

On the Diagnosis of Cyber-Physical Production Systems: State-of-the-Art and Research Agenda

Oliver Niggemann and Volker Lohweg

inIT–Institute for Industrial IT, University of Applied Sciences OWL, 32657 Lemgo, Germany

email: {oliver.niggemann, volker.lohweg}@hs-owl.de

Abstract

Cyber-Physical Production Systems (CPPSs) are in the focus of research, industry and politics: By applying new IT and new computer science solutions, production systems will become more adaptable, more resource efficient and more user friendly. The analysis and diagnosis of such systems is a major part of this trend: Plants should detect automatically wear, faults and suboptimal configurations. This paper reflects the current state-of-the-art in diagnosis against the requirements of CPPSs, identifies three main gaps and gives application scenarios to outline first ideas for potential solutions to close these gaps.

1 Introduction

The diagnosis of distributed production systems has gained new attention due to research agendas such as Cyber-physical Production Systems (CPPSs) (Lee 2008; Rajkumar et al. 2010) or its German pendant “Industrie 4.0”. In these agendas, a major focus is on the self-diagnosis capabilities for complex and distributed CPPSs. Typical goals of such self-diagnosis approaches are the detection of anomalies, suboptimal energy consumptions, error causes in large plants or wear (Christiansen et al. 2011; Isermann 2004; Windmann et al. 2013).

All these approaches have in common that they (partially) capture the state of the environment and generate fused environment information. More technically, an environment model is generated based on several sources such as sensors, engineering models and experts. In many cases, the information captured from the environment may be imprecise, incomplete or inconsistent (Li and Lohweg 2008; Lohweg and Mönks 2010). Furthermore, signal sources may be not reliable (Ayyub and Klir 2006). Therefore, it is necessary to apply sensor fusion concepts as a first step before any diagnosis is computed (Mönks and Lohweg 2013). Usually the main problems in sensor fusion can be described as follows: Too much data, poor models, bad or too many features, and improperly analysed applications (Hall and Steinberg JAN 2001). One major misbelief is that machine diagnosis can be handled only based on the observed data, in reality

knowledge about the technical, physical, chemical or other processes are indispensable for modelling diagnosis systems (Glock et al. 2011).

The term “Diagnosis” summarizes all activities which identify anomalous behavior and compute the root causes of those anomalies. Often (Benjamins and Jansweijer 1994; Reiter 1987; Grastien 2013), the task of diagnosis is split into the four steps shown in figure 1: Based on the environ-




Figure 1: Typical steps of a diagnosis process.

ment model computed by means of sensor fusion, anomalous or conspicuous observations are identified. These observations are called symptoms—this second step is therefore often also called “symptom detection”.

In many cases (Peischl and Wotawa 2003; Gregson, Li, and Fu 2003; Maier, Tack, and Niggemann 2012), the process ends here and the user is informed about the symptoms. This is because in modern production plants, the main problem is that symptoms are overlooked due to a too high number of signals. Insofar, an anomaly detection process must create a low amount of warnings which a user can handle easily.

In a third step, hypotheses for the root causes are generated (Klar, Huhn, and Gruhser 2011; de Kleer et al. 2013; Pan et al. 2012). Normally, the root cause can not be identified unambiguously—often because we lack information, that is, not all intrinsically necessary knowledge is available (Epistemic Uncertainty).

Therefore, in the fourth step, heuristics or additional observations are used to discriminate between the hypotheses (Berjaga, Pallares, and Melendez 2009; Stern et al. 2013). As a result, a root cause is disclosed to the user.

Please note that this paper neither covers the topic of acquiring the observations for CPPSs (Pethig and Niggemann

2012) nor the generation of repair instructions.

Anomaly Detection

Generally speaking, as shown in figure 2, two classes of algorithmic approaches exist for the detection of anomalous situations:

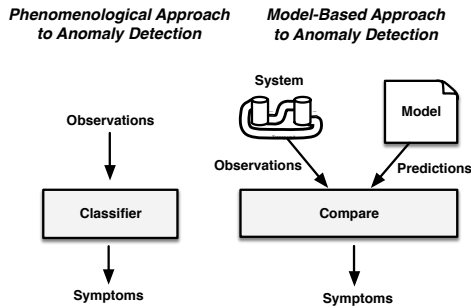


Figure 2: Different kinds of anomaly detection algorithms.

Phenomenological Approach: Here, the system observations are directly classified as correct or anomalous (Nieves et al. 2011; Ferracuti et al. 2011). Traditionally, the classification know-how is often modeled manually, e.g. in form of rules (expert systems). For fast-changing CPPSs, the classifier is trained using supervised machine learning algorithms (Matias et al. 2013; Goernitz et al. 2013).

