Proceedings of the 2005 Winter Simulation Conference
M. E. Kuhl, N. M. Steiger, F. B. Armstrong, and J. A. Joines, eds.

QUEUING MODELS OF VEHICLE-BASED AUTOMATED MATERIAL HANDLING SYSTEMS IN SEMICONDUCTOR FABS

Dima Nazzal
Leon F. McGinnis

School of Industrial & Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332, U.S.A.

ABSTRACT

This research explores analytical models useful in the design of vehicle-based Automated Material Handling Systems (AMHS) to support semiconductor manufacturing. The objective is to correctly estimate the throughput and move request delay. The analysis approach is based on queuing network models, but taking into account details of the operation of the AMHS. We analyze the vehicles movement in the system using a Markov chain. This analysis provides the essential parameters such as the blocking probabilities in order to estimate the performance measures. A numerical example is provided to demonstrate and validate the queuing model.

1 INTRODUCTION

Automated Material Handling Systems (AMHS) play a central role in today's 300mm wafer fabrication facilities (fabs). Material handling operations are becoming more complex, with strong demand on these systems to support the production system and function efficiently and robustly.

A typical 300mm fab has a spine layout. This layout has a central aisle with bays branching on both sides. Production equipment is located in the bays. Wafers travel in a lot carried by a Front Opening Unified Pod (FOUP). A lot moving from a tool in one bay to a tool in a different bay must travel through the main aisle, and therefore a transportation system is installed along the main aisle to accommodate this interbay traffic. Similarly, a transportation system is installed within each bay to facilitate the intrabay traffic.

Storage units in the 300mm fabs are known as stockers. A stocker is an Automated Storage/Retrieval System for lot exchange with the transport system. Stockers are usually located at the head of each bay, and they usually serve two purposes: temporary storage for Work-in-Process (WIP) lots, and transfer mechanism between the bay transport system (intrabay system) and the main aisle transport system (interbay system).

The AMHS in a typical 300mm fab is an overhead transport system (OHT) that consists of a collection of track segments forming a loop, Overhead Hoist Vehicles (OHV), and input/output load ports. The input/output load ports are the interfaces where carriers are picked up or delivered from and to production equipment, as well as to storage stockers. Vehicles are suspended from ceiling-mounted rail mechanisms and are capable of delivering to/retrieving from stocker ports and process tools from directly overhead. An intrabay system is illustrated in Figure 1.

Because of the space restrictions in the 300mm wafer fab, OHV travel is on a unidirectional closed loop without the ability to pass each other even when a vehicle stops to drop-off/pick-up a lot from the input/output buffer of a processor or the stocker.

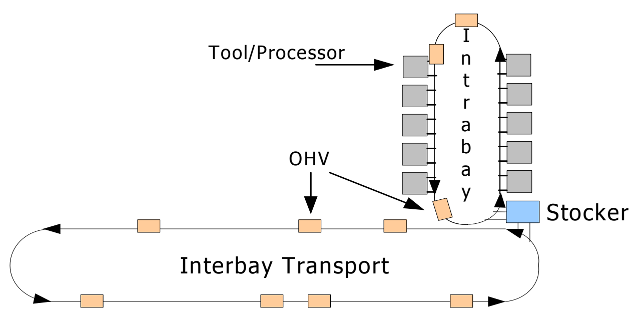

Figure 1: Closed-Loop Unidirectional Interbay and Intrabay AMHSs

Fundamentally, the role of the AMHS is to serve the production system, their interaction occurs at the transfer points in the facility. Essentially, the AMHS should be able to handle the move requests generated by the production system. Once this basic level of service is achieved, AMHSs are distinguished from each other through their cost and performance. Tradeoffs are usually between these two measures. AMHS Performance is usually measured in terms of its throughput (number of moves per unit time) and response time; a MHS that responds faster to the move requests is generally more preferable. We say "generally" because it might be the case that a faster delivery will not affect the lot's cycle time, if its destination tool is not ready to start processing it as soon as it arrives.

