*Proceedings of the 1998 Winter Simulation Conference*
*D.J. Medeiros, E.F. Watson, J.S. Carson and M.S. Manivannan, eds.*

# MOBILE ROBOT SIMULATION OF CLINICAL LABORATORY DELIVERIES

Manuel D. Rossetti, Ph.D.

Department of Systems Engineering
University of Virginia
Thornton Hall
Charlottesville, VA 22903 U.S.A.

Amit Kumar

Department of Systems Engineering
University of Virginia
Thornton Hall
Charlottesville, VA 22903 U.S.A.

Robin A. Felder, Ph.D.

Medical Automation Research Center
The University of Virginia
Pathology-Box 168,
Charlottesville, VA 22908 U.S.A.

## ABSTRACT

Flexible automation in the form of mobile robots holds the potential for decreasing operating costs while improving delivery performance in mid-size hospital delivery systems. This paper discusses the use of simulation modeling to analyze the costs, benefits, and performance tradeoffs related to the installation and use of a fleet of mobile robots within mid-size hospital facilities. The results of this study enable a better understanding of the delivery and transportation requirements of mid-sized hospitals and how a fleet of mobile robots can meet those requirements. We show that for clinical laboratory deliveries a fleet of 6 mobile robots can achieve significant performance gains over the current system of 3 human couriers while still remaining cost effective. The 6-robot alternative reduces the annual cost by approximately 56% and improves turn-around time performance by 33%.

## 1 INTRODUCTION

In this study, the University of Virginia's clinical laboratory and pharmacy delivery processes are used to examine the use of mobile robots within a mid-sized hospital facility. Mid-size hospitals use many different transportation modalities to deliver supplies to and from the service units within a hospital. The University of Virginia Hospital employs human couriers, point-to-point pneumatic tubes, tack-mounted carts, and mobile robotics. This multifaceted transportation system provides a variety of delivery options for the medical staff; however, the current system's over reliance on human couriers for deliveries has inherent disadvantages in terms of cost and delivery reliability.

Helpmate Robotics has developed a robotic courier for applications within a hospital environment. The robot is designed to meet a variety of delivery missions in a fully autonomous fashion. The robot is able to make round trip deliveries, one way trips, one-way trips with stops, and rounds with multiple stops. The robot uses a hierarchical control mechanism with a topological map of the hospital embedded into its knowledge base for navigation. Autonomous operation is enabled through the use of multiple sensing modes for including odometer based navigation, sonar, infrared and vision sensors. Additional navigational assistance is also available through the use of specialized reflective tape mounted to the ceiling. A supervisory computer with radio links to the robots is used in multiple robot applications to prevent deadlock around elevators and in hallways. The robots use specialized algorithms in order to navigate and avoid obstacles within crowded hallways. To allow full access to the hospital, elevator and door actuators must be installed. Typical applications include delivering late meal trays, sterile supplies, medications, specimens and medical records. For more information concerning the capabilities of the robot, we refer the interested reader to references Evans (1994) and Evans et al. (1992).

Automatic guided vehicles (AGV) transport material between pre-specified locations in a facility. They may be used as transporters or as a part fixture capable of holding the parts during the processing operation. Fundamentally, an AGV system is specified by:

1. the location of pickup and drop-off points,
2. the path between pickup and drop-off points,
3. the number of vehicles, and
4. the routing and scheduling of vehicles between pickup and drop-off points.

These objectives compete to trade-off cost and system performance in complex ways. Methods that have been used to design and analyze AGV systems include optimization methods (Gaskin and Tanchoco (1987)), heuristic methods (Park, Raman, and Shaw (1989)), simulation methods, and artificial intelligence methods (Thesen and Lei (1986)). Some authors have considered procedures by which the number of AGVs can be determined. Egbelu, (1987) proposed four analytical techniques that can be used to determine the number of AGVs required in a particular setting. For example, the CAN-Q method recommended by Tanchoco, Egbelu, and Tagaboni (1987) helps in determining the starting points for the number of vehicles to be used in a simulation experiment. We refer the interested reader to the references for more information on these topics.

