



Australian Government
Department of Defence
Defence Science and
Technology Organisation

Modelling Logistics Chains for Optimal Supply and Maximum Throughput

Scott Wheeler, Alex Ryan and Tristan Turner

Land Operations Division
Defence Science and Technology Organisation

DSTO-TR-1807

ABSTRACT

This report presents an analysis of a generalised logistics supply chain without back orders. Two methods are proposed: a responsive but robust delivery system based on maintaining set holdings levels; and a technique from Control Theory which pushes stock through the chain in anticipation of demand over both time and space. Furthermore, a heuristic is proposed to set the policy for holdings levels using a hybrid of statistical analysis, Simulated Annealing and Lagrangian Relaxation. Finally, a comparison between methods under uncertainty, with error in prediction and correlated demand, is conducted. Each method was found to be useful in different contexts. The impact of uncertainty and correlated consumption was quantified for a set scenario and both were found to be significant factors in the performance of the supply chain.

RELEASE LIMITATION

Approved for public release

Published by

*Land Operations Division
DSTO Defence Science and Technology Organisation
PO Box 1500
Edinburgh South Australia 5111 Australia*

*Telephone: (08) 8259 5555
Fax: (08) 8259 6567*

*© Commonwealth of Australia 2005
AR 013-547*

*Submitted: June 2005
Published: November 2005*

APPROVED FOR PUBLIC RELEASE

A Study of Logistics Chains for Optimal Supply and Maximum Throughput

Executive Summary

Conceptual incentives in logistics management and control, such as the United State's Office of Force Transformation Sense and Respond Logistics and the Australian Defence Force's Future Joint Logistics Concept, promote investment in rapid distribution supply networks based on adaptive processes across the Joint domain. It is proposed that streamlined "faster" and "leaner" methods of resupply, which trade off capacity and volume against rate of resupply, have the potential to reduce logistics footprints, improve distribution times and lower the risk of shortfalls in supply.

This work primarily supports the defence capability Land 121 - Field Vehicle and Trailer Fleet Modernisation project with additional input to the JP 126 - Joint Theatre Distribution System, JP 2059 - Bulk Liquid Distribution and JP 2077 - Improved Logistics Information System projects. It examines the practicality of the new just-in-time precision logistics concepts by contrasting a purely responsive demand-orientated model against a deliberately planned supply-orientated model.

It is determined that no single method of logistical supply is a panacea for designing, managing and controlling logistics systems. The relevance and appropriateness of each are tied to the predictability of the logistics system and its behaviour. Identifying when each approach is useful is then more important than promoting any one approach above others. Furthermore, techniques for responsive systems, planned systems and adaptive systems are not necessarily mutually exclusive and can be applied in sequence or in concurrence for some systems. The benefits of one particular concept or philosophy over another becomes a moot point as each is useful in context.

Correlations between consumption at different nodes in a supply chain were found to significantly reduce the performance of the system. This highlights the need to incorporate a level of buffering in warehousing policies to counter this effect. Uncertainty in predicting future consumption, and different probability density functions for consumption schedules were also found to significantly affect the performance of the system for the scenario under investigation. Together, these factors point to the need to maintain local capacity to achieve robust performance in the face of unexpected events.

The results of this study do not directly contradict the ideas presented in the Australian Defence Force's Future Joint Logistics Concept, but instead serve as a warning that responsive systems, simple or complex, will not and cannot entirely replace traditional planned systems, and that the problems associated with the design, implementation and analysis of adaptive precision systems are non-trivial. This study does support the conclusion that enhanced capabilities, emerging technologies and network-enabled warfare will ultimately facilitate better management and control of logistics systems. However, it does so on the basis that these new developments will reduce uncertainty across the system and therein have an indirect effect on the system rather than directly supporting better logistics supply.

Authors

Dr Scott Wheeler

Land Operations Division

Scott Wheeler completed a Bachelor of Science (Mathematics and Computer Science) at the University of Adelaide with a major in Computer Science and Applied Mathematics in 1996. He received an academic scholarship during his honours year and graduated top of his class in 1997. In 1998 he enrolled in a Ph.D. under an Australian Postgraduate Award and completed his studies in 2002. Scott joined DSTO Edinburgh as a Research Scientist in 2003 and currently works in Land Operations Division as the task manager of the Complex Adaptive Systems Long Range Research task.

Scott's research interests include random search algorithms, linear and network optimisation, mathematical programming and operations research, economic modelling and games theory, and stochastic processes.

Mr Alex Ryan

Land Operations Division

Alex Ryan is currently carrying out a PhD in the Department of Applied Mathematics at the University of Adelaide, researching the design of distributed complex systems. He received first class honours in Applied Mathematics from the University of Adelaide in 2000 and a BSc in Mathematics and Computing from Curtin University in 1998. Since joining the Land Operations Division of DSTO in 1999 Alex has managed the Complex Adaptive Systems long range research task, designed decision support software concept demonstrators, established new techniques of systems analysis and developed new models of combat.

Mr Tristan Turner
Land Operations Division

Tristan Turner completed a Bachelor of Computer and Information Science at the University of South Australia. He received a scholarship researching Logistical Simulations at the Australian Defence Force Academy and graduated with a Masters in Operation Research and Statistics. Tristan joined the Defence Science and Technology Organisation in 2005 and currently works in the Land Operations Division.

Tristan's research interests include military theory, complex adaptive systems, operations research, systems theory, optimisation techniques including heuristics, decision science and artificial intelligence.

Contents

1	Introduction	1
2	Logistics Modelling	2
2.1	Logistics Supply Network	3
2.2	Logistics Supply Chain	4
2.3	Heuristics for Time-Invariant Holding Policies	5
2.4	Cascades in Logistics Supply Chains	11
3	Control Theory	15
3.1	Multistage Decision Process	15
3.2	Dynamic Programming	16
3.3	Logistics Supply Chain	18
4	Comparison of Systems under Uncertainty and Correlation	21
4.1	Degradation of Optimality under Uncertainty	21
4.2	Degradation of Optimality under Correlation	23
5	Discussion	25
6	Conclusions	27
7	Acknowledgements	28

Appendices

A	Geometry of the State Space	30
B	Simulated Annealing	37
C	Error Component	38

1 Introduction

The nature of Land warfare is becoming increasingly complex, diverse, diffuse and lethal (Commonwealth of Australia, 2004). However, complex terrain and the changing contemporary conflict environment offers new opportunities as well as new threats for the Australian Defence Force to overcome. Facilitated by enhanced capabilities, emerging technologies and network-enabled warfare, the *Future Joint Logistics Concept* (Commonwealth of Australia, 2002) provides a conceptual foundation for the evolution of logistics management and control to meet the challenge of future warfighting.

The Australian Defence Force's *Future Joint Logistics Concept* promotes an agile logistics network which is focused on rapid distribution of goods and services. This shift from static inventory levels and stockpiling policies in favour of faster and leaner methods of resupply will not necessarily eliminate or reduce the requirement for contingency stockpiling policies but has the potential to tradeoff capacity and volume for increased rate of resupply. It is proposed that adaptive and anticipatory logistics networks streamlined for rapid distribution have the potential to reduce logistics footprints, improve distribution times and lower the risk of shortfalls in supply.

In both the civilian and military sectors, the justification for investing in Information Technology improvements for logistics supply chains has traditionally been in terms of greater efficiency. Potential improvements include "better information and prediction about consumption rates, faster transport, cheaper transport, more efficient inter-modal transfer, reduced stockpiling, quicker delivery or shuttle platforms, better information and prediction about transfer rates, better visibility into critical items and better understanding of item sequencing" (CAPT Lewandowski, L. and J.R. Cares, 2005). From this perspective, information technology plays an important role in process improvements by gathering more information on consumption levels to enable better prediction of future consumption. Clearly, the efficiency gains are dependent on the ability to predict future consumption with sufficient accuracy and lead time to evaluate and implement the optimised supply schedule.

Sense and Respond (SR) logistics is an aspirational goal for future military logistics systems promoted by the US Office of Force Transformation (United States Department of Defense, 2004) and is also reflected in the Australian Future Joint Logistics Concept's desire for Network Enabled Logistics (Commonwealth of Australia, 2002). SR logistics "is envisioned as an approach that yields adaptive, responsive demand and support for force capability sustainment that operates in situation-conditioned structures that recognize operational context, coherence, and coordination" (United States Department of Defense, 2005). Thus, the role of information technology is shifted from purely facilitating global optimisation, to enabling robust responses to local conditions throughout the logistics network.

This study examines the practicality of the new just-in-time precision logistics concepts by developing and subsequently analysing two logistics systems, each modelled on the opposing but not necessarily contradictory viewpoints introduced above.

In Section 2, a responsive, demand-orientated model is formulated. This model reacts to its environment with fixed responsive dynamics without prediction of future demand or adaptation and learning. This simple model captures the local responsiveness and

robustness of the SR concept, although the role of prediction and adaptation are not modelled. Using this model, we investigate the nature of stockholding policies in linear supply chains under our interpretation of the SR policy for resupply. The ability to use warehouses to adequately manage the risk of stockouts and shortfalls in supply throughout the system is investigated - that is, we attempt to minimise the probability of ever running out of stock and failing to meet demand. Under SR dynamics, the resupply policy is fixed. We then develop a method to determine stockholding policies for warehouses throughout the chain subject to known constraints using a hybrid technique combining aspects of statistics, Simulated Annealing (Kirkpatrick et al., 1983; Cerny, 1985) and Lagrangian Relaxation (Everett, 1963; Fisher, 1985). Hence, we investigate the premise that a logistics network can be streamlined, maintaining low holdings and capacity, while still maintaining adequate distribution and throughput.

In Section 3, a deliberately planned supply-orientated model is formulated. This model is fundamentally different to the SR model of Section 2. It reacts to its environment using predictive dynamics under the assumption of perfect foresight with compensation for error and uncertainty. In this model, we fix the holdings policies throughout the system and develop a method to determine an optimal routing or schedule for goods using a technique from Control Theory called Dynamic Programming (DP) (Bellman, 1952). The DP model instantiates the optimisation concept for our linear supply chain. The problem is highly abstracted, since the motivation is to understand the dynamics of efficient supply chains rather than build an accurate model of contemporary military logistics systems. Nevertheless the DP model provides us with a sufficient representation of the defining characteristics of the optimised logistics supply chain concept to quantify some aspects of the performance of these systems.

Finally, we relax the assumption of perfect foresight and introduce stochastic noise into predictions of future consumption as well as correlating consumption to introduce occasional catastrophic failures in the chains. This allows us to explore some of the problems associated with the practical application of DP and to compare it against the fixed SR policy introduced in Section 4.1.

A discussion of the results and further research is provided in Section 5. Conclusions are drawn in Section 6.

2 Logistics Modelling

In this section we introduce the concept of a generalised logistics network. A logistics chain is then defined as a special instance of this network. SR dynamics are introduced to describe the manner in which goods are resupplied throughout the chain. With this model, we present a technique to determine stockholding policies for warehouses throughout the chain with the ultimate goal to minimise the risk of stockouts and shortfalls in consumption. This work primarily provides an analytical conceptual framework for the study of SR dynamics, including the idea that volume and capacity can be reduced while adequate distribution and throughput are still maintained. An illustrative example of how the framework might be applied in practice is also provided. However, no empirical data

is presented from military supply chains in the real world so all examples are indicative of potential applications rather than demonstratively and provably practical systems.

2.1 Logistics Supply Network

A logistics supply network is modelled as a set of network nodes \mathcal{N} and a set of directed arcs \mathcal{A} joining these nodes. Nodes in the network are uniquely identified by an ordinal labelling over the natural numbers \mathbb{N} . For simplicity, we assume that two nodes in the network are connected by at most one arc. Hence, directed arcs are uniquely identified by the ordered node tuple (i, j) , $i, j \in \mathcal{N}$. These nodes and arcs represent entities in the real world such as warehouses, consumers or manufacturers and roads, pipelines or telephone lines, for example. A logistics supply network is then formally defined as the directed graph $G = (\mathcal{N}, \mathcal{A})$.

Every node $i \in \mathcal{N}$ at the instance in time $t \in \mathcal{T}$ is attributed with the following characteristics:

- a current holdings level h_i^t , which denotes the amount of stock held by the node i at time t ;
- a desired holdings level H_i^t , which denotes the policy for the preferred level of stock to be held at node i at time t ;
- a warehouse capacity level W_i^t , which denotes the maximum possible amount of stock that may be held at node i at time t ;
- a current consumption level c_i^t , which denotes the expenditure or usage of stock at node i at time t ; and
- a desired consumption level C_i^t , which denotes the demand or preferred level of consumption for stock at node i at time t .

Every arc $(i, j) \in \mathcal{A}$ at the instance in time $t \in \mathcal{T}$ is attributed with the following characteristics:

- a current throughput level $\delta_{i,j}^t$, which denotes the amount of stock transferred from the node i at the time t to the node j at time $t + 1$; and
- a maximum throughput $\Delta_{i,j}^t$, which denotes the maximum possible amount of stock that may be transferred from the node i at time t to the node j at time $t + 1$.

The consumption at time t for all nodes i is given by

$$c_i^t = \begin{cases} h_i^{t-1}, & C_i^t > h_i^{t-1}, \\ C_i^t, & C_i^t \leq h_i^{t-1}. \end{cases} \quad (1)$$

Hence, nodes in the network have no choice but to meet the preferred level of consumption C_i^t at time t using the stock holdings carried over from the previous time period $t - 1$. That is, nodes may not willingly hold back stock when it is in demand.

The holdings level h_i^t for every node $i \in \mathcal{N}$ at the instance in time $t \in \mathcal{T}$ is then the stock carried over from the previous time period $t - 1$, less the consumption at t , less the

difference between outgoing stock transfers and incoming stock transfers, given that the maximum warehouse capacity W_i^t cannot be exceeded. Formally,

$$h_i^t = \min\{W_i^t, h_i^{t-1} - c_i^t - \sum_{j \neq i} \delta_{i,j}^t + \sum_{j \neq i} \delta_{j,i}^t\}. \quad (2)$$

Here, we expect the warehouse capacity W_i^t to be constant in time at each node in most cases. However, we allow the capacity to vary in time to allow, for example, policy restrictions preventing the entire warehouse being utilised and expansions to the warehouse capacity.

To prohibit nodes from transferring more stock than is held by those nodes during time $t - 1$,

$$\sum_{j \neq i} \delta_{i,j}^t \leq h_i^{t-1} - c_i^t. \quad (3)$$

Hence, nodes may only transfer stock carried over from the previous time step $t - 1$ in excess of the consumption level c_i^t .

In addition, we define the quantity Q as the base multiple or unit size for stock. Stock may only be transferred between nodes in multiples of Q so that

$$\delta_{i,j}^t = Q\gamma_{i,j}^t, \quad (4)$$

for $\gamma_{i,j}^t \in \mathbb{N}$.

We propose a penalty function which is based on two terms: a penalty term for not maintaining the desired holdings level; and a penalty term for not supplying enough stock for consumption. We define this, for the nodes $i \in \mathcal{N}/\{0\}$, as

$$\pi(i, t) = A|H_i^t - h_i^t|^\alpha + B|C_i^t - c_i^t|^\beta, \quad (5)$$

where A , B , α and β are constants. Here, many different forms of penalty functions could be used as the generality of the techniques we apply does not depend on the specific form of the function. However, outcomes of particular applications of those techniques will ultimately differ based on the forms chosen.

The node 0 is considered to be the source of all stock. This node does not consume stock, has no warehouse limit and always has enough stock to transfer.

2.2 Logistics Supply Chain

A logistics chain is a particular instance of a logistics network, as given in Section 2.1. We formally define the chain linear supply chain of length $L + 1$ as the directed graph $G = (\mathcal{N}, \mathcal{A})$ in which $\mathcal{N} = \{0, \dots, L\}$ and $\mathcal{A} \equiv \{(i - 1, i) | i \in \mathcal{N}/\{0\}\}$.

In general, there are many possible ways to schedule the flow of goods within the chain, each potentially differing in value as determined by equation (5). Although the number of distinct ways to schedule the flow of goods in our logistics chain is finite for each given chain, enumeration over the space of all candidate schedules for a set of optimal solutions is impractical for all but the simplest of chains. Furthermore, an optimal schedule for

some particular instance of a chain tells us little about the optimal schedule for another. It is more important to identify general techniques or algorithms to allow us to resupply the chain across all possible chains. The aim then is to identify the best possible algorithm across a range of chains or at least to identify one or more useful and practical algorithms which assist us to understand the properties of the chain.

