System-Level CBM/PBM Aggregation: A Unified Framework for Proactive and Reactive Metrics in Redundant Architectures

Franck DESSERTENNE[1]

[1] *'Data Driven Services & Maintenance' Department*

Airbus Operations SAS, 316 route de Bayonne, 31060 TOULOUSE, FRANCE

franck.dessertenne@airbus.com

Abstract

Assessing the in-service performance of Condition-Based Maintenance (CBM) and Predictive-Based Maintenance (PBM) solutions remains a challenge, as traditional reliability formulas often fail to capture the true nature of proactive operations. First, this paper formally proves that applying classic reactive formulas to CBM/PBM KPIs (Key Performance Indicators) leads to a quantifiable underestimation of the actual Mean Time Between Failures (MTBF). While previous work established aggregation laws for "in-series" equipment configurations, this study extends the mathematical framework to "in-parallel" architectures. Indeed, such specific aggregation laws are essential for evaluating the global performance of CBM/PBM solutions applied to fault-tolerant systems involving redundancy.

Subsequently, this paper introduces a unified theory based on two distinct frames of reference: the "Operational Timeline," using the Mean Time Between Proactive Removals (MTBPR) to measure logistical workload, and the "Effectiveness Timeline," using the Mean Lifetime Reduction (MLR) correction term to account for the residual life lost due to early removal. Following a review of the laws governing MTBF and the reliability function, we demonstrate the laws applicable to key CBM/PBM indicators, including Recall and Precision-based metrics.

Finally, this framework elevates equipment-level performance into system-wide strategic indicators. Ultimately, this enables a holistic evaluation of the global performance of a CBM/PBM solution at the system or aircraft level by balancing the Proactive Aggregated Recall (PAR) for detection coverage, and the Proactive Aggregated Precision (PAP) for detection confidence.

1. Introduction

The transition from reactive to Condition-Based Maintenance (CBM) and Predictive-Based Maintenance (PBM) introduces new challenges in assessing in-service performance. Specifically, applying standard formulas designed for reactive maintenance to evaluate the Mean Time Between Failures (MTBF) in a proactive context inherently leads to an underestimation. Despite this mathematical reality, this assertion sometimes remains debated among industry stakeholders, some of whom even suspect it may lead to an overestimation.

To resolve this ambiguity, the first objective of this paper is to formally demonstrate and quantify the systematic underestimation of MTBF that occurs when legacy reliability formulas are applied without correction.

Building upon this foundation, the paper highlights the necessity of a tailored mathematical framework for CBM/PBM evaluation. While previous work (Dessertenne, 2025) established the foundational aggregation laws for "in-series" configurations, it is necessary to complete the picture with the "in-parallel" topology. The "in-series" model is the workhorse at the equipment level - where most internal components represent single points of failure - but it is insufficient for complex aerospace architectures. To address redundancies introduced by design for reliability or safety requirements, the "in-parallel" configuration is strictly essential.

Analyzing CBM/PBM performance through both architectural lenses is critical to capture the true value of predictive technologies. At the equipment level, the goal is to evaluate the algorithm's efficiency in anticipating individual component failures. However, at the system or aircraft level, where critical functions (e.g., hydraulic systems, flight control computers) are protected by redundancies, the "in-parallel" model becomes the key to unlocking the solution's overall operational effectiveness. Incorporating this parallel framework is not merely a mathematical curiosity; it is the only way to provide a quantifiable, reliable assessment of a CBM/PBM solution in a complex, fault-tolerant system.

Ultimately, this paper presents a complete, end-to-end mathematical framework that rigorously connects a system's architectural design (series versus parallel) to CBM/PBM operational KPIs (MTBPR, NDF, Recall, etc.) that are critical for maintenance organizations assessing in-service performance.

Furthermore, to evaluate the global performance of redundant architectures, this paper introduces specific system-level metrics, namely the Proactive Aggregated Recall (PAR) and the Proactive Aggregated Precision (PAP), allowing for a balanced assessment of operational efficiency.

The subsequent sections are structured as follows: Section 2 outlines the problem statement. Section 3 formally proves the underestimation of the MTBF using classical formulas and quantifies the gap. Section 4 introduces the "in-parallel" framework, providing formally proven aggregative formulas applicable to MTBF-like metrics. Section 5 extends these formulas to other end-to-end performance KPIs and highlights the value of strategic aggregative indicators. Section 6 illustrates these concepts through a system-level use case using a standard architectural example. Section 7 provides a comprehensive summary of all applicable aggregation laws. Finally, Section 8 presents the conclusion and next steps.

2. PROBLEM STATEMENT

A fundamental challenge in assessing modern aircraft maintenance performance stems from the inherent tension between classical reliability theory and the operational reality of predictive maintenance. Traditional industry standards, such as MSG-3 or conventional MTBF calculations, are founded on a reactive paradigm where system states are strictly binary: functional or failed. As the aviation industry transitions toward next-generation, data-driven maintenance frameworks (often conceptualized within the spirit of MSG-4), reconciling legacy metrics with proactive capabilities has become a critical bottleneck.

This binary logic falters when applied to Condition-Based or Predictive-Based Maintenance. In these contexts, proactive removals intentionally truncate a component's operational life to preempt failure. "Such proactive interventions introduce a 'Missing Life' bias, wherein the Mean Lifetime Reduction (MLR) is misinterpreted as an intrinsic loss of reliability rather than the deliberate operational trade-off accepted to prevent unscheduled failures. This results in a systematic underestimation of the true MTBF.

To contextualize this, the following figure from (Dessertenne, 2025) illustrates the relationship between these key CBM/PBM operational metrics:

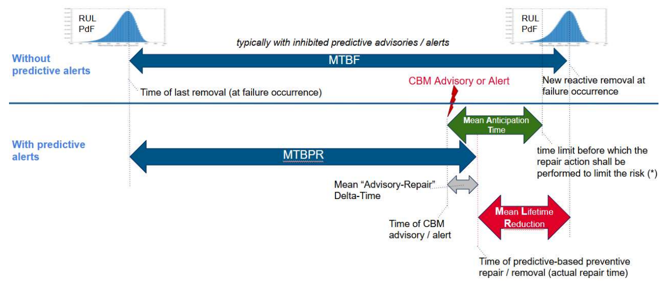

Figure 1. Links between operational metrics in a perfect proactive situation (Recall = Precision = 100%)

This discrepancy is severely amplified within fault-tolerant architectures - such as the redundant hydraulic, pneumatic, or avionics systems of commercial aircraft - where parallel configurations mask individual component health via system-level resilience. Indeed, the operational value of a proactive removal fundamentally depends on the architecture. While a proactive action in a series configuration directly averts a functional system failure, the first proactive intervention in a redundant architecture prevents a loss of redundancy. In these latter contexts, the value of CBM/PBM extends beyond preventing immediate total system loss to the 'Restoration of Redundancy': the capability to identify and rectify the initial failure in a parallel branch before it evolves into a latent risk.