Model-based Approach: In order to detect anomalies automatically, a model-based approach can be used (Struss and Ertl 2009; Christiansen et al. 2011; Niggemann et al. 2012). A model is used to simulate the normal behavior of a plant or normal product features. For this, the simulation model needs most inputs concerning the plant and the product, e.g. plant configuration, plant status, required product features, etc. If the system observations vary significantly from the model predictions, the system or product is classified as anomalous. Such models come in different flavors: Statistical (Ferracuti et al. 2011) and state-based models (Windmann et al. 2013) or physical first principle models (de Kleer et al. 2013). So model-based approaches capture the normal situation while phenomenological approaches capture the differences between normal and anomalous situations. The main challenge for this approach is high engineering efforts for the creation of such models.

While phenomenological approaches are often more straight-forward and do not require a system model, they have one major inherent drawback: They must deduce against the direction of causality since they deduce from observations to anomalies. For complex distributed systems with their high number of interdependencies between components and their complex causalities, this is a hard task because a high number of classification rules is needed to discriminate between all possible combinations of symptoms. Model-based approaches do not have this problem since system models take all inputs and compute the outputs, i.e. they work in the direction of the physical causality. So in general, phenomenological approaches are chosen for local compact devices while model-based approaches are chosen for complex, distributed plants.

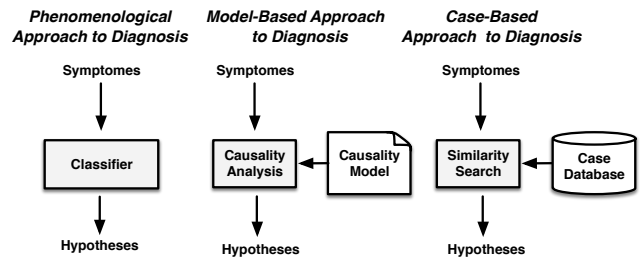


Figure 3: Different kinds of hypotheses generation algorithms.

So far, the majority of papers and projects for model-based anomaly detection (and also for model-based diagnosis as a whole) has relied on precise and component-granular models of the plant's physics (Isermann, Kimmich, and Schwarte 2004; de Kleer et al. 2013; Klar, Huhn, and Gruhser 2011; Mertens and Eppele 2009): If a drive is used, this drive is modeled, e.g. using characteristic maps. If a reactor is installed, the chemical and physical process is modeled, e.g. using differential equations. Of course, the modeling levels and modeling formalisms change, but in most cases they are chosen according to the system type—disregarding the aspect of model acquisition. Other approaches to a model-based anomaly detection use qualitative modeling formalisms such as automata (Su et al. 2002) or simplified physical models (Struss, Sachenbacher, and Dumortier 1997).

Hypotheses Generation

As shown in figure 3, there are in general three main approaches to hypotheses generation:

Phenomenological (heuristic) approaches use pre-defined associations to deduce from symptoms to root causes (Azarian, Siadat, and Martin 2011; Pan et al. 2012)). There are many techniques to represent such associations, ranging from thresholds up to machine learning algorithms. In general, it has the same drawbacks as outlines before for the phenomenological approach to anomaly detection.

Model-based approach: This approach uses a model of the causalities between the building components of the production plant (de Kleer et al. 2013; Grastien 2013; Klar, Huhn, and Gruhser 2011). By analyzing these causalities, candidates for the root cause are identified. Main drawbacks of this approach are due to the exponential number of causes and the difficulty of obtaining verified models.

Case-based approaches use a database of previous cases and identify hypotheses by exploiting similarities between the cases (Berjaga, Pallares, and Melendez 2009). This approach is still seldom used in CPPSs, mainly due to a lack of recorded cases and the unknown generalization capability of those cases.

Hypotheses Discrimination

To discriminate between hypotheses for the root causes, mainly heuristic approaches are used (Alippi, Ntalampiras, and Roveri 2013), e.g. a-priori probabilities. An exception are model-based approaches where a hypothesis is temporarily inserted into the model: if the differences between ob-

servations and model predictions disappear, the hypothesis is accepted. Otherwise, the hypothesis is removed from the model and another hypothesis is tested (de Kleer et al. 2013; Struss and Ertl 2009). For this, models must be able to simulate all relevant erroneous behaviors, in general a requirement difficult to implement in reality—and especially for CPPSs.

Sometimes, additional observations, e.g. by the user, are employed to differentiate between hypotheses (Azarian, Siadat, and Martin 2011).

