*Accepted for presentation and publication in Proceedings of 2005 IEEE Aerospace Conference*

# Framework for Post-Prognostic Decision Support

Naresh Iyer, Kai Goebel, Piero Bonissone
GE Global Research
One Research Circle
Niskayuna, NY 12309
518-387-4194
{iyerna, goebelk, bonissone}@crd.ge.com

*Abstract*—This paper describes a decision support system (DSS) for use in operational decision making with PHM-specific data. Challenges arise from the large amount of different information pieces upon which a decision maker has to act. Conflicting information from on-board and off-board PHM modules, seemingly contradictory and changing requirements from operations as well as maintenance for a multitude of different systems within strict time constraints make operational decision-making a difficult undertaking. The DSS will enable the user to make optimal decisions based on his expression of rigorous trade-offs between different prognostic and external information sources. This is accomplished through guided evaluation of different optimal decision alternatives under operational boundary conditions using user-specific and interactive collaboration. We present some preliminary results of the use of such a DSS for post-prognostics decision-making.

### TABLE OF CONTENTS

## INTRODUCTION

Prognostics and Health Management (PHM) is a requirement for advanced military aircraft such as the JSF. PHM systems provide two functions: firstly, sensor-based diagnostics detect impending faults and trigger prognostic algorithms to determine remaining component life; secondly, health management systems determine appropriate maintenance actions in response to prognostic predictions,.

Short-term in-flight and on-board prognostics evaluate the impact on the ongoing mission and support actions required to maximize success of current mission as well as aircraft and crew safety. In contrast, Off-Board PHM (OBPHM) algorithms take a larger, system-wide view, beyond immediate mission success. Additionally, the design of all these systems should be in the context of, at least, two operational modes -wartime and peacetime- each of which has a different set of priorities and primary drivers. The ideal and eventual goal is for all of the information to quickly propagate throughout the logistics infrastructure, ultimately leading to mission success, enhanced flight safety, increased sortie generation rate, eliminated false alarms, optimal supply chain management and reduced life cycle costs.

PHM promises significant benefits through decreased downtime, and reduced operational and maintenance logistics footprints [Littles and Buczek, 2004]. However, these benefits are also strongly tied to the decision-making that follows the assimilation and interpretation of prognostics information. Determining the best course(s) of action based on the convoluted set of information to address the goals described above is a non-trivial task. This is because of the large volume of information from different sources that exists and which could be partially conflicting. Additionally, there might be a need to address uncertainty associated with the pieces of information. Thus, attending to the pieces of information before their application in operational decision-making process can itself be a challenging task; it minimally entails the management of multiple sources, uncertainty, inaccuracy and inconsistency.

Also, the presence of conflicting performance metrics for the evaluation of individual courses of action can potentially result in a large possible set of actions. Choice from this potentially large set of actions requires the further application of factors like human insight based on prior experience, knowledge of prevailing conditions, many of which are too dynamic or situational to automate within the OBPHM without human intervention.

---

[1] 0-7803-9546-8/06/$20.00© 2006 IEEE
[2] IEEEAC paper #254, Version 4, Updated January 10, 2006

However, even the most proficient human decision makers are limited in their cognitive capacity to assimilate and analyze large quantities of information. Faced with a large and complex problem space, such as that described above, humans have been traditionally [Tversky, 1969] known to simplify the decision-making process in ways that are not guaranteed to result in the choice of a globally optimal course of action. Hence, the mere introduction of an human expert in the information-rich world of operational decision- making does not ensure good quality decisions; a systematic approach that supports the elicitation of human preferences, information querying and information load-sharing is essential to the optimal course of action being selected from the potential ones.

In other words, what is needed is a decision support system that also ensures that the users of the OBPHM system will make “sound” decisions despite their limited cognitive capacity in handling large quantities of information. Such a decision support system will also provide a mechanism for the discovery and evaluation of optimal decision alternatives subject to operational boundary conditions. It will enable the elicitation and application of various kinds of user preferences and constraints to a large space of multicriterially evaluated scenario-alternatives in order to find optimal alternatives that best satisfy those preferences. Also, it needs to do this while taking into account different prognostic and information sources, equipment status, other variables and constraints related to system logistics, maintenance, and operations. Because of its interactive nature, the decision support system will allow the OBPHM user to collaborate in the decision-making process while driving the selection and evaluation of operational scenarios and plans.

Post-prognostic operational decision making entails making use of prognostic information from both these systems in the evaluation of actions within various logistics platforms, including maintenance, supply chain management, mission planning and mission allocation. The ideal goal of such decision making is the optimal use of all available information in the discovery of the set of actions across the various logistics platforms such that global, system-wide metrics like lifecycle costs, mission success rate, turnaround times, safety are optimized. This involves having to deal with a number of analytical complexities: potential interactions between operation and scheduling of multiple aircraft, interactions between different subsystems (airframe, vehicle systems and mission systems) within an aircraft and, process and resource interactions that exist between maintenance, operations, supply chain management, and life management.

