IOSR Journal of Mathematics (IOSR-JM)
e-ISSN: 2278-0661, p-ISSN: 2278-8727, Volume 14, Issue 3, Ser. 1 (May. – June. 2018), PP 93-100
www.iosrjournals.org

Smart Factory Maintenance: Balancing Spare Inventory and Repair Time

K.P.S. Baghel
Govt. Degree College Manikpur, Chitrakoot (U.P.)

Abstract
A smart factory can see a failure coming, but it can't repair one without a part and a technician. That is the tension this article works through. Spare units sitting on a shelf cost money. A thin repair crew leaves broken machines waiting. Both choices pull on the same quantity, which is the fraction of time a production line has enough working machines to meet its target. I treat the problem with finite-source queueing theory. A birth–death model of M operating machines, S cold spares and R repairmen gives closed-form steady-state probabilities, availability, throughput, and mean time a failed unit spends in the repair system. I then extend the model with balking, reneging, two failure modes, a threshold repairman and an N-policy start-up rule. These extensions follow work on machine repair problems by Jain and co-authors. On the inventory side, Palm's theorem and the base-stock logic of Sherbrooke's METRIC model give a marginal rule: add a spare while the expected-backorder reduction is worth more than its holding cost. The two views join in a single cost function over (S, R). A small worked example, computed from the model rather than taken from plant records, shows that the first spare lifts availability from 0.769 to 0.897, a second repairman lifts it to only 0.800, and a second spare reaches 0.951. The same example shows a side effect that is easy to miss: more spares raise the load on the repair shop, so each failed unit waits longer. I close by showing how prognostic information from condition monitoring can enter the model as a state-dependent failure rate. That would let the spare level and the repair capacity follow the health of the machines rather than stay fixed.

Keywords: *machine repair problem, spare provisioning, finite-source queue, repair capacity, smart manufacturing, availability*

I. Introduction

1.1 Background

Maintenance is a big line item in manufacturing, and it has been for a long time. Mobley (2002) reports that maintenance accounts for roughly 15 to 60 percent of the cost of goods produced, depending on the industry. That range is wide because it hides two very different kinds of waste. One is the money tied up in parts nobody has used yet. The other is the money lost while a machine sits dead.

Smart-factory programs promise to cut both. The German Industrie 4.0 working group framed the goal as networked machines that report their own condition and plan around it (Kagermann et al., 2013). Lee et al. (2014) described the analytics layer that would turn that stream of data into decisions. The condition-monitoring literature has been building the sensing side for decades (Jardine et al., 2006).

Better data doesn't remove the physical constraint, though. When a spindle bearing fails, someone has to have a replacement bearing and the hands to fit it. Prediction moves the failure from a surprise to a scheduled event. The inventory and the labor still have to exist.

1.2 The Problem and What This Article Does

I want to put the trade-off in plain mathematical terms. Let S be the number of spare units held and R the number of repairmen. Raising S costs holding capital. Raising R costs wages. Both reduce downtime, but by different mechanisms and with different diminishing returns. A spare substitutes for a broken machine immediately. A repairman shortens the time until the broken machine comes back.

Finite-source queueing is the natural language for this. The population of machines is small, so the arrival rate of failures depends on how many are currently running. Haque and Armstrong (2007) survey this "machine interference" literature, and the Indian school of machine-repair modeling has pushed it far (Chikara et al., 2006; Jain, 2001; Jain & Baghel, 2001; Jain et al., 2000, 2002, 2004).

The plan is as follows. Section 2 reviews the models I build on. Section 3 states the base model and derives its performance measures. Section 4 adds realistic features. Section 5 treats the spare pool as an inventory problem. Section 6 combines the two and works a numerical example. Section 7 connects the result to prognostics.

II. Related Work

2.1 Machine Repair and Interference

The classical machine repair problem has a fixed set of machines, each failing independently, and one or more repairmen. Machines wait when all repairmen are busy. Gross et al. (2008) give the standard treatment. Haque and Armstrong (2007) trace how the problem grew from textile looms to modern flexible systems, and they note how many variants remain without closed forms.

Jain and co-authors extended the base model along several lines. Jain et al. (2000) added balking, reneging, spares and an additional repairman to an M/M/c/K/N system. Jain and Baghel (2001) handled a multi-component system with state-dependent rates. Jain (2001) allowed two modes of failure. Jain et al. (2002) predicted the performance of a flexible manufacturing system. Jain et al. (2004) moved to M/G/1 repair with mixed standby types, degraded failure, Bernoulli feedback and an N-policy. Chikara et al. (2006) made arrival rates controllable and interdependent.

