Appeared in: Frisk, E. et al. (eds.), 20th International Workshop on Principles of Diagnosis, Stockholm, 2009

Diagnosis of Bottling Plants – First Success and Challenges

Peter Struss, Benjamin Ertl

Comp. Sci. Dept., Technische Universität München, Garching, Germany
(e-mail: struss@in.tum.de, ertl@in.tum.de)

Abstract: The paper describes an application of component-oriented consistency-based diagnosis to the domain of bottle-filling plants. The task is to localize the causes for stops of the central aggregate, the filler, based on recorded operation data of a plant. A model-based solution is challenging in several respects, especially due to high uncertainty in the transportation processes to be modeled, the nature of the available data, and the relevance of numerical temporal information. We give a short description of the application and its requirements and summarize essential characteristics of the solution. We focus on the evaluation of the first demonstrator and a discussion of some challenges for future work, which include questioning the "classical" notion of a fault in component-oriented diagnosis.

1. INTRODUCTION

Most of the academic work in model-based diagnosis ([Struss 08]) is based on a number of simplifying assumptions and restrictions, including, as the most fundamental ones, that the system to be diagnosed is composed of a fixed set of components in a fixed structure and that the reasons for a misbehavior are faults in individual components. Besides this, some of the most popular ones are

- The behavior of the components is deterministic and can easily be modeled.
- The available observations are reliable and accurate.
- Time can be ignored (because the system is static or in a steady state), or evolution over time can be adequately described using an ordering on discrete time steps.

This is why digital circuits establish a highly popular "application" domain for this kind of research. While such a divide-and-conquer tactics regarding the problem dimensions of diagnosis can be justifiable for research purposes, it becomes an ostrich-like tactics when trying to solve diagnosis problems in the real world.

In our project on fault localization in bottling plants[1], we found the abovementioned and some more assumptions violated. In such plants, objects, especially bottles, are transported and processed in ways that are not precisely predictable. Data are sparse, represent mainly indirect, interpreted observation of behavior, and are inaccurate, often unreliable and even erroneous. The objects, and also effects of disturbances, such as tailbacks due to a blockage, are propagated with delays of unrestricted and uncertain temporal extent.

We addressed some of these challenges by exploiting qualitative behavior models stated at a high level of abstraction, as described in [Struss et al. 08a]. The architecture and principles of the diagnostic solution were outlined in [Struss et al. 08b]. The first demonstrator aims at solving a major class of diagnostic problems in this application, namely localization of hard machine failures causing an interruption of the filling process. In this paper, we present more details about this solution and the results of its evaluation on data from three plants. Besides the success part, which already triggered activities in commercial exploitation of the project results, we present some of the problems and requirements that will be addressed in a follow-on project. Some of them challenge quite fundamental concepts in model-based diagnosis, including even the notion of a component fault!

We give a brief characterization of the application task and its requirements in the following section and then describe the first demonstrator solution. Section 4 summarizes the evaluation results. Finally, we discuss those challenges for future work that we consider of general interest.

2. DISTURBANCES IN BOTTLING PLANTS

A bottling plant for filling beverages into returnable bottles is an assembly of a number of quite different types of specialized machines and conveyors that automatically handles the complete process from supplied pallets with crates containing empty bottles to the final output of pallets with (cleaned) crates and filled and labeled bottles. The plants can be large, distributed over several big halls and have a complex three-dimensional layout, as illustrated by Figure 1.

From an abstract point of view, which reflects flows and manipulation of different types of objects, the basic schematic topology can be simplified as indicated in Figure 2: there are lines for primary packaging (beverages into bottles), the top section in Figure 2, secondary packaging (into crates), shown as the middle section, and tertiary packaging (e. g. pallets), all organized as an automated branching, but directed flow. Certain backward loops, such

[1] The LineMod project was partially funded by the German Federal Minister of Economics and Technology (BMWi) through the German Federation of Industrial Research Associations (AiF) under FV-Nr.233 ZBG

Figure 1. *Conveyors of a bottling plant for returnables*

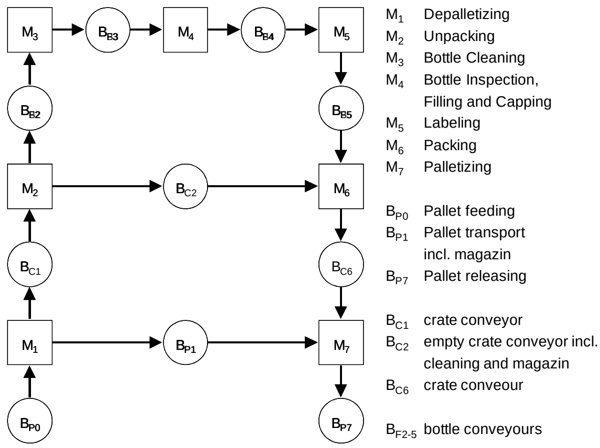

Figure 2. *Generic structure of a bottling plant for returnable bottles*

as re-submission of improperly cleaned or filled bottles to conveyors and the central filling aggregate, there are machines for de-palletizing and unpacking of returnable bottles, cleaning, inspection, and sorting out of improper objects.