Simulation modeling for automated guided vehicles for industries has been covered to a very large extent within various literatures. Ülgen, and Kedia (1990) use simulation to design a cellular assembly plant employing AGVs. Prasad and Rangaswami (1988) use simulation to analyze the control systems associated with an AGV system in an integrated circuit board manufacturing application. Newton (1985) discusses the use of simulation to determine the appropriate number of AGVs in a manufacturing setting.

This paper first presents an overview of the hospital delivery system under study. We then present a description of the simulation models used to analyze the system. Next, we present the alternatives under consideration, the experimentation process, and the results of the analysis. Finally, we conclude with recommendations and future extensions for this work.

## 2 HOSPITAL DELIVERY SYSTEM

The University of Virginia is a 683-bed facility. Each floor of the hospital is connected to each other and to the basement by two banks of elevators and two stairwells. One elevator bank is located on the West Side of the hospital while the other is located on the East Side. Each bank of elevators consists of two rows of three elevators each. For each elevator bank, one row of three is reserved for visitors and the other row is reserved for hospital personnel. Figure 1 illustrates the typical layout of a floor within the hospital.

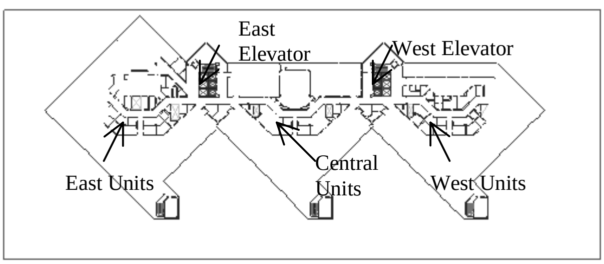


Figure 1: Generic Floor Plan for the 3rd-8th Floors

The clinical laboratory process collects specimens that are placed on floors 3 to 8 from the 29 medical units of the hospital. The clinical laboratory delivery service is divided into STAT and routine deliveries. An activity cycle diagram of this process is presented in Figure 2. For routine pick-ups and deliveries, the courier follows a predefined route. Each courier is assigned two floors: one person for the 3rd and 4th floors, a second person for the 5th and the 6th floors, and a third person for the 7th and 8th floors.

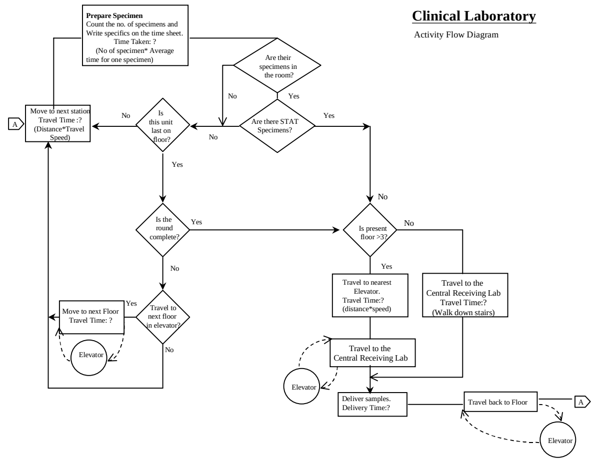


Figure 2: Example Activity Diagram

Couriers wait in the personnel lounge until it is time to start the shift. At the beginning of the shift, couriers make their way to the top floor of their route and visit each unit assigned to their route on their way to the clinical laboratory. If they have picked up items during the route, they deliver the items to the clinical laboratory; otherwise, they repeat their route. During the operation of a shift there are three breaks that are scheduled for couriers. There are 2 breaks of 15 minutes each and 1 break of 30 minutes. If the break occurs, when the courier has items to deliver, the items are first delivered before the break commences.

