4. STRUCTURE OF THE PROPOSED APPROACH

Our evaluation of strategic options for the design and control of the spare parts logistics system has focused on three levels of analysis. At the first level we consider alternative configurations of distribution centers, parts stations, outside locations and central warehouses. At the second level we consider optimal, part class specific, stocking levels at each available stocking location. This stocking level analysis is carried out for each alternative design structure to be considered. Finally, at the third level of analysis, we consider alternative control strategies for inside/outside locations; stock allocation to all stocking points, pooling (or "vanning") of assigned (outside location) stock and use of new technologies (advance diagnostics, customer engineer telecommunication and automatic dispatching). The overall hierarchy of models and their inter-relationships are summarized in Figure 4.

It is important to note the linkages between decisions and outputs of the various models in the hierarchy. The evaluation of alternative logistic system structures can only be undertaken if the inventory cost and service related consequences of using that structure are taken into account. In order to do so, the part (or part class) specific inventory stocking model must be used. In our structure, the inventory stocking model takes as input the set of available stocking locations for the configuration in question. Cost parameters which must be specified include holding, normal

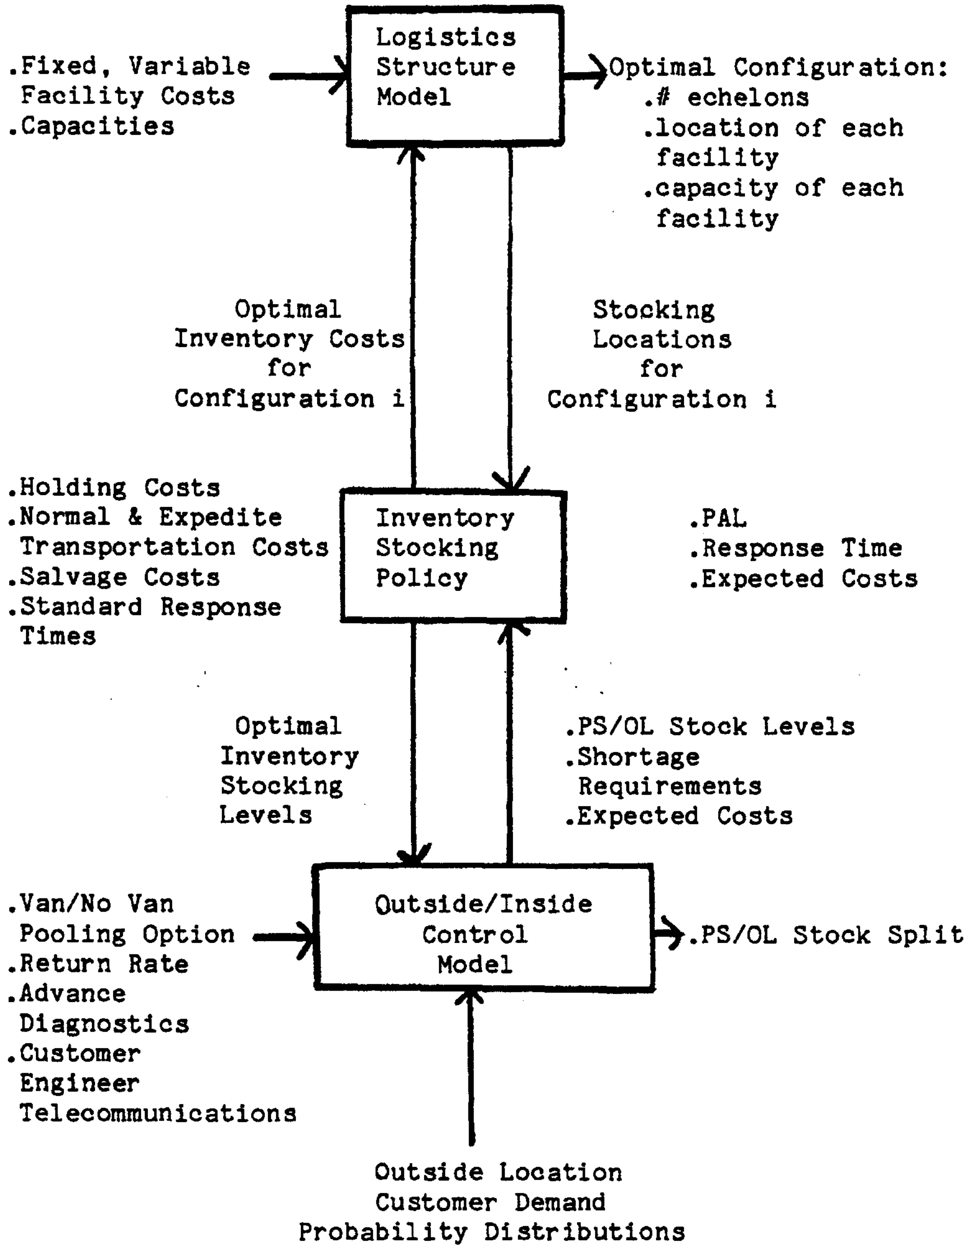

Figure 4: IBM Parts Logistics System Model Hierarchy

transportation and expedite transportation costs. The inventory stocking model also takes as input standard response times between all relevant pairs of stocking locations under both normal and expedite modes. The consequences of the outside/inside pooling and control arrangements are generated as outputs from the outside/inside control model which are then input as well into the inventory stocking model. Outputs of this detailed model include optimal stocking levels for various outside locations and parts stations, the expected shortage requirements under these stocking levels for the given part and stocking configuration, and a number of service related measures.