2 Diagnosis for Cyber-physical Production Systems: Challenges

Above, a short overview of diagnosis has been given. In the following, we will outline specific features of CPPSs and we will then derive three main theses concerning future challenges in the field of diagnosis for CPPSs. Sections 3-6 will then underline the theses by case studies.

But first, to derive CPPS-specific challenges for diagnosis, CPPSs must be characterized. A CPPS is a holistic conception of modern, often distributed, production systems: It treats mechanical, computational and external aspects (e.g. users, markets) as a inextricable system whose complexity can only be handled if the CPPS comes with a set of intrinsic cognitive capabilities such as self-diagnosis, self-configuration, self-optimization and intelligent user interaction.

Such a CPPS is characterized by a set of features which differentiate it from classical production plants:

Hybrid Systems: CPPSs are hybrid systems which comprise discrete signals, time- and value-continuous signals and structured data. Often, discrete signals such as opening a valve or turning off a robot trigger mode changes (Buede 2009), i.e. they abruptly change the system behavior. Continuous signals comprise e.g. energy consumptions and other resource requirements. Structured data comprises product and raw material information, Enterprise Resource Management (ERP) data and plant configurations.

Timed Systems: CPPSs are distributed systems which show complex causalities and therefore a complex timing behavior. I.e. unlike local devices, sophisticated behavior changes over time render such systems complex and difficult to handle. E.g. a wear in a drive or the blockage of a pipe in some part of the system will slow down some transportation mechanisms and may cause behavior changes in several other parts of the system—but only gradually and after a period of time.

Distributed Systems: CPPSs are distributed systems of systems: They encompass sensors/actuators, controller, business software, mechanical components and of course the users. And plants are often modularized into separate subsystems which use heterogeneous automation systems. Furthermore, production and logistic processes comprise locations all over the world which use different technologies and employ people with different educational background. Only by assigning tasks, now solved manually by experts, to cognitive systems in the CPPSs, these systems will remain manageable.

What follows from these CPPS features for the task of diagnosis? In the following, three theses will outline the main differences between the state-of-the-art and a future diagnosis for CPPSs.

Thesis 1: Anomaly detection will use generic physical models which will be learned

As explained in section 1, phenomenological approaches to anomaly detection are not suitable for system-level diagnosis tasks: Since they deduce from observations to anomalies and root causes, they have to contain rules to differentiate between all combinations of signals—and especially between all observation values over time. I.e. with a growing number of observations and a growing relevant time interval the number of necessary rules—which must be modeled or learned—grows also exponentially. This restricts the application of phenomenological approaches to small local devices such as single drives, reactors, etc.

This leaves model-based anomaly detection. So far, as outlined in section 1, such approaches have relied on physical models, with a different level of details, of the system behavior and of its components, models which must be created during the engineering process manually—often using object-oriented approaches to ease the modelling task. But the last 20 years have clearly shown that such models hardly ever exist and prevent the usage of model-based algorithms. So CPPSs will use more generic modeling formalisms which can be learned automatically based on observations—learning here refers to the complete learning of models (Niggemann et al. 2012) and not to the parameterization of manually created models (Isermann 2004; Zhang 2012). The reader may note that these generic modelling formalisms are still physical models, since they work in the direction of causality—unlike phenomenological solutions. But instead of using model formalisms which require high engineering efforts (e.g. a conveyor belt modeled using differential equations), in the future learnable generic formalisms will be used which require no manual efforts (e.g. timed automata for all components). An example, using timed automata, can be found in sections 3 and 5.

Thesis 2: Anomaly detection will use hybrid models which also capture the system's timing

Analyzing section 1, it becomes obvious that most model-based anomaly detection solutions either work with discrete models (e.g. propositional logics, automata) or with a static view onto the system. This does not match with the dynamic, time-based and hybrid nature of CPPS. So hybrid models, which also capture the dynamic behavior of time, will be used. Details are given in section 4

Thesis 3: Instead of root cause identification, the focus will be on a component-level anomaly detection and on predictive maintenance

Identifying root causes requires models of the causality between system components. So far, these models can not be learned and must be created manually. As outlined before, such manual modeling tasks form bottlenecks in the process. Furthermore, since phenomenological approach suffer from the problems shown above, this leaves case-based root cause analysis. But section 1 shows that such approaches are still

