

Performance Analysis of Lot Release and Dispatching Decisions in Wafer Fabrication

Pyung-Hoi Koo, Shie-Gheun Koh and Min-Jung Park

Department of Systems Management & Engineering, Pukyong National University
Yongsoro 45, Namgu, Busan, 48513, South Korea.

Abstract

Controlling semiconductor wafer fabrications (wafer fabs) is a difficult task because the wafer fabs have several distinct characteristics such as a large number of processes involved, long lead time, and re-entrant product flow. Typical operational control issues in the wafer fabs include job releasing, dispatching and batch loading decisions. Most of the research works on operational control problems in the fab studied so far address these control problems separately. This paper is motivated by the belief that an operational control decision affects the performance of the other control decisions, and the existence of bottleneck machines is a core factor for this mutual influence of the control schemes. To examine the impact of the control rules on the system-wide performance, we present bottleneck machine-focused control rules and construct simulation models where these control rules are incorporated. Through simulation studies, the performance of the presented control policies is evaluated, and some insights on the performance of the control policies and their interactions are discussed.

Keywords: semiconductor wafer fabrication; lot release; dispatching; batch loading; bottleneck, simulation analysis

INTRODUCTION

In semiconductor wafer fabrications (wafer fabs), electrical circuits are repeatedly built on the layers on the wafers. Figure 1 shows basic process steps of a typical wafer fab. Manufacturing process at the wafer fab starts with raw wafers, thin and round slices usually made of silicon. The wafers are released to the wafer fab in lots each of which generally consists of 20~25 individual wafers, and they move between the workstations as a wafer lot. To produce as many semiconductor chips as possible from a wafer, the chips need to be highly miniaturized and so the circuits are built up on multiple layers. There are about 40 circuit layers for advanced semiconductor chips. Each layer requires the basic operation steps including oxidation/diffusion, photolithography, etching, deposition, ion implantation and so on. This means that each wafer lot visits (reenters) the same workstation several times.

In general, a wafer visits 300-700 process steps on hundreds of different machines [8]. Because of the long sequences of operations required for the wafers, most wafer fabs suffer from high work-in-process (WIP) inventory and long lead times (more than one month). In addition, the cost of building a wafer fab is enormous, often requiring more than ten billion dollars. Hence, it is very important for the wafer fab to be utilized in a full capacity.

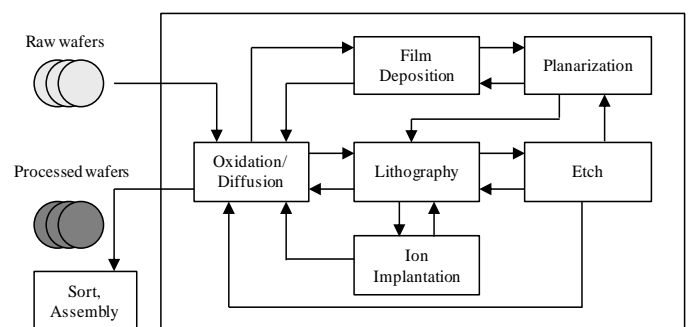


Figure 1: Operations in a wafer fab (Monch et al. 2013)

Another distinct feature of the wafer fab is the existence of batch processors (BPs). The processors with capacity more than one are referred to as batch processors (BPs) whereas the processors with unit capacity are referred as discrete processors (DPs). The BP can process several wafer lots simultaneously as long as the batch size does not exceed its capacity. Diffusion furnaces are a typical example of the BPs. In many wafer fab, more than one-third of the wafer fab operations are performed on BP stations [8]. The BPs have some distinct characteristics from operational perspectives. Due to the chemical nature of the process, it may be impossible to process the jobs for different layers together in the same batch. In general the batch processing times are very long, roughly of order of 10 hours as opposed to one or two hours for the discrete processes in wafer fabrication facilities. Since most of their upstream machines are discrete processes, the wafer lots arriving at the BPs should be formed as a batch before being processed on a BP. After being processed, the

wafer lots are split and individually transferred to the downstream machines at the same time. The non-smooth product flow leads to a formation of high WIP inventories in front of the BPs and downstream DPs. The long processing time and non-smooth product flow increase system variability in the wafer fab, and hence the batch operations have a great effect on system performance in terms of throughput, WIP inventory, cycle time and on-time delivery. A number of research works have addressed BP control problems. However, most of them have presented solutions to minimize waiting time of the jobs at the BP workstations, ignoring the impact of the batching decision on the system performance, which is a major interest in this paper.

