the environmental disturbances expected. In an organisation presented as in Figure 1, dynamics arise largely from the interactions between asynchronous processes in the organisation, notably lead times, congestion effects, and the stores of information kept by the company.

3.3 Importance of inertia and delays on system stability

Most controlled systems have some form of inertia, which means that it takes time for a change in the control parameters to be manifested and to become visible in the output. Forms of inertia in a business process include the work that builds up in queues when insufficient resource is available, information stores that the process contains, and the base of ‘integrative’ resources, such as infrastructure that must be maintained. If the target values or the inputs change, most inertial systems governed by feedback control will undergo oscillations before gradually settling close to the desired value. A system which does not approach the desired value more closely as time progresses is said to be unstable. One way to quantify system stability is to measure the time required for the output oscillations to settle within a given percentage of the desired value.

In general, the stability of a system is reduced by increased delays between cause and effect. In an organisation, regulatory processes with long lead times thus tend to cause excess oscillations. Oscillatory effects can be seen, for instance, when too much resource is allocated too late to deal with a temporary increase in workload, only to be left with excess capacity that must be paid for once the job is finished. Generally, the problems associated with instability can be minimised either by reducing lead times or by taking better account of the system dynamics in the control laws.

3.4 Principle of requisite variety

This general principle states that, to respond to a certain level of variety in its environmental conditions, a system must have requisite variety in its internal structure and/or control laws. In the context of this paper, this can be interpreted as stating that dealing with additional sorts of disturbance in the process environment requires additional flows and considerations within the business.

To illustrate, consider the examples given in Section 1: if no variation in demand for a service exists, no measures and balances are required in the organisation; a very simple process is sufficient, operating like a production line with defined work rates and steady flow. However, as the structure of the environment becomes more complex, characterised by different patterns in the demand, additional resource allocation loops and other controls are required.

3.5 Summary

Viewing an organisation as a cybernetic system highlights that each business unit and process plays a different regulatory role. These roles can be considered to construct a simulation of that organisation. The major regulation decisions (such as when, where and how much resource should be added) can be identified and included in the model. The broader context can be considered to identify what additional information could be used to assist these decisions. The impact of those proposed flows can be explored through simulation to assess their values under different operating conditions.

4. Levels of response to uncertainty

A conceptual framework was developed to structure the case study analysis. The framework, summarised in Figure 2, proposes that an organisation may be capable of four increasingly sophisticated levels of response to uncertainty and unexpected change in operating environment.

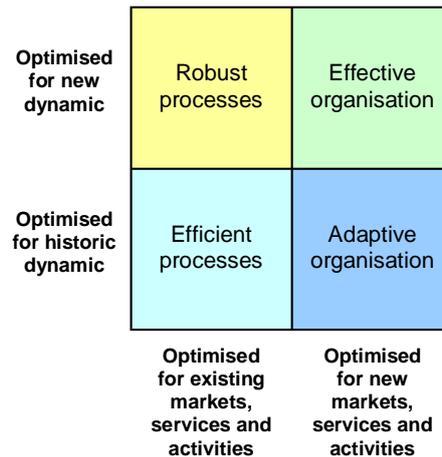


Figure 2. Categories of environmental dynamic and ideal organisation responses

4.1 Efficient processes

A process system may be designed to operate efficiently under a given operating regime, where uncertainties exist but can be localised and quantified. For instance, although a service operation is usually subject to variable demand, historical data (of which much is available in BT) allows the numbers and types of customer request to be modelled using probability distributions. Using simulation, it is possible to tune the processes to handle that pattern of demand as efficiently as possible, to reach a desired trade-off between performance indicators such as lead time and cost.