To prevent oxygen intake or microbiological contamination of the beverage, a major objective is to prevent interruptions of the filling process. Besides for internal reasons, the filling aggregate will stop operation if there is a lack of input, i.e. bottles (apart from the beverage itself and caps), or a tailback of filled bottles preventing further output, i.e. disturbances caused by other machines downstream or upstream, resp. Because of the high speeds and output rates (up to 100.000 packages per hour), machines and conveyors are failure-sensitive with an availability degree of 92-98 percent. In order to avoid that each disturbance of single machines in the line results in a filler stop, several conveyor belts are designed as buffers, which can provide a continuous supply and output of other machines and, in particular, the filler, in conjunction with a general operation principle: machines and conveyors upstream and downstream from the filler operate at higher throughput rates than the filler. This principle is usually the only global one; there is **no global control**, and the machines are controlled individually (or, sometimes, as small aggregates).

However, in practice, these measures cannot guarantee avoidance of unwanted idle time of the filler, and (unplanned) downtime of the plant can lie in the range of 10-30 percent. Taking steps to reduce downtime by identifying frequent causes requires statistics and an analysis based on the recorded operating data supplied by (some of) the machines. Because of the interlaced flows of the various object types, time offsets, the large scale of the plants, and the amount and often fragmentary nature of the data this can be difficult and time-consuming. In consequence, bottle filling and packaging industries is highly interested in an automated diagnosis tool for their plants that produces information that helps to identify bottlenecks and weaknesses in the plant, related to both the physical performance and configuration and the control principles and parameters. Providing such a tool is the goal of the LineMod project described here and its follow-on project.

previous steps, are omitted in the figure. Besides the Because many of the potential end users, e.g. breweries, are small or medium enterprises, which could not afford spending many resources on the establishment or adaptation of a tailored diagnosis system for their plant, a **model-based solution** to diagnosis was chosen, which allows performing adaptation to a plant or modification simply by (re-)specifying the plant structure. This also provides a flexible solution that derives the best diagnosis from **whatever data is available** (in contrast, for instance, to decision trees based on a fixed set of observables), which is important, because usually a plant is a combination of machines from various manufacturers with different instrumentation and available data and there may be temporarily missing data due to technical problems.

Heterogeneity and changes of the set of machines also establishes a requirement on the model: firstly, it has to be machine-centered and **compositional**; secondly, it has to be stated at a level of abstraction that covers **types of machines**, independently of specificities and the manufacturer.

3. POST-MORTEM DIAGNOSIS OF BOTTLING PLANTS BASED – A MODEL-BASED SOLUTION

The first demonstrator of the tool addresses the needs explained above by localizing those interruptions of transportation that caused downtime of the filler based on the available recorded data of the machines (over a period of days to months) collected in a data base. In this section, we discuss aspects of this data base, present the architecture and the modules of the analysis system, and summarize the modeling approach.

3.1 The Available Data

As we pointed out before, a plant usually comprises machines from different suppliers, all coming with their own control system and specific data. Since it is not feasible for end users to generate a homogeneous set of data from theses different sources, there has been an effort undertaken by one of the project partners to establish a standard for data for production data acquisition (PDA) in bottling plants, the so-called "Weihenstephaner Standard 2005" (WS2005, [Kather et al. 05]), which has meanwhile been widely accepted. This was an essential pre-requisite for a successful solution, and one of the project's objectives was an extension of this standard for diagnosis purposes.

The extended WS2005 standard includes data points like:

- Operating states of machines according to WS2005
- Maximal output rate setting of machines
- Object counters (input, output, or rejection)
- Information about mechanical barriers
- Assignment of jam switches or buffer filling degrees, if available
- Transportation or machine speeds, if available.

The standard specifies an optimal set of data, and in practice, many machines provide only a small subset or nothing at all, different suppliers may interpret and handle the data differently (e.g. counters in a cumulative or incremental way), and even erroneous data may be present.

3.2 Architecture of the Solution

The data base is large and will contain reflections of many disturbances and machine stops that are irrelevant for the overall performance of the plant. Feeding all data to a diagnostic analysis is prohibitive and not interesting. Rather, the diagnostic tool scans the data base for the relevant disturbances to be explained (see Figure 3). In the first version, the only symptom considered is a stop of the filler.