2 LITERATURE REVIEW

In the literature, Analytical models of Material Handling Systems (MHS) are usually based on mathematical models and queuing models. The former fails to capture queuing in the system which is essential to accurately estimate the key performance measures including the waiting times. Quite a few papers use queuing models to analyze MHS but their modeling assumptions decrease the accuracy of the approximations, Vis (2004) provide a survey of work in this area. Johnson (2001), Johnson and Brandeau (1993, 1994), and Kobza et al. (1998) develop the queuing models using M/G/c approximations (where c represents the number of vehicles), which gives good approximation provided that vehicle assignments are based on a First Come First Served (FCFS) discipline, but deviate considerably from simulation model results when the vehicle dispatching is system state-dependent, such as Nearest Vehicle Rule NVR. Curry et al. (2003) propose a more accurate service-dependent queuing network model that generates approximations that are close to the simulation results but will the time to solve the analytic model grows exponentially with the number of vehicles. Srinivasan et al. (1994) propose a single-vehicle queuing model to estimate the throughput of the vehicle where the vehicle dispatching to move request is based on a modification of the First Come First Serve (FCFS) rule, basically after the vehicle delivers a load at the input buffer of a station, its searches for a move request at the output buffer of that station, if one or more requests are found, the vehicle is assigned to the oldest one. In Bozer (1994), the authors use the throughput approximation in order to estimate the waiting time of move requests at each station, their results are quite close to the simulation models'. The authors propose an extension of their model to multi-vehicle systems by adjusting the travel times assuming that a *K*-vehicles AMHS can be replaced by a single device that travels *K* times faster. Their results indicate that this works well to estimate throughput but cannot be used to estimate waiting times because it does not account for congestion and blocking delays. In this research, we extend their analytic model to multi-vehicle systems using a closed queuing network approach to estimate the blocking delays.

The intention of our work is to explore the use of queuing network models to estimate the throughput and move request delays of the AMHS. These models for design and control of AMHS are scarce, and the semiconductor industry would benefit from the development and use of analytic modeling tools that have not been previously explored.

We will develop an analytical queuing model of an OHT system that is dispatched based on the First-Encountered-First-Served (FEFS) rule for vehicle/job dispatch. FEFS is a decentralized policy, first presented by Bartholdi and Platzman (1989); it is simple and efficient in the case of a discrete vehicle-based AMHS operating in a simple closed loop as long as the vehicles are not allowed to pass each other, which is the case in the 300mm fab due the space restrictions. In the FEFS, an empty OHV circulating the loop inspects the output buffer of a station (stocker or processor), if there is a lot waiting, it picks it up and delivers it to its destination. If the output buffer is empty, the vehicle travels to the next station and so forth until it encounters a waiting lot. An OHV carrying a lot (loaded/full) might pass other input and output buffers and experience delays if it has to wait while other vehicles drop off or pick up loads at those buffers. Our goal is to estimate these blocking delays in order to get a good approximation of the actual throughput of the OHT system and average response time to move requests. Based on FEFS dispatching, Bozer and Srinivasan (1991) developed a single-vehicle analytical model to approximate the utilization and the throughput capacity of the vehicle, given from-to move requirements and the distance matrix. They assume the move requests arrivals follow a Poisson process. Our model differs because it is developed for is for multi-vehicles, where queuing and blocking of vehicles is possible.

3 VEHICLE-BASED CLOSED-LOOP AMHS DESCRIPTION

3.1 Physical System

Figure 2 illustrates an example of a closed loop overhead transport system. This transportation system is used to serve the material handling requirements of the stocker(s) and the processors (also referred to as production equipment or production tools) in the bay. Flow is unidirectional along the loop. Multiple vehicles can be traversing each segment of the track simultaneously but they cannot pass each other.

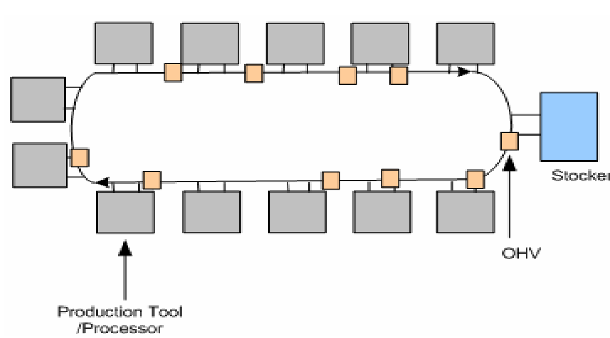

Figure 2: Unidirectional Closed-Loop Overhead Transport System

In order to accurately estimate the throughput, we need to estimate the blocking and queuing delays at each station. These delays are a result of the limited space for vehicles at stations. The main objective of the model is to quantify the length of these two types of delays at each station as a function of the layout of the transportation system, the demand rates, the speed of the vehicles and the number of vehicles circulating the loop. The analysis is complex because the service rate of stations is state-dependent, where the state is a function of where the vehicles are located and whether or not they are empty or loaded. It is not possible to analyze each station independently because the number of vehicles queued at a station impacts the service rate of another station.

