Cost-effective Inspection and Maintenance Rule for Train Control Beacons

Pierre Dersin[1,2], Maxime Lousth[3]

[1] *Luleå University of Technology, Luleå, 97187, Sweden*
pierre.dersin@ltu.se
[2] *Eumetry, Louveciennes, France*
pierre.dersin@gmail.com

[3] *SNCF, Paris, France*
maxime.lousth@sncf.fr

Abstract

A substantial part of the French railway network is equipped with ground-based beacons that control train speed. This "KVB" system (French initials for beacon-based speed control) plays an important role in ensuring train operation safety. When a beacon is failed, maximum speed compliance over the section where failure occurred is potentially at risk. Therefore, detecting failed beacons and replacing them in a timely manner but not too often is fundamental to guarantee safe operation of the network while keeping maintenance costs under control. This is accomplished by means of regular commercial trains which detect failed beacons (albeit imperfectly). Upon a number p of successive detections of the same beacon as failed by different trains over a certain period (typically one day), this beacon is signaled as failed to the maintenance control center and is replaced as soon as possible. The question therefore arises of optimizing the signaling rule, i.e., determining the best number p of apparent detections that should trigger a maintenance intervention. Three main contributions are reported in this paper: 1) Determination of steady-state operational availability as a function of the failure rate, the headway, the mean time to restore and the fault detection probability; 2) A method for the optimal choice of the number p in the signaling rule; 3) An algorithm for diagnosing whether a beacon or a train is defective, thereby reducing detection time and false positives. Several sensitivity analyses are also conducted, both of the availability and the total cost, with respect to the various relevant parameters. Generalization to other train control systems, such as the European-wide ETCS (European Train Control System) should be straightforward. To the authors' knowledge, it is the first time that such algorithms for decision optimization under uncertainty are applied in the context of train control system maintenance.

Key words: availability, maintenance costs, failure detection, optimization, decision support.

1. Introduction

The KVB is an essential system to ensure train safety in the French rail network. Its two functions are (i) to prevent the driver from passing signals at danger; (ii) to prevent the driver from overriding the speed limits ((Schön, Benso, Larraufie, Moëns, & Poré, 2019)).The trackside part of the KVB system consists of beacons distributed along the tracks, which are read by the train antenna. As a result of KVB system installation from 1990 to 2006, the safety level has increased significantly. There are more than 100 000 beacons installed throughout the country. A failure of a beacon must be detected as soon as possible, and replacement performed. From the PHM point of view, we are focusing in this paper on detection and diagnostics. Condition-based maintenance is not relevant since statistical records have shown that the beacons have a constant failure rate. Detection of failed beacons is performed by trains, but it is subject to errors: false positives and false negatives. As the cost of sending a maintenance agent to the field is fairly high, care must be taken to avoid doing so upon spurious failure detection. Hence the "common sense" maintenance rule currently in place: a maintenance action on a beacon shall take place only if the beacon has been detected as faulty at least twice during a revenue service day. However, this rule may not be optimal, for instance in view of the safety and availability consequences of a failure; and also, given the fact that, on low-traffic lines, it will take more time before a beacon fault is detected twice. There is therefore a trade-off between maintenance cost and availability. In addition, contractual agreements impose a minimum availability level on the beacons. This situation is a classical problem in maintenance decision support, but the difficulty here is compounded by the existence of intermittent failures, which make fault detection more complex, and by the fact that one is dealing with a distributed system involving a network, mobile assets (trains) and fixed assets (the beacons) interacting with each other, and that the operating conditions—summarized by traffic density, or equivalently, train headway (the time between two successive trains)—are variable from one line to another. We study the general policy, defined as "ordering a maintenance intervention when a beacon has been detected defective p times" (the current rule corresponds to $p = 2$), and attempt to optimize p. First, a model is presented, in Section 2, for beacon availability as a function of the signaling rule. Subsequently, in Section 3, the maintenance cost model is introduced and its optimization is dealt with by means of stochastic dynamic programming (Bellman), with emphasis on steady-state policies. The current policy ($p = 2$) is compared with the one where a maintenance action would be triggered by the first observation ($p = 1$), and a generalization to any value of p is hinted at. Then, in Section 4, we present a probabilistic diagnosis model aiming at dealing with beacon intermittent failures. The technique used to that end is that of "belief update". After an explanation of the interaction between the three models in Section 5, the conclusion is presented in Section 6.

2. STEADY-STATE OPERATIONAL AVAILABILITY

We shall now express the beacons' steady-state availability as a function of the operational headway (i), the number of detections (p) required to trigger a maintenance intervention, and the Mean Corrective Maintenance Time (MCMT), the mean time needed to perform that intervention. First, in a simplified model, the fact that the successive detections must take place over a finite duration is not taken into account. Subsequently, that constraint is introduced (detection over a horizon of n revenue service days). Letting n go to infinity leads back to the expression of the simplified model.

2.1. Simplified Model

Availability depends on the following quantities:

- Failure rate (assumed constant according to field experience), $\mathrm{MTTF} = 1/\lambda$;
- Mean fault detection time (MFDT);
- Mean corrective maintenance time after detection (MCMT).

As illustrated by Figure 1, after the occurrence of a beacon failure, several trains pass by (following each other at an interval of duration i), which may or may not detect the fault. Let us denote by k_p the p-th train which detects it. At time $k_p i$, the maintenance intervention is started, and lasts MCMT on average. On average, the first train that arrives after failure occurrence does so after a time $i/2$. The average down time caused by that failure is therefore equal to

$$\mathrm{MDT} = \mathrm{MFDT} + \mathrm{MCMT} \tag{1}$$

where MFDT, "mean fault detection time", is the average time elapsed before fault detection, and MCMT includes both logistics times and replacement. There follows:

$$\mathrm{MFDT} = i\,E(k_p) - \frac{i}{2} \tag{2}$$

The term $i/2$ is deducted because the first train after the failure arrives on average after a time $i/2$. Thus, for $p = 1$, if $k_1 = 1$, then $\mathrm{MFDT} = i/2$. Steady-state (asymptotic) availability D is then given by:

$$D = \frac{\mathrm{MTTF}}{\mathrm{MTTF} + \mathrm{MDT}} \tag{3}$$

with MDT given by (1). One must therefore express $E(k_p)$. Note that the detections by the successive trains are independent events. The detection probability is the same for all the trains: TP. The sequence of events (detection or not) is therefore a Bernoulli process. Notice that k_p is, within a Bernoulli process, the number of the p-th trial that results in a success. Its probability distribution is a Pascal distribution (special case of negative binomial (Ross, 2010)). In our situation, "success" means detection, therefore the success probability at each trial is TP. Consequently:

$$P[k_p = k] = \begin{cases} \binom{k-1}{p-1} \cdot \mathrm{TP}^p \cdot (1-\mathrm{TP})^{k-p}, & \text{if } p \leq k \\ 0, & \text{if } p > k \end{cases} \tag{4}$$

And thus

$$E(k_p) = \frac{p}{\mathrm{TP}} \tag{5}$$

As a constant failure rate λ is assumed (consistently with field experience), there follows an exponential distribution for times between failures,

$$\mathrm{MTTF} = \frac{1}{\lambda} \tag{6}$$

In conclusion, the following holds:

$$\begin{aligned} D &= \frac{1}{\lambda\left(\frac{1}{\lambda} + \frac{ip}{\mathrm{TP}} - \frac{i}{2} + \mathrm{MCMT}\right)} \\ &= \frac{1}{1 + \lambda\left(\mathrm{MCMT} + i\left(\frac{p}{\mathrm{TP}} - \frac{1}{2}\right)\right)} \end{aligned} \tag{7}$$