Whenever the presence of a symptom is confirmed by the data for some time period, a run of the diagnosis module is performed using the symptom (and its temporal extent) as the initial observation. Diagnosis is performed in a **cycle of prediction, observation, and computation of diagnoses** (using RAZ'R ([RAZ'R 09]), a consistency-based diagnosis engine ([Struss 08])). Since the control stopped the filler either due to a lack at the input or because of a tailback at the output, this information is (usually) represented in the data (otherwise, both hypotheses would have to be analyzed). This information is propagated to the adjacent component models and triggers queries to the data base in order to find data that confirm a disturbance of the respective component or refute it. In the latter case, the model of correct behavior generates predictions concerning the behavior of the next machine(s), which are then checked against data retrieved from the data base. Propagation is also continued if there is no information available or if it is too weak.

Whenever refuting or confirming observations are found, they and their temporal extensions are used for the next prediction step, replacing the original predictions. This is important, since model-based temporal prediction has to be very conservative, i.e. generate large intervals due to the uncertainty in the delays of propagated effects in order to guarantee that no evidence is missed. Using the **observed time periods** wherever available **restricts** widening of the **time intervals** significantly. Performing a complete prediction would quickly lead to large intervals and create overly ambiguous and spurious results.

Figure 3 shows a **data interpreter** mediating between the diagnosis module and the data base. This module is essential to the objective of creating a flexible, extensible, and adaptable solution. It organizes two important mappings:

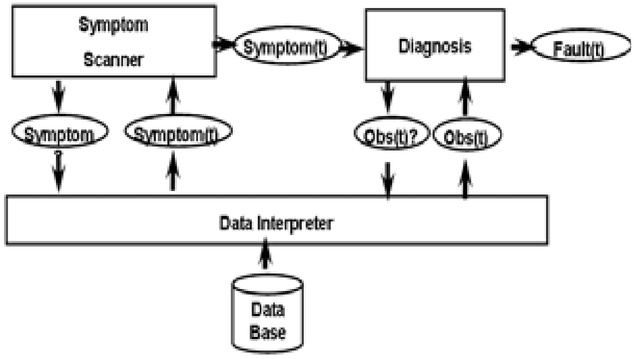

Figure 3. *The architecture of the tool*

- It maps fairly abstract and high-level concepts (called features) the diagnostic model uses to characterize behavior (such as flow, buffer filling degree, lack, tailback) and state queries to the basic and generic quantities represented by the operation data supplied by the machines (see section 3.1).
- In a second step, these generic operation data (such as machine states, counter data, speeds) are mapped to the specific records in the data base, which is needed to generate the queries and ensures the adaptation of the tool to a customer-specific data base by simply specifying this mapping, without having to touch the definition of higher-level features in the first mapping.

In the simplest case, high-level features may correspond directly to (conjunctions or disjunctions of) values assigned to operation data records, e.g. "machineX.lack=T" maps to "machineX.status=Lack", and this in turn is mapped to "machineId.typeId=<code for lack>". As discussed in section 5, more complex features will be needed in an extended solution, e.g. "stuttering" of a machine, i.e. frequent oscillation between operating and stops.

3.3 Models

The model used for diagnosis has to be

- **compositional and component-oriented** to allow for an easy configuration of the model according to the specific plant topology,
- **abstract** in order to cover general classes of machines, rather than a zoo of different variants from different suppliers,
- **qualitative** in order to cope with the uncertainties and non-determinism of the behavior.

They mainly capture the flows of objects, their separation and combination and how interruptions of the flow propagate across the machines. We developed a numerical model, implemented and validated it as a Matlab model, and then abstracted it to a qualitative diagnosis model. For more details, see [Struss et al. 08]. Figure 4 displays the model of two basic elements of the model, a transportation element with a buffer and the virtual component that connects two machines.

Transportation Element with Buffer

State variables

$B(t)$ # objects in buffer

$B_{out}(t)$ # objects buffered for immediate output

$v_{in}(t)$ velocity of input transportation means

$v_{out}(t)$ velocity of output transportation means

$t_d(t)$ minimal transportation time

Parameters

d_0 diameter of transported object (in transportation plain)

C Capacity (as number of objects)

Interface variables

$in.q_{pot}(t)$ potential inflow [objects/s]

$out.q_{pot}(t)$ potential outflow [objects/s]

$in.q_{act}(t)$ actual inflow [objects/s]

$out.q_{act}(t)$ actual outflow [objects/s]

