

A continuum model for a re-entrant factory

Dieter Armbruster, Dan Marthaler, Christian Ringhofer,
Department of Mathematics, Arizona State University,
Karl Kempf, Decision Technologies, Intel Corporation, and
Tae-Chang Jo, Mathematics Department
New Mexico Institute of Mining and Technology

March 23, 2005

Abstract

High volume, multi stage continuous production flow through a re-entrant factory is modeled through a conservation law for a continuous density variable on a continuous production line augmented by a state equation for the speed of the production along the production line. The resulting nonlinear nonlocal hyperbolic conservation law allows fast and accurate simulations. Little's law is built into the model. It is argued that the state equation for a re-entrant factory should be nonlinear. Comparisons of simulations of the PDE model and discrete event simulations are presented. A general analysis of the model shows that for any nonlinear state equation there exist two steady states of production below a critical start rate: A high volume - high throughput time state and a low volume - low throughput time state. The stability of the low volume state is proved. Output is controlled by adjusting the start rate to a changed demand rate. Two linear factories and a re-entrant factory, each one modeled by a hyperbolic conservation law, are linked to provide proof of concept for efficient supply chain simulations. Instantaneous density and flux through the supply chain as well as Work In Progress (WIP) and output as a function of time are presented. Extensions to include multiple product flows and preference rules for products and dispatch rules for re-entrant choices are discussed.

1 Introduction

In recent years fast scalable simulations of production flows in a supply chain have become an important research topic [34]. While the long term goal is to optimize production across the whole supply chain, an intermediate goal is to generate simulation tools that support the exploration of business questions and to pose what if? questions on these simulations. Since most production deals with individual parts and the processes that these parts undergo, the natural method of choice for accurate simulations is discrete event simulators.

While they have been very successful on the factory level [12], e.g. to simulate semiconductor production lines, they are relatively slow, and it seems obvious that they are not scalable to a full supply chain.

This paper proposes to develop a new approach to the simulation of production flow in analogy to traffic flow. Specifically, we will model the dynamics of material flow through a factory via a hyperbolic conservation law. The main variable that is modeled will be the density of product in a factory. The specifics of the production process will enter into a state equation relating the velocity of the product to the density of the material in the factory. The focus in our model is on re-entrant factories because they represent the most important part of arguably the most efficient current production systems, the production of semiconductor devices. However, we maintain that our approach can be tailored to many other production strategies and processes. The major advantage of hyperbolic conservation law models is the existence of a large body of numerical codes to solve such equations efficiently.

The rest of the paper is organized in the following way: Section 2 discusses the basic model and its relationship to traffic modeling. The shape of the state equation for acyclic and re-entrant flows is discussed. It shows that a hyperbolic conservation law can be written that corresponds to Little's law. Section 3 briefly discusses the numerical method to simulate our re-entrant factory model. We discuss the applicability and limitations of our model and compare it to several discrete event simulations in Section 4. Section 5 discusses the existence of multiple equilibria for a nonlinear state equation with constant start rates into a factory. There exist two equilibria: an equilibrium where a high product density moves slowly and another one where a low product density moves fast such that the product of speed times density gives the same production rate in both cases. We will show that only the low density equilibrium is stable under perturbations. Section 6 deals with the issue of controlling the output of the factory by adjusting the start rate to a changed demand rate. Section 7 shows proof of concept for the original goal of fast supply chain simulations: We link three nodes, two linear factories and a re-entrant factory each one modeled by a hyperbolic conservation law. The model of the linear factories has been discussed in [2]. Instantaneous density as well as flux through the supply chain for several scenarios is shown. Work In Progress (WIP) and output as a function of time are presented. We conclude with a discussion of the influence of dispatch rules and an outlook on attempts to derive the continuum models from first principles.

2 The Model

There are currently three major approaches to simulate production flows: Discrete Event Simulations (DES), fluid networks and conventional queueing network models. DES has successfully been used in large simulations for semiconductor factories (for instance [7], [12]) but typically is very time consuming. Fluid models [9], [23] come from traffic theory and were introduced by Newell

[32, 33] to approximately solve queueing problems. They consider the length of a queue $q(t)$ as a continuous variable whose rate of change is given by:

$$\frac{dq}{dt} = \begin{cases} \lambda(t) - \mu(t) & \text{for } q(t) \neq 0 \\ 0 & \text{for } q(t) = 0, \end{cases} \quad (1)$$

where $\lambda(t)$ is the arrival rate and $\mu(t)$ the processing rate of the queue. This basic building block for a queue can be connected to a *work-conserving fluid model* by feeding the outflux of each queue into other queues. Dai and Weiss [9] have analyzed the relationship between the stability of the fluid model and the stability of scheduling policies for the associated queueing networks. Here stability of the fluid model is represented by the boundedness of the fluid variables for a given influx $\lambda(t)$ which is assumed to be less than the smallest processing rate $\mu_i(t)$ of a queue in the network. Stability in the queueing theory sense is given by a unique stationary distribution ψ for the underlying stochastic process describing the queueing network. Dai and Weiss [9] in particular showed that a queueing discipline is stable if the corresponding fluid model is stable. Queueing network models have successfully been used for instance to develop input and priority sequencing policies in two station networks [39].

Fluid models have several distinct shortcomings:

- They do not model stochasticity very well. Either $\lambda(t)$ and $\mu(t)$ are mean rates in which case Equation (1) is a fully deterministic system and stochasticity is not modeled at all or, $\lambda(t)$ and $\mu(t)$ are stochastic processes which allows to do some theoretical analysis but drastically diminishes the advantages of a continuum model as a simulation tool.
- Assuming constant production rates, fluid models are too rigid. If they are stable there will never be any WIP waiting in a queue and if they are unstable the queues will explode. Similarly, fluctuations in one queue can never travel downstream (or upstream) to the next queue unless the queue actually becomes zero.
- While a fluid network models the work in progress with a continuous variable, it models machines as individual discrete queues. With several hundred production steps for a typical chip it is reasonable to approximate the production steps also along a continuum.

Systems dynamics [13] is related in spirit to a fluid model. It will derive ordinary differential equations that model the flow from one machine to another and analyse the resulting large system of ODEs. However, it is unclear how the correct throughput time will be incorporated in such a system. Either the ODEs will become delay equations (which makes them much harder to solve) or intermediate states will have to be introduced [18]. In any case, the changes in throughput time due to increased load in a factory will not be included.

Queueing network models have their own limitations:

- While they can deal with several re-entrant steps by introducing different customer classes, their complexity explodes as the number of machines increases.

- They solve control type problems for a queueing network in steady state. I.e. all their result are valid in the limit as $t \rightarrow \infty$.
- They have to ignore all incidence in the factories that cannot be modeled by queueing type of behavior (see below for examples).

In a modern semiconductor factory we are interested to model and simulate on the order of 250 production steps executed on about 100 machines, with a re-entrant part of the production line that cycles about 15- 20 times. In addition, the lifecycle of a product is of the order of one year, whereas the throughput time lies between 40 and 60 days. Hence it is unlikely that the factory is ever run for any longer amount of time in steady state and we are especially interested in *transient* behavior of such systems. A model that addresses these issues can be derived from traffic modelling: The simplest model for traffic on a one lane highway without on and off ramps is a hyperbolic conservation law of the form

$$\frac{\partial \rho}{\partial t} + \frac{\partial(v(\rho)\rho)}{\partial x} = 0 \quad (2)$$

where subscripts denote partial differentiation. The velocity $v(\rho)$ is described by a state equation relating vehicle velocity v and vehicle density ρ through [27, 17]

$$v(\rho(x, t)) = v_0 \left(1 - \frac{\rho(x, t)}{\rho_m}\right) \quad (3)$$

with ρ_m the capacity of the highway. This is called the Lighthill-Whitham model and is the first step in a hierachy of traffic models [17].

We propose a full continuum model for the production flow through a re-entrant factory with a similar structure: Let x be a continous variable representing completion of the product, i.e. product at $x = 0$ denotes a raw product and parts at $x = 1$ denote a finished product. While $x = 0$ corresponds to the entry into the factory and $x = 1$ corresponds to the exit from the factory, there is no one to one correspondence between factory floor space and x . However, there is a unique production process assigned to every x -value. Assuming a high volume, many stage factory we model production flow with a continuum variable on a continuum domain. We write $u(x, t)$ for the density of product at stage x and at time t . Assuming a unique entry and exit for the factory, i.e. all product enters at $x = 0$ and leaves at $x = 1$, and assuming a 100 % yield the density must satisfy

$$\frac{\partial u}{\partial t} + \frac{\partial(v(u)u)}{\partial x} = 0 \quad (4)$$

where

$$v(u) = \phi(u) \quad (5)$$

is the state equation relating the speed of the product moving through the factory to the amount of product in the factory, i.e. to WIP. Note that the

units of u are $[parts][stage]$ and the units of $v(u)u$ are $[parts][time]$. This suggests that in conventional nomenclature of process control and performance simulation $u(x, t)$ represents local WIP density and the flux $u(x, t)v(x, t)$ is the local throughput at stage x at time t , respectively. Recently, based on earlier work of Graves [14] and Karmarkar [20], Asmundsson et al [5] have developed a similar idea. In order to model the nonlinear response of the throughput to an increase in WIP they introduce the idea of *nonlinear clearing functions* of the form

$$\text{Throughput} = \alpha(\text{WIP}) \text{WIP}. \quad (6)$$

Since throughput is a flux, $\alpha(\text{WIP})$ is a velocity and hence equivalent to our state equation $\phi(u)$ (Equation 5). Asmundsson et al choose specific nonlinear clearing functions motivated by theory of queueing networks but emphasize that only empirical data will be able to decide the *correct* clearing function for a particular problem. They do not incorporate their clearing function idea into a differential equation. Instead, they use the clearing function to generate a linearization of the supply chain planning problem that can be solved via an LP algorithm.