The logistics structure model considers the system wide consequences of each configuration by summing the inventory related costs generated by the stocking model over all parts. The total inventory cost can then be compared with fixed and variable facility costs for each configuration in order to arrive at an optimal structure for the parts logistics system. The outside/inside control model of the hierarchy is actually an adjunct to the inventory stocking model. It is clear that the specific operating parts use) characteristics of Customer Engineers is critically important in determining the demands placed on the parts inventory. At our current stage of analysis, cost functions and optimization procedures are used to determine the split between outside and inside inventory locations given a total stock allocation to the parts stations: Two generic strategies related to stock pooling ("vanning" and "no vanning") are also considered. A further key input is a part specific return rate which is used to capture the difference between disbursements and usage as a result of the role of parts as diagnostic tools. At a later stage of development we plan to use this model to analyze the impact of technological changes such as advance diagnosis or enhanced customer engineer telecommunication and dispatching systems as well as the impact of product life cycle and salvage cost considerations.

In the remainder of this section we consider the specifics and rationale for the various models used in the hierarchy. In particular, assumptions, decisions, inputs, outputs and structural relationships will be described. We begin with a brief description of the model development process.

i) Model Development:

The analysis of service parts inventory systems is extremely sensitive to the nature of the processes which generate Customer Engineer demand for parts. The actual use of parts is, of course, complicated by: 1) the fact that parts serve as tools; 2) the underlying reliability of machines in the field which is, for the most part, quite high; and 3) the accountability of CEs for parts inventory is unclear at best. Notwithstanding this complex set of phenomena which together comprise the determinants of parts demand, it was clear upon our analysis of IBM parts utilization and disbursement data that demand for parts can be broken down into two, quite different, classes. The most common, and as we have seen, the most problematic parts class is associated with extremely low usage (national average annual demand $\leq$ 20 units per month). At the other extreme, there are a minority of parts with relatively medium to high demand rates whose activity account for up to 15% of total parts numbers. Our modelling approach has led to the specification of two models for each part demand class.

The first model, which was implemented and used in the current phase of the analysis, is based on the key assumption that the probability of two or more successive demands for the part occurring in a period of time which is less than the normal replenishment time is very low. In these situations it is accurate to consider on hand inventory to be greater than or equal to the order-up-to stocking levels at each stocking point. These levels are dictated by the inventory stocking policy in force. In a sample of 108 parts we verified that the on-order position at all locations except for Mechanicsburg is virtually non-existent. As noted previously, Mechanicsburg's on-order position reflects their buffering tactics in face of an eight month manufacturing delivery lead time. Additional analysis of nationwide disbursement patterns also indicated that the time between successive demands at a customer location was greater than one month in the vast majority of cases for the medium to low demand classes. This evidence suggests that the assumption for on-hand inventory levels is indeed valid down to the customer machine level for the majority of parts numbers and/or dollars in the system.

The second model is based on traditional multi-echclon inventory theory and practice which is in use in military and relatively high volume commercial logistics systems. In these situations the demand rate is sufficiently high so as to render demand probability family identification and parameter estimation at a feasible parts station level. In most cases, documented in the inventory literature, Poisson or variants of the compound Poisson distribution family are found to apply. The key factor to take note of in these types of demand environments is that available inventory (stock-on-hand) is a random variable. Indeed, within a typical replenishment cycle a random number of demands will be realized before the current on-order stock quantity arrives. While shortages may or may not occur in a particular instance, expected availability of stock is definitely affected by the phenomenon of multiple demand arrivals during the normal replenishment lead time. In multi-echelon systems the flows of stock and demand are in general quite complex. If one makes the additional assumption of one-for-one replenishment quantities, then this system can be analyzed in terms of a queueing theory analog. In particular, steady state results can be derived to relate stocking level decisions to the probability distribution for different levels of stock-on-hand. This distribution can then be used to develop functions relating inventory stocking levels to expected inventory cost and service levels. The adaptation of this approach to the IBM parts maintenance system has been initiated.

As noted above, we have implemented the first (low demand) model since the majority of dollars and items are to be found in the low usage categories. It is also apparent that these items are the ones for which control is most difficult both due to the very nature of intermittent demand systems and the paucity of knowledge about recommended controls at both the level of theory and practice. Finally, our analysis of the current system indicated that the procedures and systems currently in place (PIMS) reflect a multi-echelon stocking philosophy based on procedures more appropriate to high usage environments. The further development of all model structures (high and low demand and outside/inside control) will be discussed in some detail in the section on recommended follow-on studies to be found in Section 7.

ii) Models Structure:

A. Logistics Structure Model

The logistics structure model requires specification of a discrete set of possible design options. Associated with each option are: 1) the number of echelons; 2) the location of each stocking facility within each echelon; and 3) the capacity for storage of parts at each facility. Associated with each facility are annual fixed and variable costs. Our current analysis has led to the estimation of these parameters for the current distribution centers (see Section 5). Evaluation of these costs at the parts station level is hampered by the unavailability of parts station specific allocated costs in the branch office accounts. Estimates based on a step-down of distribution center costs can be used in the analysis. The model seeks to minimize total (sum over all parts) inventory plus operating costs by selecting the optimum configuration of facilities. Since the inventory stocking model can be optimized for each possible design alternative, both expected inventory costs and optimum stock levels are known. Model constraints include a provision for every part to be covered by facilities and total facility capacity limits.

B. Inventory Stocking Model

The inventory stocking model can best be understood by referring to Figure 5. Demand for parts is generated by parts failures at the customer machine level. Once a part is detected as being needed the MPLS will respond to fill that need. The model depicts the interaction of random demands, positioned stock levels at each location, and both the normal and expedite delivery modes by considering a typical time period of sufficient length to exceed the maximum lead time in the entire system. As noted previously, this time period will still be short enough so that all on-order levels are assumed to be zero. The following dynamics are considered within the time period.

1. Pre-position stock to each stocking location so as to bring the inventory level up to a specified stocking level by an optimal routing of shipments from all possible sources to all destinations (i.e., solve a linear programming transportation problem).