Equations

(1) $in.q_{pot}(t) = v_{in}(t) / d_0$ if $B(t)<C$

$in.q_{pot}(t) = \min (v_{in}(t) /d_0, out.q_{act}(t))$ if $B(t)=C$

(2) $dB/dt = in.q_{act}(t) - out.q_{act}(t)$

(3) $out.q_{pot}(t)= v_{out}(t) /d_0$ if $B_{out}(t)\geq 1$

$out.q_{pot}(t)= \min (in.q_{act}(t - t_d) , v_{out}(t) /d_0)$ else

(4) $dB_{out}(t) /dt = in.q_{act}(t - t_d) - out.q_{act}(t)$

Connector between Transportation Elements

Interface variables

$TE_{n+1}.in.q_{pot}(t)$ potential inflow of upstream element TE_{n+1}

$TE_n.out.q_{pot}(t)$ potential outflow of downstream element TE_n

$TE_{n+1}.in.q_{act}(t)$ actual inflow of upstream element TE_{n+1}

$TE_n.out.q_{act}(t)$ actual outflow of downstream element TE_n

Equations

(5) $TE_n.out.q_{act}(t) = \min (TE_{n+1}.in.q_{pot}(t) , TE_n.out.q_{pot}(t))$

$TE_n.out.q_{act}(t) = TE_{n+1}.in.q_{act}(t)$

Figure 4. *Equations of buffer and connector*

3.4 Generating an Application System

A major requirement was to limit the efforts needed for adaptation to a new or modified plant. This is met by the model-based solution. Firstly, the topology has to be entered, mapping the existing machines to component types, which have an associated model in the library. Due to a common typology of machines based on their basic function (washers, inspectors, buffers …) and the genericity of the models in the library, this has been found to be straightforward.

Additionally, some parameters have to be obtained:

- parameters of single machines, as nominal output rate, capacity, time needed to run empty or full, or reaction times
- parameters of single transporters, such as capacities, reaction times, or information about their geometry, if available,

which are currently exploited for temporal prediction only.

4. RESULTS OF THE FIRST DEMONSTRATOR

4.1 Evaluation Method

We followed two ways to evaluate the diagnostic results of the tool. One was by **comparison** with the fault localization on **two real plants** performed by **human experts**. Since it is

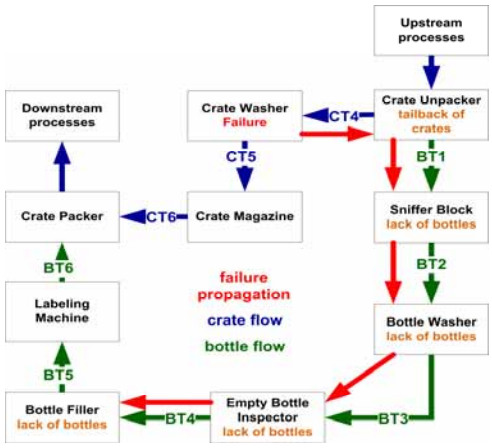

Figure 5. *Failure propagation (red arrows) via sub-branch of production line causing a lack of bottles at the Filler.*

difficult and time-consuming for them to do this based on the recorded data (which is actually what motivated the project!), the reference diagnoses were obtained by direct on-site observation: a number of experts observed the plant and immediately tried to identify the causes for filler stops. Thus, a protocol was created that contains disturbances of the Bottle Filler as symptoms S_{ref} with relevant time intervals, and with manually associated causes $C_{ref}(s)$ for each symptom $s \in S_{ref}$. This protocol serves as reference diagnosis result $result_{ref}$ for assessing the quality of the diagnosis core result $result_{core}$ produced by the LineMod tool. Obviously, this could not be done for longer periods than a few days, but, nevertheless, produced a set of several hundred relevant cases for the evaluation.

Because one cannot expect to encounter all kinds of interesting cases in a limited time (and in the chosen plants), the second way of evaluation was performed by **simulation**. For this purpose, a model of one of the real plants was created for discrete event simulation, which generated a set of 80 cases of disturbances for comparison with the output $result_{core}$ of the tool. It consists of a list of detected symptoms S_{core} with relevant time intervals and with an associated list of pairs of possible causes $C_{core}(s)$ and relevant time intervals for each symptom $s \in S_{core}$ generated by the diagnosis solution.

The causes per symptom instance $s \in S_{ref}$ are called real causes $C_{real}(s)$, which were always singletons in our case.

The evaluation checked for what we call **compliance**, i.e. whether exactly the relevant symptoms (filler stops) were detected and whether the real (observed or simulated) causes were included in the diagnostic result:

$$S_{core} = S_{ref}$$

and

$$\forall\ s \in S_{core} : C_{real}(s) \subseteq C_{core}(s) .$$

Except for individual cases, we could not perform a complete check for false positives, again because this would have required a detailed, time-consuming analysis by experts. However, in general, additional diagnoses generated by the tool were caused by incomplete discrimination due to lacking operation data in some areas of the plants.