The basic issue for our model is to come up with a useful state equation $\phi(u)$. As a starting point the model has to respect the fundamental law of factory physics, Little's law [28]. For a single factory Little's law may be written as

$$N = \tau\lambda \quad (7)$$

where N is the time averaged load of the factory (WIP), τ is the mean cycle or throughput time over all outputs (TPT), and λ is the start rate. Little's law is fundamentally a deterministic law and results from mass conservation. However, by amending it with a description of the stochastic processes in a factory we can generate a state equation characterizing the factory. In its simplest case the stochastic process is represented by its means. For instance, modeling a factory as a single linear queue with Markov arrival and processing rates (an $M/M/1$ queue) the mean throughput time τ can be determined as a function of the start rate λ and processing rate μ [16] to be

$$\tau = \frac{1}{\mu - \lambda}. \quad (8)$$

Therefore, the relationship between WIP and throughput time becomes

$$\tau = \frac{1}{\mu}(1 + N). \quad (9)$$

Equation (8) shows that μ becomes the maximal or critical start rate. With $v = \frac{1}{\tau}$ we have the desired state equation. Equation (9) shows that WIP is a linear function of the throughput time with a slope μ . The linear relationship between TPT and WIP is intuitively obvious: a part entering the queue will have to wait until the queue is served (N/μ timeunits) plus its own processing time, $1/\mu$ time units.

A linear relationship as in Equation (9) is true for a queueing network that has product form [30], i.e. the whole network can be replaced by an effective queue. BCMP networks [6] and Kelly networks [30] are of product type. Such networks have quite restrictive conditions on the stochastic processes of the network. In general, the Pollaczek Khintchine formula [16] suggests that the mean throughput time depends on the first and second order moments (the means and the variances) of all the stochastic processes involved and hence in general may lead to a *nonlinear* state equation $\phi(\rho)$. Specifically, any process that increases variance as the load in the factory is increased will lead to a nonlinear state equation. For instance, the detailed empirical analysis of Chen et al [7] of a re-entrant semiconductor factory shows that a discrete event simulation based on a product network predicts throughput times within about 10% of the actual throughput times. It depends on your point of view whether 10% accuracy is a large or a small number for describing a factory that is a billion dollar investment. Lu et al [29] show that through dispatch and scheduling rules one can reduce the variance of throughput times inside semiconductor manufacturing plants. Hence, there are several key factors that will lead to a queueing network that cannot be approximated by a product network with constant production rates:

- Chen et al [7] report that 35% of the throughput time of a chip in the HP factory studied was used up on *engineering holds*. Basically those are the queues in front of a human operator who has to determine whether the previous process worked properly or whether re-working is in order. While Chen et al's work analyzes a development factory which may result in a higher incidence of human operator interactions than in a regular production fab, it is nevertheless true for all human operators that they do not have a constant production rate. For instance, it is known that bank tellers work *faster* if the queues in front of them are larger. However, they often do that by cutting corners and reducing quality of service. For a model of a semiconductor factory with constant yield this may lead to the paradoxical effect that the operators work faster but due to errors the mean throughput time of the *good* chips will actually increase.
- Recent simulations of fully automated 300 mm wafer fabs [36] show that increased loading may lead to crowding effects of the transportation system, in turn leading to nonlinear dependence of throughput on WIP.
- Scheduling and dispatch policies play a major role: Kumar [23] discusses a mixed push/pull policy that leads to a blowup in inventory even though all start rates are smaller than all production rates.
- Similarly, the effects of hot lots and other priority rules for multiproduct flows will destroy the product property of a queueing network. DeJong and Wu [11], for instance, show that priority lots have an exponential impact on the throughput of the non-priority lots.

- Even for the restricted models of fluid networks (which do not include operator interactions and multiproduct flows) Dai and Weiss [9] showed that dispatch and scheduling policies play a major role for the stability of the queuing regime.
- Daganzo [8], in discussing the application of traffic flow models on controls and release strategies for a whole supply chain, analyzes situations where the state equation is less determined by any physical constraints but determined by the re-order policies in the supply chain.
- Finally, there is anecdotal evidence from real re-entrant factories [21] that show a stronger than linear increase in the average throughput time τ as the loading of the factory is increased.

Unfortunately, the true state equation for a re-entrant factory is not known, as large factories are rarely (if ever) run in equilibrium and controlled experiments are obviously impossible. It is worth noting here that standard discrete event simulators may not be a good substitute for real experiments, specifically if human interaction is part of the cause for the increase in variance [38].

For future reference we restate the full model, with $f(x)$ an initial density distribution, and taking into account that v does not depend on x :

$$\begin{aligned} \frac{\partial u}{\partial t} + v(u) \frac{\partial u}{\partial x} &= 0 \\ u(x, 0) &= f(x) \\ u(0, t)v(t) &= \lambda(t) \end{aligned} \tag{10}$$

with a state equation

$$v(u) = \phi(u). \tag{11}$$

As in the above examples, the velocity given by the state equation is a functional of u and typically depends on the total WIP in the factory, i.e. on $\int_0^1 u(x, t) dx$ or on suitably weighted WIP distributions. Notice that the start rate $\lambda(t)$ into the factory enters as the boundary condition for the local throughput at $x = 0$.

Our model shares many features with thermodynamic transport equations. In particular, modeling the relationship between density and velocity through a state equation is typically called an adiabatic approximation: Factory flow is modeled as if it were always in equilibrium, following adiabatically the state equation (11). In the thermodynamics of gases this corresponds to the ideal gas equations. Elementary physics tells us that the ideal gas equations are strictly valid only for infinitesimally slow perturbations. Nevertheless, they are a good approximation for many experiments. Similarly, we can expect that fast transients will be modeled poorly but that slow transients and averages will be modeled very well. An adiabatic model is the first closure of a hierarchy of moment expansions for the time evolution of the probability distribution of the parts moving through a factory. A more elaborate model including the first order moment leading to a *dynamic* equation for the evolution of the velocity is presented in [2]. Moment expansions have the advantage of finding the

time evolution of average quantities of a stochastic process via the solution of deterministic equations instead of through repeated simulations.

3 Simulations

3.1 The Method

A typical numerical method to solve Equation (10) is a first-order conservative up-winding scheme. A conservative numerical method is one in which the total mass in the computational domain will be conserved, or at the very least will vary correctly, provided the boundary conditions are properly imposed [25].

To implement this method, we first partition the spatial space into equal subintervals of width h . Let $0 = x_0, x_1, \dots, x_N = 1$ be the endpoints of these subintervals. To advance the solution on these endpoints forward in time, the first requirement to be met is that the time step k satisfy the CFL condition:

$$0 \leq \frac{v(t_n)k}{h} \leq 1$$

where t_n is the current time [26]. This condition prevents the numerical solution from traveling faster than the true solution. We compute the velocity, $v(t_n) = \phi(\int_0^1 u(x, t_n) dx)$ via an extended Simpson's rule quadrature.

Obtaining the time step, we may advance the solution at each grid point x_j , $j = 0, \dots, N$ by:

$$u(x_j, t_n + k) = u(x_j, t_n) - \frac{k}{h} v(t_n) [u(x_j, t_n) - u(x_{j-1}, t_n)]$$

where k is the time step, and the spatial step, h is given by $h = 1/(N - 1)$. To advance the solution at the left boundary we use the boundary condition for the start rate λ in Equation (10) and the propagation scheme

$$u(x_0, t_n + k) = u(x_0, t_n) - \frac{k}{h} [v(t_n)u(x_0, t_n) - \lambda(t_n)].$$

3.2 A Basic Simulation

Utilizing the numerical method of section 3.1, we run the initial boundary value problem (10) for a state equation of the Lighthill-Whitham traffic model of the form

$$v(u) = v_0 \left(1 - \frac{\bar{u}(t)}{L} \right). \tag{12}$$

with

$$\bar{u}(t) = \int_0^1 u(x, t) dx.$$




Figure 1: Snapshot of the density $u(x)$ for the solution of system (10 and 12) at time $t = 1.2641$. Notice that the true solution is discontinuous, while numerical dissipation smoothes the corners.

Here v_0 is the speed for the empty factory and L is the maximal load (capacity of the factory). The influx is given by

$$\lambda(t) = \begin{cases} 2.016 & t < 0 \\ 2.139 & t > 0 \end{cases}$$

This corresponds to a switch from steady state $u_1 = 2.8$ to a new steady state $u_2 = 3.1$. (see below). The parameter values used are: $v_0 = 1$, $L = 10$, $f(x) = 2.8$. Figures 1 and 2 show a snapshot of the solution at $t = 1.2641$ and the outflux as a function of time, respectively. We see that the density is asymptotically approaching a stable steady state density of $u_2 = 3.1$.

Figure 2 shows an important concept that is typical for re-entrant manufacturing systems which has so far received little attention in the literature. The outflux $\omega(t) = u(1, t)v(t)$ initially *declines* when the influx increases. The intuitive reason for this in the context of re-entrant manufacturing systems is the fact that an increase in material at the beginning of the production line competes for machine availability with all the other requests. In particular, an increase at the front of the production line may slow down wafers at the end of the production line as long as they still have to go through the same re-entrant machine. This effect is known as an *inverse response* in control theory and is most pronounced for push and FIFO policies. After the initial WIP profile has left the factory (which happens at about $t = 1.5$), the outflux increases drastically since the density jump entered into the factory at $t = 0$ has reached the end $x = 1$. We see that asymptotically a steady outflux of approximately 2.14




Figure 2: Outflux as a function of time for the basic simulation. Notice the transient time for the solution to reach steady state.

is reached.