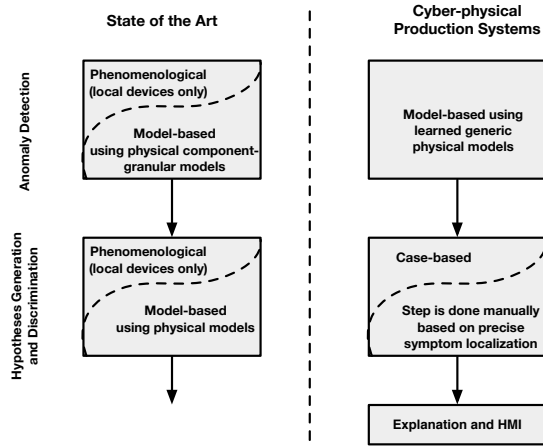


Figure 4: Comparison between the state-of-the-art in diagnosis and the diagnosis of CPPSs.

the exception—mainly because cases are hard to find and because cases can often not be generalized to new situations.

But in most cases, the root cause analysis is not the main problem in CPPSs. Instead, anomalies must be localized as precise as possible. E.g. if an anomalous energy consumption of a production line can be traced back to one or two specific drives, the root cause is easily identified. In CPPSs, users are able to identify root causes easily if symptoms are computed on such a component-level. To a large extend, the challenge for users lies in the automatic early detection of such anomalies, currently users are not able to do this due to the over-whelming number of signals. Details can be found in section 6.

These differences can also be seen in figure 4: The left hand side shows the state of the art in diagnosis, the right hand the diagnosis approach for CPPS.) After sections 3-6

have shown details on these theses, section 7 will derive a corresponding research agenda.

3 Analyzing discrete Systems and their Timing Behavior by Means of Model Learning

In contrast to the timing behavior of continuous systems, the modeling and learning of discrete timing behavior is a less mature field—especially if we want to learn such models. Significant work mainly has been done in using the formalisms of Timed Automata and Timed Petri Nets. Choosing an appropriate timing modeling formalism some key issues have to be considered, such as (i) discrete or dense time domain, (ii) explicit or implicit modeling of time, (iii) one clock or many clocks and (iv) concurrency and composition.

Due to the intuitive interpretation, Timed Automata are well-suited to model the timing behavior of CPPS. The timing information can be represented using time ranges or preferably using probability density distribution functions (PDF) over time.




Figure 5: Learned Timed Automata for a manufacturing plant.

Another key issue is the ability for automatic learnability. Several algorithms have been introduced to learn a Timed Automaton based on observations of the normal behavior only. RTI+ (Verwer 2010) and BUTLA (Niggemann et al. 2012) learn in an offline manner, i.e. first the data is acquired and stored and then the automaton is learned. However, using the algorithms in the context of CPPS, observations cannot be stored and therefore, an online learning algorithm is desirable. OTALA (Maier 2014) is an extension of BUTLA and learns a Timed Automaton in an online manner.

The corresponding anomaly detection (e.g. algorithm ANODA (Vodenčarević et al. 2011)) works according to figure 2: Any difference of the observed behavior to the predictions of the learned model hints at an anomaly. Here, the sequential behavior of the observed events as well as the timing behavior have to be checked.

Figure 5 shows learned automata for a demonstration manufacturing plant: The models correspond to modules of the plants, transitions are triggered by a control signals and are annotated with a learned timing interval.

In another project using a plant from process industry, the anomaly detection with the learned Timed Automata has been evaluated on data of a real production plant and compared with the results of the models learned by neural networks and decision tree learning.

The results are given in the confusion matrix (according to (Tan, Steinbach, and Kumar 2005)) in table 1.

Table 1: Results for the experiment on production plants.

	true positive	true negative	Accuracy
Timed Automata	100%	97%	98,5%
Neural Networks	40.2%	42.2%	41.2%
Decision Trees	90.5%	83.3%	86.9%

Using learned Timed Automata, the anomaly detection algorithms were able to detect all anomalies. The number of false positives is also a good value, only some outliers (mainly timing deviations) have not been learned. In total, a

high accuracy of 98.5% is achieved. In contrast the neural networks and decision trees achieved an accuracy of only 41% and 87% respectively.

Relation to Theses 1 & 2 (section 2): Here, models of plants from both manufacturing and process industry are learned automatically (Thesis 1). For this, a Timed Automaton is used which is a generic formalism to model timing and which is also still learnable (Thesis 2).

4 Energy Analysis

Analyzing energy data in industrial system has some special challenges: The energy consumption must be analyzed with respect to the current system's status, e.g. whether a valve is open or whether a robot is turned on. The system's status is usually defined by the history of discrete control signals. I.e. the resource consumption of plants must be modeled using hybrid model formalisms.