Without a unified aggregation layer, it is impossible to reconcile the 'Operational Timeline' observed by the airline (the physical removal of LRUs) with the 'Effectiveness Timeline' reflecting intrinsic system health (the physical degradation of internal components). Consequently, operators cannot obtain a clear, holistic view of the Return on Investment (ROI) of predictive suites. To bridge this gap, this paper introduces a unified mathematical framework designed to extend Failure Mode and Equipment-level (LRU) performance into complex, redundant architectures.

3. DEMONSTRATION OF THE MTBF UNDERESTIMATION USING A CLASSIC MEASUREMENT STRATEGY

3.1. Formal Proof

Previous work (Dessertenne, 2025) established the necessary framework for CBM/PBM operational metrics. Adopting the nomenclature from this study, we specifically distinguish the 'No Fault Found' (NFF) rate for reactive removals and the 'No Degradation Found' (NDF) rate for proactive removals. Consequently, the False Alarm Ratio (FAR) is derived as an aggregated metric covering all incorrect advisories. This composite rate integrates the proportion of incorrect advisories from the diagnostic solution (expressed via the NFF rate) and the CBM/PBM solution (expressed via the NDF rate). Consequently, the FAR serves as a global indicator, combining the populations of proactive and reactive removals.

Fundamentally, we calculate the FAR (also named FDR -False Discovery Rate in statistics and data science-) as a weighted average of the NDF and NFF rates, proportional to the volume of events generated within each category (i.e., the specific count of CBM/PBM advisories linked to proactive removals versus the count of reactive removals).

We first define the constituent operational rates:

R_p: the total rate of proactive removals, where:

$$R_p = 1 / MTBPR$$

R_r: the total rate of reactive removals, where:

$$R_r = 1 / MTBRR$$

<u>Note</u>: MTBRR (Mean Time Between Reactive Removals) corresponds to the traditional (historical) usage of MTBUR. As discussed in the previous framework, the term ‘MTBUR’ is inherently misleading within a CBM/PBM context.

Based on these rates, the False Alarm Ratio (FAR) is derived as the ratio of unconfirmed removals to the total removal flux:

FAR = Rate of unconfirmed removals / total rate of removals

$$FAR = \frac{rate\ of\ removals\ with\ unconfirmed\ faults\ or\ degradation}{total\ rate\ of\ removals}$$

Consequently, the FAR is defined as the volume-weighted False Alarm Ratio, representing the weighted average of maintenance 'waste' (NDF and NFF) relative to the total removal volume:

$$FAR = \frac{NDF\,.\,R_p + NFF\,.\,R_r}{R_p + R_r} \quad (1)$$

Replacing R_p and R_r by their definition, we have:

$$FAR = \frac{NDF\frac{1}{MTBPR} + NFF\frac{1}{MTBRR}}{\frac{1}{MTBPR} + \frac{1}{MTBRR}} \quad (2)$$

Multiplying both the numerator and the denominator by the product of MTBPR and MTBRR yields the following simplified form:

$$FAR = \frac{NDF\,.\,MTBRR + NFF\,.\,MTBPR}{MTBRR + MTBPR} \quad (3)$$

It is important to address a common initial intuition - albeit a biased one - which suggests applying the standard formulation used in reactive maintenance contexts. This approach typically relies on the Mean Time Between Unscheduled Removals (MTBUR) in the sense of Mean Time Between Reactive Removals (MTBRR).

$$\frac{1}{MTBF} = \frac{1-NFF}{MTBUR} \quad (4)$$

Thus, there is a natural temptation to consider this classic formula (4) to estimate the MTBF in a CBM/PBM context, using the following (in fact biased) equation:

$$\frac{1}{MTBF_{estimated}} = \frac{1-FAR}{MTBR} \quad (5)$$

using MTBR: the Mean Time Between Removals.

Because proactive interventions are inherently unscheduled, the MTBR effectively functions as a global 'MTBUR' in a CBM/PBM context, conflating reactive and proactive removals. To resolve this ambiguity, it is necessary to use the specific term MTBRR to strictly denote reactive removals, as established previously.

Crucially, analysis reveals that the formulation in (5) introduces an inherent mathematical bias into the reliability assessment. To substantiate this, we apply the fundamental principle of rate summation for Mutually Exclusive and Collectively Exhaustive (MECE) events. The total removal rate is defined as the sum of the constituent rates:

$$\frac{1}{MTBR} = \frac{1}{MTBPR} + \frac{1}{MTBRR} \quad (6)$$

This is equivalent to:

$$MTBR = \frac{MTBPR\,.\,MTBRR}{MTBPR + MTBRR} \quad (7)$$

Here, MTBR represents the effective mean interval between removals, aggregating all removal sources (both reactive and proactive).

To account for the 'Missing Life' bias, we introduce a corrected metric, $MTBPR^*$, which restores the consumed life to the proactive interval:

$MTBPR^* = MTBPR + MLR$

Referring to the MECE principle on ratios, we define $MTBR^*$ based on (6) as follows:

$$\frac{1}{MTBR^*} = \frac{1}{MTBPR^*} + \frac{1}{MTBRR} \quad (8)$$

This corresponds to the following definition:

$$MTBR^* = \frac{MTBPR^* \cdot MTBRR}{MTBPR^* + MTBRR} \quad (9)$$

Finally, we define by symmetry the corrected False Alarm Ratio, FAR*, to satisfy the following relationship:

$$\frac{1}{MTBF} = \frac{1-FAR^*}{MTBR^*} \quad (10)$$

Combining (3) and (7) leads to:

$$\frac{FAR}{MTBR} = \frac{NDF}{MTBPR} + \frac{NFF}{MTBRR} \quad (11)$$

Moreover (5) can be expressed as:

$$\frac{1}{MTBF_{estimated}} = \frac{1}{MTBR} - \frac{FAR}{MTBR} \quad (12)$$

Replacing $\frac{1}{MTBR}$ using (6) and $\frac{FAR}{MTBR}$ using (11), we obtain:

$$\frac{1}{MTBF_{estimated}} = \frac{1-NDF}{MTBPR} + \frac{1-NFF}{MTBRR} \quad (13)$$

This needs to be compared with (20) from (Dessertenne, 2025):

$$\frac{1}{MTBF} = \frac{1-NDF}{MTBPR+MLR} + \frac{1-NFF}{MTBRR} \quad (14)$$

'$MLR \geq 0$' is obviously a true assertion, since the Mean Lifetime Reduction (MLR) is inherently non-negative.

