Distributed Agents for Artificial Immunity in Modern Manufacturing

By Dragan Djurdjanovic
Department of Mechanical Engineering
University of Texas at Austin
dragand@me.utexas.edu

1. Introduction

Decades of traditional condition monitoring research in both academia and industry led to precedent-based approaches, i.e. approaches centered on recognizing faulty behavior modes for which a model of the fault indications is created beforehand, based on the engineering knowledge or adequate training data. In the case of highly sophisticated systems that we encounter in manufacturing today, such as robotic fixtures, networked CNC machines or advanced semiconductor tools, the precedent-based approach becomes excessively cumbersome. First, one needs to train the condition monitoring procedures to recognize a large number of potential faults, some of which cannot be anticipated during the design stage. Furthermore, modern manufacturing equipment can perform various operations and thus display highly dynamic behavior. Since faults could manifest themselves very differently under different operating conditions, training of diagnostic units for all possible conditions and all possible faults is practically impossible.

One does not have to look too far to find examples of manufacturers scrambling to deal with unprecedented faulty situations. E.g., intermittent connections in communication networks used to connect various stations in automotive assembly lines [1][2] occur commonly due to contacts compromised by moving robots and workpieces, coolants and improper installation. They happen to be inherently unprecedented faults because every intermittent connection in every plant looks different due to different network configurations, different usage patterns and different fault severities, which means that there are no fault signatures available *a priori*. Just one instance of an intermittent connection on a DeviceNet network in a major automotive assembly plant led to a 4.5 hour downtime caused by the inability to diagnose the problem and find the node that caused it. One should note that a minute of downtime in an automotive assembly line could cost well over $20K [3].

As one moves on to more sophisticated manufacturing processes and systems, consequences of one's inability to cope with unprecedented situations become more frequent and more impacting. For example, in the dataset considered in [4], which corresponds to more than 6 months of operation of a Plasma Enhanced Chemical Vapor Deposition (PECVD) in a major domestic 300mm fab, one could observe 2 downtime instances that lasted more than a week each and each cost close to $1M just in scrapped wafers (lost production damage is probably even bigger). The tool kept messing the wafers, while emitting signatures that were totally new and no one could fix the problem. One of those downtimes was finally resolved after a teleconference with a quantum physics Ph.D., who led the control design and development of the PECVD tool in question. The author saw the same type of a tool causing similar issues in 2 other semiconductor manufacturing companies, again because for a long time no one could find the root causes and do appropriate repairs.

In both examples described above, the manufacturers ended with unknown situations, searching blindly for sources of problems, devising *ad hoc* repair procedures and wasting significant resources. Such “learning by doing” is inexcusably frequent, especially in highly sophisticated, high-value manufacturing areas, which leads to costly equipment downtimes and ineffective repairs, costing U.S. manufacturers hundreds of billions of dollars annually [5]. It seems evident that the traditional precedent-based diagnostic paradigm has reached its limits and a radically new approach is needed to deal with the ever-increasing complexity of modern manufacturing. Furthermore, planning and scheduling of manufacturing operations in the environment plagued by unknown, unprecedented situations represents an exciting research opportunity with potentially enormous positive impact.

2. Coping with Unprecedented Conditions

As described above, the traditional, precedent-based diagnostic approach is based on characterizing and recognizing specific faults that can subsequently be remedied via maintenance and repair actions. Such a paradigm requires a database of fault models based on elaborate training data and/or modeling efforts describing the system behavior in the presence of those faults. Each time an abnormality is detected, the system behavior is checked against those fault models to identify the root cause of the anomaly and determine the appropriate maintenance actions. Faults whose models do not exist in that database lead to lengthy root cause identification efforts, relying heavily on human expertise and costly trial and error procedures.