In manufacturing flow line with a variety of processes, the system capacity is determined by some processes, which is called bottleneck processes. The bottleneck processes not only limit the system capacity but also have a large impact on the other system performance criteria such as lead time and WIP inventory. In general, photolithography tools are identified as a bottleneck resource in the wafer fabs. The bottleneck machines are operated by its full capacity to produce as many products as possible within a given production time. The loss of operation time at the bottleneck machines can be viewed as the time loss of the whole system. On the other hand, the loss of operation time at the non-bottleneck machines may have little effect on the system performance because they have spare capacity that can be used to make up the time loss. Hence, most existing studies on dispatching decisions focus on bottleneck machines [4, 6]. There are a large number of research works on the control issues for the bottleneck machines. However, little research works have been done on the control issues for the non-bottleneck machines.

A large amount of research works have showed that the performance of the wafer fabs is strongly affected by the operational control decisions [1, 7, 13]. In this paper, we attempt to examine operational issues of the wafer fabs by answering some questions regarding performance of various control schemes and their relationships. In most wafer fabs, the major objective of operational control decisions is to minimize production flow time while maintaining a target throughput rate. The control problems in the wafer fab can be classified into two major problems, lot release and dispatching. Lot release involves a fab-level decision that determines when and how many wafer lots should be released to the wafer fab for the first process while dispatching is a workstation-level decision that determines which lot should be processed next on an idle machine. Most of the research works on operational control problems in the fab studied so far address the job release and dispatching problems separately. This paper considers these control problems at the same time. In addition, new control rules are presented for bottleneck stations, non-bottleneck stations and batch processing stations, respectively, in which bottleneck resources are a major consideration in the decision making. A simulation model is constructed where all the control policies are incorporated. Through simulation

experiments, the performance of the control rules is evaluated and their relationships are discussed.

OPERATIONAL CONTROL RULES IN WAFER FAB

The wafers are released to a wafer fab following a lot release rule and the released wafers are loaded and processed on the machines based on dispatching rules. The lot release rules can be classified into open-loop release policies and closed-loop release policies. The open-loop release policies release wafer lots into the fab based on static production scheduling regardless of the current system status while the closed-loop release policies consider wafer fab status in release decision making. One of the simple open-loop release policies is DETERM where wafer lots are released into the fab in a constant interval, for example every four hours. With the DETERM lot release rule, the production rate can be anticipated with certainty by the release interval if the scheduled utilization of the bottleneck machines is less than 100%. However, the WIP inventories are variable depending on the operational control decisions at the workstations. A simple closed-loop release policy is CONWIP [11] where a constant number of WIPs are maintained in the fab. In CONWIP release rule, whenever a completed wafer lot leaves the wafer fab, a new wafer lot is introduced to the system. Contrary to the open-loop control systems, the throughput rates are variable depending on the operational control decisions at the workstations while WIP levels are constantly maintained. The other widely known closed-loop release policies include workload regulating (WR) rule [13] and starvation avoidance (SA) rule [2]. These two rules consider bottleneck workstations in the fab when lot release decisions are made. In the WR rule, when the expected amount of work for the bottleneck station drops to a specified level, a new lot is released to the wafer fab. In the SA rule, when virtual inventory at the bottleneck station falls down to a predetermined value, a new wafer lot is released. The virtual inventory at the bottleneck station is estimated by the sum of the actual inventory at the workstation and WIP inventory at upstream workstations that is expected to arrive at the bottleneck station within the expected lead time to the next bottleneck station. In both WR and SA rules, the throughput is controlled by changing the critical values. A lot of closed-loop lot release rules have been proposed later, which are mostly variations of WR and SA rules [7, 9].

When a machine completes processing a wafer lot and there is more than one wafer lot waiting in its queue, a dispatching decision is made which lot should be loaded for being processed next. Most of the existing dispatching approaches are based on priorities that are set by using product information such as FCFS (first come first served), SRPT (shortest remaining processing time) and EDD (earliest due date) as well as system status including workload and WIP levels in downstream and/or upstream workstations. A wide variety of dispatching policies have been studied, and they are