In (Kroll et al. 2014), an energy anomaly detection system is described which analyzes two demonstration plants and one industry plant. For this, the algorithms for learning Timed Automata are extended to learn hybrid timed automata. I.e. in addition to the analysis of events and discrete signals, also continuous data such as energy consumptions was used for anomaly detection. In figure 6, a pump is modeled by means of hybrid timed automata using the flow rate and switching signals. The three states S_0 to S_2 are separating the continuous function into three linear pieces which could be learned automated by the described algorithms.




Figure 6: A learned hybrid automaton modeling a pump.

The results show, that a model accuracy of over 80% is possible, with signal variances of 2,5%. The results shown in table 2 give an overview on each tested platform.

Table 2: Result overview

Platform	Model accuracy
Demonstration Plant	~ 100%, var > 2,5%
Intralogistic Test Platform	~ 80%
Chemical Test Platform	~ 90%

Figure 7 shows a typical learned energy consumption.

Relation to Theses 1 & 2 (section 2): Again, as already outlined in section 3, generic behavior models are learned based

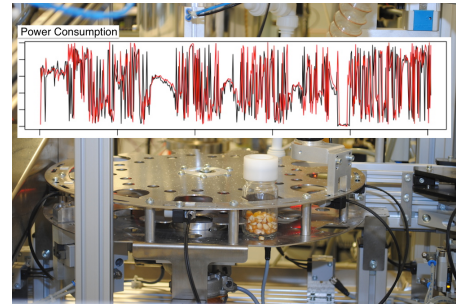


Figure 7: A measured (black line) and a learned power consumption (red line).

on system observations. But as outlined in thesis 2, here hybrid models, typical for CPPSSs, are first learned and then used for the system analysis task.

5 Model-Based Anomaly Detection

A possible use case for model-based anomaly detection is wind energy, where minimization of maintenance and overall downtimes improves energy yield of wind power stations (WPS). By learning the behavior of a WPS from historical data, an anomaly detection system is able to notify an operator in case of problems or failures. A joint research project with a term of 1.5 years aims at such an anomaly detection system where the used models are learned using unsupervised learning algorithms.

Historical data from a WPS forms the database for this project. It contains data of a normal WPS operation over a couple of months. A principal component analysis (PCA) was then used to learn a model of the system's behavior. As input for the PCA system descriptive continuous signals were used. Wind speed and temperature were also taken in account from the database. The PCA returns a dimensional reduced description of the WPS normal state of work.

For evaluation purposes, new data is presented and transformed into the learned PCA-space. In a next step, the distance of this transformed new data to the normal space is calculated. Using mexican hat density functions, a membership probability for new data is computed. Figure 8 shows a learned model: The green nodes denote correct system situations while red notes correspond to anomalous situations.




Figure 8: Learned model and anomalies.

Table 3: Evaluation results of wind power station data.

total data	evaluation	predicted anomaly	predicted normal work state
40257			
real anomaly		1165 (true positiv)	286 (false positiv)
real normal work state		38794 (true negativ)	12 (false negativ)

To measure the performance, a database of 44568 samples was used. 43117 data points describe the normal state of work, where the other 1451 are errors. The training data contains 10% of the OK-data, the remaining 90% are used as evaluation data. Table 3 shows the confusion matrix (see (Tan, Steinbach, and Kumar 2005)) as a result of the evaluation. For this use case, the F1-score is used to analyze the system's performance. A sensitivity of 80.29%, specificity of 99.96%, precision of 98.98%, balanced accuracy of 90.13% and a F1-score of 88.66% were achieved by using the above described method. So with a promising accuracy of 90.13% and a F1-score of 88.66% this method can monitor a continuous system like a wind power station in order to explicitly plan maintenance to reduce downtimes.

Relation to Theses 1 & 3 (section 2): Due to the high variety of WPS, any solution requiring manual modeling efforts is unrealistic. Instead, here a model is learned automatically, e.g. this approach is suitable to adaptable CPPSs. And, as outlined in thesis 3, the experts are satisfied with an anomaly detection since they can normally easily identify the root cause once the symptoms are known.

6 Computing Component-based Anomalies

The algorithm *MoSDA* (Niggemann et al. 2014) exploits a situation typical for CPPSs: Anomalies often refer to anomalous observations on the system-level, e.g., an anomalous energy consumption of the production system as a whole. *MoSDA* leverages on the fact that these observations create an over-determined system and *MoSDA* can therefore break down—heuristically—the system-related anomalies to component-related anomalies. The reader may note that continuous signals in plants are often only measured on the system-level only because sensors for power consumption, water consumption, output, fluidic pressure can not be installed for each component. Starting from such component-related anomalies, heuristics can be easily developed that compute a set of possible root causes.