One of the simplest techniques for restocking nodes in the logistics supply chain is to adopt a SR regime. Simply put, nodes in the chain consume goods following equation (1) and subsequently place an order o_i^t for stock with node $i - 1$ according to the recurrence relation

$$o_i^t = o_{i+1}^t + Q(\min\{H_i^t - (h_i^{t-1} - c_i^t), h_{i-1}^{t-1} - c_{i-1}^t, \Delta_{i-1,i}^t\} \bmod Q), \quad (6)$$

for $i \in \mathcal{N}/\{0, L\}$ and

$$o_L^t = Q(\min\{H_L^t - (h_L^{t-1} - c_L^t), h_{L-1}^{t-1} - c_{L-1}^t, \Delta_{L-1,L}^t\} \bmod Q), \quad (7)$$

otherwise.

Although equations (6) and (7) seem complicated, they are actually quite straightforward. For each $i \in \mathcal{N}/\{0\}$, these equations merely recursively compute the cumulative order across the sub-chain $j \geq i$ in multiples of Q . The node 0, having unlimited stock, places no orders. Then

$$\delta_{i-1,i}^t = o_i^t, \quad (8)$$

describes the dynamics of the supply scheduling. Here we assume instantaneous throughput.

2.3 Heuristics for Time-Invariant Holding Policies

Let $X_i \equiv X_i^t$, $i \in \mathcal{N}$, $t \in \mathcal{T}$, denote the random variables whose outcomes set the consumption values C_i^t . We make the assumption of time-invariance for simplicity and tractability.

For a fixed linear supply chain G , we wish to guarantee that the likelihood $\mathbb{P}(h_i^t \leq x_i)$ of ever having holdings less than or equal to the fixed amount x_i is below some stipulated acceptable risk probability ρ . This section presents a heuristic to set constant warehouse policy levels $H_i \equiv H_i^t$ in order to meet the criterion

$$\mathbb{P}(h_i^t \leq x_i) \leq \rho_i, \quad (9)$$

for constants x_i and ρ . Alternatively, we may wish to guarantee that the likelihood $\mathbb{P}(C_i^t - c_i^t > y_i)$ of a shortfall in desired consumption C_i^t of more than the fixed amount y_i is below some stipulated acceptable risk probability ρ_i . That is

$$\mathbb{P}(C_i^t - c_i^t > y_i) \leq \rho_i. \quad (10)$$

The assumption is made that we are able to sample instances or access data about the underlying distributions of X_i but do not know the true distributions of X_i . For example, it is a reasonably simple exercise to measure the average rainfall on a given day at some

fixed location. However, it is a non-trivial exercise to develop a rainfall forecasting model or even to fit sample data to a distribution with total certainty. The techniques we apply are then practical, within reasonable limits, and do not make unnecessary demands on knowing all aspects of the chain. In simulating the chain, we set the distributions for X_i and sample those distributions a number of times to form unbiased estimators, which approach X_i as n approaches infinity, in order to provide examples of our techniques. These examples are not to be confused with the general technique we propose.

The ability to ever satisfy inequality (9) for $x \equiv 0$, $\rho_i \equiv 0$ depends on the distribution of X_i . For example, if X_i is uniformly distributed over the integers in $[0, 10]$ then it is certainly possible to guarantee that node 1 never runs out of stock simply by setting the holdings policies across the chain to sufficiently large integers. However, if X_i is normally distributed then it is not possible to guarantee that the nodes will never run out of stock for any finite holdings policy H_i because the upper tail end of the normal distribution is not bounded. The probability of running out of stock quickly diminishes as H_i increases so that setting H_i to an arbitrarily large number is an easy way to satisfy our requirements for all practical purposes. However, in reality H_i is limited by physical constraints such as warehouse sizes, monetary considerations or other issues which prevent us from having arbitrarily large warehouse holdings policies. In this situation, it makes sense to minimise the total holding policies across the chain subject to the constraints (9) or (10).

$$\begin{aligned} & \text{minimise} && \sum H_i, \\ & \text{such that} && \mathbb{P}(h_i^t \leq x_i) \leq \rho_i \text{ or } \mathbb{P}(C_i^t - c_i^t > y_i) \leq \varrho_i. \end{aligned} \quad (11)$$

Alternatively, define a maximum H_{max} for the holdings policies across the chain and minimise the total system failure probabilities across the chain.

$$\begin{aligned} & \text{minimise} && \sum \mathbb{P}(h_i^t \leq x_i) \text{ or } \sum \mathbb{P}(C_i^t - c_i^t > y_i), \\ & \text{such that} && \sum H_i \leq H_{max}. \end{aligned} \quad (12)$$

Both systems (11) and (12) are difficult to solve as they stand. We use Lagrangian Relaxation (Everett, 1963; Fisher, 1985) to translate our original problem into a simpler form.

Considering now only cases for equation (9), the program (11) becomes

$$\begin{aligned} & \text{maximise} && \mathcal{L}(\phi) = \min\{\sum H_i + \sum(\phi_i(\mathbb{P}(h_i^t \leq x_i) - \rho_i))\}, \\ & \text{such that} && \phi_i \geq 0, \end{aligned} \quad (13)$$

and the program (12) becomes

$$\begin{aligned} & \text{maximise} && \mathcal{L}(\psi) = \min\{\sum \mathbb{P}(h_i^t \leq x_i) + \psi(\sum H_i - H_{max})\}, \\ & \text{such that} && \psi \geq 0, \end{aligned} \quad (14)$$

where \mathcal{L} is the objective function, $\phi = (\phi_i)$, $i \in \mathcal{N}$ is the vector of Lagrangian multipliers and ψ is a scalar Lagrangian multiplier. The cases for equation (10) follow *mutatis mutandis*.

To explain how Lagrangian Relaxation works in our case, consider the objective function $\mathcal{L}(\psi)$ in the program (14). For $\psi \rightarrow 0$, the value of the objective function is dominated

by the first term $\sum \mathbb{P}(h_i^t \leq x_i)$ so that the minimisation will set the holdings H_i as large as possible and the probability of holding less than x_i becomes arbitrarily small, assuming also that $\psi \sum H_i \rightarrow 0$ as $\psi \rightarrow 0$, which is reasonable when $\sum H_i$ is bounded. The value of the objective function approaches 0 from above. On the other hand, as $\psi \rightarrow \infty$ the value of the objective function is dominated by the second term $\psi(\sum H_i - H_{max})$ since the first term can add to at most L when $\mathbb{P}(h_i^t \leq x_i) = 1$ for every $i = 1 \dots L$. As the holding policies for all nodes is generally expected to be at least 1 parcel of size Q then H_{max} should be set by the user to at least QL and is usually significantly greater. However, we allow $H_i = 0$ even though this makes little sense in practice. The value of the objective function approaches $(L - \psi H_{max}) \rightarrow -\infty$ as $\psi \rightarrow \infty$. It is not difficult to see that, in the minimisation, there exists some finite $\psi > 0$ which causes $\sum H_i$ to equal values other than arbitrarily large numbers or 0. For example, if such a ψ were to cause $\sum H_i$ to equal H_{max} then the second term in the objective function is eliminated but the first term in the objective function is likely to exceed 0 depending on whether X_i is bounded to some finite amount, in which case there might be no gain in raising $\sum H_i$ beyond some fixed amount which suffices to ensure $\sum \mathbb{P}(h_i^t \leq x_i) = 0$. Otherwise, the value of the objective function is strictly positive. It is a ψ of this kind which solves our problem. This appeals to reason because observing the initial problem (12) we can intuitively see that solutions where $\sum H_i$ is close to H_{max} are likely to be optimal. The Lagrangian Relaxation (13) works in a more complicated but analogous way to (14). To solve our Lagrangian Relaxation problems, we propose to use Simulated Annealing (SA) (Kirkpatrick et al., 1983; Cerny, 1985) as explained in Appendix B.

The implementation of the SA algorithm for our problem starts by identifying the state space, which, given that we make no undue assumptions about the distributions of X_i , $i \in \mathcal{N}$, is not trivial. We have already indicated that, for values of ρ_i in the limit approaching zero, the H_i may become arbitrarily large. Hence, we require an approach which bounds H_i based on some measurement of statistical likelihood. The technique we employ samples the distributions of X_i , $T = 1000$ independent times. These samples are used to construct an approximation to the distribution of the random variable Y_i , whose outcome describes the average of the sum of the X_j over $j \geq i$. The reasoning here is that on average the node i needs to hold enough to meet its own consumption and also to fill any stock order made by the node $i + 1$. Node $i + 1$ orders stock based on its own consumption and the order made by the node $i + 2$. With recursion, we propose that on average the node i is required to hold enough to fill its own consumption requirements and also enough that the consumption requirements of all nodes $j > i$ are met. This embraces the concept of throughput over capacity. It is not necessarily true that node i needs to meet the total desired consumption of the sub-chain $j \geq i$. It could be that nodes stockpile against the likelihood of not being supplied. In our systems, we wish to reduce the stockpiling levels and concentrate on the benefits of increasing throughput with appropriate holding policies. Hence, our approach is reasonable within this context.

A histogram for the sample outcomes of Y_i , $i \in \mathcal{N}$ is obtained using Monte Carlo sampling. This histogram is used to set upper bound for H_i based on the tail probability ζ intended to capture the upper ζ percentile of Y_i . Figures 1, 2 and 3 display the average outcomes of Y_i over 100 independent simulations each lasting 1000 time steps for nodes 1, 2 and 3 respectively in a 4 node chain. The error bars denote 95% confidence intervals using a t -statistic. In these figures, the upper tail end of the histograms are shaded in red to show

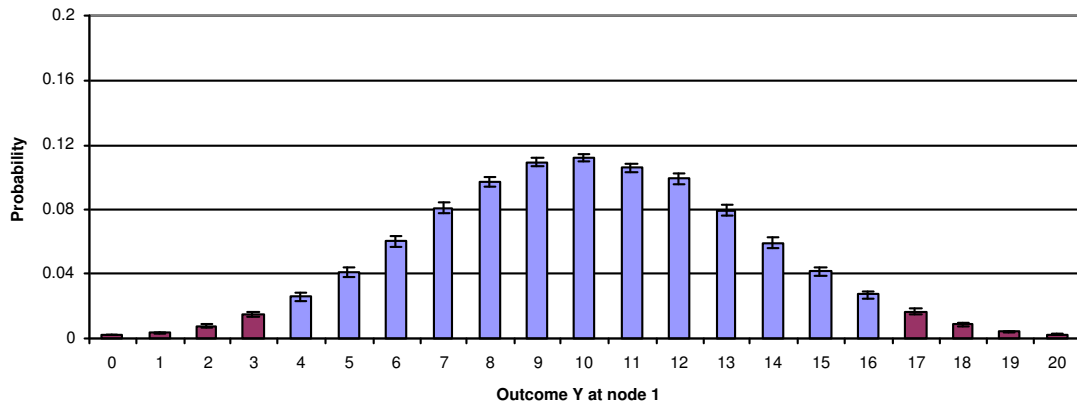


Figure 1: Histogram of sample outcomes of Y at node 1 over 100 independent simulations

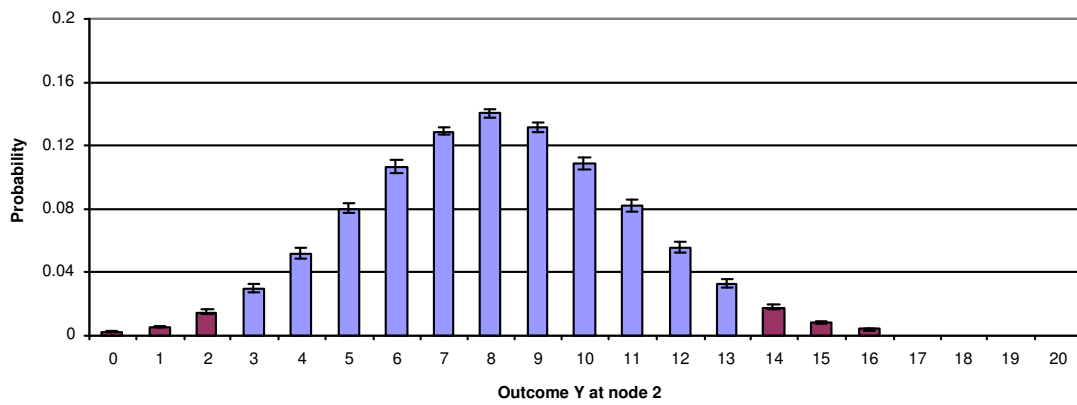


Figure 2: Histogram of sample outcomes of Y at node 2 over 100 independent simulations

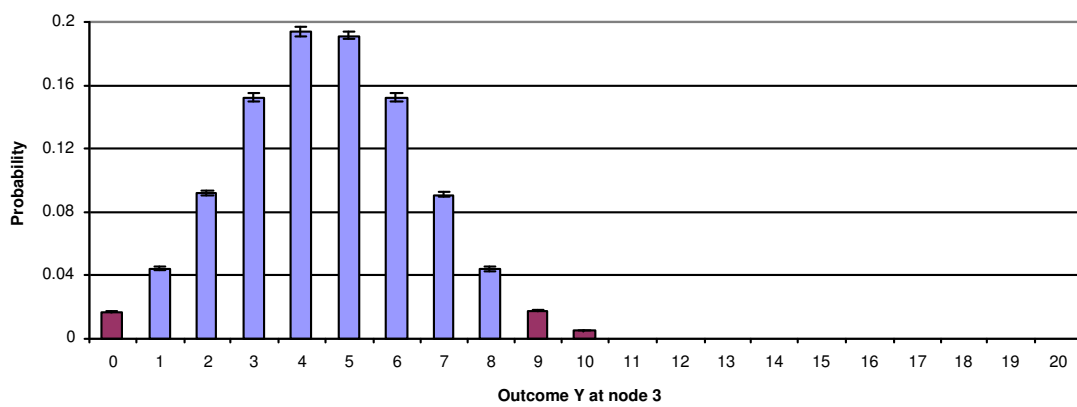


Figure 3: Histogram of sample outcomes of Y at node 3 over 100 independent simulations

the upper $\varsigma = 0.05\%$ of the histograms. The bound of zero is used as the lowest value for H_i in this model. However, the lower ς percentile of Y_i is also shaded red because setting a lower bound for H_i is useful in many practical situations.

The initial state of the system is set to

$$H_i^t \equiv H_i = \left[\frac{1}{T} \sum_{t=1}^T \sum_{j \geq i} C_i^t \right], \quad (15)$$

for $i \in \mathcal{N}/\{0\}$. The H_i values set in this step then denote the sum of the sample means for consumption of the node i itself and all nodes that node i supplies and is an unbiased approximation of the expected value $E[Y_i]$.

To implement the SA algorithm, we need to be able to perturbate the system - that is: for each temperature level $T \in \mathcal{T}$ define the disturbance functions $\xi_T : \mathcal{S} \rightarrow \mathcal{S}$ which takes as input the current state of the system and generates a new state. Let H_i^{high} and H_i^{low} denote the maximum and minimum values for node $i \in \mathcal{N}/\{0\}$ as determined by the upper and lower ς percentiles of Y_i . Again, we use $H_i^{low} \equiv 0$ for reasons that will become apparent shortly. Then, the disturbance function ξ is implemented to select a node at random, excluding the source node 0, and add or subtract a uniformly distributed random number. This number is generated between 1 and $\lambda(H_i^{high} - H_i^{low})$ with equal likelihood from its holding policy to generate a new state, where $\lambda = 0.1$ is a proportionality constant. The outcome of the algorithm does depend somewhat on λ . The value of 0.1 is chosen because it was observed to be effective - that is; the algorithm converged to a ‘‘better’’ solution more quickly than other values we tested. We use the cooling schedule

$$T_j = T_{j-1}/2, \quad (16)$$

for $j = 1, \dots, 7$ with initial temperature $T_0 = 1$.

To demonstrate our method, we compare two supply chains of length four. In the first instance, SA is used to provide an approximate solution to a Lagrangian Relaxation problem of the type (14) with $H_{max} = 25$, where the second term in the objective function from (12) is used - that is, $\sum \mathbb{P}(C_i^t - c_i^t > y_i)$, for $y_i = 0 \forall i$. In the second instance, Brute Force (BF) enumeration over all feasible solutions is used to solve the same Lagrangian Relaxation problem optimally. Desired consumption is normally distributed with a mean of 5 and a standard deviation of 3 so that the underlying simulations of the logistics chains are stochastic and optimality is itself subject to random stochastic noise. A single simulation of each chain lasts 1000 steps in time. The inner Metropolis loop of the SA algorithm perturbates the system $K = 100$ times at each temperature level. Each of the two approaches are repeated a total of 100 independent times. Results are provided in Table 1. In this table, the column labelled BF describes the best value of the objective function $\mathcal{L}(\phi)$ found by completely enumerating over the state space, the column labelled SA* describes the best value of the objective function found over all 100 independent trials and the \bar{SA} describes the mean value.