The time to do these computations is not an issue here. All the Figures in this paper were generated on a regular PC with a code written in C within seconds to a couple of minutes processing time.

4 Discrete event simulations vs. PDE models

In order to compare the PDE models with standard discrete event simulations we have performed two sets of experiments: We have generated a discrete event simulation in χ [19] and studied the transient behavior of throughput and throughput time for a step up experiment in the influx. We also simulate a temporary overload for a reverse production line. In a second experiment we have used a large scale model of an INTEL factory and compare the discrete event simulation of a four level step up experiment to the corresponding PDE simulation.

4.1 A relatively small re-entrant factory

A PDE model like Equation 10 is based on a continuum approximation for the amount of material in the factory (leading to a density u) and a continuum approximation for the number of stages in a factory (leading to an independent variable x describing the degree of completion of the product). Hence, although it is tempting to compare the PDE simulations with queuing models like Tandem queues [31] or other simple queueing networks [10], such comparisons are a new

research project altogether: Typically a queueing problem that is solvable is based on a small number of machines and a small number of steps. That is exactly the limit in which we do not expect the PDE to behave very well. It would be interesting to study how the queueing model approaches a PDE model in the limit of many machines and many steps but this is clearly beyond the current state of our understanding (but see Lefebvre [24]). In the meantime what can be done is to compare the PDE model with discrete event simulations. In order to do so, we first need to reflect on the parametrization of the PDE model: Much like the clearing function approach, the basic PDE model treats the whole factory as a black box whose input-output characteristics is given by the state equation for the velocity (Eq.11). Hence, for a given dispatch policy and a given product mix inside the factory we can either experimentally or via simulations find a state equation characterizing the steady state behavior of the factory under such circumstances. The PDE will then allow us to *simulate* the behavior of the factory for non-steady behavior such as startup ramps and increases or decreases of the start rates. We do not at present study questions of optimizing dispatch policies or product mixes.

For our first experiment the simulated network consists of 5 machines and is re-entrant, i.e. the production recipe requires that each lot will have to go through the 5 machines four times before it exits. We model only two types of stochastic processes: A stochastic arrival process into the factory and a stochastic exit process for every machine. Both processes are represented by exponential distributions: The arrival process has an exponentially distributed interarrival time whose mean we vary, the machine processes have exponentially distributed processing times with a mean of 0.12 for the first machine and a mean of 0.10 for the other four machines. We typically start up with an empty factory. To determine the state equation representing the steady state behavior we simulate a discrete event simulation run for between 500 and 4000 time units - long enough such that any trace of the initial transient has disappeared by about 1/3 of the simulation time interval. We then use the last half of the simulation time interval to determine average throughput, average throughput time and average WIP. Averaging over about 50 runs constitutes an experiment. We then repeat the experiments until we reach a predetermined confidence interval for the calculated averages (typically 5%). Figure 3 shows seven data points for throughput time vs. WIP for the above experiment with a push policy. The data look almost linear and suggest a state equation for an equivalent single M/M/1 queue, i.e. we use a least squares approximation of the form $\phi(u) = a\bar{u} + b$, where \bar{u} is the total WIP in the factory. The fact that the factory seems to be represented by an equivalent queue is presumably a result of the fact that we did not include any stochastic processes in the discrete event simulation that may generate crowding behavior as discussed in Section 2.

We use this state equation to generate PDE simulations. Trivially, the PDE simulation is very good in predicting the steady state throughput and throughput time for a discrete event simulation experiment with an arbitrary steady influx, confirming the linear least square fit for the data. Comparing the PDE simulation to step up experiments shows some nice agreement and some oppor-





Figure 3: Seven datapoints for a state equation describing the relationship between throughput time and WIP.


tunities for further improvement: We increase the arrival rate from $1/0.9$ and $1/0.8$, respectively to $1/0.7$. With a critical capacity of about $\lambda = 1/0.6$ this corresponds to an step from 65% and 75% capacity to 85% capacity, respectively. We present four figures: Figure 4 shows time series for the throughput and the throughput time as a function of time for a ramp up experiment from 75% capacity to 85% capacity with pull policy. The noisy curve is the average of the discrete event simulations, the continuous curve is the PDE simulation with the state equation generated from Figure 3.

Figure 5 zooms into the transition time series for the throughput and the throughput time of a ramp up experiment from 65% capacity to 85% capacity with push policy. Notice that the throughput times for push and pull policies are very different and hence we have to use a different state equation for each policy. While the transients in both cases are not perfectly resolved, the agreement is not bad either. The large downward spike in the throughput for the PDE simulation results from the fact that in our model the velocity is spatially uniform and depends on the total WIP in the factory. Hence any increase in WIP (through e.g. an increase in influx) will lead to an instantaneous inverse response, i.e. a reduction in velocity and hence to an instantaneous reduction in outflux. Obviously, our simulation factory is not re-entrant enough for such a strong reaction.

We have also run the same discrete event simulation with a reverse routing: Lots visit the machines in the order $S_1, S_2, \dots, S_5, S_5, S_4, \dots, S_1$ repeating this loop four times. Again, we simulate about 7 steady state influx values to generate a state equation for a push and pull policy, respectively. To challenge the

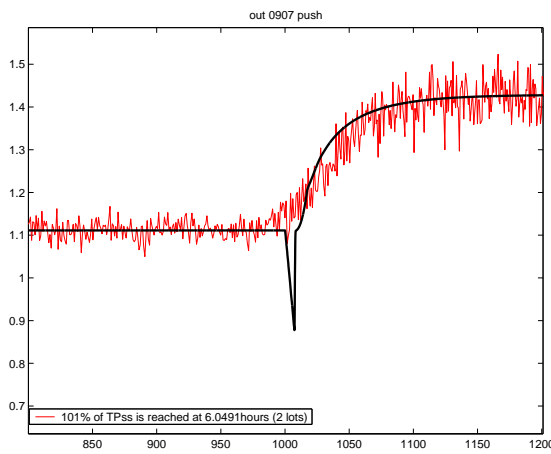


(a)

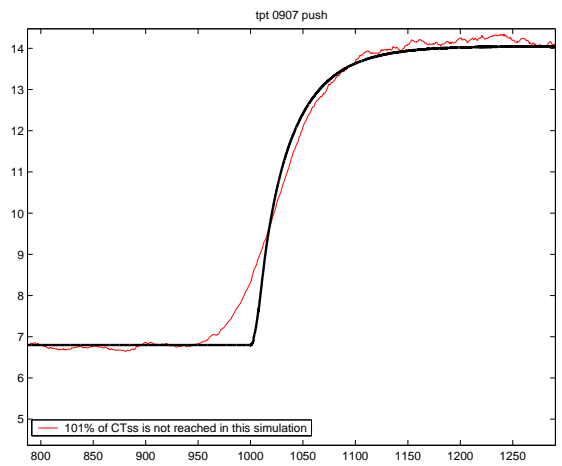


(b)

Figure 4: A ramp up experiment showing a discrete event simulation and PDE model: a) Outflux and b) TPT as a function of time



(a)



(b)

Figure 5: Like figure 4 with a push policy and a zoom into the transition a) Output; b) Throughput time

model we are modeling a temporary overload influx going from about 75% of capacity to 110% of the capacity for $2000 < t < 3000$ before it goes back to 75%. Figure 6 shows the comparison for the throughput of the discrete event simulation under a pull policy with the PDE simulations. The horizontal line shows the capacity limit in steady state. The PDE simulation is not perfect but has qualitatively the correct behavior: It predicts quite well the throughput during the overload interval, it also has a reasonably good representation of the upswing transient. The PDE simulation again has a large inverse response that is not in the discrete event simulation and it also relaxes much faster than the discrete event simulation in the downswing phase.

Trying the same comparison for a the push policy fails: The discrete event simulation shows a huge inverse response spread over a very long time such that the outflux is reduced for much of the overload interval. In addition the outflux never becomes steady. In comparison the PDE simulation restricts the inverse response to the immediate vicinity of the increase in influx and shows a constant outflux within about 250 time units. It seems that a push policy for such a topology leads to extreme variations in the locations of the WIP and hence a description based on an average WIP level in the factory will have to fail badly. This is corroborated by the fact that the discrete event simulations are much more noisy than the previous simulations. This suggests that a model based on higher moments and several stages might do much better (see section 8).