(Note: sometimes, the Pascal distribution is defined, equivalently, as the number of unsuccessful trials before achieving p successes). The probability of achieving at least two detec-

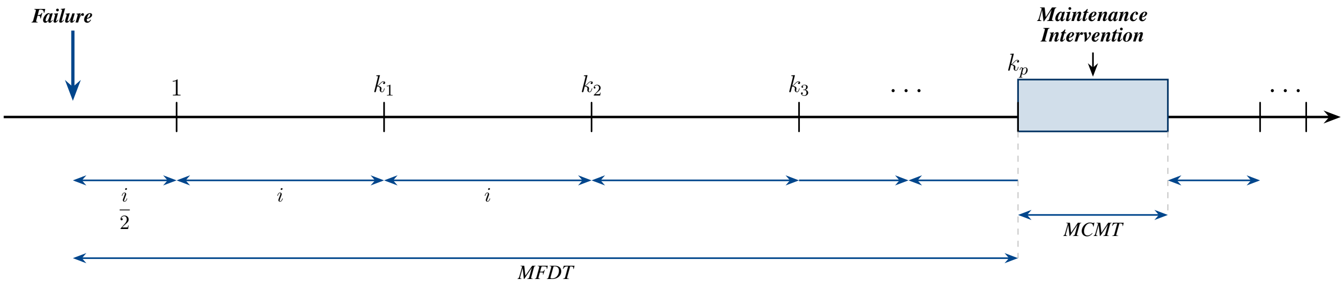

Figure 1. Availability timeline: fault detection and maintenance intervention sequence.

tions over one day (with ten train missions, i.e., $M = 10$) is illustrated in Figure 2 as a function of the true positive probability TP.

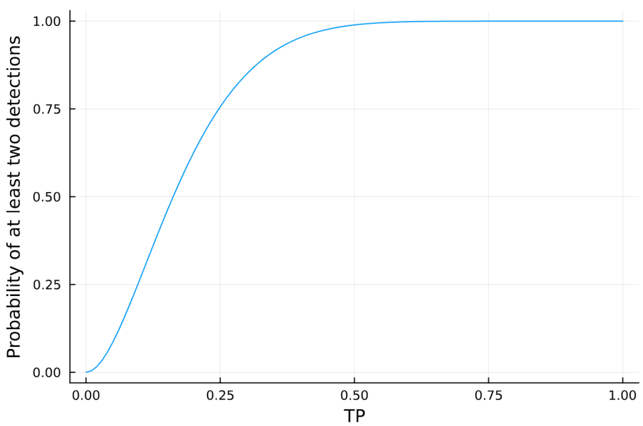

Figure 2. Probability of at least two detections over one day

2.2. Full Model: Availability as a Function of Signaling Rule

In the previous subsection, it was implicitly assumed that as soon as the beacon was detected to be faulty for the p-th time (by definition, at instant $k_p\, i$), the maintenance process (replacement) would start. This led to overestimating availability since the signaling rule consists of transmitting the "beacon faulty" notification if the fault is detected for the p-th time over a period of given duration. That duration is equal to n days of operation of duration T, which corresponds to nM missions. (The reasoning of the above subsection consists of taking n to be infinite.) The notification is sent at the end of the detection period. According to the current rule, $p = 2$ and $n = 1$. Let us study the sensitivity of availability to the choice of p and n. The situation is summarized in Figure 3. Availability is an increasing function of n since having a longer period nM increases the probability of detection and therefore decreases MFDT. On the other hand, availability decreases if p increases because on average the beacon stays down longer.

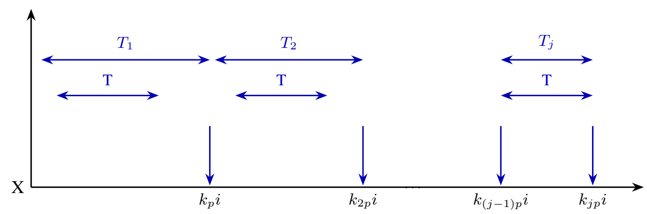

Figure 3. Successive faulty beacon detections

In Figure 3, $k_p i$ denotes, as previously, the first instant when p detections have taken place; then, k_{2p}, k_{3p}, and so on, correspond respectively to $2p$, $3p$, ... detections. The random variables T_1, T_2, and so on, are the lengths of the corresponding periods: for instance T_j is the time elapsed between the $(j-1)p$-th detection and the jp-th detection. T is the deterministic duration that is part of the signaling rule: the notification is sent as soon as p successive detections have occurred over a period of duration $T = nMi$, i.e., as soon as there is an integer j such that $T_j < T$. In the numerical example, T has been taken equal to 16 h. One is in fact interested in j^* defined by

$$j^* = \min\{j : T_j < T\}. \tag{8}$$

We wish to express $E(k_{j^*p})$. Indeed:

$$\mathrm{MFDT} = i\, E(k_{j^*p}) \tag{9}$$

The random variables $T_1, T_2, \ldots, T_j$ are independent, identically distributed (iid) with a Pascal (negative binomial) (Ross, 2010) distribution with parameters p and TP. Therefore it is easy to find an expression for $P[T_j < T]$, which will be denoted ϖ.

$$P(T_j < T) = \varpi \tag{10}$$

In the sequence of random events $(T_1 < T)$, $(T_2 < T)$, ..., $(T_k < T)$, the rank of the "first success", i.e., the smallest j such that $T_j < T$, follows a geometric distribution:

$$P[j^* = j] = (1-\varpi)^{j-1}\,\varpi \tag{11}$$

therefore

$$E(j^*) = \frac{1}{\varpi} \tag{12}$$

Let us denote by $\langle j^* \rangle$ the expectation $E(j^*)$. In conclusion, the MFDT can be expressed as follows:

$$\mathrm{MFDT} = i\, E\big(k_{\langle j^* \rangle\, p}\big) - \frac{i}{2} \tag{13}$$

To obtain MFDT, it suffices, taking (12) into account, to determine the detection probability ϖ. But

$$\varpi = P[T_j < T] = P[T_1 < T] = P[k_p \leq nM] \tag{14}$$

is the probability of at least p detections over n successive periods of operation. That expression is immediately derived from (4):

$$\varpi = \sum_{k=p}^{nM} \binom{k-1}{p-1} \mathrm{TP}^p (1-\mathrm{TP})^{k-p} \tag{15}$$

There follows that:

$$\mathrm{MFDT} = i\, E\big(k_{\langle j^* \rangle\, p}\big) - \frac{i}{2} = i\, \frac{p}{\varpi \cdot \mathrm{TP}} - \frac{i}{2} \tag{16}$$

From (15) and (16), it is verified that:

(i) if p increases, MFDT increases (ϖ decreases since k_p is increasing with p);

(ii) if n increases, ϖ increases, so MFDT decreases;

(iii) if TP increases, MFDT decreases.

Numerical examples.