2.2 Spares as Inventory

The inventory view started separately. Palm (1938) showed that when units arrive for repair as a Poisson stream and there are always enough servers, the number in repair is Poisson. Feeney and Sherbrooke (1966) generalized the (S−1, S) policy to compound Poisson demand. Sherbrooke (1968) built METRIC, a multi-echelon model for recoverable items, and Muckstadt (1973) extended it to multiple indenture levels. Barlow and Proschan (1965) supply the reliability foundations underneath all of this (Baghel, 2013; 2014; 2016; 2017).

Table 1 sets the main models side by side. It shows what each one adds and which decision variable it speaks to.

Table 1: Summary of Machine-Repair and Spare-Provisioning Models Used in This Article

Source	Model type	Feature added	Decision variable addressed
Palm (1938)	Infinite-server repair	Poisson number in repair	Spare level S
Feeney & Sherbrooke (1966)	(S−1, S) inventory	Compound Poisson demand	Spare level S
Sherbrooke (1968)	METRIC, multi-echelon	Expected backorders as objective	Spare level S
Jain et al. (2000)	M/M/c/K/N	Balking, reneging, spares, extra repairman	S and R
Jain & Baghel (2001)	Multi-component	State-dependent rates	S
Jain (2001)	Finite source	Two failure modes	R, repair mix
Jain et al. (2002)	Flexible manufacturing system	Performance prediction	S and R
Jain et al. (2004)	M/G/1	Mixed standbys, N-policy, feedback	R, start-up threshold
Chikara et al. (2006)	Interdependent finite source	Controllable arrival rates	S

Reading the table, one gap stands out. The inventory models treat repair capacity as given, and the repair models often treat spares as an afterthought. The cost function in Section 6 tries to close that gap.

III. The Base Model

3.1 Assumptions and State Space

Take a production line that needs M machines running to meet its target. The plant keeps S identical cold spares, so the total population is N = M + S. A cold spare does not fail while idle. Each operating machine fails at rate λ, with exponential times between failures. R repairmen each repair at rate μ, with exponential repair times. A failed unit joins the repair queue and, once fixed, goes back into the pool, becoming operating or spare as needed.

Let n be the number of units in the repair system, running from 0 to N. Then N − n units are working or idle in stock, and the number actually operating is min(M, N − n). Figure 1 shows the resulting birth–death chain, which is the object everything else in this section hangs on.

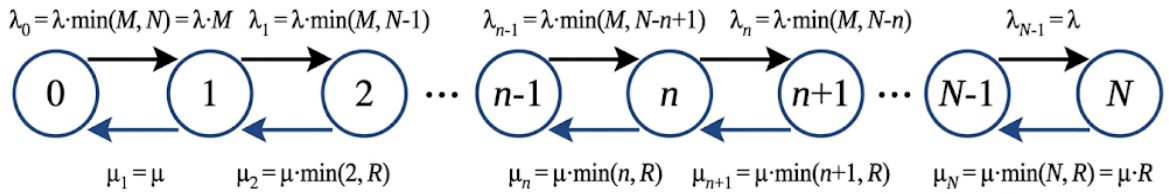


Figure 1: State-Transition Diagram of the Finite-Source Machine Repair Model with Spares

The figure shows states 0 through N arranged in a row, where state n means n units are failed and in the repair system. Arrows pointing right carry the failure rate λ_n = λ·min(M, N − n), and arrows pointing left carry the repair rate μ_n = μ·min(n, R). The key insight is that failure rates stay flat until the spare pool is used up, after which they fall step by step, while repair rates plateau once all R repairmen are busy. This is the standard birth–death structure for finite-source queues as presented in Gross et al. (2008).

