

This article was downloaded by: [EBSCOHost EJS Content Distribution - Superseded by 916427733]

On: 9 February 2010

Access details: Access Details: [subscription number 911724993]

Publisher Taylor & Francis

Informa Ltd Registered in England and Wales Registered Number: 1072954 Registered office: Mortimer House, 37-41 Mortimer Street, London W1T 3JH, UK



## International Journal of Production Research

Publication details, including instructions for authors and subscription information:

<http://www.informaworld.com/smpp/title~content=t713696255>

### Controlling a re-entrant manufacturing line via the push-pull point

Dominique Perdaen <sup>a</sup>; Dieter Armbruster <sup>ab</sup>; Karl Kempf <sup>c</sup>; Erjen Lefeber <sup>a</sup>

<sup>a</sup> Department of Mechanical Engineering, Eindhoven University of Technology, NL-5600 MB

Eindhoven, The Netherlands <sup>b</sup> Department of Mathematics and Statistics, Arizona State University,

Tempe, AZ 85287-1804, USA <sup>c</sup> Decision Technologies, Technology and Manufacturing Group, Intel

Corporation, Chandler, AZ 85226, USA

**To cite this Article** Perdaen, Dominique, Armbruster, Dieter, Kempf, Karl and Lefeber, Erjen(2008) 'Controlling a re-entrant manufacturing line via the push-pull point', International Journal of Production Research, 46: 16, 4521 – 4536

**To link to this Article:** DOI: 10.1080/00207540701258051

**URL:** <http://dx.doi.org/10.1080/00207540701258051>

PLEASE SCROLL DOWN FOR ARTICLE

Full terms and conditions of use: <http://www.informaworld.com/terms-and-conditions-of-access.pdf>

This article may be used for research, teaching and private study purposes. Any substantial or systematic reproduction, re-distribution, re-selling, loan or sub-licensing, systematic supply or distribution in any form to anyone is expressly forbidden.

The publisher does not give any warranty express or implied or make any representation that the contents will be complete or accurate or up to date. The accuracy of any instructions, formulae and drug doses should be independently verified with primary sources. The publisher shall not be liable for any loss, actions, claims, proceedings, demand or costs or damages whatsoever or howsoever caused arising directly or indirectly in connection with or arising out of the use of this material.

## Controlling a re-entrant manufacturing line via the push–pull point

DOMINIQUE PERDAEN†, DIETER ARMBRUSTER\*†‡, KARL KEMPF§ and ERJEN LEFEBER†

†Department of Mechanical Engineering, Eindhoven University of Technology,  
P.O. Box 513, NL-5600 MB Eindhoven, The Netherlands

‡Department of Mathematics and Statistics, Arizona State University,  
Tempe, AZ 85287-1804, USA

§Decision Technologies, Technology and Manufacturing Group,  
Intel Corporation, 5000 W. Chandler Boulevard, Chandler, AZ 85226, USA

*(Revision received January 2007)*

A reduced model of a re-entrant semiconductor factory exhibiting all the important features is simulated, applying a push dispatch policy at the beginning of the line and a pull dispatch policy at the end of the line. A commonly used dispatching policy that deals with short-term fluctuations in demand involves moving the transition point between both policies, the push–pull point (PPP), around. It is shown that, with a mean demand starts policy, moving the PPP by itself does not improve the performance of the production line significantly over policies that use a pure push or a pure pull dispatch policy, or a CONWIP starts policy with pure pull dispatch policy. However, when the PPP control is coupled with a CONWIP starts policy, then for high demand with high variance, the improvement becomes approximately a factor of 4. The unexpected success of a PPP policy with CONWIP is explained using concepts from fluid dynamics that predict that this policy will not work for perishable demand. The prediction is verified through additional simulations.

*Keywords:* Re-entrant production; CONWIP; Dispatch policy

### 1. Introduction

A very important feature of the production of semiconductor wafers is the re-entrant line: wafers are produced in layers and hence after one layer is finished a wafer returns to the same set of machines for processing of the next layer. Modern semiconductors may have on the order of 20 to 30 such layers. It is typical for wafers to spend several weeks in such a re-entrant production line, much of the time waiting for available machines. Process control in such long production lines with thousands of wafers and hundreds of processing steps making tens of different products is a special challenge. Most of the time the demand fluctuates on a much faster timescale

---

\*Corresponding author. Email: armbruster@asu.edu

than the factory cycle time, making it very difficult to use starts policies to react to the demand fluctuations. Typically, for a product with a constant mean demand, the mean demand is started. Due to stochasticity in the production and due to variation in the demand there is nevertheless a large mismatch in daily outputs and demand. In practice, to reduce the mismatch, production targets over a certain time horizon are given and wafers at the end of the production process are sped up or slowed down using dispatch policies. We are not concerned here with longer and larger fluctuations that might require an adjustment of the starting rate to cover changes of the desired WIP level as discussed by Sterman (2000).

The combination of lot release and dispatching strategies is called Workload (or Flow) Control. An overview of state-of-the-art published research on workload control as applied to semiconductor industry is provided by Fowler *et al.* (2002). A thorough overview of the literature on order release as a flow control is provided by Bergamaschi *et al.* (1997), whereas Panwalkar and Iskander (1977) and Blackstone *et al.* (1982) are two thorough surveys of the dispatching literature. Commonly used dispatching policies include: First-In, First-Out (FIFO), Earliest Due Date (EDD), Weighted Shortest Processing Time (WSPT), Least Slack (LS), and Least Setup Cost (LSC). In the seminal paper (Wein 1988) many of these lot sequencing rules as well as a variety of input controls were evaluated using simulation models of representative but fictitious semiconductor fabs. The main conclusion was that order release is more important than dispatching (30–40% change versus less than 10%), although there is an important connection between these decisions. Dynamic scheduling studies were carried out by Aytug *et al.* (1994), who implemented learning of dispatch rules in their simulation environment. Pure push and pull dispatch policies were studied by Atherton and Dayhoff (1986).