The best solution found for H_i across the chain occurs at $\psi = 0.06$ with value of $\mathcal{L}(\phi) = 0.0783$. For values of ψ between 0.55 and 0.65 in steps of 0.001 a better solution at $\psi = 0.059$ with value $\mathcal{L}(\phi) = 0.814$ is located. Random variation in the results impedes

Table 1: Results of simulations for BF enumeration and SA in a four node chain for values of ψ between 0.03 and 0.10

ψ	BF	(H_i)	$\sum H_i$	SA*	(H_i)	$\sum H_i$	SA	SDev
0.03	0.527	(17, 13, 10)	40	0.533	(17, 13, 9)	39	0.615	0.065
0.04	0.665	(17, 12, 8)	37	0.628	(16, 12, 9)	37	0.753	0.093
0.05	0.752	(15, 12, 7)	34	0.753	(15, 11, 7)	33	0.897	0.120
0.06	0.783	(10, 7, 0)	17	0.755	(11, 8, 0)	19	0.890	0.099
0.07	0.718	(10, 7, 0)	17	0.696	(11, 7, 0)	17	0.782	0.068
0.08	0.610	(7, 0, 0)	7	0.595	(7, 0, 0)	7	0.673	0.070
0.09	0.423	(6, 0, 0)	6	0.415	(6, 0, 0)	6	0.529	0.071
0.10	0.266	(5, 0, 0)	5	0.241	(7, 0, 0)	7	0.379	0.090

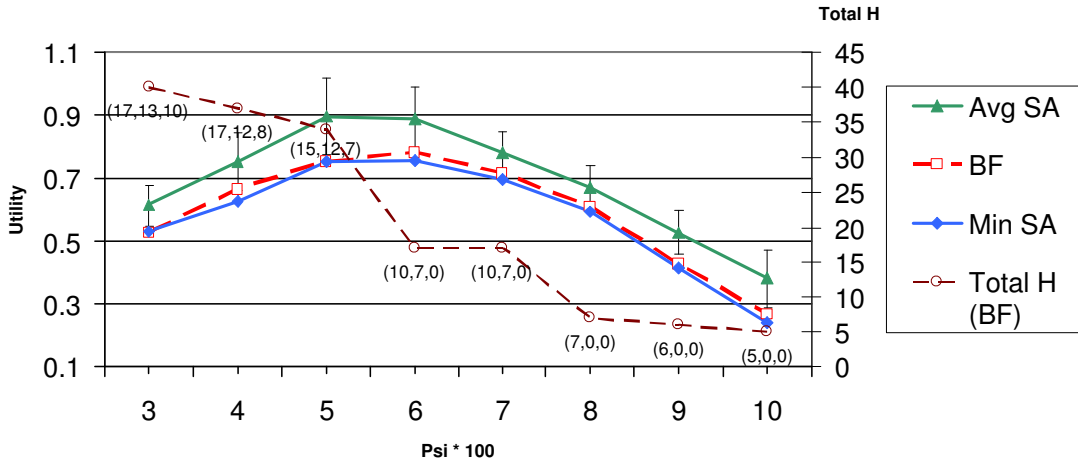


Figure 4: Graph of the total of the holdings policies in the chain and the utility $\mathcal{L}(\phi)$ for BF enumeration and SA plotted against ψ

our ability to accurately calculate the objective function beyond this sensitivity. A phase type behaviour is observed between $\psi = 0.057$ and $\psi = 0.058$ where the optimal holdings policies change from (14, 10, 4) to (10, 6, 0). Overall, the dynamics appear to contain two phase transitions in the BF $\sum H_i$ plot of Figure 4. A shift in policies between those in which $H_3 > 0$ and $H_3 = 0$ is important because if at any time $H_3 > 0$ then there must be sufficient capacity and throughput in all other nodes for the investment to be worthwhile. When $H_3 = 0$, an increased penalty is invoked in not meeting desired consumption at node 3 but the holdings policies of all other nodes are reduced to offset the loss. To demonstrate this principle, we ran a BF enumeration with the additional constraint $H_3 > 0$. The optimal policy for the system shifted from (10, 6, 0) to (9, 6, 3). In this situation one might have assumed that H_3 would be set to the lowest feasible value, namely unity. Instead, it was set to a value of 3 due to the higher-order non-linear dynamics of the chain. Obviously, by enforcing the constraint $H_3 > 0$ we are forcing the system to stay in one phase. This is important because it clearly demonstrates that discontinuities in the chain exist; that is, points at which two distinct solutions exist with equal value. Care must then be taken

in the neighborhood of transition points.

Figure 4 graphs the results of Table 1, where the rightmost ordinate axis (labelled ‘Utility’) measures the value of the objective function $\mathcal{L}(\phi)$ and the leftmost ordinate axis (labelled ‘Total H’) measures the $\sum H_i$. In this figure, the curve depicting the minimum value of $\mathcal{L}(\phi)$, obtained for each fixed ψ in the SA algorithm, actually lies below that obtained using brute force enumeration. This is clearly an aberration in our method caused by random variation across the consumption generated in each independent simulation. The reason is one or more of the 100 independent SA results are simply better than the result of the single simulation using brute force enumeration. The concept is somewhat similar to rolling a die once and then trying to beat that role by rolling 100 dice. In any case, it is clear that the best result of the SA algorithm is as close as one can reasonably get to optimal, given the random variation across the chain. The effect of this random variation can be reduced in practice by running each simulation for longer. In our example, 1000 time steps is considered sufficient.

2.4 Cascades in Logistics Supply Chains

We are interested in the stability of the logistic supply chain to perturbations. We seek to determine if small changes cascade through the system or whether the logistics supply chain is robust to perturbation. Under the SR policy, we investigate the stability of the network to a perturbation of size $\varepsilon = mQ$ in consumption at the end node L , where $m \in \mathbb{N}$. That is, ε is a non-negative multiple of the unit size for stock Q . If the desired consumption C_L^t results in current consumption c_L^t and in a resupply order o_L^t , then the desired consumption $\tilde{C}_L^t \equiv C_L^t + \varepsilon$ produces a current consumption

$$\tilde{c}_L^t = \begin{cases} h_L^{t-1}, & \tilde{C}_L^t > h_L^{t-1} \\ \tilde{C}_L^t, & \tilde{C}_L^t \leq h_L^{t-1} \end{cases} \quad (17)$$

and $\tilde{c}_i^t = c_i^t$, for $i \in \mathbb{N}/\{L\}$.

The resupply order to node $L - 1$ is

$$\begin{aligned} \tilde{o}_L^t &= Q(\min\{H_L^t - (h_L^{t-1} - \tilde{c}_L^t), h_{L-1}^{t-1} - \tilde{c}_{L-1}^t, \Delta_{L-1,L}^t\} \bmod Q), \\ &\leq Q(\min\{H_L^t - (h_L^{t-1} - (C_L^t + \varepsilon)), h_{L-1}^{t-1} - C_{L-1}^t, \Delta_{L-1,L}^t\} \bmod Q), \\ &\leq o_L^t + \varepsilon. \end{aligned} \quad (18)$$

For the recurrence relation,

$$\begin{aligned} \tilde{o}_i^t &= \tilde{o}_{i+1}^t + Q(\min\{H_i^t - (h_i^{t-1} - c_i^t), h_{i-1}^{t-1} - c_{i-1}^t, \Delta_{i-1,i}^t\} \bmod Q), \\ &\leq o_{i+1}^t + \varepsilon + Q(\min\{H_i^t - (h_i^{t-1} - c_i^t), h_{i-1}^{t-1} - c_{i-1}^t, \Delta_{i-1,i}^t\} \bmod Q), \\ &\leq o_i^t + \varepsilon, \end{aligned} \quad (19)$$

for $i \in \mathcal{N}/\{0, L\}$. Therefore, orders do not grow along the logistics supply chain when the system is perturbed. We now consider the state of the system at time $t + 1$ after a perturbation at node L at time t . The holdings at node L is

$$\tilde{h}_L^{t+1} = \min\{W_L^{t+1}, \tilde{h}_L^t - c_L^{t+1} + \delta_{L-1,L}^{t+1}\}, \quad (20)$$

where

$$\tilde{h}_L^t = \min\{W_L^t, h_L^{t-1} - \tilde{c}_L^t + \tilde{\delta}_{L-1,L}^t\}. \quad (21)$$

Note that

$$\tilde{\delta}_{L-1,L}^t - \delta_{L-1,L}^t \leq \tilde{c}_L^t - c_L^t \Rightarrow \tilde{h}_L^t \leq h_L^t. \quad (22)$$

Suppose W_L^{t+1} is the minimal term in equation (20). Then

$$\begin{aligned} W_L^{t+1} &\leq \tilde{h}_L^t - c_L^{t+1} + \delta_{L-1,L}^{t+1} \\ &\leq h_L^t - c_L^{t+1} + \delta_{L-1,L}^{t+1}. \end{aligned} \quad (23)$$

Therefore,

$$\tilde{h}_L^{t+1} = h_L^{t+1} = W_L^{t+1}. \quad (24)$$

That is, the current holdings level equals the warehouse capacity regardless of whether the system is perturbed, meaning the perturbation dies out at node L after one time step.

Suppose instead that $W_L^{t+1} > \tilde{h}_L^t - c_L^{t+1} + \delta_{L-1,L}^{t+1}$. Then

$$\begin{aligned} \tilde{o}_L^{t+1} &= Q(\min\{H_L^{t+1} - (\tilde{h}_L^t - c_L^{t+1}), \tilde{h}_{L-1}^t - c_{L-1}^{t+1}, \Delta_{L-1,L}^{t+1}\} \bmod Q), \\ &\leq Q(\min\{H_L^{t+1} - (h_L^t - \varepsilon - c_L^{t+1}), h_{L-1}^t - c_{L-1}^{t+1}, \Delta_{L-1,L}^{t+1}\} \bmod Q), \\ &\leq o_L^{t+1} + \varepsilon. \end{aligned} \quad (25)$$

This shows that the SR ordering policy is stable, since any perturbation at node L cannot grow either along the chain, or over time. The above proof can be easily extended for nodes other than L . Interestingly, linear supply chains are not always stable. Bender (2004) has shown that if orders are placed to satisfy a stockpiling policy $H_i^t \equiv a_i + C_i^t$, where a_i is constant, the logistics supply chain is unstable and any perturbation is amplified exponentially. Hence, it is possible that the logistics supply chain will exhibit chaotic dynamics under different ordering policies other than SR.

We also investigate the interdependencies between failures at different nodes in the supply chain. For the brute force enumeration in Appendix A, the expected probability of failure $E_i[\mathbb{P}(C_i^t - c_i^t \leq -1)]$ averaged over all 1296 holding policy combinations for each node i is 0.0000, 0.0006, 0.0294, and 0.2669 respectively. On average, most of the failures (90%) occur at the fourth node in the supply chain. However, while this symptom is exhibited at node four, it is not clear that the holding policy at node four is responsible for all of the failures. One way to investigate interdependencies between nodes is to remove the first node from the supply chain and simulate the remaining nodes with the same consumption schedules, with node 2 now connected directly to the source node 0. Because this removes the effect of the first node's consumption and holding policy, the reduction in $E_j[\mathbb{P}(C_j^t - c_j^t \leq -1)]$, $j \in \{2, 3, 4\}$ can be considered to be the component of failures at node j that are caused by the holding policy and consumption at node 1. This process can be applied recursively to identify the impact of each node on downstream failures.

Figure 5 depicts $E_i[\mathbb{P}(C_i^t - c_i^t \leq -1)]$ for each node. Each column in the graph is segmented into regions that represent the causal influence of upstream nodes. At node 4, even when all other nodes are removed from the logistics supply chain, $C_4^t - c_4^t \leq -1$ occurs 0.166 of

the time. In comparison, the activities of node 3 add 0.055 to $E_4[\mathbb{P}(C_4^t - c_4^t \leq -1)]$, while the contributions of nodes 2 and 1 are 0.030 and 0.016 respectively. For node 3, in the absence of nodes 1 and 2, $C_3^t - c_3^t \leq -1$ occurs less than 0.006 of the time, with consumption and holdings in nodes 2 and 1 contributing to 0.012 and 0.010 of failures respectively. In summary, the holding policy at node 4 is responsible for 62% of the failures at node 4, whereas 77% of node 3's failures are due to interdependencies with upstream nodes.

This process is also performed in reverse to identify how consumption by nodes downstream can cause failures to meet consumption upstream in later time steps. When node 4 is removed, $E_3[\mathbb{P}(C_3^t - c_3^t \leq -1)]$ is reduced from 0.029 to 0.008, so the consumption at node 4 causes 0.021 of failures at node 3. $E_i[\mathbb{P}(C_i^t - c_i^t \leq -1)]$ is less than 0.001 for nodes 1 and 2, so these values are too small to reliably estimate the causal influence of downstream nodes. Of course, the nodes can be removed in any other permutation, which will show different causal relations. However, the motivation for recursively removing from one end of the supply chain is to determine the causal relation compared to the distance between two nodes, which is successively reduced by one when the nodes are removed in order.

The above analysis shows the impact of upstream consumption can be greatly magnified downstream. However, the five node supply chain is not long enough to determine whether the impact continues to grow or eventually dies out. Figure 6 shows simulation results from a 20 node logistics supply chain. The results are averaged over the set of holding policies within one unit distance of the solution found by SA, $H^* = (84, 87, 79, 94, 76, 84, 84, 62, 57, 48, 49, 46, 30, 36, 31, 24, 12, 18, 7)$. Only nodes 11 through 19 are graphed, since $\mathbb{P}(C_i^t - c_i^t \leq y)$ is less than 0.0001 for nodes 1 through 10, and node 0 is the source node. This graph shows that the proportion of failures due to node 1 across nodes 10 through 17 is significantly greater than any other node and that if this node were to be removed from the simulation then the majority of the failures at nodes 10 through 17 would be eliminated. However, nodes 18 and 19 show that the effect of node 1 reduces, both in absolute terms and as a percentage of total failures. Node 13 also causes magnified downstream failures, which peak at node 18 and are smaller at node 19. The impact of upstream consumption has not continued to grow unabated along the chain. We hypothesize this is due to the stochastic nature of consumption in the supply chain. Although an above average consumption at some time step upstream will have an adverse effect downstream, the further downstream a node is, the greater the probability is that intermediate nodes will have below average consumption. The further downstream a node is, the more fluctuations at the start of the chain are lost in the noise of the fluctuations at intermediate nodes. This would act to balance the apparent short range magnifications to produce a peak that eventually dies out.