Nevertheless, practical information needed by maintenance personnel in a manufacturing facility is to localize the source of the fault to the so-called Field Replaceable Units (FRUs), and the specific character of the fault is usually secondary. For example, a maintenance worker on a PECVD tool needs to know if a pendulum valve on the tool is anomalous or not. The information about whether the anomaly was caused by a faulty actuator, or accumulation build up, or any other reason is secondary. The remedial action is the same – replace the pendulum valve. Hence, rather than making positive identification of various faulty behavior regimes, one could focus on localizing the source of anomalous behavior using the paradigm of distributed anomaly detection. Essentially, monitoring can be realized through a set of anomaly detectors (ADs) covering the target system, with each detector that perceives an anomaly splitting into a set of ADs monitoring the pertinent subsystems. Such cascading proliferation of ADs would continue until the FRUs that caused the abnormality are identified, as seen in [6]-[8].

Rather than a database of detectors of specific faulty behavioral modes, this approach requires a database of *only* normal behavior models of the target system and its subsystems, with each AD using the corresponding normal behavior model to detect anomalies as statistically significant departures from that model. Let us show how this approach leads to precedent-free localization of subsystems causing anomalies in the Exhaust Gas Recirculation (EGR) system of an automotive diesel engine (taken from [8]). Initially, an AD monitors the entire EGR system based on the dynamic model of its normal behavior, as indicated in plot (a) of Fig. 1. Once it detects an anomaly, 5 ADs are distributed, as shown in plot (b) of Fig. 1, each using the relevant model of normal behavior of the target subsystem to monitor it.

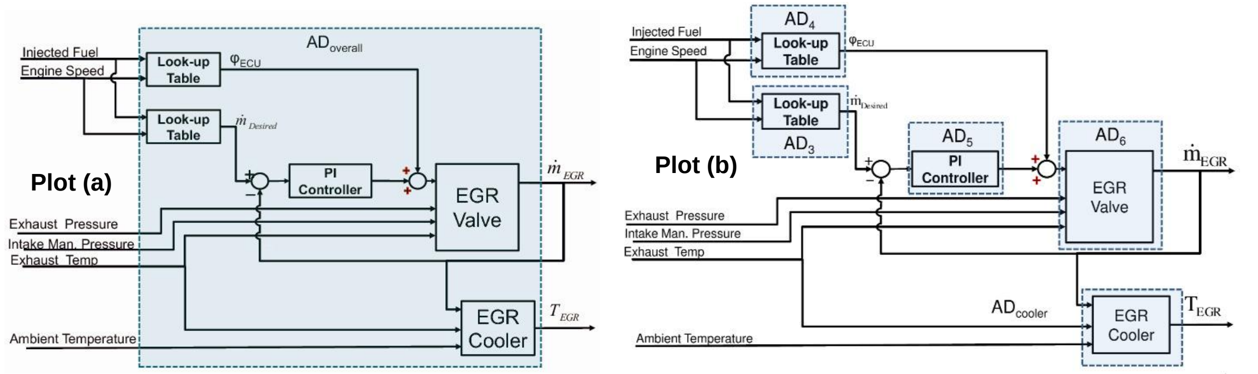

Fig. 1: Distribution of ADs in a diesel engine EGR system. Once an anomaly is detected, the overall AD (depicted in plot a) splits into ADs monitoring pertinent subsystems, as illustrated in plot (b).

The "normalness" of behavior of each of these systems is assessed through the concept of Confidence Values (CV), expressing the overlap of modeling residuals observed during normal behavior and the most recently observed modeling residuals. This quantity fluctuates between 0 and 1, with 1 indicating a perfect match (performance identical to normal) and small values indicating anomalies (using the terminology of the literature on artificial immune systems, CVs are analogous to each detector's "affinity" to the normal behavior of their respective target systems [9]). Fig. 2 shows CVs from the relevant ADs when progressively more severe clogging was simulated in the EGR valve 1350 sec. after the start of the simulation. Interpretation of CVs clearly points to the culprit subsystem (EGR valve), even though no fault signatures were ever collected (only models of normal