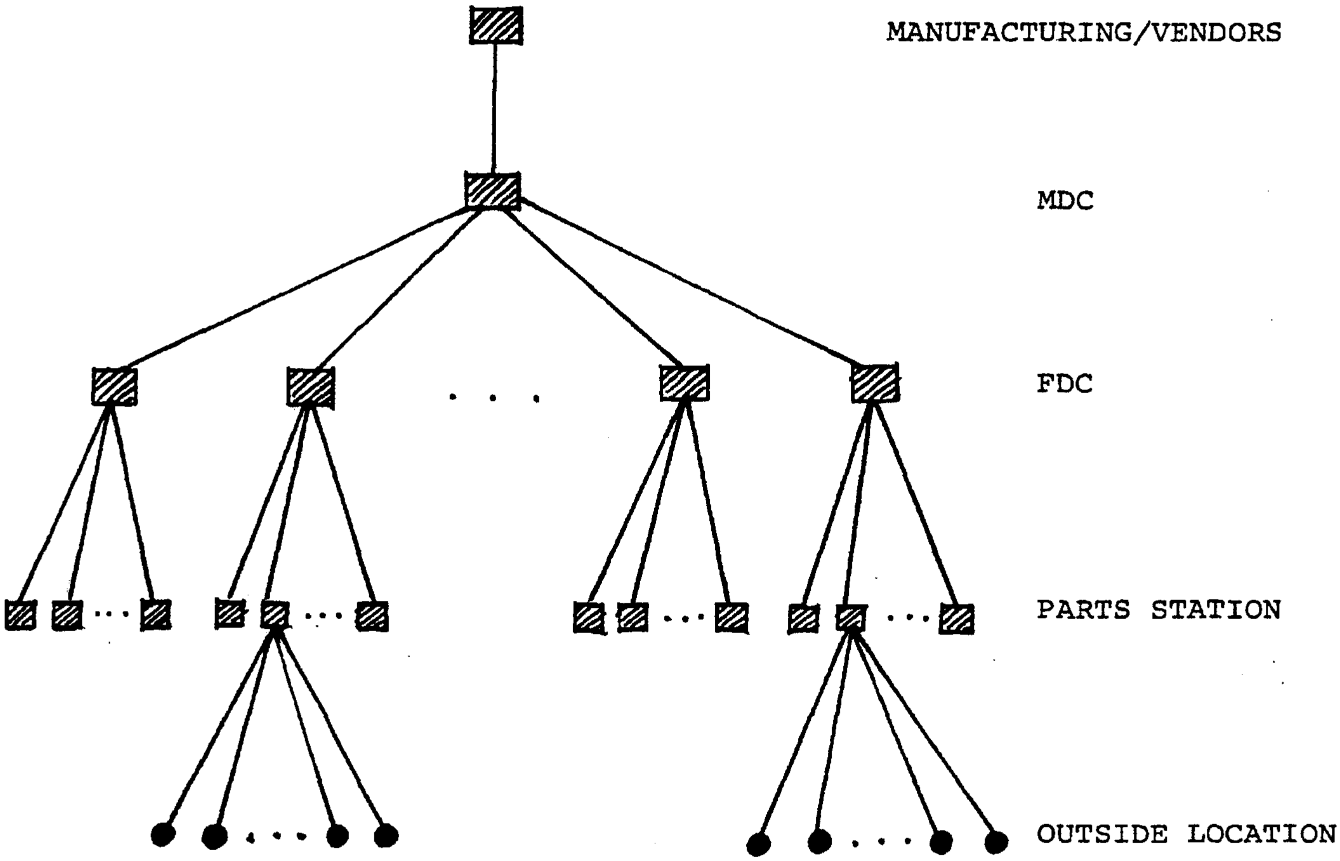

Figure 5a. Inventory Stocking Policy Model

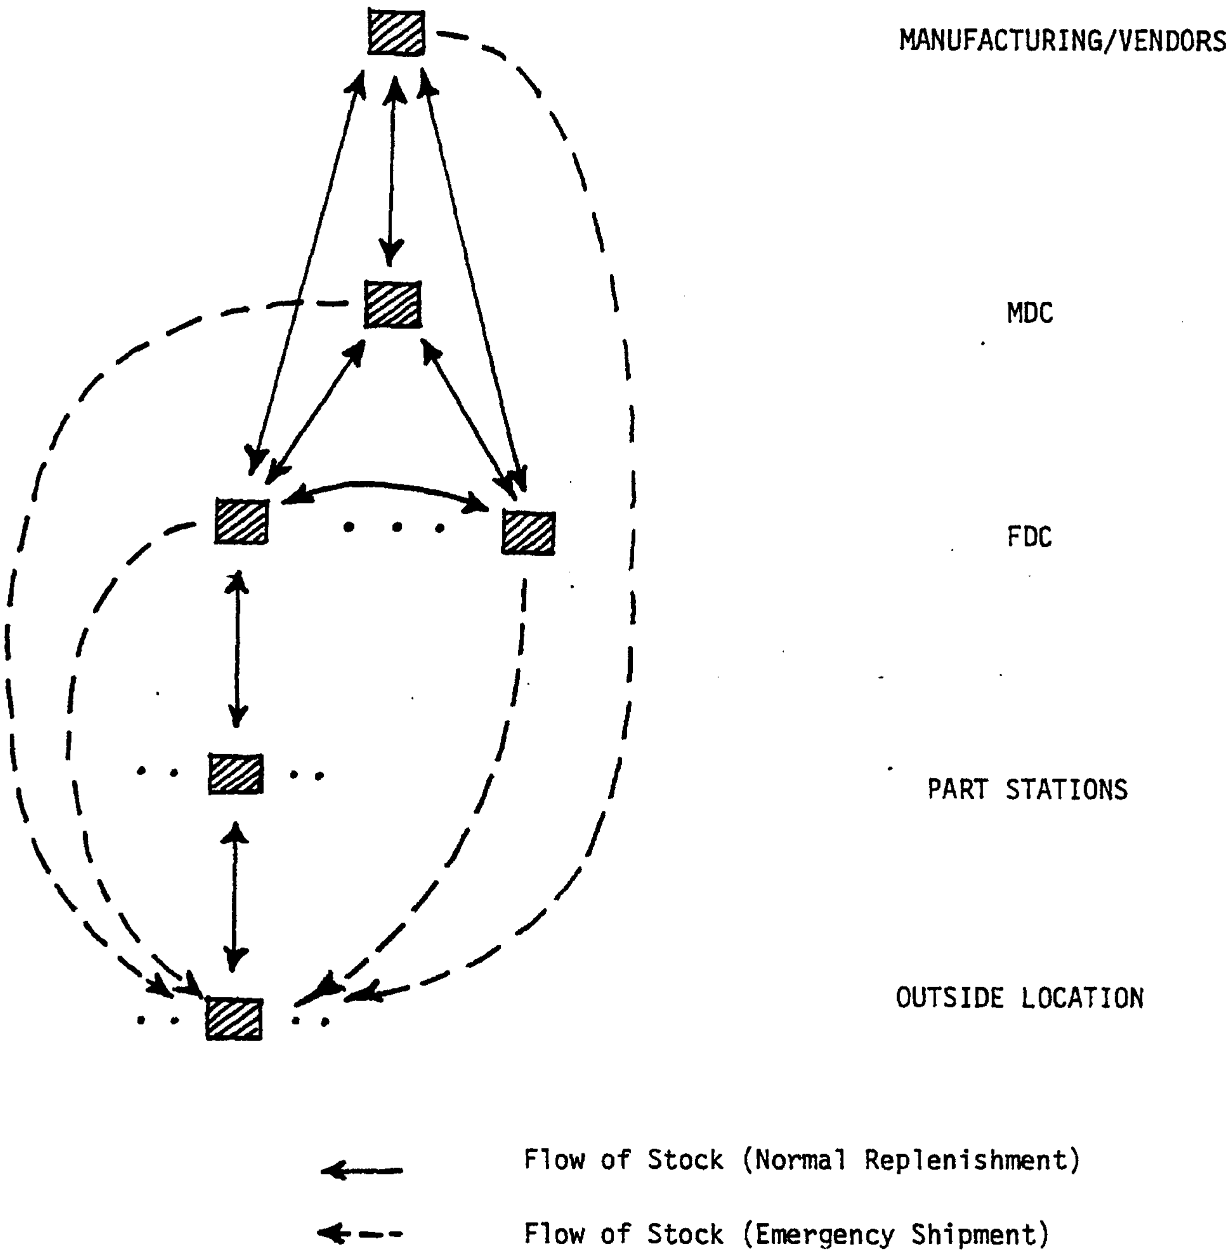

Figure 5b. Inventory Stocking Policy with Flows

2. Fill demand at a level by level basis, beginning at the outside location (OL) and ending at the manufacturer/vendor. If the OL cannot meet demand with pre-positioned stock on hand, a series of expedite shipments are considered which include:

    - Parts Station (PS) to OL
    - Other OL to OL (under Vanning)
    - Distribution Center (DC) of OL's region to OL
    - Other regions' DCs to OL
    - Mechanicsburg (MDC) to OL
    - Manufacturer/Vendor to OL

    The ability of each higher echelon stocking point to meet downstream and laterally generated expedite requests is determined by the pre-positioned stock levels at each point and total demands in the system. In general, a least cost sourcing algorithm is used with allocation rules based on relative demand size.

3. After all demands for the period are met and CE returns are received, ending inventory levels are computed. Based on the expected value of these levels, the shortfall between desired and on-hand levels can be computed. These quantities are the inputs to the pre-positioning step 1.

4. Expected costs for a fixed set of inventory stocking levels are computed by a program which cycles through all possible demand realizations.

5. Optimal stocking levels are determined by searching through the range of possible stock levels at each stocking point so as to minimize expected costs. A five level branch and bound algorithm is used at this step.

6. Given the optimal stock levels for the configuration, an output program computes optimal PAL, response times and cost components.

C. Outside/Inside Control Model