When a specimen requires STAT delivery, the courier picks up the specimen and then takes the best direct route to the clinical laboratory for delivery. Any items that have already been picked up along the route are also dropped off at the laboratory. The courier then travels back to the unit that was next on the route before they responded to the STAT delivery. The determination of whether or not a specimen is STAT is dependent on the nurses or the doctors and their determination of the patients medical needs. No specific STAT delivery time requirement has been specified by Distribution Services although the response should be as immediate as possible and typically less than 15 minutes.

## 3 MODEL DEVELOPMENT

The major objective of the simulation models was to develop an understanding of the trade-off between cost and system performance, including utilization of vehicles, amount of work in process, system throughput, delivery turn around time, and delivery variability. This was accomplished by modeling the use of fleets of mobile robots in performing delivery services under realistic hospital demand situations.

To analyze the clinical laboratory delivery process, we developed two models. The first model describes the system as it currently operates using three couriers. The second model describes the operation of the system with mobile robots serving as the primary delivery mechanism. The mobile robot model acts essentially the same as the courier model except for minor changes to accommodate the speed of the robots, their dwell at the hospital units, and elevator interactions.

Within the model, entities are used to represent:

- delivery items for clinical laboratory and pharmacy
- control logic for the couriers and mobile robots, and
- logical entities for initializing the model and generating specific arrival processes

We used automated guided vehicle (AGV) movement systems to model the transportation processes associated with both human couriers and mobile robots. For simplicity, let us refer to the transportation device, either human or robot, as a transporter. We create an entity to control the movement of the transporter. This controller entity follows a process that describes the routes used by couriers or robots within the hospital. Within the movement system, we define the possible paths for the delivery mechanisms between each of the hospital units. These paths include movement between units on the various floors and movement between floors using the elevator. A processing station is associated with each hospital unit. A network of links and intersections describe the paths available to the transporters between the stations. A link describes the path between two intersections. Links can be unidirectional, bi-directional, or spur. Spurs enable the modeling of dead-end links. Intersections are associated with each hospital station and with hallways where multiple links intersect. The elevator travel between floors was modeled with additional links and intersections associated with beginning and ending of an elevator trip.

At the beginning of the simulation, we create a controller entity for each transporter and send the controller entity to the appropriate movement system. A transporter within the movement system picks up the controller entity and then follows the defined route for that transporter.

We generate entities representing clinical laboratory delivery items according to a non-stationary Poisson process and send them to the appropriate units for pick-up. Figure 3 illustrates the non-stationary arrival pattern. Figure 4 presents the probability of a demand across the units. Notice that the probability is higher for the intensive care units. The items that arrive to a hospital unit for pick-up then wait in a queue associated with the current hospital unit until a transporter (courier or robot) arrives for transport. When the entity controlling the transporter arrives to a hospital unit, it picks up any waiting items from the hospital unit's queue and delays for any material handling required for the items. A robot arriving to a station will announce itself and request that any items be loaded into its cargo hold. If the robot does not get a response, the robot will dwell at the unit for a pre-specified period of time. Although the dwell can be caused by the lack of a nurse to load the robot, the dwell can also be triggered by the lack of items for delivery. Since nurses are assigned to the pick-up/drop-off stations, we assume some nurse will always be available to load the robot and that the dwell is only invoked if no items are needed.