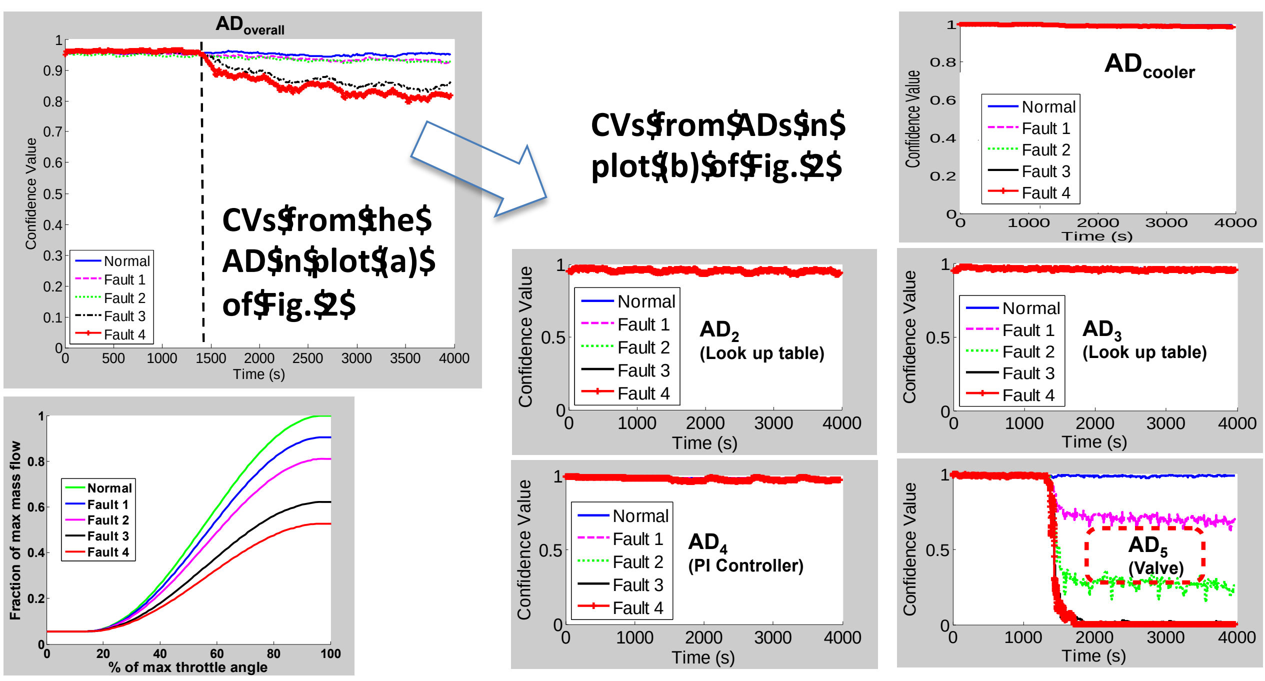

Fig. 2: CVs from ADs that proliferate through the EGR system as progressively severe faults are inserted into the EGR valve. Interpretation of CVs clearly identifies the valve as the source of the problem.

behavior of the EGR system and its subsystems were available). The same concept was also used in [8] to isolate a series faults in the controller and EGR cooler, using only normal models of the pertinent systems.

Similar, precedent-free fault localization has *not yet* been employed in manufacturing. The major issue why is the immense complexity of modern manufacturing machines in which the aforementioned approach is really needed. One aspect of the difficulties lies in the challenges associated with the inherent lack of observability of the underlying phenomena in such machines. E.g., the state of plasma in a PECVD or an etch tool is inherently unobservable, as will be elaborated a bit later. Another reason why we do not see this approach in manufacturing yet lies in the fact that distributed anomaly detection requires understanding of interactions between various subsystems (one needs to know what variables affect the performance of a given system as inputs, and what variables act as its outputs, potentially affecting other systems around it). In an automotive engine considered in [8], this information is known from the very design of the control system. However, in a modern lithography tool, or a modern etcher, hundreds of subsystems and components operate in very different physical domains and at very different time-scales. These systems emit thousands of variables, with interactions between them not being fully understood even by the domain experts.