- For $p = 1$, $n = 1$: $\varpi = 1 - (1 - \mathrm{TP})^M$.
- For $p = 2$, $n = 1$ (current strategy), upon expanding:

$$\varpi = 1 - \left[(1-\mathrm{TP})^{nM} + nM\, \mathrm{TP}\, (1-\mathrm{TP})^{nM-1}\right] \tag{17}$$

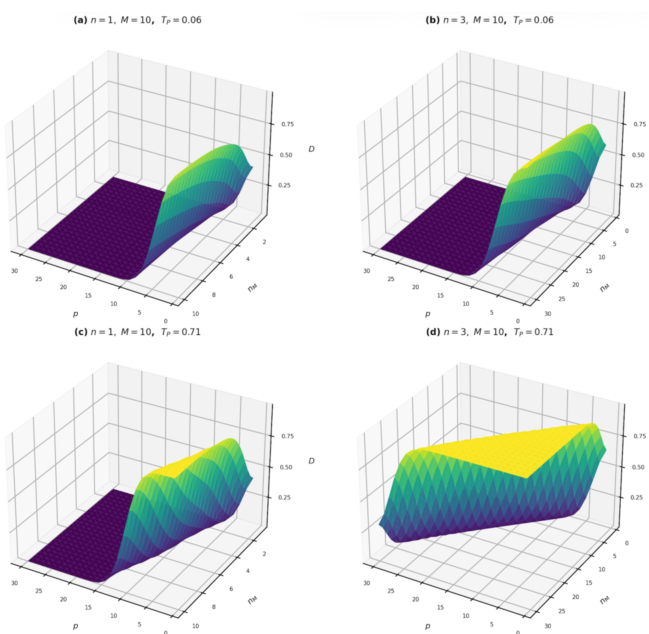

Figure 4. Sensitivity analysis of availability to n, M and p

Let us look at some limit cases.

- **Limit when** $n \to \infty$**:** It is easily verified that, when n goes to infinity, $\lim_{n\to\infty} \varpi = 1$, therefore the right-hand side of (12) goes to $i\left(\frac{p}{\mathrm{TP}} - \frac{1}{2}\right)$. It follows from the definition that $\lim_{n\to\infty} P[k_p \leq nM] = 1$. The expression for MFDT in the simplified model of Subsection 2.1 is found again, as that model in fact corresponds to a period of infinite duration.
- **Limit when** $i \to 0$**:** When the headway i goes to 0 (which would correspond to an infinite traffic density), $\lim_{i\to 0} \mathrm{MFDT} = 0$.
- **Limit when** $\mathrm{TP} \to 1$**:** When TP goes to 1 (which means detection is certain):

  $$\lim_{\mathrm{TP}\to 1} \mathrm{MFDT} = ip - \frac{i}{2} \tag{18}$$

  Indeed, in that case, $k_p = p$ with certainty (the p-th train performs the p-th detection), and $\varpi = 1$, provided $p < nM$. Figure 4 illustrates the sensitivity of availability to p, the number of detections required, and n, the number of days during which p detections are required to send the notification.

Availability is an increasing function of n because a larger n means more opportunities to detect the failed beacon (higher ϖ and thus lower MFDT). On the other hand, it is a decreasing function of p (it takes longer to perform p detections, hence a larger MFDT).

3. COST OPTIMIZATION

We shall now model the maintenance costs and determine the maintenance policy that minimizes the total cost, by a suitable selection of the fault notification policy. The method retained is Bellman's stochastic dynamic programming (Bertsekas, 2012). From observations, actions are decided in such a way as to minimize the total maintenance costs. System state is estimated from observations and depends on past actions. System evolution from one state to another is modeled by a continuous-time Markov chain (Fig. 5). The decision process is then Markovian. The system (i.e., the beacon) evolves between an operational state 0 and a fault state 1. The transition from 0 to 1 is a failure, characterized by a failure rate λ, and the transition from 1 to 0, a replacement characterized by a restoration rate μ. The system is observed at discrete instants, indexed by k, which correspond to the passages of the successive trains on the beacon. The k-th observation, Z_k (corresponding to the passage of train k), depends on the actual state of the beacon, x_k, but this state can be observed with an error (probability FP of false positive, or FN of false negative). The decision U_k is based on observations (at stage k, but possibly also at previous stages). There are two possible decisions: to act, i.e., to replace the beacon: $u_k = 1$; or to do nothing: $u_k = 0$. The costs, or rewards, correspond to the decisions.

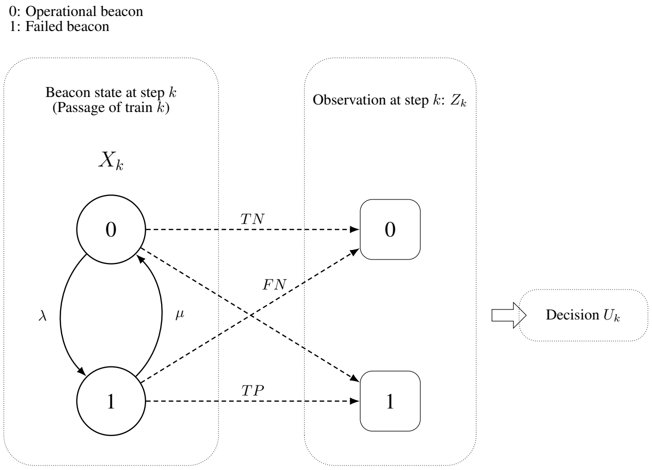

Figure 5. Bellman model: beacon state X_k, observation Z_k, and decision U_k at each train passage k.

Specifically, we distinguish different types of costs, which are included or not depending on the state and the decision:

- **An intervention cost**, depending on the observation z_k (which prompts, or not, an intervention) and on the actual state of the beacon, x_k, which determines whether it must be replaced or not.Two cases may occur: either a false alarm (the cost is then denoted C_1'); or a true positive: the cost is then denoted C_3. C_3 is the maintenance and operation cost corresponding to the resolution of a true fault that is detected.

- **An operational cost, or a reduction in quality of service (QS), resulting from a non-intervention on a beacon that is really faulty:** C_2. This is the cost of beacon downtime. For instance, quality of service may be measured by some key performance indicator such as average train punctuality. This in an indirect cost that may result from a reduced speed when driving in manual mode or if the failed beacon keeps sending restrictive speed information.

Let us denote $g_k(x_k, z_k, u_k)$ the sum "cost minus reward" during the period $(k, k+1)$, which depends on the policy selected. The latter results in decision u_k (to intervene or not), which depends on the observations (z_k). The cost function g_k can be described by a truth table. In the case of Policy P1 (to intervene if the alarm is positive, i.e., if $z_k = 1$), it is summarized in Table 1.

3.1. Boolean Expression for Cost Function

Equivalently, the cost function g_k can be expressed by a Boolean equation:

Table 1. Cost Function in Policy P1

Variables	Configurations			
x_k (actual state)	0	0	1	1
z_k (Observation)	0	1	0	1
u_k (Action)	0	1	0	1
g_k (Cost P1)	0	C_1'	C_2	C_3

$$g_k(\text{P1}; x_k, z_k) = (1 - x_k) z_k C_1' + x_k(1 - z_k) C_2 + x_k z_k C_3 \tag{19}$$

Note: to take the rewards (quality of service) into account, one may replace C_i with $(C_i - \Delta QS_i)$. Since the cost function is homogeneous in the various unit costs, optimal decisions depend only on the ratios to one of the unit costs, for instance C_i/C_3, not the absolute values.