MoSDA has been applied, among others, to a high storage system. The system comprises real world components such as drives, conveyer belts and automation devices. Typically, wear and faults occur mainly in the drives, causing too high energy consumptions, too long transport durations and finally plant downtimes. To prevent this, any wear in a drive must be detected as early as possible.

The overall power consumption, observed for the whole plant, of the high storage system is split by *MoSDA* into the individual power consumptions of the 7 drives. The results of the *MoSDA* algorithm are shown in figure 9: The first row of figure 9 shows the 7 accelerations, the second row the 7 velocities and the last row shows the measured overall power

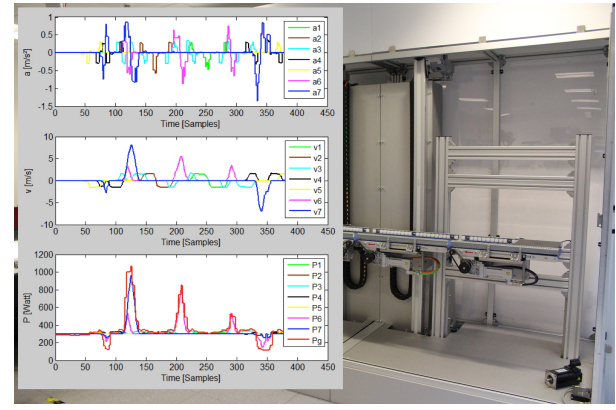


Figure 9: The results of the *MoSDA* algorithm.

consumption P_g and the 7 separated power consumptions as computed by the *MoSDA* algorithm. In the research high storage system (unlike in real plants), the separated power consumptions can be compared to the real power consumptions: 10 cycles have been measured, the average prediction error was between 1.7% and 7.4%. I.e. the separate power consumptions can be learned effectively. Therefore, in this example, deviations of more than 7.4% from the normal behavior of the single drives can be used to detect wear and faults.

Relation to Theses 3 (section 2): *MoSDA* does not implement a root cause analysis—root cause analysis being the holy grail of traditional diagnosis. Instead, symptoms and anomalies are computed for individual components. The reader may note the difference to the state of the art: Currently, sensors such as energy sensors capture not the behavior of individual components but of whole subsystems—mainly due to costly and elaborate installation and maintenance processes for sensors. And since sensors correspond to subsystems, symptoms are computed on the subsystem level.

So breaking down the symptoms to the level of individual components is a significant improvement: In modern production plants, users are overwhelmed by the sheer number of signals. If anomaly detection systems can focus the users attention onto unusual behaviors, and if they can pinpoint the anomalous components, the main problems are solved. E.g. if among 1000 drives—no unusual number—4 are marked as anomalous, most havoc is avoided.

7 Conclusion and Research Agenda

Based on the structure for CPPS diagnosis developed in section 1, section 2 phrases three main theses: 1. We need learnable generic formalisms for behavior models. 2. We need model formalisms which capture both hybrid and timing aspects. 3. Instead of focusing on root cause analysis, we should work on component-level anomaly detection algorithms. Sections 3-6 then outlined the theses using several application scenarios.

From the three theses follows directly a corresponding research agenda:

First of all, data-driven approaches to modeling must move into the research focus, replacing manual engineering approaches. I.e. the fields of machine learning and diagnosis must work more closely together. As outlined in this paper, this requires different formalisms which may capture less system aspects than complex manually engineered models but are learnable, i.e. we sacrifice precision for learnability.

Second, research should focus on formalisms which support hybrid aspects and which model a systems' timing explicitly. Hybrid aspects, today often neglected, must move into the focus for CPPSs. Furthermore, time must be modeled explicitly in order to support a better analysis of a plant's timing behavior—timing being crucial for distributed systems such as CPPSs.

Last, diagnosis solutions must take the cognitive capabilities of human experts into consideration: How do experts identify anomalies and root causes and where do they need assistance? For CPPSs, the traditional challenge of root cause analysis must be complemented by an identification of symptoms on the level of single components.

In general, diagnosis challenges for CPPSs must be defined more from an application point of view, i.e. the challenges lie in an efficient support of the human expert in the plant.

8 Acknowledgments

We would like to thank our colleagues A. Maier, S. Windmann, J. Jasperneite, B. Kroll, J. Eickmeyer and S. Volgmann for their work in the field of CPPSs and also for valuable discussions. We also thank the German ministry BMBF for funding the projects AGATA, itsowl-EE and AVA and the German ministry BMWi for funding the projects PrognosBrain and Anubis.