Most of the time, demand fluctuates on a much faster timescale than the factory cycle time. Unfortunately, almost no literature exists on how to deal with the impact of a production surge or short-term increase in wafer starts that occurs when unexpected orders are received by a fab that is operating close to its designed capacity. Preliminary investigations into the surge problem have been performed by McKiddie (1995), Kato (1996) and Dummler (2000).

In order to deal with these short-term variations in demand we consider a dispatching policy which to the authors' knowledge has not been considered in the literature before, but which is used in practice. We simulate a reduced model of a re-entrant semiconductor factory exhibiting all the important features, applying a push (dispatch) policy at the beginning of the line and a pull (dispatch) policy at the end of the line. Here a push (pull) policy refers to the fact that a machine that is able to process more than one step gives priority to the earlier (later) step. Push policies are also known as first buffer first served and pull policies are known as shortest-expected-remaining-process-time policies. We use a push policy upstream and a pull policy downstream. The step at which we switch from a push to a pull policy is called the push–pull point (PPP). Its dynamics is the control variable. Our objective (metric) is to reduce the mismatch between daily outputs and demand over a long time interval. We assume that over that time interval the demand has a constant mean demand and varies stochastically around the mean. By focussing on the output, this study complements the important work by Lu *et al.* (1994), who were not concerned with output but with the behaviour of the mean and variance of the cycle times as a function of different scheduling policies.

We show that with a policy that starts the mean demand, moving the PPP by itself does not improve the performance of the production line significantly over a pure push, a pure pull policy, or a pure CONWIP starts policy (Spearman *et al.* 1994) with pure pull dispatch. However, when the PPP dispatch control is coupled with a CONWIP starts policy, then for high demand with high variance, the improvement becomes approximately a factor of 4. We explain the unexpected success of a PPP policy with CONWIP using concepts from fluid dynamics that predict that this policy will not work for perishable demand. We verify this prediction.

## 2. The factory model

Our basic factory model consists of 26 production steps executed on nine machine sets. Table 1 contains all the specifications of this model. The first six machines are called diff 1, diff 2, litho 1, etch clean, etch 1, and ion impl, corresponding to production steps associated with diffusion, photolithography, etching and ion implantation, respectively. They are associated with the transistor section of the production line and a wafer performs four loops through these machines in a specific order as indicated in table 1. The last three machine sets are called metal dep, litho 2 and etch 2, corresponding to production steps that generate metal layers for interconnection of the transistors. The wafer loops through the metallization section of the production line twice. The transistor and metal loops are completely disjoint and do not share equipment. Rows 1–26 of table 1 correspond to the 26 production steps. The entries in each row indicate the machine set that performs the step and the processing time spent in a machine in the set. For instance, step 3, 6, 10 and 14 are all performed on the photolithography machine litho1 with cycle times of 1, 1.25, 1 and 1.25 h, respectively.

The second part of table 1 is a spreadsheet calculation to determine the required number of machines (tools) to have a production target of 200 lots per week, given availability rates of the machines and desired levels of constraints for a given machine set. Consider, for instance, the last eight rows in the column litho1: a wafer spends a total of 4.5 h in litho1. Hence to produce 200 wafers per week we need 900 h per week of machine time. Assuming that a litho1 machine is 90% available and a work week of 168 h, this machine works for 151.2 h per week and hence we need 5.95 machines of that type. Since this is a very expensive machine, it is planned to be the bottleneck and hence has a constraint factor of 1.0. As a result six machines will be installed. Taking into account that the diffusion machines batch four wafers per machine cycle we reach the installation targets in the last row in a similar way for all columns.

This model is implemented as a discrete event simulation in  $\chi$  (Hofkamp and Rooda 2002, Vervoort and Rooda 2003), a specification language developed at the Eindhoven University of Technology. Stochasticity enters the simulation at various levels: the time that a machine is in service, and the time that it is not, is distributed by a Weibull distribution (Hofkamp and Rooda 2002) with a mean 'in service' time of 10 process times and a variance of 50%. The demand is randomly generated and is fixed for a simulation.

Table 1. Factory model.

Step	diff 1	diff 2	litho 1	etch clean	etch 1	ion impl	metal dep	litho 2	etch 2	Station
	0	1	2	3	4	5	6	7	8	
1				0.25						Clean wafer
2	8.00									Grow a layer
3			1.00		1.00					Pattern it
4										Etch some away
5		6.00								Grow a layer
6			1.25							Pattern it
7						2.50				Implant ions
8				0.50						Remove mask
9	7.00									Grow a layer
10			1.00							Pattern it
11					1.00					Etch some away
12				0.25						Clean wafer
13		5.00								Grow a layer
14			1.25							Pattern it
15						3.50				Implant ions
16				0.50						Remove mask



The actual processing times are pulled out of an exponential distribution (Hofkamp and Rooda 2002) with the mean equal to the process times in table 1. Note that, while the raw processing times of semiconductor processing machines are narrowly distributed, the unloading of machines depends on the availability of human operators and is highly variable. Nevertheless, using an exponential distribution probably constitutes a worst case scenario for a practical model. Overall, the stochastic parameters are fixed in such a way that simulations of the model generate an outflux variance of 20% around the nominal influx of 200 per week, i.e. the throughput varies between 160 and 240 wafers per week.

