In Proceedings of ASI '98, Bremen, Germany, June, 1998

Disturbance Measurement in Manufacturing Production Systems

G. Frizelle [1], D. McFarlane [1], L. Bongaerts [2]

Summary

This paper considers the problems of measuring and modelling production disturbances in discrete manufacturing operations, in conjunction with the design and operation of manufacturing control systems, capable of handling both internal and external production disturbances. The approach given here, is that being developed by the MASCADA Consortium [MASCADA, 1997] within a multi-agent systems environment. This paper is organised as follows. Section 1 sets the problem in the context of approaches to this type of scheduling, section 2 formalises the concepts, a mathematical model of propagation is developed in 3, while the final section considers the implications for providing robustness in discrete-event scheduling and control solutions.

1. Introduction.

Manufacturing systems, because of their complexity, are quite difficult to schedule and control. One reason is that manufacturing systems are prone to suffer disturbances, like machine breakdowns, operation delays and rush orders. Traditional approaches involving some form of hierarchical scheduling procedure, have proved to be too rigid to address disturbances adequately. New approaches, like reactive scheduling [Smith, 1995] have made significant contributions, but do not completely solve the problem.

A very promising approach to cope with changes and disturbances are agent based systems [Duffie, 1986]. Their presumed abilities to cope with disturbances stem from the following ideas:

- Multi-agent systems sometimes show emergent behaviour: an interesting and non-trivial overall system behaviour resulting from the interaction of local agents following simple local rules. The important point of these emergent properties is that they are not apparent at agent level.

- Agents can be linked to machines, thereby allowing for a quick response to a disturbance.

- Having agents work in parallel increases the calculation speed to react to disturbances.

However, looking at existing agent system implementations, one can see several factors limiting their success.

Experience with agent based systems indicate that they can sometimes over react to minor perturbations. This causes causing starvation and trashing - the tendency of a system to spend more time changing its mind than acting. In manufacturing, this results

[1] Manufacturing Engineering Group, University of Cambridge, Mill Lane, Cambridge, UK, CB2 1RX. Email: {dcm,gdmf}@eng.cam.ac.uk. Tel: +44 1223 338069. Fax: +44 1223 338076.
[2] Mechanical Engineering Department, Katholieke Universiteit Leuven, Celestijnenlaan 300B, B-3001 Leuven, Belgium. E-mail: Luc.Bongaerts@mech.kuleuven.ac.be. Tel: 32-16-32 24 80.

in an excessive numbers of set-ups and transportation activities at the expense of production. Another problem is that, by only reacting to disturbances, important opportunities are lost. There may already be some in-built robustness in the system that can be used to lessen the impact of disturbances. For instance, a delay on one machine may be confined to that work centre and have no overall impact on the order.

The paper attempts to put the issue of disturbances on a formal footing by considering how they propagate within a manufacturing system. From that a more precise way to look at and measure robustness can be presented. The paper does not address the issue of nervousness directly but some of the conclusions do have implications for how it might be handled.

Up to now, few authors appear to have addressed how the schedule is affected under changing conditions. Bongaerts et al [1998] did, but they only uses this approach for the evaluation of possible corrective actions and the selection of the optimal reaction strategy. He did not investigate the possibilities for (1) calculating a more robust schedule; (2) controlling the nervousness of the reactive system; and (3) reacting in a more robust way. Szelke [1995] defined the concept of *proactive scheduling*, as a way to timely prevent the impact of anticipated disturbances, as early as they are foreseeable from monitored performance trends. However, to our knowledge, there are no scheduling algorithms available yet for proactive scheduling.

The paper is structured as follows. First, the concept and range of production disturbance conditions is examined and the nature of disturbance propagation through a production operation is explained. The paper then presents a measure for the size of a disturbance and defines the concept of a recovery time to define a measure of robustness. There follows a presentation of a formal model for the disturbance propagation. The final section explains how these concepts can be used for the design, evaluation and optimisation of good reactive and proactive control strategies.