4.2 Robust processes

A process system may also be designed to cope with hitherto unexperienced variations, which occur when demand moves outside its normal window of operation. For instance, the expected pattern of demand might be disturbed by an unexpected and unplanned event, such as a significant increase in orders or a reduction in resourcing levels. Several strategies might be used to cope with the possibility of such events, including:

- 1. Relax the performance targets.** Goals will become easier to meet [Chalupnik et al. 2009]. This may have knock-on consequences, for instance by reducing the attractiveness of a service to its customers. However, this is not always the case: for instance, a customer who has to wait at home for an installation may be mostly concerned that an engineer should arrive on the designated day; adding a few days buffer time when the order is placed might reduce risk of late delivery with little negative impact on satisfaction.
- 2. Make the process less efficient.** Tuning a process system to be less efficient, for instance by having extra resources on stand-by, allows more headroom to cope with unexpected needs. Variations on this have been studied in the context of organisational slack [Lawson, 2001].
- 3. Make the organisation re-tune itself after a problem is spotted (feedback).** If a process system contains more ‘intelligence’, it may be able to identify changes in operating conditions and re-tune its own behaviour. For instance, consider an unplanned increase in demand, which may create a backlog of work; if the backlog is monitored, the increase can be spotted and additional resource assigned to reduce it. On a higher level, if these adjustments repeatedly prove ineffective, the policy for allocating resource might be changed.
- 4. Feed information between business units to anticipate problems (feed-forward).** If a change in operating conditions can be anticipated, for instance by the sales department who have some ability to foresee change in customer demand, appropriate actions can be put in place to compensate in good time and major disturbances can be avoided.

4.3 Adaptive organisation

An adaptive organisation is able to make decisions at all levels that consider issues that span multiple products and timeframe concerns in a portfolio. For instance, it may rationally decide to allocate more

resource to improve support processes, recognising the ultimate impact on many primary processes and on specific product areas. An adaptive organisation can readily make trade-offs between short-term and long-term goals, for instance, by one process relinquishing resource now so that it can be used to improve performance later. It is able to rebalance the mix of what is done, rather than simply tuning the existing activities. It is also able to consider how it defines the market for its services through its own interactions with customers.

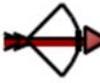
4.4 Effective organisation

An effective organisation is able to re-design itself in response to major events that are outside prior experience and which fall outside the scope of the first three levels. Typically, this cannot be easily achieved through a process-oriented analysis. Efforts towards an effective organisation might focus on how the team can dynamically create new ways of working appropriate to truly unexpected situations. Such analysis is outside the scope of this paper, although is being considered in our ongoing research.

5. Simulating service operations

Analysis of BT processes with respect to these levels of response to uncertainty was undertaken using discrete-event simulation. A graphical notation was developed to allow models capturing uncertainty, information flows and regulatory dynamics to be tied to a picture of the organisation recognisable to the business. This notation is summarised in Table 1. A new toolbox for the Cambridge Advanced Modeller software [Wynn 2007] was created to develop and simulate the models.

Table 1. Types of element provided in the modelling framework

Name	Icon	Properties	Description and usage
External event		Cycle time Output variables	An event from outside the boundary of the modelled system, such as a customer request, that can initiate a process. Events have a certain cycle time (the delay between successive firings).
Proactive function		Cycle time Output variables	A regular event inside the business that creates a message, initiating a process. For instance, a meeting held every three weeks.
Triggered function		Cycle time Trigger condition Output variables	Model(s) are observed and the contents considered on a regular basis. If certain conditions are met, a message is created and will flow into one or more reactive functions.
Reactive function		Lead time Channel count Output variables	Represents a job that must be done to process a message. Triggered when messages are received from another function. Each message takes a certain lead time to process. The function has a certain number of channels, each of which can process one message at a time. Messages are stored in a queue until a channel is available to process them. When a message is processed, an output message is created and flows into one or more downstream reactive functions. Or, a model is updated directly.
Switch		Output selection	Routes each message down one or more output channels, according to some user-defined functions. A switch has no time delay and infinite processing capacity.
Message channel		List of variables	Represents a channel of infinite capacity, from which messages can flow from the upstream function that creates them into one or more reactive functions that process them.
Model		List of variables Initial values	A store of information. Models contain variables that can be updated by the completion of an event, function or trigger that feeds into them. Model variables can be used to influence the properties of any element.

Two workshops were held to elicit the structure of business units, processes and organisation flows in the BT data service using this modelling toolbox.

