Hospital Delivery System Comparison Via Computer Simulation

Manuel D. Rossetti, Ph.D.
Department of Industrial Engineering
Room 4207 Bell Engineering Center
University of Arkansas
Fayetteville, AR 72701, U.S.A.

Francesco Selandari
University of Milan
Via Giulio Uberti, 12
20129 Milano, ITALY

Abstract

This paper focuses on clinical laboratory and pharmacy deliveries within middle to large size hospitals in order to evaluate whether or not a fleet of mobile robots can replace a traditional, human-based delivery system. Two computer simulation models were developed and the alternative solutions were compared in terms of technical and economic performance. The final results of this research enable a better understanding of the delivery and transportation requirements of middle to large size hospitals and how a fleet of mobile robots can meet these requirements. Our results are based on a case study of the University of Virginia Health Sciences Center and the use of HelpMate® mobile robots. Our methodology and final conclusions should have applicability to other middle to large size hospitals with similar delivery requirements.

Keywords
Simulation, health care, mobile robots

1. Introduction

The University of Virginia Medical Center is a 591-bed facility, with an average daily census of 454. In the fiscal year between July 1, 1996 and June 30, 1997 the system cared for 29,189 inpatients, including newborns. Outpatient and emergency visits totaled 514,307 for the year. The University of Virginia Hospital is an eight-floor complex plus an additional underground floor. It has roughly the shape of a W and the floors are connected to each other by two banks of elevators located respectively in the East and in the West side of the building. Each elevator bank consists of six elevators in two rows of three. One-half of each elevator bank is intended for staff use while the other half is intended solely for the use of visitors and patients. Currently, human couriers perform delivery tasks. Three couriers per shift are involved in specimen delivery, while two are used in pharmacy delivery. Three shifts of eight hours each ensure 24-hour full time service.

HelpMate® Robotics Inc. (HRI), in Danbury, Connecticut, USA, has developed a mobile robot of commercial success in the emerging field of service robots. HelpMate® is a fully autonomous robot capable of carrying out delivery missions between hospital departments and nursing stations. HelpMate® robots use a specific world model for both mission planning and local navigation. The world is represented as a network of links (hallways) and an elemental move for the robot is navigating in a single hallway, avoiding people and other obstacles. In situations where more than one robot is present, a system supervisor properly spaces the robots along the hallways, since they compete for space and for the elevators. HelpMate® robots have two main elements for the human interface: a screen and a keypad on the control console and voice output. We refer the interested reader to [2] and [3] or to www.helpmaterobotics.com for further information concerning the HelpMate® Robot.

This paper first presents a description of the simulation models used to assess the technical performance measures in the alternative systems. Next, the simulation output is discussed and the differences in terms of cost are highlighted. The paper concludes with recommendations and potential areas for further study.

2. Simulation Model

Two simulation models were developed: one for the robotic, the other for the human-based delivery system. The language utilized was MODSIM III, by CACI Products Company, La Jolla, CA, USA. MODSIM III, see reference [4] is a general-purpose, modular, block-structured, high-level programming language, which provides direct support for object-oriented programming and discrete-event simulation; therefore, we describe the features of the main objects interacting in the system.

2.1 Model Building

Three main objects compose the delivery system simulation models: hospital, robot and elevator. The *Hospital* is modeled as a set of floors. Each floor consists of a network of links and nodes. Links and nodes have a capacity, limiting the maximum number of robots allowed to occupy them at any time. A careful choice of link and node capacity makes it possible to avoid deadlocks in floors where more than one robot can travel simultaneously; moreover, it enables the modeling of a space cushion between two consecutive robots in the same hallway. The length of the space cushion can be set by choosing the appropriate minimum length for each link. Hospital units, clinical laboratories and the pharmacy warehouse, if present, are associated, through a reference variable, to a hospital node.

An identification number marks each robot and each robot is associated with a route. A route is completely specified by the sequence of the nodes and floors that the robot must visit. From a generic position, the robot determines the next node of destination by checking the data contained in its route. Robots are allowed to proceed if both the next node and link are available. In this way, if nodes and links have unitary capacity and if hallways are one-way, deadlock is avoided.