3.2 Steady-State Probabilities

The state-dependent rates are

$$\lambda_n = \lambda \cdot \min(M,;N-n)\,, \qquad \mu_n = \mu \cdot \min(n,;R)\,.$$

Balance across each cut between state n−1 and state n gives $\lambda_{n-1}p_{n-1} = \mu_n p_n$. Iterating,

$$p_n = p_0 \prod_{k=1}^{n} \frac{\lambda_{k-1}}{\mu_k}, \qquad n = 1, \dots, N.$$

Normalization fixes the remaining constant:

$$p_0 = \left[,1 + \sum_{n=1}^{N} \prod_{k=1}^{n} \frac{\lambda_{k-1}}{\mu_k}\right]^{-1}.$$

It helps to split the product into regimes. For $n \le min(S,R)$ the product is $(M\lambda/\mu)^{\mathrm{n}}/n!$. Once n passes S, the factor Mλ is replaced by $(N-k+1)\lambda$, so the chain starts to thin out. Once n passes R, the denominator stops growing like n! and grows like $R! \cdot R^{\mathrm{n-R}}$instead. Those two kinks are exactly where a spare or a repairman begins to matter.

3.3 Performance Measures

Four quantities answer most plant questions. The expected number of operating machines is

$$E[O] = \sum_{n=0}^{N} \min(M,;N-n)\,,p_n,$$

and the line availability, meaning the average fraction of required capacity actually running, is

$$A = \frac{E[O]}{M}.$$

The expected shortage, in machines below target, is

$$E[B] = \sum_{n=S+1}^{N} (n-S)\,,p_n = M(1-A).$$

The effective failure rate, which equals throughput of the repair shop, is

$$\lambda_{\text{eff}} = \sum_{n=0}^{N} \lambda_n\, p_n = \mu \sum_{n=1}^{N} \min(n,R)\,,p_n.$$

With L = Σ n p_n, Little's (1961) law gives the mean time a failed unit spends in the repair system, queueing plus service:

$$W = \frac{L}{\lambda_{\text{eff}}}.$$

I'd flag W as the number most plants forget to track. It is repair time as the machine experiences it, and it is not the same as the technician's hands-on time. Section 6 will show it can move in the opposite direction from availability.

IV. Extensions Toward Realistic Shop Floors

4.1 Balking, Reneging and a Threshold Repairman

Real maintenance queues leak. A failed unit may be diverted to an outside vendor at once (balking) or pulled out of the queue after waiting too long (reneging). Jain et al. (2000) modeled both. Let $b_n$be the probability that a failure arriving in state n actually enters the queue, and let ξ be the reneging rate per waiting unit. Then

$$\lambda_n = b_n, \lambda \min(M, ; N - n), \qquad \mu_n = \mu \min(n, R) + \xi, (n - R)^+.$$

The same paper adds an extra repairman who joins the crew whenever the queue reaches a threshold K. In rates,

$$\mu_n = \{\mu \min(n, R), \quad n < K, \mu \min(n, R + 1), \quad n \geq K.$$

The product-form solution of Section 3.2 survives untouched. Only the rates inside the product change. That is the practical appeal of the birth–death approach. A plant engineer can swap in a measured $b_n$curve and rerun the same recursion.

4.2 Two Modes of Failure

Many machines fail in two quite different ways. A tool breaks in minutes to fix, or a gearbox fails and takes days. Jain (2001) treated a system with two failure modes, with rates λ_1 and λ_2 and repair rates μ_1 and μ_2. The total failure rate is $\lambda = \lambda_1 + \lambda_2$. A failed unit is of type 1 with probability $q = \lambda_1/\lambda$, so the mean repair time is the mixture

$$E[T] = \frac{q}{\mu_1} + \frac{1 - q}{\mu_2}.$$

The second moment matters just as much, because repair-time variability drives queueing delay:

$$E[T^2] = \frac{2q}{\mu_1^2} + \frac{2(1 - q)}{\mu_2^2}.$$

If μ_1 is ten times μ_2, the squared coefficient of variation of T is well above 1, so an exponential repair assumption will understate waiting. This is the strongest argument for moving to the M/G/1 form below.