So we have: $$\frac{1}{MTBF_{estimated}} \geq \frac{1}{MTBF} \quad (15)$$

So $MTBF \geq MTBF_{estimated}$.

Conclusion: This formally proves that estimating MTBF without bias correction - specifically, without accounting for the MLR impact inherent to the CBM/PBM context - results in a systematic underestimation of the true in-service MTBF.

Besides, to quantify this systematic underestimation and provide actionable tracking metrics for industry applications, the corrected and estimated False Alarm Ratios (actual and estimated) can be directly isolated by rearranging equations (10) and (5):

$$FAR^* = 1 - \frac{MTBR^*}{MTBF} \quad (16)$$

$$FAR = 1 - \frac{MTBR}{MTBF_{estimated}} \quad (17)$$

To align these metrics, we introduce a correction factor defined as the ratio of the corrected to the uncorrected False Alarm Ratio:

$$corr_{factor} = \frac{FAR^*}{FAR} \quad (18)$$

This $corr_{factor}$ serves to quantify the explicit leverage of the proactive-based maintenance context on traditional operational metrics. By establishing the ratio between unbiased and biased False Alarm Ratios, it provides a single, actionable coefficient to evaluate how severely early component removals warp perceived system reliability. This analytical link is explored conceptually below and quantified under realistic operational scenarios.

3.2. Explaining with timelines considerations

Proactive removals inherently truncate a component's operational life, introducing a 'Missing Life' bias quantifiable as the Mean Lifetime Reduction (MLR). Consequently, applying standard reactive reliability formulas without correcting for this MLR results in a systematic underestimation of the true MTBF.

To resolve this fundamental temporal paradox, the proposed framework mathematically reconciles two distinct frames of reference (cf. Figure 2) to align differing operational perspectives. On one hand, the Effectiveness Timeline (involving $MTBPR^*$ and $MTBR^*$) captures the system's intrinsic health by mathematically restoring the "Missing Life" to prevent statistical bias; this theoretical frame allows the data scientist to objectively evaluate algorithmic robustness and compute the true False Alarm Ratio (FAR^*). On the other hand, the Operational Timeline (involving MTBPR and MTBR) reflects the empirical reality of the airline's schedule; it aggregates all physical removal events to provide the maintenance planner with an accurate measure of the logistical burden, frequency of interventions, and overall resource utilization.

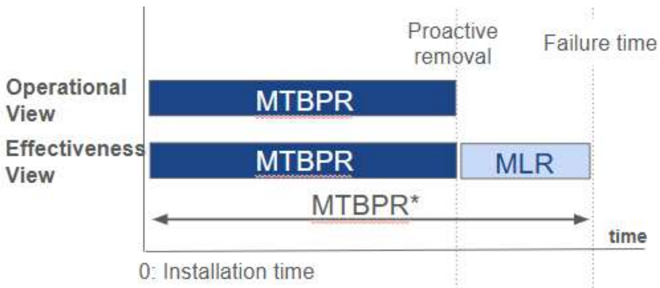

Figure 2: Comparison of the Operational Timeline and the Effectiveness Timeline .

3.3. Numeric application:

To illustrate these concepts, consider a scenario governed by the following operational assumptions:

-Immediate Action: Maintenance is performed immediately upon the triggering of a CBM/PBM alert.

-Ground Truth Validation: Degradations and operational faults are confirmed via a highly sensitive test bench, applying appropriate detection thresholds for proactive and reactive events respectively.

-Numerical Case Study: let us adopt the following baseline parameters:

- Actual MTBF = 25000 FH
- NFF rate = 10%
- NDF rate = 15%
- Measured (actual) MTBRR = 43600 FH
- Measured (actual) MTBPR = 38800 FH

This leads, by computation, to following results:

- MTBR = 20530 FH
- Recall = 48.4% *cf formula (43) in (Dessertenne, 2025)*
- FAR = 12.6% *cf formula (3)*

To determine the Mean Lifetime Reduction (MLR), we utilize again (14). Rearranging the terms to solve for MLR gives:

$$MLR = \frac{1-NDF}{\frac{1}{MTBF} - \frac{1-NFF}{MTBRR}} - MTBPR \qquad (19)$$

Applying the example values results in:

- MLR = 5110 FH
- Ratio MLR/MTBF = 20.4%

However, applying the classical estimation formula (5) without correction leads to the following discrepancy:

- Underestimated MTBF = 23502 FH
- $corr_{factor}$ = 98.8%
- MTBF underestimation = 6%

Intermediate Conclusion: by reducing the measured MTBR (which effectively acts as the MTBUR in a mixed maintenance environment), the CBM/PBM strategy introduces a mathematical penalty. As demonstrated, utilizing classical formulas without accounting for this distortion results in a quantifiable underestimation of the MTBF.

Financial and Operational Implications: while a 6% distortion (in this example) may appear negligible in a theoretical context, a systematic underestimation of MTBF carries significant financial and legal consequences:

- Spare Parts Provisioning (CAPEX): it drives the over-provisioning of inventory, freezing millions of dollars in unnecessary capital to support a failure rate that does not exist in reality.
- Direct Maintenance Costs (OPEX): it artificially inflates projected maintenance budgets, forcing operators to allocate financial reserves for phantom failure events.
- Contractual Penalties & Pricing:
  - For operators: it weakens positions in Power-by-the-Hour (PBH) pricing negotiations by suggesting lower reliability, potentially leading to unwarranted risk premiums.
  - For suppliers: it exposes OEMs to unjustified warranty claims or penalties if the measured (underestimated) MTBF erroneously falls below the minimum design guarantees required by aircraft integrators.