3.2. Influence of the Decision on State Evolution

Action u_k influences system state at the next stage $(k + 1)$. Denote, as before, by i the time between two successive inspections, i.e., the passage of two successive trains.

In case of a maintenance action ($u_k = 1$), the beacon will be operational at the next stage $(k+1)$, except if it fails between k and $k + 1$. Those state transitions can be expressed by a Markov chain.

3.3. Optimality Equation (Bellman)—Policy P1

We now use stochastic dynamic programming (Bertsekas, 2012). Denoting by V_k the value function, i.e., the optimum of the cost function at stage k, over a horizon of N stages, Bellman's recursive equation is written as follows:

$$V_k = \min\big[E\, g_k(\text{P1}; x_k, z_k) + V_{k+1}\big] \tag{20}$$

Minimization is performed on u_k, i.e., the action performed at stage k; the expectation is calculated over the state x_k and the observation z_k. If the horizon consists of N stages, the algorithm must be initialized with $k = N - 1$. In the present application, we have

$$\begin{aligned} E\, g_k(\text{P1}; x_k, z_k) = P(x_k = 0)\, E[g_k(\text{P1}; 0, z_k) \mid x_k = 0] \\ + P(x_k = 1)\, E[g_k(\text{P1}; 1, z_k) \mid x_k = 1] \end{aligned} \tag{21}$$

$$E(g_k(\text{P1}; 0, z_k) \mid x_k = 0) = E(z_k \mid x_k = 0) \cdot C_1' \tag{22}$$

Referring to Fig. 5,

$$E(z_k \mid x_k = 0) = \text{TN} \cdot 0 + \text{FP} \cdot 1 = \text{FP} \tag{23}$$

Therefore, from (22),

$$E(g_k(\text{P1}; 0, z_k) \mid x_k = 0) = \text{FP} \cdot C_1' \tag{24}$$

Likewise,

$$E[g_k(\mathrm{P1}; 1, z_k) \mid x_k = 1] = [1 - E(z_k \mid x_k = 1)] \cdot C_2 + E(z_k \mid x_k = 1) \cdot C_3 \quad (25)$$

From Fig. 5, there follows: $E(z_k \mid x_k = 1) = \mathrm{TP}$, whence

$$E[g_k(\mathrm{P1}; 1, z_k) \mid x_k = 1] = (1 - \mathrm{TP})C_2 + \mathrm{TP}\, C_3 \quad (26)$$

From (24) and (26),

$$E\, g_k(\mathrm{P1}; x_k, z_k) = P(x_k = 0) \cdot \mathrm{FP} \cdot C_1' + P(x_k = 1)\, [(1 - \mathrm{TP})C_2 + \mathrm{TP}\, C_3] \quad (27)$$

One must therefore introduce that expression into the Bellman equation (20) and solve the latter recursively by back-propagation.

3.3.1. Steady-State Regime

If the length N of the horizon goes to infinity, $\mathrm{Prob}[x_k = 0]$ and $\mathrm{Prob}[x_k = 1]$ converge respectively to the system's asymptotic availability and unavailability. The sequence of states x_k follows a discrete-time Markov chain (indexed by k). The probabilities for the system to be respectively in state 0 (normal operation) or 1 (faulty) at stages k and $k+1$ are derived from the transition equations of the Markov chain (transition matrix P). That transition matrix depends on the train headway i. In vector form, denoting by p_k' the vector (transposed) with components $p_k(0)$ and $p_k(1)$,

$$p_k(0) = \mathrm{Prob}[x_k = 0] \quad (28)$$

$$p_k(1) = \mathrm{Prob}[x_k = 1] \quad (29)$$

In steady state ($k \to \infty$), p_k converges to a limit vector p, solution of

$$p' = p\, P(i) \quad (30)$$

whose components are the steady-state availability and unavailability respectively. Therefore, asymptotically,

$$C_\infty = \lim_{k \to \infty} E[g_k(\mathrm{P1}; x_k, z_k)] = p(0)\, \mathrm{FP} \cdot C_1' + p(1)\, [(1 - \mathrm{TP})C_2 + \mathrm{TP}\, C_3] \quad (31)$$

This expression represents the average cost per unit of time. From now on we denote by D_i the asymptotic availability corresponding to a train headway i,

$$D_i = p(0) = 1 - p(1) \quad (32)$$

In the limiting case of a "perfectly testable system" ($\mathrm{TP} = 1$ and $\mathrm{FP} = 0$), this cost reduces to the average cost per stage of a replacement (a fault is always detected):

$$C_\infty = (1 - D_i)\, C_3 \quad (33)$$

In the general case, $\mathrm{FP} > 0$ and $\mathrm{TP} < 1$, so that two additional terms occur in Eq. (31). The first term, weighted by availability, depends only on frequency and cost of unwarranted interventions, hence on false positives; it would be the only term if availability were 100%. The second term, weighted by unavailability, depends only on the cost resulting from an action or the absence of it in case of a fault; thus it is a function of true positives. Availability D_i is a function of TP (see Section 2), but in practice the sensitivity to TP is very weak; accordingly it has been taken as constant.

3.3.2. Interpretation

Policy P1 consists of sending a maintainer on the field at the first detection of a faulty beacon. Intuitively, it is clear that this policy is not satisfactory if the false positive rate is high (there will be too many useless replacements). On the other hand, if the intervention cost C_3 is less than the cost of operation and maintenance C_2 resulting from non-intervention upon faulty beacon detection, it is preferable to intervene more often in order to avoid field failures; therefore Policy P1 can be justified, provided, however, that the false negative rate (lack of alarm in case of a fault state) is high enough (the false alarm rate occurs in (31) through $(1 - \mathrm{TP})$)[1].

3.4. Optimality Equation (Bellman)—Policy P2

Recall the definition of Policy P2: intervention takes place only if the beacon is detected as faulty at least twice over an operation period. An approximate expression of the cost function is obtained with a simplified definition of P2: intervention if the fault is detected by two successive trains: $(k - 1)$ and k. The truth table of Policy P1 is then replaced by Table 2 to reflect the fact that the decision u_k depends on the last two observations.