2. Production Disturbance Conditions

In this section we briefly describe the range of production disturbance conditions that can be considered in the context of the model developed. A production disturbance in this context is considered to be an unpredictable disruption (internal or external) to production. Typical examples are given below [Matson, McFarlane, 1998]:

- **Upstream disturbances**

  Materials quality problems, supplier production problems, materials delivery delays, material property variations, incorrect deliveries.

- **Internal disturbances**

  <u>Information, control and decision-making</u>

  Control and communication system failures, operator errors and omissions, recording / communication errors, materials ordering errors, materials stock control problems.

  <u>Production Equipment and Labour</u>

  Machine breakdowns, variability in machine performance (quality, cost, production rate), unavailability of labour

Material Handling and Flow

Blockages, handling equipment failure.

- **Downstream disturbances**

  Make to order:

  Rush orders, changes to orders (quantity, due date), quantity and mix variations (e.g. due to demand variations in customer's business), customer production problems.

  Make to stock:

  Demand variations (e.g. due to seasonality, marketing activity, competitor activity), forecasting errors, finished goods delivery delays, lost stock, poor stock monitoring,

In this paper we are concerned only with how a disturbance first affects the production process rather than its particular source, and it is noted that this initial effect can be a) local to a single machine or operation or b) distributed across a number of operations. In section 4 we examine a mechanism for measuring the propagation of this initial effect.

A critical issue in understanding the disturbance is in terms of its impact on production. Impact can be linked to the affect of the disturbance on production goals (as indicated by appropriate performance measures). Impact measured in this way provides for an objective comparison of production disturbances and can be used as a guide for developing differing recovery strategies. The model developed in the next section proposes a specific method for quantifying impact through disturbance propagation and the subsequent recovery time needed.

3. Formalising the Production Disturbance Description

In this section we provide some formalisation for describing the manner in which the disturbances introduced in the previous section affect a production operation. We note that the descriptions used here are not necessarily universally applicable. In particular the analysis has been based around the concept of a flow shop although broader interpretations are possible.

The manufacturing system. In this paper a manufacturing system, S, is characterised as a collection of resources and jobs execute by the resources and having a well defined boundary B. It is an input/output system with work coming in across B from 'suppliers' and work passing out across B to 'customers'. In practice, suppliers and customers may simply be further stages of production.

Within B resources are linked to one another by the passage of the jobs between them. The pattern of work is determined by some sort of schedule, considered to be a sequencing over time of a number of jobs N, where each job consists of a sequence of activities, onto M resources that execute them. There is no assumption about whether the schedule is efficient or not. This formalisation is felt to be general enough to encompass both traditional manufacturing systems and those based on agent technology.

Every job does not have to visit every resource. However each job follows a path through the M resources visiting some but not others. Such paths are often called Routings. An activity is the transformation process of a particular job by a particular resource. The sequence in which the activities are carried out on one resource may not be the same sequence as on a second resource. Thus if the activities associated with jobs 1 and 2 are carried out in that order on Resource (A), the model allows for the fact that they may be executed in the sequence 2 and 1 on Resource (B). The model excludes assembly and disassembly operations.

Each activity on a resource is characterised by a duration. It may include set up times as well as make times. There will also be some slack, which may be zero, between it and the next activity on the resource. If all these slacks are zero then the resource is a bottleneck. Transport time is ignored as its inclusion adds little to the development of the ideas.

However if the job visits a second resource, there will also be a second (distinct) slack between the completion of the job (the end of the activity) on the first resource and the start of the same job on the next resource. A feasible schedule is one where the slacks are all zero or positive.

The schedule starts at 'today' at time zero and finishes at some time horizon, usually the completion of the last job to be scheduled.

Disturbances. Deviations from steady operating conditions are characterised by the incidence of disturbances in our terminology originating from one of the different sources outlined in Section 2. These are considered here to be random events, in the Bernoulli sense, that occur from "time to time". In this work we only consider those disturbances that actively influence production and we assume that they cause a disruption of manufacturing for a duration Δ. In the model, it is assumed that the operation is halted. However it is easy to extend the formulation to include situations where the operation merely slows. Disturbances can strike at one resource, at several or at all. At this stage we can infer, and will shortly establish, that the effect of the former can be purely localised or can radiate from the point of occurrence. The first type will be called 'local' and the second 'propagated'.