Sensitivity Analysis: The "Paradox of Efficiency"

To further illustrate the impact of the 'Missing Life' bias, it is instructive to consider a scenario driven by a more aggressive CBM/PBM strategy (characterized by earlier anticipation). In this second case, we retain the same intrinsic reliability (True MTBF = 25000 FH, NFF rate = 10%, NDF rate = 15%).

However, we modify the following operational metrics to reflect high-intensity prevention:

- MTBPR = 20000 FH
- MTBRR = 100000 FH

Applying the corrected calculation methodology yields the following results:

- Recall = 77.5%
- MLR = 7419 FH
- Ratio MLR/MTBF = 29.7%
- Underestimated MTBF = 19418 FH
- MTBF underestimation = 22.3%

This comparison highlights a critical paradox: the more aggressive the CBM/PBM strategy is in pre-empting failures (specifically when achieving this via earlier removals, thus increasing the MLR), the more severe the statistical penalty becomes if uncorrected formulas are used.

Having formally demonstrated that standard formulas lead to a systematic MTBF underestimation at the equipment level, we have established the necessary 'clean' inputs for higher-level analysis. This correction is not merely theoretical; it is a mandatory prerequisite. Indeed, attempting to aggregate biased component metrics into a system-level view would only amplify the initial error. With this unbiased MTBF baseline now established, we can confidently proceed to the second challenge: aggregating within redundant 'in-parallel' architectures to evaluate true system coverage.

4. AGGREGATIVE LAWS FOR MTBF-LIKE METRICS IN PARALLEL CONFIGURATION

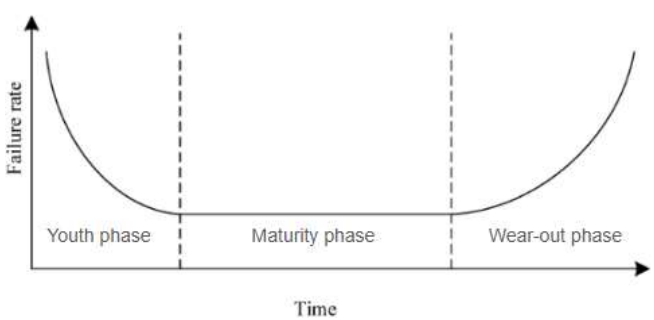

Figure 3. bathtub reliability curve

In aeronautical engineering, the constant failure rate approximation is widely adopted because operation is generally restricted to the flat region of the 'Bathtub Curve'. This is achieved by:

- Eliminating infant mortality: the initial non-linear phase is mitigated through robust 'burn-in' or stress testing to suppress early-life failures.
- Preempting wear-out**:** The final non-linear phase is avoided by enforcing strict safe-life limits, ensuring parts are retired before age-related degradation accelerates.

Consequently, assuming constant failure rates λ_i, the Reliability function $R_i(t)$ is expressed as:

$$R_i(t) = e^{-\lambda_i t} \quad (20)$$

where $\lambda_i = \frac{1}{MTBF_i}$

The Reliability function R(t) - also referred to as the Survival function S(t) - denotes the probability that a component remains functional at time t. It is related to the system's failure probability, defined by the Cumulative Distribution Function F(t), through the following principle:

$$R_i(t) = 1 - F_i(t) \quad (21)$$

For simplicity, we restrict this analysis to a dual-component parallel configuration. Note that complex system architectures can be modeled by iteratively applying series and parallel laws to build an aggregated system-level view.

In a parallel configuration of two independent assets, the system fails only if both components fail. Thus, the unreliability of the parallel system, $F(t)_{//}$, is the product of individual failure probabilities:

$$F(t)_{//} = F_1(t) \times F_2(t) \quad (22)$$

Substituting the reliability terms:

$$F(t)_{//} = (1 - R_1(t)) \times (1 - R_2(t)) \quad (23)$$

The system reliability is the complement of this failure probability:

$$R(t)_{//} = 1 - F(t)_{//} \quad (24)$$

Which expands to:

$$R(t)_{//} = R_1(t) + R_2(t) - R_1(t) \times R_2(t) \quad (25)$$

Finally, by definition, the MTBF is the expected value of the time to failure, calculated as the integral of the reliability function over time:

$$MTBF = \int_0^\infty R(t)\, dt \quad (26)$$

So we have:

$$MTBF_{//} = MTBF_1 + MTBF_2 - \int_0^{\infty} e^{-(\lambda_1+\lambda_2)t}\, dt \quad (27)$$

Recognizing that $\lambda = \lambda_1 + \lambda_2$ represents the combined failure rate of a series configuration (where $MTBF_{series} = \frac{1}{\lambda_1+\lambda_2}$), the expression simplifies to:

$$MTBF_{//} = MTBF_1 + MTBF_2 - MTBF_{series} \quad (28)$$

This parallel aggregation law can be directly extended to the proactive and reactive metrics, MTBPR and MTBRR. This extension relies on the equivalent stochastic behavior of their associated events. While individual component degradation is inherently time-dependent (wear-out), standard reliability analysis at the macroscopic level (system or fleet) approximates the aggregated stream of independent events as a Poisson process, supported by the Palm-Khintchine superposition theorem. This foundational assumption ensures analytical tractability, allowing the inter-arrival times for failures, proactive alerts, and reactive removals to be modeled using exponential distributions with respective means MTBF, MTBPR, and MTBRR.

Consequently, by mathematical analogy, we can define a 'Proactive Survival Function', $S_p(t)$, representing the probability that a component has not undergone a proactive removal by time t:

$$S_p(t) = e^{-\frac{t}{MTBPR}} \quad (29)$$

Since the underlying stochastic process - characterized by independent events and constant rates - mirrors that of reactive failures, the mathematical form of the distribution is identical. Consequently, the derivations previously applied to the reliability function R(t) are equally valid for the proactive survival function S(t).

$$MTBPR_{//} = MTBPR_1 + MTBPR_2 - MTBPR_{series} \quad (30)$$

Extending this logic to the metrics MTBRR, we conclude that for a parallel configuration with independent failure modes and constant rates, the following aggregation laws apply:

$$MTBRR_{//} = MTBRR_1 + MTBRR_2 - MTBRR_{series} \quad (31)$$

To fully leverage this unified mathematical framework, it is crucial to delineate the architectural levels at which the stochastic and logistic aggregation laws apply.