The problem we are tackling involves answering the following question: given that we have reliable algorithms to perform asset-level prognostics, what is the process by which one uses this information to influence logistics across a *fleet of assets*? This is a critical problem since prognostics by itself only produces information which enables the making of logistical decisions that actually address the global goals of improving metrics like availability, costs, reliability etc. across the fleet. We show that this decision-making problem has its own subtleties and complexities and describe our approach to tackle the decision problem.

## DSS FOR POST-PROGNOSTIC OPERATIONAL DECISION MAKING

Figure 1 shows the process flow and modules that are envisioned for the human centered DSS: As numerous heterogeneous data arrive from the On-board PHM, models from the OBPHM, and supporting information, the information gets processed in the Prognostics Decision Support Mechanism (PDSM) which consists of two modules – the Information Processor (IP) that processes the information to make it more usable for decision-making, and a multi-objective decision support system (MODSS).

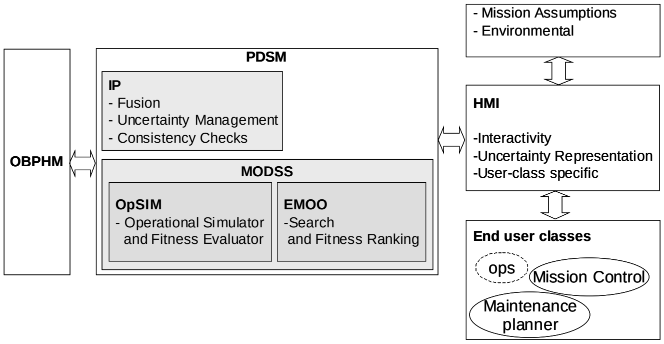

**Figure 1 -** Decision Support System for OBPHM

The IP further consists of modules that aggregate redundant and partially redundant information, check for consistency, and deal with the inherent uncertainty present in the information. The MODSS module explores the space of competing operational scenarios and plans in order to select the optimal set of actions from within these competing alternatives, given the set objectives along which to optimize. The two central elements of MODSS are the *Operational Simulator* module (OpSIM), and the *evolutionary multiobjective optimization module* (EMOO).

The Operational Simulator allows us to model the impact of hypothetical operational decisions on other operational and prognostic variables thereby allowing the performance evaluation of the corresponding operational decisions or actions. The EMOO module, on the other hand, performs an efficient search of the potentially huge space of actions/decisions that would be expensive to search exhaustively. The EMOO module uses the performance evaluation produced by the OpSIM module to rate the fitness of a generated decision or course of action. By using this fitness as means to eliminate the consideration of inferior decisions, the EMOO module in conjunction with the OpSIM module produces only the subset of superior courses of operational actions with highly desirable values for the multiple objectives.

Finally, the information is represented to the user through a human-centered HMI. Whereas, we emphasize the criticality of the two peripheral modules of the DSS shown in Figure 1, namely the IP and the HMI modules, our focus for the current paper remains on the MODSS module since it is the core module dealing with the decision-making functionality of the overall system.

The premise of this paper is that prognostic information – remaining useful component life – ultimately needs to be used as the basis for decision making. As indicated earlier, the decision making spans a wide range of different areas from mission planning to maintenance. In that context, it has to be recognized that remaining life is really expressed as a distribution with which one can associate some user-defined lower and upper bounds as illustrated in Figure 2a. Using a fictitious sequence of missions k through k+l (each one with its own uncertainty due to unknown exact duration and speed/load requirements) a mission reliability can be expressed as shown in Figure 2b. While we do not claim to be able to precisely express the reliability/remaining life relationship at this point (and while it is probably not linear), we will in the following take advantage of this conceptual relationship in the decisioning process.

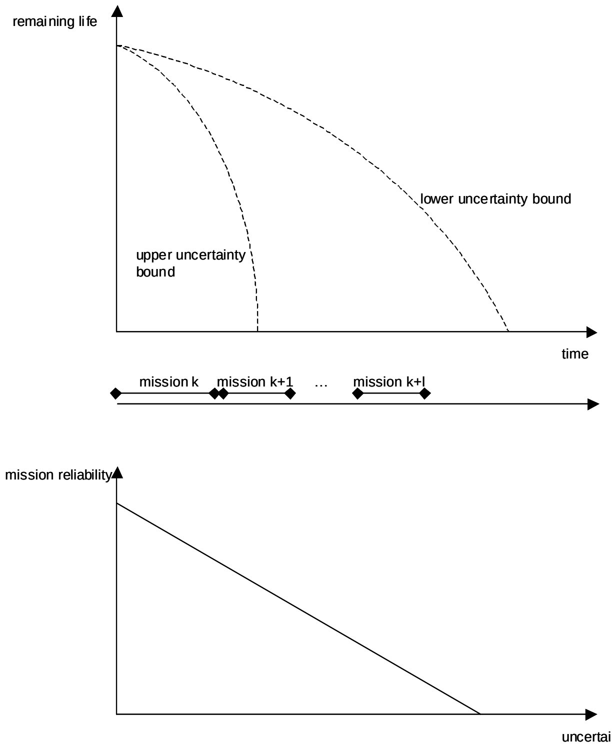

**Figure 2 –** Relationship between remaining life uncertainty and mission reliability

## MULTIOBJECTIVE DECISION SUPPORT

The MODSS module uses the information created by the Information Processor (IP) module to support the end-user in the exploration and evaluation of competing operational scenarios and plans, and in the selection of the optimal set from within these competing alternatives. Since, we expect for the competing alternatives to be evaluated along more than a single performance criterion, we refer to this module as the *Multiobjective Decision Support System* (MODSS). As stated previously, the MODSS module is composed of two modules: the *Operational Simulator* module (*OpSIM*), and the *evolutionary multiobjective optimization* module (EMOO).