We shall call the disturbance 'distributed' where the impact is on all resources in S, simultaneously .

However disturbances that are propagated might still be contained within S. Of particular interest are those that are not contained as they will impact customers and suppliers. Understanding the impact of disturbances is central to this paper [Matson et al 1998].

Recovery Time. A further issue critical to the understanding the disturbance is response effectiveness [Bongaerts 1993]. Recovery time is a function of how responsive the system is. Responsiveness can be achieved by one of two complementary properties: robustness and flexibility . Robustness is provided to the system by the presence of buffers. There are two types usually considered; 'stock' buffers and 'time' buffers. The former, as the name implies, are quantities of stock built into the system at various points to make up production short-falls. The second takes the form of slack available at resources. However both are time related.

Therefore robustness is a time based property. It is the time available for the system to recover, even though that time may appear as inventory. It is to be expected that different resources will have different levels of robustness. Therefore the robustness of the system will be the robustness of the critical resource i.e. the resource with the minimum ability to recover.

Flexibility can be thought of as a more subtle form of responsiveness. If a machine breaks down, then it may be possible to complete the job on a second machine. Alternatively, if an operator is ill, a second may have been trained to cover for his/her absence. However such a means of recovery is only possible where there is capacity available on resources - i.e. there is time available on the machine, the replacement worker can be spared, etc.

To proceed further, the concept of recovery time needs to be defined. At the resource level, it is merely the time for the operation to attain its original output rate following a disruption. In control theory, the change is usually modelled as some sort of impulse or step change. For a manufacturing plant, then the impulse would be the instantaneous breakdown of a machine or absence of a worker say, or the sudden interruption of supplies from a vendor. The recovery time would be the time it takes for that machine or worker (say) to achieve the budgeted speed of working: for a raw material supplier, how long before he is again supplying at the necessary rate.

This simple model does not hold at the systems' level. The reason is that the system S will be 'weakened' in some way as a result of a disturbance. It means that the system, if it has had insufficient time to recover, is more vulnerable to a second disturbance of the same size. In the model this weakness shows itself by stock excesses and shortages. The recovery time will be the time for the inventories to regain their planned levels. However if the initial disturbance has also impacted the customer, it will be the time for the customer's deliveries to be back on track.

4. The Disturbance Propagation Model.

There are in fact two models that need to be considered in establishing the response to a disturbance. The first, the Steady State Model describes steady state feasible schedules. The total schedule can be generated from the individual activities on the resources by the solution of difference equations. There are two sets of equations, looking forward in time (the Forward Equations) and looking backwards in time (the Backwards Equations). They provide equivalent answers. They are useful for modelling disturbances that propagate forwards and backwards, respectively. Note that the schedule is a 'given' for the model. No assumption is made about its effectiveness.

The second is the Perturbation model which examines what happens when a disturbance is applied to the schedule. Here there are four models. The first two look at forward propagation of a disturbance, the downstream effects and the upstream effects are modelled. The third and fourth look at backward propagation of disturbances, once again looking at the upstream and downstream effects.

4.1 Steady State Model.

We consider a manufacturing schedule s to be the mapping of a number of jobs , i = 1,.., N, where each job i consists of a number of activities, onto j = 1,..,M resources. Hence:

$$s:N \rightarrow M$$

followed by a sequencing of the jobs on each resource, for a set of given delivery dates to the customers of S and release dates for the start of work.

The activities are executed on the resources. However not every job visits every resource. The passage of each job i through the M resources forms a simple directed path, whose vertices represent the activities and whose edges represent the links between the activities. Path i has M_i vertices. Such paths are often called Routings. The direction of each of the N paths is in terms of increasing j. Note that the requirement for the path to be simple excludes assembly and disassembly from the current model.