This model considers the relationship between stock levels at a Parts Station and stock levels at a variety of OLs. As demand at each OL_i is realized, the possibility of pooling stock is considered under the "vanning" option in which the following sequence of events is considered:

Vanning:

1. Fill demand at OL_i by stock at that location.

2. Try to fill excess demand from PS.

3. Search stock at all other OLs reporting to the PS before sending an expedite request up to the FDC.

4. Use the FDC, MDC and manufacturer/vendor in the appropriate manner as indicated by the inventory stocking model.

Under the "No Vanning" option the above sequence of events is altered as follows:

No Vanning:

1. Fill demand at OL_i by stock at that location.
2. Try to fill excess demand from PS.
3. Pass total PS excess demand up to FDC, MDC, etc., without using other OL stock.

These two options can be viewed as representing extremes in the pooling of stock at the customer level. The added fixed cost for maintaining a vanned operation (vehicles, dispatchers, telecommunications, etc.) must, of course, be calculated and compared with the expected savings predicted by the models.

5. ANALYSIS AND REPRESENTATIVE RESULTS

In order to explore the nature of optimal stocking and logistics structure policies for the IBM parts maintenance system, a series of analyses were carried out for a representative distribution center and its reporting stocking locations for a number of specific part categories. Our objective at this stage of system development was to 1) validate the models, 2) to develop hypotheses on the nature of the relationship between optimal stocking policy and parts characteristics, and 3) to explore the nature of optimal system structure by identifying the total dollar value of inventory and facility related costs which can be traded off against various measures of service.