### 3. The push–pull point algorithm

The goal of the PPP policy is to reduce the mismatch between fluctuating demands and the stochastically varying outflux of the factory. This policy divides the production line into two parts. Upstream of the PPP, priorities are assigned using a push strategy, downstream they are assigned according to a pull strategy. In conflicts across the PPP we always give priority to the steps in the pull part. Figure 1 shows a typical priority assignment.

The PPP is moved depending on the demand: given a demand period and a distribution of the work in progress (WIP) over the queues of all production steps

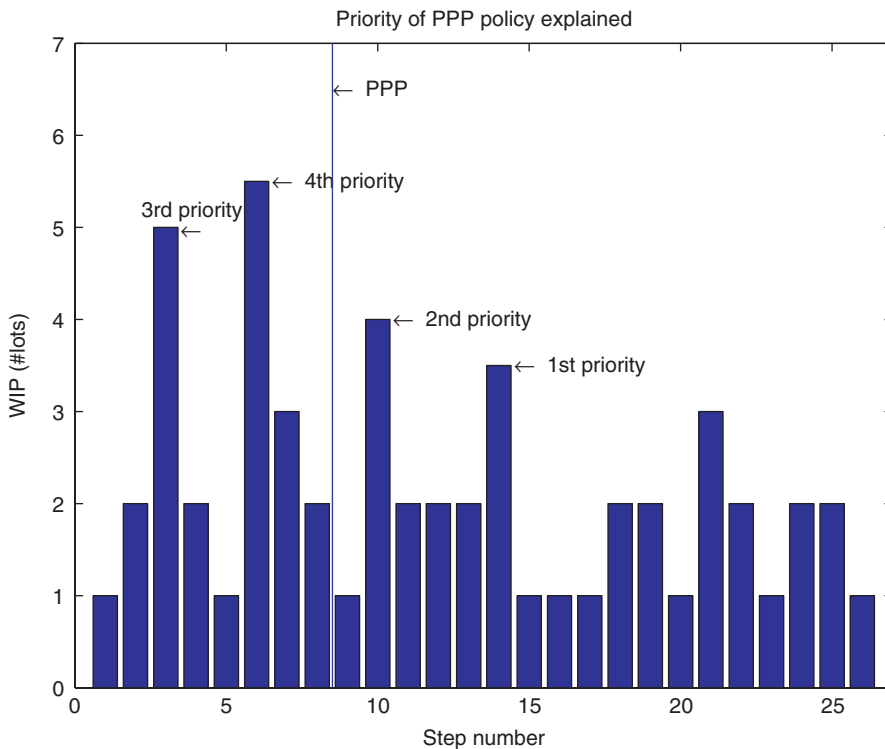


Figure 1. Priority distribution when the PPP policy is used.

(the WIP profile), we place the PPP at such a point that the WIP downstream from the PPP is equal to the demand in the chosen demand period. When the demand increases, more products have to be pulled out of the line, moving the PPP upstream. When the demand decreases, the PPP will shift downstream.

The possible success of such a strategy is based on three important facts.

- The clearing function (Asmundsson *et al.* 2003), i.e. the throughput as a function of the load in the factory in steady state, is significantly higher for a production line run completely with a push dispatch policy than for one run completely with a pull dispatch policy. Hence by increasing or decreasing the part of the production line that is run in pull policy we should temporarily increase or decrease the outflux. We show below the details of this effect for our model production line.
- The location of the push–pull point determines the average shape of the WIP profile in steady state. In particular, on average, WIP decreases in the queues downstream of the PPP and increases upstream from the PPP. Figure 1 shows this schematically for the queues in front of the photolithography machines for a fixed PPP point. Figure 2 shows that this is true to a large

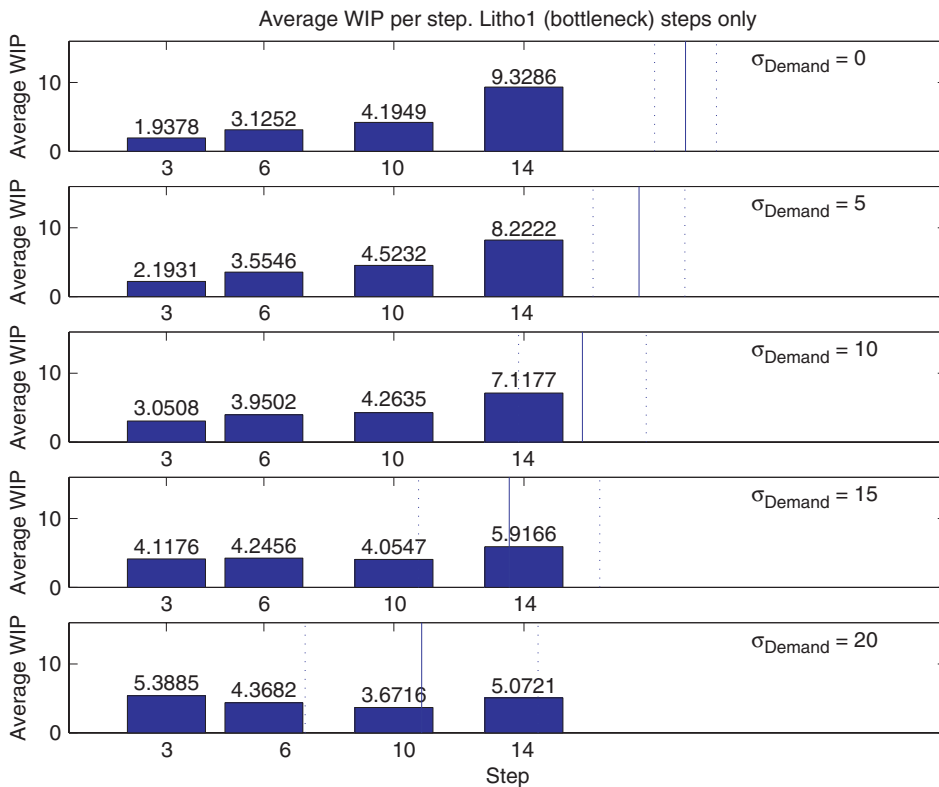


Figure 2. Average queue length at the litho 1 steps. The vertical line and its two dashed side bars are the average position of the PPP plus/minus  $2\sigma$ . As the PPP point moves upstream the WIP in the last two photolithography steps decreases and the WIP in the first two photolithography steps increases.

extent for simulations on average, even when the PPP point is dynamically moved.