Robots use the elevators to move from one floor to another. One elevator in each bank is dedicated to the robots. No other users are allowed to travel in the elevator if occupied by a robot: when a request for the elevator from a robot arrives, all passengers are unloaded before the elevator becomes available for the robot. If more than one robot is requesting the elevator, the robots are served according to FIFO logic, regardless of their current floor. When a robot approaches a hospital unit, it drops off the pharmacy items and it loads lab specimens. Loading specimens implies a delay on a per item basis (from 30 to 60 seconds) since information must be recorded for each specimen. If no specimens are present, then the robot dwells in the unit for one minute. In the clinical lab each robot drops off all the specimens collected, while in the pharmacy warehouse it loads the pharmacy items requested by the hospital units along its route; if no requests are present, then the robot dwells in the pharmacy warehouse for one minute. Pharmacy item loading and unloading operations require a variable amount of time from 25 to 35 seconds, since the nurse or employee assisting the robot is not required to perform any additional recording tasks but only to move a tray or a bag.

Only two relevant differences distinguish *Human Couriers* from robots: human couriers do not compete for space while they are walking along the hospital hallways and during an eight-hour shift each courier is allowed to take three breaks: two breaks are short (15 minutes), one break is long (30 minutes). Human couriers use the elevators as any other human passenger since they do not have a special priority. All the other human courier behaviors are modeled the same as robots'.

Elevators are modeled as resources. If no users are requesting the elevators, the elevator dwells on a floor that can be specified by the modeler (home base). The arrival of a user, be it a passenger or a robot, activates the elevator. When the elevator is travelling upwards, it updates its destination as it reaches a new floor. The destination floor is determined as the highest among the current floor, the highest destination floor of the current passengers or the highest floor where another passenger requested an elevator. After reaching the top of its trip, the elevator starts travelling downwards. Similarly to the previous case, the lowest floor of its trip is updated at each new floor. It is the lowest among the current floor, the lowest destination floor of the current passengers or the lowest floor where another passenger is requesting the elevator. The model accurately predicts elevator travel time. The time contribution due to acceleration and deceleration and to door opening, which applies only if the elevator stops on a floor, was separated from the contribution due to the actual travel time. Moreover, an additional delay is introduced to model person or robot boarding. The model is able to account optionally for robot and elevator *availability*. When a robot or the elevator designated for robot transportation has a breakdown, human couriers intervene. As soon as the failure has been corrected, the robots start to operate again.

2.2 Performance Measures

In the comparison of the systems, the effect of the robot fleet on elevator performance was carefully analyzed. The elevator represents a potential bottleneck in the system because visitors and hospital employees currently heavily use the elevators. The technical comparison of the systems is based on the following time-based performance measures:

- *Specimen Turn Around Time* is defined as the amount of time elapsing from specimen creation, coinciding with its depositing in a hospital unit buffer, and specimen delivery at the clinical laboratory.
- *Pharmacy Item Turn Around Time* is defined as the amount of time elapsing from transmission of the request to the pharmacy warehouse and pharmacy item delivery at the hospital unit.
- *Robot (Human Courier) Cycle Time* is defined as the amount of time that a robot (human) takes to entirely complete its route once.
- *Visitor Elevator Wait Time* is defined as the amount of time a visitor spends in queue for the visitor elevators.
- *Staff Elevator Wait Time* is defined as the amount of time a hospital employee spends in queue for the staff elevators.
- *Robot (Human Courier) Elevator Wait Time* is defined as the amount of time a robot (human courier) spends in queue for the elevator.

2.3 Data Collection

The simulation model required a large amount of very inhomogeneous data that were obtained from a variety of sources. Distances and layout data were determined through hospital blueprints and actual measurements. Robot technical performance was discussed with HelpMate® engineers and designers. Robot speed modeling deserves a particular mention: it was modeled as a triangular random variable to account for possible interference with humans in the hospital hallways (if a robot is blocked it automatically stops) and for the different robot speed in curves and in rectilinear tracts. Let S_{min}, S_{mod}, S_{max} be respectively the minimum, mode, and maximum value that robot speed S can assume. A new value of S is assigned every time the robot covers a new link, according to $S = k * Triangular(S_{min}, S_{mod}, S_{max})$. For the purposes of this paper, k will assume the value of 1. The values chosen for S_{min}, S_{mod}, S_{max} are: 0.274 m/s, 0.508 m/s, 0.63 m/s based on vendor data.

All the data concerning the elevator were assessed through actual observation, with the exception of mean time to failure and mean time to repair, which were provided by OTIS Elevator. The travel time between floors to another was assessed as 3 seconds. The acceleration/deceleration time and door opening took approximately 6 seconds, and an extra delay of 1 second was added for each person boarding the elevator. The time required by a robot to board the elevator was assessed as 35 seconds. The mean time to failure and the mean time to repair were set respectively to 1500 hours and 5 hours and modeled with a negative exponential distribution.