The other interesting feature of Figure 6 is that only nodes 1, 13, 17 and 19 cause significant percentages of the failures in this supply chain. All of these nodes share the property that $H_i^t < H_{i+1}^t$, despite the fact that they must supply at least as much as node $i + 1$ every turn. This suggests that these nodes are weak points in the chain that may provide the greatest marginal return for increases in H_i^t , although we also note that other nodes share this property yet do not appear to be weak points. Therefore, this condition may be necessary but not sufficient for identifying potential weak points. Because the SA policy after a finite number of iterations is not guaranteed to be optimal, the causal analysis

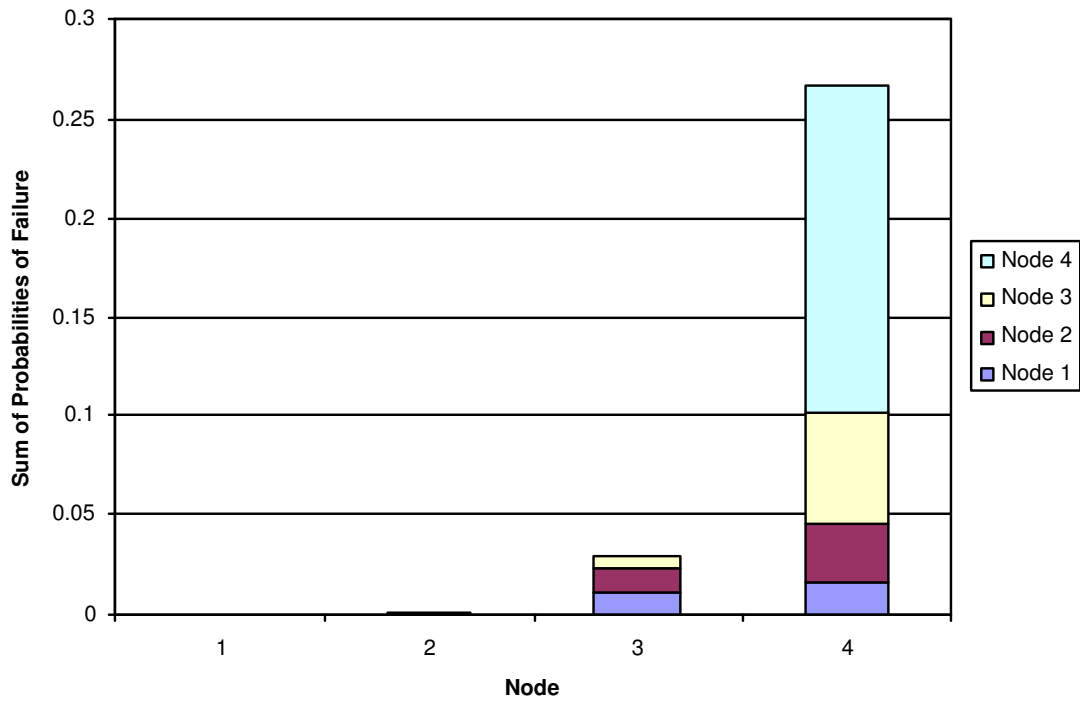


Figure 5: Average proportion of failures for a five node supply chain as nodes are removed from the start of the chain

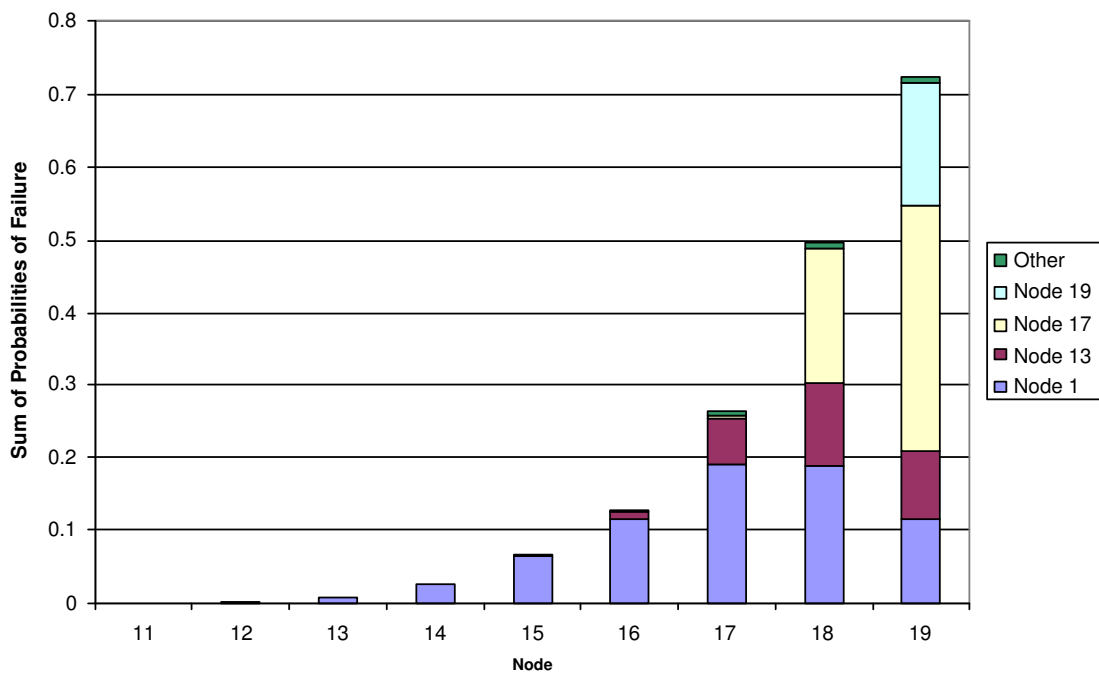


Figure 6: Average proportion of failures for a 20 node supply chain as nodes are removed from the start of the chain

method we have developed could potentially offer a constructive technique for improving on the SA policy.

3 Control Theory

In the previous section, the manner in which goods are routed or scheduled throughout the system is fixed, according to equations (6) and (7). Under this fixed scheduling, the holdings policies H_i^t are varied, subject to known constraints, with the objective to minimise stockouts and shortfalls in supply. In this section, we propose to fix the holdings policies H_i^t and determine the optimal schedule for goods using Control Theory under the assumption of perfect foresight in consumption. The performance of both techniques for scheduling goods are then tested under a relaxation of the perfect foresight assumption by incorporating stochastic noise into predictions for consumption throughout the system in both time and space. Additional comparisons with correlated consumption are also conducted to stress the system.

3.1 Multistage Decision Process

Consider a discrete process on the state space \mathcal{S} with a finite number of stages $\mathcal{N} = \{1, \dots, N\}$; that is, a multistage process with finite horizon. At each stage in the process $n \in \mathcal{N}$, a control decision $d_n \in \mathcal{D}$ is applied. The process obeys the transition function $\tau_n : \mathcal{S} \times \mathcal{D} \rightarrow \mathcal{S}$ such that

$$s_{n+1} = \tau_n(s_n, d_n), \quad s_n \in \mathcal{S}, d_n \in \mathcal{D}, n \in \mathcal{N}. \quad (26)$$

This describes a deterministic Markov decision process because the transitions between states are not stochastic and the state at each stage of the process depends only on the state of the system and a control decision at the previous stage. A logistics network in general could also be modelled to depend on an extended history of the process, beyond the previous stage. Such requirements can be built into the decision process under the transition function (26), for dependence on an extended history of finite length, with a suitable interpretation of the state space of the system which takes into account the physical condition of the system across multiple instances in time.

Of course, the set of feasible control decisions $\mathcal{F}_{n,s}$ at any given stage $n \in \mathcal{N}$ also depends on the state of the process $s \in \mathcal{S}$ so that $\mathcal{F}_{n,s} \subseteq \mathcal{D}$. Hence, we further require not only that $d_n \in \mathcal{D}$ but also that $d_n \in \mathcal{F}_{n,s}$, $s \in \mathcal{S}$ for every $n \in \mathcal{N}$. Also note that we can define $\mathcal{D} \equiv \cup_{n \in \mathcal{N}, s \in \mathcal{S}} \mathcal{F}_{n,s}$ so that the set of control decisions \mathcal{D} contains only feasible states.

At each stage $n \in \mathcal{N}$ the process is associated with a return function $v_n : \mathcal{S} \times \mathcal{D} \rightarrow \mathbb{R}$. The value

$$\nu = \sum_{n \in \mathcal{N}} v_n(s_n, d_n), \quad (27)$$

denotes the total return of the process over all stages.

Refer to Figure 7 for an illustrative description of the interactions in a single stage of the multistage decision process.

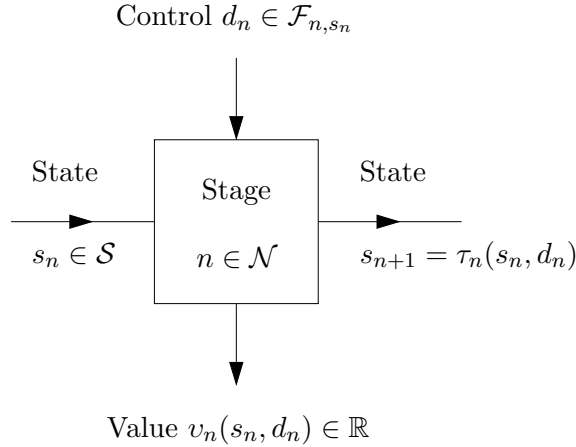


Figure 7: Multistage decision process

A deterministic multistage decision process (Sniedovich, 1992, pp.31–33) is then defined as the 6-tuple $D = (\mathcal{N}, \mathcal{S}, \mathcal{F}, \tau, v, s)$ where s denotes the initial state of the process with the objective to

$$\begin{aligned}
 & \text{maximise} && \nu &= & \sum_{n \in \mathcal{N}} v_n(s_n, d_n), \\
 & \text{subject to} && d_n &\in & \mathcal{F}_{n,s_n}, && n \in \mathcal{N}, \\
 & && s_{n+1} &= & \tau_n(s_n, d_n), && n \in \mathcal{N}, \\
 & && s_1 &= & s.
 \end{aligned} \tag{28}$$

That is, given a multistage decision process D , on a finite horizon of N stages with initial state s , determine the sequence of controls d_i , $i = 1 \dots N$, that maximises the value of the process v . We define this sequence of controls as an optimal policy of D .

3.2 Dynamic Programming

Bellman's Principal of Optimality (Bather, 2000, p.18) states that

An optimal policy has the property that, whatever the initial state and initial decision are, the remaining decisions must constitute an optimal policy with regard to the state resulting from the first decision.

This result means that, if the initial control decision of a multistage decision process is contained within an optimal policy for that process, then the remaining control decisions form an optimal policy for the remainder of the multistage decision process with the input state resulting from that initial control decision. Bellman (1952, 1957) used this principle to define a method of solving the multistage decision process using recursion. That is, starting at the final stage and last state of the process, work backwards to identify an optimal policy for the process. Hence, Bellman solved the multistage decision process as a number of sub-processes in which each process extended the optimal policy of the previous process until the optimal policy of the original multistage decision process had been calculated. This procedure is called Dynamic Programming, denoted earlier DP. Refer to Boudarel et al. (1971, pp.11–12) for a direct derivation of mathematics of the method.

To apply Bellman's method, recursively define the optimal value $\hat{v}_n(s_n)$ and optimal control \hat{d}_n , over the stages n through N , for the multistage decision process D , as

$$\hat{v}_n(s_n) = \max_{d_n \in \mathcal{F}_{n,s_n}} (v_n(s_n, d_n) + \hat{v}_{n+1}(\tau_n(s_n, d_n))), \quad n \in \mathcal{N}, \quad (29)$$

and

$$\hat{d}_n = \arg \max_{d_n \in \mathcal{F}_{n,s_n}} (v_n(s_n, d_n) + \hat{v}_{n+1}(\tau_n(s_n, d_n))), \quad n \in \mathcal{N}, \quad (30)$$

respectively, where $\hat{v}_{N+1}(s_n) \equiv 0$.

Then, to compute an optimal policy of D , we need only to recursively compute \hat{d}_n , s_n and \hat{v}_n .

For example, suppose that a supply node is provided with a schedule of consumption x_n over the next four months $n = 1 \dots 4$. The node initially has no stock in hand but does have an arbitrarily large warehouse capacity. However, the node is penalised for carrying over stock from previous months. A monthly penalty is invoked for holding excess unused stock. Furthermore, the supply node is unable to regulate changes in deliveries easily. Agencies delivering stock to the supply node charge a penalty for varying the delivery schedule between months. Then our Dynamic Program D becomes a minimisation problem and we associate the return function with the penalty function.

Let the consumption in the first, second, third and fourth months be $x_1 = 3$, $x_2 = 0$, $x_3 = 1$ and $x_4 = 1$ units of goods respectively. Let the state of the process during stage n be defined by the tuple (h_n, o_{n-1}) , where h_n denotes the stock holdings carried over in the warehouse from the stage $n - 1$ and o_{n-1} denotes the delivery received by the warehouse during stage $n - 1$. No stocks are carried over from $n = 0$ and no deliveries are received in $n = 0$. Then, let

$$v_n = 100h_n + 300|o_n - o_{n-1}|, \quad (31)$$

so that

$$\hat{v}_n(h_n, o_{n-1}) = \min_{y \geq x_n - h_n} (100h_n + 300|y - o_{n-1}| + \hat{v}_{n+1}(h_n + y - x_n, y)). \quad (32)$$

Table 2: Calculations for example Dynamic Program

Stage 4 (4th month)				Stage 3 (3rd month)				Stage 2 (2nd month)				Stage 1 (1st month)			
h_4	o_3	\hat{v}_4	\hat{d}_4	h_3	o_2	\hat{v}_3	\hat{d}_3	h_2	o_1	\hat{v}_2	\hat{d}_2	h_1	o_0	\hat{v}_1	\hat{d}_1
0	0	300	1	0	0	300	1	0	3	1100	1	0	0	2000	3
0	1	0	1	1	0	400	0	1	4	1600	1				
1	0	100	0	1	1	500	1	2	5	2000	0				
1	1	400	0	2	0	300	0								
1	2	700	0	2	1	600	0								
				2	2	900	0								

The optimal policy of our example is displayed in Table 2 in bold italic numerals. This policy is interpreted as:

- ordering 3 units goods in the first month, supplying 3 units goods and emptying the warehouse;

- ordering 1 unit goods in the second month and holding this unit in the warehouse;
- ordering 1 unit goods in the third month, supplying 1 unit goods and holding 1 unit in the warehouse; and
- supplying 1 unit goods and emptying the warehouse in the fourth and final month.

3.3 Logistics Supply Chain

The multistage decision process, for this general logistics supply network of Section 2.1, is posed as follows. Given H_i^t , W_i^t , C_i^t , Q , and $\Delta_{i,j}^t$ determine the $\delta_{i,j}^t$, $\gamma_{i,j}^t$, h_i^t , and c_i^t to

$$\begin{aligned}
\min \quad & \sum_{i \in \mathcal{N}/\{0\}, t \in \mathcal{T}} \pi(i, t) = A|H_i^t - h_i^t|^\alpha + B|C_i^t - c_i^t|^\beta, \\
\text{subject to} \quad & h_i^t = \min\{W_i^t, h_i^{t-1} - c_i^t - \sum_{j \neq i} \delta_{i,j}^t + \sum_{j \neq i} \delta_{j,i}^t\}, \quad i \in \mathcal{N}/\{0\}, \quad t \in \mathcal{T}, \\
& c_i^t = \begin{cases} h_i^{t-1}, & C_i^t > h_i^{t-1}, \\ C_i^t, & C_i^t \leq h_i^{t-1}, \end{cases} \quad i \in \mathcal{N}/\{0\}, \quad t \in \mathcal{T}, \\
& \sum_{j \neq i} \delta_{i,j}^t \leq h_i^{t-1} - c_i^t, \quad i \in \mathcal{N}/\{0\}, \quad t \in \mathcal{T}, \\
& 0 \leq \delta_{i,j}^t \leq \Delta_{i,j}^t, \quad i, j \in \mathcal{N}, \quad t \in \mathcal{T}, \\
& \delta_{i,j}^t = Q\gamma_{i,j}^t, \quad i, j \in \mathcal{N}, \quad t \in \mathcal{T}, \\
& h_i^0 = \lfloor H_i^0 \rfloor, \quad i \in \mathcal{N}, \\
& \gamma_{i,j}^t \in \mathbb{N}, \quad i, j \in \mathcal{N}, \quad t \in \mathcal{T}.
\end{aligned} \tag{33}$$

DP is well-suited to calculating an optimal policy for the controls in a multistage decision process for a linear supply chain - that is, solving (33) for the network topology of a linear supply chain. By the very definition of optimality, DP finds policies that are both temporally and spatially efficient. The DP model anticipates future consumption for goods and holds those goods dispersed along the chain, appropriately displaced in time, to best meet the consumption and hence maximise the difference in the penalties accrued for holding excess stock and the penalties accrued for failing to meet consumption.

To better examine the impact of the penalty function on the optimal policy for the DP model, consider a four node network with warehouse capacity of ten and holdings policy of four, over four stages, with a desired consumption of ten unit goods in the final two stages of the process at node three. For this experiment, the value of the optimal policy is not important and only one optimal policy is provided, although the nature of all optimal policies is discussed.

Let $h^t = (h_i^t)$, $i = 1, \dots, 3$ and let $h = (h^t)$, $t = 1, \dots, 4$. Equation (5) with $A = B = 1$ and $\alpha = \beta = 2$ results in three solutions, one of which is

$$h = ((5, 4, 4), (4, 6, 7), (4, 4, 6), (4, 4, 4)). \tag{34}$$

This solution is displayed in Figures 8 through 11. In these figures, the arcs denote goods, either transferred between nodes, consumption or holdings carried from the previous stage. These figures are not to be confused with Figure 7 in which arcs between nodes denote stages, the input arc denotes a control decision, and the output arc denotes a return value.