Table 2. Truth Table in Policy P2

z_{k-1}	0	0	0	0	1	1	1	1
z_k	0	0	1	1	0	0	1	1
x_k	0	1	0	1	0	1	0	1
u_k	0	0	0	0	0	0	1	1
$g_k(\mathrm{P2})$	0	C_2	0	C_2	0	C_2	C_1'	C_3

Accordingly, the Boolean expression for the cost function,

[1] In actuality, the notations FP, TP, FN, TN used in this paper refer to rates: true positive rate, false positive rate, etc. Rigorously they ought to be denoted FPR, TPR, FNR, TNR. If FP (respectively TP) denotes the number of false positives (resp. true positives) over time, i.e. integers, the relations between those integers and the rates are as follows: $\mathrm{FPR} = \dfrac{\mathrm{FP}}{\mathrm{FP} + \mathrm{TN}}$ (proportion of actual negatives incorrectly flagged positive), $\mathrm{TPR} = \dfrac{\mathrm{TP}}{\mathrm{TP} + \mathrm{FN}}$ (proportion of actual positives correctly detected), $\mathrm{FNR} = \dfrac{\mathrm{FN}}{\mathrm{FN} + \mathrm{TP}}$ (proportion of actual positives missed). Thus $\mathrm{FNR} = 1 - \mathrm{TPR}$. To simplify notations, we have omitted the R at the end.

and the expression for the expected cost, are modified as well. The steady-state expected cost is obtained as follows:

$$C_\infty(\mathrm{FP},\mathrm{TP}) = C_2(1-D) + (C_3 - C_2)\big[(1-A)(1-D)\,\mathrm{FP}\,\mathrm{TP} + B(1-D)\,\mathrm{TP}^2\big] + C_1'\big[A\,D\,\mathrm{FP}^2 + (1-B)(1-D)\,\mathrm{FP}\,\mathrm{TP}\big] \quad (34)$$

where A is a function of beacon failure rate and train headway, and B is a function of beacon restoration rate and train headway. The proofs are omitted for conciseness; they use the truth table and the Markov transition matrix. Contrary to the P1 case, the cost function is now a quadratic function of the parameters FP and TP. If a plausible representation can be found for the receiver operating characteristic (ROC) (Bewick V. Cheek, 2004), i.e., $\mathrm{TP} = f(\mathrm{FP})$, then only FP appears in the cost function. This is what has been done in order to compare the two policies.

3.5. Comparing Policies P1 and P2

As shown in Section 2, the choice of p influences availability. However, in the usually encountered range of values, that sensitivity is not very strong. On the other hand, a higher MCMT, essentially the reaction time once a defective beacon has been detected, can lower availability considerably. This is why we consider here the influence of the signaling rule (P1 or P2) in two different availability contexts: standard availability ($D = 0.976$) and low availability ($D = 0.797$). Those numbers are given for illustrative purposes and are not meant to reflect actual field data, for confidentiality reasons. Figures 6 and 7 illustrate four numerical examples, corresponding to those two different values for D, and two values for C_2, with the costs expressed in terms of FP in both policies. They both show the asymptotic cost C_∞ as a function of the false positive rate FP. Equations (31) and (34) are applied. In Fig. 6, the cost C_2 of a failure is equal to 2.1, just slightly above the cost C_3 of an intervention (equal to 2). The low availability case ($D = 0.797$) is illustrated by the top curve for Policy P1 and the curve immediately below it for Policy P2. The standard availability case ($D = 0.976$) is illustrated by the two bottom curves (with the one corresponding to P1 above the one for P2). In Figure 7, the cost of a failure C_2 is equal to 100, much higher than the intervention cost C_3. Again, the two top curves correspond to the low availability case and the two bottom curves to the standard availability, but now Policy P2 entails a higher cost than P1, in both cases.

Discussion Let us compare the two policies from the minimum cost perspective. In the case $C_2 > C_3$ (the most frequent), expressions (31) and (34) show that, for large enough C_2, Policy P2 is more costly than Policy P1, while the opposite holds if the difference between C_2 and C_3 is small. This observation is confirmed by Fig. 6: for $C_2 = 2.1$ (minute difference with C_3), Policy P2 is less costly than P1; whereas, for $C_2 = 100$ (very large difference), the opposite situation holds. Those results are readily understandable: when the cost of a field failure is high as compared to that of an intervention, it is important to avoid that failure, and that is less easy if one waits for two alarms (detection by two trains, P2) than if one reacts after the first alarm (P1). In contrast, if the difference is small, the effect of fewer interventions (P2) dominates. The effect is more important for a lower availability .The figures show that, except for high values of FP, a higher availability corresponds to a lower maintenance cost: indeed higher availability usually means fewer interventions. For higher values of FP, the cost of useless interventions may predominate. Also note that, for the low availability case ($D = 0.797$), when the (C2-C3) difference is large, the costs are decreasing functions of FP. This is explained by the fact that TP increases with FP and, the larger TP, the more likely it is that expensive field failures will be avoided.

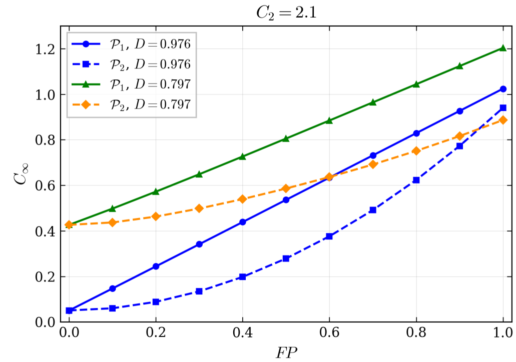

Figure 6. Cost comparison between the two policies ($C_2 = 2.1$)

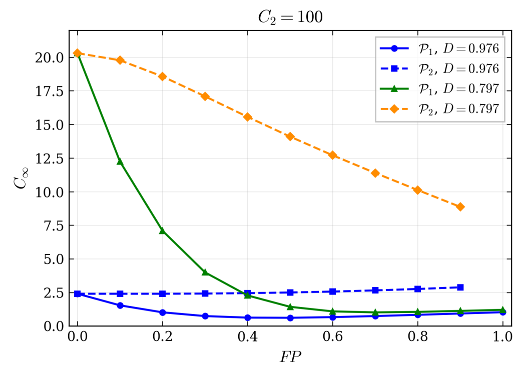

Figure 7. Cost comparison between the two policies ($C_2 = 100$)

3.6. General Case: Policy P_p

In the general case of a policy P_p, defined by the rule "notification sent after p detections of the beacon as faulty", it is still possible to obtain a general expression for the asymptotic cost as a function of TP, FP and p. It is obtained by calculating the p-th power of a 2×2 matrix function of FP and TP. The function is polynomial of degree p in TP and FP.

4. PROBABILISTIC DIAGNOSIS

When a communication error between train antenna and beacon occurs, an alarm is raised. However, a single alarm is inherently ambiguous: it could be due to a defective beacon, a defective train antenna, or environmental noise. Deterministic diagnosis is often impossible given the sparse connectivity of the observation network.

We propose a probabilistic approach using Graphical Models (Koller & Friedman, 2009). By treating beacons and trains as random variables and observations as probabilistic factors, we can infer the most likely state of the network. The core idea, derived from Pearl's seminal work(Pearl, 1988), is to propagate *messages* between variables through observation factors to compute updated beliefs.

4.1. System Modeling

We formalize the diagnosis problem within the framework of probabilistic graphical models. The key insight is that the structure of the train-beacon interaction naturally induces a bipartite factor graph (Kschischang, Frey, & Loeliger, 2001).

4.1.1. State Space Definition

Definition 1 (Component States)**.** *We define the following discrete state spaces:*

- ***Beacon states:*** $\mathcal{S}_\mathcal{B} = \{1, 2, 3\}$ *corresponding to* $\{$*Functional, Intermittent Failure, Hard Failure*$\}$
- ***Train states:*** $\mathcal{S}_\mathcal{T} = \{1, 2\}$ *corresponding to* $\{$*Functional, Faulty*$\}$
- ***Observations:*** $\mathcal{O} = \{0, 1\}$ *corresponding to* $\{$*No Alarm, Alarm*$\}$

The choice of three beacon states reflects operational reality: beacons may exhibit apparent self-recovering failures, which we classify as 'intermittent failures'. These are often due to environmental conditions including for instance, vibration (cause by train speed) or temperature. This behavior is highly stochastic and unpredictable. Ultimately, this intermediate state represents the key insight to be gained from our study.