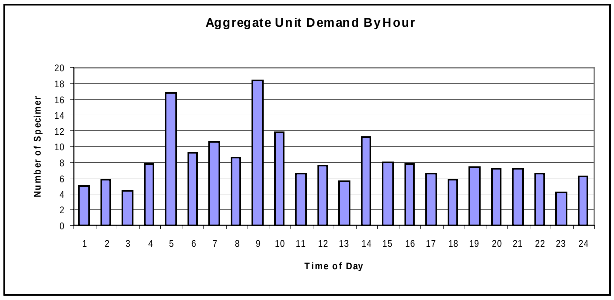


Figure 3: Non-stationary Arrival Pattern

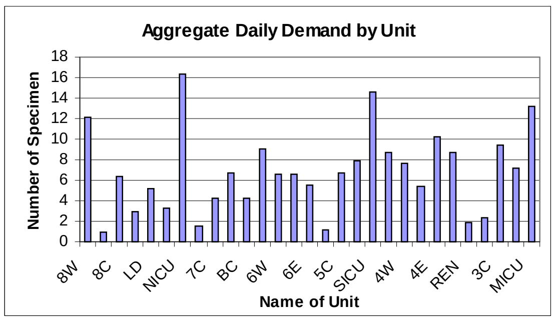


Figure 4: Hospital Unit Probability of Demand

The controller entity then checks to see if any of the picked up items are STAT deliveries. If no STAT deliveries are required, the controller entity continues to the next unit on its route. If any STAT deliveries are present, the controller entity takes the best path to the clinical laboratory for delivery. Based on the demand data supplied by Distribution Services, the probability of a STAT delivery was approximately 25%. At the clinical laboratory, the controller entity delivers the items and then returns via the best path to the next hospital unit on its route. The controlling entity repeats this process at each hospital unit on its route.

The elevators were modeled as resources in the model. Associated with elevator access points on each floor are two stations. An elevator begin-station represents a location where the transporter may seize the elevator before travelling to the next floor. The begin-station also represents a place where the transporter can be positioned to prevent deadlock if another transporter is travelling to the same floor. A courier will experience a delay for the elevator to approximate the resource contention associated with other uses of the elevator. Based on observations of the time until the arrival of an elevator after requesting service, we modeled the elevator delay with a Gamma Distribution with parameters $\beta$ = 0.575 and $\alpha$= 2.47 based on a statistical best fit of the data.

Before traveling to the elevator, the robot will call ahead to the elevator to ensure that no people are using the elevator during its time in the elevator. Because the robot has the ability to call ahead, we assumed that any delay due to other contention for the elevator would be included in the robots travel time to the elevator. An elevator end-station represents the destination of a transporter traveling to a floor via the elevator. This station allows an exit point for the transporter to prevent deadlock and provides a place where the elevator can be released for any other waiting transporters.

When traveling within the hospital, couriers do not block each other's paths. To model this situation using AGV constructs, the couriers were modeled as zero length transporters. This allows passing on the links and mitigates any need for deadlock avoidance or zone control. When travelling, the mobile robots contend for space within the hallways. Two robots should not be permitted to travel down a bi-directional hallway, and a distance of approximately 2 meters should be maintained between the robots. This situation is handled via the use of zone control, properly directed links, and the allocation of waiting zones (such as the elevator begin and end stations). In the mobile robot model, because routes are defined to cover a number of floors, no two robots will ever be on a floor associated with a route at the same time; however, when visiting the clinical laboratory multiple robots may be on the second floor at the same time. Figure 5 illustrates the paths on the second floor to avoid deadlock.

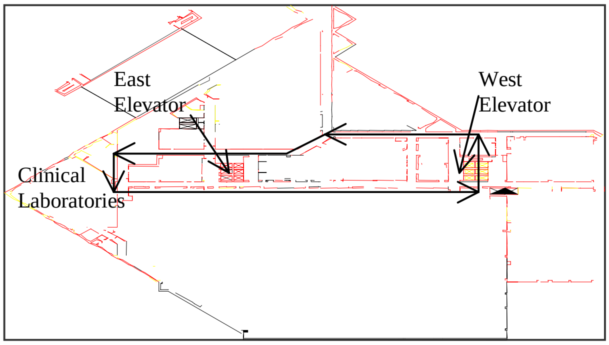


Figure 5: Clinical Laboratory Floor Layout

## 4 MODEL VERIFICATION & VALIDATION

Verification is concerned with building the model right. It is used in the comparison of the conceptual model to the computer representation. The input parameters and logical structure of the model have to be correctly represented. The verification procedure for our model followed a questions/answer format.

1. Has someone other than the developers checked the computerized representation of the model? Yes, Distribution Services checked the computerized representation.