Figure 1 shows the MODSS module as a part of the PDSM module in the overall architecture for prognostic decision support. The simulator module models different operational scenarios and produces evaluations for the performance objectives along which each solution or alternative scenario is to be evaluated. For a given flight schedule assignment and its associated maintenance plan, the EMOO module then conducts evolutionary search to find solutions that are optimal in the sense of being *non-dominated* according to the stated performance objectives. We next describe each of the two modules of MODSS in greater detail.

### *Operational Simulation*

The Operational Simulator allows us to model the impact of operational decisions on other operational and prognostic variables. Some inputs that the simulator will operate on include:

1. Prognostic inputs from the OBPHM for components related to aircraft systems like the airframe, safety systems, and mission capability systems,
2. Maintenance capabilities including repair resources, costs and times
3. Mission profile, sorties rates, and other field variables, and
4. Decision thresholds on mission readiness, capability, and safety.

Some outputs produced by the simulator would be: assessment of mission capability and readiness, flight safety, health status, resource availability, requirements of cost, time, and space, evaluation of what-if scenarios with indications of confidence and uncertainty.

The operational simulator could have the following specific capabilities:

1. **Prognostic simulation:** This simulation produces posterior prognostic information under assumptions of a given set of inputs including mission profile.

2. **System simulation:** This simulation can be used to simulate operational events pertaining to the entire system like operational scheduling, maintenance scheduling, supply chain scheduling, repair actions, and various maintenance events in order to study their impact on variables like space, time, cost, and availability. These studies can be used to further drive decision-making towards optimizing the above variables.

3. **User-driven simulation:** This simulation will have the added capability of allowing the decision-maker to schedule various “what-if” scenarios by changing variables like cost, reliability, capability in order to study the impact of those on operational variables of interest to him. He could potentially use such a simulation to perform a cost-benefit analysis while comparing two scenarios with different variable settings in order to help him distinguish the impact of one from the other.

Many of these simulation capabilities are exemplified by the work described in [Iyer, Aragones, Hansen, 2002a] and in [Iyer, Aragones, Hansen, 2002b]. Any or all of these simulation capabilities will allow a given flight schedule assignment and maintenance plan to be *evaluated* along potentially many objectives. The EMOO module can then perform a multiobjective search along these objectives to find the best solution. Each added capability of the above kind makes the simulator more powerful.

An additional dimension along which the power and realism of the simulator can be measured is the level of detail or complexity of the simulator. We identify three levels of resolution that would impart additional realism and power to the operational simulator:

**Low Resolution:** A simulator with low resolution will capture the main system dependencies such as system-subsystem decomposition, static and deterministic relationships, which is sufficient enough to coarsely reflect the primary interactions and relationships between various events and variables of the system respectively. An example of such a simulator would be a script that executes a simple and deterministic set of rules on a given set of inputs to produce deterministic outputs.

**Medium Resolution**: A simulator with medium resolution will be able to additionally address the uncertainty associated with both its inputs and its outputs. Also, it will have the capability to handle modifiable external parameters like mission tempo, weather, etc.

**High Resolution**: A high-resolution simulator will also include stochastic elements and dynamic behavior related to the variables and processes in the system.

In general, as the level of resolution increases, the fidelity of the model improves and we expect the results to be reliable not only qualitatively but also in terms of their relative quantitative differences. By initially using a low-resolution simulator, one can exercise the primary requirements, from beginning to end, and ensure that one has captured the essence of the decision criteria. Subsequently more details can be added thereby increasing the resolution of the simulator and resulting in more realistic evaluations. It is to be noted that there is no *a priori* reason to categorize the various levels of fidelity as above; rather, this categorization is motivated by considerations of how difficult it is to introduce the various elements of fidelity like stochasticity, dynamic behavior, uncertainty modeling.

As mentioned earlier, the operational simulator is only one piece of the MODSS that allows for various alternative operational plans to be modeled and evaluated along multiple criteria of interest to the decision-maker. An additional module is needed that will make use of these evaluations to locate the plans that are optimal along different objectives of interest to the decision maker. This module is the EMOO module that performs multiobjective search to find the best solutions, and is described next.