For each resource, j, there will be N_j = N activities assigned to it. Each job can have, at most, one activity per resource. We further assume that all N jobs have activities on Resource 1. The sequence on Resource 1 is in terms of increasing i. However it cannot be assumed that the sequence will be maintained on resource j. Therefore for Resource, we shall refer to the activity in position k as activity k, which will be the job i; where i may not be equal to k.

Associated with each activity k on resource j, there will be a duration d_{kj}, where d_{kj} = 0. There will also be slack s_{kj} between activity k and activity k+1. Further there will be slack $\sigma_{j,j+l}$ between job i on resource j, and job i on resource j+l; (j+1,.., j+l-1 = 0) i.e. the next resource that job i visits.

A schedule will be called feasible if s_{kj} , $\sigma_{j,j+l}$ = 0 for all s_{kj} , $\sigma_{j,j+l.}$ Otherwise it will be called infeasible. We shall only deal with feasible schedules.

Forward Equations.
We let t_{kj} = 0, be the start time of the activity in position k on resource j, and t_{ij} be the start time of job i on resource j. We can write the two following recursive equations for a feasible schedule, which we shall call Forward Equations:

The capacity equation

$$t_{kj} = d_{k-1j} + s_{k-1j} + t_{k-1j} \quad (1)$$

Observation: if s_{kj-1} = 0 for all k on j, then j is a bottleneck.

The sequence equation

$$t_{ij} = d_{ij-1} + \sigma_{ij-1} + t_{ij-1} \quad (2)$$

Observation: j only appears as a variable in equation (2), therefore in this model the resources are only linked by the jobs. A resource would be independent of all other resources if it had no jobs in common.

There is also the boundary condition for B, that t_{0j}, d_{0j} , s_{0j} σ_{0j} = 0 and the initial condition that $t_{11} = 0$.

Solving (1) recursively gives:

$$t_{N_j j} = \sum_{k=1}^{N_j-1} \left(d_{kj} + s_{kj} \right) + t_{1j} \qquad (3)$$

Solving (2) recursively gives:

$$t_{iM_i} = \sum_{j=1}^{M_i-1} \left(d_{ij} + \sigma_{ij} \right) + t_{i1} \qquad (4)$$

Looking first at equation (4). By virtue of the boundary condition, job i must start on resource 1 where it occupies position i, by definition. Therefore from (3) we may write t_{i1} as

$$t_{i1} = \sum_{k=1}^{i-1} \left(d_{k1} + s_{k1} \right) + t_{11}$$

so that

$$t_{iM_i} = t_{N_j M_i} = \sum_{j=1}^{M_i-1} \left(d_{ij} + \sigma_{ij} \right) + \sum_{k=1}^{i-1} \left(d_{k1} + s_{k1} \right) \qquad (5)$$

since $t_{11} = 0$ and we are at liberty to set $i = N_j$, the last job on M_i.

If we now turn to (3), assume that job i* occupies the first position on resource M_i*(=j). From (5) we have that:

$$t_{1M_i} = t_{1M_i*} = \sum_{j=1}^{M_i*-1} \left(d_{ij} + \sigma_{ij} \right) + \sum_{k=1}^{i*-1} \left(d_{k1} + s_{k1} \right)$$

so that the complete solution to (2) is

$$t_{N_j M_i*} = \sum_{k=1}^{N_j-1} \left(d_{kj} + s_{kj} \right) + \sum_{j=1}^{M_i*-1} \left(d_{ij} + \sigma_{ij} \right) + \sum_{k=1}^{i*-1} \left(d_{k1} + s_{k1} \right) \qquad (6)$$

by substituting (5) into (3).

Equation (5) just represents a path that traces back through the jobs progress on previous resources to the first resource. Once it reaches the first resource it tracks back along it to the start time. (6) represents a path that traces back to the first activity on the resource, then back to the first resource and then back to the start. (5) has two legs, (6) has three legs. Of course (5) and (6) are equal since we are at liberty to choose M* = M. This is true for all start times on all resources. Therefore all simple paths from t_{11} to any start time t_{kj} are of equal length. There are up to NM - 1 possible paths.

We can use this property of equal path length to define a steady state schedule.

Backward Equations.