The first workshop lasted 4.5 hours. Participants were asked to create a detailed model of the service provision by labelling sticky notes and placing them on a large sheet of paper. The emphasis was placed on building a picture of information flows across business units, and accordingly the workshop participants only decomposed a given process when required to highlight such flows. The second workshop was similar in duration to the first, but focused on validating the process structures. Two methods were used; firstly, talking through the map; and secondly, looking for “models” that were used to inform decisions, but that were not updated by any processes. The workshop participants were also asked to estimate numeric data that would later be used to populate and calibrate the simulation. This included the lead time of each function as well as the frequency of certain operations. Some calibration data were obtained from company databases. Other information was estimated by the participants who were familiar with the process.

A key concern while estimating numeric data was to identify the root causes of uncertainty in duration of each job (ie, each reactive function). In each case, workshop participants were asked to consider if variability could be attributed to properties of the work at hand. One property thus identified was the order ‘complexity’. A complex order would require more validation steps, and take longer to design, than a straightforward ‘off-the-shelf’ installation. In the model, this parameter would be defined when each order was received, and then influence the duration of several steps throughout the L2C process.

Finally, the participants were asked to identify key performance indicators (KPIs) for the data service provision and estimate expected values for each KPI; this information was later used to help calibrate the simulation.

The resulting model, shown in Figure 3, provides a simplified representation of the organisation and its processes. It was felt to capture the most pertinent issues with respect to the research objectives, and was agreed by the workshop participants to provide a good foundation for simulation.

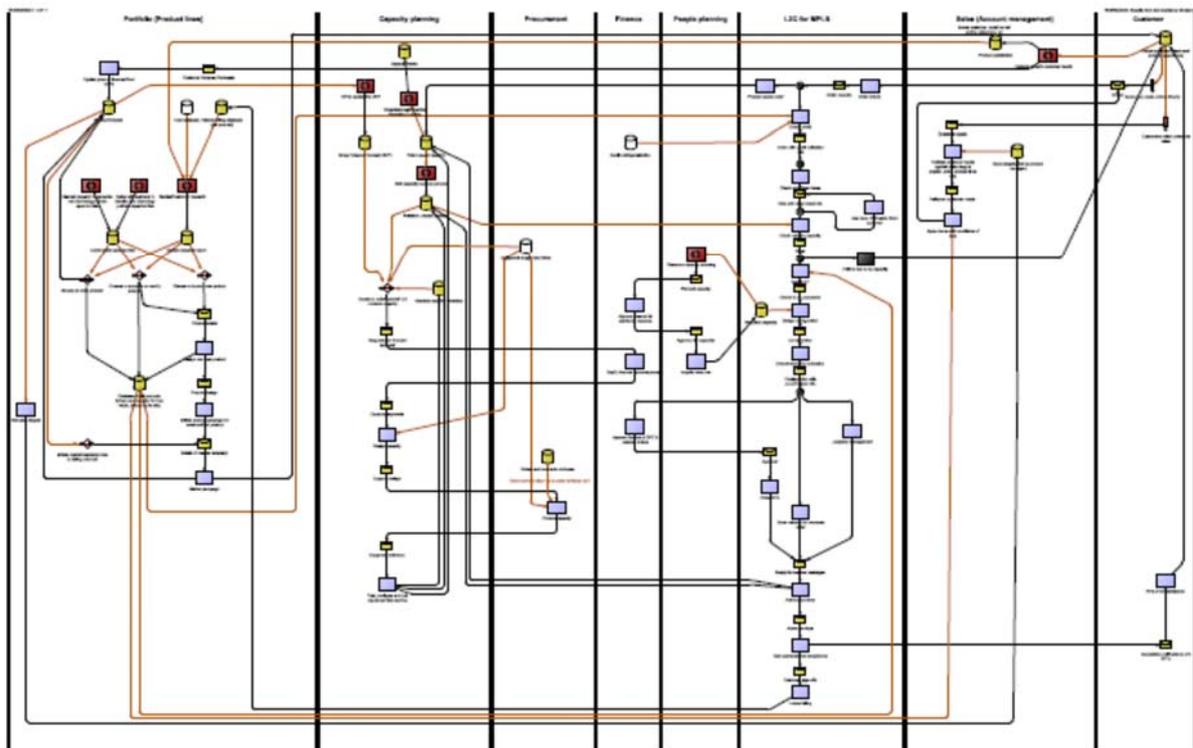


Figure 3. The simulation model of data service operations. Swimlanes delineate business units