In order to carry out a meaningful analysis of a particular distribution center region, it is of course necessary to capture the interactions of that region with the entire parts maintenance network. We defined an aggregated model structure to perform this analysis in a relatively efficient manner. The aggregated structure is used to identify key policy constraints and the nature of optimal stocking policies. In this section the process of analysis is detailed for a specific part category and the Detroit FDC region. We also consider logistics system structure by identifying a number of structural design options. The model and aggregated structure are then used to quantify cost and stocking policy differences for each such option. Finally, we also present some specific results in this section of a facility operating cost analysis exercise.

5.1. Aggregated Model Structure:

Figure 6 is a diagram of the initial aggregated model structure used in the analysis of inventory stocking policy for Detroit and for its reporting Parts Stations and Outside Locations. Detroit was selected for analysis since it is representative of a typical distribution center in terms of size and other operating characteristics which are approximately equal to the national mean value of these attributes. The particular aggregate model structure was chosen for a variety of statistical and computational reasons discussed below. The basic model structure described in section 4 can, of course, capture model structures of arbitrary complexity.

A representative analysis of the Detroit region necessitates characterization of the facilities which both sit above or are parallel to the Detroit FDC as well as those which sit below it in the multi-echelon facility network structure. Thus, as noted in Figure 6, Mechanicsburg is explicitly included in the model structure. All stocking points reporting to Mechanicsburg and not reporting directly to the Detroit FDC were grouped together into the "Other FDC" facility. Since Detroit represents approximately 1/21 of the country's volume, the "Other" facility was scaled to be twenty times the size of Detroit.

It is also necessary to account for the collection of Parts Stations and Outside Locations reporting to Detroit in order to complete our picture of the flow of parts within the system. Due to the extremely low rate of usage for the majority of parts, it is difficult to carry out statistical analysis of parts demand for a

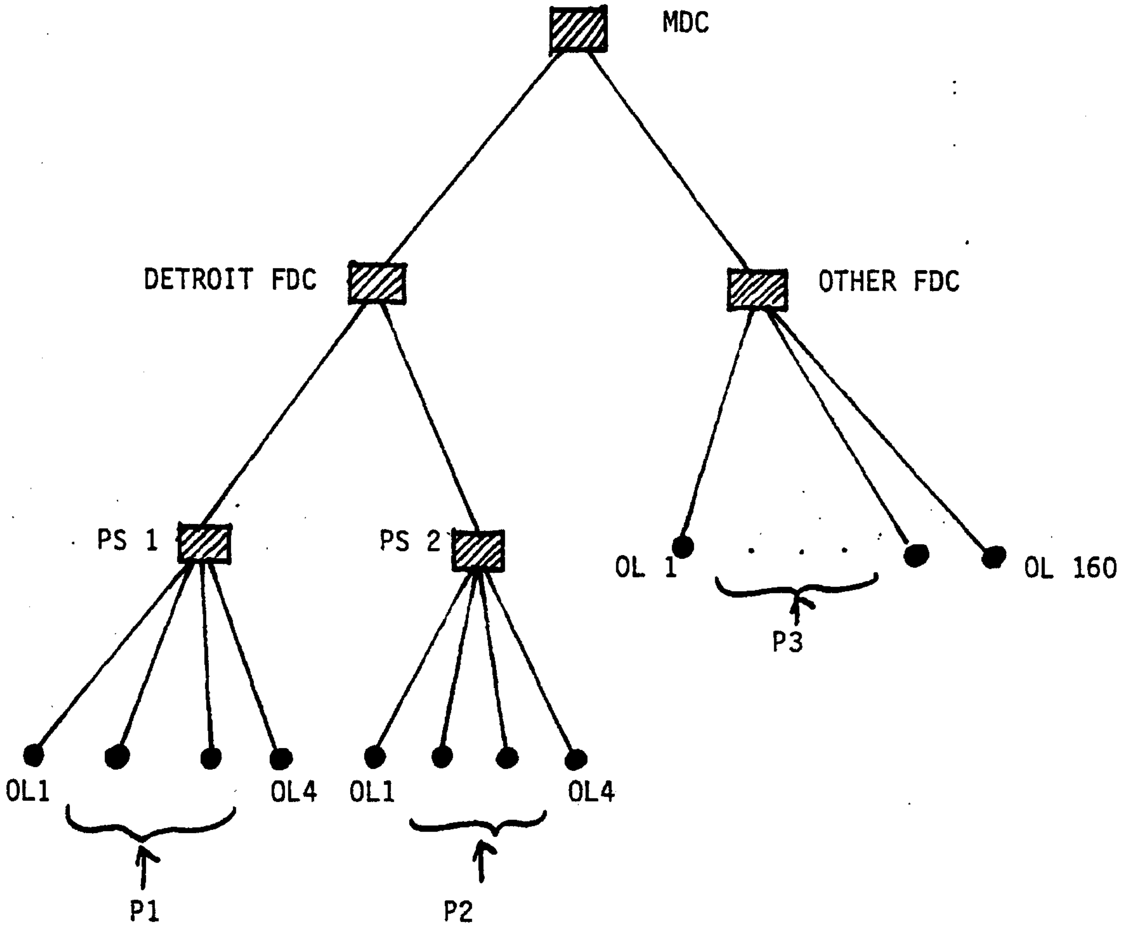

Figure 6. Aggregated Model Structure

particular part category at the OL or even PS level. Yet a realistic picture of the random demand process requires specification of location specific demand probabilities. An aggregated estimation procedure was therefore used to develop estimates of demand probabilities. Further development of this procedure will be carried out in the recommended follow-on study of demand forecasting.