One may regret the dependency of the evaluation on the expert diagnostics and their potential faults. Indeed, the automatically obtained results raised doubts for some of the reference diagnostics. Nevertheless, compliance with existing expertise is probably the most important criterion for the quality and acceptance of the tool. This issue is actually deeper than one may expect, as discussed in section 5.

4.2 Sample Scenarios

To illustrate the task, we present two scenarios from the evaluation, which are considered as non-trivial by the experts. While blockage of machines or transportation elements directly downstream or upstream from the filler are frequent causes of lacks or tailbacks, disturbances may propagate via several machines and transportation elements not only on the main-branch of the production line (involving the bottles), but also via its sub-branches.
In a first scenario taken from the simulated plant ("plant C"), a **disturbance at the crate washer** was the origin of a lack of bottles at the bottle filler via the following causal chain displayed in Figure 5 (since, for the real plant A, there is no operation data available for components upstream from the crate unpacker and downstream from the crate packer and, hence, no detailed localization was possible in these parts of the plant, they are represented as aggregates upstream processes and downstream processes, resp.):

- the disturbance at the crate washer causes a **tailback** of crates, which ultimately reaches and stops the **crate unpacker**
- the stopped crate unpacker interrupts the flow of bottles, causing a **lack** which propagates successively to sniffer block, bottle washer, empty bottle inspector and, finally, **bottle filler**.

Figure 6 shows a gantt chart of the failure scenario. This is the main support for the experts in the diagnosis process. Every line represents the operating states of a certain machine within a defined time interval. Especially the lines marked with arrows are of interest, because they show the failure propagation through the plant.

The manual localization is considered complicated because the disturbance originated from the crate line up to the bottle filler. The diagnosis demonstrator correctly localized the crate washer as the origin of the symptom.
The second scenario (shown in Figure 7) is taken from the diagnostic data of the real plant A, which provided quite bad data quality. It illustrates the capabilities of the model-based diagnosis solution even under such conditions. A defective checkmat (a crate inspection machine) upstream from the crate unpacker caused a tailback of bottles at the bottle filler as follows:

- a **lack** of crates propagates, successively stopping crate unpacker, crate washer, crate magazine, and **crate packer**
- the stopped crate packer produces a **tailback** of bottles, which propagates via labeling machine to **bottle filler.**

The respective gantt chart is shown in Figure 8. Despite the completely non-observable crate line between crate unpacker and crate packer (grey-shaded components in the figure), the system model produced predictions for this part and enabled the diagnosis system to produce a proper fault localization, namely the upstream processes component, which contains the defective checkmat, but does not allow for further discrimination due to the lack of data. This example shows that the diagnosis solution provides best possible results, even if the available data is incomplete or not available for some components.

4.3 Statistics

Plant A was not designed according to the data standard WS2005, but according to the predecessor, and produced a fairly insufficient data base. Only 65% of the original amount of symptoms was considered valid, altogether 49 disturbances of the Bottle Filler, whereof 30 were self-caused faults and 19 were caused by other machines or transportation elements (see Figure 9). Bad data included even erroneous state messages of machines, which were eliminated from the reference diagnosis results.