The understanding generated from the workshops suggested how the process regulation should operate, but it was not practicable to elicit concrete rules by which every decision is made. For instance, no rules were identified to explain exactly how much resource would be sought to

compensate for a given build-up of work in operations. It was therefore decided to represent the regulation functions using simple and generic control laws intended to capture the essence of the decision-making as had been described during the workshops. These laws were verified to give a sensible representation of the process dynamics, by using simulations to check that the processes would respond in a common-sense way to a range of operating conditions. Three main regulation loops were implemented:

- 1. Resource regulation.** The decision regarding how much resource to add to the L2C process following an observed increase in congestion was implemented as a PD control law. The proportional and derivative gains were chosen by manual tuning, using trial-and-error to find parameters that brought queue length under control quickly and with minimal oscillation following a disturbance. The chosen parameters were trialled and shown to give acceptable results under a range of conditions: flat mean demand (numbers of orders per day); steps in mean demand of different magnitudes; and ramps in mean demand of different slopes.
- 2. Network capacity extension (“Capex”).** Capacity is required to handle the number of services that are installed in total, and their bandwidth requirements. It is thus ‘used up’ as each new order is handled. In the model available capacity is observed regularly; if it drops below a threshold, a defined ‘increment’ of new capacity is ordered. After some delay, this becomes available for use by L2C.
- 3. Demand regulation.** The main regulation loop implemented in portfolio management is responsible for balancing orders by initiating a market campaign once the number of orders being placed falls below a certain threshold. After some delay, this is assumed to cause a jump in demand, which then gradually tails off at a constant rate. In the model, this effect only comes into play during simulations with long timescales spanning several years.

Interactions between these regulation loops and other process logic, combined with the general objective to reduce unused capacity and resource usage without jeopardising the customer experience by causing delays to service installation, contribute significantly to the complexity of the process dynamics. For instance, if there is insufficient network capacity to install a service in the model, that order will ‘block’, apparently reducing the demand on L2C resources. This complexity is further increased by the variability in demand and changing operating conditions for the system as a whole.

Due to the complex interactions and many variables involved in the model, its configuration was not an easy task. After some consideration of different ways to break down the problem, the model was progressively configured for simulation by building in the processes and regulation logics outlined above one-at-a-time, starting with the L2C process and working leftwards across the process map. Many simulation runs were executed under different patterns of demand, the results were studied in each case, and problems in the emerging configuration were progressively identified and corrected.

6. Analysis of data circuit provision

The model enabled application of the conceptual framework described in earlier sections to explore opportunities for improvement of the data service provision. A set of analyses was run to consider the first three of the four levels of demand dynamic discussed in Section 4. This section discusses the analysis undertaken to explore robust processes (level 2), which is representative of the others.

After configuring and calibrating the model as discussed above, it could be used to test the process’ response to different operating conditions. To begin the analysis of the design for robustness, an unanticipated step in demand was assumed. Figure 4 shows a single simulation run to illustrate how the balancing loops are able to account for the difference in throughput before and after this step. The bottom timeseries in this figure shows the demand changing over time. This is effectively the ‘input signal’ for the process system, in this case exhibiting normal variability overlaid with a step function.

After the step, queues start to build up at the L2C bottleneck causing an increase in lead time. This can be seen in the top two plots of Figure 4, which show respectively the queue length at the L2C bottleneck at a given time and the mean lead time for processing each order. During simulation, the resource monitoring process observes the increase in queue size, and adds resource capacity as shown in the 3rd-from-top timeseries. The additional resource capacity allows the queue to be reduced, after which resource is rebalanced and all the processes reach a new steady-state of processing the greater

number of orders. The Capex process self-adjusts and occurs more frequently as shown in the 4th-from-top timeseries. Similar behaviour was shown to occur for various step sizes, amounts of noise, and control parameter combinations tried.