Applicability Scope of the Time-Based Frameworks:

Stochastic aggregation laws for time-based metrics (MTBF, MTBPR, MTBRR) apply primarily from the Failure Mode to the Equipment (LRU) level. Reliability engineers rely on teardown feedback to aggregate localized rates and evaluate LRU performance, making both series and parallel laws indispensable for modeling internal redundancies. Conversely, scaling from LRU to System renders these aggregated time-based metrics operationally irrelevant. Since maintenance physically targets individual LRUs regardless of system-level fault tolerance, the global logistic framework prevails over stochastic times at this architectural level.

5. Other Aggregative laws and Proactive Aggregated Recall (PAR)

Having established these foundational relationships, we now turn our attention to other critical performance indicators. In the following section, we derive the parallel aggregation laws for Recall, Precision, NFF, and NDF.

5.1. Demonstration of Aggregative Laws in Parallel Configurations

5.1.1. Normalization Coefficients

Building upon the formalism established in (Dessertenne, 2025), we define the space of all possible Failure Modes FM_i for the considered equipment or system using the following sets:

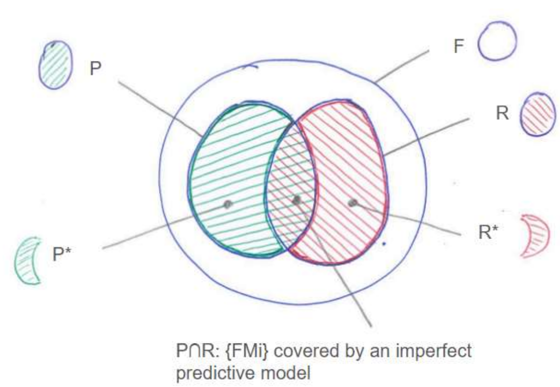

Figure 4. Sets of Failure Modes

- **F**: The complete set of all possible Failure Modes $\{FM_i\}$ for the considered equipment or system
- **P**: The set of Failure Modes (at equipment of system level) that are covered by a predictive monitoring function, therefore associated with predictive and reactive removals.
- **P***: The subset of P containing Failure Modes for which only predictive removals are performed (detection rate = 100%).
- **R**: The set of Failure Modes (at equipment or system level) that can result in a failure leading to reactive removals. Such failures can originate from two sources:
    - Failure Modes that have no predictive monitoring coverage.
    - Failure Modes where the predictive monitoring is imperfect (Recall < 100%).
- **R***: The subset of R associated with only reactive removals. This subset contains only the Failure Modes that are not covered by any predictive monitoring.
- **P∩R**: This intersection represents the set of Failure Modes covered by an imperfect predictive model. It corresponds to failures that are "missed" by the monitoring system, resulting in reactive removals.

Let us define normalisation coefficients respectively on F, P and R:

$$\alpha_i = \frac{\frac{1}{MTBF_i}}{\sum_{k\in F}\frac{1}{MTBF_k}} \quad over\ F \quad (32)$$

We have obviously $\sum_{i\in F}\alpha_i = 1$

$$\beta_i = \frac{\frac{1}{MTBPR_i}}{\sum_{k\in P}\frac{1}{MTBPR_k}} \quad over\ P \quad (33)$$

We have obviously: $\sum_{i\in P}\beta_i = 1$

$$\gamma_i = \frac{\frac{1}{MTBRR_i}}{\sum_{k\in R}\frac{1}{MTBRR_k}} \quad over\ R \quad (34)$$

We have obviously: $\sum_{i\in R}\gamma_i = 1$

5.1.2. Aggregated Recall

Deriving system Recall for parallel configurations follows the logic of redundancy used in reliability engineering. Just as system reliability is the complement of the probability of total failure:

$$R_{Sys} = 1 - \Pi(1 - R_i) \quad (35)$$

In a redundant (parallel) architecture, a CBM/PBM system is considered successful if it detects a fault on at least one branch before a total system functional loss occurs. To calculate this system-level Recall without bias, individual branch detection rates cannot be treated equally; they must be weighted by their relative contribution to the total failure rate, denoted as α_i.

The term "$\alpha_i . Recall_i$" represents the effective probability that a system-level failure originating from branch *i* is successfully detected by the predictive monitoring. Consequently, within the global failure space F, the probability that branch *i* does not provide this detection is its complement. Therefore, the probability of a miss for branch *i* is:

$$P_i = 1 - \alpha_i . Recall_i \quad (36)$$

Under the assumption that both the physical failure modes and their corresponding detection algorithms are statistically independent, the joint probability of a total system miss - where all redundant branches simultaneously fail to trigger an alert - is equal to the product of their individual miss probabilities:

$$P_{sys} = \Pi P_i = \Pi (1 - \alpha_i . Recall_i) \quad (37)$$

Finally, when evaluating functional survival in a parallel configuration, the Proactive Aggregated Recall (PAR) represents the probability that the system avoids a total miss. We apply the complement rule to the result above to define the parallel PAR:

$$PAR_{//} = 1 - \Pi(1 - \alpha_i . Recall_i) \quad (38)$$

For a dual-redundant system, this yields:

$$PAR_{//} = 1 - (1-\alpha_1.Recall_1).(1-\alpha_2.Recall_2) \quad (39)$$

Conceptually, parallel redundancy grants the CBM/PBM solution multiple independent opportunities to intercept a degradation process before a total system loss. Equation (38) translates this joint probability of independent misses through the complement rule. This confirms that parallel configurations yield an aggregated $PAR_{//}$ higher than each component's isolated effective contribution, explicitly quantifying the mathematical benefits of redundant monitoring.

However, it is crucial to note that applying this parallel formulation to compute the PAR depends strictly on the evaluation objective. When aggregating from Failure Modes to a single Equipment (LRU) featuring internal redundancy, the parallel law is strictly required to assess the survival of the equipment's function. Conversely, when scaling from LRUs to a System, the choice of the aggregation law dictates the KPI's philosophy. If the objective is to evaluate the Functional Reliability (preventing a total loss of system redundancy), this parallel law remains appropriate. If the objective is to evaluate the Maintenance Coverage (a flat coverage of all physical assets regardless of system fault-tolerance), the series aggregation law must be applied instead to act as a volume-weighted average of physical anticipations.