4.2 A large discrete event simulation - an INTEL factory

We compare the simulation of the PDE with a large scale discrete event simulation of a real INTEL Corporation factory. For obvious reasons, we have changed throughput times as well as WIP numbers. The model contains approximately 100 machines, simulates 250 steps for a product mix of 10 products. We are running this model for 6 different start rates and determine the associated average WIP levels and throughput times in steady state. This results in Figure 7. The interpolating curve in the figure follows a suggestion by Asmundsson [5] and is a least square fit of the 6 data points to

$$\tau = \frac{W}{\alpha(1 - e^{-\beta W})} \quad (13)$$

with W the WIP in the factory and α and β are determined by the fit. With $v = \frac{1}{\tau}$ as steady state equation we now study the behavior of the factory to a successive ramp up of the start rate over about a 1000 days. Figure 8 shows the start rate increasing from 500 per day to a 1000 per day in 4 different plateaus. We simulate the PDE with a deterministic start rate that is constant on a plateau but follows the average start rates of the discrete event simulation. We compare the output of the discrete event simulation and the PDE in Figure 9(a). While a single discrete event simulation run takes 1 hr/year on a standard desktop computer, the PDE simulation takes seconds. While the output of the discrete event simulation varies dramatically day by day, the PDE simulation

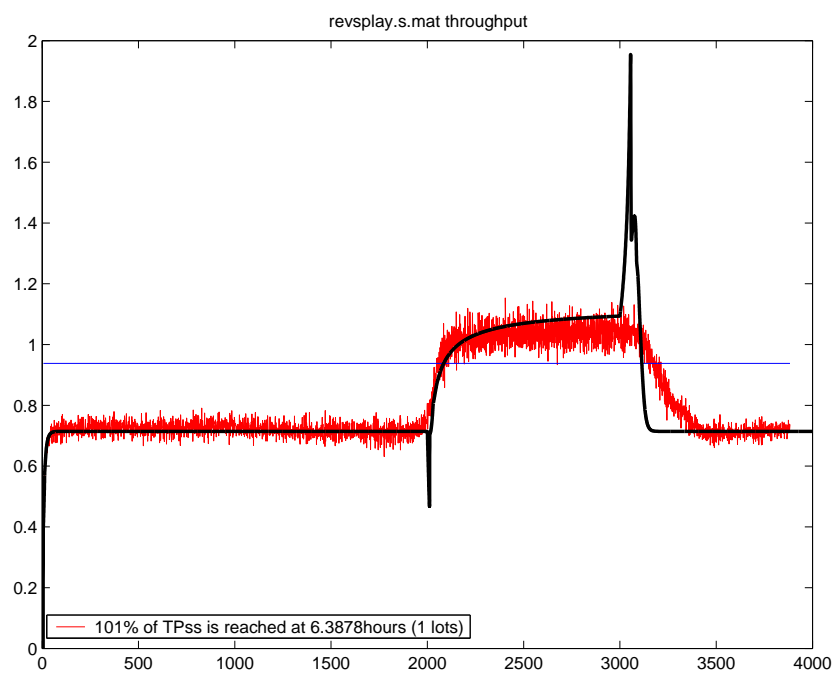


Figure 6: Throughput for the reverse production line with pull policy

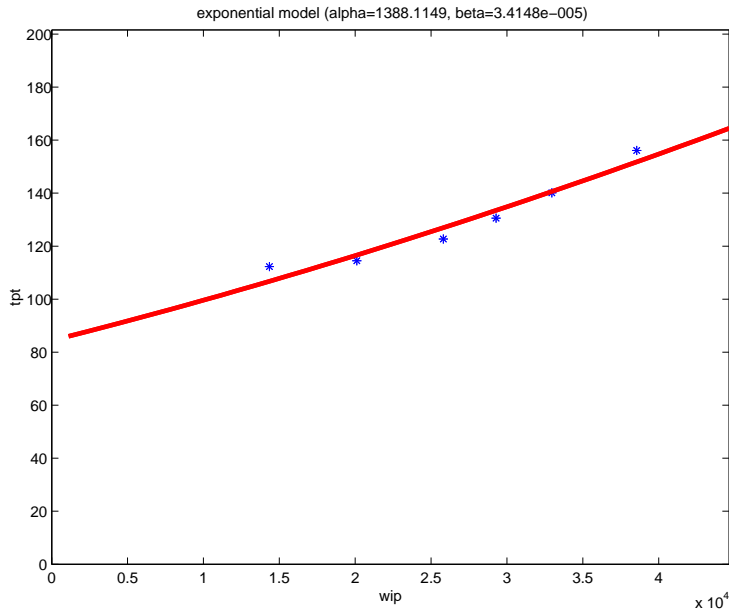


Figure 7: Throughput time versus WIP levels in steady state for a large scale discrete event simulation. The throughput time has an arbitrary linear scale, chosen such that the intersection of the fit is at approximately 100%.

is deterministic and seems to be more or less centered on the average of the output. To illustrate further how the discrete event simulation and the PDE simulation are related we smooth the output by averaging the output over 3 days in the future and 3 days in the past, i.e. a moving 7 day average. Figure 9(b) shows the averaged outs and the PDE simulation again. Figure 9(c) shows the outs averaged over a moving window of 21 days. In [1] we discuss, that closing the hierarchy of moment equations at the zero order level via a state equation implicitly defines a closure for the variance of the throughput time. It implies that the coefficient of variation stays constant. The dotted lines in Figure 9 are calculated by determining the variance for a given plateau of the outputs. Assuming then that the coefficient of variation stays constant, the dotted lines are the mean \pm the standard deviation. It seems that the PDE and its associated standard deviation are a good description of a smoothed transient process.

This comparison allows us to make an interesting observation that we most likely would not expect to find if we are doing only a pure discrete event simulation. As discussed previously, for a re-entrant factory an increase in the start rate leads to an inverse response in the output, i.e the output initially *drops* before it increases to the new level. The direct discrete event simulation Figure 9(a) shows no such inverse response, if it is there it is masked in the raw outputs by the daily variation in the outputs. It is also not commonly found




Figure 8: Start rate ramp up for the discrete event simulations

in the reality of the factory - due in part to the fact that operators are paid to *maintain or increase* the outputs and hence they will work hard to *speed up* WIP at the end of the production line [21]. However, the smoothed outputs in Figure 9(b) and (c) indicate that, without the change of output policies resulting from operator interference, the inverse response can be found in the discrete event simulation too and it follows quite well the PDE simulation. The reason why the inverse response for the discrete event simulations and the PDE for the INTEL factory agree much better than for the simulations in the previous section (4.1) presumably depends on the fact that the INTEL factory is much more re-entrant and that the relative change in its load is much lower than in the other simulations.


5 Equilibria

For the rest of the paper we *assume* a nonlinear state equation analogous to the Lighthill-Whitham traffic model (Eq. 12) repeated here for convenience


$$v(u) = v_0 \left(1 - \frac{\bar{u}(t)}{L} \right). \quad (14)$$

Note however, that all the analysis in the following chapters can easily be adapted to fit any other state equation. We begin by determining the steady state solutions.

Proposition 1 For a given (constant) influx $\lambda(t) = \lambda$ the model (10 and 14)



(a)



(b)

Figure 9: Discrete event simulation and PDE simulation: a) raw outputs, b) outputs smoothed over 7 days, c) outputs smoothed over 21 days.




Figure 13: c) Discrete event simulation and PDE simulation: outputs smoothed over 21 days.

has two steady states:

$$u_+ = \frac{L}{2} + \sqrt{\frac{L^2}{4} - \frac{\lambda L}{v_0}} \quad \text{and} \quad u_- = \frac{L}{2} - \sqrt{\frac{L^2}{4} - \frac{\lambda L}{v_0}}$$

provided $\frac{4\lambda}{v_0} < L$.

Proof: Let u_{ss} represent a constant solution to equation (10). Then the boundary condition becomes a quadratic equation

$$u_{ss}^2 - Lu_{ss} + \frac{\lambda L}{v_0} = 0 \tag{15}$$

and the solutions to this quadratic equation are u_+ and u_- . \square

This implies that a re-entrant factory has two modes of steady operations that generate the same throughput: A high-WIP-high-TPT state and a low-WIP-low-TPT state. Note that any state equation that grows stronger than linear will show two equilibria.

Proposition 2 The high-WIP-high-TPT state, u_+ , is unstable while the low-WIP-low-TPT, u_- , is stable to perturbations.

Proof: Perturbing the solution near a steady state u_{ss} we write $u(x, t) = u_{ss} + p(x, t)$ and insert into system (10)

$$\begin{aligned} \frac{\partial p}{\partial t}(x, t) + v_0 \left(1 - \frac{1}{L} \int_0^1 (u_{ss} + p(x, t)) dx\right) \frac{\partial p}{\partial x}(x, t) &= 0 \\ u(x, 0) &= u_{ss} + p(x, 0) \\ (u_{ss} + p(x, 0))v_0 \left(1 - \frac{1}{L} \int_0^1 (u_{ss} + p(x, t)) dx\right) &= \lambda. \end{aligned} \quad (16)$$

Linearizing the PDE and the boundary condition by discarding higher order terms of $p(x, t)$ we get a linear unidirectional wave equation with constant speed $v_{ss} := v(u_{ss})$ and an integral boundary condition.

$$\begin{aligned} \frac{\partial p}{\partial t}(x, t) + v(u_{ss}) \frac{\partial p}{\partial x}(x, t) &= 0 \\ u(x, 0) &= u_{ss} + p(x, 0) \\ -\frac{u_{ss}v_0}{L} \int_0^1 p(x, t) dx + v_{ss}p(0, t) &= 0. \end{aligned} \quad (17)$$

Taking the time derivative of the boundary condition and solving for $\frac{d}{dt}p(0, t)$ we find

$$\begin{aligned} \frac{d}{dt}p(0, t) &= \frac{u_{ss}v_0}{Lv_{ss}} \int_0^1 \frac{\partial p}{\partial t}(x, t) dx \\ &= -\frac{u_{ss}v_0}{Lv_{ss}} \int_0^1 v_{ss} \frac{\partial p}{\partial x}(x, t) dx \\ &= \frac{u_{ss}v_0}{L} (p(0, t) - p(1, t)). \end{aligned}$$

To go from the second to the third equation, the wave equation for $\frac{\partial p}{\partial t}(x, t)$ (17) was used to replace $\frac{\partial p}{\partial t}$ with $-v(u_{ss})\frac{\partial p}{\partial x}$. Since Equation (17) has a constant wavespeed v_{ss} , the solution at the right boundary is $p(1, t) = p(0, t - \frac{1}{v_{ss}})$. Hence with $\tau = \frac{1}{v_{ss}}$, $\gamma = \frac{u_{ss}v_0}{L}$, and $z(t) = p(0, t)$ stability of the steady state u_{ss} is determined by the stability of the trivial solution of the delay differential equation [22]

$$\frac{d}{dt}z(t) - \gamma z(t) + \gamma z(t - \tau) = 0. \quad (18)$$

Its characteristic equation is

$$\xi - \gamma + \gamma e^{-\xi\tau} = 0.$$

Obviously, $\xi = 0$ is a solution. This corresponds to the fact that the system is a conservation law and hence is neutrally stable towards all perturbations that preserve the load (this reflects the fact that any factory can be run at any load with a constant WIP policy (CONWIP)) [37].