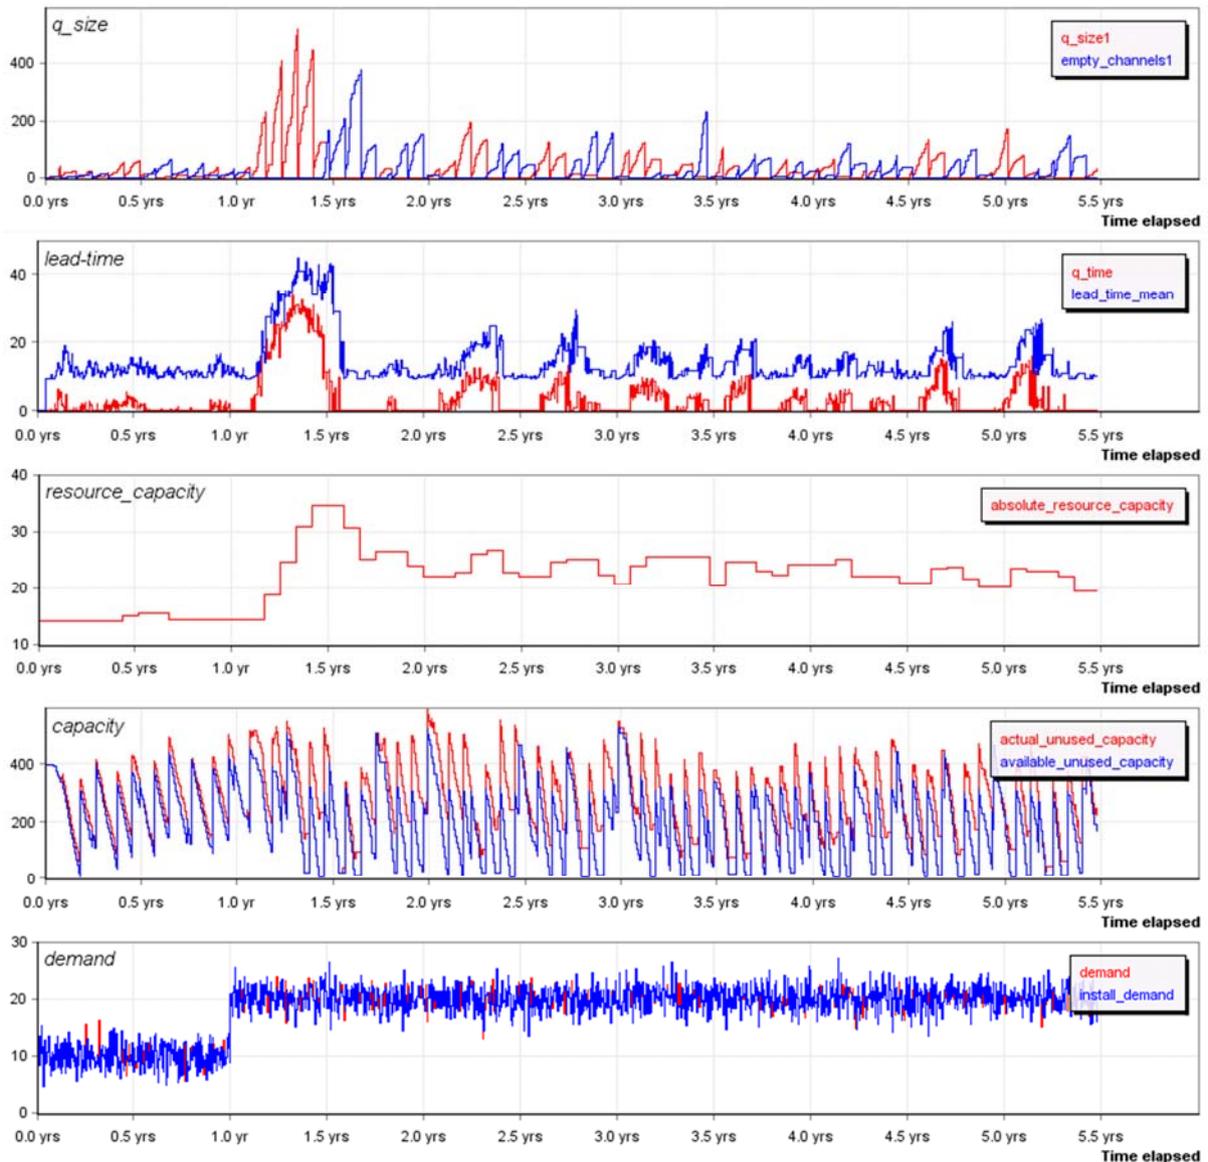


Figure 4. Simulated process operation under noise and a step in demand

Analyses were run to find an efficient operating policy that is able to cope with these unexpected steps in demand, ie. that is robust to this kind of variation. The first step was to identify whether such a setup could be created through ‘policy tweaks’ within the existing process structure. Five policy levers were identified, such as running the capacity recovery process more frequently. These represent the control parameters of the process system.

