A PRACTICAL MODEL OF HEINEKEN'S BOTTLE FILLING LINE WITH DEPENDENT FAILURES

ROGER M. COOKE, ATZE BOSMA, AND FRANK HÄRTE

Abstract. This article describes a mass balance production line model based on constant machine rates, fixed finite buffer size and stochastic failure and repair behavior. Failure and repair processes need not be exponential and need not be independent. The model is easy to initialize with field data and we report favorable results from a validation exercise at Heineken Brewery. This article presents a new diagnostic for distance to infinite buffer equilibrium and a new test for independence of failure-repair processes are derived. It is shown that the stochastic fluctuations in a typical line are large, and machine unreliability can perturb the line's design.

1. Introduction

Production lines or transfer lines, or manufacturing flow lines have been studied for at least 40 years. Underlying these models is the theory of repairable systems, which has been studied by [22, 14, 9, 18, 20]. In references [22, 18] it is proved that the total downtime for repairable systems with independent failure and repair processes is asymptotically normal, [20] discusses the covariance structure, and more recently [21] derives the total downtime distribution under arbitrary failure-repair processes using the theory of point processes.

Less is known about the behavior of systems of repairable components with buffers, as encountered in production lines. A classic reference is [25]. For a review of mathematical models, see [5, 15, 6]. Exact mathematical results typically concern equilibrium behavior of the line and introduce strong restrictions on the number of buffers [13, 7, 23, 25, 16] or on the machine's production rates [17]. Reference [1] derives bounds on steady state behavior. All such approaches assume that machine failures and repairs are independent. Dallery and Gershwin [6], p. 75 remark that the variance in production rates and buffer levels has been entirely neglected. The same might be said of the problem of diagnosing whether a line ever attains equilibrium. A result in this direction for the infinite buffer equilibrium is given in section 4. Reference [21] gives the joint downtime distribution for two independent failure repair processes, which is the basis for the statistical test for independence described below. Reference [8, 24, 11, 10] describe bottling lines from an engineer's perspective.

Sophisticated software models are available for analyzing the behavior of production lines. The transport function of conveyors and variable machine production rates are among the features such models capture. As data required to run a sophisticated model are not easily obtained from the field, validation of model predictions

Date: 4-12-03.

Key words and phrases. Production/scheduling, a production line with dependent machine failures, joint unavailability, failure repair process, uncertainty analysis.

is difficult. While such models are helpful in understanding detailed behavior, they are often unable to provide insights into the behavior of the line as a whole.

Line managers typically face questions like : Is the buffer capacity necessary? Is this capacity big enough? Would an upgrade in machine X improve line performance? Is there a significant difference between the day-to-day performance of production teams? etc.

In this paper we present a practical production model based on mass balance equations. The model has modest data requirements, and can be initialized with machine mean up- and down times and joint unavailabilities, in addition to fixed line parameters. The independence assumptions mentioned above are not supported by data and are not invoked. Because of its simplicity, the model can be validated; favorable validation results at Heineken Brewery are reported. While the model is not suitable to study detailed variations in machine or conveyor design, it has given line managers some macroscopic insights into line behavior.

The next section describes the one-way filling line. The mathematical model is presented in section 3. The zero- and infinite-buffer equilibria are discussed in section 4, and a simple diagnostic is given for distance to the infinite buffer equilibrium. Sections 5 and 6 discuss validation, and section 7 discusses an application of the model to week aggregated data. A concluding section describes the ways in which experience with the model supported line managers decision making.

2. The One-Way Filling line

A number of machines function in series to produce filled bottles of beer. Empty bottles arrive on pallets. The depalletizer takes the bottles from the pallets and puts them on a conveyor. The rinser/filler washes the bottles and fills them with beer. Next the pasteurizer heats the bottles for 45 minutes. The labeller applies the labels and caps, the packer places the bottles in crates, and the palletizer places the crates on pallets. Between each machine there are conveyors. Conveyors have both a transport function and a buffer function, that is, they convey bottles from one machine to the next, and they store bottles so that a given machine may keep running even if the upstream or downstream machine is temporarily down. Each machine has a *nominal rate* of production, but can be run faster or slower to compensate for breakdowns. The *core machine* is the machine with the slowest rate (in this case the rinser/filler with 667 bottles per minute). The core machine is typically the most expensive and most difficult to upgrade. The nominal rates and buffer sizes for the line analyzed are shown in Table 1.

The major simplification of the model introduced below is that the transport function of the conveyors is neglected. All machines are modelled to run at their nominal rates if they are up and run at rate zero if they are down (or starved or blocked, see below).

If the machines were perfectly reliable, then the line would produce at the rate of the core machine and the buffer function would be unnecessary. The effective rate of a machine is the nominal rate times the availability, where availability is the probability that the machine is not down. The effective rate of the core machine in the data analyzed below is 473 bottles per minute. Evidently, machine reliability is an important factor in determining line performance. When a given machine

Rates and Buffer capacities			
Machine	Rate (bottles/min)	Downstream Buffer	Capacity (bottles)
depalletizer	900	b1	9000
rinser/filler	667	b2	4400
pasteur	667	b3	3900
labeller	800	b4	6700
packer	912	b5	9600
palletizer	900	OUT	∞

TABLE 1. Filling line rates and buffer capacities with V design

goes down, the upstream machines will continue to function as long as buffer space is available. These buffers will eventually fill up, at which time the upstream machine(s) must stop. A machine which is forced to stop in this way is called *blocked*. Machines downstream from the broken machine will continue to function as long as there are bottles in the appropriate buffers. These buffers will eventually become empty, at which point the downstream machine(s) must stop. A machine which is forced to stop in this way is called *starved*. Periods of blockage or starvation are *not* unavailabilities. Machine unreliability thus leads to large and expensive buffers.

A filling line should be designed in such a way that the core machine is never starved or blocked. This results in the 'V design' shown in Figure 1 below. Each machine is represented as a pump whose vertical size is proportional to the machine's nominal rate. The vertical size of each buffer is proportional to the buffer capacity. Figure 1 is based on the data from Table 1 and shows clearly the "V" with apex at the core machine indicating that the filling line studied here has been properly designed. Of course, if effective rates differ from nominal rates, then this might disturb the V design and degrade performance. (Interestingly in a similar though not identical context, [4], p. 432, reports that "there is a slight advantage in the first and last stage having somewhat higher breakdown rates" than that of the other machines.)