Let $L(n)$ refer to the OHT directed loop with n vehicles. We use the term *machine* to refer to either the stocker or the production tool, let M be the set of machines in L(n). Each machine has two load ports: input port where loads are dropped off by the vehicle and output port where loads are picked up by the vehicle to be delivered to their destination. Each port can accommodate one vehicle at a time. These pickup and drop-off ports are all modeled as stations that the vehicles visit while traveling on the loop, and we use the term *station* to refer to the input and output ports of the machines. Thus, a loop serving m machines consist of $s = 2m$ stations. Let m_i denote machine i, then m_i has two stations: the drop-off station s_i^d, and the pick-station s_i^p. Figure 3 illustrates a network representation of the system in Figure 2 as a directed loop with stations represented as nodes, and the tracks as directed arcs.

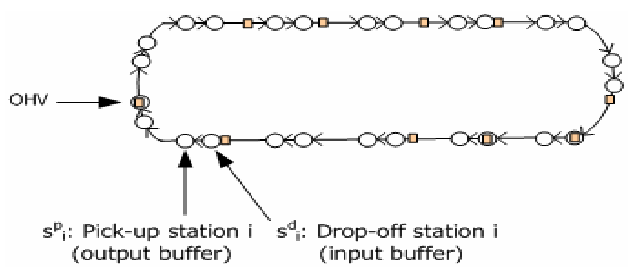

Figure 3: Network Representation of the Closed Loop AMHS

3.2 Logic Description

The vehicles are constantly circulating the unidirectional loop. The vehicle encounters the unloading port (input) before the loading port (output) of a machine m_i (Figure 3). As an empty OHV approaches m_i, it passes through the drop-off station s_i^d, then travels to the pick-up station s_i^p. If there is a load (job) waiting at s_i^p, it picks it up, which requires time delay l for loading the job and then delivers it to its destination, say machine j, visiting machines i+1, i+2,…, j-1, and finally the load's destination the drop-off station of m_m, denoted by s_j^d. The vehicle does not stop at machines i+1, i+2,…, j-1 unless it is blocked by other vehicles. If the output port s_i^p is empty, the vehicle travels to s_{i+1}^d, then inspects the output port s_{i+1}^p and so forth until it encounters a waiting lot.

3.3 The Probabilistic Model

In this section, we develop a probabilistic model of the system L(n) described above. Demands for transportation occur according to a Poisson process that depends on the production rates at the machines and the routing sequences for jobs. Jobs that queue at the machines requesting transportation are processed in first-come first-serve (FCFS) order. A vehicle can experience two types of delays:

1. Queuing delays that occur at pickup and drop-off stations due to the time needed for the other vehicles to clear the station, illustrated in Figure 4, where vehicle 2 (v_2) is queuing behind v_1 until it finishes dropping off a load at s_i^d.

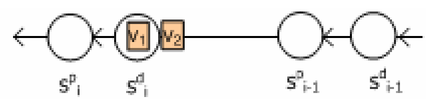

Figure 4: Queuing delay illustration

2. Blocking delays that occur when a vehicle has finished its service at its current station but cannot move because the downstream station does not have any space to accommodate the vehicle. This type of delay is illustrated in Figure 5, v_2 finished its service at s_i^d then got blocked by v_1 while it finishes picking up a load from s_i^p.

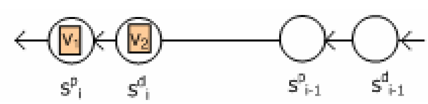

Figure 5: Blocking Delay Illustration

4 THE PROBABILISTIC MODEL ANALYSIS

4.1 Notation

M: set of tools and stockers in the system.

S: set of stations, including all the output and input buffers of the tools and the stockers, $|S|=2|M|$.

S^p : set of pick-up stations $|S^p|=|M|$.

S^d : set of drop-off stations $|S^d|=|M|$.

m_i : machine i, which could be either the output buffer of the input buffer of a tool or a stocker $m_i \in M$.

s_i^p : pick-up station of m_i. $s_i^p \in S^p$; $m_i \in M$

s_i^d : drop-off station of m_i. $s_i^d \in S^d$; $m_i \in M$

$t_{i,i+1}$: time to travel from s_i^p to s_{i+1}^d ; $m_i \in M$

t_i : time to travel from s_i^d to s_i^p ; $m_i \in M$

u: unloading time at s_i^d ; $i \in S^d$

l: loading time at s_i^p ; $i \in S^p$

p_{ij} : probability that a load which is picked up from s_i^p is destined to s_j^d .