A full-factorial simulation experiment was designed to evaluate different combinations of the 5 control parameters under both of two operating conditions: (1) noise and (2) noise plus a step. All five control parameters were varied in four steps. Per configuration, 20 simulation runs were executed, allowing the effects of the noise to be very roughly indicated. For each of the two operating conditions, the results were analysed to identify the combination of the 5 levers that minimises customer response lead time without requiring very high resourcing levels. This analysis was performed graphically in the CAM software, using an interactive selector on the parallel co-ordinate plots shown in Figure 5.

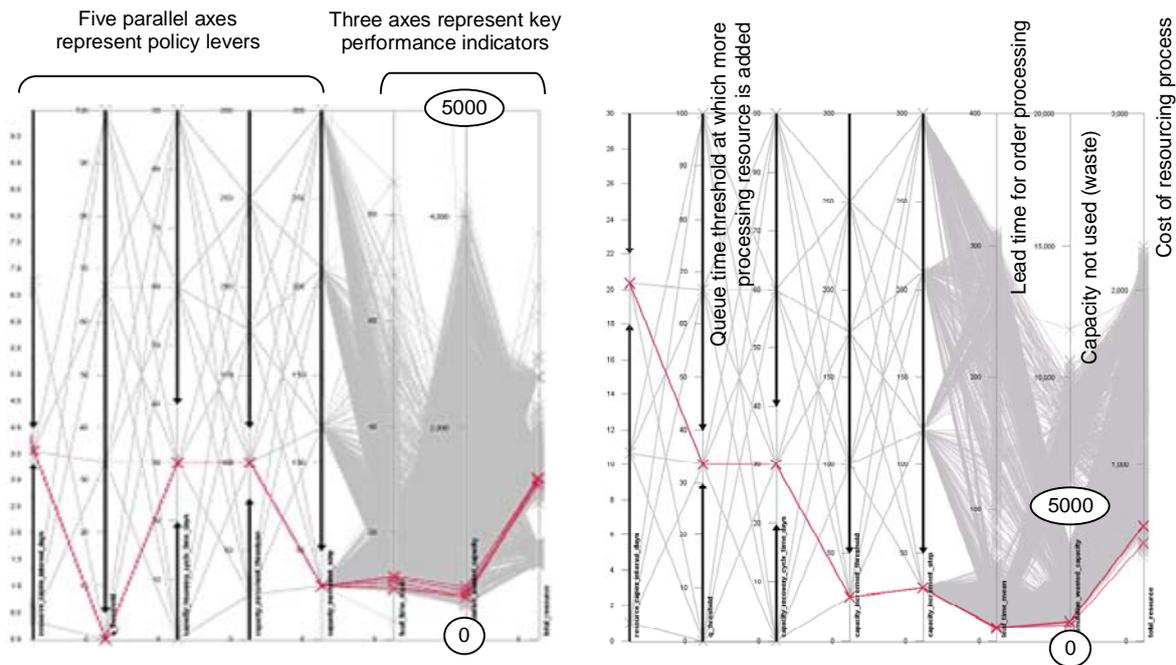


Figure 5. “Optimal” policies for operating under noise (left) and noise with step (right)

Comparison of the two plots shows that a different policy is required to give good performance in the case of unplanned disturbances (right plot) than for the case of noisy, yet predictable demand (left plot). Comparing the policies depicted in the plots, the 1st and 2nd policy levers are set higher, and the fourth is set lower in the right-hand plot. In terms of KPIs resulting from these two ideal policies, it was found that total resource requirement and wasted resource must be significantly higher to cope effectively with the step in demand (this is shown by the difference in scale on the labelled axis on the two plots).