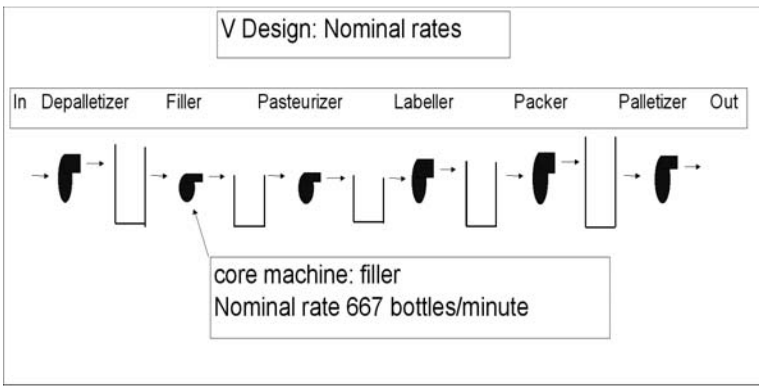

FIGURE 1. *Filling line with V design*

Times of machine failures and repairs are recorded and can be recovered at the end of a shift. The movement of bottles is monitored for throttling production rates.

The actual status of the buffers and the thruput of each machine are not recoverable at present (though further automation is underway). A typical production run lasts 5 hours; runs are terminated by product changes and (sometimes) by shift changes.

3. Discrete time filling line model

The filling line is modelled as a continuous production line with buffers, constant production rates, discrete time and stochastic availability. "Continuous production line" means that production output is real valued - fractional bottles of beer are allowed. With a production of hundreds of bottles per minute, the error introduced by modelling production as a continuous variable is small. We assume that all failure and repair events occur at discrete times $t = 0, 1, 2, 3...$, with time step size δt which is small relative to the failure and repair rates of the various machines. Let $X_i(t) = 1(0)$ if machine i is up (down) at time t. We wish to produce a sequence of 0's and 1's which model the up-down behavior of machine i, then we express the performance of the line as a function of the $X_i(t)$'s.

Let $U_i(t)$ be uniformly distributed on $[0, 1]$, and let $U_i(t), U_i(t')$ be independent for $t \neq t'$. For $r \in [0, 1]$, the event $\{U_i(t) < r\}$ has probability r. We set $X_i(0) = 1$ (machines start in the up state), and set:

$$X_i(t) = \left(\begin{array}{ll} 1 & \text{if } X_i(t-1) = 1 \text{ and } U_i(t) > e^{-\lambda_i} \\ 1 & \text{if } X_i(t-1) = 0 \text{ and } U_i(t) < 1 - e^{-\mu_i} \\ 0 & \text{otherwise} \end{array} \right.$$

$U_i(t)$ and $U_j(t), i \neq j$, may be independent or may be assigned positive or negative dependence on the basis of observed unavailabilities (see below). The failure rate λ_i and repair rate μ_i above are constant in time. The failure rate is the inverse mean time to failure (MTTF), and the repair rate is the inverse mean time to repair (MTTR).

Of course dependence in time can easily be modelled by choosing $\lambda_i = \lambda_i(t), \mu_i = \mu_i(t)$ and correlating $U_i(t), U_j(t')$, but it is very difficult to collect data for estimating such time dependencies. Uncertainty in the failure and repair rates can be modelled by drawing the $\lambda_i' s$ and $\mu_i' s$ from appropriate distributions (either at the inception of each run or at the inception of each time step).

The availability of machine i is the probability that $X_i = 1$. In general this is a function of time and depends on the initial state of X_i when time begins. However, for large time values the initial state is "forgotten" and the "equilibrium availability" of i is found by setting $P(X_i(t) = 1) = P(X_i(t-1) = 1)$ in:

$$P(X_i(t) = 1) = P(X_i(t) = 1 | X_i(t-1) = 1) P(X_i(t-1) = 1) \tag{3.1}$$

$$+ \; P(X_i(t) = 1 | X_i(t-1) = 0) P(X_i(t-1) = 0); \tag{3.2}$$

$$P(X_i(t) = 1) = \frac{1 - e^{-\mu_i}}{(1 - e^{-\mu_i}) + (1 - e^{-\lambda_i})}. \tag{3.3}$$

If $\mu_i < 0.1$, and $\lambda_i < 0.1$, then

$$P(X_i(t) = 1) \approx \frac{\mu_i}{\lambda_i + \mu_i}$$

which is the equilibrium availability for a continuous time exponential failure-repair process.

The computation of performance is explained below with five buffers:

$$\begin{array}{ccccccccccccc} \text{IN} \rightarrow & b_1 & \rightarrow & b_2 & \rightarrow & b_3 & \rightarrow & b_4 & \rightarrow & b_5 & \rightarrow & \text{OUT} \\ v_1 & & v_2 & & v_3 & & v_4 & & v_5 & & v_6 & \end{array}$$

The first machine draws bottles from an infinite buffer IN; the last machine places the finished product in an infinite buffer OUT. v_i is the nominal production rate of machine i; that is, v_i bottles can be processed in one time step. $b_i(t)$ is the actual amount in buffer i at time t; At the first time step, $b_i = 0$. k_i is the maximal amount in buffer b_i. At the first time step we assume every machine is working. Further, we introduce for each time step the thruput for machine j:

$$th_j(t) = \text{number of units moved from } b_{j-1} \text{ to } b_j \text{ in step } t$$

where $b_0 = \text{IN}$ and $b_6 = \text{OUT}$.