reviewed extensively in Sarin et. al. [10]. For the batch processors, the operational problems are involved on both which wafer lot to process next (dispatching decision) and how many lots to put in the batch (loading decision). The dispatching decision refers to the prioritization of the wafer lots that are put together in a batch while the loading decision considers a trade-off between starting the batch immediately and waiting for more lots to arrive. A commonly used operational rule for the batch processors is minimum batch size (MBS) rule where a batch of wafer lots is loaded when current batch size is more than or equal to a predefined MBS. The MBS control policies are easy to implement on the shop floor because it requires minimal computation with only local information such as current wafer lots waiting in queue. However the MBS rule may lose a chance to improve system performance by utilizing real-time shop floor information that is often available in a modern manufacturing system. Today, many semiconductor shop floor attempt to apply look-ahead batching rules where downstream and/or upstream workstation status and/or near-future product arrival are considered. For example, if we have some information about future product arrivals at the time that a machine is idle and there are wafer lots in queue smaller than the machine capacity, we may use this information to determine whether to start the process right away (thus utilizing the machine below its capacity) or wait for the future job arrivals to increase the fullness of a certain batch (but delaying all jobs currently in queue). A number of look-ahead batching rules have been proposed in an attempt to minimize waiting time or tardiness. Extensive review on batching decisions is found in van der Zee [12] and Koo & Moon [5].

In order to analyze the performance of the operational decisions, we present operational control rules for the wafer fabs. As discussed earlier, the system capacity of the production line is determined by bottleneck stations. The new control rules take the idea of Theory of Constraints (TOC) popularized by Goldratt [3] in operational decision process. To examine and compare the impact of the bottleneck and non-bottleneck stations, new dispatching policies are presented separately for bottleneck and non-bottleneck stations. In addition, a batch loading policy is also presented. The bottleneck machines are a core part of the decision processes. For the lot releasing, since the existing rules WR and SA consider the workload of the bottleneck machines when the decision is made, they are incorporated in our models. The following summarizes dispatching for non-bottleneck stations, dispatching for bottleneck station, and batch loading for batch processing station.

Dispatching for non-bottleneck DP stations

Dispatching decisions at the non-bottleneck machines is expected to affect the performance of downstream bottleneck machines. Hence, the operational control in the non-bottleneck machines should focus on the performance of the

downstream bottleneck machines. We present FCFS_BM and SRPT_BM rules where the decision is made in the way that the downstream bottleneck machines are not starved. In these rules, the control decision is made with FCFS and SRPT rules, respectively, in an ordinary situation. However, if the downstream bottleneck machine of a wafer lot is anticipated to be idle shortly, the wafer lot is loaded with higher priority. The new dispatching rule can be viewed as a variation of the rule proposed by Zhou and Rose [14] where non-bottleneck stations are an immediate predecessor of the bottleneck station.

Dispatching for bottleneck DP stations

In the wafer fab, it is important to produce the products evenly over time because the large fluctuation of the production output over time may result in high inventory or stockout. Hence, we present LOOP_LVL rule for the bottleneck machines where the workload of each process loop is leveled. Here the process loop is defined a series of processes between two bottleneck stations. In LOOP_LVL, a target WIP is first calculated for each process loop, proportionate to the average and variability of the flow time. The average and variability of WIP levels for each process loop are obtained through preliminary simulation experiments. A wafer lot with lower actual/target WIP ratio has a process priority.

Batch loading for BP stations

Most of the existing batching rules address isolated problems for optimizing the local performance of batch processing machines [5]. In this paper, we present a batch loading rule, MBS_BM, where downstream bottleneck machines are considered in loading decision. In MBS_BM rule, the loading decision is made with MBS rule in an ordinary situation. However, if the downstream bottleneck machine is expected to be idle shortly, the wafer lots are loaded with higher priority, even with batch size less than MBS.

SIMULATION MODEL AND PERFORMANCE ANALYSIS FOR WAFER FAB

To examine the performance of various operational control strategies, simulation experiments are carried out on a wafer fab. The configuration of the wafer fab in this study is obtained from Wein [13] with a slight modification in equipment type and processing time. The fab consists of 24 single or multi-server workstations subject to unexpected facility downtime. Each wafer lot has a process flow with 172 operation steps at 24 different workstations. Some workstations (station 1 through 4) are batch processing workstations. Each job visits the bottleneck workstation (workstation 14) twelve times whose utilization is 91.9%. Table 1 shows the control rules considered in the simulation studies.