Data collected from direct observations were used to model delivery item arrival rates. For the whole hospital, the specimen arrival rates ranged from 12 to 55 items per hour, while pharmacy requests ranged from 6 to 46 per hour. The higher rates were observed in the early hours of the morning, from 7 a.m. to 11 a.m. The hourly rate pattern is available from the authors upon request. A non-stationary Poisson process was used to model these arrival processes with the arrival rate dependent on the hour of the day.

Person arrival rates were assessed through actual observations. From 10 to 180 elevator requests in one hour were observed. The higher rates are observed approximately from 7 a.m. to 5 p.m., while the lower approximately from 9 p.m. to 7 a.m., as reported in Table . Since people usually arrive to the elevator banks in small groups, we modeled the elevator system using batch arrivals. Each group has from one to four persons. Table 1 reports the probability associated with each group size.

Table 1 - Probability of each size of visitor or staff groups.

Size of groups	Probability (Visitors)	Probability (Staff)
1	0.7	0.8
2	0.2	0.1
3	0.07	0.07
4	0.03	0.03

Human courier speed was modeled in the same way as robot speed to account for courier unscheduled breaks or fatigue. The appropriate values of S_{min}, S_{mod}, S_{max} for the human couriers are: 0.381 m/s, 0.762 m/s, 0.875 m/s based on standardized data and observation of the human couriers. The model data and the model were verified and validated using standard statistical simulation techniques discussed in Banks, Carson, and Nelson [1].

3. Output Analysis

Five factors were recognized in the system under consideration: speed of robots or couriers, person arrival rates, delivery item arrival rates, elevator availability, and robot availability. A 2^k factorial design was performed on the above factors, but due to space limitation, we report only the results for the nominal case. The nominal case consists of the speed factor setting of k=1, person and delivery arrival rates at observed mean values, and elevator or robot breakdowns possible. Table 2 provides a brief summary of the performance measures of the system in the nominal case. The standard deviation is reported in brackets. The data are in minutes. A steady-state simulation experiment was run for the robot and human courier cases. A warm-up period of 1 day, during which no statistics were recorded, was largely sufficient to account for the initial transient period. Statistics were collected on a daily basis for 64 days. We ensured through a statistical analysis that the observations were not affected by autocorrelation. Table 7 explains the acronyms used in reporting the results.

The delivery performance measures are almost equivalent in terms of the averages. Specimen turn around time is better in the human-based system, while the opposite is true for pharmacy item turn around time, although the difference in this case is less than 5 percent. On the other hand, the robotic system guarantees significantly lower delivery performance variability. We analyzed three delivery performance measures ($j = 1,2,3$)

1. Specimen Turn Around Time.
2. Pharmacy Item Turn Around Time.
3. Cycle Time.

Let $s_j^{COURIER}$ and s_j^{ROBOT} be the expected value of the standard deviation of the j^th delivery performance measure, respectively in the human-based delivery system and in the case of robot introduction. Let Δ_j be their difference and let m_{Δ_j} be the expected value of Δ_j. We performed a Paired t Test using the observations of the delivery performance measures collected in the nominal case. For the hypotheses, $H_0 : m_{\Delta_j} > 2$ minutes versus $H_1 : m_{\Delta_j} \leq 2$ minutes, we failed to reject H_0 at the a = 95% level. We felt that at a minimum a 2-minute difference in the delivery variability would be perceptible by staff. The main reason for the difference between the standard deviations is the absence of breaks in robot schedules. Since the standard deviation is generally considered more important than the average for delivery performance measures by health care professionals and lab staff, the robotic system seems to be preferable as far as delivery performance is concerned.

The effects of robot introduction on the elevator system in three different time ranges, defined in Table 3, were analyzed. The theoretical distributions best fit to the simulation output were determined in each time range during the day. The results are summarized in Table 4.

Table 2. System response in the nominal case.

	ROBOTIC SYSTEM	HUMAN-BASED SYSTEM
STATm	26.03 (0.613)	25.133 (0.715)
STATsd	9.7035 (0.273)	12.115 (0.825)
PITATm	24.373 (0.822)	32.115 (0.637)
PITATsd	9.632 (0.2442)	13.966 (0.539)
CTm	27.802 (0.218)	26.323 (0.461)
CTsd	2.4728 (0.1816)	9.7632 (0.5775)
VEWTm	1.1079 (0.032)	1.1088 (0.029)
VEWTsd	0.91964 (0.027)	0.91927 (0.025)
SEWTm	1.0989 (0.0814)	0.88225 (0.0298)
SEWTsd	0.9167 (0.07425)	0.75677 (0.0291)
R/C EWTm	0.9118 (0.0549)	0.74315 (0.028)
R/C EWTsd	0.8202 (0.0967)	0.6558 (0.03)