Let us first tackle the problem of anomaly detection in the case of inherently unobservable phenomena, such as plasma used in CVD, etch and lithography processes. Namely, plasma is a phenomenon described by 3-dimensional fields of magnetic induction, pressures and temperatures. From these fields, one senses only one or a few temperature and pressure points, as well as a few characteristics of the magnetic field in 1-2 selected points. Thus, the state of plasma in between those special points is inherently invisible, though it can be inferred with more or less confidence from the available sensor readings. In addition, variability of operating conditions that modern manufacturing machines experience in the highly flexible and reconfigurable manufacturing environments means that degradation dynamics and fault signatures, which, as mentioned above, are often only probabilistically visible, change with the operating regimes. Assuming that the underlying condition of the system is related to one or several sensor readings in this context, which is the basis of the decades-long direction pursued in the area of condition monitoring, cannot cope with such complexity. Advanced signal processing, statistical analysis and time-series modeling, which worked so well with rotating machinery and traditional manufacturing, do not help in this case.

Recently, a new mathematical construct was devised for degradation modeling and anomaly detection in inherently unobservable processes undergoing variable operating regimes. The new method models the degradation process through a collection of operating regime-specific Hidden Markov Models (HMMs) [4][10]. In this context, the equipment conditions are hidden states that are stochastically related to the observable variables representing the available sensor readings. The stochastic models relating observable variables (sensor readings) with the hidden states (equipment conditions), as well as the dynamics of progression of hidden states are made to be operating regime dependent, which enables context-dependent operating regime-specific degradation modeling.

Fig. 3 illustrates the new concept of degradation modeling and anomaly detection, where interconnected HMMs model the observations and the dynamics of degradation processes corresponding to various operations executed on the monitored system. Following the continuity of degradation, each of the underlying HMMs is made to be unidirectional (indicating that without maintenance, condition of the monitored system keeps progressing towards increasingly degraded states), and probability distributions of hidden states at the end of one operation become initial state probabilities for the next operation. Based on these assumptions, a Genetic Algorithm based procedure is introduced in [4] to identify parameters of the ensemble of degradation-describing HMMs using observations (sensor readings) emitted by a system operating in arbitrarily mixed operating regimes.

Once the operating regime specific degradation HMMs are identified, conditional log-likelihoods of the newly arrived sensor observations, given the operating mode specific degradation HMMs, can be used to detect abnormalities. Namely, these log-likelihoods drop linearly with the length of an observation sequence, with slopes corresponding to the HMM dynamics. Normalizing these slopes by removing means of operating mode dependents slopes and scaling them with operating mode dependent variances, enables one to detect "unusual" slopes (i.e. observation sequences inconsistent with the degradation HMMs observed during normal system behavior) using simple Statistical Process Control methods, such as Exponentially Weighted Moving Average (EWMA) chart. More details about methods for identifying operating mode specific degradation HMMs, as well as various methods for anomaly detection based on such models of degradation can be found in [4].

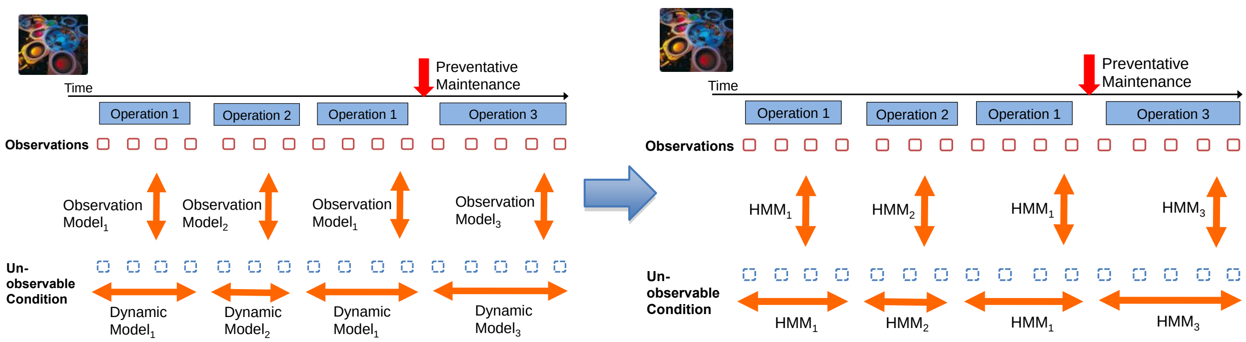

Fig. 3: Illustration of degradation modeling based on the concept of interconnected, operating-regime specific Hidden Markov Models (HMMs) of equipment degradation.