In our aggregate structure, the four Parts Stations reporting to Detroit were treated as one stocking point (PS #1) as were all of the other Branch offices which use Detroit FDC as their Parts Station (PS #2). The Outside Locations were modelled by grouping the set of all OLs reporting to each PS into four OL concentrations, each of which faces an equal probability distribution of demand for a part. As noted above, the model structure can accommodate an arbitrary number of stocking points at each level with unique demand distributions. The aggregation of OLs into four concentrations with equally likely disbursement requirements will of course limit the computational burden of the Branch and Bound optimization algorithm. We note that additional data on machine populations at customer locations could be used in the follow-on forecasting study to develop more representative OL groupings. Since "Other FDC" is an aggregate which is twenty times Detroit in size, we modelled its demand process as being generated by 160 OLs (20 x (4 + 4)). Our subsequent, more extensive, analysis of the Detroit region increased the number of OLs per PS to 8 (see Section 6). It will be shown that the number of OL groupings used in the aggregated structure reflects the maximum number of demands considered in the model time period from all customer locations reporting to a PS.

5.2. An Example of Stocking Policy Analysis:

The analysis in our initial set of runs was restricted to specific parts classes which represented examples of high, medium and low demand rates and high, medium and low unit costs. National monthly demand rate varied from one to 50 (units/month) and unit cost ranged from $2 to $2000. In order to illustrate the workings of the models and the data input analyses required, we consider, in some detail, the case of a single part category with a national monthly average demand rate of 1 unit and a unit cost average of $2 (low demand, low cost). Demand for a typical part in this class was estimated by computing the national average disbursement rate for all parts in the class. This average was computed by taking a time weighted average of the ratio (total disbursements/total part numbers). Data was collected for each Parts Station grouping in the Detroit region. It is necessary to allocate this average disbursement rate over the four OL groupings reporting to each PS. The average disbursement rate was therefore equally divided among the four OLs. It was assumed that the resulting number represented the probability of demand for one part in one period at the OL grouping level. The actual value of demand observed can then be for either zero or one unit at the OL and is distributed as a Bernoulli (0-1) random variable. Maximum demand generated at the PS is therefore four (the number of OL groupings). Demand at all non OL levels (which represent expedite requests) will be computed by appropriate aggregation as indicated by the multi-echelon structure and the stock levels in place. The demand for parts to "Other FDC" was computed by taking the average of the two PS probabilities and assuming that the resultant average disbursement rate is shared by the 160 OLs reporting to Other. These demand probability estimates and process specifications are admittedly crude, but they represented a reasonable approach given initial data availability and the need for preliminary analysis of stocking policies. The resulting probability estimates for the part class in question are indicated in Table 5.

The next set of data inputs considered is transportation costs and times. Both normal and expedite costs are required. Normal transportation costs refer to the cost per part per shipment when the part is a part of a shipment made to replenish stock so as to bring inventory at the destination stocking point up to its required level. Clearly source/destination part specific costs, which incorporate standard modes, traffic control policies, shipment sizing and warehousing procedures are required. Our first attempts at analyzing transportation cost data indicated that there are major problems in the current data base related to these transactions. A uniform charge of $0.35/part/shipment (normal shipments) was estimated and applied to all possible source/destination pairs in the multi-echelon network. Expedite or emergency shipment costs are required by the model for each possible sourcing point (manufacturer, Mechanicsburg, Other FDC, Detroit FDC, own PS, other PS, other OL) under the assumption that all such shipments are made directly to the OL level. These costs include special handling at source and destination, uneconomical shipment quantities and use of faster and more expensive delivery modes. The current system has a wide range of delivery options with varying degrees of urgency. We chose, at this stage of the analysis, to define one emergency mode. Costs were estimated and the actual values

1. Unit Cost: $2
2. Holding Cost: 2% (monthly)
3. Demand Data:

	PS #1	PS #2	Other
a. Probability of Demand at an OL grouping	.00673	.00048	.00361
b. Return Rate to OL grouping	.2	.2	.2

4. Transportation Costs and Times:

Source-Destination	Emergency Mode unit transportation cost ($/part/shipment)	Emergency Mode expedite time (hours)	Normal Mode unit transportation cost ($/part/shipment)
MFG-Detroit OL	5.50	24	0.35
MDC-Detroit OL	6.00	18	0.35
FDC Detroit-Detroit OL	3.50	6	0.35
FDC Other-Detroit OL	4.25	24	0.35
FDC Other-Detroit FDC	4.25	24	0.35
Detroit PS-Detroit OL	4.00	2	0.35
Detroit OL-Detroit OL	4.50	2	0.35
All Others	——	—	0.35

Table 5: Data Inputs for Example

are indicated in Table 5. Due to the large degree of uncertainty associated with these costs, we also considered a second set of "high" expedite costs with values scaled up by a factor of ten.

The model assumes that normal shipment response time falls well within the monthly time period and thus will not affect parts availability. Expedite reponse times are of course critical and they must reflect every possible sourcing point. Based on our discussions with I & D, response times ranging from 2 hours from the local PS to 24 hours from the manufacturer were used. Further study is required to refine these numbers. Other model inputs include an inventory holding cost (2% per month was used) and a return rates of .2.

Runs of the stocking model were carried out for the following cases:

1. Vanning, Low Emergency Cost
2. Vanning, High Emergency Cost
3. No Vanning, Low Emergency Cost
4. No Vanning, High Emergency Cost

Table 6 indicates the various model outputs generated for each of these cases. For each case the optimal stocking level at each location in the aggregated model structure is determined. Thus for this low demand, low value part and the case of Vanning, Low Emergency Cost, it is optimal to stock only one unit in the Detroit region at one of the high demand OL groupings reportins to PS 1. Three other units are to be distributed at locations reporting to the Other FDC. This configuration will give a very high PAL but a relatively poor response time. This difference is due to the manner in which the two measures of service are defined. PAL is defined here to be the probability that a shortage will <u>not</u> occur. Hence the preponderance of cases where demand is zero do not cause a shortage and a high PAL is assured. Expected response time was calculated by computing the weighted average of all response times generated by part demands for only those cases where demand was not zero. While other measures of service can be defined, it is clear that this configuration will lead to excessive delays for a number of customers in those relatively rare instances when a part is actually needed.