Define $h(\xi, \tau) = (\xi - \gamma)e^{\xi\tau} + \gamma$ and consider $h(\xi, \tau)$ as a function of real ξ . Then:

$$h(0, \tau) = 0, \quad h(\gamma, \tau) = \gamma > 0$$

and

$$\begin{aligned}\frac{\partial h}{\partial \xi}(\xi, \tau) &= e^{\xi\tau} + (\xi - \gamma)\tau e^{\xi\tau} \\ \frac{\partial h}{\partial \xi}(0, \tau) &= 1 - \tau\gamma.\end{aligned}$$

Hence, if $\tau > \gamma^{-1}$, then $\frac{\partial h}{\partial \xi}(0, \tau) < 0$, which implies that there is a $\delta > 0$ such that, when $0 < \xi \leq \delta$, $h(\xi, \tau) < 0$. Since $h(\gamma, \tau) > 0$, there is at least one $\bar{\xi}$, $\delta \leq \bar{\xi}\gamma$, with $h(\bar{\xi}, \tau) = 0$. Hence the characteristic equation always has a positive root, signifying an unstable trivial solution. Conversely, if $\tau < \gamma^{-1}$, then $h(\xi, \tau) > 0$ everywhere. Hence there are no $\lambda > 0$ that solve the characteristic equation and the trivial solution is stable.

It is easy to see that u_+ implies that $\tau > \frac{1}{\gamma}$ while u_- implies $\tau < \frac{1}{\gamma}$ which proves the claim. \square

Note that for a state equation $v(u) = \frac{u_0}{1+u}$ (Equation 9) which describes a simple linear queue (or a product network), there exists only one stable equilibrium. Furthermore it is easy to show that for any state equation $v(u) = \phi(u)$ that slows down *faster* with increasing load than a simple queue there exist two equilibria and that the high speed - low WIP equilibrium is stable while the low speed - high WIP equilibrium is unstable. In particular, the singularity at $\bar{u} = L$ is not a necessary condition for the existence of the two equilibria.

The relationship of these results to standard queueing theory is the following: Consider a small WIP perturbation of a steady state - for instance a localized WIP increase but keep the start rate the same throughout. The perturbation will travel downstream and will eventually leave the factory. In the stable steady state, the perturbation will have generated secondary perturbations that are smaller than the initial perturbation and hence eventually the system will return to equilibrium. In the unstable steady state, an increase in WIP will slow down the system even more leading to an even bigger increase in WIP - i.e. the secondary perturbations are bigger than the original one and hence eventually WIP will increase to infinity. For traditional queueing systems this correspond to the examples of scheduling policies for re-entrant flows [9, 29] that will lead to growth of WIP without bound although the utilization (measured as the start rate relative to the processing rate of the bottleneck) stays below one.

6 Control

The simplest control problem associated with running a factory is to change the production flow of the factory from one steady state, corresponding to outflux that meets a specific constant demand, to another one. Let us assume that the demand $d(t)$ changes in the following way:

$$d(t) = \begin{cases} d_1 & t < 0 \\ d_2 & t > 0 \end{cases}$$

The control problem is to design an influx $\lambda(t)$ that would move the system from the equilibrium u_1 corresponding to a production rate of d_1 to a new equilibrium u_2 corresponding to a production rate d_2 . The outflux and the equilibrium densities u_i are related as

$$d_i = u_i v_0 \left(1 - \frac{u_i}{L}\right). \quad (19)$$

The influx $\lambda(t)$ is uniquely determined by the boundary condition at the left: $\lambda(t) = u(0, t)v(t)$ where the velocity depends on the global load. Hence to generate a WIP profile that has a discontinuous step from an equilibrium $u(x) = u_1$ to new value u_2 at time $t = 0$ we choose the following influx:

$$\lambda(t) = \begin{cases} u_1 v_0 \left(1 - \frac{u_1}{L}\right) & t < 0 \\ u_2 v_0 \left(1 - \frac{1}{L} \left[\int_0^{x(t)} u_2 ds + \int_{x(t)}^1 u_1 ds \right] \right) & 0 < t < \tau \end{cases} \quad (20)$$

where $x(t)$ is the characteristic for the discontinuity emanating from $x = 0, t = 0$ in the xt -plane. At $t = \tau$ the discontinuity is at $x(\tau) = 1$ and we have a new equilibrium profile $u(x) = u_2$. Since u_1 and u_2 are constants, equation (20) reduces to

$$\lambda(t) = \begin{cases} u_1 v_0 \left(1 - \frac{u_1}{L}\right) & t < 0 \\ u_2 v_0 \left(1 - \frac{1}{L} [u_1 + (u_2 - u_1)x(t)]\right) & 0 < t < \tau \\ u_2 v_0 \left(1 - \frac{u_2}{L}\right) & \tau < t. \end{cases} \quad (21)$$

The characteristic $x(t)$ is the solution of the initial value problem: $\dot{x}(t) = v(u(x(t)), t)$, i.e.

$$\begin{aligned} \dot{x}(t) &= v_0 \left(1 - \frac{1}{L} \left[\int_0^{x(t)} u_2 ds + \int_{x(t)}^1 u_1 ds \right] \right) \\ x(0) &= 0. \end{aligned} \quad (22)$$

Since u_1 and u_2 are constants the integrals can be evaluated and Equation (22) reduces to a simple linear ODE. Its solution is

$$x(t) = \frac{L - u_1}{u_2 - u_1} \left(1 - \exp\left(\frac{v_0(u_1 - u_2)t}{L}\right)\right), \quad (23)$$

and hence $\lambda(t)$ yields:

$$\lambda(t) = \begin{cases} u_1 v_0 \left(1 - \frac{u_1}{L}\right) & t < 0 \\ \frac{v_0 u_2 (L - u_1)}{L} \exp\left(\frac{v_0(u_1 - u_2)t}{L}\right) & 0 < t < \tau \\ u_2 v_0 \left(1 - \frac{u_2}{L}\right) & \tau < t. \end{cases} \quad (24)$$

Solving $x(t) = 1$ for t gives

$$\tau = \frac{L}{u_1 - u_2} \ln \left(\frac{L - u_2}{L - u_1} \right). \quad (25)$$




Figure 14: a) Snapshot of the density $u(x)$ for the solution of Equation (10). The influx is given by Equation (24). Notice that with this start rate, the density is piecewise constant. b) Outflux; notice that the flux is constant after the transient time of moving between the steady states.

Using the numerical method described in Section 3.1, the system was run with the parameter values $v_0 = 1$, $u_1 = 2.8$, $u_2 = 3.1$, $L = 10$. With these values we get $\tau = 1.42$. For $t > \tau$ we continue the steady influx $\lambda = u_2 v_0 (1 - \frac{u_2}{L})$. Figure 6(a) shows a WIP profile and Figure 6(b) shows the resulting outflux: Upon the introduction of u_2 into the factory, the total load increases and hence the velocity decreases. The velocity continues to decrease for $t < \tau$. Therefore, the outflux given by $\omega(t) = u_1 v(t)$ decreases as well until the new density $u = u_2$ reaches the end of the production line.

If the change in demand happens without sufficient notice then there exists a time lag during which the factory will not produce according to the new demand, as can be seen in Figure 6. During the transition between the equilibria we produce an outflux $\omega(t) = u(1, t)v(t)$ while we wished to produce an outflux of d_2 . The difference between actual and desired production up to time t is called the *backlog*. It may be written as:

$$b(t) = u_2 \left(1 - \frac{u_2}{L}\right) t - \int_0^t \omega(s) ds. \quad (26)$$

An optimal control problem can now be defined that asks for the input function $\lambda(t)$ that moves the factory from one steady state to another in shortest time, subject to the constraint of a zero backlog. Formally we write:

min T
subject to

$$\begin{aligned} \frac{\partial u}{\partial t} + \frac{\partial(v(u)u)}{\partial x} &= 0 \\ v(u) &= v_0 \left(1 - \frac{\bar{u}(t)}{L}\right) \quad \text{for } t \leq T \\ u(x, 0) &= u_1 \\ v(u)u(0, t) &= \lambda(t) \end{aligned} \tag{27}$$

and

$$b(t) = \int_0^t (d_2 - v(s)u(1, s))ds = 0 \quad \text{for } t \geq T. \tag{28}$$

Standard optimal control approaches in the production and inventory modeling context (e.g. [15]) cannot solve our problem here: They are based on ODEs that cannot take into account the slowdown of the factory as the load in the factory increases due to the control actions. Lefebvre [24] has suggested an approach based on control theory of delay systems. Whether such an approach will work for the re-entrant factory with its large delays remains to be proven.