Mass balance equations say that the amount in any buffer b_j at time t equals the amount in b_j at $t-1$ minus what leaves in step t, plus what arrives in step t. Further, the thruput of machine i in a given time step is v_i if X_i, is up, if at least v_i bottles are available in the upstream buffer, and there is room for v_i bottles in the downstream buffer. Where Δ OUT denotes the increment in output, the mass balance equations may be written as:

$$\begin{array}{lcl} \Delta\text{OUT}(t) & = & \min\{v_6 * X_6(t), b_5(t-1)\} \\ b_5(t) & = & \min\{k_5, b_5(t-1) - \Delta\text{OUT}(t) + min\{v_5 * X_5(t), b_4(t-1)\}\} \\ b_4(t) & = & \min\{k_4, b_4(t-1) - th_5(t) + min\{v_4 * X_4(t), b_3(t-1)\}\} \\ \cdots & = & \cdots \\ b_1(t) & = & \min\{k_1, b_1(t-1) - th_2(t) + v_1 * X_1(t)\}. \end{array}$$

The thruput quantities can be expressed as:

$$\begin{array}{lcl} th_5(t) & = & b_5(t) - (b_5(t-1) - \Delta\text{OUT}(t)) \\ th_4(t) & = & b_4(t) - (b_4(t-1) - th_5(t)) \\ \cdots & = & \cdots \\ th_1(t) & = & b_1(t) - (b_1(t-1) - th_2(t)). \end{array}$$

It will be observed that $b_i(t)$ can be computed on the basis of $X_i(t), b_{i+1}(t), ...b_5(t)$, and previous values of the buffers. Although much useful information may be extracted from a simulation, we focus on one output variable:

$$\text{RATE}(t) = (\text{cumulative output up to time t})/t$$

The simulation algorithm for $t = 0, 1, ...T$ may be described as follows:

(1) START $X_i(0) = 1; i = 1, ..6; b_j(0) = 0, j = 1, ...5.$
(2) at time step t, sample $U_i(t), i = 1, ...6;$
(3) compute $X_i(t), th_i(t), i = 1, ..6, b_j(t), j = 1, ..5$, and $\Delta\text{OUT}(t)$
(4) compute TIME := TIME+1; OUT := OUT+ΔOUT, RATE = OUT/TIME
(5) if $t = T$, stop; otherwise set $t := t + 1$ and goto step 2.

This model does not describe variations caused by throttling machine running rates. Nor does it model the processing time for each machine. The pasteurizer for example, requires about 45 minutes to process one bottle. Once the pasteurizer is full of course, bottles enter and leave at the same rate and the the processing time is effectively zero. However, following each product change and buffer drainage, 45 minutes of run time will be consumed in simply filling the pasteurizer.

The model can accommodate dependencies between machine unavailabilities. One might anticipate that unavailabilities of adjacent machines would be negatively correlated, since machines will not fail while they are starved or blocked due to unavailability of a neighbor. Further one would expect such negative dependence to be strongest immediately following a buffer evacuation, and to decrease in strength as buffer contents increase. The model, however, cannot handle the temporal behavior of correlations, it can only replicate average correlations over time. One might anticipate positive correlations if machine breakdowns shared a common cause. This would also be expected to affect adjacent machines most strongly. Dependence is modelled using Markov trees and maximum entropy distributions under correlation and marginal constraints, as described in [19, 12, 2].

We remark that the model randomly draws a machine state at each time step. Hence in comparing model predictions to data, we must take account of the statistical fluctuations of model predictions. A spread of values can be generated by running the model several times with the same rates and availabilities, but with different random seeds.

Simulations have been carried out with the generic simulation package [19].

4. Zero- and Infinite-Buffer equilibria

The performance of the line, i.e. production rate, can be theoretically bounded by the zero-buffer and the infinite buffer limits [3].

With zero buffers, the line behaves as a series system: the failure of one machine brings the entire line down, and the availability of the line is the product of the individual availabilities. This assumes that the machines can fail independently; in particular, it assumes that one machine can fail while another is down. The rate of the line when all machines are up is the rate of the slowest machine. Letting P_i denote the equilibrium availability of machine i, the probability that all machines are available is (assuming independence) $\prod P_i$, and:

$$\text{Zero-buffer limit } = \min\{v_j\} \prod P_i.$$

With infinite buffers in equilibrium, the effectively slowest machine (we assume there is a unique effectively slowest machine) will never be starved or blocked; the adjacent upstream buffer will be infinitely full, and the adjacent downstream buffer will be empty. Hence:

$$\text{Infinite-buffer limit } = \min\{v_i P_i\}.$$

The definition of equilibrium is based on a Markov representation of the system. A state of the system is an assignment of values to the X_i and the b_i. The transitions from one state to another are Markov because of the exponentiality assumptions for failure and repair. Under reasonable assumptions a unique equilibrium distribution over the possible states exists independent of the starting values. This means the following:

Consider a very large number of identical copies of the line, and assign them states according to the equilibrium distribution. Now run these lines for any given number of minutes. The distribution of states over the lines after running will be statistically the same as the initial distribution.

With regard to a single line started at time 0, we say that it has reached equilibrium when the relative frequency with which it visits the possible states corresponds to the equilibrium distribution. This notion of equilibrium is intractable for us. The total number of states for the line studied here is in the order $6.3 \cdot 10^{28}$.

There is a feeling among line managers that the older filling lines have been over-designed with respect to buffer capacity. It is therefore of interest to have a simple diagnostic for equilibrium near the infinite buffer limit. We show that the correlation between the state of the effectively slowest machine and its thruput may serve as a diagnostic for the infinite buffer equilibrium. Suppressing indices, let $X(t), th(t)$ denote the state and thruput of the (unique) effectively slowest machine and let v denote the rate of X.

Proposition 4.1. *Assume that equilibrium exists and at equilibrium, $P(th(t) = v) > 0$. Let ρ denote the product moment correlation between th and X at equilibrium. Then $\rho = 1$ if and only if the equilibrium probability that X is starved or blocked is zero.*

Proof. We consider equilibrium and drop the time variable. Let E denote expectation, p denote $P(X = 1)$.

$$E(th) = E(th|X = 1)p + E(th|X = 0)(1 - p) = E(th|X = 1)p$$

since $th(t) = 0$ if $X(t) = 0$.

$$\begin{aligned} E(th \cdot X) &= E(th|X = 1)p = E(th). \\ E(th \cdot X) - E(th)E(X) &= E(th)(1 - p). \\ \sigma_{th}^2 &= E(th^2) - E(th)^2 \\ &= E(th^2|X = 1)p - E(th|X = 1)^2p^2. \\ \sigma_x^2 &= p(1 - p). \end{aligned}$$

Combining these, we find

$$\begin{aligned} \rho^2 &= \frac{(E(th \cdot X) - E(th)E(X))^2}{\sigma_{th}^2 \sigma_X^2} \\ &= \frac{E(th)^2(1 - p)^2}{(E(th^2|X = 1)p - E(th|X = 1)^2p^2)p(1 - p)} \\ &= \frac{1 - p}{\left(\frac{E(th^2|X=1)}{E(th|X=1)^2 p} - 1\right) p} \\ &= \frac{1 - p}{\frac{E(th^2|X=1)}{E(th|X=1)^2} - p} \end{aligned}$$