2. Has a flow chart of each logical possible action a system can take when an event occurs been made? Yes, activity flow diagrams for the services were prepared and were verified by a supervisor in Distribution Services.

3. Verify the animation of the computerized model? Yes, the model was shown to a supervisor in the Distribution Services and has been verified.

4. Has the output parameters of the model been verified by someone other than the model developers? Yes, a supervisor in Distribution Services verified the outputs.

For example, we validated the cycle times generated from the model by discussing the outputs with a supervisor in Distribution Services. Cycle time is the time required to complete one round that starts from the first unit and ends at the first unit. Cycle time includes the time taken to drop-off items at the central laboratory. The staff of Distribution Services estimates the cycle time for the clinical laboratory routes to be between 20 and 30 minutes.

To confirm this estimate, we ran our simulation model of the current system under varying demand conditions and estimated the average cycle time for the routes.

Table 1: Average Cycle Time

| | |
|---|---|
| Low Demand | 15.01 |
| Medium Demand | 18.32 |
| High Demand | 20.91 |

The results from the Table 1 indicate that the cycle time is dependent upon the demand rate, but that the values roughly confirm the intuitive analysis of the staff. We also walked the route to confirm the times estimated from the model. The difference can be attributed to the fact that our model does not account for unscheduled breaks taken by the courier.

Validation is concerned with building the right model. It is utilized to determine that a model is an accurate representation of the real system. Validation is usually achieved through an iterative process and by determining the discrepancies and the insight gained to improve the model. Validation is the overall process of comparing the model and its behavior to the real system. For the validation procedure, we are using a widely followed approach formulated by Naylor and Finger (1967). The technique has the following three steps:

1) Build a model which has high face validity: This involves being in constant touch with the model users and others who are knowledgeable about the system. To this end we have been in constant touch with supervisors in Distribution Services who have found the model to be reasonable. We have been in touch with the representatives of Helpmate Inc. to gather more information about their robot and its behavior so as to incorporate that functionality into our model.

2) Validation of Model Assumptions: Modeling assumptions fall into two general classes.

    a) <u>Structural Assumptions</u>: Structural assumptions consist of how the system operates and some simplification and abstraction. An example of this is the waiting time for elevators. We modeled the elevator as a resource rather than model the complicated control logic associated with the timing and movement of elevators between floors.

    b) <u>Data Assumptions</u>: Data assumptions should be based on the collection of reliable data and correct statistical analysis. Data such as the time of collection of specimens were collected by the personnel of the Distribution Services; therefore, we have assumed that the source of the data is reliable. The statistical tests that were performed on the data have been carried out in the input analyzer of Arena.

3) Validating Input-Output transformations: For validating the Input-Output transformations of the model Naylor and Finger (1967) proposed performing sensitivity analysis on the model.

To analyze the sensitivity of the input parameters we performed a $2^k$ factorial experimental analysis on the courier and robot models. We investigated the effect of varying the arrival generation rate, the elevator delay distribution, the STAT/regular specimen distribution and the dwell time of the robots at the hospital units. The factors were varied by +/- 20% from the nominal values. The response variables examined were the turn-around time, the delivery variability, the cycle time, and the utilization. The change in response variables behaved as expected for the selected range of inputs. This indicates that the model is well behaved over the range of these factors and that subsequent analysis can be performed using the nominal levels of the factors. We refer the interested reader to Kumar (1998) for a complete description and analysis of the experiments.

## 5 TRADE-OFF ANALYSIS

Trade-off analysis involves determining multiple criteria for decision making and the formation of a decision function that yields an objective value for the best alternative. In this study, we made a comparison between the existing system with three couriers and mobile robot alternatives using 2, 3, and 6 robots. We developed an objective function to incorporate a metric for each of the competing objectives. The performance metrics of interest were:

1. Cost: Cost is defined as the equivalent annualized cost of the alternatives over a 5 year planning horizon using a 6% discount rate. The alternative with lower cost is considered better.