Comparison among the four cases indicates some interesting responses of the optimal stocking policy to emergency costs and pooling arrangement changes. The high emergency cost (ten times low) cases stock considerably more units in the OLs in both the Vanning and No Vanning cases. Response time approaches its minimum value (2 hours) and parts availability is essentially 100%. Costs have, of course, increased significantly. In the two No Vanning cases we note that optimal stocking levels move away from the PS/OL levels and towards the FDC (or MFG) levels when compared to the Vanning cases. This increase in centralization is due to the fact that under No Vanning stock at an OL or PS cannot be shared. Rather than replicate stock at each such location, it is optimal to move the stock up into those locations where it can be shared. While the cost implications of No Vanning are not in this case great, the degredation of service (especially response time) is severe.

1. Vanning/Low Emergency Costs

a. Optimal Stocking Levels:

MDC			0	
FDC (Detroit, Other)		0		3
PS (1, 2-Detroit)	0		0	
OL (1, 2-Detroit)	0		0	

b. Expected Costs ($/month):

Holding	.1504
Emergency Shipment	.1193
Normal Shipment	.1680
Total	.4377

c. Parts Availability Levels (PAL):

MDC or below	.9977
FDC or below	.9980
PS and OL	.9980

d. Expected Response Time (hours):

3.564

<u>Table 6(1)</u>

2. Vanning/High Emergency Costs:

a. Optimal Stocking Levels

MDC			0	
FDC (Detroit, Other)		0		4
PS (1, 2-Detroit)	0		0	
OL (1, 2-Detroit)	4		1	

b. Expected Costs ($/month):

Holding	.3503
Emergency Shipment	.0839
Normal Shipment	.1697
Total	.6039

c. Parts Availability Levels (PAL):

MDC or below	.9999
FDC or below	.9999
PS and OL	.9999

d. Expected Response Time (hours):

2.001

Table 6(2)

3. No Vanning/Low Emergency Cost:

a. Optimal Stocking Levels

MDC			0	
FDC (Detroit, Other)		1		3
PS (1, 2-Detroit)	0		0	
OL (1, 2-Detroit)	0		0	

b. Expected Costs ($/month):

Holding	.1504
Emergency Shipment	.1192
Normal Shipment	.1705
Total	.4401

c. Parts Availability (PAL):

MDC or below	.9991
FDC or below	.9998
PS and OL	.9714

d. Expected Response Time (hours):

6.101

Table 6(3)

4. No Vanning/High Emergency Cost

a. Optimal Stocking Levels

MDC			0	
FDC (Detroit, Other)		0		4
PS (1, 2-Detroit)	0		0	
OL (1, 2-Detroit)	4		0	

b. Expected Costs ($/month):

Holding	.3103
Emergency Shipment	.1247
Normal Shipment	.1691
Total	.6041

c. Parts Availability Levels (PAL):

MDC or below	.9999
FDC or below	.9999
PS and OL	.9999

d. Expected Response Time (hours):

2.001

Table 6(4)

5.3. Optimal Stocking Policy

The analysis of the preceding section was repeated for the remaining eight parts classes with the same cost data and structure. These sets of runs yielded optimal stocking quantities for 36 cases (9 parts classes x 2 emergency cost classes x 2 Vanning options). These runs yielded a solution for a system unconstrained by response time requirements and space limitations. The impact of response time constraints will be considered in Section 6.

The general conclusions on the nature of optimal stocking policies can be summarized as follows:

i) It is never optimal to stock a part at all stocking locations.

ii) Every facility will stock some parts.

iii) Recommended inventory levels at a designated stocking point increase as

  - unit cost decreases
  - demand rate increases
  - emergency costs increase

iv) No Vanning (pooling)

  - increases centralization of stock
  - increases total stock
  - increases emergency shipment costs
  - increases response time

v) Centralization of stock increases as

  - unit cost increases
  - demand rate decreases

- emergency costs decrease

Figure 7 summarizes, in part, these conclusions. A line of centralization can be defined to delineate the choice of stocking locations. Below the line of decentralization, optimal stocking policy positions units at central locations (MFG, MDC or FDC). Above the line stock is positioned at more local levels (FDC, PS, OL). The line of centralization will move in the directions indicated in response to changes in transportation costs, Vanning or inclusion of response time constraints. Further analysis of this characterization of stocking policy will be discussed in Section 6.

5.4. Analysis of Logistics Structure

Figure 4 illustrates the hierarchy of models which relates analysis of logistics structure decisions to inventory stocking policy. The general procedure requires specification of alternative configuration for which inventory costs (at optimal stocking quantities) are required. These costs are then traded off against fixed and variable facility costs by a model which will sum over all parts. Seven structural options or configurations were chosen. For each option the inventory stocking policy model was run for a variety of parts. The seven options considered are:

1. MDC only

2. MDC, FDCs only

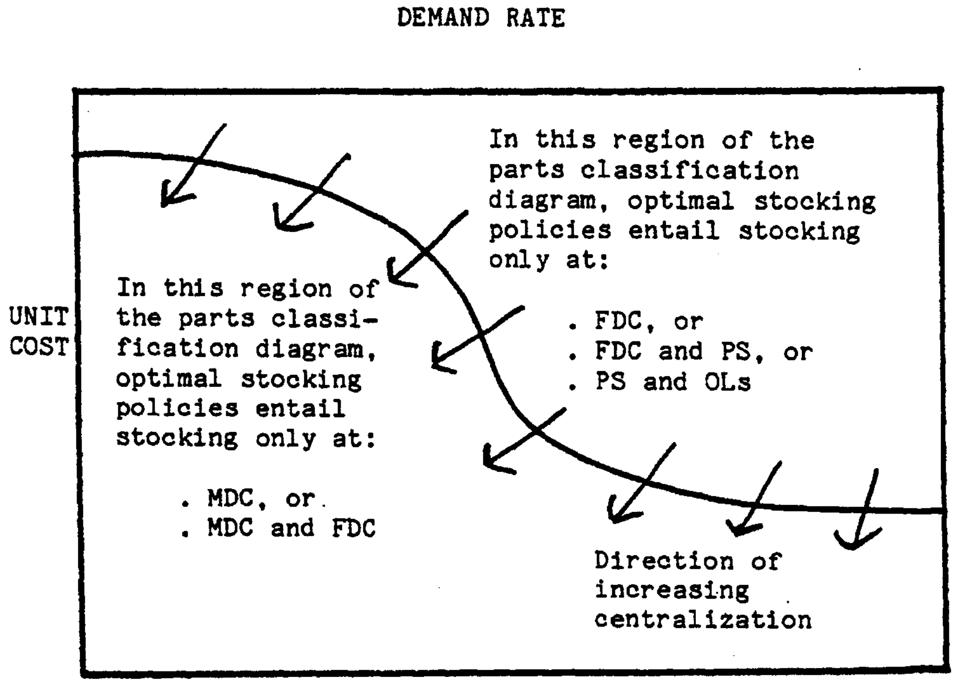

Figure 7: Generic Stocking Policies

Structure*	Transp. Cost	Stocking Levels (MDC+FDC, PS+OL)	Holding Cost	Transp. Cost	% PAL (at PS)
1	Low High	(3, 0) (4, 0)	0.03 0.04	1.60 15.21	98.80 98.80
2	Low High	(4, 0)	0.04	0.11	98.80
3	Low High	(3, 0)	0.05	0.11	99.99
4	Low High	(4, 0) (4, 8)	0.04 0.12	0.12 0.07	98.80 100.00
5	Low High	(3, 0)	0.05	0.11	99.99
6	Low High	(4, 0)	0.04	0.12	98.80
7	Low High	(4, 0)	0.04	0.12	98.80

(Vanning)

Part Class: 1

Demand: 1

Cost: $2

*Structure Options

1 MDC only
2 MDC, FDC only
3 MDC, OL only
4 no MDC
5 no FDC-Detroit
6 no PS-Detroit
7 unconstrained

Table 7: Generic Stocking Policies for Example

3. MDC, OLs only

4. no MDC

5. no FDC-Detroit

6. no PS-Detroit

7. no constraints on structure

Option 7, no constraints on structure, was used to generate the results discussed so far. The remaining six options act to remove certain stocking locations from the system. The results for the seven options for part 1 are summarized in Table 7. These results were run with different transportation costs and not all of the runs were completed. These partial results indicate that for this part the only restrictive case in terms of cost penalties is option 1 which eliminates all FDCs, PSs and OLs. The implication of this output is that within a wide range of possible system structures which make some use of a multi-echelon structure, the cost of optimal stocking policies is relatively insensitive. Conversely, eliminating all stocking points below MDC causes an enormous cost penalty to be incurred. Moreover, the response time performance of this option would be completely unacceptable. These results were replicated with the other part categories considered.

A further issue to consider in the analysis of facility configurations and logistics system structure are the fixed and variable costs for operating the various facilities. Our stocking policy analysis, so far, indicated that 1) some parts are stocked at every level, and 2) within a wide range of multi-echelon structures, total inventory costs are relatively insensitive. It should be noted that the second point may not apply in an analysis where response time constraints are introduced. In any case, before conclusions can be drawn on the potential savings for re-configuring the logistics structure, the costs associated with structure must be estimated.

Data were collected from the annual accounting reporting system for MDC and each of the 21 FDCs. These data indicated such things as transportation costs, facility (occupancy costs) equipment costs and other variable operating costs. Additional data were collected on total average daily transactions, FE transactions, OPD transactions and year end dollar value of FE related inventory for each FDC. A regression analysis was then carried out to explain the variance in total annual variable cost (total minus fixed cost) at each FDC as a function scale of operations as indicated by daily transactions and inventory investment. The results were indicated that a very good (adjusted $R^2 = .94$) fit can be achieved with the following equation.

$$\begin{aligned}\underset{(\$000)}{\text{Total Cost}} &= \text{Fixed Cost} + 2.12\ (\text{FE Transactions (000's/day)})\\ &\qquad\qquad + 0.99\ (\text{OP Transactions (000's/day)})\end{aligned}$$

which indicates the total annual operating cost for an FDC. It is interesting to note that inventory level is not a significant explanatory variable and that the unit cost of an FE transaction is more than double the unit cost of an OPD transaction. More analysis is required to explore the reasons for this rather large difference.

It is interesting to note that the data also indicated that the average annual total operating cost for an FDC was $1.14 million, of which $.47 can be allocated to the sum Fixed Cost categories and the constant term in the regression. Thus a maximum of $9.84 million could be saved by closing down all FDCs.

Additional data on I & D's annual budget indicate a further interesting cost breakdown. The key tradeoff in logistic structure will be facility cost vs holding and transportation costs. Our preliminary analysis of the figures indicates that the fixed component of FDC facilities is on the order of $10 million, holding is $98 million, and transportation is $32 million. It is of key importance to note that the fixed operating costs at Mechanicsburg and the PSs are not known to us at this time. Nonetheless, it is evident that a first priority should be given to the holding/transportation cost tradeoff. For a wide variety of reasonable logistics system structures this tradeoff is essentially determined by optimal stocking policy. Before conclusions can be drawn on the nature of such stocking policy, our results of this section also indicate that the inclusion of response time constraints is of critical importance.

I & D 1981 Cost Analysis	
. Transportation	$32 million
. Other Operating	<u>$46 million</u>
. Total I & D	$78 million
. Inventory Holding ($408 million @ 24%)	$98 million
Total Operating + Inventory	$176 million