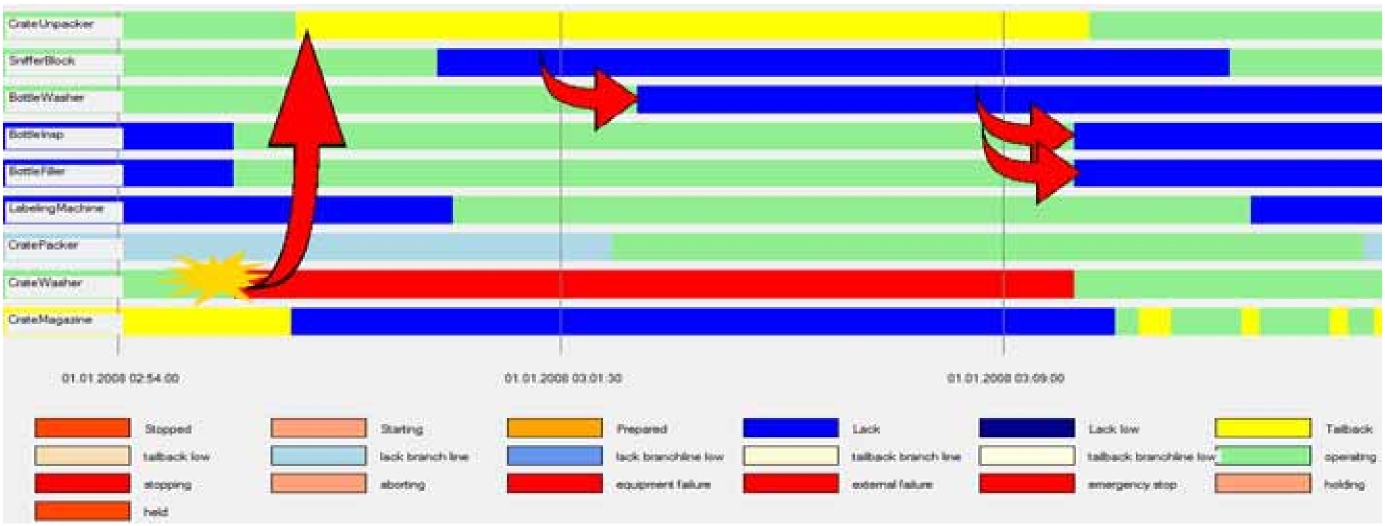

Figure 6. Gantt chart of failure scenario1. Legend: red = failure, yellow = tailback, blue = lack

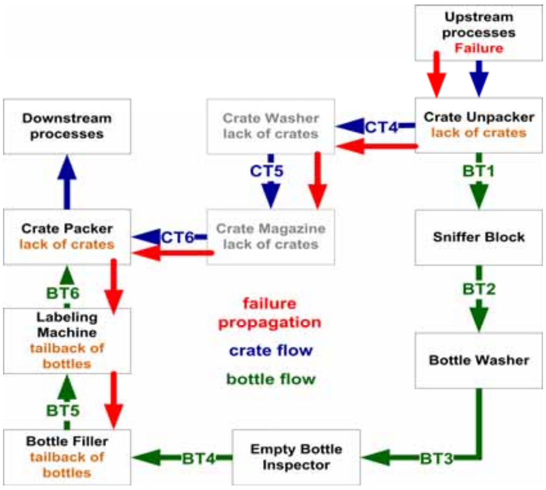

Figure 7. *Failure propagation via sub-branch of production line causing a tailback of bottles at the Filler*

Despite this, Figure 10 shows compliance of 88% of the diagnosis results $result_{core}$ with the reference diagnosis results $result_{ref}$. The remaining 12% can mainly be ascribed to missing data making a correct diagnosis impossible.

Plant B was designed according to the data standard WS2005, which resulted in much better statistical evaluation of the diagnosis results. Nevertheless, even for this plant the amount of bad data is significant (17%), as shown in Figure 9. Most of the bad data deleted from $result_{ref}$ were imprecision or questionable results in the manually obtained localization, or unusable data produced during maintenance, cleaning, or grade change. After deletion of bad data, 416 symptoms were left in $result_{ref}$, whereof 265 disturbances were self-caused faults, and 151 disturbances were caused by other machines or transportation elements.

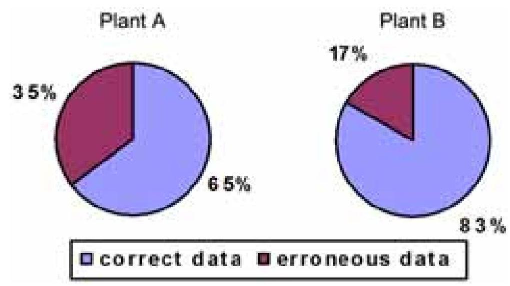

Figure 9. *Data quality of the diagnostic data of plant A (left) and B (right)*

Because of the available data of four production days, the statistical evaluation of the diagnosis results of plant B can be seen as a representative example indicating the quality of results based on an appropriate data base. Figure 10 shows the statistical evaluation of the diagnosis results. The diagram shows that $result_{core}$ complies with $result_{ref}$ in most of the cases.

All of the 80 simulated disturbance cases of **plant C** were correctly diagnosed.

Still there are misdiagnoses, and in case of missing or corrupted data, their number may not be small. However, one has to keep the purpose of the analysis in mind: providing evidence for potential deficiencies of the plant and, hence, possible steps to improving the overall performance. A certain error in the statistics will not trigger a major remedial measure and usually be validated by a more detailed observation and analysis, anyway. This is quite different from, say, diagnosis in a workshop, which is guided by the objective of precisely and unambiguously identifying broken components in order to replace them.