- The cycle time through the factory and the time between readjustments of the PPP have to be related. In particular, if adjusting the PPP according to demand on average places the PPP approximately in the middle of the production line, adjusting to higher and lower demand by changing the PPP should be feasible.

#### 4. Results

To determine the effectiveness of the PPP strategy we compare it with simulations with a starts policy of the mean demand and dispatch policies of pure push, pure pull as well as a CONWIP starts strategy using a dispatch policy of pure pull. We have also combined the PPP strategy with CONWIP as a starts policy. In all simulations we employ a FIFO policy within a given queue for a given production step. We run 500 simulations per data point. The demand  $d(t)$  for each simulation is generated independently by choosing a demand for a 2 day period out of a normal distribution (throwing away the rare events that gave negative demands) with an average of 180 lots per week. The demand is not perishable, which means that the backlog or the inventory of the previous demand period is taken into account for the present demand period. The PPP is adjusted every 2 days (one demand period). Since the cycle time for our simulation factory is on the order of 5 days, the 2 day readjustment time places the PPP well inside the production line. The simulation time for every single run is 144 weeks. The different control strategies are compared using the absolute value of the mismatch between output and demand over each demand period. Mismatch  $m(t)$  and costs are given as

$$m(0) = 0, \quad (1)$$

$$m(i) = m(i - 1) + d(i) - o(i), \quad (2)$$

$$\text{cost}(t) = \sum_i^t |m(i)|. \quad (3)$$

Here,  $o(t)$  is the output of the factory plus backlog and storage, i.e. over and under-production cost the same \$1 per lot per demand interval (2 days).

Figure 3 shows the average costs over 500 simulations as a function of the variance in the demand for all the different strategies. Table 2 shows the variances for the nine simulation points in figure 3.

The results are surprising: pure push, pure pull, regular PPP (all with mean demand starts policy) and a CONWIP starts policy (pure pull dispatch policy) with a WIP level of 119 lots all increase monotonically with the demand variation and have very similar average cost. In contrast, a policy that combines the starts policy of a CONWIP rule and a WIP of 150 lots with the PPP control policy has almost constant costs over a wide range of demand variations. In addition, the costs for high demand variations are significantly lower for the PPP with CONWIP than for the other policies—\$50 vs. more than \$200.

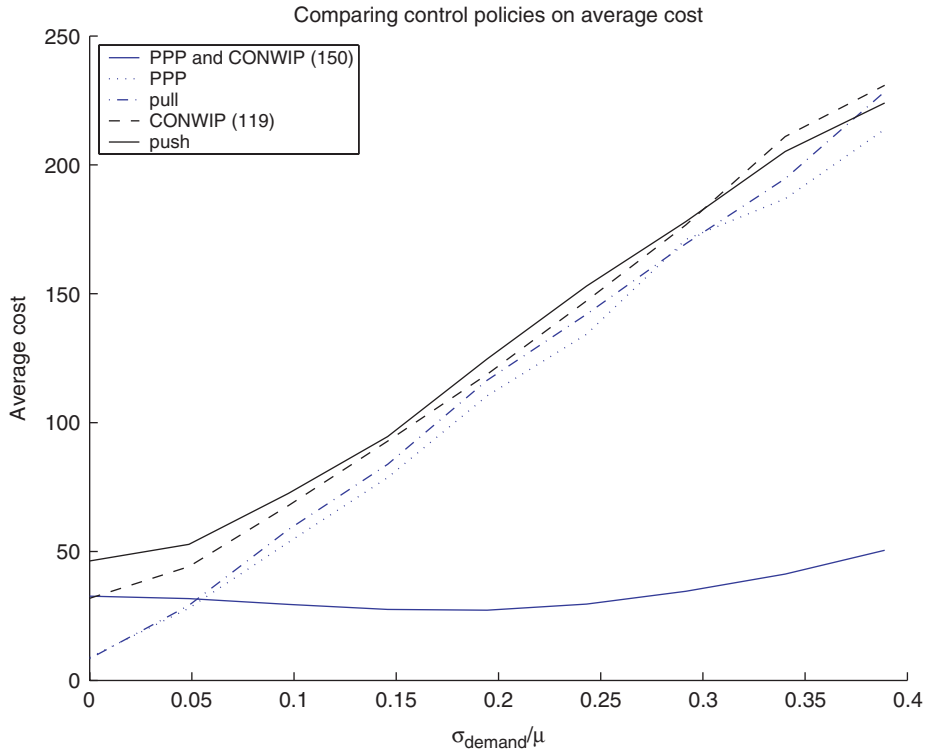


Figure 3. Average costs per simulation for different control strategies as a function of the coefficient of variation of the demand.

Table 2. Variance of cost as a function of the variation of demand.