λ_i : mean arrival rate of move requests picked up from s_i^p

Λ_i : mean arrival rate of move requests dropped at s_i^d

θ : arrival rate of vehicles to stations, also referred to as the throughput of the AMHS.

As we mentioned earlier the objective is to quantify the length of queuing and blocking delays at each station. Due to the limited space for vehicles at stations and particularly at the pick-up stations, which can accommodate one vehicle at a time, blocking of upstream stations is likely to occur. The service rate of stations is therefore state-dependent, where the *state* is a function of where every vehicle is located and its type. The location of a vehicle specifies the station the vehicle is receiving service or arriving at. The type of the vehicle specifies whether it is loaded, empty, blocked while empty, blocked while loaded, or receiving service (picking up or dropping off a load), we use *f*, *e*, *b*, *k*, and *s* to denote each of these states, respectively. Moreover, the order of the vehicle in the queue is important to determine the transition between the states. Since only the first vehicles in stations' queues can be blocked or receiving service, only these vehicles can be in states *e*, *f*, *b*, *k*, or *s*, while the other vehicles are either in state *e* or *f*. Also, a vehicle can be blocked only if the immediately downstream station has reached its capacity.

Let b_i denote the capacity of s_i including the loadport, then the total number of possible vehicle locations in the loop is $B = \sum_{i=1}^{S} b_i$. We represent a state by a string of size ***B***, where each character in the string represents a possible location. For example, consider a system with four station, s_1-s_4, where b_1=1, b_2=2, b_3=1, and b_4=2. State (0 *s f* 0 0 0) indicates that there are two vehicles at the second station, the first one is receiving service (unloading) and the other is arriving loaded.

We assume that travel times, and loading and unloading times are deterministic, however, the model is probabilistic because the Poisson arrival process of move requests determines the transitions between system states. Based on the FEFS dispatching discipline, Bozer and Srinivasan (1991) developed a single-vehicle analytical model to approximate the utilization and the throughput capacity of the vehicle, given from-to move requirements and the distance matrix. They assume the move requests arrivals follow a Poisson process. Our model differs because it is developed for is for multi-vehicles, where queuing and blocking of vehicles is possible, and we assume that the vehicle travels from s_i to s_{i+1} with probability 1.

We propose to analyze the system using a Markov Chain. Since the move request (loads) arrivals follow a Poisson process, by the PASTA property (Poisson Arrivals See Time Averages), we can assume that the instant at which a vehicle arrives to s_i^p is a random point in time, and the system transition between the possible states is Markovian.

Consider the transition matrix **R**, which specifies the movement of the system between the states. The position and type every vehicle (*e*, *f*, *b*, *k*, or *s)* determines the system transitions, but the vehicles that are at the head of the stations' queues control the probabilistic transitions. For instance, if there are two empty vehicles arriving at some pickup station s_i^p , the transition to the next state depends on whether or not the first vehicle will encounter a waiting load; in any case, the next state will have the second vehicle arriving empty at s_i^p .

An empty vehicle arriving to a drop-off station will certainly move to the next pick-up station, and so is the case for a loaded vehicle arriving to a pick-up station. However, a loaded vehicle approaching a drop-off station s_i^d will leave empty if the load was not destined to that station, this happens with probability (1-r_i) which depends on the rate of moves destined to this station and also on the rate of loaded vehicle arrivals that pass through the station. Similarly, an empty vehicle approaching a pick-up station s_i^p will leave empty if there was no load waiting at s_i^p , this happens with probability q_i which depends on the rate of moves originating at this station and on the rate of empty vehicle arrivals to the station. A vehicle receiving service at s_i^d must be dropping off its load, and thus it will certainly move empty to s_i^p , and a vehicle receiving service at s_i^p must be picking a load and thus it will certainly move loaded to s_{i+1}^d .

4.2 Transition Probabilities

Consider again the transition matrix R, which specifies the movement of the system between the states. We demonstrate the transitions through the following example.

Example: consider a closed-loop OHT system with two vehicles (n=2) and four stations (Figure 6). The pickup stations (s_1^p, and s_2^p) have capacity for one vehicle each, while the drop-off stations (s_1^d and s_2^d) have capacity for two vehicles each.