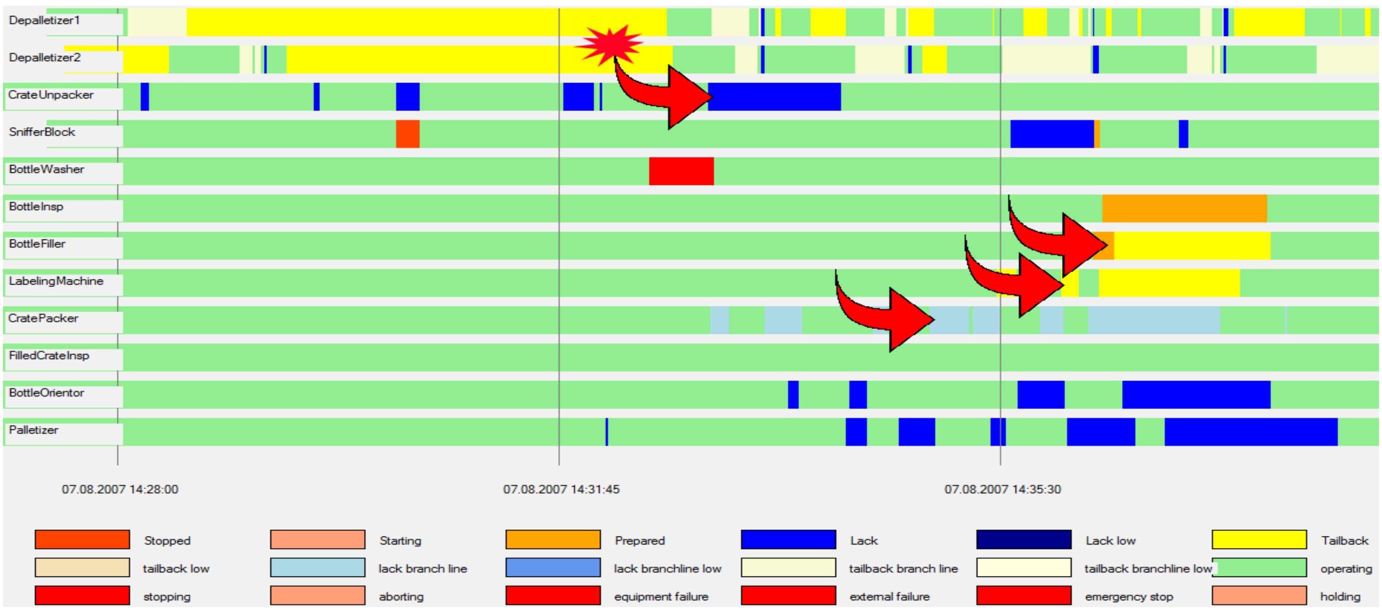

Figure 8. *Gantt chart of failure scenario 2. Legend: yellow = tailback, blue = lack, light blue = lack branch line (crate line)*

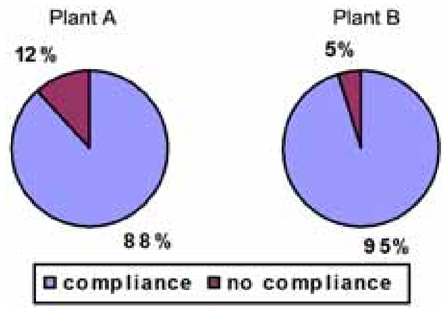

Figure 10. *Compliance of core diagnosis results with corrected reference diagnosis results of plant A (left) and B (right)*

4.4 Commercial Exploitation

The results obtained from the first demonstrator have triggered significant interest of both end users (breweries etc.) and suppliers. Two of the latter have started to include results from the project in their analysis tools, which are essentially based on decision trees. A third one is forming a joint venture with a software company to extend its tools by exploiting the model-based solution described in this paper. It is expected that even under the limited scope of the current demonstrator (diagnosing "hard" reasons for filler stops only), a commercial benefit will be obtained.
The following section discusses the practical importance of widening this scope and the resulting challenges to model-based diagnosis, which has lead to a follow-on project.

5. CHALLENGES

5.1 Indexing with Metric Temporal Scopes

As discussed in section 3, the data as well as the high-level features used in the models need to be indexed with intervals of absolute time. Without obtaining and processing this temporal information, no proper diagnosis can be obtained. The problem is that there is no finite (or relatively small) universe of relevant time points that enables the exploitation of previous qualitative temporal indexing schemes (e.g. [Williams 86], [Dressler-Freitag 94]), because the temporal granularity of data is small compared to the time span between a bottle entering and leaving the plant (seconds to an hour or so).
The demonstrator used an ad-hoc solution that separated a specialized model for the temporal inferences from a model with time stripped off, which is used in the consistency-based diagnosis engine. For a future system, a principled solution has to be developed that avoids the double work of maintaining two closely related models.