By Jensen's inequality, this expression equals 1 if and only if th is constant. Since $P(th = v) > 0$, this holds if and only if $P(th = v|X = 1) = 1$. This means that if X is up, the thruput takes its maximal value v. The slowest machine is never starved or blocked. □

Note that this result assumes the existence of equilibrium, and a unique slowest machine, but makes no assumption on the joint failure-repair distribution, does not require the machine states to be independent, and does not require the values of v_i to be known. If states of different machines are independent, then at infinite buffer equilibrium, the thruput of the effectively slowest machine is independent of the other machine states:

Proposition 4.2. *With the notation and assumptions of proposition 4.1, let the states of machines Y and X be independent, and let the line be at infinite buffer equilibrium, then Y and th are independent.*

Proof. As in the proof of proposition 1, at infinite buffer equilibrium,

$$\begin{aligned} P(th = 0 \text{ or } th = v) &= 1, \\ \{th = v\} &= \{X = 1\}. \end{aligned}$$

Hence $P(Y = 1 \text{ and } th = v) = P(Y = 1 \text{ and } X = 1) = P(Y = 1)P(X = 1)$ □

When the line is first started with buffers empty, it behaves somewhat like a series system. As soon as one machine is down, the entire line goes down and there is no output. If all machines are up the production rate is equal to the lowest production rate. However, the availability of the system is the product of the availabilities of the individual machines. If these availabilities were at equilibrium, then the expected output per time step is the zero buffer limit. If the machines are assumed to start in the up state then availabilities near the time origin will be higher than in the zero-buffer limit where the initial state has been "forgotten".

As a general rule, if the buffers are large enough to make the line behave like the infinite buffer line at equilibrium, then we could improve the machine with the slowest effective rate by increasing its speed, increasing its repair rate or decreasing its failure rate. However, improving the slowest effective machine pays off only up to the point where its effective rate equals that of the next-to-slowest machine. Beyond this, we must improve both machines, and so on. This general rule does not hold if the line does not reach equilibrium, or if the equilibrium is not the infinite buffer equilibrium.

5. Validation data for the filling line model

A production line model is validated by comparing model prediction with line performance. In addition to the data in table 1, prediction requires the MTTF and MTTR for each machine, the state of the buffers at t = 0, and the exact run time. The machine states are automatically logged. We compute the MTTF and MTTR for each machine as:

$$\text{MTTR} = \text{average duration of breakdown}$$
$$\text{MTTF} = \frac{(\text{total duration - MTTR} \cdot \# \text{ breakdowns})}{(\# \text{ breakdowns} + 1)}.$$

Under the exponential model, for each machine i, $\lambda_i = 1/\text{MTTF}_i$, and $\mu_i = 1/\text{MTTR}_i$.

The most difficult parameter to capture is the exact run time. Product changes, crew breaks, planned maintenance and unintended interruptions make it difficult to register run time exactly.

Validation data were gathered from a production run lasting 480 minutes with 45 minute crew break and no product change or other interruption. The Pasteurizer was full at the inception of the run, thus the run is effectively 435 minutes. The total production during this run was 188,000 bottles, or 432 bottles/minute. In addition we retrieved the joint unavailabilities of adjacent machines in order to estimate dependence. The data is shown in Tables 2, and 3. The Theoretical Availability shown in Table 2 is computed with formula 3.3. The fact that the machines begin in the up state explains why the observed availablities are a little larger than the theoretical. In general, the exponential failure-repair model adequately explains the observed availabilities. For adjacent machines, observed joint unavailabilities and the joint unavailabilities assuming independence (i.e. the product of the observed unavailatilities in Table 2) are shown in Table 3 .

At equilibrium, simple significance test can be applied to test the hypothesis that adjacent failure-repair processes are independent. The variance derived below for the joint down time is the equilibrium version of the variance derived in [21].

We consider a fixed, long observation period T of the line at equilibrium. Let P_i denote the unavailability of machine i, let N, D denote the number and total duration of joint unavailabilities of i, j, and let N_i denote the number of failures of machine i which occur during a down period of Machine $j, j \neq i$. Then $N = N_1 + N_2$ and $EN = P_1 T \lambda_2 + P_2 T \lambda_1$. If the joint unavailabilities are small in relation to T, then the processes N_1, N_2 are almost independent Poission processes. It follows that

$$VAR(N) \approx VAR(N_1) + VAR(N_2) = EN_1 + EN_2 = EN;$$
$$E(N^2) = VAR(N) + (EN)^2 \approx EN + (EN)^2.$$

Assuming independence and equilibrium, we exit a joint unavailiability with repair rate $\mu_i + \mu_j$. One can calculate

$$\begin{aligned} D &= \sum_{\#\text{ joint unavailabilities}} \text{joint repair times}; \\ ED &= \frac{EN}{\mu_1 + \mu_2}; \\ ED^2 &= \frac{2EN}{(\mu_1 + \mu_2)^2} + \frac{E(N^2 - N)}{(\mu_1 + \mu_2)^2}, \\ VAR(D) &\approx \frac{2EN}{(\mu_i + \mu_j)^2}. \end{aligned}$$

The difference between the theoretical and observed joint unavailabilities, multiplied by the total run time and divided by the standard deviation of D are shown in the last column of Table 3. The second and fourth differences are significant. [1].

We remark that the filler has the lowest effective rate: 473 bottles per minute. If the buffers are appropriately designed, the line at equilibrium should be able to produce at nearly this rate.