As we currently cannot solve the optimal control problem in its general setting, we are going to solve the following simplified problem: We assume we want to go from an equilibrium solution u_1 to an equilibrium solution u_2 through *one* intermediate constant density $u = u_3$. We will find the optimal level u_3 and the optimal time that the system will stay in that intermediate equilibrium. While this constraint mostly is dictated by the fact that we cannot solve optimal control problems for nonlinear hyperbolic equations, it is not entirely unreasonable to have this additional constraint: We are trying to find the input sequence that solves the optimal control problem and, in doing so, limits the disruption of the WIP profile to two step functions. Clearly, from a resources and management point of view it is desirable to have a product density in the factory that is as homogenous as possible. Furthermore, we hope that this approach can be the bases for a numerical algorithm that solves the original optimal control problem (work in progress).

We therefore assume that the solution at $x = 0$ takes the following form:

$$u(0, t) = \begin{cases} u_1 & t < 0 \\ u_3 & 0 < t < T \\ u_2 & T < t, \end{cases} \tag{29}$$

i.e. we enter the intermediate density u_3 for T time units. Note that once we enter the desired density u_2 there is still the transition time τ before the system is completely in equilibrium. Under the assumption that v_0 , u_1 , u_2 , and L are known, there are two parameters, u_3 , and T to choose. Therefore, the requirements for a (u_3, T) pair to satisfy are

$$u_2 \left(1 - \frac{u_2}{L}\right) (T + \tau) - \int_0^{T+\tau} \omega(s)ds = 0. \tag{30}$$

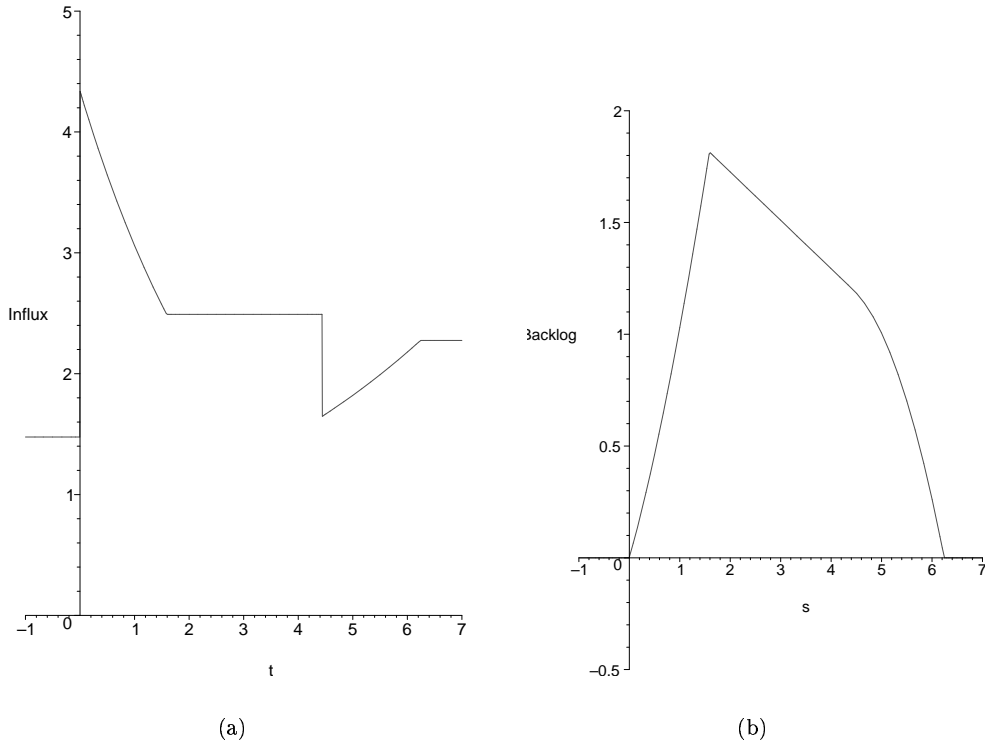


Figure 15: a) Start rate $\lambda(t)$ and b) backlog $b(t)$ for the optimal parameters discussed in Section 6.1. Notice that after $t \approx 6.2421$, the influx is at the new steady state and the backlog stays zero.

There are two cases to consider:

6.1 u_3 saturates the factory

Here we seek to move to an intermediate equilibrium solution, u_3 , and remain on it for some time before moving again to the final equilibrium solution u_2 . The method of Lagrange multipliers and simple algebra allows us to find the optimal values for the intermediate level u_3 and the time T to stay on that intermediate level. As an example, for $v_0 = 1$, $u_1 = 1.8$, $u_2 = 3.5$, and $L = 10$, we find a minimum time for $u_3 = 5.29321$, $T = 2.85286$. Figures 6.1(a) and 6.1(b) show the start rate and backlog for the optimal solution with these parameter values. The drop at around $t = 4$ is due to the inverse response again. As we lower the start rate, we reduce WIP in the factory and hence speed up production and subsequently increase the output. In order not to overproduce, we need to start less for a short amount of time.

6.2 Small backlog- u_3 does not saturate

Here we input the intermediate state u_3 only for a time $t < \tau_{u_3}$ such that the factory never saturates on u_3 . There will be three stages to the transient in moving from u_1 to u_2 : The first stage will be introducing the density u_3 into the factory. The discontinuity from u_1 to u_3 moves through the factory along the characteristic $x_1(t)$. This will yield a decreasing velocity since the load keeps increasing. This will last until $t = T$, at which time the introduction of u_2 will begin. The discontinuity from u_3 to u_2 will move along a characteristic called $x_2(t)$. This second stage will consist of the time needed to purge the factory of u_1 , so that only u_2 and u_3 are left. This time will be referred to as $T + T_1$. The third and final stage will be the expulsion of u_3 from the factory at which time the new equilibrium is reached. We call this time $t = T + T_1 + \tau$ and we are seeking a minimum time subject to the backlog constraint (30) by varying the two free parameters u_3 and T_1 . Again, straightforward algebra and the method of Lagrange multipliers allows us to find the minima.

We find, that for every choice of u_1, u_2 and L that we checked numerically, for which there exists a solution for the backlog constraint, the *shortest* possible solution was to make up the backlog through an instantaneous δ -impulse with a strength of the maximal backlog. While we can prove the optimality of the δ -impulse for special cases analytically, we have not been able to do so for arbitrary u_1, u_2 and L . We have however confirmed it for many cases numerically. Figure 16(a) shows a WIP profile for a jump from $u_1 = 2.8$ to $u_2 = 3.5$ via a maximal density value of $u_{3_{max}} = 17$. Figure 16(b) shows the associated temporal evolution of the backlog.

There is an easy to understand limit for this case: Assume that a backlog of h needs to be made up. If we enter a δ -impulse with strength h into the factory, then the factory will have its maximal load just before the δ -impulse leaves the factory. At that time the load will be $u_2 + h$, assuming wlog. $u_2 > u_1$. Hence, $h < L - u_2$ since otherwise the factory is overloaded and the PDE model breaks down. There is also the practical issue of a δ -impulse in the density: Clearly this is not a realistic model for any factory. However, if we assume that our factory is characterized by a maximal density $u_{3_{max}}$, then our analysis provides a useful heuristic:

If the jump in demand is small (leading to a small enough backlog) then the optimal strategy will be to get the factory to its maximal density instantly and keep it there for the right amount of time, such that when this extra product leaves the factory, the backlog is zero.

We conjecture that the optimal solution for the general optimal control problem (27) for a small enough demand jump is a δ -impulse.

7 Supply Chains

The major advantage of the PDE model for a re-entrant manufacturing flow is speed and scalability. Specifically, using a model like (10) as a node we can




Figure 16: (a) The WIP profile for an approximation of a δ -pulse; (b) backlog as a function of time.

easily link many nodes together to form a supply chain that would be impossible to simulate via a discrete event simulator within a reasonable timeframe. For a network of conservative flows, the challenge for a simulation is not computational time but rather display, analysis and control of such a network. The recent book by Daganzo [8] advocates similar ideas. As a prototype example we are studying a 3-node chain: An acyclic factory feeding into a re-entrant factory feeding into another acyclic factory.

The model for the acyclic (or queueing) factory has been discussed in [2]. It is based on a generalization of ideas from gas dynamics to queueing networks. Consider a job arriving at a queue with processing rate μ . Its throughput time depends on the number of jobs waiting before it

$$\tau = \frac{1}{\mu}(1 + N). \quad (31)$$

Using this relationship as a boundary condition we derive a set of coupled hyperbolic conservation laws for the WIP density $u(x, t)$ and the velocity $v(x, t)$ of the form:

$$\begin{aligned} \frac{\partial u}{\partial t} + \frac{\partial(vu)}{\partial x} &= 0 \\ \frac{\partial v}{\partial t} + v \frac{\partial v}{\partial x} &= 0 \\ v(0, t) &= \frac{\mu}{1 + \bar{u}(t)} \\ v(0, t)u(0, t) &= \lambda(t), \end{aligned} \quad (32)$$

where λ is again the arrival rate and $\bar{u}(t) = \int_0^1 u(x, t) dx$ is again the total load (WIP). The first equation in system (32) has again the form of a hyperbolic

conservation law. However, instead of the heuristic state equation used in the re-entrant model (10) that relates the velocity to the density, we have a Burger’s equation for the time evolution of the velocity. Note that the boundary condition for v provides a nonlocal feedback of the history of the production on the current velocity through the length of the queue in front of an incoming part. This model for the queueing system is more sophisticated than the model for the re-entrant system since it allows for non-adiabatic relaxation of the velocity fields and hence a better modeling of transient phenomena. In addition, in contrast to the heuristics of section 2 it has been rigorously derived in [2].