The HMM-based degradation modeling and monitoring methods were applied to the PECVD tool operating in a major domestic 300mm semiconductor manufacturing facility. A number of sensor signals were collected at a 10Hz sampling rate from the gas flow, pressure, radio-frequency (RF) power generation system, chamber pressure and digital sensors mounted on the tool over a period of more than 6 months, corresponding to more than 110,000 wafers (due to the proprietary nature of the data, the exact timeframes and wafer numbers must be omitted). During that period, the tool was depositing films on standard 300mm silicon wafers of the same chemistry but of 4 varying thicknesses. A set of features such as signal rise times, overshoots, time-durations and amplitudes of various events during the deposition process were extracted from the sensor signals[1].

The EWMA chart of normalized log-likelihoods of observations and the corresponding 4-σ control limits can be seen in Fig. 4. The dashed line labeled "Training" indicates the end of the data used to identify the parameters of the regime specific HMMs. A number of out-of-control events of interest are labeled and will be discussed in conjunction with the available maintenance and metrology data.

[1] For more details on the feature extraction procedure, the reader is referred to [11].

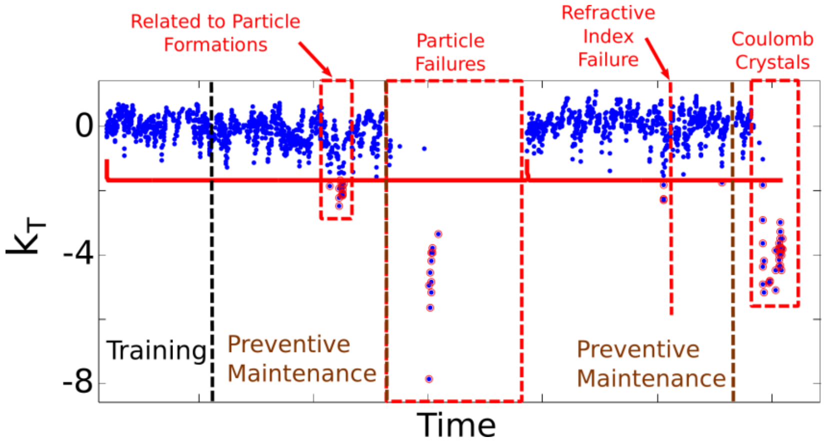

Fig. 4: EWMA control chart of normalized log-likelihood slopes k_T, with indications of the times of four abnormal behavior events observed during the monitoring period. *Please note that each dot in this figure corresponds to a batch of 25-100 wafers.*

Event 1: The first event is a number of points before the first preventive maintenance. During this period, particle-monitoring wafers[2] showed a significant increase in particle counts. These anomalous events lasted several days and were consistent with particle contamination from within the chamber itself, leading to a number of scrapped dies on the contaminated wafers.

Event 2: The out-of-control points in the box labeled "Particle Failures" in Fig. 4 correspond to the weeks-long tool downtime caused by defects seen on the particle monitoring wafers. The length of the downtime was so big simply because of the inability to find the source of particles and various subsystems of the tool were overhauled until the problem went away (suspicion is that an improper preventive maintenance event before those events led to all the problems).

Event 3: An out-of-control cluster of points immediately before a refractive index failure is reported. Process control adjustments were sufficient to mitigate the problem and consequently there was no prolonged downtime of the tool.