### *Evolutionary Multiobjective Optimization (EMOO)*

Many real-world problems are often characterized by multiple measures of performance, all of which need to be satisfied simultaneously. The task facing a decision-maker (DM) is typically one of finding an alternative that best satisfies his preferences along all objectives of interest to him For instance, a DM might be asked to generate and assess the design for a new product according to the following (conflicting) objectives: 1) minimize production cost; 2) reduce cycle time; 3) maximize reliability and quality. Since, the objectives are potentially conflicting, the selection of the alternative that best satisfies the decision-maker involves the elicitation of various kinds of preferential information (like the expression of tradeoffs between the criteria, and other preferences like aggregation along the multiple criteria, and so on) from the decision-maker, and in their application to the decision-making process. Depending upon the temporal relationship between elicitation of user preferences and their application to the potential set of alternatives, [Horn, 1997] identifies the following three models of decision-making:

- All of the decision-maker's preferences are elicited *a priori* before applying them to search the decision space.
- All of the decision-maker's preferences are applied after obtaining the set of alternatives. This second approach postpones tradeoffs until large numbers of inferior, dominated solutions are eliminated and the efficient Pareto frontier has been identified.
- The search for the best solution and the elicitation and application of user preferences occurs iteratively. This third approach starts with a multi-criteria search that provides the DM with a preliminary idea of possible tradeoffs. The DM can then make multi-criteria decisions, limiting the search space dimensionality. A new search is then performed in this region of the solution space.

In all of the three cases, we need to perform a search in the objective(s) space and we need to elicit user preferences.

For PHM, the goal is to support the decision-maker(s) in exploring the multiple alternatives that apply, given the operational requirements and constraints, in evaluating them, and in further selecting the optimal operational decisions. The primary issue in searching through a space of numerous alternatives is the combinatorial complexity involved in considering all, or most, of them in the quest to find the optimal ones. The computational complexity of this search becomes a crucial factor when we consider the fact that for each such alternative, we might be required to run a simulation in order to assess the "goodness" of the alternative. An additional complexity arises due to the potential non-linearity and discontinuity of the search space (incremental changes to inputs can produce significantly different outputs); it rules out the application of traditional, gradient-based optimization techniques. What is required is a search mechanism that can efficiently search through the multidimensional, potentially non-linear and discontinuous, space of numerous alternatives. *Evolutionary Algorithms or EAs* are known to be best suited for a searching such spaces.

For multiobjective problems, EAs are typically used towards generating the alternatives that lie on the *Pareto frontier* of the objective values. For a 2-dimensional problem dealing with *Yield* and *Quality* as the performance objectives, Figure 2 indicates the notion of optimality in the sense of belonging to the Pareto frontier. Alternatives that belong to the Pareto frontier are also equivalently referred to as *non-dominated* alternatives.

In accordance with the three models of decision-making identified above, a decision-making technique may obtain alternatives in the Pareto frontier before, after, or during the time that the decision-maker expresses preferences related to which criteria matter to him, and what specific region of the Pareto frontier satisfies his subjective requirements the best. One approach, as exemplified in [Josephson et al., 1998], involves the interactive elicitation, expression, and application of tradeoff preferences, between different objectives, by the decision-maker *after* having obtained alternatives lying on the Pareto frontier (independent of how the search space was sampled).

For an EMOO-based search, one could use the power of EAs to create a well-sampled representation of the Pareto frontier and then use trade-off preferences of the decision-maker to further rank or select alternatives lying on the frontier.

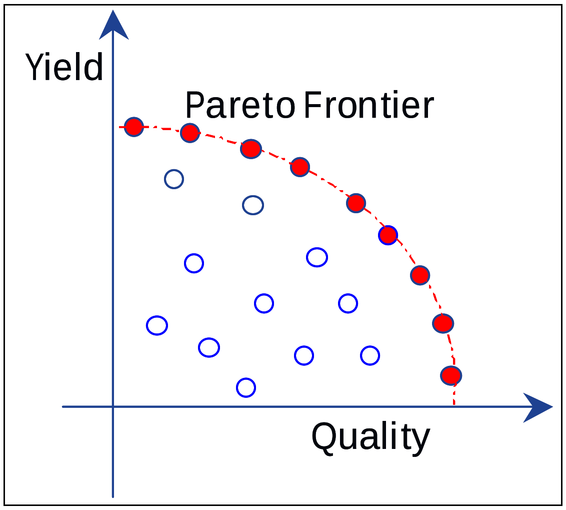

**Figure 3 -** Indicating the idea of a Pareto (non-dominated) frontier along two dimensions

### *EMOO in Prognostics Decision Support*

The use of EMOO will always lead to the identification of *alternative flight schedules* and *maintenance plans* that are *non-dominated* and therefore logically superior to the other, dominated counterparts considered. Figure 4 shows a *hypothetical* example-display of non-dominated, alternative operational plans for a group of 8 aircraft, generated by running EMOO, with the simulator evaluating each generated plan. Each point in the 2-D scatter-plot indicates the *percent-availability* and the aggregated *mission-capability* numbers for a generated operational plan. The aggregated *mission-capability* number for a plan can be obtained from the individual *mission-capability* status of each of the individual aircraft for that plan as evaluated by the simulator.