6. Validation Results

Simulation results giving RATE as a function of time are shown in Figure 2. Since the model involves stochastic failure and repair behavior, we must run the model

[1]This test is not possible for the data in Table 6 as the run times are not known

machine	MTTR (min)	MTTF (min)	rate (bottles/min)	Observed Availability	Observed Effective rate	Theoretical Availability
depalletizer	2.63	31.03	900	0.922	830	0.909
Rinser/filler	2.19	5.34	667	0.709	473	0.682
Pasteurizer	4.05	25.22	667	0.862	575	0.849
Labeller	1.23	19.89	800	0.899	719	0.919
Packer	0.80	1.59	912	0.665	607	0.604
Palletizer	0.52	9.85	900	0.95	855	0.898

TABLE 2. Validation data

machines	Observed joint unavailability	Joint unavailability predicted with independence	standardized difference of observed and predicted downtime
(depal,filler)	0.0085	0.0227	-0.71
(filler,pasteur)	0.0926	0.0402	2.30
(pasteur,labeller)	0.0025	0.014	-1.09
(labeller,packer)	0.0674	0.034	1.74
(packer,pallet'r)	0.0120	0.0168	-0.362

TABLE 3. joint unavailabilities

several times to get a picture of the statistical fluctuations. 20 runs are shown. Each run uses the same MTTR and MTTF but uses a different random seed. On each run the buffers are started empty. The fluctuations in the production rate $OUT(T)/T$ are caused by the random appearance of machine unavailabilities. The spread of $OUT(T)/T$ grows narrower as time increases. If the line were near equilibrium after T minutes, than the value of $OUT(T)/T$ should be nearly independent of the choice of random seed. Production rates are very low near the origin. At 435 minutes, the spread comprises 20% of the average rate. The average rate over the 20 simulated runs is 433 bottles/minute; the observed rate on this production run was 432 bottles/minute.

It will be noted that the observed rate 432 bottles/minute is lower than the maximally attainable rate with large buffers: 473 bottles/minute. This seems to suggest that the buffer sizes could be profitably increased; however Table 4 indicates that this is false. Table 4 shows the "nominal" production rate with random seed = 1 and the rates with this seed and buffers as indicated. "Filler always up" uses the nominal values, except that the failure rate of the filler is set equal to zero. We see that increasing all buffers to 10,000 has only marginal effect. Behavior near the origin causes the observed rate to fall below the infinite buffer limit. The correlation $\rho(X_2, th_2)$ is 0.99.

Figures 3 and 4 show the contents of each buffer as a function of time. Figure 3 corresponds to the nominal run in table 4. Figure 4 shows the same information, but with the unavailabilities assumed to be independent. In the infinite buffer equilibrium the buffer upstream of the core machine (buffer b_1)would be full, and

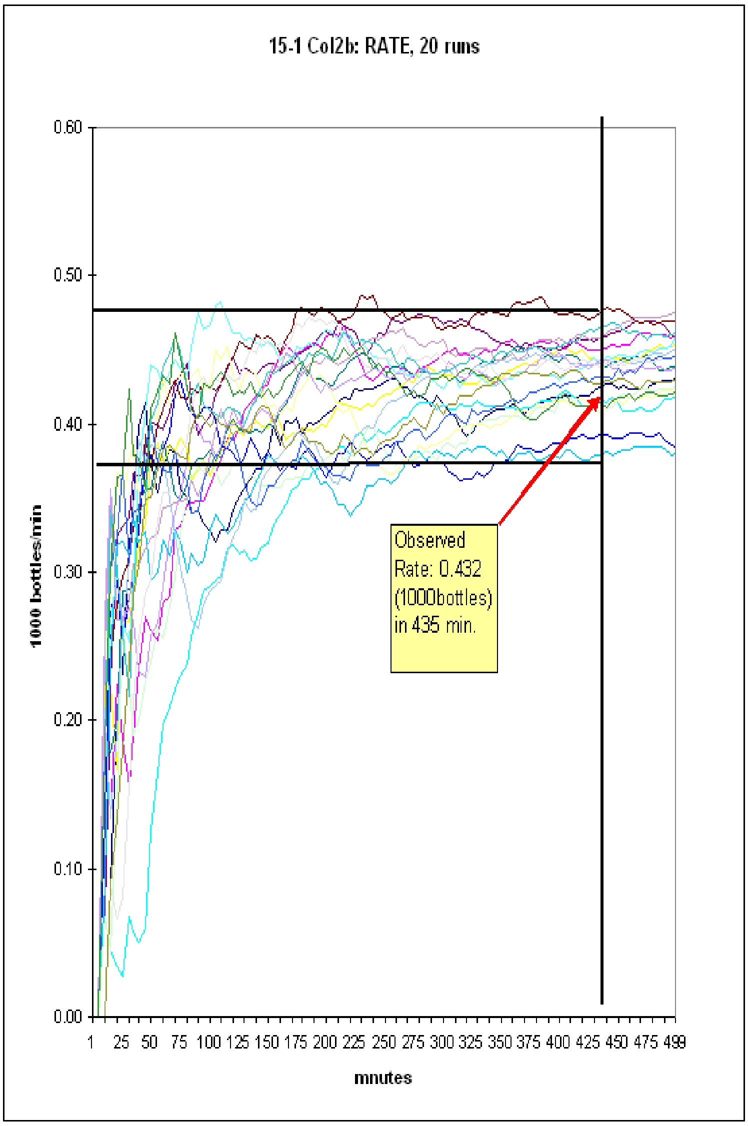

FIGURE 2. *20 simulation runs with different random seeds for validation data*

the downstream buffers would be empty. Figure 4 is markedly closer to this situation than Figure 3.

7. Week aggregate data

In practice, product changes, maintenance actions, crew breaks, meetings, etc make it difficult to assess the effective run time. It is difficult, with current data processing capabilities, to use the above filling line model in a continuous on-line fashion. It is therefore of interest to extract insights from the model applied to aggregate data. The data in Table 5 is based on one week. As we do not have

Case	Production (bottles/min)
nominal random seed	426
All buffers = 10,000	431
All buffers = 1000	328
All buffers = 3900	424
All buffers reduced by 1/2	411
Filler always up	536