Table 1: Operational control policies applied for simulation experiments

Control		rules
Lot release		DETERM, CONWIP, WR,
Dispatching	Non-BM	FCFS, SRPT, FCFS_BM, SRPT-BM
	BM	FCFS, SRPT, FCFS_BM, SRPT-BM, LOOP_LVL
	Batch	MBS, MBS-BM

For lot release rules, a open-loop control policy, DETERM, and two closed-loop control policies, CONWIP and WR(workload regulation), are examined. The inter-arrival time of the DETERM rule is 12 hours (leading to 91.9% utilization for the bottleneck station). The parameters for the closed-loop control rules are chosen through preliminary experiments so that the average throughput rate is roughly the same as the average throughput rate of the DETERM lot release cases. Here, the number of lots maintained in the fab is set to 68 lots in CONWIP rule, and the total workload of the bottleneck station is set to 895 hours in WR rule. For dispatching the non-bottleneck station, FCFS, SRPT, FCFS_BM and SRPT_BM are modeled. For bottleneck machines, Loop_LVL rule is examined in addition to the dispatching rules of the non-bottleneck stations. For batch processing machines, MBS and MBS_BM are investigated. A set of control rules (DETERM, FCFS, FCFS, MBS) is considered as a base case where for this 4-tuple (r1,r2,r3,r4) control rules, r1 denotes lot release, r2 dispatching for non-BM, r3 dispatching for BM, and r4 loading for batching machines. The simulation runs are repeated 20 times for each ruleset. Simulation experiments run for two years, and the statistics only from the second year are used to see the system performance in a steady state.

In order to evaluate the proposed control schemes where the workload of the downstream bottleneck machine is taken into account, experiments are performed with different control rule sets. From the simulation experiments, we come up with some interesting findings on the effect of the operational control decisions on the system performance. Table 2 summarizes the results of different control schemes. Here, among two important performance measures, throughput rate and lead time, the lead time is only compared because the throughput rate is deterministic in DETERM lot release rule when the system utilization is less than 100%. With the base case, the average lead time is 816.4 hours and the production rate is 728.1. The results show that a set of the control schemes considering bottleneck machines, RA0, provides better performance (less lead time) than the base case. The lead time is reduced by 22.7 hours (almost one day) from 816.4 to 793.7. When we take a close look at the results from rulesets RA1, RA2 and RA3, we find that operational controls for the bottleneck, non-bottleneck, and batch processors give different impact on the system performance: i.e., among these

three rule sets, the RA3 focusing on batch loading decision gives the lowest lead time. The large impact of the batch loading decision can be explained by the distinct feature of batch processors including long processing time (five to ten times longer than the discrete type machines) and repetitive batching/splitting activities leading to non-smooth product flow. Hence, the experimental results say that the decisions on batch processing operations are an important issue from managerial perspective to achieve good fab-wide performance.

Table 2: Simulation results with different dispatching rules

Ruleset case	Operational control rules				Experimental results	
	Lot release	Dispatching for Non-BM	Dispatching for BM	Loading for batch processor	Lead time (hours)	Throughput rate(lots)
Base case	DETERM	FCFS	FCFS	MBS	816.4	728.1
RA0	DETERM	FCFS_BM	FCFS_BM	MBS_BM	793.7	733.0
RA1	DETERM	FCFS_BM	FCFS	MBS	807.7	733.5
RA2	DETERM	FCFS	FCFS_BM	MBS	817.0	732.2
RA3	DETERM	FCFS	FCFS	MBS_BM	795.4	731.0

The workstations in a production line can be classified as bottleneck and non-bottleneck stations. Since the bottleneck stations determine the capacity of the whole system, they are controlled with much attention. On the other hand, the non-bottleneck workstations are paid less attention. Most existing control decisions in non-bottleneck machines use a simple control rule like FCFS rule (This is a general control rule for non-bottleneck machines in TOC implementations). To examine the impact of the operational controls at bottleneck stations and non-bottleneck stations, a variety of control rule combinations have been simulated. The experimental results are shown in Table 3. In rulesets RA1.1 and RA1.2, the new operational control rules are applied to the non-bottleneck stations while in rulesets RA2.1, RA2.2 and RA2.3, the new control rules are applied to the bottleneck stations. It is seen from the table that the controls at the non-bottleneck stations have a great influence on the performance as much as the control at the bottleneck stations. Among the five control rules, the RA1.2 in which SRPT_BM is applied to the non-bottleneck stations gives the least lead time. From these results, we insist that the non-bottleneck stations should be also carefully controlled for higher system-wide performance.