Event 4: During this event, a number of particle failures and "plasma formations" were noted in the chamber. Again, weeks of downtime ensued, with fruitless attempts to clean the chamber of the source(s) of particles. After consultation with experts from the company that made the tool, these events were found to be the result of improper evaporation of the deposition product, which caused a phenomenon known as Coulomb Crystals [12]. In other words, the source of the problem was not even close to the chamber – it was in the gas delivery system and that failure just led to symptoms that looked as if the chamber was bad.

3. Potential Approaches for Localizing the Sources of Problems in Complex Manufacturing Machines

Though the newly introduced degradation modeling approach enables anomaly detection even in partially observable, complex processes, the problem of automatically localizing the subsystems and components that led to the anomaly remains a challenge. HMM based anomaly detection could warn the manufacturer using the PECVD tool dealt with in the previous section, thus precluding the damage from wasted wafers. Nevertheless, the downtimes

[2] Non-production wafers used to "sense" the particle contamination in the chamber by essentially undergoing a deposition process, after which they are sent to a metrology tool that counts the number of particles on the wafer

could not be reduced because the manufacturer would still need to go through a trial and error procedure until the root cause of the problem was found.

In order to automate the root cause identification procedure using distributed anomaly detection, i.e. precedent free, as was done previously in the automotive applications, one needs to understand the causal interactions between various subsystems of the monitored system. Simply said, one needs to know what variables describe each subsystem and FRU of the PECVD tool examined in the previous section, as well as through what variables and in what ways these systems interact while the tool operates. In the case of the automotive EGR system, all relevant variables were adequately sensed and one knew exactly how various subsystems of the EGR system interacted (what were the inputs and what were the outputs of each subsystem and component). In the case of many types of modern manufacturing equipment, such causal topology of interactions between various subsystems (what affects what) may be much harder to obtain. Seemingly, everything could affect everything, which is especially true for highly complex and integrated systems, like PECVD or lithography tools. Based on intense interaction with the tool experts, Fig. 5 shows a graph of causal interactions for a PECVD tool. However, even those tool experts, who designed and manufactured the tool, are not sure if this graph really represents all its interactions. In addition, the consensus is that a lithography tool is even less understood.

Formal and systematic identification of causalities in the manufacturing equipment must be addressed to understand how FRUs interact with each other. Optimal model can perhaps be identified using some metaheuristic topological search, such as GA or Tabu Search, with model evaluations using Akaike Information Criterion [13] or Minimal Description Length criterion [14]. Such model discovery methods were attempted in hot rolling in the past [15][16], but never in anything even remotely as complex, as a semiconductor manufacturing tool, where the number of variables is several orders of magnitude larger, modes of operation are much more diverse and boundaries between (definitions of) FRUs and subsystems are much more blurred.

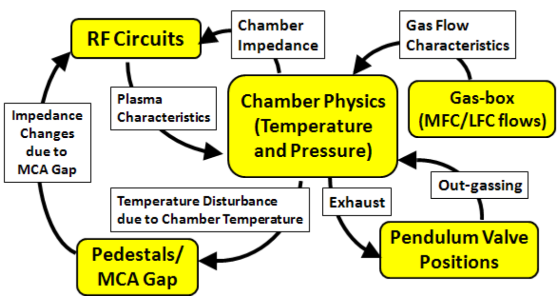

Fig. 5: Causal graph of PECVD tool FRUs, with brief explanations of interaction mechanisms.

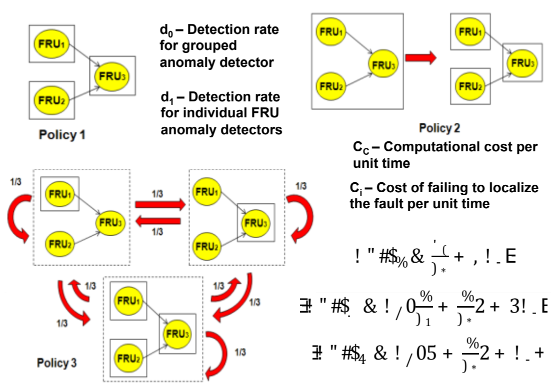

Fig. 6: Example of tradeoffs in agent distribution policies for distributed anomaly detection.