5.1.3. Proactive Aggregated Precision (PAP) and Derived Metrics

In our previous foundational work on "in-series" aggregation (Dessertenne, 2025), the concept of the Effectiveness Timeline was introduced to preserve the physical MTBF calculation. At that stage, this timeline was uniformly applied to Precision-based metrics. However, extending this framework to redundant architectures reveals a necessary refinement: while Coverage (Recall) mandates the Effectiveness Timeline, Precision metrics strictly evaluate the actual logistical burden experienced by the shop. Therefore, Precision and its associated aggregation weights must be fundamentally decoupled from the theoretical MLR and evaluated strictly within the Operational Timeline, using raw uncompensated rates.

Based on this corrected consideration, using only the operational baseline, let Nr_i denote the operational volume of all proactive removals generated by branch *i* over a normalized period *T*, based on its proactive removal rate:

$$Nr_i = T / MTBPR_i \quad (40)$$

Let Nc_i denote the theoretical volume of confirmed defects intrinsically intercepted by the algorithm, based on its successful rate, using the Mean Time Between Proactive Successful Removal (MTBPSR):

$$Nc_i = T / MTBSPR_i \quad (41)$$

The system-level theoretical totals are:

$$Nr_{sys} = Nr_1 + Nr_2 \quad (42)$$

$$Nc_{sys} = Nc_1 + Nc_2 \quad (43)$$

Strictly grounded in the logistic framework, the Proactive Aggregated Precision (PAP) is defined as the ratio of valid proactive alerts to the total volume of proactive removals:

$$PAP = \frac{Nc_{sys}}{Nr_{sys}} \quad (44)$$

$$PAP = \frac{Nc_1 + Nc_2}{Nr_1 + Nr_2} \quad (45)$$

By definition, the intrinsic Precision of each branch is $Precision_i = MTBPR_i \ / \ MTBSPR_i$

Substituting Nc_i:

$$PAP = \frac{Precision_1 . Nr_1 + Precision_2 . Nr_2}{Nr_1 + Nr_2} \quad (46)$$

$$PAP = Precision_1 . \frac{Nr_1}{Nr_1 + Nr_2} + Precision_2 . \frac{Nr_2}{Nr_1 + Nr_2} \quad (47)$$

By replacing Nr_i with its definition $T / MTBPR_i$, the term T cancels out, and the removal ratio perfectly equates to the theoretical weighting coefficient β_i defined in Eq. (33). Thus, we obtain the exact closed-form expression:

$$PAP = (\beta_1 . Precision_1) + (\beta_2 . Precision_2) \quad (48)$$

Considering that:

$$Precision_i = 1\text{-}NDF_i \quad (49)$$

The NDF is therefore aggregated using the exact same operational weights:

$$NDF = (\beta_1 . NDF_1) + (\beta_2 . NDF_2) \quad (50)$$

By equivalence, in the reactive scope, the No Fault Found (NFF) rate is similarly aggregated using reactive removal weights:

$$NFF = (\gamma_1 . NFF_1) + (\gamma_2 . NFF_2) \quad (51)$$

Furthermore, since Reactive Precision equals 1 - NFF, it aggregates according to the same law using the γ_i weighting coefficients.

This strictly operational definition of the β_i weights refines the initial framework presented in (Dessertenne, 2025) by preventing any MLR-induced temporal distortion within the Precision aggregation formula. While this paradigm shift preserves the core operational equations from previous work, it consequently supersedes the previously established serial MLR aggregation law. In this updated framework, the system-level MLR is no longer governed by a dedicated aggregation law, but is inherently derived as the direct mathematical solution to the global timeline equation, cf. (19).

5.2. A Strategic Aggregation Coverage KPI

Evaluating parallel, fault-tolerant configurations requires a system-level indicator to assess detection coverage. The Proactive Aggregated Recall (PAR) fulfills this role by establishing the probability that the integrated architecture triggers a predictive alert prior to a functional system failure. By synthesizing component-level performance into a single metric, the PAR serves two strategic functions: performance benchmarking and risk assessment.

- Performance benchmarking: It offers operators and system integrators a standardized global indicator to objectively evaluate and compare the effectiveness of different CBM/PBM solutions, synthesizing complex algorithmic performance into a single coverage metric.
- Risk assessment: It acts also as a unified 'Confidence Index' for decision support systems. When relying on aggregated Health Indicators (HIs) or RULs, considering this system-level Proactive Aggregated Recall (PAR) is essential for managing the risk of missed detections in mission-critical scenarios.

6. ILLUSTRATION & CONSIDERATIONS

6.1. EXAMPLE

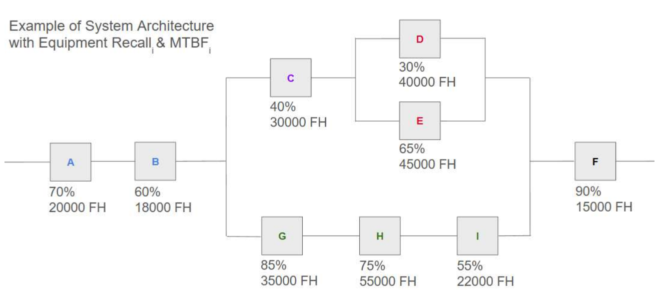

Figure 5. System Architecture Example

The system architecture illustrated above, comprising a hybrid set of series and parallel configurations, serves as a representative case study with assigned Recall and MTBF values for each component.

By iteratively applying the pairwise ('2-to-2') aggregation laws derived previously, we reduce the complex system into simpler blocks. This process yields the first-stage intermediate architecture shown below:

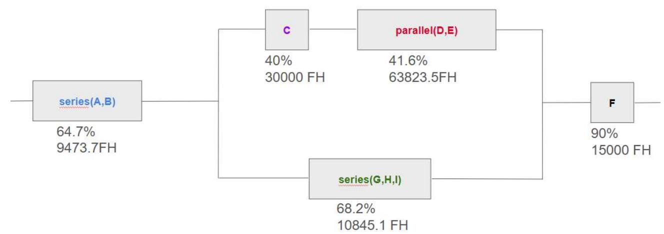

Figure 6. First Aggregation of Previous Example

The final iteration yields the system configuration depicted below:

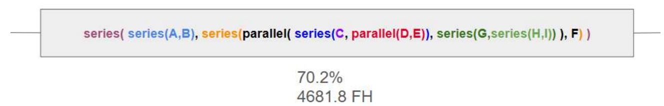

Figure 7. Last Aggregation of Previous Example

Technical Note: when aggregating a 'virtual equipment' (e.g., the composite subsystem resulting from the parallel combination of D and E in Figure 6), it is imperative to renormalize the weighting coefficients at each iteration step against the total failure rate of the new super-system. This ensures that the normalization condition $\sum_{Virtual\ Eqpt} \alpha_i = 1$ is satisfied at each iteration.