2. Turn-Around-Time (TAT): Turn around time is defined as the time lapsed between the generation of the specimen and its subsequent delivery to the Clinical Laboratories. Lower turn around time is better.

3. Delivery Variability (DV): Delivery variability is defined as the standard deviation of the turn around time. Delivery variability gives an indication of the consistency of the delivery process.

4. Cycle Time (CT): Cycle time is the time taken by the courier or the robot to complete one round of the assigned route. Cycle time takes into account the time that is spent delivering both STAT and regular specimens. Lower cycle times are better.

5. Utilization (UTIL): Utilization is defined as the ratio of the total time spent by a courier or a robot carrying specimens to the total available time for delivery. Higher utilization is better.

The objective function is an equation, which incorporates each of the Indices of Performances. Each Index of Performance is weighted by the decision-maker, to describe the importance the decision-maker gives to the objective. After evaluating the objective function for each of our alternatives we can then decide which alternative is better. The objective function is defined as follows:

$$IP = \sum_{i=1}^{5} w_i IP_i$$

where $IP_i$ is the Index of Performance for the $i$th objective and $w_i$ is the corresponding weight that the decision-maker attaches to each index of performance. The weights must sum to 1.

Since the indices of performance in the objective function have different units of measure, a linear scaling method was used to convert the observed average values into comparable units of measure. The linear scaling method scales the individual observed values to a scale of 0-100, where 100 is mapped to a high value and 0 corresponds to a low value.

Table 2 summarizes the comparison obtained for the alternative delivery mechanism in the hospital. The table presents the values obtained from the simulation model for each of the robot and courier models averaged across 50 replications under the nominal parameter settings. The standard deviations for those performance measures estimated from the simulation are given in parenthesis.

The salaries paid to the couriers are the primary factor in the courier system. To analyze the costs associated with the deployment of robots, we used the same structure/methodology used by the staff of HelpMate, Inc. For example, the following is a summary of the information obtained for the 3 robot cost analysis.

1) Robot Support Equipment Requirement:
   a) Annunciators 17
   b) Door Sensors 1
   c) Door Sensors and actuators 1

2) Robot Requirement:
   a) Number of Robots 3
   b) Number of Backpacks 3
   c) Number of Radios 2

3) Cost of Robot
   a) Cost of Equipment $301,800
   b) Cost of Installation $ 37,100
   c) Annual Service Contract $24,114
   d) Cost of Courier Service $407,613

Table 2: Summary of Performance Measures

| | Two Robots | Three Robots | Six Robots | Courier |
|---|---|---|---|---|
| COST | $81,110 | $107,605 | $178,027 | $407,614 |
| TAT | 47.28 min (1.97) | 33.54 min (1.07) | 18.9 min. (0.44) | 28.08 min. (2.16) |
| DV | 24.77 min. (1.87) | 16.67 min. (0.82) | 8.63 min. (0.04) | 20.72 min. (2.83) |
| CT | 67.03 min. (2.01) | 42.25 min. (0.87) | 20.72 min. (0.33) | 26.3 min. (1.57) |
| UTIL | 92.50% (0.44) | 91.90% (0.63) | 81.70% (1.52) | 88.33% (0.68) |

The cost of the courier system is based on a loaded hourly rate of $10.26/hr for 24 hours/day and 365 days/year. In order to obtain full yearly coverage over sick days, vacations, etc. 1 person is considered equivalent to 1.4 FTE.

To perform sensitivity analysis on the objective function, the weights associated with each IP must be varied over a range of values. Since cost and turn around time tend to be the most important performance measures, we varied the weights on these responses and fixed the weights for the other performance measures. This also allows for a simpler analysis. Table 3 presents the three weighting schemas used in our analysis. Table 4 presents the average objective function value under the three weighting schemas.