Equations (1), (2), (8) and (9) can similarly be written starting at the other end of the problem. However we now need to define $t^{(f)}_{kj} = 0$, be the *finish* time of the activity in position k on resource j, and $t^{(f)}_{ij}$ be the *finish* time of job i on resource j. Thereafter the equations are solved as before.

The solution will therefore work forwards in time, starting from the beginning. The resulting equations are equivalent but written in terms of the start time t_{01}.

4.2 The Disturbance Model.

We now consider the situation where a disturbance occurs and model its impact on the steady state operations:

Forward Propagation; downstream effects. Suppose now that a disturbance of size Δ, $-\infty < \Delta < +\infty$ is applied at t_{11}.

Define $t_{11}' = \Delta + t_{11} = \Delta$. We shall only consider positive and negative values of Δ as the case of $\Delta = 0$ is trivial.

Let us further suppose that the disturbance has been attenuated to size Δ_{kj} at time t_{kj}.

We now say that the disturbance is propagated along resource j if $\Delta_{kj} > s_{kj}$. Otherwise it is absorbed. Clearly the larger the values s_{kj} , the smaller will be the propagated disturbance Δ_{kj} and the sooner it will be absorbed.

For propagation then:

$$\Delta_{kj} - s_{kj} > 0 \text{ or } \Delta_{kj} - s_{kj} = \Delta_{k+1j} \text{ the size of the disturbance propagated.} \qquad (7)$$

It can immediately be seen from (7) that no propagation can occur if $\Delta < 0$. Moreover, we would require $t_{11} > 0$, to ensure $t'_{11} = 0$. In other words a negative value of Δ only makes sense where the start of the schedule is some way in the future, and the additional slack generated only affects the first operation. That is not to say that we cannot readjust the other slacks to suit ourselves, subsequently.

From here on we shall only consider $\Delta > 0$ Rewriting (1)

$$t_{kj} + \Delta_{k-1j} = d_{k-1j} + s_{k-1j} + t_{k-1j} + \Delta_{k-1j}$$

substituting (7) into the above and letting $t_{k-1j} + \Delta_{k-1j} = t'_{k-1j}$, then

$$t_{kj} + \Delta_{kj} = d_{k-1j} + t'_{k-1j} \text{ so we equivalently can call } t_{kj} + \Delta_{kj} = t'_{kj}$$

So that (1) becomes

$$t'_{kj} = d_{k-1j} + t'_{k-1j} \qquad (8)$$

It means that the disturbance is propagated along the resource by pushing each activity against its successor and absorbing the slack in the process but also attenuating the disturbance.

Observation: by virtue of the definition of absorption, the equation at the point of absorption must be of the form $t_{kj} + \Delta_{kj} = d_{k-1j} + s'_{k-1j} + t'_{k-1j}$, since by definition there has to be some positive slack left. Moreover since there is no propagation $\Delta_{kj} = 0$ because Δ_{kj} can only have positive or zero value. Therefore, as would be expected the start time of activity k is unaltered. Of more importance, all the subsequent start times will be unaltered. Therefore there is a unique absorption point. The disturbance cannot disappear and then reappear somewhere else.

A similar definition of propagation holds for the sequence equation. We can therefore argue analogously that (2) can be rewritten

$$t'_{ij} = d_{ij-l} + t'_{ij-l} \qquad (9)$$

where $t'_{ij} = \Delta_{ij} + t_{ij}$

Now $t'_{kj} = t_{kj} + \Delta_{kj}$ but $t'_{ij-l} = t_{ij} + \Delta_{ij}$, so there is no reason to suppose that $t'_{kj} = t'_{ij}$. In other words the two possible paths to this point are no longer necessarily of equal length.