It is straightforward to link the two models together to form a supply chain. We assume no buffers in between factories and hence the outflux of one factory becomes the influx of the next. This implies, if we extend our completion variable x to cover the whole chain, that the flux will be a continuous variable of x while the density will show discontinuous jumps between factories.

In all experiments since the PDE models are deterministic moment equations, steady state modeling is not interesting: Either on the average the influx is below capacity in which case we find a stable equilibrium given by the state equation or the influx is above capacity in which case WIP will explode. However, the PDE models allow us to focus simulations driven by a time dependent influx $\lambda(t)$.

7.1 Experiment 1: short time overload of the queueing factory

Fig. 17 shows snapshots of a steady state density and flux for a linear chain of three factories - two identical queue modules surrounding a re-entrant module. Each module makes up one-third of the total x -axis. The raw velocity ($\frac{1}{\mu}$) for each queue is 1.2 and for the re-entrant module (v_0), it is 2.0. The capacity for the re-entrant module is $L = 12$. We start the experiment from an empty supply chain. The snapshot is taken at approximately $t = 32$ at which time the transients have disappeared. The influx is periodic, given by $\lambda(t) = 0.75 + 0.5 \sin(2\pi t)$. The starts are constructed so that they exceed the threshold allowable for an equilibrium to exist in the queue. This occurs for only a short part of a period, on average they stay below. Figure 18 shows the outflux of this experiment while Figure 19 shows WIP profiles for each module. Interesting observations are:

- transient oscillations in the output that are much larger than the oscillations in equilibrium - notice the “bulge” in the outs between $t = 10$ and $t = 25$;
- extremely long transients (on the order of 30 throughput times) for the queueing factories;
- the re-entrant factory equilibrates much faster than the queueing factories;
- the re-entrant factory acts as a damping device reducing the amplitude of the oscillations in WIP in the downstream factory.




Figure 17: Snapshot of the steady state density and flux in the 3-node chain.

7.2 Experiment 2: stronger overload of the queuing factory

We increase the starts to $\lambda(t) = 1.1 + 0.5 \sin(2\pi t)$, i.e. the mean influx of 1.1 is just below the threshold of $\mu = 1.2$ allowable for a constant steady state to exist in the queues. Figures 20 and 21 show the outflux, and WIP for all three factories for this experiment. The major difference to the previous experiment are the much longer transients and the absence of the bulge in the outflux.

7.3 Experiment 3: Short time overload of the re-entrant part

We return to the starts profile of experiment 1: $\lambda(t) = .75 + 0.5 \sin(2\pi t)$. However, we now reduce the capacity for the re-entrant module to $L = 4$ leading to a critical influx for the re-entrant part of $\lambda_c = 1$. This influx profile violates the threshold for constant steady states to exist for the queues (for a very short time) and for the re-entrant module (for a longer time). The time series for outflux and WIP for this experiment are very similar to those in experiment 1. In particular, the "bulge" in the transient of the outflux re-appeared. The most striking difference can be seen in the snapshot of the density and flux of the




Figure 18: The outflux of the experiment of ramping up the three node chain from zero

re-entrant factory: Since the re-entrant part is partially overloaded its velocity will decrease and hence the overall WIP will increase from an average of $u \approx 0.5$ in Figure 17 to an average of $u(x) = 1$ in Figure 22. In addition we find that the flux and density waves travelling through the factory have a much shorter wavelength Λ than in the previous experiments suggesting a dispersion relation $\Lambda \propto v$.

8 Conclusions and Further Work

This study can be viewed as a response to the call for scalable simulation of supply chains [34]. We have shown that high volume production in linear and re-entrant factories can successfully be modeled via hyperbolic conservation laws. This idea opens up the vast literature on traffic modeling and hydrodynamic transport equations to be adjusted for the details of production flow. In particular, standard numerical schemes allow fast and accurate simulation of production networks. Obviously a lot of future work needs to be done to make such an approach a real competitor to the standard discrete event simulators.

An issue that needs clarification is the issue of network topology and the related issue of prioritization or dispatch rules. The appeal of the current formulation is that it treats the dynamics of a full factory as one unit whose parametrization is completely determined by its state equation. While this allows to discuss experiments changing the input of the factory and incorporating such PDEs as building blocks in supply chains (see Section 7) it does not allow us

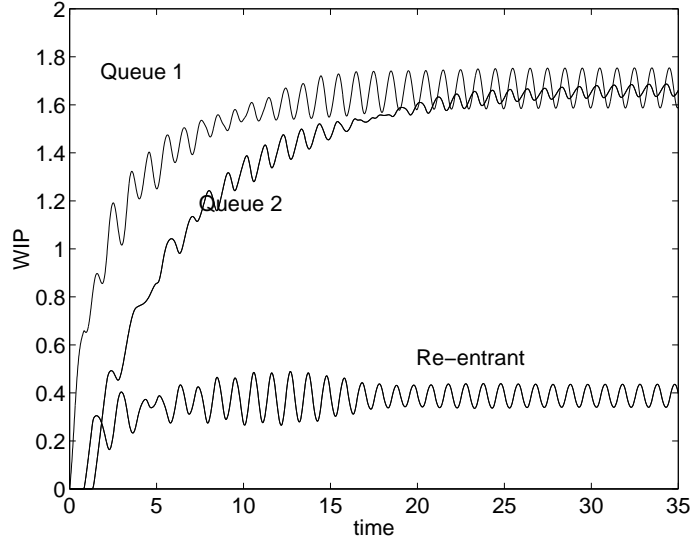


Figure 19: WIP profiles from 7.1 for each module. Each module began empty. Notice the length of the transient to equilibria: the queues took far longer than the re-entrant module to reach steady state.

to study easily the changes due to changing dispatch rules or changing product routing. If we want to do this we need to give up the concept of a homogeneous model (in completion space) of the factory and introduce completion dependent parameters. In the simplest case this can be done by coupling two PDE models together where one feeds into the other. For instance, the transistor and the metalization segments in chip processing would naturally be separated like that. Each part is re-entrant but not across the transistor/metalization boundary.

A single state equation, yet a more accurate model, uses appropriate weight functions $w(x, s)$ to represent the influence of WIP at stage s on the velocity of WIP at stage x

$$N(x, t) = \int_0^1 w(x, s)u(s, t)ds.$$

For instance

$$N_{pull}(x, t) = \int_x^1 u(s, t)ds.$$

restricts the WIP relevant for the velocity at point x to the WIP in front of x , whereas

$$N_{push}(x, t) = \int_0^x u(s, t)ds.$$

restricts it to the WIP behind x - mimicking pull and push policies, respectively.

We are currently pursuing several first principle models and heuristic extensions of these ideas:




Figure 20: The outflux for experiment 2. Note the absence of the “bulge” seen in figure 18.

- While it is extremely hard to characterize in any meaningful way the actual stochastic processes involved we know that there exist some physical limits. In particular, a fixed production line has a maximal production capacity at every machine. No matter what policies, the recipe for a certain chip expects it to stay x hours in a diffusion oven. Hence we can define a maximal capacity function $C(\xi)$ that describes the capacities of all machines that are involved in the production process, where ξ now is a variable that does not describe the stages but the sequence of machines. The resulting quasi-static model then becomes

$$\begin{aligned}
 \frac{\partial u}{\partial t} + \frac{\partial F}{\partial x} &= 0 \\
 F(x) &= \min\{uv_{eq}, \mu(x)\} \\
 v_{eq} &= \Phi(u)
 \end{aligned}
 \tag{33}$$

where the maximally available capacity at stage $\mu(x)$ depends on the dispatch policy: Assume that there are n layers that are produced on the same machines leading to n loops through those machines. The map between machine position and production stage is then given by modular division: The stage variable x acquires a layer index i such that $x_i = \frac{i-1}{n} + \xi$ for $i = 0..n - 1$ describes a production stage in the $i + 1$ th loop at the machine position ξ . At any particular machine, flux requests from all n loops may compete for the maximally available capacity $C(\xi)$ leading to a distribution of the maximally available capacities at stage i , $\mu(x_i)$, depending on the fluxes $F(x_j)$ in the other loops. For instance for a PUSH policy, capacities are allocated from front to end. Hence we can iterate,

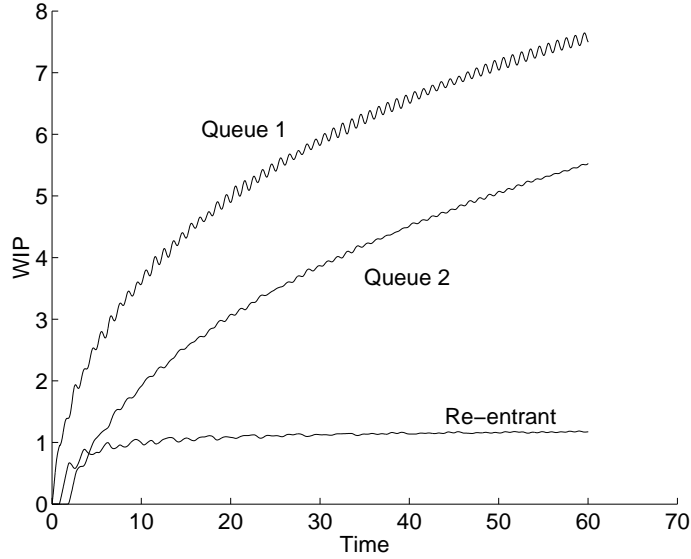


Figure 21: WIP profiles for experiment 2. The queues are clearly not yet in equilibrium.

for $i = 0..n - 1$ the following scheme to find the capacity distribution $\mu(x)$.

$$\begin{aligned}
 \mu(x_0) &= C(\xi) \\
 F(x_i) &= \min\{\mu(x_i), uv_{eq}|_{x_i}\} \\
 \mu(x_i) &= \max\{0, \mu(x_{i-1}) - F(x_{i-1})\}
 \end{aligned}$$

Such a model will lead to the formation of bottlenecks and hence δ -distributions in the density variable for any influx that temporarily exceeds the total capacity.