References

- Alippi, C.; Ntalampiras, S.; and Roveri, M. 2013. A cognitive fault diagnosis system for distributed sensor networks. *Neural Networks and Learning Systems, IEEE Transactions on* 24(8):1213–1226.
- Ayyub, B. M., and Klir, G. J. 2006. *Uncertainty modeling and analysis in engineering and the sciences*. Boca Raton and FL: Chapman & Hall/CRC.
- Azarian, A.; Siadat, A.; and Martin, P. 2011. Optimization of a knowledge-based system by a meta-heuristic approach for the automotive diagnosis. In *Industrial Engineering and Engineering Management (IEEM), 2011 IEEE International Conference on*.
- Benjamins, R., and Jansweijer, W. 1994. Toward a competence theory of diagnosis. *IEEE Expert* 9(5):43–52.
- Berjaga, X.; Pallares, A.; and Melendez, J. 2009. A framework for case-based diagnosis of batch processes in the principal components space. In *Emerging Technologies Factory Automation, 2009. ETFA 2009. IEEE Conference on*, 1–9.
- Buede, D. 2009. *The Engineering Design of Systems: Models and Methods*. John Wiley & Sons.
- Christiansen, L.; Fay, A.; Opgenoorth, B.; and Neidig, J. 2011. Improved diagnosis by combining structural and process knowledge. In *Emerging Technologies Factory Automation (ETFA), 2011 IEEE 16th Conference on*.
- de Kleer, J.; Janssen, B.; Bobrow, D. G.; Saha, T. K. B.; Moore, N. R.; and Sutharshana, S. 2013. Fault augmented modelica models. *The 24th International Workshop on Principles of Diagnosis* 71–78.
- Ferracuti, F.; Giantomassi, A.; Longhi, S.; and Bergantino, N. 2011. Multi-scale pca based fault diagnosis on a paper mill plant. In *Emerging Technologies Factory Automation (ETFA), 2011 IEEE 16th Conference on*, 1–8.
- Glock, S.; Voth, K.; Schaede, J.; and Lohweg, V. 2011. A Framework for Possibilistic Multi-source Data Fusion with Monitoring of Sensor Reliability. In Yager, R. R., and Reformat, M. Z., eds., *World Conference on Soft Computing*.
- Goernitz, N.; Kloft, M. M.; Rieck, K.; and Brefeld, U. 2013. Toward supervised anomaly detection. In *Journal Of Artificial Intelligence Research*, volume 46, 235–262.
- Grastien, A. 2013. A spectrum of diagnosis approaches. *The 24th International Workshop on Principles of Diagnosis* 130–135.
- Gregson, D.; Li, K.; and Fu, W. 2003. Fault detection using phenomenological models. In *Electrical and Computer Engineering, 2003. IEEE CCECE 2003. Canadian Conference on*, volume 2, 855–859 vol.2.
- Hall, D. L., and Steinberg, A. JAN 2001. *Dirty Secrets in Multisensor Data Fusion*. Ft. Belvoir: Defense Technical Information Center.
- Isermann, R.; Kimmich, F.; and Schwarte, A. 2004. Model Based Fault Detection of Diesel Engines. ETAS Conference.
- Isermann, R. 2004. Model-based fault detection and diagnosis - status and applications. In *16th IFAC Symposium on Automatic Control in Aerospace*.
- Klar, D.; Huhn, M.; and Gruhser, J. 2011. Symptom propagation and transformation analysis: A pragmatic model for system-level diagnosis of large automation systems. In *Emerging Technologies Factory Automation (ETFA), 2011 IEEE 16th Conference on*, 1–9.
- Kroll, B.; Schaffranek, D.; Schriegel, S.; and Niggemann, O. 2014. System modeling based on machine learning for anomaly detection and predictive maintenance in industrial plants. In *19th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA)*.
- Lee, E. 2008. Cyber physical systems: Design challenges. In *Object Oriented Real-Time Distributed Computing (ISORC), 2008 11th IEEE International Symposium on*, 363–369.
- Li, R., and Lohweg, V. 2008. A Novel Data Fusion Approach using Two-Layer Conflict Solving. In *International Workshop on Cognitive Information Processing (CIP 2008)*, 132–136. IEEE.
- Lohweg, V., and Mönks, U. 2010. Sensor Fusion by Two-layer Conflict Solving. In *2nd International Workshop on Cognitive Information Processing (CIP 2010)*, 370–375. IEEE.
- Maier, A.; Tack, T.; and Niggemann, O. 2012. Visual anomaly detection in production plants. In *9th International Conference on Informatics in Control, Automation and Robotics (ICINCO)*. Rome, Italy, Jul 2012.
- Maier, A. 2014. Online passive learning of timed automata for cyber-physical production systems. In *The 12th IEEE Inter-*