In the robotic system, an increment of elevator waiting time for hospital employees of about 10 seconds is observed. Let X and Y represent random variables from the distribution associated with hospital staff waiting time, respectively in the case of robot introduction and human-based delivery system. Let Δ be their difference and let m_Δ be the expected value of Δ. We performed the hypothesis test, $H_0 : m_\Delta > 10$ seconds versus $H_1 : m_\Delta \leq 10$ seconds, with waiting time observations collected during peak hours. We failed to reject H_0 at the a = 95% level. Thus, we are 95% confident that the robot system increases the mean elevator waiting time by more than 10 seconds.

Table 3. Periods of the day

NIGHTTIME	12 - 7 a.m. and from 9 p.m. to 12 a.m.
PEAK HOURS	7 a.m. to 5 p.m.
OFF PEAK HOURS	5 p.m. to 9 p.m.

Table 4. Elevator waiting time best-fit analytical distributions

	DISTRIBUTION	MEAN	VARIANCE
Human based system Nighttime	LOGNORMAL	0.572	0.437
Automatic system Nighttime	LOGNORMAL	0.724	0.902
Human based system Peak hours	BETA	1.106	0.674
Automatic system Peak hours	BETA	1.321	0.935
Human based system Off peak hours	LOGNORMAL	0.64	0.56
Automatic system Off peak hours	GAMMA	0.94	0.665

Robot introduction does affect the parameters and the shape of the distribution of the waiting time. The Lognormal distribution is a good model for the waiting time of an elevator user in the night. In a generic elevator bank, few users are waiting at the same time and it is likely that each user can board the first arriving elevator. Therefore, the distribution has an initial peak roughly in correspondence of the expected value and presents a moderate asymmetric dispersion. The Beta distribution instead models well the daytime and the peak hours. Since many people can be waiting at the same time and the elevators have a bounded capacity, they may not be able to catch the first arriving elevator. Therefore, because of the wider variety of relatively high waiting times, the distribution is well described by the typical shape of the Beta.

In the intermediate level of elevator use characterizing *off-peak hours*, the system has some of the characteristics of *nighttime* and others of *peak hours*. Not very many users should be waiting at the same time, but because of the interaction between the robots' presence and the higher person arrival rates (with respect to *nighttime*); it is likely that more users could miss the first elevator. Therefore, the waiting time distribution presents a wider dispersion and a higher probability of relatively high waiting time. Robot introduction tends to move the distribution from a Lognormal to a Gamma, having intermediate characteristics, in terms of shape, between a Lognormal and a Beta.

We complete the analysis of robot impact on the elevator system by determining the probability of worse case scenarios with the theoretical best-fitted distributions reported in Table 5. The table shows an increment of the probability of elevator waiting greater than three minutes of only 4%. The probability of waiting more than one minute instead presents a more significant increment, reaching 7.8% in the peak hours and 18.6% in the off-peak hours. Furthermore, the high probability of elevator waiting time greater than one minute indicates that the elevator system in the UVA-HSC is inadequate, although robot introduction does not worsen significantly the situation (OTIS Elevator recommends that the probability of elevator waiting time greater than 1 minute be lower than 10%). The probability of elevator waiting time greater of 30 seconds instead presents a significant difference only in off-peak hours (21.7%).

Table 5. Waiting Time Tail Probabilities

Period	**System**	**P[SEWT>½m]**	**P[SEWT>1m]**	**P[SEWT>3m]**
Nighttime	**Human**	0.3766	0.1429	0.0119
Nighttime	**Robotic**	0.4482	0.2053	0.0273
Peak	**Human**	0.7172	0.4693	0.0277
Peak	**Robotic**	0.7648	0.5479	0.0689
Off-Peak	**Human**	0.4214	0.1723	0.0166
Off-Peak	**Robotic**	0.6388	0.3584	0.0277

4. Cost Analysis

The human-based delivery system requires the full time presence of five human couriers three shifts a day, 365 days a year. The global number of couriers, N_{FTE} is obtained by multiplying the daily requirement for a full time equivalent factor, f_{FTE}, to account for scheduled holidays as well as unpredictable absenteeism. The University of Virginia Medical Center assesses f_{FTE} as 1.4, therefore, $N_{FTE} = 21$. Since the average wage of a courier including benefits is $21,344.19 per year, the actual cost of the human based delivery system, is currently $448,228.