4.1.2. Factor Graph Representation

Definition 2 (Factor Graph (Kschischang et al., 2001))**.** *A factor graph is a bipartite graph* $G = (V, F, E)$ *where* V *is a set of variable nodes,* F *is a set of factor nodes, and edges* E *connect variables to factors.*

In our setting, the joint posterior distribution factorizes as:

$$P(\boldsymbol{\beta}, \boldsymbol{\tau} \mid \mathbf{O}) \propto \prod_i \phi_\mathcal{B}(\beta_i) \prod_j \phi_\mathcal{T}(\tau_j) \times \prod_{(i,j) \in \mathcal{E}} \psi_{ij}(\beta_i, \tau_j) \quad (35)$$

where $\mathcal{E}$ is the set of observed train-beacon encounters, $\phi_\mathcal{B}$ and $\phi_\mathcal{T}$ are prior potentials, and ψ_{ij} encodes the likelihood $P(O_{ij} \mid \beta_i, \tau_j)$.

Figure 8 presents a simplified view of this bipartite structure for a minimal example with 3 beacons and 2 trains, while Figure 9 shows a realistic KVB topology and its corresponding factor graph.

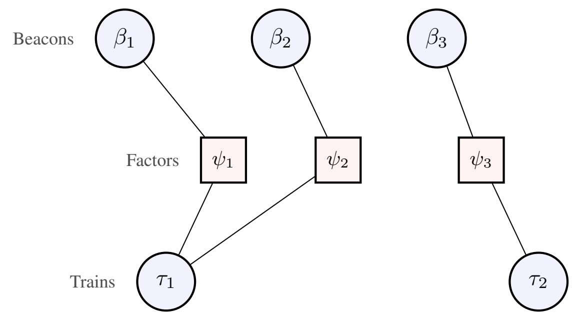

Figure 8. Simplified factor graph illustrating the bipartite structure. Variable nodes (circles) represent beacons and trains (hidden state); factor nodes (squares) encode observation potentials ψ_{ij}.

4.1.3. Observation Model

The observation model $P(O_{ij} = 1 \mid \beta_i, \tau_j)$ quantifies the probability of an alarm given component states. We parameterize this as a conditional probability table, calibrated from operational data and expert knowledge; a representative example of these emission probabilities is shown in Table 3. To complete the observation model, we also require the transition matrices for the beacon and the train—specifically, a 3×3 and a 2×2 matrix, respectively. Parameter estimation from the observation data was carried out within the Factorial HMM framework, as described in (Ghahramani & Jordan, 1997).

Table 3. Conditional probability table $P(\text{Alarm} \mid \beta_i, \tau_j)$.

	$\tau_j =$ Funct.	$\tau_j =$ Faulty
$\beta_i =$ Functional	0.01	0.95
$\beta_i =$ Intermittent	0.50	0.98
$\beta_i =$ Hard Failure	0.90	0.99

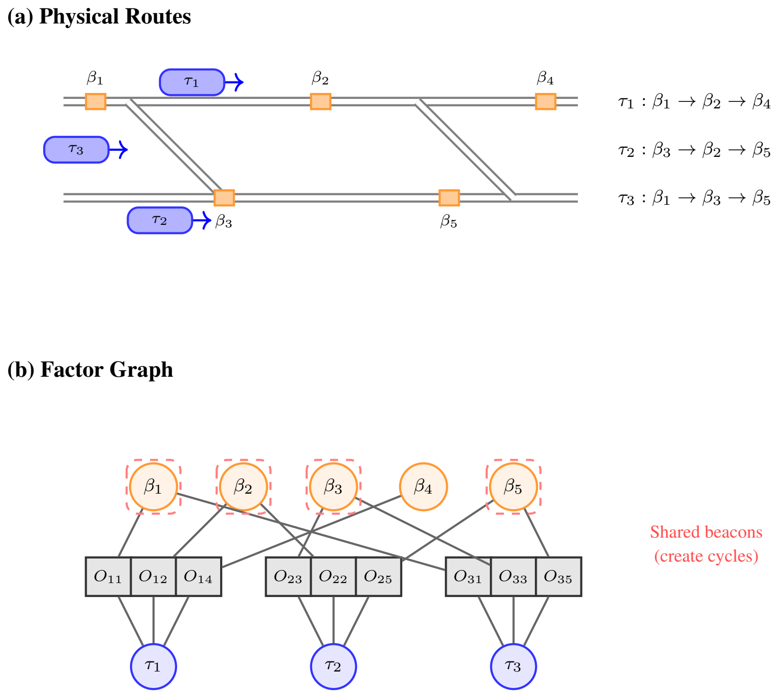

Figure 9. System topology. (a) Physical railway layout with three train routes. (b) Corresponding factor graph with observation factors O_{ij}. Beacons shared across routes (highlighted with dashed borders) create cycles that preclude exact inference.

4.2. Inference Algorithm

Exact inference in graphical models with cycles is computationally intractable (Cooper, 1990). We employ **Loopy Belief Propagation** (LBP), an iterative approximation that applies the sum-product algorithm (Kschischang et al., 2001) despite the presence of cycles.

4.2.1. Sum-Product Algorithm on Factor Graphs

The sum-product algorithm computes marginals by passing messages between variable and factor nodes. On tree-structured graphs, a single upward-downward pass yields exact marginals (Pearl, 1988).

Definition 3 (Variable-to-Factor Message)**.** *The message from variable x to factor f aggregates all incoming messages except from f:*

$$\mu_{x\to f}(x) = \prod_{g\in\mathcal{N}(x)\setminus\{f\}} \mu_{g\to x}(x) \tag{36}$$

Definition 4 (Factor-to-Variable Message)**.** *The message from factor f to variable x marginalizes over all other variables connected to f:*

$$\mu_{f\to x}(x) = \sum_{\mathbf{x}_f\setminus x} f(\mathbf{x}_f) \times \prod_{y\in\mathcal{N}(f)\setminus\{x\}} \mu_{y\to f}(y) \tag{37}$$

Definition 5 (Belief Computation)**.** *The belief (marginal estimate) at variable x is the normalized product of all incoming messages:*

$$b(x) \propto \prod_{f\in\mathcal{N}(x)} \mu_{f\to x}(x) \tag{38}$$

The core concept of Belief Propagation is the exchange of "local opinions" between nodes. A node sends a message representing its belief about a neighbor's state, based on all information received from its other neighbors. Intuitively: "I have heard from ψ_1 and ψ_2 that you are likely in state X. Therefore, I am telling ψ_{target} that you are in state X." Figure 10 illustrates these update rules graphically.