$\sigma_{\text{demand}}/\mu$	0	0.05	0.1	0.15	0.2	0.25	0.3	0.35	0.4
$\sigma_{\text{cost}}^2$	3.0	3.7	6.6	7.9	9.7	27.8	75.1	216.6	578.1

## 5. Analysis of the PPP–CONWIP policy

Figure 4 begins to explain the success of the PPP–CONWIP policies. It shows the clearing functions for CONWIP policies with different fixed push–pull points. The curve indicated by  $ppp = 0$ , corresponding to a pure pull dispatch policy, gives the highest throughput of all possible policies. The curve labeled  $ppp = 27$  is a pure push dispatch policy that gives the lowest throughput of all. The intermediate curves indicated by  $ppp = x$  denote a dispatch policy where the push–pull point has been fixed at step  $x$ . Note that, for a complete push policy, the throughput actually decreases with an increase in WIP. This is the result of the interplay between the back-loaded WIP distribution of the push policy and the batching in the diffusion steps. Figure 4 also explains the choice of a CONWIP starts policy with a WIP level of 119 lots for the pure pull dispatch policy used in figure 3: the top curve in figure 4 represents a pure pull dispatch policy. The associated WIP level in steady state for

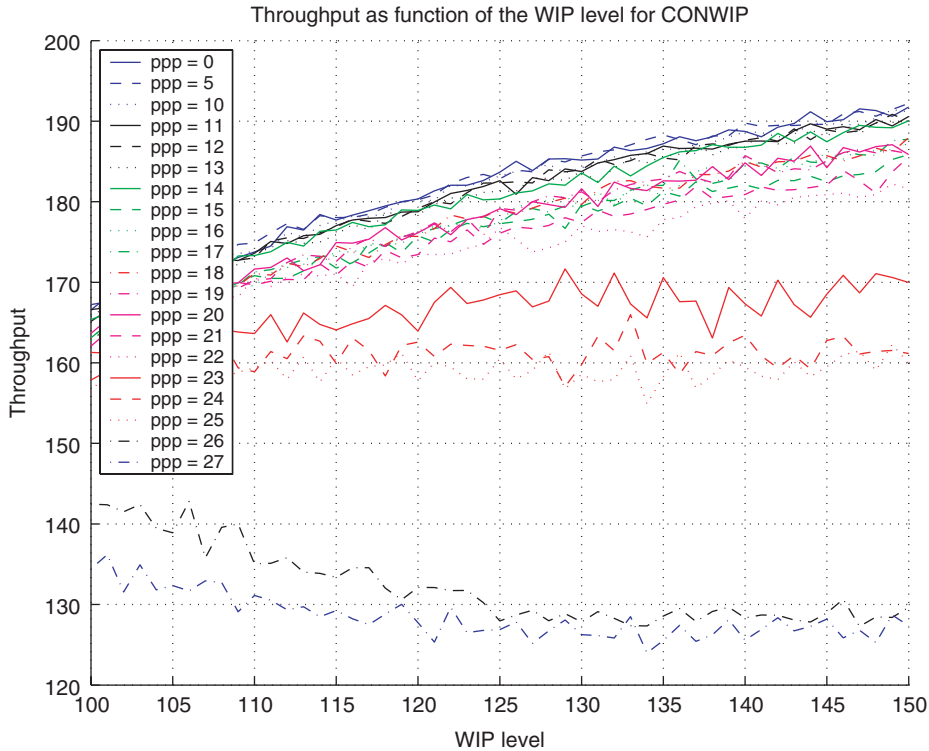


Figure 4. Throughput as a function of total WIP for CONWIP policies with fixed push-pull points.

a throughput of 180 lots/week is 119 lots, which we use as the desired WIP level (Sterman 2000).

These clearing functions suggests one reason for the success of the PPP-CONWIP policy: by using a CONWIP starts policy with a high WIP level and switching the PPP, we can change the outflux in the factory by a significant amount. For instance, for a WIP level of 150 lots we can obtain throughputs between approximately 130 and 190 per week. Note also that there is no good push-pull point for a WIP level of 150 that creates a throughput of 180 per week that we are using for our simulations. A PPP at stages 1-15 creates a throughput much higher and a PPP at stages 15-26 creates a throughput much lower than 180 per week. As a result, a completely deterministic demand cannot use a fixed PPP even though the demand is constant and hence has to jump back and forth, creating extra backlog or overproduction cost. This is the reason for the slight increase in cost for the PPP-CONWIP policy with WIP level 150 in figure 3 for low demand variation.

A different issue explains the failure of the pure PPP dispatch policy to be much better than a regular pull dispatch policy. Assume a push-pull point in the middle of the production line and an increase in demand. In response we will move the PPP upstream and clear out more of the WIP than we usually do over the demand period. However, we will only *start the average* amount. Consequently, WIP goes down and a second increase in demand will move the PPP rapidly further upstream. As a result we easily reach the point where the PPP is at the beginning of the line and the policy

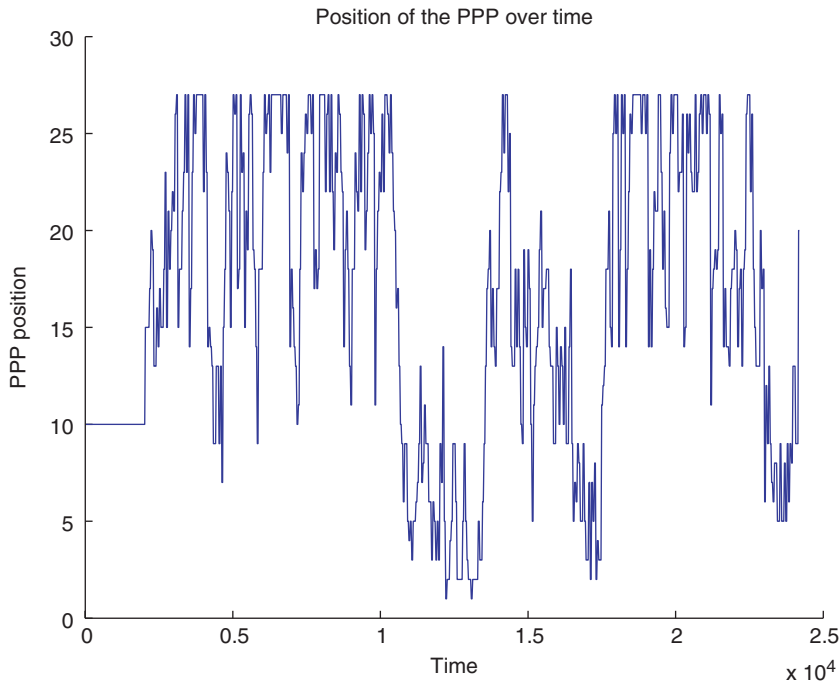


Figure 5. Time evolution of the push–pull point as a function of time for a PPP–CONWIP policy with a WIP level of 130.