We therefore have to define the start time for activity k on resource j as

$$t_{kj} = \max\left[(d_{k-1j} + t'_{k-1j}), (d_{ij-l} + t'_{ij-l})\right] \qquad (12)$$

By inspection the solution to this equation can be seen to be, remembering that $t_{11}' = \Delta + t_{11} = \Delta$.

$$\tau_{N_j M_i} = \max\left[\left(\sum_{k=1}^{N_j-1} d_{kj} + \sum_{j=1}^{M_i*-1} d_{ij} + \sum_{k=1}^{i*-1} d_{k1} + \Delta\right), \left(\sum_{j=1}^{M_i-1} d_{ij} + \sum_{k=1}^{i-1} d_{k1} + \Delta\right)\right] \qquad (13)$$

Observations. Because (13) holds for all start times other than the first, one path the disturbance takes will always be absorbed (unless the two paths are of equal length by chance). Thus we can talk of the major path the disturbance follows which 'swallows' minor paths. Moreover there are again up to MN-1 possible paths the disturbance can follow. It is hardly surprising that companies find it so hard to link cause and effect, when a disturbance hits.

Also note that when the disturbance is finally absorbed, there is a unique start time that is unchanged, so the path lengths again become equal, absorption is about equalling path lengths and hence stabilising the schedule.

Suppose that on resource 1, for all positions k the path along the resource is longer than the path length through the sequence. That means, because of the absorption, the disturbance cannot be transmitted to other resources, since the link is via the jobs. This type of disturbance was earlier defined as a Local Disturbance.

Forward equations ;upstream effects.

Assume now, for simplicity, that the disturbance strikes at resource m. What happens at resources m-1, m-2,...,1. It is clear from (8) and (9) that there is no direct backwards propagation, since the direction of propagation is in increasing k and j.

In other words resources m-1, etc. carry on working. However suppose that there is a buffer of size $C_{m,m-1}$ between m and m-1. Suppose further that at time t = 0, when the disturbance impacted m, there were i_m jobs in the buffer. Therefore the buffer will have an available capacity of $C^A_{m,m-1}$ where

$$C^A_{m,m-1} = C_{m,m-1} - i_m. \qquad (14)$$

Therefore the buffer will continue to receive work until the available capacity is full. This can be expressed as:

$$n_{k',m-1} - n_{k,m} = C^A_{m,m-1} = C_{m,m-1} - i_m \qquad (15)$$

where k' are the number of jobs processed on m-1 and going into the buffer and k are the number of jobs processed by m and withdrawn from the buffer.

Solving (15) recursively gives:

$$n_{k1} - n_{km} = \sum_{j=1}^{m-1} C^A_{m-j,m-j-1}$$

where n_{k1} is taken to mean the activity occupying the k_1 position on resource 1, n_{km} is the activity occupying the k_m position on m (not necessarily the same relative position as on 1) and $C^A_{0,1}$ is the buffer between the input boundary and the supplier.

Moreover we have the initial condition that $n_{km} = 0$, since the disturbance is assumed to strike at time t = 0.

Therefore we finish with the very simple equation

$$n_{k1} = \sum_{j=1}^{m-1} C^A_{m-j,m-j-1} \qquad (16)$$

Now (16) is non decreasing for increasing values of m, since the right hand side terms are all positive or zero. It means that the further back we move upstream, the longer it will be before the resource halts. In other words, the resources m-1, m-2,...,1 will shut down progressively. Notice also that the values n_{kj} will be maximised if $C^A_{m,m-1} = C_{m,m-1}$ in (14). Therefore the upstream buffers will run on for the longest periods if they start by being empty.

Backward equations.

A disturbance can also be applied at the end time of the schedule T, which pushes everything backwards. Using the backwards equations to solve for the resulting perturbation, results in an equation similar to (13), except expressed in terms of end dates. The same logic is followed. However there is now the possibility of getting an infeasible solution, where Δ is sufficiently large not to be absorbed, so that it impinges on t_{11}.

5. Towards Robust Control - Discussion and Conclusions

The model developed in Section 4 provides some useful guidelines for the development of manufacturing control methods that are robust or resilient to classes of production disturbances.

In particular, it gives the basis for defining robustness; this will be the distance travelled (both backwards and forwards) by a given disturbance of size Δ. Clearly that definition, though precise, is too general to be of much use. Better first to identify critical resources. We shall use equation (3) for that purpose and define the critical resource as the one (or more) where (17) holds:

$$\sum_{k=1}^{N_{j*}-1} s_{k,j*} = \min_{j \in M} \left\{ \sum_{k=1}^{N_j-1} s_{k,j} \right\} \qquad (17)$$

In other words it is the resource(s) having minimal capacity. If the left hand side is zero, then this is a bottleneck, within the time horizon defined by N_{j*}.