This optimal policy presents a demonstration of the temporal and spatial properties of the system. Notice that equation (34) results in a stockpile of 7 units at node 4 in stage 2 to meet the consumption of 10 (with a shortfall of 3 units) in the subsequent time period. This choice of stockpiling 7 units rather than 10 is a direct result of the penalty function. It costs $(7 - 4)^2 = 9$ units to hold the 7 units (an excess of 3 units on the desired holdings level) over a stage and it costs $(10 - 7)^2 = 9$ units in failing to meet the target quota for consumption of 10 units (a shortfall of 3 units on the desired level). Notice that these two cost terms are actually invoked in different stages. The total of these two terms is 18 units with an additional once off cost of 1 unit due to the holdings at the 1st stage. It is easy to see that this is the optimum behaviour because it costs $(8 - 4)^2 + (10 - 8)^2 = 20$ units to stockpile 8 units instead of 7 and $(6 - 4)^2 + (10 - 6)^2 = 20$ units to stockpile 6 units instead of 7. Both of these options, and indeed all other options, are worse than the optimum. Next, equation (34) indicates that 6 units of goods are stockpiled at node 3 (for node 4) in stage 2 (for consumption in stage 4). This costs $2^2 + 2^2 = 8$ units in holdings over two stages and $4^2 = 16$ units in failing to meet target consumption levels. The total cost is then 24 units. Again, stockpiling any amount of goods other than 6 units results in a higher total cost (eg stockpiling 7 goods results in a total cost of $3^2 + 3^2 + 3^2 = 27$ and stockpiling 5 goods results in a total cost of $1^2 + 1^2 + 5^2 = 27$).

In a military context, small stockpiles of goods are desirable because they have the effect of reducing the logistics footprint compared with the larger footprint under increased stockpiling. Infrastructure such as warehouses is also costly to establish and maintain. Furthermore, keeping large inventory levels with idle stock is wasteful of resources. Hence, A , B , α and β in equation (5) are important parameters in determining the penalty invoked for stockpiling. We provide three examples to illustrate the types of policies obtained under different parameterisations of the model.

Suppose that, $A = B = 1$ and $\alpha = \beta = 1$ in equation (5). Then there are five solutions, one of which is

$$h = ((4, 4, 4), (4, 4, 4), (4, 4, 4), (4, 4, 4)). \quad (35)$$

This optimal policy simply meets the warehouse holding policy levels of 4 units at each stage. The shortfall in consumption is not ignored, there is merely no advantage to be gained in stockpiling goods because the penalty invoked in holding excess goods exactly equals the penalty invoked in failing to meet desired consumption. Likewise, there is no advantage to be gained in not stockpiling goods. Hence, the other four solutions have a holdings level of 1, 2, 3, and 4 units respectively at node 4 in the 2nd stage (to be consumed in the 3rd stage). It is not efficient to stockpile goods over 2 stages, as it was in the previous example, because any penalty x is invoked twice for holdings above 4 units but the shortfall in consumption $10 - 4 - x$ is invoked only once so that the net penalty actually increases by x units. Hence, none of the five optimal policies attempts to stockpile any additional goods for the consumption at node 4 in stage 4.

When $A = B = 1$, $\alpha = 2$ and $\beta = 1$ the solution (35) is again optimal. However, in this example there is only one additional optimal policy

$$h = ((4, 4, 4), (4, 4, 5), (4, 4, 4), (4, 4, 4)). \quad (36)$$

Here it is costly to stockpile goods, hence the best policy is to simply meet target warehouse levels. The solution (36) is optimal because $1^2 = 1$ so the squared term for holding excess

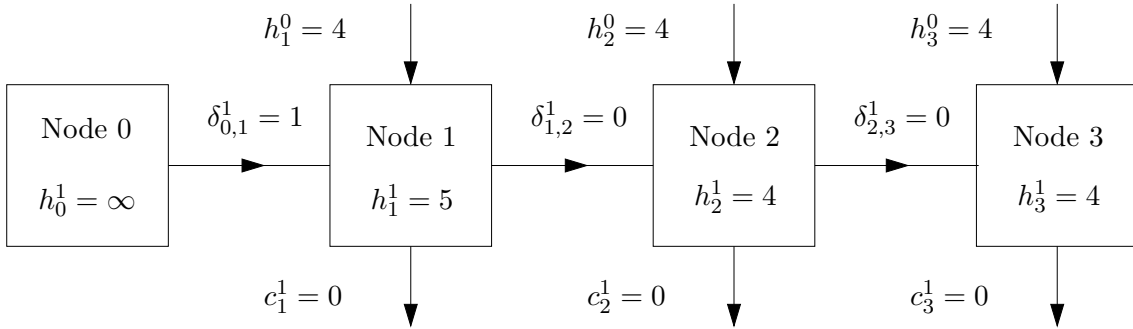


Figure 8: Linear supply chain stage $t = 1$

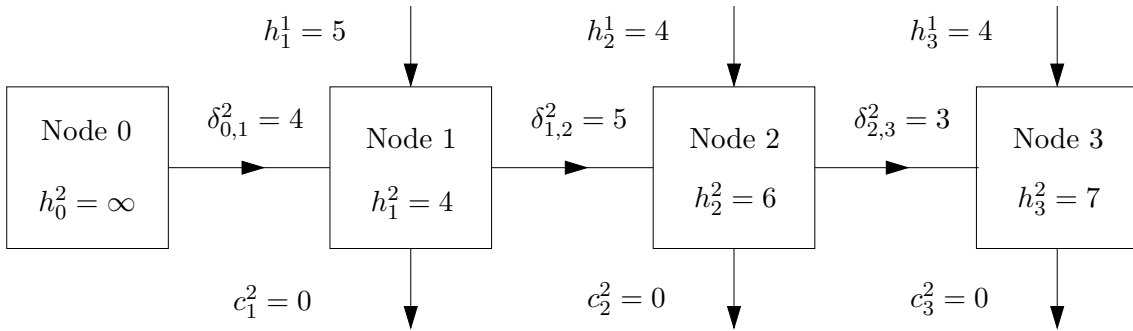


Figure 9: Linear supply chain stage $t = 2$

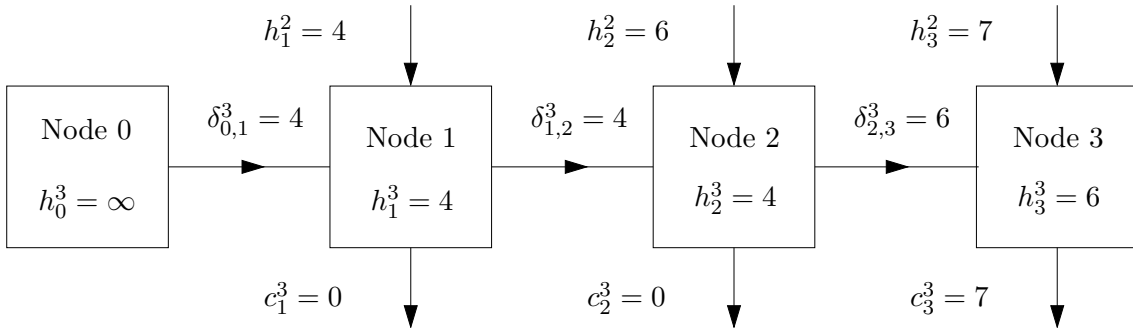


Figure 10: Linear supply chain stage $t = 3$

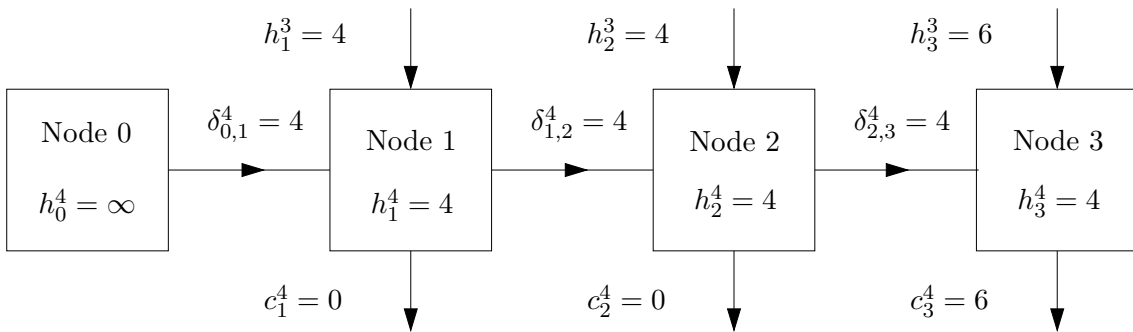


Figure 11: Linear supply chain stage $t = 4$

stock is exactly equal to the linear term for shortfall in consumption only when the amount of goods is unity.

If $A = B = 1$, $\alpha = 1$ and $\beta = 2$ we have twenty-eight solutions, including

$$h = ((8, 5, 4), (4, 8, 9), (4, 4, 8), (4, 4, 4)). \quad (37)$$

All solutions involve large transfers of goods. In this scenario, the cost of failing to meet consumption dominates the cost of overstocking the warehouses. However, even here no node stocks more than 9 units of goods at any stage.

4 Comparison of Systems under Uncertainty and Correlation

4.1 Degradation of Optimality under Uncertainty

In this section, we propose that the consumption schedule, C_i^t , $i \in \mathcal{N}$, $t \in \mathcal{T}$, is uncertain. Instead, we have an *a priori* estimate \hat{C}_i^t of the true consumption C_i^t . In the last section, this estimate coincided with the realisation so that $\hat{C}_i^t = C_i^t$. Here we use

$$\hat{C}_i^t = \max\{0, C_i^t + \varepsilon_i^t\} \quad (38)$$

where ε_i^t is an error term.

In this study, we generate the error term ε_i^t according to a Bernoulli sequence as follows. First, ε_i^t is initialised to 0. Define the random variable Z with two outcomes, either *success* with fixed probability p or *failure* with probability $1 - p$. Then, a perturbation is repeatedly introduced to ε_i^t by repeatedly adding a uniformly distributed random integer over -3 to 3 so long as the random variable $Z = \textit{success}$. Refer to Appendix C for an illustrative depiction of the relationship between the value of the error term $\varepsilon_i^t \equiv \varepsilon$ and the probability that it is observed for values of p between 0.1 and 0.6.

Under uncertainty, the T stage DP model is solved in multiple instances. Initially, one complete T stage optimal policy is computed using the inexact approximations \hat{C}_i^t to C_i^t . The true values of C_i^1 are then revealed. With this information, an optimal policy for the remaining $T - 1$ stages is re-computed using an appropriately adjusted Dynamic Program and the true values of C_i^2 are revealed. This step process of solution and re-computation is repeated until all values C_i^1, \dots, C_i^T are known.

Six logistics chains are compared in Figure 12 each of length 4 with consumption C_i^t generated on the integers between 0 and 3. In each case, the three numbers after the initial designator DP and SR describe the holdings policies of nodes 1 through 3 respectively. Each chain is independently simulated 100 times for 15 time steps. The cost or penalty of each system is measured according to equation (5) with $A = B = 1$ and $\alpha = \beta = 2$. Figure 12 illustrates the outcome of this comparison for values of p between 0 and 0.6 in increments of 0.1. The error bars describe 99% confidence intervals.

There are infinitely many possible logistics networks and ways to introduce error and uncertainty into those systems. Hence, Figure 12 does not itself tell us much about the

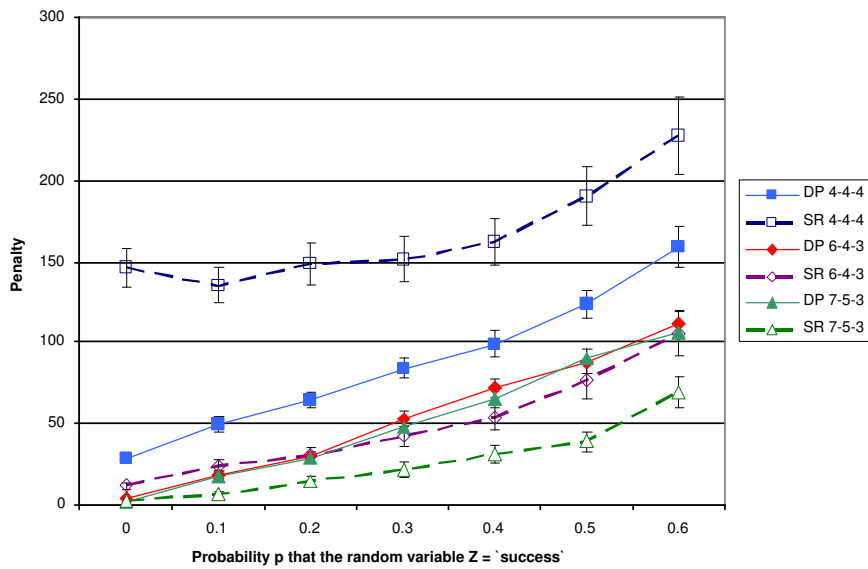


Figure 12: Comparison of average penalty for SR and DP algorithms over the level of error in prediction

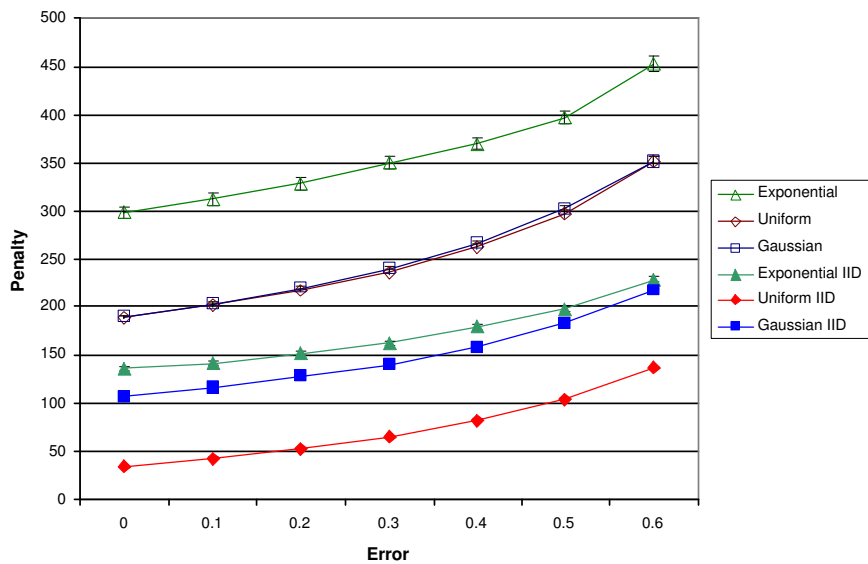


Figure 13: Comparison of average penalty for SR algorithm with independent and correlated consumption generated from Uniform, Normal and Exponential distributions

interplay between the two approaches across all possible logistics networks. However, certain enduring features do arise. The DP model is optimal only when $p \equiv 0$. The utility of this approach decreases as uncertainty increases. In comparison, the SR system of Section 2.2 has no concept of estimating future consumption and is in that sense robust to uncertainty. For all logistics systems studied, one is presented with the problem of determining which of the two approaches is most appropriate given the level of uncertainty in future consumption. In general, we observe that the SR system performs best when the holdings policies are set relatively large. In fact, when running DP with policies (7, 5, 3), there is little visible difference in curves between simulations with policies (6, 4, 3) while the performance of the SR (7, 5, 3) system is substantially improved. On the other hand, DP handles situations in which the holdings policies are relatively low substantially better than SR dynamics. While both approaches benefit from increased holdings policies, the marginal utility gained by increasing holdings policies is different. In both approaches, where the relative benefit of increasing holdings policies are negligible, the system has reached a “saturation” point in which desired consumption is always or almost always satisfied and any penalty incurred is due to a failure to resupply the system to the desired holdings levels.

4.2 Degradation of Optimality under Correlation

In this section, the uncertain consumption schedule C_i^t , $i \in \mathcal{N}$, $t \in \mathcal{T}$ is further relaxed to disregard the assumption of independence of consumption at different nodes. This is motivated by the observation that in complex systems, correlations between nodes in a network can result in unexpected behaviour. One source of correlation may be that the consumption at each node in the supply chain occurs on the same battlefield. Then during a major attack, high consumption at one node could be experienced at the same time as high consumption at other nodes. There exists both historical and simulation evidence that power laws exist in casualty statistics of warfare (Richardson, 1941, 1960; Roberts and Turcotte, 1998; Lauren, 2001), which could indicate correlation in the activities of combat units. Since these units are responsible for setting the consumption schedule it is important to investigate the effect of correlations on the performance of the logistics supply chain.

We compare the performance of the DP and SR algorithms under independent and correlated consumption generated by Uniform, Normal and Exponential distributions. The independent consumption schedules with uncertainty are calculated as per Section 4.1, where each node draws a random number from a distribution and adds error ε_i^t . The three cases investigated are $X_i \sim \{U[0, 5], N(2.5, 2), E(2.5)\}$, which have equal means. The correlated consumption is generated by a single distribution for the entire chain, that when divided amongst the L nodes is indistinguishable from the independent consumption schedules, with the exception that correlation between any two nodes is almost one¹. For $L = 3$ this gives the distributions $X \sim \{U[0, 15], N(7.5, \sqrt{12}), E(7.5)\}$. Total consumption when x^t is drawn from X is $x^t + \varepsilon^t$, and consumption at node i equals

¹The correlation is not exactly 1 because L may not divide $x^t \in X$ exactly.

$$C_i^t = \begin{cases} \lfloor \frac{x^t + \varepsilon^t}{L} \rfloor, & i \notin \mathcal{R}, \\ \lceil \frac{x^t + \varepsilon^t}{L} \rceil, & i \in \mathcal{R}, \end{cases} \quad (39)$$

where \mathcal{R} is the set of randomly chosen nodes that are allocated an extra unit of consumption from the remainder of the Integer division.