becomes a pure push dispatch policy. We cannot increase the outflux further than that. Similarly, a demand signal that has several periods below average will eventually move the PPP to the end of the factory and hence constitute a pull policy. We cannot reduce the outflux further than that. A CONWIP starts policy reduces the instances where the push–pull point is at one of the extremes of the production line by instantaneously starting more when more is pulled out of the factory and starting less if more is left in the factory. Figures 5 and 6 show the position of the PPP as a function of time for a PPP–CONWIP and a free PPP policy, respectively. Clearly, the free PPP policy is locked into pure push or pure pull policies much more often than the PPP–CONWIP.

We can illustrate the difference between free PPP and PPP–CONWIP policies with the following illustration based on fluid flows. For the purpose of this illustration let us consider the average behaviour of a large number of lots as they move through the factory. We assume that the average speed  $v(t)$  of a lot for a factory that is in steady state is constant over all production steps and depends on the dispatch policy. In particular, the average cycle time for a lot under a pull (dispatch) policy is shorter than for a lot produced under a push (dispatch) policy. Hence the associated average velocity for a pull policy is higher than for a push policy. Let us consider a continuum of production steps and a continuum of lots such that we can define a WIP density  $\rho(x, t)$  that describes the density of lots at stage  $x$  at time  $t$ . Then the throughput of the factory becomes  $\lambda(x, t) = \rho(x, t)v$ . In steady state, the throughput is constant and hence we obtain a constant WIP profile  $\rho(x) = \lambda/v$  that does not depend on  $t$  because we are looking at steady state, and does not depend on  $x$ ,

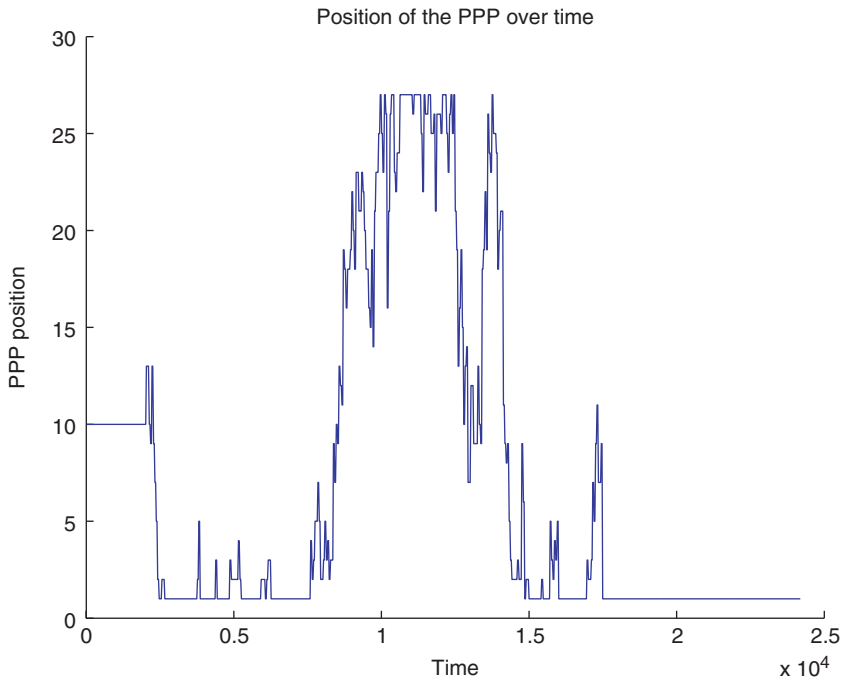


Figure 6. Time evolution of the push–pull point as a function of time for a free PPP policy.

because we assume  $v$  to be constant. This is certainly not exactly true but a good approximation for the purpose of this illustration. Now, for a PPP policy we can consider the upstream part of the production line as a homogeneous push line and the downstream part as a homogeneous pull line, each with its own constant velocity with  $v_{\text{push}} < v_{\text{pull}}$ . Since the throughput is the same everywhere and since  $\rho v = \lambda$  has to hold, we obtain a jump in the WIP profile at the push–pull point by the amount

$$\frac{\rho_{\text{push}}}{\rho_{\text{pull}}} = \frac{v_{\text{pull}}}{v_{\text{push}}}. \quad (4)$$

Figure 7(a) shows the constant throughput and the discontinuous WIP profile.