- national Conference on Industrial Informatics (INDIN 2014)*. Porto Alegre, Brazil.
- Matias, T.; Gabriel, D.; Souza, F.; Araujo, R.; and Costa Pereira, J. 2013. Fault detection and replacement of a temperature sensor in a cement rotary kiln. In *Emerging Technologies Factory Automation (ETFA), 2013 IEEE 18th Conference on*, 1–8.
- Mertens, M., and Epple, U. 2009. Plant asset management functions driven by property models. In *Emerging Technologies Factory Automation, 2009. ETFA 2009. IEEE Conference on*, 1–8.
- Mönks, U., and Lohweg, V. 2013. Machine Conditioning by Importance Controlled Information Fusion. In *18th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA 2013)*, 1–8.
- Nieves, J.; Santos, I.; Ugarte-Pedrero, X.; and Bringas, P. 2011. Anomaly detection for high precision foundries. In *Industrial Informatics (INDIN), 2011 9th IEEE International Conference on*, 169–174. IEEE.
- Niggemann, O.; Stein, B.; Vodenčarević, A.; Maier, A.; and Kleine Büning, H. 2012. Learning behavior models for hybrid timed systems. In *Twenty-Sixth Conference on Artificial Intelligence (AAAI-12)*, 1083–1090.
- Niggemann, O.; Windmann, S.; Volgmann, S.; Bunte, A.; and Stein, B. 2014. Using learned models for the root cause analysis of cyber-physical production systems. In *International Workshop on the Principles of Diagnosis (DX)*. Graz, Austria.
- Pan, Y.; Loi, Q.; Ren, Z.; Dong, L.; and Zhang, X. 2012. Test point optimization for model-based fault diagnosis expert system. In *Automatic Control and Artificial Intelligence (ACAI 2012), International Conference on*.
- Peischl, B., and Wotawa, F. 2003. Model-based diagnosis or reasoning from first principles. *IEEE Intelligent Systems* 18(3):32–37.
- Pethig, F., and Niggemann, O. 2012. A process data acquisition architecture for distributed industrial networks. In *Embedded World Conference*.
- Rajkumar, R. R.; Lee, I.; Sha, L.; and Stankovic, J. 2010. Cyber-physical systems: The next computing revolution. In *Proceedings of the 47th Design Automation Conference, DAC '10*, 731–736. New York, NY, USA: ACM.
- Reiter, R. 1987. A theory of diagnosis from first principles. *Artif. Intell.* 32 1:57–95.
- Stern, R.; Kalech, M.; Feldman, A.; Rogov, S.; and Zamir, T. 2013. Finding all diagnoses is redundant. *The 24th International Workshop on Principles of Diagnosis* 216–221.
- Struss, P., and Ertl, B. 2009. Diagnosis of bottling plants - first success and challenges. In *20th International Workshop on Principles of Diagnosis, Stockholm*.
- Struss, P.; Sachenbacher, M.; and Dummert, F. 1997. Diagnosing a Dynamic System with (almost) no Observations. In *Working Papers of the 11th International Workshop on Qualitative Reasoning (QR-97)*.
- Su, R.; Wonham, W.; Kurien, J.; and Koutsoukos, X. 2002. Distributed Diagnosis for qualitative Systems. In Silva, M.; Giua, A.; and Colom, J., eds., *Proceedings of the 6th International Workshop on Discrete Event Systems, WODES-02*, 169–174.
- Tan, P.-N.; Steinbach, M.; and Kumar, V. 2005. *Introduction to Data Mining, (First Edition)*. Boston, MA, USA: Addison-Wesley Longman Publishing Co., Inc.
- Verwer, S. 2010. *Efficient Identification of Timed Automata: Theory and Practice*. Ph.D. Dissertation, Delft University of Technology.
- Vodenčarević, A.; Büning, H. K.; Niggemann, O.; and Maier, A. 2011. Using behavior models for anomaly detection in hybrid systems. In *Proceedings of the 23rd International Symposium on Information, Communication and Automation Technologies-ICAT 2011*.
- Windmann, S.; Jiao, S.; Niggemann, O.; and Borchertding, H. 2013. A stochastic method for the detection of anomalous energy consumption in hybrid industrial systems. In *INDIN*.
- Zhang, D. Y. H. H. G. Z. X. 2012. Anomaly detection for continuous sequence based compression process. In *Software Engineering and Service Science (ICSESS), 2012 IEEE 3rd International Conference on*, 364 – 367.