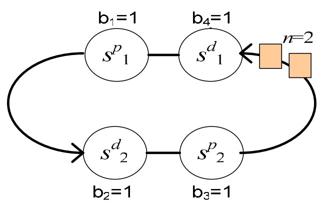

Figure 6: A 4-Station 2-Vehicle Example

This simple system has two machines each sends loads to the other, thus $p_{12} = p_{21} = 1$. Also, $\Lambda_1 = \lambda_2 p_{21} = \lambda_2$, and $\Lambda_2 = \lambda_1 p_{12} = \lambda_1$.

Each state is defined by the string that specifies the type of vehicle at each location, the first character refers to the single buffer of s_1^p, the second and third characters refer to the two buffers of s_1^d, the fourth character refers to the single buffer of s_2^p, and the last two characters refer to the two buffers of s_2^d. State : (*e e 0 0 0 0*) indicates that there is one empty vehicle at s_1^p and one empty vehicle at s_1^d. The transitions from this state depends on whether or not the first vehicle finds a load at s_1^p, which happens with probability 1-q_1. If the vehicle finds a load it starts receiving its service at s_1^p while the second vehicle is blocked while empty, and the system enters state (*s b 0 0 0 0*). The first vehicle does not find a load waiting with probability q_1, and in this case, the first vehicle moves empty to s_2^d while the second vehicle moves empty to s_1^p, and the system enters state (*e 0 0 0 0 e*). Consider the transition from state (*e 0 0 0 f 0*) to state (*s 0 0 0 s 0*): in the first state, one vehicle is arriving empty to s_1^p, and the other is arriving loaded to s_2^d, in the second state, both vehicles are receiving service at the same stations they were at. This transition happens if the first vehicle finds a load waiting at s_1^p with probability 1-q_1, and the second vehicle drops off its load at s_2^d with probability r_2. Part of the 54×54 transition matrix for this example is provided in Figure A-1.

Note that in the above example r_1=r_2=1, because the two machines send loads only to each other.

If we can estimate the steady-state probability of each state, we can estimate the blocking probabilities and the system throughput.

Let C_r be the expected time between two consecutive visits to state r, r=1,…,|R|. Without loss of generality, assume that C_1 is the expected time between two visits to state (1). Let **v** ={v_r} r=1, …, |R|, where v_r denote the visit ratio to state r, which is the number of times the system visits state r between two successive visits to state (1), by this definition, v_1=1, and the visit ratios are uniquely obtained from:

$$vR = v \tag{1}$$

$$v_1 = 1 \tag{2}$$

The number of unknowns in the above equations is |R| unknowns for the visit ratios, plus |m| for the unknown probabilities q_i's.

Let E_i be the set of states where the first vehicle in queue at pickup station s_i^p is arriving empty i=1, …, S^p, Let v_{E_i} denote the visit ratio to state set E_i, which, conceptually, is the number of times empty vehicles visit s_i^p in a cycle of length C_1. Let e_i^p denote the arrival rate of empty vehicles to s_i^p, therefore:

$$v_{E_i} = e_i^p C_1, \ \forall i \in S^P \tag{3}$$

For a stable system, the rate of pickups from s_i^p must equal the rate of move requests generated at s_i^p, which is λ_i. The rate of pickups from s_i^p is the rate of empty vehicle arrivals to s_i^p (denoted by e_i^p) multiplied by the probability of finding a load waiting at s_i^p, thus the stability conditions are:

$$q_i = 1 - \frac{\lambda_i C_1}{v_{E_i}}, \forall i \in S^P \tag{4}$$

This set of equations provide |m| equations but introduces one unknown, C_1.

4.3 Expected Transition Times

We can develop an expression for C_1 by considering the transition times for the system to move from one state to the next. Let T_r denote the expected time from the instant the system enters state r until the instant it enters the next state. C_1 defined as the expected time between two successive visits to state (1), can be obtained from:

$$C_1 = \sum_{\forall r \in R} v_r T_r \tag{5}$$

The terms T_r, $r \in R$ can be determined similar to the transition probabilities based on the positions and types of the vehicles that are at the head of the stations' queues. For instance, if there are two empty vehicles arriving at some pickup station s_i^p, the time the system spends in this state depends on whether or not the first vehicle will encounter a waiting load; with probability q_i, the first vehicle will move to the next drop-off station s_{i+1}^d and the expected transition time is simply $t_{i,i+1}$. With probability $1-q_i$, the vehicle will find a load and in the next state this vehicle will be receiving a service at s_i^p and thus the transition time is zero. The expressions for T_r, $r \in R$ become more complicated when more than one vehicle is able to change its position and/or its status.