As seen in the set of non-dominated operational plans, some plans result in good *percent-availability*; however they can result in very few aircraft being fully mission capable (indicated as *FMC* as opposed to *PMC* which represent *partially mission-capable* aircraft). Conversely, certain other operational plans can result in lot more capable aircraft but with low availability numbers. Depending upon present and future requirements, the end-user can potentially use such a plot to select the operational plan that best meets his field requirements. We expect that the user will indicate preferences for the various tradeoffs to rank the alternatives and inspect the results through the HMI following the ranked order. Alternatively the user could filter the original Pareto frontier using constraints and imprecise preferences, and select a subset of alternatives (neighbors in a region of the Pareto) for an iterative selection process. We also foresee the user selecting a portion of a solution (for instance 4 out of 12 tail numbers assigned to a flight schedule) and asking the PDSM system to find alternative selections for the remaining portion of the problem (e.g., the other 8 platforms).

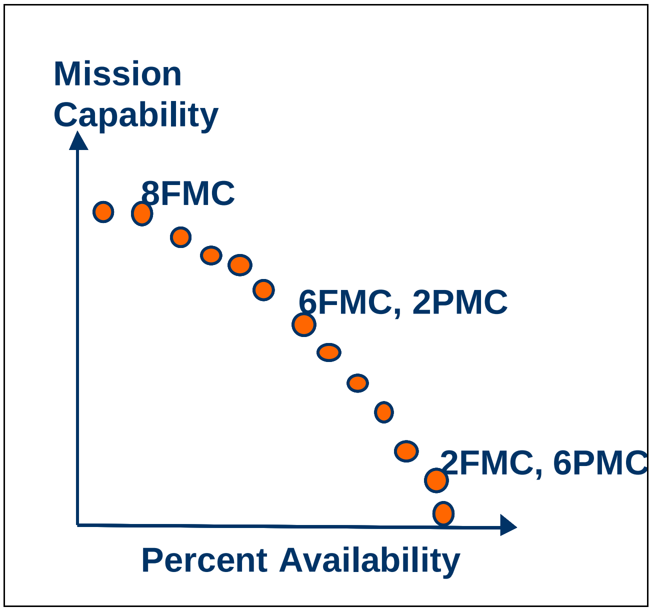

**Figure 4 -** Example showing Multiobjective Trade-Offs in PHM

The EMOO will re-run a reduced version of the problem, with smaller degrees of freedom, to represent a solution that incorporates the user's handpicked selections. If there are feasible non-dominated solutions, the EMOO will generate the corresponding Pareto Frontier and complete the assignments. This interactive method would also allow the user to play "what-if" situations. This will permit the user to manually test the robustness of the solution, thus increasing user's acceptability of the overall system. In addition, many elements of decision-making in an operational environment like interactions between user and the system, state representation and awareness, preference elicitation, query answering, uncertainty representation are all better addressed by means of an interactive module.

## PRELIMINARY EXPERIMENTS

In this section, we present some preliminary results showing some of the features of a DSS that helps in decision-making following the assimilation of diagnostic data. In the future, we plan to extend this system to be able to additionally support the use of prognostics information as well. The essential features of the DSS remain fairly similar, which is why we consider the preliminary round of results as sufficient to exemplify the primary aims of the paper.

### *Problem Formulation*

For our initial experiments, we formulated the post-prognostics decision problem as the following one:
For a time horizon T at a given instant t,
Suppose,

$M_T(t) = \{m_1, m_2, m_3 \ldots\}$ is a set of Missions to be satisfied in time horizon $T$ where,

$m_i = (r_i, c_i, C_i)$ where,

$r_i$ is desired Mission Reliability,

$c_i$ is the Mission Capability and

$C_i$ is a set of constraints related to the time within which mission $m_i$ is to be met

$A = \{a_1, a_2, a_3 \ldots\}$ is a set of available assets where,

$a_j = \{p_1^j, p_2^j, p_3^j \ldots\}$ where $p_i^j$ is part $i$ in asset $j$

$P(t) = \{ (p_1, n_1, c_1, t_1), (p_2, n_2, c_2, t_2), (p_3, n_3, c_3, t_3), \ldots\}$ is an inventory of parts available at time $t$ for use in repair where,

$(p_k, n_k, c_k, t_k)$ is the current inventory with $n_k$ units of availability of the part $p_k$ with of cost of each part being $c_k$ and repair or replacement time associated with part being $t_k$

Then the post-prognostics decision problem is:

What is the best set of assignments from
$P \to A$ (we refer to this as part allocation)
$A \to M_T(t)$ (we refer to this as asset allocation)

such that

$M_T(t)$ is maximally satisfied

while minimizing total cost, part usage, and total time to repair?

*Input Data*

The input data for our preliminary experiments comprises of probabilistic evaluations of expected mission reliability obtained from using an asset based on its current state of health. The current state of health of a given asset is tied to the state of health of its individual parts for which diagnostic information is available. This diagnostic information indicates the reliability of each part of an asset at the current point in time. Since asset reliability is available on a part-by-part basis, overall asset reliability is computed by probabilistic aggregation of the reliabilities of the individual parts taking into account the repair actions conducted on the asset.

We contend that an optimal solution to the post-prognostics decision problem is one which solves both the part allocation problem and the mission allocation problem globally, since both these are inter-related in terms of their impact on overall performance metrics for the fleet of assets. We do this by exhaustively considering all possible combinations of parts-to-asset and asset-to-mission while making sure that all considered parts-asset allocations are constrained by the number of available parts of each kind. Similarly, we make sure that each asset is allocated to one and only one mission.