To consider the value of the approach, look first at the idea of the recovery of the entire system S. (13) and (16) give the distances the disturbance has travelled. We can specify that we wish to know that distance when a defined disturbance of size Δ is applied to the critical resource(s). Of particular interest is when that distance is just to the boundary B but does not cross it.

This then provides an alternative view of robustness; the size of disturbance that, when applied to the critical resource(s) propagates to the closest boundary but does not cross it. Usually manufacturing operations are most interested in not having the disturbance affect the customer; they tend to be more relaxed about suppliers.

This then has implications for the recovery time of the system S. Clearly if there is no spare capacity on the critical resource, then downstream the buffers can never recover. Upstream the recovery time i.e. when the all the buffers contain a level of inventory appropriate to normal working (business as usual), will be determined by the spare capacity on the resource with least slack.

Equally, if there is spare capacity on all resources, then the recovery time will be the time for the buffer(s) down stream of the resource with minimum spare capacity, to recover.

It is useful that the mathematical structure of (13) and (16) is linear. It means that the ideas developed here should be applicable to agent based systems. The linearity of the functions suggests that they should be capable of being distributed amongst the agents in a fairly simple way.

The problem of nervousness has not been directly addressed. Suffice it to say that, one can speculate that the more robust the system, the less likely is it to exhibit this rather unfortunate property. The reasoning is that the more robust the system, the longer time there is before any decision needs to be implemented - about how to react to the disturbance. Indeed one can envisage having an 'inner' boundary. One approach might be to delay implementing any decision until that inner boundary (decision boundary) has been broached. This reduces the number of decisions being made and hence the nervousness. However such thinking is speculative at this stage.

Nor does the model look at how recovery might be actively pursued i.e. what decisions to take. Rather it concentrates on the in-built robustness of the system to recover. What is striking is that the 'optimal' policy for buffering with maximum robustness is different downstream to upstream. Upstream the requirement is to leave buffers as empty as possible, downstream as full as possible. In addition there is still the issue of what corrective action to take and when, the reactive aspect of disturbance management. This is the subject of on-going research. A further paper will present details on the specification of robust control algorithms currently under development [MASCADA, 1997], [Bongaerts et al., 1998].

6. References

Bongaerts, L, 1993. Discrete Event Systems: performance evaluation for flexible manufacturing systems, Technical Report PMA 93R35, Mechanical Engineering Dept, K.U.Leuven, 1993.

Bongaerts, L, Van Brussel H., Valckenaers P. 1998, Schedule Execution using Perturbation Analysis, IEEE International Conference on Robotics and Automation, Leuven, Belgium, 16-21 May 1998.

Duffie, N. A. And R. S. Piper, 1986, *Non-Hierarchical Control Of Manufacturing Systems*, Journal Of Manufacturing Systems, Vol. 5, No. 2.

MASCADA 1997, *Manufacturing Control Systems Capable of Handling Change and Disturbance* - An Application of Autonomous, Co-operative Agents, ESPRIT Long Term Research Project No 22728, 1997-2000.

Matson, J, McFarlane, D, 1998. Tools for Assessing the Responsiveness of Manufacturing Systems, IEE Workshop on Responsiveness in Manufacturing, London, 1998

Smith S. F., 1995. "Reactive Scheduling Systems", Appeared In "Intelligent Scheduling Systems", Brown, Donald E. (Ed), Kluwer Academic Publ. Boston (Mass.), (Isbn/Issn : 0-7923-9515-8).

Szelke E., L. Monostori, 1995 (Sept. 11-15), Reactive And Proactive Scheduling With Learning In Reactive Operation management, Proceedings of IFIP International Working Conference on Managing Concurrent Manufacturing to Improve Industrial Performance, Seattle, Washington, USA, pp. 456-483.