In addition, agent distribution policies that optimize tradeoffs between computational resources needed for each AD, their sensitivity to anomaly detection and speed of localization of the sources of anomalous behavior need to be explored. A simple example in Fig. 6 illustrates how increased utilization of computational resources can lead to increased fault sensitivity and speed of localization, and *vice versa*. A study like this in any realistic setting has never been conducted.

4. Humans as Distributed Agents for Removal of Faults ("Antigens") in a Manufacturing System

The precedent-free fault localization process described above resembles to a degree to the process in which a natural immune system labels an antigen by coating it with appropriately generated antibodies (the diagnostic system described here "coats" faulty subsystems and FRUs with alarming ADs). Once the natural immune system labels an antigen in such a way, while blood cells dispose of that intrusion by effectively killing anything that is coated with antibodies [17] (if the body mistakenly labels normal body cells, the person ends up having an autoimmune disorder, such as multiple sclerosis, or celiac disease). The job of "antigen removal" in a manufacturing system is performed by maintenance practitioners, who effectively act as leukocytes. However, unlike the natural leukocytes programmed to kill anything antibodies label as "non-self," manufacturing "leukocytes" (people) can think, learn and forget over time. There is a tremendous need for innovative methods for modeling and matching of dynamically evolving human skills with maintenance and operational jobs, including (especially) those corresponding to unprecedented faults.

In this talk, a tree-based representation of a machine faults and human skills, as illustrated in Fig. 7, is suggested as a potential solution avenue for skill matching based dispatching of practitioners, even in situations when the underlying situation is unprecedented. Furthermore, an optimization procedure is proposed for joint scheduling of operations and dispatching of operators, which takes into account the dynamics of evolving skills of the operators and interactions between various components of the manufacturing system.

This last portion of the talk obviously deals with a far less mature area of research and pertains to the possibility (need) to realize not just equipment diagnostics, but the entire realm of manufacturing system operations through the concept of distributed agents. Such distributed agents based architecture of autonomous process and operations control, cognizant and aware of "self" and "non-self" within it, may be dearly needed as the complexity of future manufacturing systems grows into the realm of something we can call *cognitive manufacturing*.

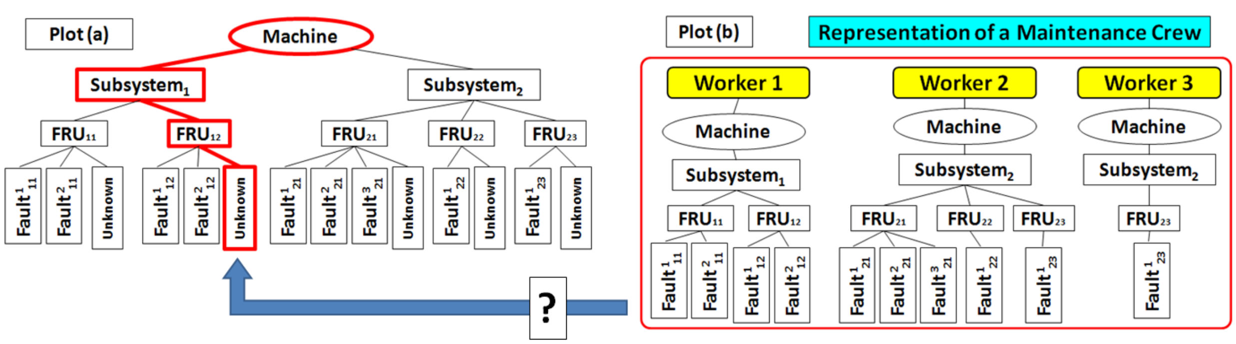

Fig. 7: Plot (a) illustrates the proposed tree-based representation of a machine and its subsystems and FRUs. It also shows a representation of a maintenance job as a portion of the machine tree "illuminated" by the diagnostic system (here, it points to an unknown fault in FRU_{12}). Plot (b) shows the tree-based representation of skills of a 3-member maintenance crew. Who should do this job?