Table 3: Simulation results with dispatching rules of bottleneck and non-bottleneck stations

Ruleset case	Lot release	Dispatching for Non-BM	Dispatching for BM	Loading for batch processor	Lead time (hours)	Throughput rate(lots)
Base case	DETERM	FCFS	FCFS	MBS	816.4	728.1
RA1.1	DETERM	FIFO_BM	FCFS	MBS	807.7	733.5
RA1.2	DETERM	SRPT_BM	FCFS	MBS	799.4	729.5
RA2.1	DETERM	FCFS	FIFO_BM	MBS	817.0	732.2
RA2.2	DETERM	FCFS	SRPT_BM	MBS	809.6	728.2
RA2.3	DETERM	FCFS	LOOP_LVL	MBS	811.0	729.2

Lot releasing and dispatching activities are often performed in a sequential way: that is, the wafer lots are released first into the wafer fab, and then the released lots are processed according to dispatching rules. Many previous research works insist that lot release rules have more effect on the system performance than dispatching rules [4, 7, 13]. In order to examine and compare the effect of lot release and dispatching rules, simulation experiments have been performed with ruleset RB1 where a bottleneck-focused lot release rule is applied to the wafer fab level and ruleset RB2 where bottleneck-focused dispatching rules are applied to the workstation level. Table 4 shows the simulation results showing impact of lot release and dispatching rules. Compared with the base case, bottleneck-focused dispatching ruleset, RB2, reduces the lead time by 22.7 hours from 816.4 to 793.7 while the closed-loop lot release rule, RB1, reduces the lead time by 20.0 hours from 816.4 to 796.4. Here, the results say that the dispatching rules have no less influence on the system performance than the lot release rule, which is contradictory to the most previous works handling lot release and dispatching rules together. One of the possible reasons may be as follows: when there are hundreds of processes that a product should visit before it leaves the system (like in case of the wafer fab), the impact of the lot release rules may be dissipated as the process goes on along the long routes. In addition, with the reentrant property of the wafer fab, for each workstation, there is a bottleneck workstation in a downstream, and so dispatching decision of a workstation affects the performance of the downstream bottleneck machines, which makes the dispatching decision at the local workstations more important.

Table 4: Simulation results showing impact of lot release and dispatching rules

Ruleset case	Lot release	Dispatching for Non-BM	Dispatching for BM	Loading for batch processor	Lead time (hours)	Throughput rate(lots)
Base case	DETERM	FCFS	FCFS	MBS	816.4	728.1
RB1	WR	FCFS	FCFS	MBS	796.4	728.2
RB2	DETERM	FCFS_BM	FCFS_BM	MBS_BM	793.7	733.0

Finally, we attempt to examine the performance of the loading decisions at the batch processors (BPs). Most research works on BP control decisions attempt to minimize the waiting time at the BP station [5]. This is based on the belief that minimizing the waiting time is important because the waiting time at the batch processors accounts for a large portion of the production lead time at the wafer fab. We have performed simulation experiments to see whether the waiting time minimization at the local batch processors is a reasonable performance criterion from the system's perspective. MBS and MBS_BM rules are used to see the performance of the loading decisions. Table 5 shows simulation results. From the simulation results, we can find that the MBS rule provides less

waiting time at the local BP stations than MBS_BM does. However, the MBS_BM provides less system lead time than MBS. For example, the ruleset (DETERM, FIFO, FIFO, MBS) gives waiting time 3.93 hours, less than the ruleset (DETERM, FIFO, FIFO, MBS_BM) by 4.4 hours. However, when we look at the system lead time, the former ruleset gives longer lead time than the latter ruleset by 21 hours. It can be concluded that even though a control rule produces a lower waiting time at BP stations, it may increase system lead time which is a major performance measure. Hence, very careful attention should be made when we develop the control schemes for BP stations.

Table 5. Simulation results with different batch loading rules

Operational control rules				Experimental results	
Lot release	Dispatching for Non-BM	Dispatching for BM	Loading for batch processor	Lead time (hours)	Waiting time at BP stations (hours)
DETERM	FIFO	FIFO	MBS	816.4	3.93
			MBS_BM	795.4	8.33
	FIFO_BM	FIFO_BM	MBS	801.9	3.84
			MBS_BM	793.7	8.15