Combining equations (1), (2), (4) and (5), we have $|R|+|m|+1$ equations and $|R|+|m|+1$ unknowns, therefore we can find the unique solution to the system of equations and calculate the visit ratio to every state, the blocking probabilities and system throughput.

To estimate the throughput, choose a pick-up station s_i^p arbitrarily, let L_i be the set of states where the first vehicle in queue at pickup station s_i^p is arriving loaded i=1, …, S^p, Let v_{L_i} denote the visit ratio to state set L_i, which, conceptually, is the number of times loaded vehicles visit s_i^p in a cycle. Recall that v_{E_i} denote the number of times empty vehicles visit s_i^p in a cycle, the throughput is the sum of empty and loaded arrivals divided by the cycle length:

$$\theta = \frac{v_{L_i} + v_{E_i}}{C_1} \tag{6}$$

The choice of s_i^p is arbitrary since the system is a closed loop and the arrival rate of vehicles should be equal for all the stations.

We now apply the above method to the system with two vehicles and four stations. Recall that the pickup stations (s_1^p, and s_2^p) have capacity for one vehicle each, while the drop-off stations (s_1^d, and s_2^d) have capacity for two vehicles each. The total number of states for this system is 54, provided in Table A-1.

We use the layout in Figure 6 to validate the estimates of system throughput using discrete event simulation. The layout was analyzed with various data sets that result in different AMHS throughput. The error percentage when comparing the throughput estimated using the queuing model to that using the simulation is reported in Figure 7, for different release rates.

The queuing model can be used to study the system behavior. For example to investigate the effect of increasing the number of vehicles on the AMHS throughput. The behavior is similar to what would be expected in a closed queueing network with finite buffers, as illustrated in Figure 8. As the fleet size increases the system throughput increases until the point where the blocking delays start to adversely affect the throughput and increasing the number of vehicles will reduce the throughput. For our example, the maximum throughput is achieved at n=3.

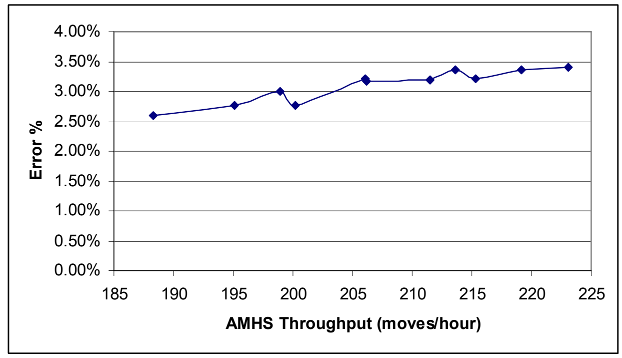

Figure 7: Comparison of Queuing and Simulation Model Throughput Estimates at Different Release Rates.

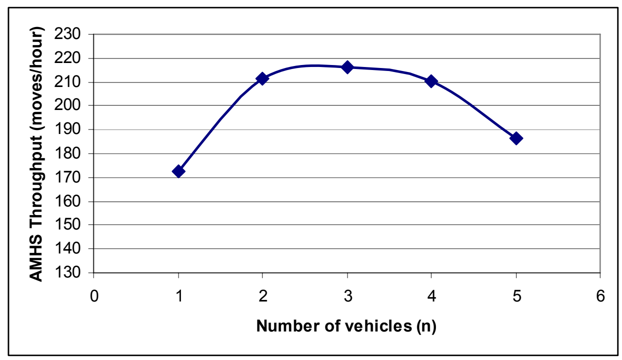

Figure 8: Impact of Increasing AMHS Fleet Size on Throughput.

5 FUTURE WORK

We are currently working on further validation of the model for more realistic systems. Later, we will develop expressions for the expected waiting time of loads for a vehicle. From there, we will explore the possibility of using the AMHS analytic model as a tool used in conjunction with simulation to provide quick performance evaluation of large-scale simulation models such as those used to simulate 300mm wafer fabrication facilities. Detailed AMHS simulation models are capable of capturing the variability of the fab, but this capability comes with the high price of excessive development, validation and running times. Specifically, we will explore the possibility of linking the AMHS queueing model with a detailed simulation model of the production equipment and lot scheduling system. This approximation will tremendously reduce the running time of the simulation model.