As stated in Section 2, a genetic search could easily replace the need to consider the exhaustive set of allocations for a given problem. However, we plan to use a GA for our next round of experiments; the use of exhaustive search for a toy problem of two assets, two missions and seven parts per asset is a means to explore the potential of our approach theoretically.

*Results*

After generating all[3] possible allocations going from available parts to assets and assets to missions, we evaluate each global allocation or plan in terms of mission satisfaction metrics like reliability, capability etc. by using asset reliability figures aggregated from the diagnostic information. Other metrics like posterior state of the repair shop, of the inventory can also be used. When prognostic information related to the times that we expect currently flying assets to visit the repair-shop becomes available, the global plan will also need to take the future parts requirement and supply chain management into account.

We provide two modes of decision-support in allowing a decision-maker to choose from the generated set of allocations. For a 2-asset, 2-missions, 7 maximum parts per asset problem, with only one available part on the inventory, we find that only 4374 (of the total 32768) allocations are feasible within part availability. The goal of the DSS is to allow the decision-maker to choose an appropriate allocation plan from the above set. Clearly, this can be a daunting task given that each of these 4374 plans are feasible.

*Constraint-Based Preference Expression*

In this mode, the decision-maker expresses desired levels of performance metrics across each of the current missions like *maximum time within which to dispatch mission, minimum reliability required to fly mission* and so on. The DSS looks through the feasible set of allocations to find the ones that satisfy the specified constraints and presents them to the decision-maker.

In the absence of any allocation that satisfies the constraints, the DSS additionally tries to find a currently infeasible allocation that could potentially satisfy the expressed constraints. Upon finding such an allocation, it indicates to the decision-maker the infeasible allocation as well as the requirements in terms of additional parts that would make this allocation feasible. Thus the DSS translates currently unattainable goals of the decision-maker into a set of actions for the decision-maker by virtue

---

[3] For a 2 asset, 2 mission, 7 maximum parts per asset problem, the exhaustive space of global allocations is 32768.

of which those goals could be attained. If none of the infeasible plans can satisfy the currently expressed constraints, then the DSS indicates that to the decision-maker and communicates the need to weaken some of the constraints.

*Interactive Decision Preference Expression*

In this mode of decision-making, the decision-maker is not in a position to express exact constraints on the performance metrics. In such a case, the decision-maker is presented with all of the available feasible allocations in an interactive fashion that allows him to identify the levels of performance that are attainable across multiple missions simultaneously.

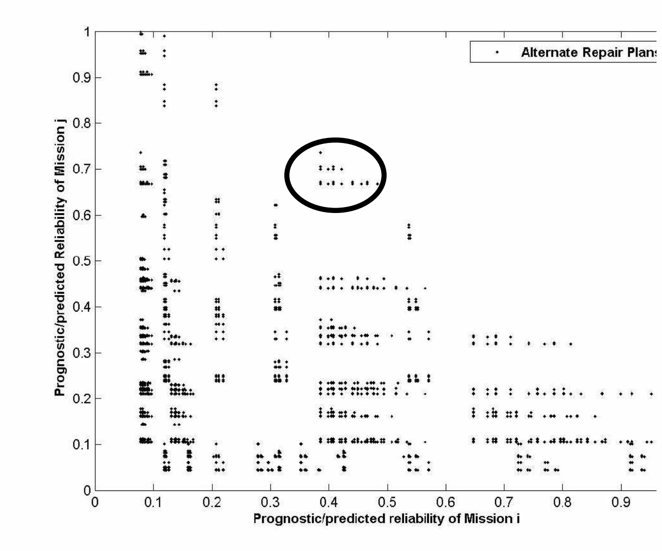

**Figure 5 -** Indicating the decision problem for selection a plan with the prognostic ability to translate each plan in expected mission reliability.

Figure 5 shows an instance of how prognostic information related to expected mission reliability (translated from prognostic health of the asset assigned to the mission) culminates into a decision-making problem, where the decision space is comprised of actions in some logistics platform. In the figure, each point represents a potential *plan* that prescribes the repair actions to be performed on each asset in the repair-shop, as well as the asset to be allocated to a mission. Each such plan is feasible only if the repair actions are constrained by the part availability in the shop (for example, if only one unit of part $P_b$ is available in the shop, a plan that allocates $P_b$ to more than one asset is infeasible[4]).

The plot of Figure 5 shows the intrinsic trade-offs present in the real-world when trying to satisfy multiple missions (*i* and *j* in this case) which compete for the same resources (parts, time, man-power). The plot shows that repair-plans with very high values of *predicted reliability* for a mission *i* are also plans that result in low *predicted reliability* values for competing mission *j* in the deck of missions to be satisfied (and vice-versa). Presenting actors in the logistics platforms with such plots confronts them with the need to understand the competing nature of the metrics they are trying to simultaneously maximize, and thereby presents them also with the opportunity to locate feasible plans that can potentially optimize along all such metrics simultaneously. For example, the set of plans labeled inside the ellipse represent one within which each plan can be expected to result in fairly superior *predicted reliability* for both competing missions[5].