Applying the defined aggregation laws, the final system-level metrics are:

- System PAR: 70.2%
- System MTBF: 4681.8 FH

6.2. GUIDING PRINCIPLES

To ensure the validity of aggregated metrics in redundant architectures, the following governing principles must be observed when using a HI or RUL aggregation layer:

- Sensitivity to False Negatives (Recall Dependency): the reliability of a system-level aggregated indicator (e.g. HI/RUL) is directly constrained by the Recall of the underlying assets. Undetected failures or degradations result in an overestimation of the system's health. In redundant configurations, this is particularly misleading, as it may lead to a "healthy" status being displayed even when the system’s fault tolerance has been critically compromised.

- Minimum Recall Level: this represents the minimum system-level coverage required for economic viability. It is the equilibrium point where the operational cost of 'missed proactive removal opportunities' (leading to reactive failures) is balanced by the positive value of optimized downstream decisions (e.g., aircraft swap, immediate repair vs. MEL application …).

These principles demonstrate that an aggregation of HI or RUL indicators is insufficient without an associated confidence indicator. This provides the fundamental justification for the PAR metric, which quantifies the actual 'proactive detection rate' of the system, ensuring that maintenance decisions are grounded in the true health state of systems considering internal redundancies.

6.3. PRACTICAL CONSIDERATIONS: MTBF ESTIMATION AND OBSERVABILITY

Regarding the practical estimation of the Mean Lifetime Reduction (MLR), it is important to clarify that MLR is analytically derived from Equation (19) once the true, unbiased MTBF is established. Therefore, the "observability paradox" raised by proactive interventions fundamentally applies to the estimation of the true MTBF itself. While previous work (Dessertenne, 2025) proposed a baseline methodology relying on a reactive "control population", this approach will not scale as CBM/PBM adoption becomes ubiquitous and unmonitored equipment disappears. To estimate an unbiased MTBF in a fully predictive ecosystem, ongoing research focuses on advanced survival analysis. Specifically, future work may explore the use of Kaplan-Meier estimators and Weibull distribution fitting, where proactive removals are rigorously treated as right-censored data. This methodological evolution is critical for fleet-wide deployment and remains a key perspective for future investigations.

7. SUMMARY OF AGGREGATION LAWS

This section summarizes the pairwise aggregation laws for two independent pieces of equipment (or failure modes). By combining the 'in-series' formulas initiated in (Dessertenne, 2025) with the 'in-parallel' framework derived in this study, these equations provide the complete fundamental building blocks required for complex system aggregation modeling.

Recall:

- In series: $(\alpha_1 \text{ Recall}_1) + (\alpha_2 \text{ Recall}_2)$ (52)
- In parallel: $1 - [(1-\alpha_1\text{Recall}_1) \times (1-\alpha_2\text{Recall}_2)]$ (53)

Proactive Precision: In series & In parallel

- $(\beta_1 \text{ Precision}_1) + (\beta_2 \text{ Precision}_2)$ (54)

NDF: In series & In parallel

- $(\beta_1 \text{ NDF}_1) + (\beta_2 \text{ NDF}_2)$ (55)

Reactive Precision: In series & In parallel

- $(\gamma_1 \text{ Precision}_1) + (\gamma_2 \text{ Precision}_2)$ (56)

NFF: In series & In parallel

- $(\gamma_1 \text{ NFF}_1) + (\gamma_2 \text{ NFF}_2)$ (57)

For MTB_X ∈ {MTBF, MTBPR, MTBRR}:

- In series: $1 / (1/\text{MTB\_X}_1 + 1/\text{MTB\_X}_2)$ (58)
- In parallel: $\text{MTB\_X}_1 + \text{MTB\_X}_2 - \text{MTB\_X}_{series}$ (59)

8. CONCLUSION AND FUTURE WORKS

This paper establishes a comprehensive mathematical framework designed to align predictive maintenance performance metrics with the operational realities of complex, fault-tolerant aerospace systems. We formally demonstrated that applying classical reactive reliability laws in a CBM/PBM context inherently ignores the 'Missing Life' bias, leading to a systematic and quantifiable underestimation of in-service MTBF. Correcting this bias is a mandatory prerequisite for accurate system-level analysis.

Building upon this corrected baseline, we extended equipment-level "in-series" aggregation laws to encompass "in-parallel" redundant architectures. Because deriving generalized closed-form expressions for these parallel configurations becomes algebraically unwieldy, we outlined an iterative, stepwise approach to aggregate fundamental performance KPIs - including MTBF, MTBPR, MLR, Recall, Precision, NDF, and NFF - up to the system and aircraft levels.

Ultimately, this unified framework paves the way for a more mature evaluation of the effectiveness of CBM/PBM solutions. Moving beyond isolated equipment metrics requires a paradigm shift in how we measure proactive maintenance across redundant architectures. By leveraging the specific 'in-parallel' aggregation laws demonstrated in this study, reliability engineers and PHM practitioners can rely on a robust triptych to assess true operational efficiency: the PAR to guarantee system-level detection coverage, and the PAP to control the maintenance burden (false alerts).

Furthermore, the proposed system-level aggregation laws rely on the assumption of statistical independence between the aggregated entities (whether failure modes or equipment) as well as monitoring branches. While this establishes a robust foundational baseline, it inherently represents a theoretical upper bound for predictive coverage. In complex real-world avionics, common-cause failures or sensing and environmental couplings may challenge this assumption. Future evolutions of this framework will therefore aim to incorporate coupling factors (e.g., beta-factor models) to refine these metrics, leveraging accumulating field data to better capture these physical interdependencies.