TABLE 4. nominal random seed with different buffer sizes

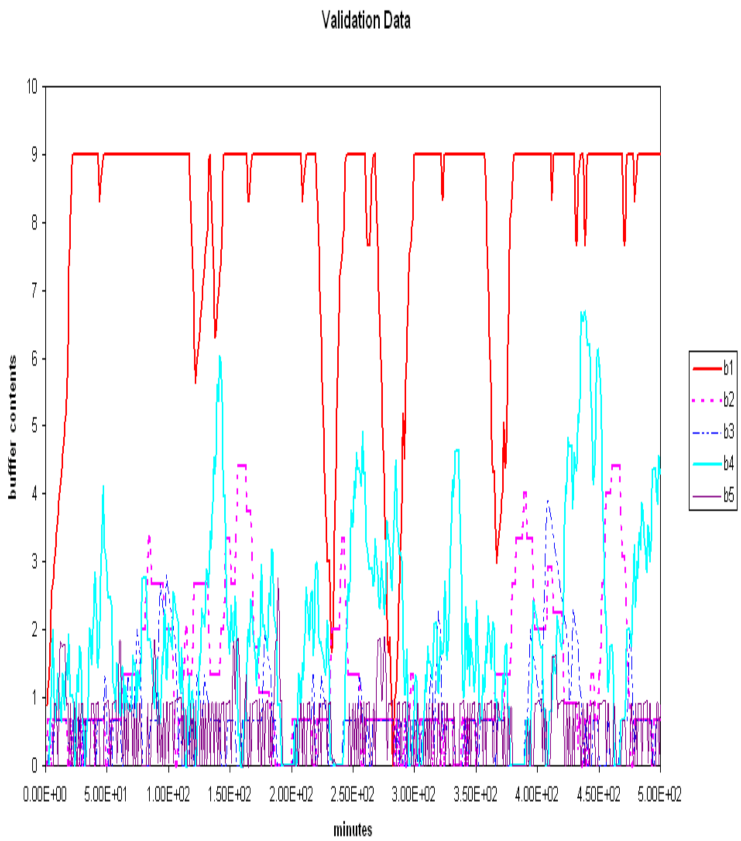

FIGURE 3. *Buffer contents against time, nominal random seed, with dependence*

information on run times during this week we cannot compare model performance with observed production rates. We note that the observed availability majorizes the theoretical availability, as in Table 2. Table 5 shows the theoretical instead of the observed effective rate, as we are interested in analysing the simulated output. Between the previous data and the aggregate data, a problem with the filler had

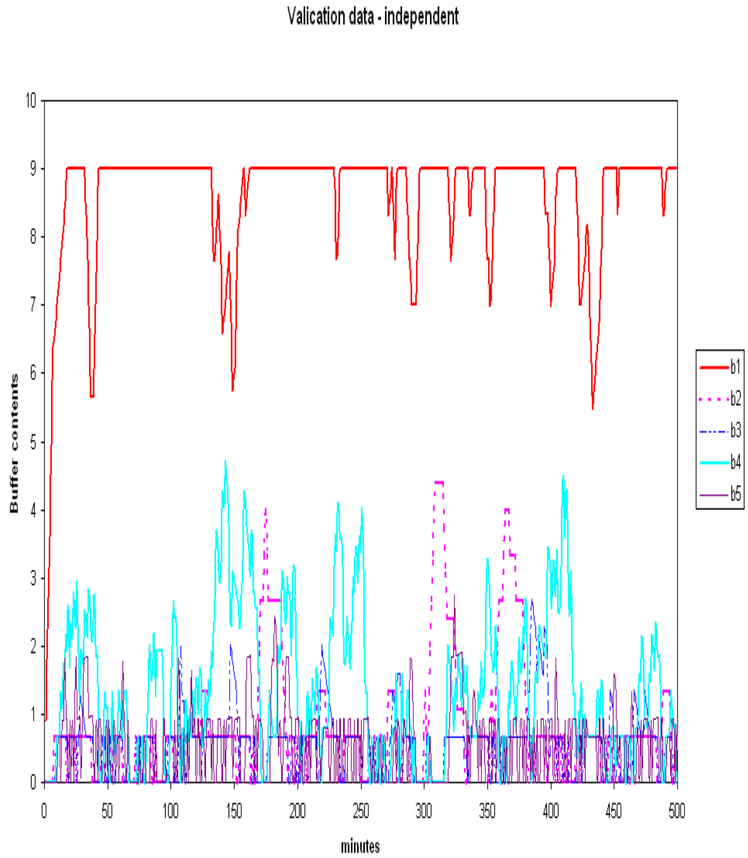

FIGURE 4. *buffer contents against time, nominal random seed, no dependence*

machine	MTTR (min)	MTTF (min)	rate (bottles/min)	Theoretical Availability	Theoretical Effective Rate	Observed Availability
depalletizer	2.51	28.06	900	0.9037	813	0.9196
Rinser/filler	1.12	9.45	667	0.8547	570	0.8960
Pasteurizer	2.45	837.6	667	0.9965	665	0.9971
Labeller	1.12	9.1	800	0.8502	680	0.8911
Packer	0.73	1.38	912	0.5913	537	0.6530
Palletizer	0.82	14.9	900	0.9156	824	0.9496

TABLE 5. Week aggregate data

been fixed. Note that the packer now has the slowest (theoretical) effective rate, and this threatens the V shape in the line's design.

	Observed joint unavailability	predicted joint unavailability with independence of observed unavailabilities
(depal,filler)	0.0045	0.00836
(filler,labeller)	0.0134	0.0113
(labeller,packer)	0.0276	0.0378
(packer,pallet'r)	0.0164	0.0175

TABLE 6. Week aggregate joint unavailabilities

Table 6 shows the observed unavailabilities. Since the pasteurizer effectively does not fail on this data set, its unavailability is independent of the other machines. In the model implemented here, a dependence has been assumed which approximates the observed joint down times of adjacent machines.

Nominal run length was chosen as 300 minutes. The production rate against time is shown for 20 simulation runs in Figure 5. Unlike the previous data set, the production rate increases slowly from 60 minutes out to 300 minutes. This suggests that the line may be slower in reaching equilibrium, although its equilibrium is higher. This can be explained by the fact that the effectively slowest machine is machine 5. Near equilibrium buffers 1...4 must all be rather full. It will take longer to fill more buffers especially since the filler's effective rate is almost as low as that of the packer. Hence, the line remains longer under the influence of the "zero buffer" regime. This is an example where unreliability disturbs the intended V design of the filling line. Figure 6 shows buffer contents against time, and should be compared to Figure 3. Note that b_2 and b_4 take some time to fill up. The packer has the lowest effective rate, and as it empties into buffer b_4, buffers b_4 and b_5 would be nearly empty in the infinite buffer equilibrium. In Figure 6, however, we see that these buffers fill to about 4000 bottles by the end of the 300 minutes. If the line's equilibrium is close to the infinite buffer equilibrium, then equilibrium is not reached in 300 minutes.