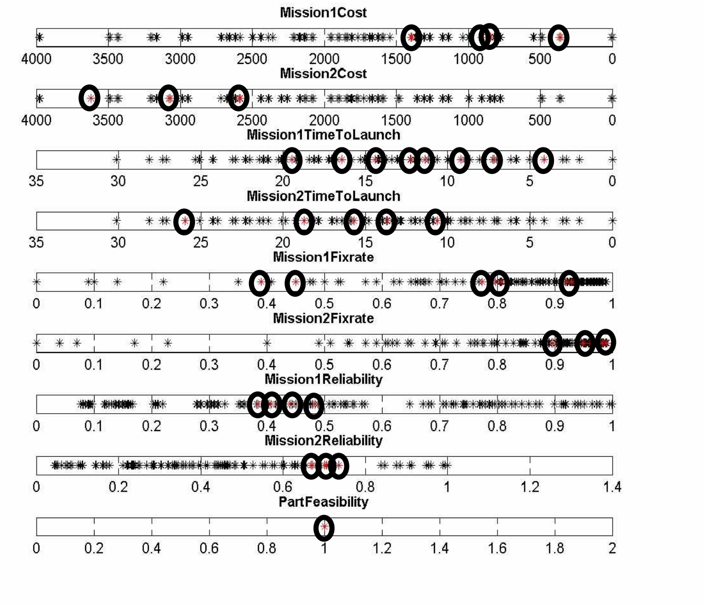

**Figure 6 –** Distribution of Points Selected in Figure 5 across additional dimensions

This gives the decision-maker an opportunity to identify what is globally attainable in the space of feasible allocations/plans. For example, by seeing that there are feasible allocations that allow a reliability metric of at least 0.7 for mission 1 and 0.55 for mission 2, he might discover this subset of allocations to be sufficient based on the mission profiles. By using interactive features like cross-linked plots as in [Josephson et al., 1998], heatmaps and other interactive tools, the user can further explore allocations of interest in the set of visible plots and use this visual exploration to drive his choice of the best plan to implement. The 1D plots[6] shown in (Figure 6) are like range plots that show the evaluation of the individual plans along single dimensions or variables each of which characterize a plan. Such painting and brushing tools are

---

[4] Of course, such infeasible plans can still be considered and translated into a set of part replenishment actions that can be feasibly executed within the mission time constraints.

[5] It is to be noted that *interior points* in the above plot are not to be interpreted as *dominated* points since the above plot is only a 2D projection of the overall metric or objective space.

[6] Several licensed visualization tools provide this capability

available in many interactive tools in the market today. Since any selection made by the user in a single plot is correspondingly indicated in the other open plots, the user is able to see the global evaluation of a plan along all dimensions of importance even if he is making selections on a projected space or set of axes. In summary, either by concrete expression of his preferences or by looking at the available options in an interactive environment, the decision-maker can be empowered to drive the decision-making so as to meet all his global performance metrics.

## SUMMARY AND CONCLUSIONS

The ideal goal of PHM is for all of the relevant prognostic information to quickly propagate throughout the logistics infrastructure, ultimately leading to mission success, enhanced flight safety, increased sortie generation rate, reduced false alarms, and reduced life cycle costs. OBPHM promises significant benefits through decreased downtime, and reduced operational and maintenance logistics footprints. However, these benefits can only be realized if actions are taken subsequent to proper assimilation and interpretation of prognostics information.

In the end, any PHM system will rely on the human decision-makers to provide the insight necessary to make the right operational decisions, based on prior experience and knowledge of conditions that systems like OBPHM cannot observe. However, the complexity of information that needs to be processed will by far exceed the cognitive, information processing capacity of human decision-makers thereby leading to the potential of making suboptimal decisions. In response, we propose a human-centered decision support system that enables and aids users of the OBPHM to make “good” decisions. Such a decision support system will provide direct aiding in the discovery and evaluation of optimal decision alternatives subject to operational boundary conditions and user preferences. We indicate the need for mechanisms that perform information fusion, uncertainty management, decision generation, evaluation, and optimization as crucial and necessary to the success of an effective PHM decision support system. We envision these modules to complement each other inside a robust decision support system for prognostic decision making that will simultaneously guide and allow the OBPHM user to collaborate in the decision-making process while driving the selection and evaluation of operational scenarios and plans that are optimal. We presented preliminary results of a DSS architecture modeled around supporting post-prognostics decision-making.

## REFERENCES

P. Bonissone and K. Goebel, “When will it break? A Hybrid Soft Computing Model to Predict Time-to-break Margins in Paper Machines”, Proceedings of SPIE 47th Annual Meeting, International Symposium on Optical Science and Technology, Vol. #4787, pp. 53-64, 2002.

A. Garga, K. McClintic, R. Campbell, C.-C. Yang; M. Lebold, T. Hay, C. Byington, “Hybrid Reasoning for Prognostic Learning in CBM Systems”, IEEE Proceedings Aerospace Conference, vol. 6, 10-17 March 2001, pp. 2957 – 2969, 2001.

K. Goebel, P. Bonanni, N. Eklund, “Towards an Integrated Reasoner for Bearings Prognostic”, Proceedings of 2005 IEEE Aerospace Conference, paper z11_0704, 2005.