For both DP and SR, for each distribution the (7, 5, 3) holding policy performs better than (6, 4, 3) and (4, 4, 4) for all values of ε^t . Therefore, we limit the following comparisons to the (7, 5, 3) holding policy. Figure 13 shows the average penalty for the SR algorithm with independent and correlated consumption distributions over 1000 replications, where each supply chain is simulated for 1000 time steps. For every distribution, the penalty increases nonlinearly with error, such that the rate of increase in penalty grows with the size of ε . All independent and identically distributed (denoted IID in Figure 13) consumption schedules perform significantly better than correlated consumption schedules. On average, the SR algorithm has a penalty of 131 for independent consumption, which more than doubles to 288 with correlated consumption. This is despite the fact that total consumption of the independent and correlated supply chains have identical distributions - it is only the way consumption is distributed to individual nodes that varies. The complete correlation modelled here represents an extreme example of correlation, and serves to quantify an upper bound on the effect of correlations between nodes in a logistics supply chain.

Figure 13 also contains insights into the effect of different types of distributions on supply chain performance. For a given level of ε , for both correlated and independent cases the following ordering holds within the 99% confidence intervals: the Uniform penalty is less than or equal to Normal penalty, which is strictly less than the Exponential penalty. This is intuitive, since a major contributor to the penalty is when the consumption level exceeds the holding level at a node. The Uniform distribution is bounded, which bounds the size of consumption at any time step. The Normal distribution has a thin tail compared to the Exponential distribution, meaning very large consumption is much more likely with an Exponential distribution. Interestingly, for independent nodes, the performance is significantly better with the Uniform distribution, while the penalty is similar between the Exponential and Normal distributions. However, when the node consumption is correlated, the penalty is not significantly different between the Uniform and Normal distributions, which are both considerably better than the Exponential distribution. Therefore, the penalty with the Normal distribution is bounded by the Uniform and Exponential distributions, and varies within this range as a function of the degree of correlation.

The performance of the DP and SR with Normal and Exponential correlated consumption schedules is compared in Figure 14. The results for the Uniform correlated distribution is not shown because there is no significant difference to the Normal correlated consumption performance for either algorithm. In all cases, DP performs significantly better than SR. The performance of both algorithms is significantly worse when the total consumption across the supply chain is generated from the Exponential distribution, even though the mean consumption is the same as the Normal and Uniform consumption schedules.

The error bars for the SR data in Figure 14 are much smaller and the curves are smoother because the results are averaged over 1000 time steps compared to only 15 time steps for DP. When the SR data is run for only 15 time steps the penalty function is slightly

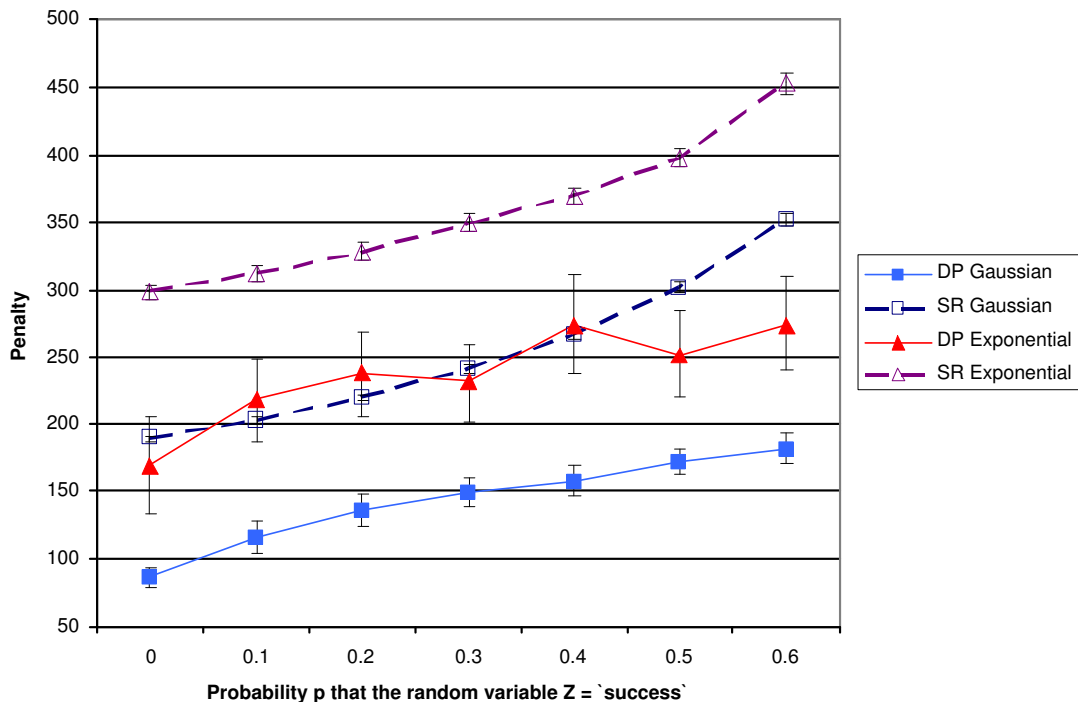


Figure 14: Comparison of average penalty for SR and DP algorithms when node consumption is correlated

underestimated (although the 1000 step results are still well within the confidence intervals for the 15 step data), since the initial conditions $h_i^0 = H_i^0 \forall i$ produces upward bias in the steady state estimate of the expected penalty. DP is only run for 15 time steps because its computational complexity prevented solving a 1000 step problem within a reasonable amount of time with the available computing resources. One way to improve the performance of DP is to consider a finite time horizon of say 15 time steps, and compute a solution to the 1000 step problem by successively solving the problem for segments of 15 steps. Another option is to consider solving the DP model for disjoint subsets of nodes, where each subset must set its ordering policy independently of nodes from other subsets. Both approaches have reduced time complexity and can be considered approximations to the full DP solution. Neither of these methods are implemented in this study, since the 15 time step data provides statistically significant results for comparisons with the SR algorithm.

5 Discussion

Time-invariance is assumed in Section 2.3 for tractability and ease of analysis. This assumption is reasonable in a subset of all possible systems in the real world, those being systems whose underlying distributions for consumption do not change in time. Cyclic phenomenon or other more complex system dynamics are not captured. For example,

consumption in a logistics system varies in time but the way in which it varies may fundamentally change between day and night. A day and night scenario can be broken into two independent time-invariant systems, one for daytime and one for nighttime. At the other extreme, a system whose distributions vary each and every new time step may to all practical purposes be sufficiently random to be adequately modelled as a time-invariant system, although this is by no means guaranteed. Between simple cyclic behaviour and total unpredictability lies system dynamics which are partially predictable and therein display trend-like behaviour yet are sufficiently complex to defy routine analysis. In such systems, adaptive learning, or a similar technique, is required to teach the system how to detect, anticipate and react to changes within the system.

The formulation for the logistics network of Section 2.1 includes a maximum link capacity $\Delta_{i,j}^t$ between nodes i and j , and a maximum warehouse capacity W_k^t , for nodes k and time t . These capacity constraints are used in this report indirectly and are somewhat peripheral to our results. Future applications of this research include a direct study of the effects of link and warehouse capacities by including additional cost terms to the penalty function (5). This future work complements our study with a focus on logistical network design rather than policy setting and optimal routing in existing systems. Further extensions include varying the stock multiplier Q , which describes the container or parcel size. The benefits and costs between various methods of resupply, which deliver stock in different base quantities, can then be gauged against the number of packages required and the frequency of delivery required.

This study identified interesting nonlinear dynamics that could be explored in further detail. In particular, the phase transitions observed in the Lagrangian relaxation were peripheral to the scope of this report but suggest the potential for further research into the general existence and cause of phase transitions in simple linear supply chains. Due to time limitations, only a 20 node supply chain was tested for correlation and downstream effects, but larger empirical studies using the methods developed in Section 2.4 may lead to more conclusive results. The SR model may show improved performance if adaptation is incorporated using a machine learning technique such as Reinforcement Learning. An adaptive model would also be a more accurate representation of the SR logistics concept. An obvious extension is to consider logistics networks rather than only linear supply chains. It is not obvious whether a network is more or less vulnerable to uncertainty and correlation, or what network topology provides the best performance in different contexts. This study would provide a useful baseline for any research into more general network models of SR logistics.

A limitation to the use of DP is the underlying assumption of perfect foresight. Put simply, the future consumption is known or estimated *a priori*. This assumption is reasonable for capturing predictable behaviour. For example, over the duration of a day, a force in combat may expend a well known quantity of munitions. It may be reasonable in such a situation to assert that the expenditure of munitions today will be similar to that of yesterday or perhaps a weighted average of the expenditure over the last week. This is called naïve and adaptive expectations respectively (Lucas, 1986). Perfect foresight is also reasonable in situations of routine delivery, such as supplying provisions for infantry, and any kind of planned activity, such as a planned offensive attack, where one can estimate the requirement for medical supplies with reasonable certainty. In any case, in conducting

a planned offensive attack, ordinance, medical supplies, provisions and other goods are ordered and stored in advance and not upon action due to the non-zero delay in logistical support. Finally, DP is reasonable to use when rationing goods. In this situation, the true real demand for goods exceeds the actual preferred consumption level set in the DP model. Consumption across the network is deliberately capped.

A three faceted approach to the study of logistics supply chains is recommended for a balanced program of work across a broader context. This report delivers one of these approaches by exploring the use of the techniques from statistics, control theory and optimisation. The remaining two recommended studies employ classical network theory and queuing theory respectively. In the first instance, transportation problems, multi-commodity flow problems and scheduling problems have the potential to generate valuable insights when appropriately modelled in a logistics context. In the second instance, metrics such as average occupancies, mean waiting times and stationary distributions are useful in measuring performance in logistics systems. The union of these three studies then comprise a well-rounded and balanced research program designed to deliver coherent, relevant and useful recommendations to the logistics community.

6 Conclusions

This study has presented a foundation for the analysis of logistics supply chains without back-orders. Two models were developed. First, a purely reactive demand-oriented model was formulated without any prediction or estimation of future demand. This model was used to determine policies for holdings levels throughout the chain which minimise the probability of stock-outs and shortfalls in supply. Second, a deliberately planned supply-oriented model was formulated. This model was used to determine the optimal supply regime which minimises the cost associated with over and under supply. Finally, a comparison between the two methods was conducted where demand is uncertain and cannot be accurately predicted.

It is concluded that both the demand oriented and supply oriented models are useful in context of the predictability and uncertainty of future demand in the system. While the demand oriented model exhibits far from optimal performance, its utility or benefit does not depend on being able to predict future demand. Conversely, the optimal supply regime in the supply oriented model explicitly requires an accurate forecast of future demand throughout the system. The utility or benefit obtained in using demand-oriented mechanics substantially decreases when the holdings policies in the system limit throughput. However, while supply oriented techniques are superior in efficiency and effectiveness in these cases, there is a point at which uncertainty in the system becomes excessive. For any real logistics system, identifying the boundary between the techniques in terms of uncertainty and sufficient throughput is necessary to determine which of the two options is most appropriate and relevant.

Correlations between demand at different nodes in a supply chain can significantly reduce the performance of the system, which highlights the need to incorporate a level of buffering in warehousing policies to counter this effect when correlations are possible. The statistical nature of demand is also an important factor in the performance of the system, and

additional buffering is recommended for logistics supply chains when statistical inference from historical data indicates demand is fat-tailed.

The results of this study do not directly contradict the ideas presented in the Australian Defence Force's *Future Joint Logistics Concept* (Commonwealth of Australia, 2002) but instead serve as a warning that responsive systems, simple or complex, will not and cannot entirely replace traditional planned systems and that the problems associated with the design, implementation and analysis of adaptive precision systems are non-trivial. This study tentatively supports the conclusion that enhanced capabilities, emerging technologies and network-enabled warfare will ultimately facilitate better management and control of logistics systems. However, it does so on the basis that these new developments will reduce uncertainty across the system and therein have an indirect effect on the system rather than directly supporting better logistics supply. Further work at the Defence Science and Technology Organisation is currently underway to explore these issues in greater detail

7 Acknowledgements

Thanks to Axel Bender, Adrian Pincombe and Maria John, Land Operations Division, for their invaluable assistance with this study.

References

- Bather, J. (2000). *Decision Theory. An Introduction to Dynamic Programming and Sequential Decisions*. Wiley-Interscience Series in Systems and Optimisation. West Sussex, England: John Wiley and Sons.
- Bellman, R. (1952). On the theory of dynamic programming. In *Proceedings of the National Academy of Sciences*, Volume 38, pp. 716–719.
- Bellman, R. (1957). *Dynamic Programming and its application to Optimal Control*. Princeton, New Jersey: Princeton University Press.
- Bender, A. (2004). Forecasting Under Logistics Uncertainty. Unpublished.
- Boudarel, R., J. Delmas, and P. Guichet (1971). *Dynamic Programming and its application to Optimal Control*, Volume 81 of *Mathematics in Science and Engineering*. New York, New York: Academic Press.
- CAPT Lewandowski, L. and J.R. Cares (2005). *Sense and Respond Logistics: The Fundamentals of Demand Networks* (Long ed.). United States of America: Office of Force Transformation.
- Cerny, V. (1985). Thermodynamical Approach to the Traveling Salesman Problem: An Efficient Simulation Algorithm. *Journal of Optimization, Theory and Applications* 45(1), 41–51.

- Commonwealth of Australia (2002). *Future Joint Logistics Concept*. Australia: Army Headquarters, Future Land Warfare Branch.
- Commonwealth of Australia (2004). *Complex Warfighting*. Australia: Army Headquarters, Future Land Warfare Branch.
- Everett, H. (1963). Generalized langrangean multiplier method for solving problems of optimum allocation of resources. *Operations Research* 2, 399–417.
- Fisher, M. L. (1985). An application oriented guide to lagrangean relaxation. *Interfaces* 15, 7–26.
- Jolliffe, I. T. (1986). *Principal Component Analysis*. New York: Springer-Verlag.
- Kirkpatrick, S., C. D. Gelatt Jr, and M. P. Vecchi (1983). Optimization by simulated annealing. *Science* 220, 671–680.
- Lauren, M. (2001). Describing Rates of Interaction between Multiple Autonomous Entities: An Example Using Combat Modelling, <http://aps.arxiv.org/pdf/nlin/0109024>, Cornell Univeristy, accessed 1 July 2005, Defence Technology Agency, New Zealand Defence Force.
- Lucas, R. E. (1986). Adaptive behaviour and economic theory. *Journal of Business* 59(4), s401–s426.
- Metropolis, N., A. Rosenbluth, A. Teller, and E. Teller (1953). Equation of state calculations by fast computing machines. *Journal of Chemical Physics* 21(6), 1087–1092.
- Richardson, L. F. (1941). Frequency of occurrence of wars and other fatal quarrels. *Nature* 148(598).
- Richardson, L. F. (1960). *Statistics of Deadly Quarrels*. Boxwood Press, Pittsburgh.
- Roberts, D. C. and D. L. Turcotte (1998). Fractality and self-organised criticality of wars. *Fractals* 6, 351–357.
- Sniedovich, M. (1992). *Dynamic Programming*. Pure and Applied Mathematics, A Program of Monographs, Textbooks, and Lecture Notes. New York, New York: Marcel Dekker.
- United States Department of Defense (2004). *Operational Sense and Respond Logistics: Coevolution of an Adaptive Enterprise Capability, Concept Document* (Long ed.). United States of America: Office of Force Transformation.
- United States Department of Defense (2005). *Operational Sense and Respond Logistics: Coevolution of an Adaptive Enterprise Capability, Metrics Overview*. United States of America: Office of Force Transformation.
- Wolpert, D. H. and W. G. Macready (1997). No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation* 1(1), 67–82.

Appendix A Geometry of the State Space

The success of any global optimisation heuristic depends on the geometry of the fitness landscape and the underlying state space. Wolpert and Macready (1997) have shown that over all fitness landscapes, no heuristic is better than any other. Therefore, if the SA algorithm we developed in Section 2.3 is to be effective, it must exploit some systemic feature of the state space. Most optimisation techniques are sensitive to the number of local minima on the fitness landscape. For very rugged landscapes, the SA algorithm would require a very slow cooling of the temperature parameter to avoid becoming trapped in one of many suboptimal local minima that may exist at a large distance (measured by the minimum number of operations of the SA neighbourhood function) from the global minimum. In this section, we develop a distance measure and use Principal Components Analysis to better understand the geometry of the state space.