Table 7 shows the correlations between the thruput of the effectively slowest machine and the states of the other machines. Since the Pasteurizer effectively does not fail on this data, its state is constant and hence independent of everything else, and it is omitted from this comparison. The nominal run using the nominal random seed is compared to runs in which the buffer sizes are halved and quartered, and the run time is extended by a factor 10 (always with the nominal random seed). Reducing buffer sizes by one half does not have a great effect on performance. The production rate does not increase noticeably after 300 minutes. Reducing buffer sizes to one quarter the nominal size, however, does have an effect on the production rate. The correlation $\rho(th_5, X_5)$ is reduced significantly. Correlations between th_5 and other machine states is present on the 300 minute run, but die off in the 3000 minute run.

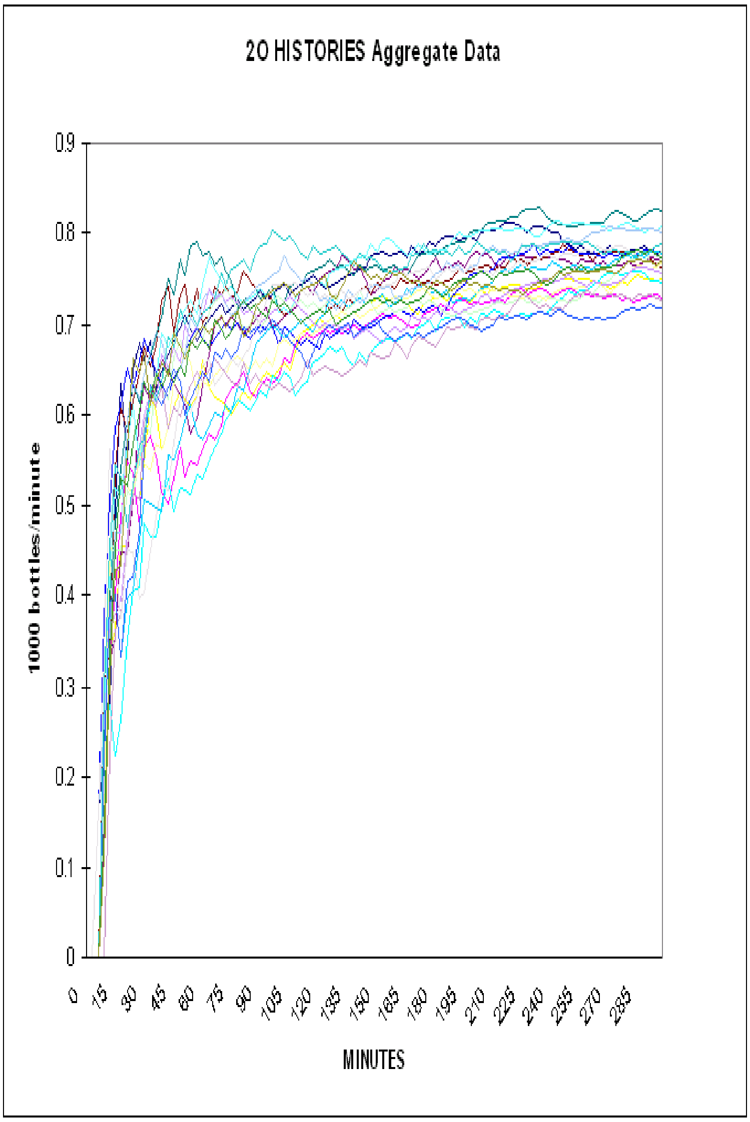

FIGURE 5. *20 simulation runs, week aggregated data*

8. Conclusions

A number of conclusions and insights obtained with this model were deemed relevant by the line managers.

(1) The average production over, say, 20 simulation runs appears to be a good predictor of line performance. Of course, this will hold only if machines run at close to their nominal rates, when they are running.

(2) The stochastic fluctuations in line output for, say, a production run of 8 hours are rather large. For $T = 435$, the spread of time-average production rates $OUT(T)/T$ in Figure 2 is about 20% of the overall average production rate. The

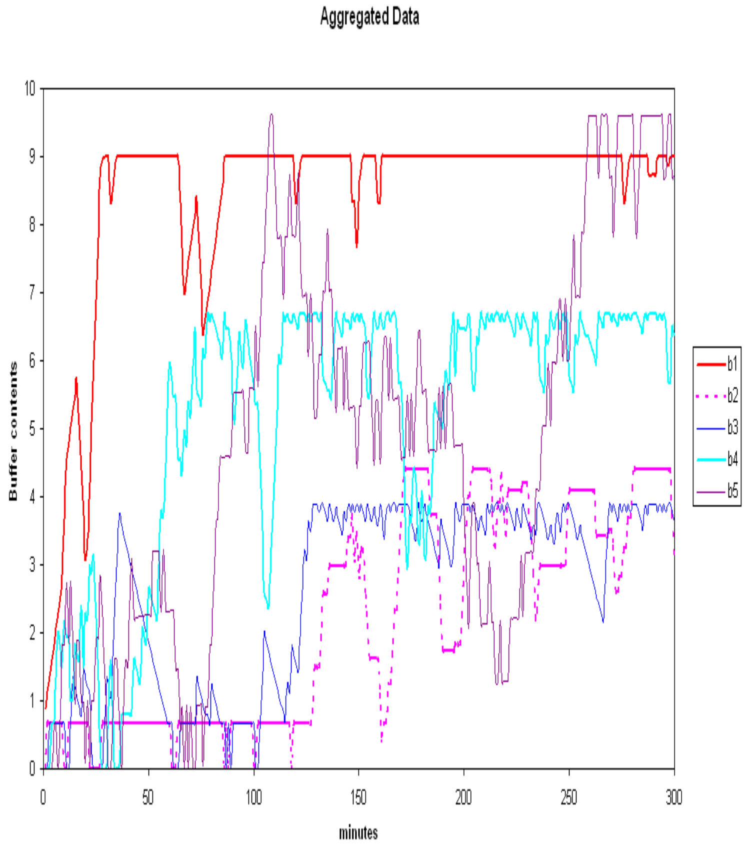

FIGURE 6. *Buffer contents against time, nominal random seed, week aggregated data*

Buffer size	$\rho(th_5, X_1)$	$\rho(th_5, X_2)$	$\rho(th_5, X_4)$	$\rho(th_5, X_5)$	$\rho(th_5, X_6)$	Rate
Nominal size, 300 min	0.05	0.03	0.11	0.98	0.05	518
Buffers half size, 300 min	-0.03	0.03	0	0.96	-0.05	520
Buffers half size, 3000 min	0.01	0	0.02	0.97	0	522
Buffers quarter size, 300 min	0.11	0.08	0.13	0.88	0.07	455
Buffers quarter size, 3000 min	0.02	0.01	0.06	0.89	0.01	465

TABLE 7. Correlations thruput slowest effective th_5 with machine states

fluctuations are largest at the beginning of a run. Appreciating the size of these fluctuations is important in assessing the performance of operator crews. Obviously, it makes no sense to reward or punish crews for fluctuations in production which fall within the noise.