- For a completely deterministic flow network we have derived the system (33) from first principles using models from gas dynamics [3]. In [4] we extend this to derive a general model supporting arbitrary policies.
- The fundamental PDE model (Eq. 10) with a state equation (Eq. 5) based on mass conservation can be shown to be the zero order moment expansion of a Boltzmann equation with appropriate closure assumptions [2]. Specifying the stochastic process that generates the randomness in the production process as a stochastically varying throughput time adds a diffusion term to the mass conservation [1].

We are currently working on the following issues




Figure 22: Snapshot density and flux profiles for experiment 3 at $t \approx 58$ seconds.

- A major source of stochasticity in factory production are machine breakdowns. Usually the distributions of time to failure and time to repair can be approximated. This will lead to a stochastic version of the capacity limited flow network discussed above.
- The optimal control problem (27) should be solved in its most general form. Specifically, extending the analysis of sections 6.1 and 6.2 to allow for more than one intermediate level should lead to a numerical algorithm for the optimal control problem.
- Parametrization and validation of PDE based models for complicated discrete event models and for real factories. The goal is to determine modifications of the simple state equation (12) model to describe more accurately the inverse response of a production system and the general form of transients and their dependence on policies. This will involve diffusion terms ([1]) and/or terms representing a relaxation time. Another goal will be to determine the discretization error associated with treating a fundamentally discrete problem as a continuous flow.
- While we showed the feasibility of linking hyperbolic conservation laws together to make a simple linear chain this is really just a proof of concept. The real test for the usefulness of our approach will be whether relevant business questions can be answered for a supply network by linking our nodes together and performing simulations. This may involve the development of an object oriented simulation interface. A related project uses the PDE based models as the predictive model for attempts to optimize

the behavior of a whole supply chain via Model Predictive Control (MPC) algorithms [35]

Acknowledgements

We thank Edward Yellig of INTEL Corporation for executing the simulation runs for us. John Fowler, Daniel Rivera and Matthias Kowski and their students helped to focus our ideas through many constructive discussions. The χ simulations are extensions of codes written by R. v.d. Berg and E. Lefeber. D.A. and C.R. are supported by NSF grants DMS-0204543. Support from Intel Corporation is also gratefully acknowledged.

References

- [1] D. Armbruster, C. Ringhofer: "Thermalized kinetic and fluid models for re-entrant supply chains", in print, *SIAM J. Multiscale Modeling and Simulation*, 2005.
- [2] D. Armbruster, D. Marthaler, C. Ringhofer: Kinetic and Fluid Model Hierarchies for Supply Chains, *SIAM J. on Multiscale Modeling* 2, pp.43-61 (2004).
- [3] D. Armbruster, P. Degond, C. Ringhofer: A Model for the Dynamics of large Queuing Networks and Supply Chains, submitted, 2004
- [4] D. Armbruster, P. Degond, C. Ringhofer: "Kinetic and fluid models for supply chains supporting policy attributes submitted, (2004).
- [5] Jakob Asmundsson, Reha Uzsoy, and Ronald L. Rardin: Compact nonlinear capacity models for supply chains: Methodology, preprint, 2002, Purdue University
- [6] F. Baskett, K.M. Chandy, R.R. Muntz and F.G. Palacios, Open, closed and mixed networks of queues with different classes of customers, *Journal of the Association of Computing Machinery*, **22**, 248-260 (1975)
- [7] H. Chen, J.M. Harrison, A.Mandelbaum, A. Van Ackere, and L.M. Wein, Empirical evaluation of a queueing network model for semiconductor wafer fabrication, *Operations Research* **36**, 202- 215 (1988)
- [8] Carlos F. Daganzo, *A Theory of Supply Chains*, Springer Verlag 2003
- [9] J. G. Dai, G. Weiss, Stability and instability of fluid models for certain re-entrant lines, *Mathematics of Operations Research*, **21**, 115-134 (1996)
- [10] J. G. Dai, J.H. Vande Vate, The stability of two-station multitype fluid networks, *Operations Research*, **48**(5), 721-744 (2000)

- [11] Chad D. DeJong, Scott P. Wu: Simulating the transport and scheduling of priority lots in Semiconductor factories, in Proceedings of the 2002 Winter Simulation Conference, E. Yücesan, C.-H. Chen, J.L.Snowdon and J.M.Charnes, eds. pp. 1387 - 1391, (2002)
- [12] A.M. Law and W.D. Kelton: *Simulation Design and Analysis*, McGraw-Hill New York, 1991
- [13] J. W. Forrester, *Industrial Dynamics* (MIT Press, 1962).
- [14] S. C. Graves: A tactical planning model for a job shop, *Operations Research* **34** 525-533 (1985)
- [15] S.B. Gershwin, r. Akella, Y.F.Choong: Short-term production scheduling of an automated manufacturing facility., *IBM J.Res. Develop.* **29**(4), 392-400 (1985)
- [16] D. Gross and C.M. Harris, *Fundamentals of Queueing Theory* (Wiley, 1985).
- [17] D. Helbing, Traffic modeling by means of physical concepts, in: *Workshop on Traffic and Granular Flow*, D.E. Wolf, M. Schreckenberg and A Bachem, eds. World Scientific, Singapore (1996)
- [18] Jim Hines, MIT, personal communication, July 2003
- [19] A.T. Hofkamp and J.E. Rooda, Chi reference manual, Technische Universiteit Eindhoven, <http://se.wtb.tue.nl/documentation/>, 11/2002
- [20] U.S. Karmarkar: Capacity loading and release planning in Work-in-Progress (WIP) and Lead-times, *J. Mfg. Oper.Mgt.*, **2**, 105 - 123 (1989)
- [21] Karl Kempf, Intel Corporation, personal communications (2001)
- [22] Yang Kuang, *Delay Differential Equations* (Academic Press, 1993).
- [23] P.R. Kumar, Re-entrant lines, *Queueing Systems*, **13**, 87-110 (1993)
- [24] Erjen Lefeber, Nonlinear Models for Control of Manufacturing Systems, to appear in Radons, Neugebauer (Eds): *Nonlinear Dynamics of Production Systems*, (Wiley-VCH), 2003
- [25] R.J. LeVeque, *Finite Difference Methods for Differential Equations*. Draft Version for use in AMath 585-6 University of Washington, (1998).
- [26] R.J. LeVeque, *Numerical Methods for Conservation laws* (Birkhäuser-Verlag, 1992).
- [27] M.J. Lighthill, G.B. Whitham, On kinematic waves II. A theory of traffic flow on long crowded roads, *Proceedings of the Royal Society, Series A*, **229**, 317-345 (1955)

- [28] John D.C. Little, A proof for the queuing formula $L = \lambda W$, Operations Research, **9**, 383-387, 1961
- [29] S.C.H. Lu, D. Ramaswamy, and P.R. Kumar, Efficient scheduling policies to reduce mean and variance of cycle-time in semiconductor manufacturing plants.
- [30] R. Nelson, *Probability, Stochastic Processes, and Queueing Theory*, Springer Verlag, 1995
- [31] G.F. Newell, *Approximate Behavior of Tandem Queues*, Springer Verlag, Berlin (1979)
- [32] G.F. Newell, Approximation methods for queues with application to the fixed-cycle traffic light. SIAM Review **7**(2) 223-240 (1965)
- [33] G.F. Newell, Scheduling, location, transportation and continuum mechanics; some simple approximations to optimization problems, SIAM Journal of Applied Mathematics **25**(3), 346-360 (1973)
- [34] Scalable Enterprise Initiative, NSF 2000
- [35] Wang, W., D. E. Rivera, K.G. Kempf and K. D. Smith, "A Model Predictive Control Strategy for Supply Chain Management in Semiconductor Manufacturing under Uncertainty," 2004 American Control Conference, Boston, MA, pgs. 4577-4582, June 30-July 2, 2004.
- [36] Sameer T. Shikalgar, David Fronckowiak, and Edward A. MacNair: 300mm wafer fabrication line simulation model in Proceedings of the 2002 Winter Simulation Conference, E. Yücesan, C.-H. Chen, J.L.Snowdon and J.M.Charnes, eds. pp. 1365- 1368, (2002)
- [37] M.L.Spearman, D.L.Woodruff, W.J.Hopp: CONWIP: a pull alternative to Kanban, International Journal of Production Research **28**(5), 879-894 (1989)
- [38] J.Spier, K. Kempf: Simulation of emergent behavior in manufacturing systems, Proc. 6th SEMI/IEEE Adv. SemiCond. Mfg. Conf. 90-94 (1995)
- [39] L.M. Wein: Scheduling networks of queues: Heavy traffic analysis of a two-station network with controllable inputs, Operations Research **38**(6), 1065 - 1078 (1990)