Consider a linear supply chain of length five (including the source node 0) where each node has consumption generated according to a Normal distribution $N(5, 3)$. Using brute force enumeration, we explore a region of integer valued solutions to the holding policy for each node. That is, for every possible combination of node holding policies within the region, the chain is simulated for 1000 time steps. The region is defined by setting

$$H_i^{low} = \sum_{j \geq i} \mu_j, \quad (\text{A1})$$

and

$$H_i^{high} = 5 + \sum_{j \geq i} \mu_j, \quad (\text{A2})$$

where μ_j is the mean of the distribution at node j . Thus the (H_i^{low}, H_i^{high}) tuples for nodes $1 \leq i \leq 4$ are $(20, 25), (15, 20), (10, 15), (5, 10)$. The minimisation problem defined in (12) with the second objective function $\sum \mathbb{P}(C_i^t - c_i^t \leq y)$ is used, with $y = -1$ and $H_{max} = 60$. Here, 60 is chosen to provide the maximum number of 146 valid combinations where $\sum H_i = H_{max}$ within the region defined by (A1) and (A2), although the following results are found to hold for other values of H_{max} . The interpretation of setting $y = -1$ is that desired consumption exceeds actual consumption, which means the supply chain is unable to meet demand. Henceforth, we refer to the objective using the abbreviated notation $\sum \mathbb{P}$.

The obtained optimal solution for a simulation of 1000 time steps is $H^* = (21, 17, 12, 10)$ which has $\sum \mathbb{P} = 0.173$. This is a boundary solution, since node four is at its maximum value. The objective function $\mathcal{L}(\psi)$ is graphed for each of the 146 combinations where $\sum H_i = H_{max}$ in Figure A.1. As can be seen from Figure A.1, the optimal solution is the 27th combination. The combinations are ordered by enumerating between H_i^{low} and the H_i^{high} for each i , incremented in reverse node order and satisfying $H_{max} = 60$. The first combination in Figure A.1 is $(20, 15, 15, 10)$ and the last is $(25, 20, 10, 5)$.

Although Figure A.1 has many local minima, this is mainly an artifact of ordering four dimensional data on one dimension of the graph. To understand the space better we measure the distance in state space between the best configurations, where the state is completely defined by the L -tuple H . The distance in state space is defined as the minimum

number of units of stock that must be switched from one warehouse to another in order to move between two stocking policies. Formally, for $\mathcal{L} : \mathbb{N}^L \times \mathbb{N}^L \rightarrow \mathbb{R}$,

$$\mathcal{L}(G, H) = \frac{1}{2} \sum_i |G_i - H_i| \quad (\text{A3})$$

where G_i, H_i represent the holding policy at node i under policies G and H respectively. Table A.1 lists the 10 best configurations, each with value less than 0.03 worse than the optimal solution, along with the five worst holding policies. The 10 best holding policies equate to points on the eight local minima that are below the 0.2 value in Figure A.1. Four of the top ten solutions are within distance one of the optimal policy H^* , which has nine holding policies within distance one in total. Further, the remaining six solutions in the top ten are all only two units from H^* . A result that is not shown in Table A.1 is that the first policy to have a distance of three or greater from H^* is the 18th best solution overall.

If we consider the five worst policies in Table A.1, we observe that two are seven units distance away from the optimal policy. Seven is the maximum distance from the optimal policy, and these are the only two policies that exist with $\mathcal{L}(H, H^*) = 7$.

Table A.1: Ten best and five worst holding policies H ranked according to $\sum \mathbb{P}$ compared to their distance to the optimal solution H^*

H	$\mathcal{L}(H, H^*)$	$\sum \mathbb{P}$
(21, 17, 12, 10)	0	0.173
(22, 16, 12, 10)	1	0.186
(21, 18, 13, 8)	2	0.186
(21, 17, 13, 9)	1	0.187
(20, 18, 12, 10)	1	0.19
(22, 17, 13, 8)	2	0.191
(20, 17, 14, 9)	2	0.193
(21, 19, 11, 9)	2	0.193
(22, 18, 12, 8)	2	0.196
(22, 17, 12, 9)	1	0.199
(22, 16, 13, 9)	2	0.203
(25, 19, 11, 5)	6	0.455
(25, 15, 15, 5)	7	0.457
(24, 20, 11, 5)	6	0.459
(24, 18, 13, 5)	5	0.463
(25, 20, 10, 5)	7	0.542

If the value of node 1 is held constant, the remaining three dimensions for nodes 2 through 4 can be visualised. Figure A.2 shows the $\sum \mathbb{P}$ as the area of the bubbles for each holding policy when node 1 equals 22. We can see that the landscape is quite smooth, since the bubbles vary in size following a regular pattern.

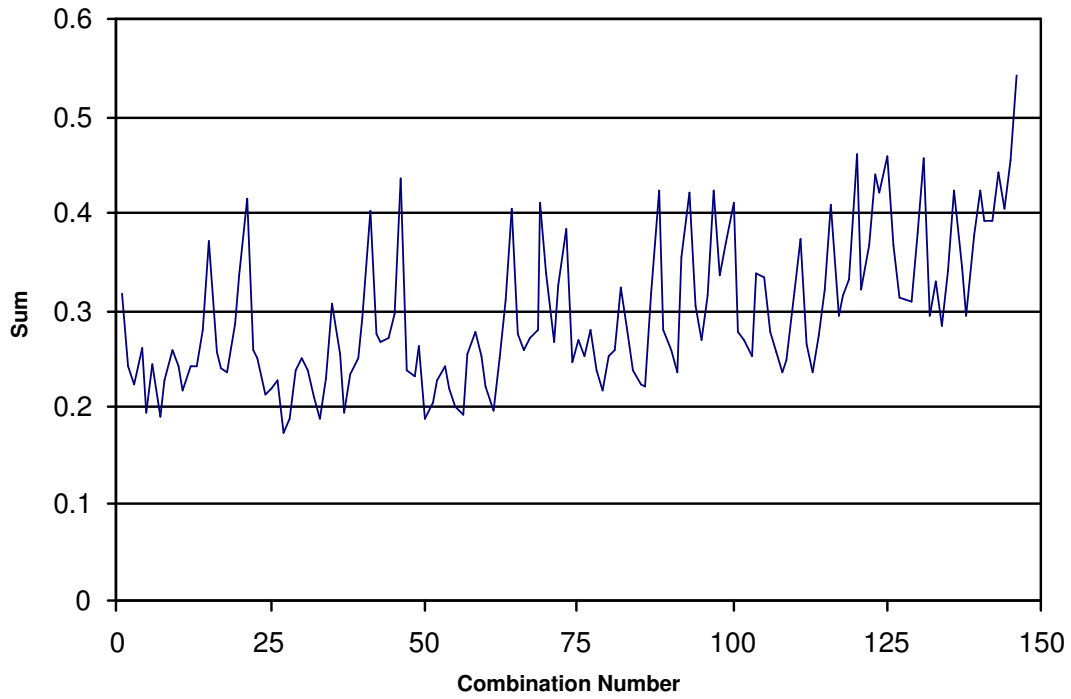


Figure A.1: $\sum \mathbb{P}$ for each brute force enumeration for a five node supply chain

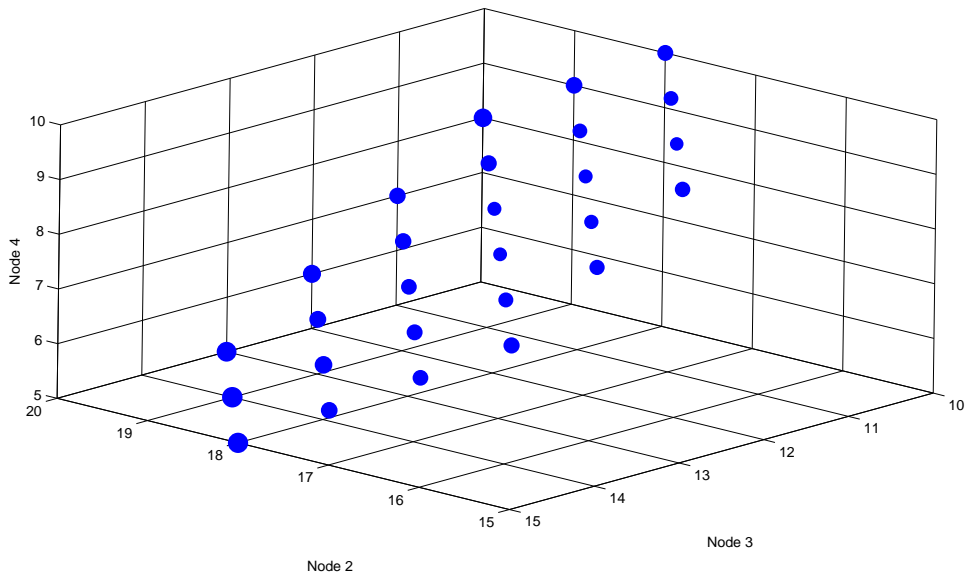


Figure A.2: Bubble scatter plot of $\sum \mathbb{P}$ for a five node supply chain when node 1 = 22

In order to visualise the landscape for more than three nodes it is necessary to perform a dimensionality reduction. Principal Components Analysis (PCA) (Jolliffe, 1986) is a factor analysis method that is used to perform a variance maximising rotation of the original state space. By selecting the p dimensions under the rotation with the largest variance, the data is projected onto p orthogonal factors with the minimum error to the original data under the ℓ^2 -norm. The principal components are given by the first p eigenvectors of the correlation matrix, ordered by the magnitude of the associated eigenvalues.

The state space includes the L node holding policies in the natural numbers, not including the source node 0, plus an additional non-negative real dimension for $\sum \mathbb{P}$. First, we consider the brute force results for the region defined by (A1) and (A2) regardless of whether $H_{max} = 60$ is satisfied. This consists of five columns of data x_i , $1 \leq i \leq 5$, one for each dimension of the state space, and 1296 rows, one for each brute force combination of node holding policies. The correlation between two columns of data is given by

$$\text{cor}(x_i, x_j) \equiv \frac{\text{cov}(x_i, x_j)}{\sigma_i \sigma_j}, \quad (\text{A4})$$

where σ_i is the standard deviation of x_i . The covariance is given by

$$\text{cov}(x_i, x_j) \equiv E((x_i - \mu_i)(x_j - \mu_j)), \quad (\text{A5})$$

where E is the mathematical expectation and μ_i is $E(x_i)$. The correlation matrix between each column is shown in Table A.2. Because the state space is completely enumerated over the 1296 observations, there exists no correlation between the nodes. We observe on average, as each node increases its holding policy, that $\sum \mathbb{P}$ decreases, and the magnitude of the correlation increases for nodes further down the supply chain. Therefore, if all nodes have the minimum stock levels $H = (20, 15, 10, 5)$ we would expect the marginal utility of adding extra holding to be greatest at node 4 and smallest at node 1.

Table A.2: Correlation matrix between nodes and $\sum \mathbb{P}$ from brute force enumeration of a five node supply chain

	Node 1	Node 2	Node 3	Node 4	$\sum \mathbb{P}$
Node 1	1				
Node 2	0	1			
Node 3	0	0	1		
Node 4	0	0	0	1	
$\sum \mathbb{P}$	-0.2055	-0.3104	-0.4193	-0.7688	1

The eigenvectors and eigenvalues are shown in Table A.3. The percentage of variation in the direction of each eigenvector is calculated from the relative size of the corresponding eigenvalue. The first eigenvector, which is a combination of all four node holding policies plus 50% of the $\sum \mathbb{P}$ dimension, captures 39% of all variation. The next three dimensions are independent of $\sum \mathbb{P}$ and account for a further 60% of the variation. The final eigenvector is a reflection on a hyperplane of the first eigenvector, which can be thought of as the error of the first eigenvector in accounting for the variation in $\sum \mathbb{P}$ as a function of the node holding policies. The final eigenvector only contributes 1% to the total variation. Therefore, the underlying data is essentially four dimensional, but the redundant

Table A.3: Percentage of variation for principal components from brute force enumeration of a five node supply chain

Eigenvectors	Eigenvalues	Percentage of Variation
(-0.1527, -0.2307, -0.3116, -0.5713, 0.7071)	1.9516	39.0%
(0.9762, -0.0559, -0.0922, -0.188, 0)	1	20.0%
(0.0173, -0.9422, 0.109, 0.3163, 0)	1	20.0%
(-0.0079, -0.0521, 0.8862, -0.4602, 0)	1	20.0%
(0.1527, 0.2307, 0.3116, 0.5713, 0.7071)	0.0484	1.0%

dimension is a combination of all five of our original dimensions, so none of the original dimensions can be completely eliminated.

The first three principal components graphed in Figure A.3 account for 80% of the total variation. This space is well behaved in the sense that changes in the second and third principal components, which are independent of $\sum \mathbb{P}$, result in smooth changes to the first principal component that captures the majority of the variation in $\sum \mathbb{P}$. The other relevant feature of Figure A.3 is the way in which the data is clustered into six layers according to their second and third principal components.

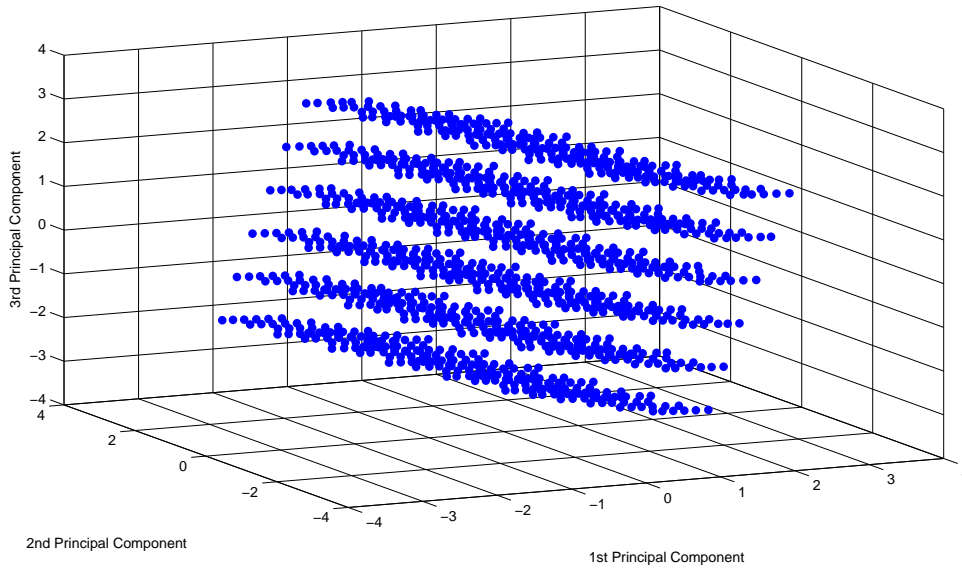


Figure A.3: First three principal components from brute force enumeration of a five node supply chain

We now consider the effect of the additional constraint $\sum H_i = 60$ on the principal components. Because $\mathbb{P}(C_i^t - c_i^t \leq y)$ is a monotonic decreasing function of $H_i \forall i$, we consider only the Pareto dominant subset of $\sum H_i \leq H_{max}$ given by $\sum H_i = H_{max}$, which contains 146 different holding policy combinations. This introduces a relation between the values of the nodes which lie on the Simplex defined by $\sum H_i = 60$. Specifically, if we know H_i for all nodes except one node j we can deduce that $H_j = 60 - \sum_{k \neq j} H_k$. The correlation matrix for the brute force enumeration given $\sum H_i = 60$ is given in Table A.4. Compared to the correlation matrix in Table A.2, it can be observed that Table A.4 contains corre-

lations between nodes due to the additional constraint. The only node to have a negative correlation with $\sum \mathbb{P}$ is node 4. The reason the other nodes are positive is because any increase in say node 1 will result in a reduction in the holding policy for some node further down the chain and so $\sum \mathbb{P}$ can increase even when $\mathbb{P}(C_i^t - c_i^t \leq y)$ decreases. The first three nodes have decreasing positive correlations with $\sum \mathbb{P}$ as the node number increases, so the sign but not the ordering of the marginal utility has changed compared to Table A.2.