4.1 Robotic Delivery System Cost

The cost associated with the robotic delivery system (six robots) can be divided into three main categories: cost of equipment (robots, spare batteries, door sensors, etc.), cost of installation, and cost of operation. The cost of the equipment is assessed as approximately $567,800, according to current charges applied by HelpMate® Robotics. The installation cost is associated with the installation of the elevator interface, annunciators, door sensors and Ethernet radios and is estimated as $44,700. The cost of operation is composed of a service contract, human courier wages, and energy costs. The service Contract, offered by HelpMate® Robotics is equivalent to an annual cost of 8% based on the initial equipment cost. We made the conservative assumption that not all of the twenty-one human couriers can be eliminated after automation introduction. One courier was left in each shift, to account for interventions in case of robot or elevator breakdowns or minor maintenance services. In consideration of the full time equivalent factor, five couriers are considered still employed yearly in the delivery system, even after automation introduction.

Energy cost for six robots working 24 hours per day was assessed as $2,400 yearly. The total costs of operation and maintenance were assessed as approximately $154,548.

4.2 Evaluation of the Investment

Table 6 summarizes the most relevant economical performance measures associated with the automation introduction, with respect to the current costs of the human based delivery system. We determined the Net Present Value (NPV), the Internal Rate of Return (IRR) and the Payback Time (t_{PB} is based on non-discounted cash flows, while t^*_{PB} is obtained using the discounted values; the discount rate used is 8%). The final value of the equipment was considered as 30% of the initial value, while a five-year period was chosen as the useful life of the investment, since those values are generally used in HelpMate® Robotics cost reports. Finally, the possibility of a gradual introduction of the automation was investigated. Table 6 shows that under the assumption of introducing three robots in the first year, then one in each of the following three years, assuming a useful life of the investment of six years, the NPV is still greater than zero.

Table 6. Summary Of Economic Performance Measures

Complete Initial Automation Introduction				
Discount Rate		6%	8%	10%
NPV		$751,888	$676,022	$606,557
Annual Cost	**Robot**	$269,732	$278,914	$288,220
	Couriers	$488,228		
IRR		42%		
t_{PB}		2 years and 1 month		
t^*_{PB}		2 years and 5 months		
Gradual Automation Introduction				
Discount Rate		6%	8%	10%
NPV		$676,359	$594,347	$520,744
Annual Cost	**Robot**	$350,682	$367,361	$382,328
	Couriers	$488,228		
IRR		37%		
t_{PB}		3 years and 2 months		
t^*_{PB}		3 years and 6 months		

5. Conclusions

Two possible implementations of a hospital delivery system, a fleet of mobile robots and human couriers, were analyzed. From the technical standpoint, three major conclusions can be drawn. The introduction the automation enables a substantial reduction of the variability of the delivery performance measures. The effect on the average, instead, is not always in favor of the robotic system, although the differences are very small. The fleet of robots does have an impact on the elevator system, in that the mean elevator waiting time increases by about 10 seconds.

From the economical standpoint it was shown that the robotic couriers enable a savings of over $600,000 in five years. Therefore, we conclude that a fleet of mobile robots can meet the requirements of specimen and pharmacy item delivery and transportation in middle to large size hospitals and it represents a valid alternative to human-couriers. Possible areas for further study include an analysis of robot interaction with nursing staff and an analysis of a robotic system used in conjunction with a pneumatic tube system.

Table 7. Acronyms used in the paper.

Acronym	Explanation
STATm	Specimen Turn Around Time Mean
STATsd	Specimen Turn Around Time Standard Deviation
PITATm	Pharmacy Item Turn Around Time Mean
PITATsd	Pharmacy Item Turn Around Time Standard Deviation
CTm	Cycle Time Mean
CTsd	Cycle Time Standard Deviation
VEWTm	Visitor Elevator Waiting Time Mean
VEWTsd	Visitor Elevator Waiting Time Standard Deviation
SEWTm	Staff Elevator Waiting Time Mean
SEWTsd	Staff Elevator Waiting Time Standard Deviation
R/C EWTm	Robot (Courier) Elevator Waiting Time Mean
R/C EWTsd	Robot (Courier) Elevator Waiting Time Standard Deviation

References

1. Banks, J., Carson J.S. II, & Nelson, B.L. (1996). Discrete-Event System Simulation. Englewood Cliffs, NJ: Prentice-Hall International.
2. Evans, J., (1994). “Helpmate: An Autonomous Mobile Robot Courier for Hospitals”, Proceedings IROS,
3. 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.
4. Mullarney, A., West, J., Belanger, R., & Rice, S. (1997). ModSim Tutorial. La Jolla, CA: CACI Products Company.