REFERENCES

[1] Y. Lei, D. Djurdjanovic, L. Barajas, G. C. Workman, J. Ni and S. Biller, "Network Health Monitoring for DeviceNet Using Physical Layer Parameters", *J. of Intelligent Manufacturing*, Vol. 22, pp. 289-299, 2011.

[2] Y. Lei, D. Djurdjanovic and J. Ni, "DeviceNet Reliability Assessment Using Physical and Data Link Layer Parameters", *Quality and Reliability Engineering International*, Vol. 26, pp. 703-715, 2010.

[3] S. A. Spiewak, R. Duggirala and K. Barnett, "Predictive Monitoring and Control of the Cold Extrusion Process," *CIRP Annals of Manufacturing Technology*, Vol. 49, No. 1, pp. 383-386, 2000.

[4] M. E. Cholette. "Performance Monitoring and Fault-Tolerant Control of Complex Systems with Variable Operating Conditions," *Ph.D. dissertation, Mech. Eng., Univ. of Texas at Austin*, Austin, TX, 2012

[5] A. Heng, S. Zhang, A. C. C. Tan and J. Mathew, "Rotating Machinery Prognostics: State of the Art, Challenges and Opportunities," *Mechanical Systems and Signal Processing*, Vol. 23, pp. 724-739, 2009.

[6] D. Djurdjanovic, J. Liu, K. A. Marko and J. Ni, "Immune Systems Inspired Approach to Anomaly Detection, Fault Localization and Diagnosis in Automotive Engines," in *Applications of Neural Networks in High Assurance Systems* (edited by Y. Liu and J. Schumann), Springer, pp. 141-164, 2010.

[7] M. Cholette and D. Djurdjanovic, "Immunity Inspired Monitoring of Systems of Interacting Dynamic Systems," *Applied Intelligence*, vol. 37, no. 1, pp. 60-79, 2012.

[8] Cholette and D. Djurdjanovic, "Precedent-Free Fault Isolation in a Diesel Engine EGR System", *Transactions of ASME, J. of Dyn. Systems, Measurement and Control*, Vol. 134, No. 3, pp. 031007-1 – 031007-11, 2012.

[9] S. Forest, A. S. Perelson, L. Allen and R. Cherukuri, "Self-Nonself Discrimination in a Computer," *Proc. of the 1994 IEEE Symposium on Research in Security and Privacy*, Los Alamitos, CA, IEEE Computer Society Press, 1994.

[10] L. R. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," *Proceedings of the IEEE*, Vol. 77, No. 2, pp. 257– 286, 1989.

[11] A. Bleakie and D. Djurdjanovic. "Dynamic Feature Monitoring Technique Applied to Thin Film Deposition Processes in an Industrial PECVD Tool." *ASME 2011 International Manuacturing Science and Engineering Conference*, Corvallis, OR, 2011.

[12] J. H. Chu and I. Lin, "Direct observation of Coulomb crystals and liquids in strongly coupled rf dusty plasmas," *Physical Review Letters,*vol. 72, no. 25, pp. 4009-4012, June, 1994

[13] H. Akaike, "A new look at the statistical model identification," *IEEE Trans. on Automatic Control*, Vol. 19, No. 6, pp. 716–723, 1974.

[14] J. Rissanen, "Modeling by shortest data description," *Automatica*, Vol. 14, No. 5, pp. 465–471, 1978.

[15] J. Li, J. Jin and J. Shi, "Causation-Based T2 Decomposition for Multivariat Process Monitorin and Diagnosis," Journal of Quality Technology, Vol. 40, No. 1, pp. 46-58, 2008.

[16] J. Li and J. Jin, "Optimal Sensor Allocation by Integrating Causal Models and Set Covering Algorithms,"*IIE Trans. on Quality and Reliability Engineering*, Vol. 42, No. 8, pp. 564-576, 2010.

[17] NaturalImmuneSystems.com, "Immune System Facts for Everyday People," http://www.natural-immune-system.com/functions-of-the-immune-system.html, accessed July 20, 2013.