Assume we now move the PPP upstream by an amount  $\Delta x$  instantaneously. The queues that were just upstream of the PPP and hence had the lowest priority on the line now move up in priority and therefore speed up. In other words, part of the WIP profile that used to be in the push region and had a high WIP level is now in the pull region. As the velocity in the pull region is higher, the product  $\rho_{\text{push}} v_{\text{pull}} > \lambda$ , i.e. we create a flux bump. Similarly, we create a flux dip by moving the PPP downstream. The flux changes are

$$q\Delta x = \lambda \frac{v_{\text{pull}}}{v_{\text{push}}}, \quad (5)$$

$$q\Delta x = \lambda \frac{v_{\text{push}}}{v_{\text{pull}}}, \quad (6)$$

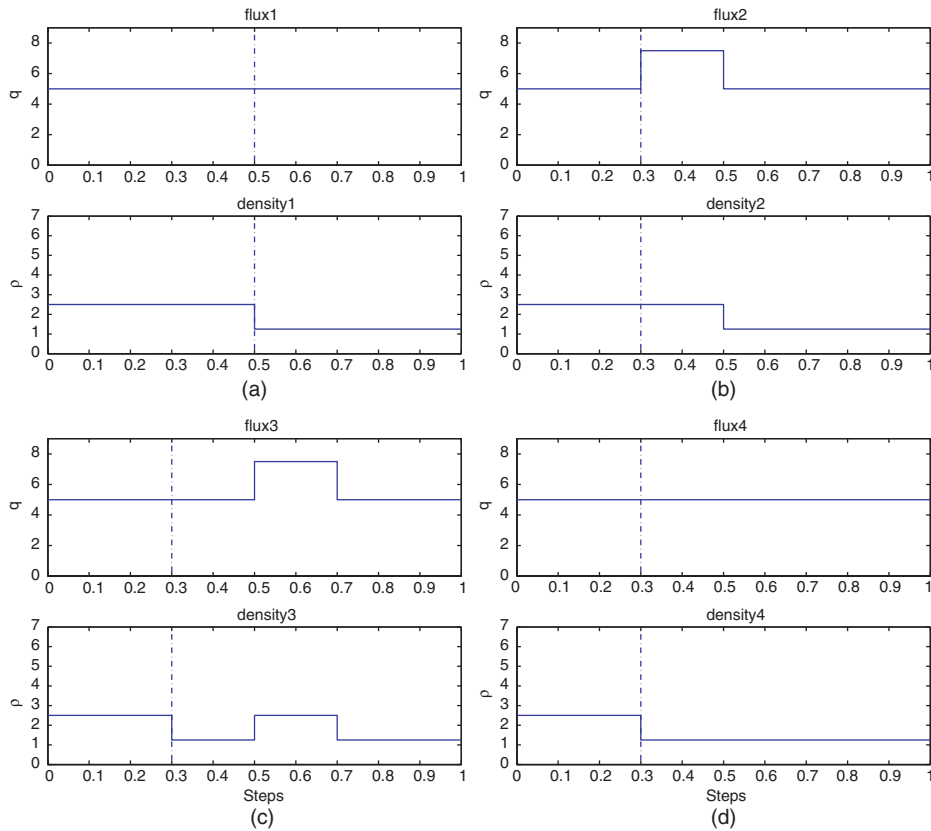


Figure 7. Stages of creating a flux bump.

for the flux bump and flux dip, respectively. Keeping the PPP at its new location the flux bump is downstream from the PPP and hence moves downstream with constant speed  $v_{\text{pull}}$  pulling a WIP bump with it until they both exit the factory. During the time they exit they will increase the outflux. Depending on the remaining processing time from the push–pull point to the end of the production line, the increase in outflux may or may not happen within the demand time interval. Figure 7(b) and (c) show this time evolution. After the WIP/flux bump has exited, the total WIP in the factory is lower and hence in order to satisfy the same demand, the push–pull point will have to move yet further upstream driving it towards the beginning of the factory.

In contrast, the time evolution of the flux bump for the PPP–CONWIP policy is illustrated in figure 8.

As the CONWIP starts policy is implemented by matching the starts to the outflux, once the WIP bump moves out of the factory, the starts will be increased to create a new WIP bump. In that way, the total throughput will stay high until the PPP point is moved downstream again. That will happen when the backlog has moved to zero and the sum of actual backlog and actual demand has decreased. In that way we have a policy that reverts all the time to a match between demand and outflux. This explanation can be checked by running the simulation with a perishable

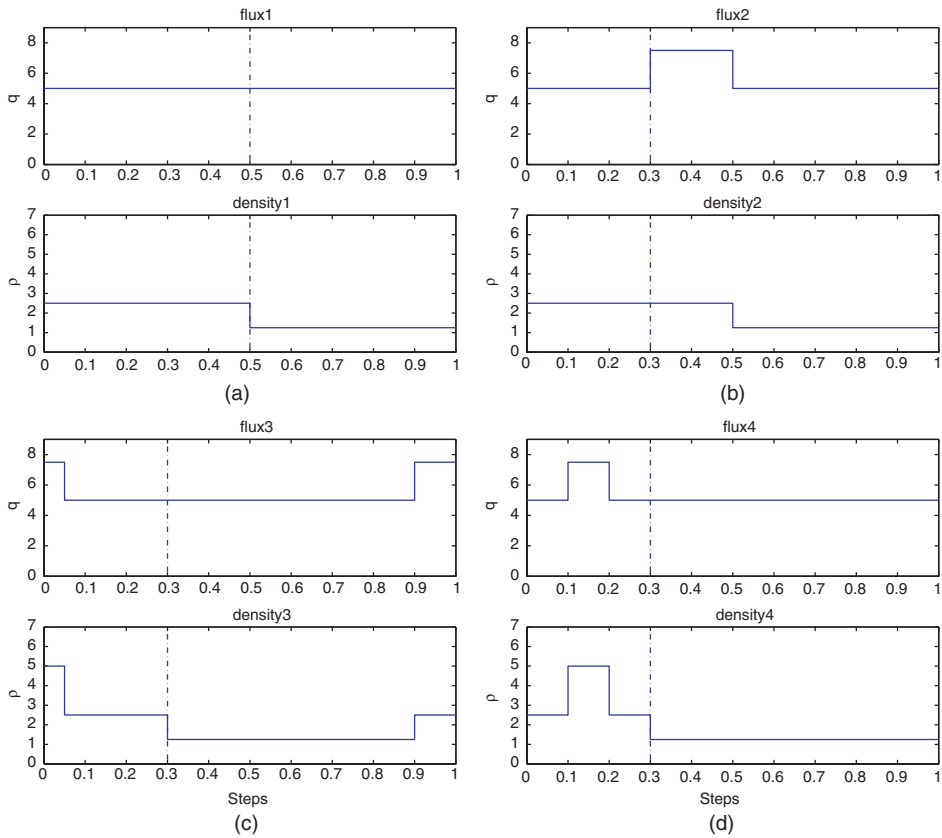


Figure 8. Stages of creating a flux bump for a PPP-CONWIP policy.