These results can be explained as follows. A process system that is tuned to the limits of its performance has fewer built-in buffers, so that when unexpected events occur there is less slack to cope with them. In terms of the system dynamics, a process carefully tuned to respond quickly to small variations is less stable and tends to overshoot when large disturbances occur. In other words, the process as modelled can either be robust to disturbance, or efficient under normal demand – but not both. Of course, the organisation would be interested in a process system that gives the best of both worlds – it would like to be both efficient and robust. The results summarised above show that this is not possible simply by “tuning” the existing process design.

An alternative approach is suggested by the principle of requisite variety that was mentioned in Section 3. This principle implies that a more complex environment can be dealt with by a more complex system. In this case, the principle suggests that additional information feeds and/or more complex regulation principles could help the organisation account efficiently for both variability and unexpected events in its operating environment. After some consideration it was determined that this could be achieved by a feed-forward flow that predicted a disturbance. Such a prediction could be sent to the resourcing process by the sales business unit, because the latter is closely connected to the customer and thus has some capability to foresee changes in the pattern of demand.

To evaluate this proposed process improvement, the model was altered to include an extra information flow from sales into capacity planning, allowing an early warning of increases in demand to ‘short-cut’ the delays introduced by the long-lead time resourcing process and thereby ensure that the L2C process has sufficient capacity to deal with the disturbance when it occurs.

Assuming that the change in demand can be perfectly predicted and the exact amount of resource put in place, simulation revealed that the process is able to smooth out the ‘bump’ with no significant effect upon the main performance parameters. However, if the prediction is not perfect, the extra resource might be put in place too early, or too late. A simulation experiment was therefore created to assess the impact of poor prediction fidelity on the performance of the most efficient design as

determined in Stage 1. The results of this experiment are shown in Figure 6. This confirms that perfect prediction of changes in demand gives better process performance than poor prediction.

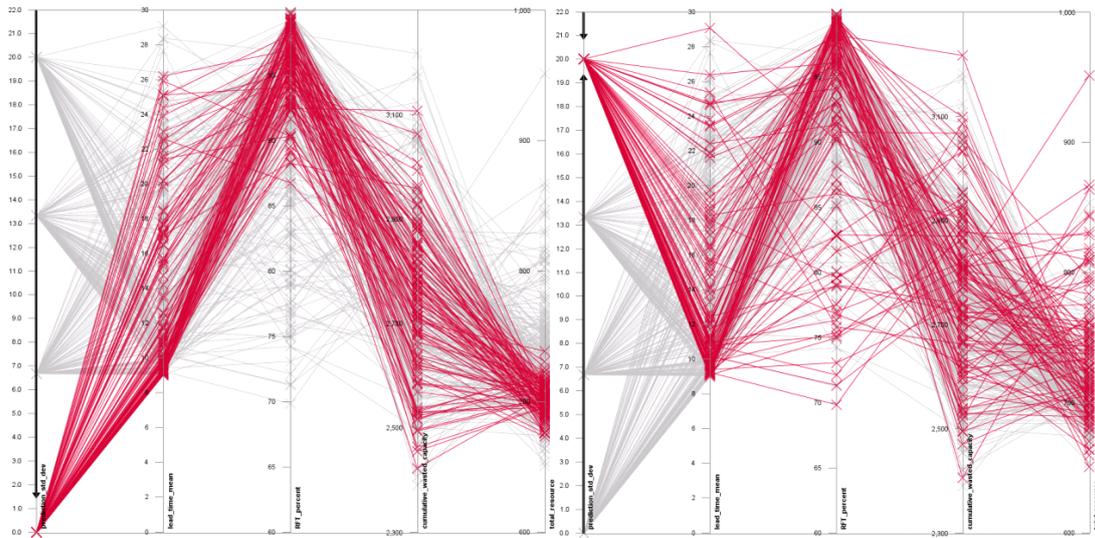


Figure 6. Perfect prediction of changes in demand (left) gives more robust process with respect to resource requirement and lead time metrics than poor prediction (right)

Further analyses were conducted along the same lines, for instance considering other types of disturbance, other control policies and investigating how accurate the predictions need to be for the additional information flow to be worthwhile. These could not be described here due to space constraints, yet are similar in principle and setup to those described above.