4.3 General Repair Times and the N-Policy

For a single repairman with general repair times, the Pollaczek–Khinchine formula gives the mean wait in queue under Poisson arrivals at rate λ (Gross et al., 2008):

$$W_q = \frac{\lambda, E[T^2]}{2, (1 - \rho)}, \qquad \rho = \lambda E[T].$$

Jain et al. (2004) added an N-policy, where the repairman stays away until N failed units have accumulated and then works until the system is empty. The mean number in the system rises by (N − 1)/2 over the ordinary M/G/1 value:

$$L(N) = \rho + \frac{\lambda^2 E[T^2]}{2(1 - \rho)} + \frac{N - 1}{2}.$$

Why would anyone choose such a policy? Because starting the repair crew has a cost, say K per busy cycle for travel, tooling or line changeover. Busy cycles occur at rate λ/N. With a holding cost h per failed unit per unit time, the part of the cost that depends on N is

$$f(N) = h, \frac{N - 1}{2} + \frac{K\lambda}{N}.$$

Setting f′(N) = 0 gives the threshold

$$N^* = \sqrt{\frac{2K\lambda}{h}}.$$

Round it to the nearest integer and check both neighbors. The formula is blunt, but it shows how batching repairs trades start-up cost against idle broken machines. In a smart factory, where a repair visit might be a robot or a technician pulled off another job, K is rarely zero.

V. The Spare Pool as an Inventory Problem

5.1 Palm's Theorem and Expected Backorders

Suppose repair capacity is large enough that nobody waits for a repairman. Then each failed unit spends an average time T in repair, and the number in repair is the number in an infinite-server queue. Palm (1938) showed this number is Poisson with mean $\rho = \lambda_f$T, where λ_f is the aggregate failure rate:

$$P(X = x) = e^{-\rho}, \frac{\rho^x}{x!}.$$

Under a base-stock (S−1, S) policy in the sense of Feeney and Sherbrooke (1966), a machine is short whenever more than S units are in repair. The expected number of backorders is therefore

$$\mathrm{EBO}(S) = \sum_{x=S+1}^{\infty} (x - S)\,, P(X = x).$$

Sherbrooke (1968) used exactly this quantity as the performance measure in METRIC.

5.2 The Marginal Rule for Spares

Take two consecutive spare levels and subtract:

$$\mathrm{EBO}(S) - \mathrm{EBO}(S + 1) = P(X > S).$$

That is neat. The benefit of the (S+1)-th spare is the probability that demand in repair exceeds S. Let $c_s$be the holding cost of one spare per unit time and c_b the cost of one backordered machine per unit time. Minimizing $c_s \cdot S + c_b \cdot EBO(S)$ gives a newsvendor-type rule. Add a spare as long as $c_b P(X > S) > c_s$, so the optimal level is the smallest S satisfying

$$P(X \le S) \ge 1 - \frac{c_s}{c_b}.$$

Notice what this does and doesn't say. It says the right amount of stock scales with the cost ratio and with the repair-time load ρ, both of which a plant can estimate. It does not say anything about R, because Palm's setting assumes repair capacity is never the bottleneck. That assumption is where the finite-source model of Section 3 earns its keep. Once R is small, repair times stretch, ρ rises, and the Poisson mean in the rule above is no longer a constant.

VI. Joint Optimization of Spares and Repair Capacity

6.1 The Cost Function

Pull the two decisions into one expression. With $c_s$per spare, c_r per repairman and c_d per machine-unit of shortage per unit time, the long-run cost rate is

$$C(S, R) = c_s S + c_r R + c_d, E[B](S, R),$$

where E[B] is computed from the steady-state distribution in Section 3. The shortage term is the only place where the queue matters, and it depends on both S and R. A discrete search over small grids of S and R is enough, since both are single-digit integers in most cells.

The first-order intuition comes from comparing marginal gains. Adding a spare changes E[B] by

$$\Delta_S = E[B](S,R) - E[B](S+1,R),$$

and adding a repairman changes it by

$$\Delta_R = E[B](S,R) - E[B](S,R+1).$$

The better lever is the one with the larger ratio of Δ to unit cost, Δ_S/c_s against Δ_R/c_r. Neither lever wins in general. It depends on how close the repair shop is to saturation.

6.2 A Worked Illustration

The numbers below are not plant data. I computed them directly from the model of Section 3 to show how the mechanics behave. Take M = 2 machines required, failure rate λ = 1 per machine per time unit, and repair rate μ = 4 per repairman. Four configurations follow.

Start with S = 0, R = 1. Here N = 2, with $\lambda_0 = 2$ and $\lambda_1 = 1$. The recursion gives unnormalized weights 1, 0.5 and 0.125, so p = (8/13, 4/13, 1/13). Then E[O] = 20/13 and A = 10/13 ≈ 0.769.

Add one spare, S = 1, R = 1. Now N = 3, with weights 1, 0.5, 0.25 and 0.0625 (sum 1.8125). The probabilities are about (0.552, 0.276, 0.138, 0.034), E[O] = 52/29 ≈ 1.793 and A ≈ 0.897.