Figure 6 presents the trade-off graphs plotting the objective function values for the 4 alternatives under the three weighting schemas. Under weighting schemas 1 and 2 the robot alternatives clearly dominate the 3-courier system. Weighting schema 1 places the highest importance on cost and low importance on turn around time. Weighting schema 2 places equal weighting between cost and turn around time. In this case, the 6-robot system is the clear winner. As more weight is placed on turn around time the 3-courier system becomes competitive against the 2 and 3 robot alternatives. This is because the 2 and 3 robot alternatives have more difficulty meeting the performance requirements of the system. The indifference weights for the 3-courier system versus the 2 and 3 robot alternatives are:

Table 3: Weighting Schemas

| | Cost $w_1$ | TAT $w_2$ | DV $w_3$ | CT $w_4$ | UTIL $w_5$ |
|---|---|---|---|---|---|
| $\vec{w}_1$ | 0.65 | 0.05 | 0.10 | 0.10 | 0.10 |
| $\vec{w}_2$ | 0.35 | 0.35 | 0.10 | 0.10 | 0.10 |
| $\vec{w}_3$ | 0.05 | 0.65 | 0.10 | 0.10 | 0.10 |

Table 4: Average Objective Function Values

| | 2 Robots | 3 Robots | 6 Robots | 3 Couriers |
|---|---|---|---|---|
| $\vec{w}_1$ | 30.69 (0.83) | 27.72 (0.36) | 31.61 (0.13) | 65.73 (1.10) |
| $\vec{w}_2$ | 49.46 (1.78) | 38.04 (0.86) | 30.39 (0.3) | 55.46 (2.16) |
| $\vec{w}_3$ | 68.24 (2.76) | 48.35 (1.39) | 29.18 (0.51) | 45.19 (3.24) |

Table 5: Indifference Weights

| | Cost $w_1$ | TAT $w_2$ | DV $w_3$ | CT $w_4$ | UTIL $w_5$ |
|---|---|---|---|---|---|
| 2 Robot | 0.2825 | 0.4175 | 0.10 | 0.10 | 0.10 |
| 3 Robot | 0.0925 | 0.6075 | 0.10 | 0.10 | 0.10 |

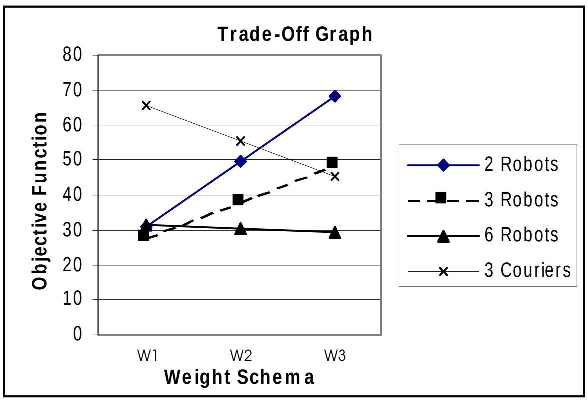


Figure 6 Trade-off Analysis Graph

Thus, cost can be weighed as low as 0.28 before the courier model becomes competitive with the 2-robot alternative. In addition, cost can be weighed as low as 0.09 before the courier system becomes competitive with the 3-robot case.

## 6 SUMMARY

Simulation modeling enabled the entire hospital clinical laboratory delivery system to be realistically modeled so that system performance could be predicted under the alternatives of 2, 3 and 6 robots. From Table 2, it is clear that even though the 2-robot alternative has lower cost it has difficulty matching the performance of the 3-courier model. A one for one replacement of the couriers with robots reduces the cost by roughly 74% with only an approximate 20% increase in turn-around time. The 6-robot alternative dominates the other alternatives by maintaining low cost and significantly improving the turn-around time and the delivery variability. Through simulation, this study was able to clearly demonstrate that fleets of mobile robots can meet the delivery requirements of mid-sized hospitals.

## REFERENCES

Egbelu Pius J., 1987. "The use of non-simulation approaches in estimating vehicle requirements in an automated guided vehicle based transportation system," Elsevier Science Publishers B. V., Amsterdam, pp. 17-32

Evans, J. 1994. "Helpmate: An Autonomous Mobile Robot Courier for Hospitals", Proceedings IROS.