The approach which was taken, and the results that are partially described here, were fed back to participants in the data service processes through a series of meetings and presentations. Reactions to the potential of this type of modelling were positive, especially since the company has collected rich data against which the simulation model can be further validated. Efforts are ongoing to extend the analysis, especially with respect to the “effective organisation” discussed in section 4, to modelling the maturation of internal regulation capabilities throughout the lifecycle of a service, and to modelling in detail the interactions between service operation policies and satisfaction of the customer base.

7. Conclusions

Cybernetic systems theory can provide a useful framework for understanding the function of information flows in an organisation. In particular, information flows can provide value by helping to regulate process performance in the face of uncertainty and a changing environment, through the mechanisms of feed-forward and feedback. The analysis in this paper suggested several approaches to improve organisation performance through tuning regulation policies and putting in place suitable information flows. These suggestions were validated through analysis of a process system in BT.

Certain qualitative insights emerged for the design of robust service operations. In particular, the framework presented in Section 3 illustrates that an organisation can have different levels of ability to respond to uncertainty. It provides a means to structure analyses of a process system, by highlighting that different flows and controls in an organisation enable response to different levels of disturbance. To summarise:

1. Regulation by feedback does not require a detailed understanding of process dynamics to correct bottlenecks once they occur. However, if there is a significant time lag between observed cause and corrective action, i.e., if a regulatory process has long lead time, this tends to cause instability and oscillations.
2. Opportunities to redesign a process system to improve its performance can thus be found by seeking regulatory processes that have long lead-time, and finding ways to ‘short-cut’ those processes, for instance through the feed-forward of information from one business unit to

another. This was considered to be an important insight by the company personnel who were involved in the research.

3. The principle of shortcutting processes was shown to be useful only if prediction of demand changes can be made with sufficient fidelity.
4. Not all information is equally helpful. Flows that cut across levels of the organisation potentially have greater impact than flows within business units, especially when considering the higher-order forms of response to uncertainty. For instance, balancing resource across steps in a process to smooth demand variations is only effective if the process has been assigned enough resource overall.

The simulation model presented in this paper is appropriate to support design of well-structured 'quasi-steady flow' processes such as service operations. It would be less suitable to analyse ambiguous, ad-hoc or one-off processes that do not have clearly defined information flows, such as those that occur during development projects. Nevertheless, the same key issues are clearly visible in both situations: the existence of asynchronous intertwined processes; the need for communication between business units to manage interactions; and the importance of regulation, feedback and feed-forward to manage changes and uncertainty in workload. Thus, while the simulation itself is best suited to well-structured and repeating processes, the cybernetic principles and qualitative findings that were derived through simulation can provide useful insight for structuring and managing one-off processes such as product development.

References

- Akkermans, H., Vos, B., "Amplification in Service Supply Chains: An exploratory case study from the telecommunications industry", *Production and Operations Management*, Vol.12, No.2., 2000, pp 204-223.
- Beer, S., "Brain of the Firm", Allen Lane, 1972.
- Chalupnik, M. J., Wynn, D. C., Clarkson, P. J., "Approaches to mitigate the impact of uncertainty in development processes", *Proceedings of the International Conference On Engineering Design- ICED 2009, Stanford, USA*.
- Delgado, C., Larsen, E. van Ackere, A., Arango, S., "Capacity adjustment in a service facility with reactive customers and delays: Simulation and experimental analysis", *Proceedings of the 29th International Conference of the System Dynamics Society, 2011*
- Lawson, N. G., "In praise of slack: Time is of the essence", *Academy of Management Executive*, Vol. 15, No. 3, 2001.
- Le, Y.N., An, L., Connors, D., "Controlling workforce in response to demand disturbances in services supply chains", *Proceedings of the 28th International Conference of the System Dynamics Society, 2010*.
- Wynn, D. C., "Model-Based Approaches to Support Process Improvement in Complex Product Development", *PhD thesis, University of Cambridge, 2007*.
- Wynn, D. C., Maier, A. M., Clarkson, P. J., "How can PD process modelling be made more useful? An exploration of factors which influence modelling utility", *Proceedings of the International Design Conference - DESIGN 2010, Dubrovnik, Croatia*

Dr. David C. Wynn
Senior Research Associate
Engineering Design Centre
Cambridge University Engineering Department
Trumpington Street, Cambridge, UK.
Email: dcw24@cam.ac.uk
URL: <http://www-edc.eng.cam.ac.uk/people/dcw24.html>