In our KVB setting, these rules specialize as follows. The message from observation factor O_{ij} to beacon β_i marginalizes over the train state:

$$\mu_{O_{ij}\to\beta_i}(b_i) = \sum_{t_j} P(O_{ij} \mid b_i, t_j) \cdot \mu_{\tau_j\to O_{ij}}(t_j) \tag{39}$$

The message from beacon β_i to observation O_{ij} is the product of all other incoming messages:

$$\mu_{\beta_i\to O_{ij}}(b_i) \propto \prod_{k\in\mathcal{N}(\beta_i)\setminus\{j\}} \mu_{O_{ik}\to\beta_i}(b_i) \tag{40}$$

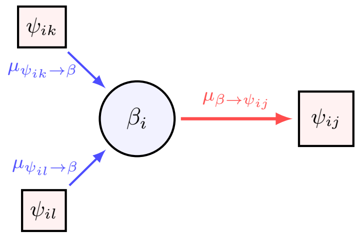

(a) Variable → Factor
$\mu_{\beta\to\psi} = \prod_{k\neq j} \mu_{\psi_{ik}\to\beta}$

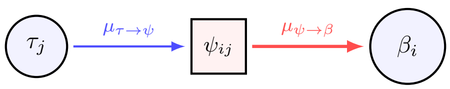

(b) Factor → Variable
$\mu_{\psi\to\beta}(b) = \sum_t \psi(b,t)\cdot\mu_{\tau\to\psi}(t)$

Figure 10. Message update rules in the sum-product algorithm.

4.2.2. The Cycle Problem

The presence of cycles fundamentally changes the behavior of belief propagation. In tree-structured graphs, each message is computed exactly once and carries independent information. In cyclic graphs, messages can circulate indefinitely, leading to "double-counting" of evidence that violates the independence assumptions underlying the sum-product derivation and can produce overconfident or oscillating beliefs.

Consider a minimal cyclic configuration with 2 beacons and 2 trains, where each train observes both beacons (Figure 11). An observation O_{11} influences beliefs through the following path: $O_{11} \xrightarrow{1} \beta_1 \xrightarrow{2} O_{21} \xrightarrow{3} \tau_2 \xrightarrow{4} O_{22} \xrightarrow{5} \beta_2 \xrightarrow{6} O_{12} \xrightarrow{7} \tau_1 \xrightarrow{8} O_{11}$. After 8 message updates, evidence from O_{11} returns to influence $\mu_{\tau_1\to O_{11}}$, effectively counting the same observation multiple times.

4.2.3. Convergence Strategies

We evaluated three approaches to obtain stable beliefs in the presence of cycles:

1. **Damping** (Murphy, Weiss, & Jordan, 1999): interpolating between old and new messages via $\mu^{(k+1)} = \alpha \cdot \tilde{\mu}^{(k+1)} + (1-\alpha)\cdot\mu^{(k)}$, where $\alpha \in (0,1]$. While damping reduces oscillations, it does not fundamentally prevent the accumulation of circular evidence.
2. **Convergence threshold**: monitoring the maximum message change $\Delta^{(k)} = \max_{(i,j)} \|\mu_{ij}^{(k)} - \mu_{ij}^{(k-1)}\|_\infty$ and terminating when $\Delta^{(k)} < \epsilon$. Empirically, $\epsilon = 10^{-2}$ provides a robust trade-off between accuracy and stability.

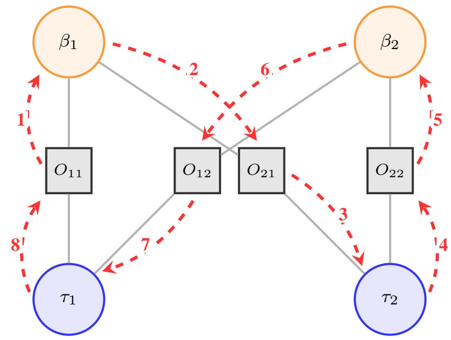

Figure 11. Information circulation in a minimal cycle (2 beacons, 2 trains). After 8 message steps, evidence from O_{11} returns to O_{11}, causing double-counting.

3. **Fixed iteration count**: limiting iterations to $K_{\max}$ based on the graph diameter. For KVB networks, $K_{\max} = 5$ to 23 iterations suffice to propagate information without excessive amplification. Without any stopping criterion, the system self-amplifies messages indefinitely, producing aberrant beliefs.

Remark 1. *The choice of stopping criterion depends on the network topology. For networks with many short intricate loops, resonance behavior appears rapidly; therefore, the damping factor should be large and the iteration count lower.*

4.3. Experimental Results

A complete algorithm was implemented in Julia to simulate the network, modeling train circulation, beacon observations, and stochastic failures. The results demonstrate that the algorithm effectively diagnoses whether a beacon or a train is defective, provided that necessary precautions are taken to ensure its convergence.

Indeed, simulations involving 20 trains and 100 beacons demonstrate that, without convergence control, the model's beliefs exhibit resonance and significant oscillations. These oscillations can be so severe that the entire system state (all 100 beacons and 20 trains) may jump from 'functional' to 'hard failure' based on a single observation. Once this unstable regime is reached, subsequent iterations fail to stabilize. As shown in Figure 12 implementing a convergence threshold $\epsilon = 10^{-2}$ eliminates this behavior, with convergence typically achieved within 8–15 iterations depending on network size.

Remark 2. *Regarding the computational complexity, each factor-to-variable message requires marginalizing over the factor's scope, yielding $\mathcal{O}(K_A^2)$, $\mathcal{O}(K_B^2)$, and $\mathcal{O}(K_AK_B)$ operations for the transition and observation factors respectively. Summing over the T time steps and I iterations,*

$$\mathcal{O}\big(I\cdot T\cdot(K_A^2 + K_B^2 + K_AK_B)\big),$$

linear in T as expected for LBP on a factorial-chain graph: the chain structure prevents any combinatorial blow-up in the sequence length.

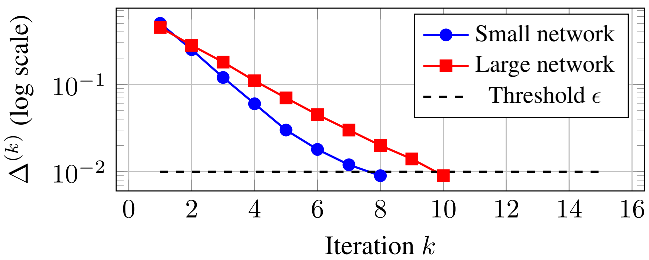

Figure 12. Convergence behavior of LBP. Message variation $\Delta^{(k)}$ decreases monotonically for typical KVB topologies, reaching the threshold $\epsilon = 10^{-2}$ within 8–10 iterations.

5. ARTICULATION OF THE THREE MODELS

As shown in Figure 13, the analysis first organized the problem parameters around availability, then related them to operational costs through the two most consequential diagnostic errors — unnecessary removal and non-detection, the latter carrying the risk of a railway incident. The study identified how maintenance policy can be optimized as a function of system parameters and scale. Safety considerations, however, must always prevail over cost optimization; moreover, the national — indeed international — perimeter of the network makes direct manipulation of maintenance schedules extremely difficult in practice.

A complementary approach therefore consists in exploiting data already reported by the train–beacon pair. These data, required for signaling and speed supervision, were not originally intended for maintenance, and the interest of the belief-propagation approach lies precisely in repurposing them: rendering them intelligible to the maintenance operator and supporting maintenance orders as accurately as possible. The algorithm chiefly targets the state which is the hardest to apprehend — intermittent beacon failures — and which generates the bulk of field disturbances and accounts for the principal share of maintenance under-performance

A natural extension of this work is to incorporate, into the joint analysis of availability and cost, a detection system whose accuracy evolves with the nature of the incoming observations. This conceptually delicate task is left to a subsequent study.