J. Horn, Multicriteria decision making and evolutionary computation. In: Back T, Fogel DB, Michalewicz Z, editors. Handbook of evolutionary computation. Bristol, UK: Institute of Physics Publishing; 1997.

Iyer N.S., Aragones, J. K., Hansen, C. H. (2002a). An Event-Driven Monte-Carlo Simulation for Engine Service Planning. Class 2 Technical Report, GE GRC, Service Algorithms Lab (Information and Decision Technologies Group).

Iyer N.S., Aragones, J. K., Hansen, C. H. (2002b). An Event-Driven Monte-Carlo Simulation for Engine Service Planning: A Design Description Document. Class 3 Technical Report, GE GRC, Service Algorithms Lab (Information and Decision Technologies Group).

John R. Josephson, B. Chandrasekaran, Mark Carroll, Naresh Iyer, Bryon Wasacz, Giorgio Rizzoni, Qingyuan Li, David A. Erb, An Architecture for Exploring Large Design Spaces. Proceedings of the 4th National Conference of the American Association for Artificial Intelligence, Madison, Wisconsin, pp. 143-150, 1998.

J. Littles and M. Buczek, “Engine System Prognosis”, Proceedings of Materials Science & Technology 2004, September 26-29, 2004, New Orleans, Louisiana, 2004.

A. Tversky (1969). Intransitivity of preferences. *Psychological Review,* 76(1), pp 31-48.

G. Vachtsevanos, P. Wang, :Fault Prognosis Using Dynamic Wavelet Neural Networks”, AUTOTESTCON Proceedings, IEEE Systems Readiness Technology Conference , 20-23 Aug. 2001, pp. 857 – 870, 2001.

R. Xu, M. Stevenson, C.-M. Kwan, and L. Haynes. “Bearing Monitoring”, Proceedings of SPIE - Component and Systems Diagnostics, Prognosis, and Health Management, P. Willett, T. Kirubarajan, (Eds), vol. 4389, pp. 216-218, July 2001.

## BIOGRAPHY

***Naresh Iyer*** Naresh Iyer received his Ph.D. in Computer and Information Science from The Ohio State University in 2001. He also has an M.S. in Computer Science and Engineering from the University of South Florida.

He is currently employed as a Senior Research Scientist at the GE Global Research Center in Schenectady, NY. He is a member of the IEEE society and his fields of expertise and interests include artificial intelligence, cognitive science, multiple criteria decision theory, multiobjective optimization and reasoning with uncertainty.

***Kai Goebel*** received the degree of Diplom-Ingenieur from the Technische Universität München, Germany in 1990. He received the M.S. and Ph.D. from the University of California at Berkeley in 1993 and 1996, respectively.

Dr. Goebel joined General Electric's Corporate Research and Development facility in Schenectady, NY in 1997 as a computer scientist after working as a visiting postdoctoral fellow at UC Berkeley from 1996 to 1997. He has carried out applied research in the areas of artificial intelligence, soft computing, and information fusion. His research interest lies in advancing these techniques for real time monitoring, diagnostics, and prognostics. He has fielded numerous applications for aircraft engines, transportation systems, medical systems, and manufacturing systems.

Dr. Goebel is an adjunct professor of the CS Department at Rensselaer Polytechnic Institute (RPI), Troy, NY, since 1998 where he teaches classes in Soft Computing and Applied Intelligent Reasoning Systems.

***Piero Bonissone*** received the BS degree in EE/ME from the University of Mexico City, in 1975, the M.S. in EECS and M.S. in ME from UC Berkeley, in 1976 and 1978, and the PhD in EECS from UC Berkeley, in 1979.

A computer scientist at the General Electric Corporate Research and Development Center (GE CRD) since 1979, Dr. Bonissone has carried out research and projects in Artificial Intelligence, expert systems, simulation, fuzzy sets, soft computing, and data mining.

In 1989 he received the Dushman Award from GE CRD for his work on reasoning with uncertainty. In 1993 he received the Coolidge Fellowship Award from GE CRD for overall technical accomplishments. In 1996 he became a Fellow of the American Association for Artificial Intelligence (AAAI). In 1999 he received the Dushman Award from GE CRD for his work on medical equipment diagnostics.

Dr. Bonissone is an adjunct professor of the ECSE Department at Rensselaer Polytechnic Institute (RPI), Troy, NY, since 1982. Since 1993 he has been the Editor-in-Chief of the International Journal of Approximate Reasoning (North-Holland Publishing Company). He co-edited the book Expert Systems in Structural Safety Assessment' (Springer-Verlag 1989), Uncertainty in Artificial Intelligence 6 (North-Holland 1991), and Uncertainty in Artificial Intelligence 7 (Morgan Kaufmann 1991). He is also one of the three Editors-in-Chief of the Handbook of Fuzzy Computation (Institute of Physics Publishing).

Dr. Bonissone has published more than ninety articles in the area of expert systems, approximate reasoning, fuzzy sets, pattern recognition, decision analysis, and soft computing. He received seven patents from the U.S. Patent Office for his work on reasoning with uncertainty and fuzzy control.