Table A.4: Correlation matrix when $\sum H_i = 60$

	Node 1	Node 2	Node 3	Node 4	$\sum \mathbb{P}$
Node 1	1				
Node 2	-0.3333	1			
Node 3	-0.3333	-0.3333	1		
Node 4	-0.3333	-0.3333	-0.3333	1	
$\sum \mathbb{P}$	0.502	0.2459	0.0144	-0.7623	1

The largest eigenvalue in Table A.5 shows that 42.5% of all variation can be captured by a combination of three nodes plus the $\sum \mathbb{P}$. Although all four nodes have non-zero values, node three's value of 0.0117 is not significant. With three dimensions, 95.9% of the variation is be represented, while four dimensions accounts for all of the data. This means the fifth dimension, which is an equal combination of all four node holding policies is redundant.

Table A.5: Percentage of variation for principal components between nodes and $\sum \mathbb{P}$ from brute force enumeration of a five node supply chain given $\sum H_i = 60$

Eigenvectors	Eigenvalues	Percentage of Variation
(0.4068, 0.1992, 0.0117, -0.6177, 0.6428)	2.1266	42.5%
(0.1012, 0.5099, -0.8266, 0.2155, 0)	1.3333	26.7%
(0.6766, -0.6499, -0.2578, 0.2311, 0)	1.3333	26.7%
(-0.3413, -0.1672, -0.0098, 0.5183, 0.766)	0.2067	4.1%
(0.5, 0.5, 0.5, 0.5, 0)	0.0000	0.0%

The first three principal components for the brute force enumeration given $\sum H_i = 60$ are graphed in Figure A.4. The Simplex relation between the nodes has reduced the six distinct layers of clustered points to a single plane. The first principal component now captures all variation due to $\sum \mathbb{P}$, and changes to the second or third principal component produces smooth changes to the first factor. Further, the minimum value along the first principal component occurs at a corner of the plane, and other low values are nearby in this space, while large values in the first factor are distant. This is the most accurate visualisation of the state space that is possible in three dimensions, and confirms the intuition we developed earlier when analysing the distances in Table A.1. Therefore, we conclude that for this particular system under the distance measure (A3) the geometry is suitably well behaved to expect good performance using SA with a relatively rapid cooling schedule.

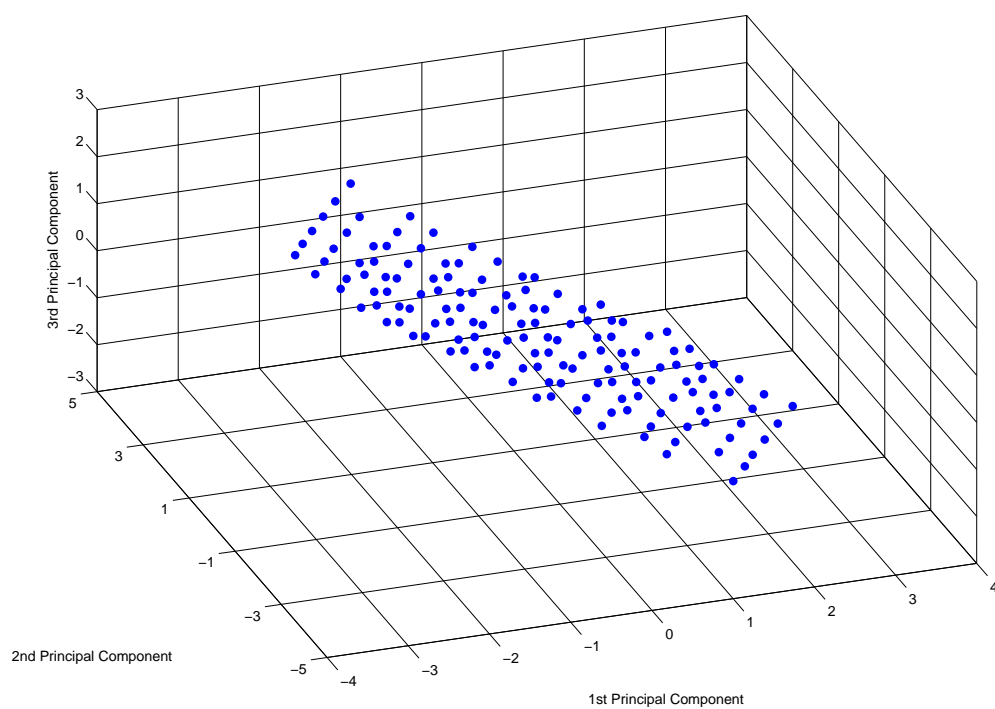


Figure A.4: First three principal components from brute force enumeration of a five node supply chain given $\sum H_i = 60$.

Appendix B Simulated Annealing

SA is a heuristic technique loosely based on the physics of thermal equilibria and the properties of materials like metal and glass. In such materials, adding heat disrupts the state of the material and through cooling the material is annealed to alter its properties. Typical examples of this process include glass blowing and the tempering of steel. The SA algorithm is conceptually modelled on this process and is applied as an heuristic optimisation technique for problems which are otherwise intractable or difficult to solve. The algorithm is traditionally posed to minimise energy levels in a system and therein stabilise it. The algorithm successively introduces one or more perturbations, disturbances or transitions into the system in an attempt to restabilise it in more desirable states. If the perturbation produces a favourable result, then the new system state is accepted. Otherwise, the new state is accepted on the basis of some pre-determined stochastic acceptance function which classically resembles the Boltzmann distribution.

Denote the space of all possible states \mathcal{S} over which the SA algorithm searches. The algorithm is initialised by identifying the initial state $s_0 \in \mathcal{S}$ of the system being annealed and the cooling schedule of temperatures (T_j) , whose elements $T_j \in \mathcal{T} \subset \mathbb{R}^+$ are monotonically decreasing, where $i = j \dots J$ corresponds to an iteration counter of the algorithm for some fixed number of annealing steps J . The probability of accepting a transition depends on the objective function, often referred to as the energy function $E_T : \mathcal{S}^2 \rightarrow \mathbb{R}$, and is defined as

$$\mathbb{P}(\text{accept new state } \hat{s} \mid \text{current state } s) = \min \left\{ 1, e^{\frac{-(E_T(\hat{s}) - E_T(s))}{kT}} \right\}, \quad (\text{B1})$$

for some scaling constant k , where $T \in \mathcal{T}$ is the constant denoting the current temperature level of the algorithm. The probability of accepting transitions leading to less desirable states potentially allows the algorithm to escape local minima.

The algorithm iteratively conducts a biased random walk over the state space \mathcal{S} at each temperature $T_j \in \mathcal{T}$, initially starting in $s_0 \in \mathcal{S}$, as follows. Set the current state of the system $s = s_0$. For each temperature T_j in order: generate a new state \hat{s} ; accept or reject the state \hat{s} as the current state s according to (B1); do this for K time steps. This loop of K steps is referred to as a Metropolis loop after Metropolis et al. (1953). The number of steps in the Metropolis loop and the cooling schedule are tuning parameters that are set based on constraints on available processor time and an evaluation of how well the problem is solved (usually this means satisfying some criteria of utility combined with evidence that the SA is able to converge when the best solution does not change over a large number of steps). We expect SA will find a near optimal solution even with a rapid cooling schedule due to the nature of the solution space (see Appendix A).

Appendix C Error Component

To illustrate the value of the error component ε in Section 4.1, we generated 50,000 independent error terms and counted the frequency of each outcome, see Figure C.1. In this figure, values outside of $[-6, 6]$ are not displayed as their frequency is too low.

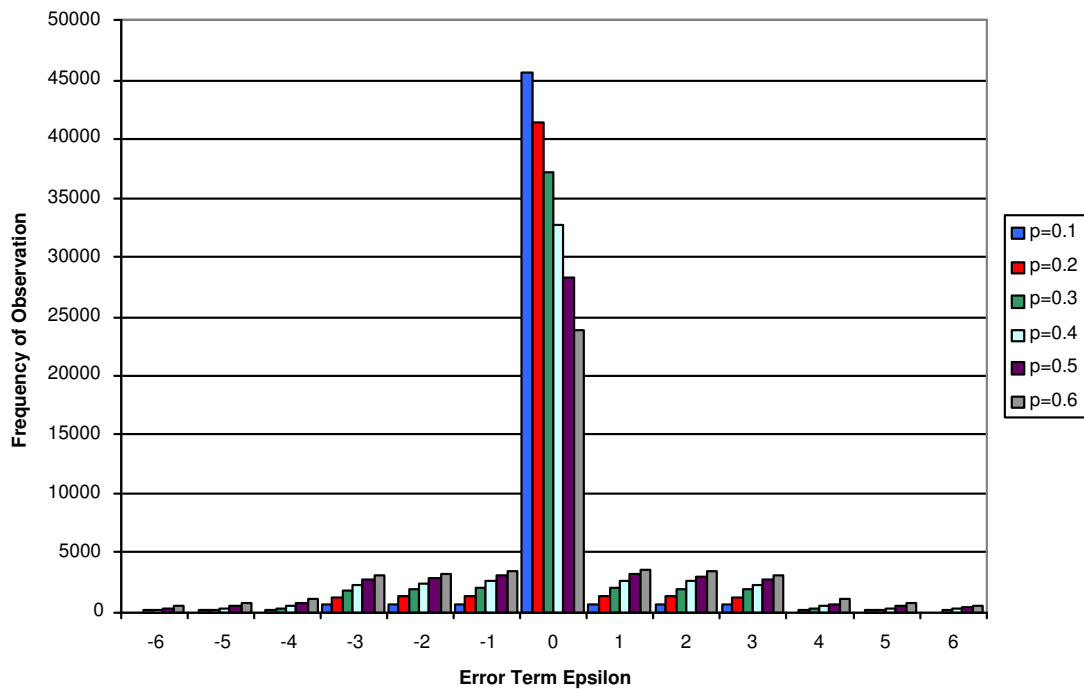


Figure C.1: Frequency of outcome against observed error in 50,000 independent trials

DISTRIBUTION LIST

Modelling Logistics Chains for Optimal Supply and Maximum Throughput

Scott Wheeler, Alex Ryan and Tristan Turner

AUSTRALIA

DEFENCE ORGANISATION

	No. of copies
Task Sponsor	
DGLD	1 printed
S&T Program	
Chief Defence Scientist	email pdf
Deputy Chief Defence Scientist Policy	email pdf
AS Science Corporate Management	email pdf
Director General Science Policy Development	email pdf
Counsellor Defence Science, London	Doc Data Sheet
Counsellor Defence Science, Washington	Doc Data Sheet
Scientific Adviser to MRDC, Thailand	Doc Data Sheet
Scientific Adviser Joint	email pdf
Navy Scientific Adviser	Doc Data Sht & Dist List
Scientific Adviser - Army	1 printed
Air Force Scientific Adviser	Doc Data Sht & Dist List
Scientific Adviser to the DMO	Doc Data Sht & Dist List

Systems Sciences Laboratory

Chief of Land Operations Division	Doc Data Sht & Dist List
Research Leader Land Systems	Doc Data Sht & Dist List
Research Leader Operations Research	1 printed
Heads:	
Lin Zhang	1 printed
Darren Reid	1 printed
Task Managers:	
Axel Bender	1 printed
Alex Ryan	1 printed
Authors:	
Scott Wheeler	1 printed
Tristan Turner	1 printed

DSTO Library and Archives

Library Edinburgh	2 printed
Defence Archives	1 printed

Capability Development Group

Director General Maritime Development	Doc Data Sheet
Director General Capability and Plans	Doc Data Sheet
Assistant Secretary Investment Analysis	Doc Data Sheet

Director Capability Plans and Programming	Doc Data Sheet
Director General Australian Defence Simulation Office	Doc Data Sheet

Chief Information Officer Group

Head Information Capability Management Division	Doc Data Sheet
AS Information Strategies and Futures	Doc Data Sheet
Director General Information Services	Doc Data Sheet

Strategy Group

Director General Military Strategy	Doc Data Sheet
Assistant Secretary Governance and Counter-Proliferation	Doc Data Sheet

Navy

Director General Navy Capability, Performance and Plans, Navy Headquarters	Doc Data Sheet
Director General Navy Strategic Policy and Futures, Navy Headquarters	Doc Data Sheet
Maritime Operational Analysis Centre, Building 89/90 Garden Island Sydney NSW	
Deputy Director (Operations)	
Deputy Director (Analysis)	shared Doc Data Sht & Dist List

Army

ABCA National Standardisation Officer, Land Warfare Development Sector, Puckapunyal	e-mailed Doc Data Sheet
SO (Science) - Land Headquarters (LHQ), Victoria Barracks NSW	Doc Data & Exec Summ
SO (Science), Deployable Joint Force Headquarters (DJFHQ) (L), Enoggera QLD	Doc Data Sheet

Air Force

SO (Science) - Headquarters Air Combat Group, RAAF Base, Williamtown NSW 2314	Doc Data Sht & Exec Summ
--	--------------------------

Joint Operations Command

Director General Joint Operations	Doc Data Sheet
Chief of Staff Headquarters Joint Operations Command	Doc Data Sheet
Commandant ADF Warfare Centre	Doc Data Sheet
Director General Strategic Logistics	Doc Data Sheet
COS Australian Defence College	Doc Data Sheet

Intelligence and Security Group

Assistant Secretary Concepts, Capability and Resources	email pdf
DGSTA DIO	1 printed
Manager, Information Centre, DIO	email pdf
Director Advanced Capabilities, DIGO	Doc Data Sheet

Defence Materiel Organisation

Deputy CEO	Doc Data Sheet
Head Aerospace Systems Division	Doc Data Sheet
Head Maritime Systems Division	Doc Data Sheet

Chief Joint Logistics Command	Doc Data Sheet
Program Manager Air Warfare Destroyer	Doc Data Sheet
Guided Weapon and Explosive Ordnance Branch	email Doc Data Sheet

OTHER ORGANISATIONS

National Library of Australia	email pdf
NASA (Canberra)	email pdf

UNIVERSITIES AND COLLEGES

Australian Defence Force Academy	
Library	email pdf
Head of Aerospace and Mechanical Engineering	email pdf
Serials Section (M list), Deakin University Library, Geelong, VIC	email pdf
Hargrave Library, Monash University	Doc Data Sheet

OUTSIDE AUSTRALIA

INTERNATIONAL DEFENCE INFORMATION CENTRES

US Defense Technical Information Center	email pdf
United Kingdom - Dstl Knowledge Services	email pdf
Canada - Defence Research Directorate R&D Knowledge & Information Management (DRDKIM)	email pdf
NZ Defence Information Centre	email pdf
Michael Lauren, Dstl, UK	1 printed

ABSTRACTING AND INFORMATION ORGANISATIONS

Library, Chemical Abstracts Reference Service	email pdf
Engineering Societies Library, US	email pdf
Materials Information, Cambridge Scientific Abstracts, US	email pdf
Documents Librarian, The Center for Research Libraries, US	email pdf

SPARES	5 printed
--------	-----------

Number of copies:

Printed	19	
PDF	20	
Total		39

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA				1. PRIVACY MARKING/CAVEAT (OF DOCUMENT)	
2. TITLE Modelling Logistics Chains for Optimal Supply and Maximum Throughput			3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION) Document (U) Title (U) Abstract (U)		
4. AUTHORS Scott Wheeler, Alex Ryan and Tristan Turner			5. CORPORATE AUTHOR Defence Science and Technology Organisation PO Box 1500 Edinburgh South Australia 5111 Australia		
6a. DSTO NUMBER DSTO-TR-1807		6b. AR NUMBER AR 013-547		6c. TYPE OF REPORT Technical Report	7. DOCUMENT DATE November 2005
8. FILE NUMBER 2005/1046841	9. TASK NUMBER ARM 04/037 LRR 02/229	10. TASK SPONSOR DGLD CLOD	11. NO. OF PAGES 40		12. NO. OF REFERENCES 23
13. URL on the World Wide Web http://www.dsto.defence.gov.au/corporate/reports/DSTO-TR-1807.pdf				14. RELEASE AUTHORITY Chief, Land Operations Division	
15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT <i>Approved for public release</i>					
OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE, PO BOX 1500, EDINBURGH, SA 5111					
16. DELIBERATE ANNOUNCEMENT No Limitations					
17. CITATION IN OTHER DOCUMENTS Yes					
18. DEFTTEST DESCRIPTORS Logistics management Logistics planning Logistics information systems Supply management Optimisation					
19. ABSTRACT This report presents an analysis of a generalised logistics supply chain without back orders. Two methods are proposed: a responsive but robust delivery system based on maintaining set holdings levels; and a technique from Control Theory which pushes stock through the chain in anticipation of demand over both time and space. Furthermore, a heuristic is proposed to set the policy for holdings levels using a hybrid of statistical analysis, Simulated Annealing and Lagrangian Relaxation. Finally, a comparison between methods under uncertainty, with error in prediction and correlated demand, is conducted. Each method was found to be useful in different contexts. The impact of uncertainty and correlated consumption was quantified for a set scenario and both were found to be significant factors in the performance of the supply chain.					