Add one repairman instead, S = 0, R = 2. The weights are 1, 0.5 and 0.0625, so p = (0.64, 0.32, 0.04). Then E[O] = 1.6 and A = 0.800.

Add two spares, S = 2, R = 1. With N = 4, the weights are 1, 0.5, 0.25, 0.125 and 0.03125 (sum 1.90625). Then E[O] ≈ 1.902 and A ≈ 0.951.

Figure 2 plots these four availabilities.

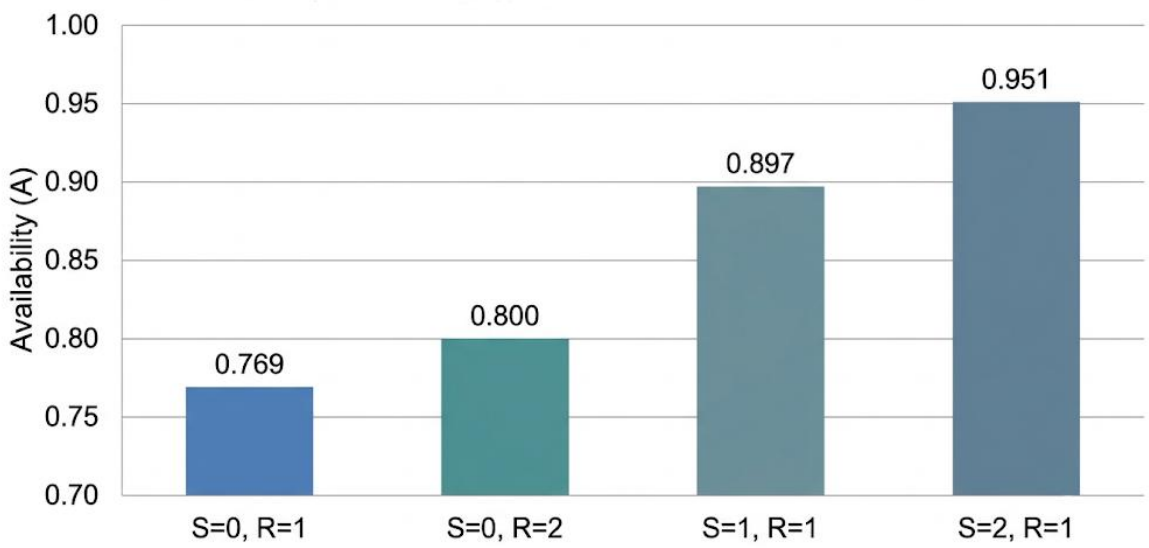


Figure 2: Computed Line Availability for Four Spare–Repairman Configurations (M = 2, λ = 1, μ = 4)

A bar chart whose horizontal axis lists four configurations (S=0,R=1; S=0,R=2; S=1,R=1; S=2,R=1) and whose vertical axis shows availability A from 0.70 to 1.00. The bars read approximately 0.769, 0.800, 0.897 and 0.951 respectively. The values are illustrative results computed by the author from the birth–death model in Gross et al. (2008) and are not measured plant data. The takeaway is that a first spare gains far more than a second repairman at this load level.

6.3 What the Example Teaches

The first spare adds about 0.128 to availability. The second repairman adds about 0.031. A second spare adds a further 0.054. For this load, with failures at a quarter of repair speed per unit, spares dominate. That outcome is not universal. If μ were close to λ, the repair queue would be the problem and repairmen would win.

There's also a side effect. Compute $W = L/\lambda_e ff$ for each case. With S = 0, R = 1, I get L = 6/13 ≈ 0.462 and W = 0.300. With S = 1, R = 1, L ≈ 0.655 and W ≈ 0.365. With S = 2, R = 1, L ≈ 0.787 and W ≈ 0.414. With S = 0, R = 2, L = 0.400 and W = 0.250.

Spares raise availability and lengthen the mean time in repair. The reason is plain once stated. A spare keeps the line producing, so more machines stay in the failure-prone pool, so the repair shop sees more traffic. If a plant measures only mean repair time, it may conclude that adding spares made maintenance worse. Availability says the opposite. I'd treat W as a diagnostic and A as the objective.