ACKNOWLEDGMENTS

We gratefully recognize the generous support by the W. M. Keck Foundation of the Keck Virtual Factory Lab.

APPENDIX: ADDITIONAL DATA

Table A-1 lists the possible states for the numerical example in Section 4.2. Part of the transition matrix is displayed in Figure A-1.

Table A-1: Possible States of the System in Figure 6

State	String	State	String	State	String	State	String	State	String	State	String
1	*0s0e00*	11	*e000s0*	21	*0s0s00*	31	*ef0000*	41	*000sk0*	51	*f00e00*
2	*0e00f0*	12	*0s00f0*	22	*e00f00*	32	*e00e00*	42	*f00s00*	52	*f00f00*
3	*ss0000*	13	*ee0000*	23	*000ef0*	33	*0e00s0*	43	*0s0f00*	53	*000ff0*
4	*e000e0*	14	*s00e00*	24	*0e0e00*	34	*0s00e0*	44	*000ss0*	54	*000ee0*
5	*s00f00*	15	*se0000*	25	*0f0e00*	35	*ff0000*	45	*s000s0*		
6	*0000sf*	16	*000sf0*	26	*f000e0*	36	*0sf000*	46	*0se000*		
7	*fe0000*	17	*f000f0*	27	*sf0000*	37	*000se0*	47	*0s00s0*		
8	*e000f0*	18	*s00s00*	28	*f000s0*	38	*e00s00*	48	*0000se*		
9	*0e00e0*	19	*0f00e0*	29	*0f00f0*	39	*sb0000*	49	*000fe0*		
10	*0e0f00*	20	*0f0f00*	30	*000sb0*	40	*0f00s0*	50	*sk0000*		