Having established these theoretical aggregation laws under the assumption of constant failure rates (the flat region of the reliability bathtub curve), future research will focus on dynamic operational environments. Next steps include validating this framework against massive, real-world airline fleet datasets to quantify such aggregated metrics in service. Additionally, further work is required to evaluate the relevance of adapting these aggregation laws for components entering their wear-out phase.

Nomenclature

CAPEX	CAPital EXpenditure
CBM	Condition Based Maintenance
DMC	Direct Maintenance Cost
EQPT	Equipment
FAR	False Alarm Ratio
FDR	False Discovery Rate
FH	Flight Hours
FM	Failure Mode
HI	Health Indicator
KPI	Key Performance Indicator
LRU	Line Replace Unit
MECE	Mutually Exclusive and Collectively Exhaustive
MEL	Minimum Equipment List
MLR	Mean Lifetime Reduction
MTBF	Mean Time Between Failures
MTBPR	Mean Time Between Proactive Removals
MTBPR*	Mean Time Between Proactive Removals added to the Mean Lifetime Reduction
MTBR	Mean Time Between Removals
MTBRR	Mean Time Between Reactive Removals
MTBSPR	Mean Time Between Successful Proactive Removals
MTBUR	Mean Time Between Unscheduled Removals
NDF	No Degradation Found, or rate of equipment with ‘No Degradation Found’ when used in a formula
NFF	No Fault Found, or rate of equipment with ‘No Fault Found’ when used in a formula
PAP	Proactive Aggregated Precision
PAR	Proactive Aggregated Recall
PBM	Predictive Based Maintenance
PDF	Probability Density Function
PHM	Prognostics and Health Management
ROI	Return On Investment
RUL	Remaining Useful Life

References

Dessertenne, F. (2025). CBM / PBM Operational Performance: Concepts & Applicable Formulas For Aligning Model Performance Metrics With Operational Ones At Failure Mode And Equipment Levels. *17th Conference of the PHM Society (US)*.

Fitzpatrick, M., & Paasch, R. K. (n.d.). Analytical Method for the Prediction of Reliability and Maintainability and Maintainability Based Lifecycle Labor Costs.

Fu, S., & Avdelidis, N. P. (2023). Prognostic and Health Management of Critical Aircraft Systems and Components: An Overview. *PMC - PubMed Central*.

Fu, S., Avdelidis, N. P., Plastropoulos, A., & Fan, I.-S. (2023). Fusion and Comparison of Prognostic Models for Remaining Useful Life of Aircraft Systems.

Goebel, K., & Rajamani, R. (2021). Policy, regulations and standards in prognostics and health management. *International Journal of Prognostics and Health Management, 12*(1).

Li, J., & Xu, H. (2012). Reliability Analysis of Aircraft Equipment Based on FMECA Method.

Luna, J. J. (2021). Metrics, Models, and Scenarios for Evaluating PHM Effects on Logistics Support.

Maataoui, S., Bencheikh, G., Ezziani, M., & Bencheikh, G. (2024). Data-Driven Predictive Analysis for Cutting Machine Failures: A Technical Report on Reliability Optimization.

Meissner, R., Rahn, A., & Wicke, K. (2021). Developing prescriptive maintenance strategies in the aviation industry based on a discrete-event simulation framework for post-prognostics decision making.

Modarres, M. (2016). *Reliability Engineering and Risk Analysis - A Practical Guide* (3rd ed.).

Romeu, J. L. (2006). Understanding Series and Parallel Systems Reliability.

SAE International. (2020). *Integrated Vehicle Health Management - The Technology*.

SAE International. (2023). *JA6268: Design & Run-Time Information Exchange for Health-Ready Components* (1st ed. 2018).

Saxena, A., Celaya, J., Balaban, E., Goebel, K., Saha, B., Saha, S., & Schwabacher, M. (2008). Metrics for Evaluating Performance of Prognostic Techniques. *International Conference on Prognostics and Health Management*.

Saxena, A., Celaya, J., Saha, B., Saha, S., & Goebel, K. (2010). Metrics for Offline Evaluation of Prognostic Performance. *International Journal of Prognostics and Health Management*.

Sen Gupta, J., Trinquier, C., Medjaher, K., & Zerhouni, N. (2015). Continuous validation of the PHM function in aircraft industry. *First International Conference on Reliability Systems Engineering (ICRSE)*, Beijing, China, pp. 1-7. https://doi.org/10.1109/ICRSE.2015.7366505

Teubert, C., Jarvis, K., Corbetta, M., Kulkarni, C., & Daigle, M. (2022). *ProgPy Python Prognostic Packages v1.4*. URL https://nasa.github.io/progpy.

Teubert, C., Pohya, A. A., & Gorospe, G. (2023). An Analysis of Barriers Preventing the Widespread Adoption of Predictive and Prescriptive Maintenance in Aviation.

Torres, C. E. (2024). Challenges in Implementing Condition-Based Maintenance. *Power-MI Blog*.

Wagner, C., & Hellingrath, B. (2020). Supporting the Implementation of Predictive Maintenance – a Process Reference Model. *International Conference on Industrial Engineering and Operations Management*.

Wei, P., Lu, Z., & Tian, L. (2013). Addition Laws of Failure Probability and their Applications in Reliability Analysis of Structural System with Multiple Failure Modes.

Biographies

Franck Dessertenne is a Technical Leader in Health Monitoring, Diagnostics, and Predictive Maintenance at Airbus Operations SAS. He earned a Master's Degree in Engineering from ISMRA–ENSICAEN and a Master's Degree in Materials Science & Radiation from the University of Caen (France), both in 1997.

Since joining Airbus in 2003, he has held several key roles. He began in software design and development for simulation products, including diagnostic tools, before becoming responsible for the technical integration of an airborne maintenance server on a military program. Subsequently, his work has focused on Condition-Based Maintenance (CBM) and Predictive-Based Maintenance (PBM). In this domain, he has served as a Health Engineer and Data Analyst, designing and implementing health monitoring and predictive solutions for various aircraft systems.

Currently, in a technical leadership role, he guides Research & Technology (R&T) activities in health monitoring, diagnostics, and prognostics. His research interests include reactive and predictive maintenance, advanced data analytics, artificial intelligence, and the application of machine learning to complex aerospace systems. He is an active contributor of the PHM Society and is also an active voting member of the International SAE HM-1 committee.