VII. Connecting the Model to Prognostics

7.1 From Fixed Rates to Health-Dependent Rates

The base model assumes a constant λ. Smart-factory sensors give something better. Barlow and Proschan (1965) describe wear-out behavior through an increasing hazard rate. A common form is the Weibull hazard

$$h(t) = \frac{\beta}{\eta}\left(\frac{t}{\eta}\right)^{\beta-1}, \qquad \beta > 1,$$

where β is the shape and η the scale. Condition-monitoring systems estimate the machine's position on that curve (Jardine et al., 2006). Wang (2002) surveys the maintenance policies that act on such estimates.

Suppose each operating machine has a health index and the plant can group machines into a few degradation classes, each with its own failure rate $\lambda^{(k)}$. The chain of Section 3 then gains a second dimension, the count of machines in each class. The size of the state space grows, but the structure stays a continuous-time Markov chain, and the same balance equations apply.

7.2 Prognostics and Spare Ordering

A forecast helps most through lead time. If a spare has replenishment lead time L_t, then an order should be triggered when the predicted time to failure drops below L_t. Let F be the cumulative distribution of the remaining useful life estimate. The probability that a machine fails before its spare arrives is $F(L_t)$. With m machines under watch, the expected number of unprotected failures during one lead time is $m \cdot F(L_t)$. The base-stock level in Section 5.2 can then be recomputed with that forecast demand in place of the long-run mean. Lee et al. (2014) describe the analytics that would supply F.

7.3 Prognostics and Repair Scheduling

Forecasts also help with labor. If failures can be predicted a few shifts ahead, repairs can be batched. That connects directly to the N-policy of Section 4.3. A start-up cost K that includes a line stoppage is smaller when the stoppage is planned. A smaller K lowers N* in the square-root rule, so the crew is called sooner and fewer failed units wait.

VIII. Practical Notes and Limits

A few cautions seem worth stating plainly.

First, exponential assumptions are convenient, not true. Real repair times have heavier tails than exponential, as the two-mode example in Section 4.2 shows. The M/G/1 results are safer for decisions about a single crew, and Jain et al. (2004) offer a framework for it.

Second, cold-spare behavior is an idealization. Spares sitting in a humid store room can degrade. Jain and Baghel (2001) handle state-dependent rates, which is one way to model warm standby, and a plant with that problem should use them.

Third, the parameters have to come from somewhere. Estimating λ and μ from maintenance logs is easy to say and hard to do, because logs often record the time a work order closed and not the time a repair ended. If the data are poor, the nice closed forms give false comfort. I'd rather run the model with a plausible range of inputs and see whether the ranking of configurations holds.

Fourth, the cost coefficients c_s, c_r and c_d are where most of the real argument sits. A mathematically optimal (S, R) under a wrong c_d is just a precisely wrong answer. Plant managers should spend more time pinning down the cost of an idle machine than polishing the queueing model.

IX. Conclusion

The question was how many spares to hold and how much repair capacity to staff. A finite-source birth–death model answers it with a few lines of algebra. The steady-state probabilities follow from the product form, and availability, shortage and mean time in repair fall out as sums over those probabilities.

Three results stand out. One is the marginal rule from Palm's theorem: the benefit of the next spare equals P(X > S), so the spare level should rise until that probability falls below the cost ratio c_s/c_b. Another is the cost function C(S, R), which forces the spare and repair decisions to be made together, because each changes the other's value. The third is the N-policy threshold $N * = \sqrt{(2K\lambda/h)}$, which shows how the start-up cost of a repair visit sets the right batch size.

The worked example computed from the model, with M = 2, λ = 1 and μ = 4, illustrates the trade-off. A first spare moved availability from 0.769 to 0.897, a second repairman reached only 0.800, and two spares reached 0.951. Mean time in the repair system rose with each added spare, so repair time alone is a misleading indicator of maintenance health. These figures are computed illustrations and aren't plant measurements.

The extensions by Jain and co-authors, covering balking, reneging, threshold repairmen, two failure modes and general repair times, keep the same structure and make the model fit real shops better. Prognostics add the missing ingredient for a smart factory, a failure rate that follows machine health. The next step I'd want to see is an empirical one. Fit λ, μ and the cost coefficients from a real plant's records, solve for (S, R), and compare against what the plant actually does today. The mathematics is ready for that test. The data usually isn't, and that is where the work remains.