(3) From the data analysed here, it seems doubtful that equilibrium is reached within typical production runs of 5 to 8 hours. If a run of T minutes were sufficient to bring the production rate $OUT(T)/T$ near its equilibrium value, then the value $OUT(T)/T$ should not depend on the random seed. Figures 2 and 5 indicate a rather strong dependence on the random seed. For production runs of this length the overall performance of the line is strongly influenced by the zero buffer regime near $T = 0$. In Figure 2 the average production rate over 20 random seeds appears stable after 180 minutes, whereas in Figure 5 the average production is increasing out to 300 minutes. This reflects the disruption in the V design caused by the packer having a lower effective rate than the filler. We can predict that upgrading the packer so as to restore the V design would enable the overall average production rate to stabilize more quickly. It must be noted that approach to equilibrium in the actual line is influenced by many factors not present in the model. In particular, the ability to adjust for up- and downstream failures by running machines above or below their nominal rates may influence the rate at which the real line approaches equilibrium.

(4) The performance of the line is somewhat less than the infinite buffer limit, but realistic increases of buffer sizes would not seem to help. Indeed, the model confirms the engineer's suspicion that the buffer sizes could be reduced without seriously degrading performance. Of course the infinite (and zero) buffer limit are influenced by the availabilities. Large gains in productivity could be achieved by enhancing the availabilities.

(5) Comparing the validation dataset to the week aggregate dataset we see significant variations in line and machine performance. Moreover, the unavailabilities of individual machines may disturb the line's V design and delay reaching equilibrium. This is further aggravated if the core machine is not near the beginning of the line.

(6) The correlation of availability to thruput of the core machine is a good index of distance to the infinite buffer equilibrium, which does not depend on the actual production rate. Using this criterion, we can compare distances to infinite-buffer equilibrium in lines with different rates.

References

1. J. Schmidt A. Soyster and M. Rohrer, *Allocation of buffer capacities for a class of fixed cycle production lines*, AIIE Trans. **11** (1979), no. 2, 140–146.
2. T. Bedford and R.M. Cooke, *Vines-a new graphical model for dependent random variables*, Annals of Statistics **30** (2002), no. 4, 1031–1068.
3. J. Buzacott, *Automatic transfer lines with buffer stocks*, Int. J. Prod. Res. **5** (1967), no. 1, 183–200.
4. ______, *The role of inventory banks in flow-line production systems*, Int J. Prod. Res. **9** (1971), no. 4, 425–436.
5. J. Buzacott and J. Shantikumar, *Stochastic models of manufacturing systems*, Prentice Hall, Englewood Cliffs, 1993.
6. Y. Dallery and S. Gershwin, *Manufacturing flow line systems: a review of models and analytical results*, Queuing Systems **12** (1992), 3–94.
7. E.Elsayed and R. Turley, *Reliability analysis of production systems with buffer storage*, Queuing Systems **12** (1980), 3–94.
8. G.Haines, *The design and layout of packaging lines*, The Brewer **March** (1995), 92–97.

9. K.Funaki and K Yoshimoto, *Distribution of total uptime during a given time interval*, IEEE Trans. Reliability **43** (1994), 145–161.
10. H. Kronseder and P. Hinterwimmer, *Problems arising within a bottle line*, MBAA Technical Quarterly **15** (1978), no. 1, 43–52.
11. H. Manning, *Computer simulation of packaging lines*, Brewery Age **August** (1985), 10–24.
12. A.M.H. Meeuwissen and T.J. Bedford, *Minimally informative distributions with given rank correlation for use in uncertainty analysis*, J. statist. comput. Simulation **57** (1997), 143–175.
13. R. Murphy, *Examining the distribution of buffer protection*, AIIE Trans. **11** (1979), no. 2, 113–120.
14. J. Muth, *A method for predicting system downtime*, IEEE Trans. Reliability **R-17** (1968), 97–102.
15. J. Muth and J. White, *Conveyer theory, a survey*, AIIE Trans. **11** (1978), no. 4, 270–277.
16. K. Okamura and H. Yamishina, *Justification for installing buffer stocks in unbalanced two stage automatic transfer lines*, AIIE Trans. **11** (1979), 308–312.
17. S. Panwalker and M. Smith, *A predictive equation for average output of k stage series systems with finite interstage queues*, AIIE Trans. **11** (1979), 136–139.
18. A. Renyi, *On the asymptotic distribution of the sum of a random number of independent random variables*, Acta Mathematica Academiae scientiarum Hungaricae **8** (1957), 193–199.
19. R.M.Cooke, *Unicorn: Methods and code for uncertainty analysis*, Tech. report, AEA Technologies, 1995.
20. R. Subramanian S.K. Scrinivasan and K.S. Ramesh, *Mixing of two renewal processes and its applications to reliability theory*, IEEE Trans. reliability **R-20** (1971), 51–55.
21. Suyono, *Renewal processes and repairable systems*, Tech. report, Dept. of Mathematics, Delft University of Technology, 1995.
22. L. Takacs, *On certain sojourn time problems in the theory of stochastic processes*, Acta Mathematica Academiae scientiarum Hungaricae **6** (1957), 169–191.
23. J. Wijngaard, *The effect of interstage buffer storage on the output of two unreliable production units in series, with different production rates*, AIIE Trans. **11** (1979), no. 1, 42–47.
24. T. Yokoi, *Designing higher productivity packaging lines*, Bottling and Canning **October** (1992), 34–38.
25. B. Zimmern, *Etudes de la propagation des arrets aleatoires dans les chaines de production*, Rev, Statist. Appl. **4** (1956), 85–104.

Department of Mathematics, Delft University of Technology, Delft,, The Netherlands
E-mail address: r.m.cooke@its.tudelft.nl

Heineken Technical Services

Heineken Technical Services