Evans, J. Krishnamurthy, B., Barrows, B., Skewis, T., Lumelsky, V. 1992. "Handling Real-World Motion Planning: A Hospital Transport Robot", IEEE Control Systems, 12, (1), Feb., 15-20.

Gaskins, R. J. and J. M. A. Tanchoco, 1987. "Flow Path Design for Automated Guided Vehicle Systems," *International Journal for Production Research*, 25,5, 667-676.

Kumar, A. 1998. Comparing Delivery Mechanisms in Mid-Sized Hospitals, M. S. Thesis, Department of Systems Engineering, University of Virginia.

Law, A. M., and W. D. Kelton 1991. *Simulation Modeling and Analysis*, 2nd Edition, McGraw-Hill, New York.

Naylor, T. H., and J. M. Finger 1967. "Verification of Computer Simulation Model," *Management Science*, 2, B92-B101

Newton, D. 1985. "Simulation Model Helps Determine How Many Automatic Guided Vehicles Are Needed," *Industrial Engineering*, 17, 2, 68-79.

Park, S. C., Raman, N., and Shaw, M. J., 1989. “Heuristic Learning for Pattern Directed Scheduling in a Flexible Manufacturing System,” In the *Proceedings of the Third ORSA/TIMS Conference on Flexible Manufacturing Systems: Operations Models and Applications*, edited by K.E. Stecke and R. Suri, Elsevier Science Publishers B. V., Amsterdam, 369-376.

Prasad, K. and M. Rangaswami, 1988. “Analysis of Different AGV Control Systems in an Integrated IC Manufacturing Facility Using Computer Simulation,” In the *Proceedings of the 1988 Winter Simulation Conference*, 568-574.

Tanchoco J.M.A., Egbelu P.J. and Taghaboni, F, 1987. “Determination of the total number of vehicles in an AGV-based material transportation system,” Material Flow, Elsevier Science Publishers B. V., Amsterdam, 33-51

Thesen, A. and L. Lei, 1986. “An ‘Expert’ System for Scheduling Robots in a Flexible Electroplating System with Dynamically Changing Workloads," In the *Proceedings of the Second ORSA/TIMS Conference on Flexible Manufacturing Systems: Operations Models and Applications*, edited by K.E. Stecke and R. Suri, Elsevier Science Publishers B. V., Amsterdam, 555-566.

Ülgen, O. M., and P. Kedia, 1990. “Using Simulation in Design of a Cellular Assembly Plant with Automatic Guided Vehicles,” in the *Proceedings of the 1990 Winter Simulation Conference*, O. Balci, R. P. Sadowski, R. E. Nance (eds.), 683-691.

## AUTHOR BIOGRAPHIES

**MANUEL D. ROSSETTI** is an Assistant Professor in the Systems Engineering Department at the University of Virginia. He received his Ph.D. in Industrial and Systems Engineering from The Ohio State University. His research interests include the design, analysis, and optimization of manufacturing, health care, and transportation systems using stochastic modeling, computer simulation, and artificial intelligence techniques. Dr. Rossetti is an Associate Member of the Institute of Industrial Engineers and a member of the IIE OR Division. Dr. Rossetti is also a member of INFORMS and SCS.

**AMIT KUMAR** is completing his M.S. graduate studies in the Department of Systems Engineering at the University of Virginia. He will be working as an Associate Consultant at Cooper’s & Lybrand LLP starting August 1998.

**ROBIN A. FELDER**, Ph.D. is Professor of Pathology at the University of Virginia Health Sciences Center and Associate Director of Clinical Chemistry and Toxicology. He serves as the Director of the Medical Automation Research Center at the University of Virginia. Dr. Felder received a B.S. degree in chemistry from the College of William and Mary and a Ph.D. in biochemistry from Georgetown University.