References

[1]. Baghel, K. P. S. (2013). Generalists vs. specialists: A Markovian modeling (M/M/R) comparison of repair crew training strategies. Journal of Research in Applied Mathematics, 1(1), 10–15. https://doi.org/10.35629/0743-01011015

[2]. Baghel, K. P. S. (2014). Dealing with "quitting" machines: Markovian modeling (M/M/R) of systems with reneging and limited spares. International Journal of Mathematics and Statistics Invention, 2(1), 56–61. https://doi.org/10.35629/4767-02015661

[3]. Baghel, K. P. S. (2016). Back in action: How spare parts keep broken factories moving. International Journal of Mathematics and Statistics Invention, 4(5), 59–66. https://doi.org/10.35629/4767-04055966

[4]. Baghel, K. P. S. (2017). Preventive vs. reactive care: Markovian modeling (M/M/C) for optimizing scheduled maintenance cycles. International Journal of Mathematics and Statistics Invention, 6(4), 26–31. https://doi.org/10.35629/4767-06042631

[5]. Barlow, R. E., &Proschan, F. (1965). *Mathematical theory of reliability*. Wiley.

[6]. Chikara, D., Jain, M., & Baghel, K. P. S. (2006). Interdependent machine repair problem with controllable arrival rates and spares. *Acta Ciencia Indica. Mathematics, 32*(2), 569–574.

[7]. Feeney, G. J., & Sherbrooke, C. C. (1966). The (S−1, S) inventory policy under compound Poisson demand. *Management Science, 12*(5), 391–411.

[8]. Gross, D., Shortle, J. F., Thompson, J. M., & Harris, C. M. (2008). *Fundamentals of queueing theory* (4th ed.). Wiley.

[9]. Haque, L., & Armstrong, M. J. (2007). A survey of the machine interference problem. *European Journal of Operational Research, 179*(2), 469–482.

[10]. Jain, M. (2001). Machine repair problem with spares, reneging, additional repairman and two modes of failure. *Journal of MACT, 33*, 69–79.

[11]. Jain, M., & Baghel, K. P. S. (2001). A multi-component repairable system with spares and state-dependent rates. *The Nepali Mathematical Sciences Report, 19*, 81–90.

[12]. Jain, M., Sharma, G. C., & Baghel, K. P. S. (2002). Performance prediction of a flexible manufacturing system. *IJE Transactions A: Basics, 15*(4), 355–368.

[13]. Jain, M., Sharma, G. C., & Baghel, K. P. S. (2004). N-policy for M/G/1 machine repair model with mixed standby components, degraded failure and Bernoulli feedback. *IJE Transactions B: Applications, 17*(3), 265–277.

[14]. Jain, M., Singh, M., & Baghel, K. P. S. (2000). M/M/c/K/N machine repair problem with balking, reneging, spares and additional repairman. *Gujarat Statistical Review, 26–27*, 49–60.

[15]. Jardine, A. K. S., Lin, D., &Banjevic, D. (2006). A review on machinery diagnostics and prognostics implemented on a condition-based maintenance. *Mechanical Systems and Signal Processing, 20*(7), 1483–1510.

[16]. Kagermann, H., Wahlster, W., & Helbig, J. (2013). *Recommendations for implementing the strategic initiative INDUSTRIE 4.0: Final report of the Industrie 4.0 Working Group*. acatech – National Academy of Science and Engineering.

[17]. Lee, J., Kao, H.-A., & Yang, S. (2014). Service innovation and smart analytics for Industry 4.0 and big data environment. *Procedia CIRP, 16*, 3–8.

[18]. Little, J. D. C. (1961). A proof for the queuing formula: L = λW. *Operations Research, 9*(3), 383–387.

[19]. Mobley, R. K. (2002). *An introduction to predictive maintenance* (2nd ed.). Butterworth-Heinemann.

[20]. Muckstadt, J. A. (1973). A model for a multi-item, multi-echelon, multi-indenture inventory system. *Management Science, 20*(4), 472–481.

[21]. Palm, C. (1938). Analysis of the Erlang traffic formula for busy-signal arrangements. *Ericsson Technics, 5*, 39–58.

[22]. Sherbrooke, C. C. (1968). METRIC: A multi-echelon technique for recoverable item control. *Operations Research, 16*(1), 122–141.

[23]. Wang, H. (2002). A survey of maintenance policies of deteriorating systems. *European Journal of Operational Research, 139*(3), 469–489.