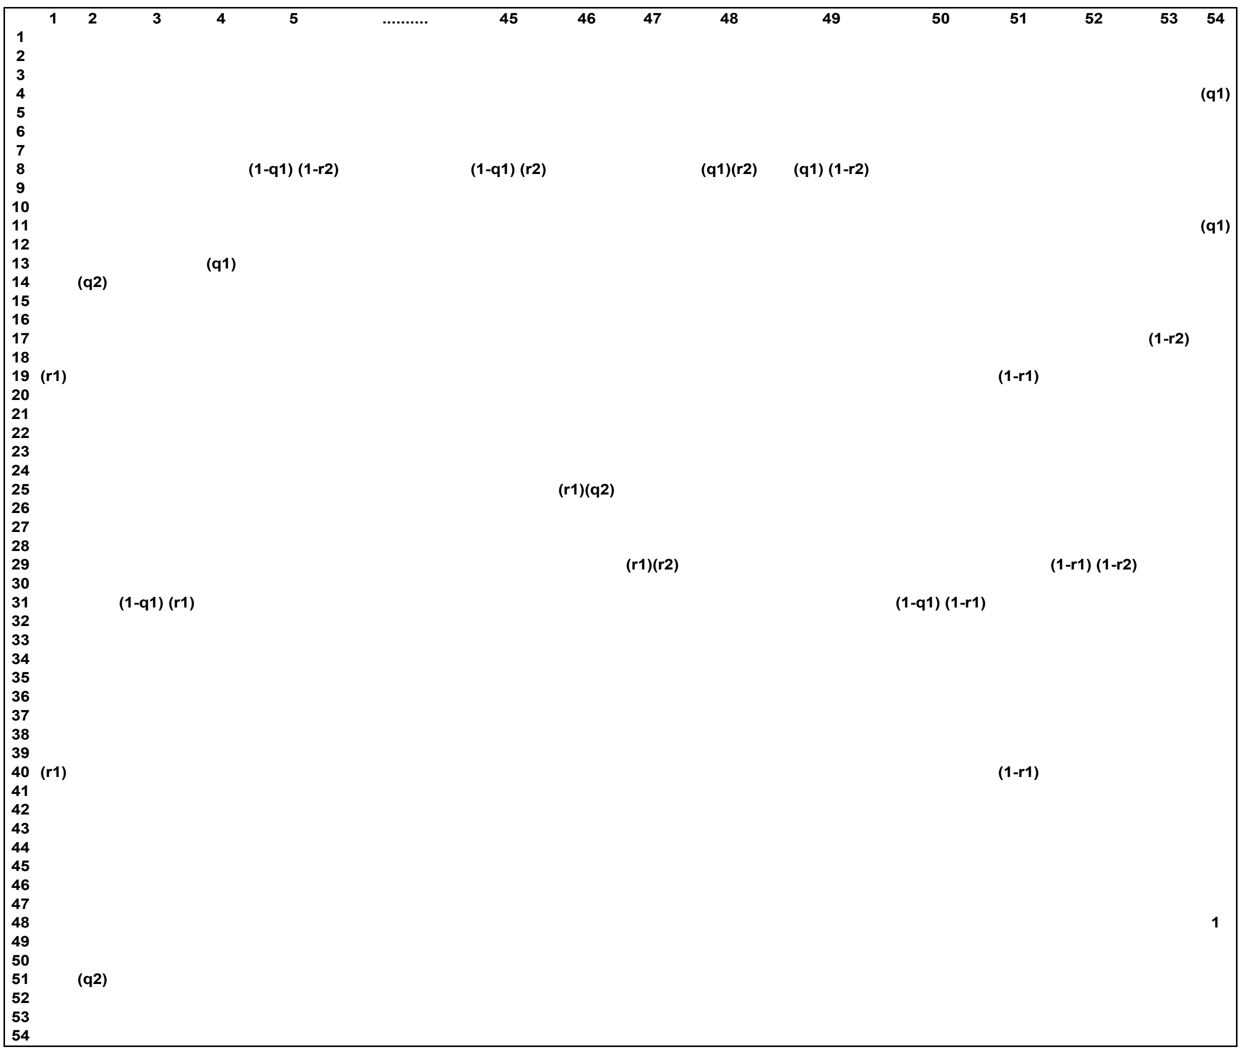

Figure A-1: Transition Matrix

REFERENCES

Bartholdi J. J. III and L. K. Platzman. 1989. Decentralized control of automated guided vehicles on a simple loop. *IIE Transactions*. 21, 76-81.

Bozer Y. A., C. Myeonsig , and M. M. Srinivasan. 1994. Expected waiting times in single-device trip-based material handling systems. *European Journal of Operational Research*, 75, 200-16.

Bozer Y. A. and M. M. Srinivasan, , and C. Myeonsi. 1991. Tandem configurations for automated guided vehicle systems and the analysis of single vehicle loops. *IIE Transactions*, 23, 72-82.

Curry G. L., B. A. Peters, and M. Lee. 2003. Queueing network model for a class of material-handling systems. *International Journal of Production Research*, 41, 3901-20.

Johnson M. E. 2001. Modeling empty vehicle traffic in AGVS design. *International Journal of Production Research,* 39, 2615-33.

Johnson M. E. and M. L. Brandeau. 1994. An analytic model for design and analysis of single-vehicle asynchronous material handling systems. *Transportation Science,* 28, 337-53.

Johnson M. E. and M. L. Brandeau. 1995. Designing multiple-load automated guided vehicle systems for delivering material from a central depot. *Transactions of the ASME. Journal of Engineering for Industry,* 117, 33-41.

Kobza J. E., S. Yu-Cheng and R. J. Reasor. 1998. A stochastic model of empty-vehicle travel time and load request service time in light-traffic material handling systems. *IIE Transactions*, 30, 133-42.

Srinivasan M. M., Y. A. Bozer, and C. Myeonsig. 1994. Trip-based material handling systems: throughput capacity analysis. *IIE Transactions*, 26, 70-89.

Vis I. F. A. 2004. Survey of research in the design and control of automated guided vehicle systems. *European Journal of Operational Research* [online]. Available via <http://www.elsevier.com/locate/dsw> [accessed November, 2004].

AUTHOR BIOGRAPHIES

DIMA NAZZAL is a Ph.D. student in the School of Industrial and Systems Engineering at the Georgia Institute of Technology. She can be contacted by email at <dnazzal@isye.gatech.edu> and her Web address is <http://www.isye.gatech.edu/~dnazzal>.

LEON F. MCGINNIS is Gwaltney Professor of Manufacturing Systems at the Georgia Institute of Technology, and Director of the Product Lifecycle Management Center of Excellence. His primary research interests are in computational models to support design of large-scale complex discrete event logistics systems. His e-mail address is <leon.mcginnis@isye.gatech.edu> and his Web address is <http://www.isye.gatech.edu/~lfm>.