6. CONCLUSION AND SUGGESTIONS FOR FUTURE RESEARCH

In this paper, we have presented innovative probabilistic models to quantify the impact of the maintenance decision rule on steady-state availability and average steady-state maintenance cost of KVB beacons, and a stochastic dynamic programming model to minimize that cost. Thus, a quantitative instrument for maintenance decision support has been provided. We have also successfully explored the use of belief update methods to discriminate between beacon defects and train defects and take intermittent failures into account. Implementation by SNCF on field data is now being considered, but the additional difficulty of data quality and validation must be confronted. Further research could address the following.

- Adding the constraint of achieving a prescribed availability level to the maintenance cost minimization problem.
- Studying maintenance cost optimization in the time-varying case (time-varying rather than steady-state Bellman equation).
- Updating dynamically the values of TP and FP and updating the ROC on the basis of observations.
- Pursuing validation of the belief propagation algorithm on real life data.

Finally, application to ETCS (European Train Control System) can be contemplated.

ACKNOWLEDGMENT

The authors are grateful to Mr. A. Tabosa for fruitful discussions over the KVB system operation, and to Mr. V. Tempé for useful insight on the processing of field data.

NOMENCLATURE

Acronyms

KVB – “contrôle de vitesse par balise” (beacon-based speed control)

ETCS – European Train Control System

MTTF – Mean (operating) Time To Failure

MFDT – Mean Fault Detection Time

MCMT – Mean Corrective Maintenance Time

MDT – Mean Down Time

ROC – Receiver Operating Characteristic

LBP – Loopy Belief Propagation

EM – Expectation-Maximization

Detection rates

TP – true positive rate

FP – false positive rate

TN – true negative rate

FN – false negative rate

Availability and signaling rule (Sec. 2)

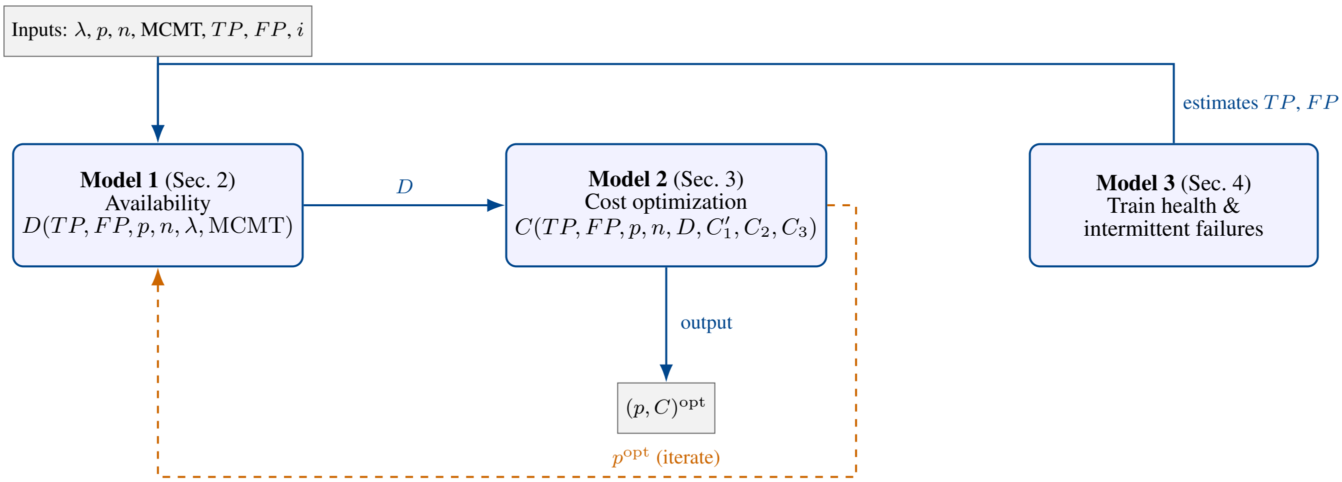

Figure 13. Articulation of the three models. Model 1 feeds the availability D to Model 2, which returns an optimal policy p^{opt} that is re-injected into Model 1. Model 3 provides estimates of TP and FP from observations.
Notes: (i) $n = 1$ is assumed throughout this study; (ii) D is treated as an exogenous parameter in Model 2 (not as a function of p), see Section 3.

i – operational headway

T – duration of daily operation period

M – number of train missions in one operation period

n – number of days over which p detections are required

p – number of "beacon faulty" detections over n successive operation periods before beacon is reported failed

λ – beacon failure rate

μ – beacon restoration rate

D – availability

D_i – availability with headway i

T_j – time between detection $(j-1)$ and detection j

ϖ – probability that $T_j < T$

k_p – order of passage of the train that detects the fault for the p-th time

j^* – smallest j such that $T_j < T$

Cost optimization (Sec. 3)

X_k, x_k – state of the beacon at stage k (passage of train k)

Z_k, z_k – observation at stage k

U_k, u_k – decision at stage k

P_p – policy p

C_3 – unit intervention cost (true failure)

g_k – cost function at stage k

V_k – value function (Bellman) at stage k

C_1' – false alarm unit cost

C_2 – unit operation cost of failure

C_∞ – asymptotic average cost per stage

Probabilistic diagnosis (Sec. 4)

β_i – state of beacon i ($\mathcal{S}_\mathcal{B} = \{$Functional, Intermittent, Hard Failure$\}$)

τ_j – state of train j ($\mathcal{S}_\mathcal{T} = \{$Functional, Faulty$\}$)

$\mathcal{B}, \mathcal{T}$ – sets of beacons and trains

O_{ij} – observation (alarm) from train j at beacon i

ψ_{ij} – observation potential (factor) for the pair (β_i, τ_j)

$\phi_\mathcal{B}, \phi_\mathcal{T}$ – prior potentials on beacons and trains

$\mu_{x\to f}$ – variable-to-factor message

$\mu_{f\to x}$ – factor-to-variable message

$b(x)$ – belief (marginal estimate) at variable x

α – damping factor of LBP

ϵ – convergence threshold on message variation

$K_{\max}$ – maximum number of LBP iterations

REFERENCES

Bertsekas, D. (2012). *Dynamic programming and optimal control*. Athena Scientific.

Bewick V. Cheek, J., L. Ball. (2004). Statistics review 13: Receiver operating characteristic curve. *Critical Care*, *8*.

Cooper, G. F. (1990). The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial Intelligence*, *42*(2–3), 393–405.

Ghahramani, Z., & Jordan, M. I. (1997). Factorial hidden markov models. *Machine Learning*, 1–31.

Koller, D., & Friedman, N. (2009). *Probabilistic graphical models: Principles and techniques*. MIT Press.

Kschischang, F. R., Frey, B. J., & Loeliger, H.-A. (2001).

Factor graphs and the sum-product algorithm. *IEEE Transactions on Information Theory*, *47*(2), 498–519.

Murphy, K. P., Weiss, Y., & Jordan, M. I. (1999). Loopy belief propagation for approximate inference: An empirical study. In *Proceedings of the fifteenth conference on uncertainty in artificial intelligence* (pp. 467–475).

Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. Morgan Kaufmann.

Ross, S. (2010). *Introduction to probability models*. Elsevier.

Schön, W., Benso, J.-N., Larraufie, G., Moëns, G., & Poré, J. (2019). *Railway signalling and automation* (Vol. 3). La Vie du Rail.