CONCLUSIONS

This paper examines the performance of operational control strategies in semiconductor wafer fab. We present shop floor control schemes where the status of bottleneck stations is a key factor of the control decisions. Simulation experiments have been carried out to evaluate the performance of the various control schemes. Through the simulation experiments, we have come up with some interesting findings. Most existing studies on dispatching decisions focus on bottleneck machines because the bottleneck machines determine the capacity of the whole system. Hence, it is a common idea that the bottleneck machines are scheduled in detail and managed with great attention while non-bottleneck machines are controlled in a simple way, for example by FCFS control rules. Our experimental results say a different story: i.e., the control decisions at the non-bottleneck machines could be as important as the control decisions at the bottleneck machine. This is because the inappropriate control of non-bottleneck machines may lead to disruptions of the bottleneck machines such as job starvation. So, we believe that the non-bottleneck machines should be also controlled with great attention. Another finding is about the importance of the batch processors. Simulation experiments say that the control decisions at the batch processors have greater effect on the system performance than those of the discrete processing stations. The characteristics of batch processors such as long processing time and non-uniform flow are the reasons that the batch processors have a great impact on the system performance. Another finding is about the relationship of lot releasing and dispatching decisions. Most existing research

works insist that the system-level lot release decision has more effects on the system performance than station-level dispatching decisions. However, in our simulation results, dispatching rules give no less impact on system performance than lot release rules. The authors think that the contradictory results are because of the long flow steps with which the impact of the job release is dissipated as the jobs flow downstream.

This paper has some rooms to be improved with additional works. The issues about sequence-dependent setups, multiple product types with due dates, make-to-order case may be considered in the possible future works. In addition, the experiments with different system configurations with different sizes are recommended to have additional insights. Finally, additional research works are called for to have ideas about how the control schemes including lot releasing and dispatching should be harmonized in wafer fabs in different manufacturing situations.

ACKNOWLEDGEMENT

This work was supported by a Research Grant of Pukyong National University (2016 year)

REFERENCES

- [1] Bahaji N, Kuhl ME. A simulation study of new multi-objective composite dispatching rules, CONWIP, and push lot release in semiconductor fabrication, *International Journal of Production Research*, 2008; 56(14): 3801-3824.
- [2] Glassey CR, Resende MGC. Closed-loop job release control for VLSI circuit manufacturing, *IEEE Transactions on Semiconductor Manufacturing*, 2008; 1(1): 36-46.
- [3] Goldratt EM., and Cox J. *The Goal: A process of ongoing improvement*, The North River Press: MA, USA.
- [4] Kim YD, Lee DH, Kim JU. A simulation study on lot release control, mask scheduling, and batch scheduling in semiconductor wafer fabrication facilities, *Journal of Manufacturing Systems*, 1998; 17(2): 107-117.
- [5] Koo PH, Moon DH. A review on control strategies of batch processing machines in semiconductor manufacturing, *IFAC-PapersOnLine, Series Title: Manufacturing Modelling, Management, and Control*, 2014; 7(1): 1690-1695.
- [6] Lee YH, Park, JW, Kim SY. Experimental study on input and bottleneck scheduling for a semiconductor fabrication line, *IIE Transactions*, 2002; 34: 179-190.
- [7] Li Y, Jiang Z, Jia W. An integrated release and dispatch policy for semiconductor wafer fabrication, *International Journal of Production Research*, 2014; 52(8): 2275-2292.
- [8] Monch L, Fowler J, Mason S. *Production Planning and Control for Semiconductor Wafer Fabrication Facilities: Modeling, Analysis and Systems*, New York, Springer, 2013.
- [9] Qi C, Sivakumar AI, Gershwin SB. An effective new job release control methodology, *International Journal of Production Research*, 2009; 47(3), 703-731.
- [10] Sarin SC, Varadarajan A. Wang L. A survey of dispatching rules for operational control in wafer fabrication, *Production Planning and Control*, 2011; 22:1, 4-24.
- [11] Spearman ML, Woodruff DL, Hopp WJ. CONWIP: a pull alternative to Kanban, *International Journal of Production Research*, 1990; 28(5): 879-894.
- [12] Van der Zee D. Look-ahead strategies for controlling batch operations in industry-an overview, *Proceedings of Winter Simulation Conference*, 2003; 1480-1487.
- [13] Wein LW. Scheduling semiconductor wafer fabrication, *IEEE Transactions on Semiconductor Manufacturing*, 1998; 1(3): 115-130.
- [14] Zhou Z, Rose O. A bottleneck detection and dynamic dispatching strategy for semiconductor wafer fabrication facilities, *Proceedings of the 2009 Winter Simulation Conference*, 2009; 1646-1656.