demand protocol: we only register whether there is a mismatch of the current outflux and the current demand but do not try to make up for that mismatch in the next time interval. For such a model the PPP-CONWIP policy should not be better than the free PPP policy. The only thing that matters is whether the flux bump or flux dip that is created arrives at the end of the factory within the demand time window. Our simulations confirm this: PPP and PPP-CONWIP policies behave very similarly and do not improve the performance of the production line appreciably with perishable demand.

## 6. Conclusion

We have studied process control in a reduced model of a re-entrant semiconductor factory using discrete event simulations. We have shown that when running a factory with a push dispatch policy at the beginning of the factory and a pull dispatch policy at the end of the factory while using an average demand starts policy, the transition point (the PPP) can be used to reduce the mismatch between stochastic outfluxes of the factory and stochastic demands.

We have two results that are of immediate practical interest.

1. A pure PPP dispatch policy that reaches into the factory from the end and pulls out the desired demand will not significantly reduce the mismatch between outflux and demand for a demand signal that has a constant average and varies stochastically around that average.
2. A PPP dispatch policy coupled with a CONWIP starts policy adjusted for a WIP level that allows maximal flux changes through moving the PPP will significantly reduce the mismatch for a production with non-perishable demand.

Process control in these re-entrant production lines is very difficult since only starts policies and dispatch rules are obvious control actuators that influence the outflux of the factory. However, as a byproduct of this study we have identified another control parameter: the actual WIP profile will be very important for the success of a PPP policy. It seems likely that very homogeneous WIP profiles are better for the control action of the PPP policy than the WIP profile that we have currently examined. Those WIP profiles are determined by the level of constraint we choose for a particular machine set. It will be an interesting further study to determine the interplay between the constraint levels and the PPP policy.

### Acknowledgements

The research of D.A. was supported by NSF grants DMS-0604986 and DMS-0204543. We thank Ton Geubbels for help in developing the  $\chi$ -model.

### References

- Asmundsson, J., Rardin, R.L. and Uzsoy, R., Tractable non-linear capacity models for aggregate production planning. Technical Report IN 47907-1287, Purdue University, School of Industrial Engineering, West Lafayette, 2003.
- Atherton, R.W. and Dayhoff, J.E., Signature analysis: simulation of inventory, cycle time and throughput tradeoffs in wafer fabrication. *IEEE Trans. Compon. Hybrids Mfg Technol.*, 1986, **9**, 498–507.
- Aytug, H., Koehler, G.J. and Snowdon, J.L., Genetic learning of dynamic scheduling within a simulation environment. *Comput. Oper. Res.*, 1994, **21**, 909–925.
- Bergamaschi, D., Cigolini, R., Perona, M. and Portioli, A., Order review and release strategies in a job shop environment: a review and classification. *Int. J. Prod. Res.*, 1997, **2**, 399–420.
- Blackstone, J.H., Phillips, D.T. and Hogg, G.L., A state-of-the-art survey of dispatching rules for job shop operations. *Int. J. Prod. Res.*, 1982, **1**, 27–45.
- Dummler, M.A., Analysis of the instationary behavior of a wafer fab during product mix changes, in *Winter Simulation Conference*, 2000, pp. 1436–1442.
- Fowler, J.W., Hogg, G.L. and Mason, S.J., Workload control in the semiconductor industry. *Prod. Plann. Contr.*, 2002, **13**, 568–578.
- Hofkamp, A.T. and Rooda, J.E.,  $\chi$  Reference Manual. Eindhoven University of Technology, Systems Engineering Group, Department of Mechanical Engineering, 2002. URL: <http://se.wtb.tue.nl/documentation>
- Kato, K., Lot start evaluation system using simulator, in *International Symposium on Semiconductor Manufacturing*, 1996, pp. 293–296.

- Lu, S., Ramaswamy, D. and Kumar, P.R., Efficient scheduling policies to reduce mean and variance of cycle time in semiconductor plants. *IEEE Trans. Semicond. Mfg.*, 1994, **7**, 374–388.
- McKiddie, R., Some ‘no-panic’ help for wafer-start surges. *Semicond. Int.*, 1995, 115–120.
- Panwalkar, S.S. and Iskander, W., Capacity planning and control. *Oper. Res.*, 1977, **1**, 45–61.
- Spearman, M.L., Woodruff, D.L. and Hopp, W.J., CONWIP: a pull alternative to kanban. *Int. J. Prod. Res.*, 1994, **28**, 879–894.
- Sterman, J.D., *Business Dynamics, Systems Thinking and Modeling for a Complex World*, 2000 (McGraw-Hill: New York).
- Vervoort, J. and Rooda, J.E., Learning  $\chi$ . Technical Report, Eindhoven University of Technology, Systems Engineering Group, Department of Mechanical Engineering, 2003. URL: <http://se.wtb.tue.nl/documentation>
- Wein, L.M., Scheduling semiconductor wafer fabrication. *IEEE Trans. Semicond. Mfg.*, 1988, **1**, 115–129.