5.2 Erroneous Data

In our application domain, erroneous data will inevitably occur. The problem for a straight-forward consistency-based solution is that wrong observations cause spurious inconsistencies, which bears a high risk of missing the proper diagnosis. The only way out is not treating observations as true propositions. When evaluating the current demonstrator, we would usually detect the presence of corrupted operation data though the generation of highly implausible or unlikely diagnoses (or none at all). In fact, one can extend the approach to include the "diagnosis of data", which has actually been proposed some 20 years ago ([Struss 88]).
A fairly principled way to realize this is by introducing components that represent either the ultimate source generating some observation (sensors) or the control software processing and supplying the data, which appears more appropriate in our application. The diagnosis algorithm generates diagnostic hypotheses that may include such data sources. As a result, a highly unlikely combination of faults of physical components may compete with the combination of a single (or even no) fault and one or more incriminated data, providing hints for erroneous data.
A quick demonstration of this approach was found highly attractive by the experts, and, hence, a tool for detecting erroneous operation data as a spin-off solution. While it may be an overload during the routine diagnostic analysis, it should be a useful tool during the installation phase of a new plant or the adaptation of the diagnostic system.

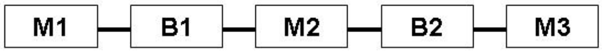

Figure 11. *Three machines M_i with buffers B_j in between*

5.3 Complex Faults

The diagnostic task solved by the current demonstrator (finding the machine whose interrupted operation caused the filler to stop) appears to be straightforward. In reality, there are cases in which such an answer is at least questionable or even misleading. This has to do with the fact that the plant achieves some tolerance with respect to limited disturbances of individual machines through buffer elements.
To illustrate this, we consider a sequence of three machines M_i with two intermediate buffers B_j (Figure 11). If buffer B_2 is filled to some extent, it will prevent that some short interruption of output from M_2 will cause missing input to M_3. Now assume that M_1 experienced a series of several small disturbances over a longer period of time ("stuttering"), none of which caused a lack of input to M_2 or M_3.
However, as an accumulated effect of this stuttering, the amount of objects buffered in B_2 has been reduced significantly, affecting its capability to compensate for further interruption of flow. As a result, a fairly short fault in M_2, which would have had no serious effect under "normal" circumstances, now causes a lack at M_3, and the algorithm described would actually detect this and blame solely M_2 for the stop of M_3.
In order to properly understand the origin of the symptom and for determining appropriate remedies, it would be necessary to reveal the role of M_1 as the originator of or, at least, one contributor to the problem. While it is possible to extend the time windows and the observed behavior patterns in the tool to detect the "stuttering", it is not obvious whether and how to distribute the blame among the components, even more if there are more complex situations of disturbances in several machines.
The core of this problem lies in the fact that there is an element of the designed plant behavior that is not captured by the local component models: although there is no global

control, there do exist some principles and intentions how to run the plant properly. For instance, as mentioned in section 2, the output rates of the machines should increase both upstream and downstream from the filler (to avoid filler lacks and tailbacks, respectively). For the same reason, the filling degree of buffers should be “neither too high nor too low”. Unfortunately, this healthy degree is not fixed and, more importantly, cannot be measured directly or appropriately estimated from other sensor data.

A solution to this problem is one of the main targets of the follow-on project. Besides a principled way to assess the contribution of several machines, this will require the representation of more complex behavior patterns, such as “stuttering” of machines. The current demonstrator has already been designed to support this by enabling the definition of complex features via more low level variables. Stuttering, for instance, can be specified as the presence of an oscillation between different status values.

Another challenge arises from the fact that, sometimes, the ultimate cause of a problem lies beyond the boundaries of the physical plant, which is captured by the current model. For instance, the logistics (e.g. delivering pallets with improper bottles or with delay), manual intervention (sorting out improper bottles), or maintenance schedule may be inappropriate. In the follow-on project, we will attempt to extend the model to include such aspects, although there are usually few related observations available.

ACKNOWLEDGEMENTS

Our work would have been impossible without the expertise and help of our colleagues from the department of food packaging technology (LVT) and the other members of the Model-based Systems and Qualitative Modeling Group (MQM). We are also indebted to the end users and suppliers for their support, which was well beyond what you can usually expect from industrial partners.

REFERENCES

[Dressler-Freitag 94] Dressler. O. and Freitag, H.: Prediction sharing across time and contexts. In: *Proceedings of the National Conference on Artificial Intelligence* (AAAI-94)

[Kather et al. 05] Kather, A.; Voigt, T.; Langowski, H.-C..Plug and Acquire – The new Weihenstephan Standards for process data acquisition in bottling plants. In: *Proceedings of the 30th EBC Congress*, Prag. Nürnberg: Hans Carl, 2005, on CD

[RAZ’R 09] RAZ’R 09 www.occm.de

[Struss 88] Struss, P. Extensions to ATMS-based diagnosis. In: J. S. Gero (Ed.) *Artificial Intelligence in Engineering: Diagnosis and Learning*, pp. 3-27, August 1988

[Struss et al. 08a] Struss, P., Kather, A., Schneider, D., and Voigt, T.: Qualitative modeling for diagnosis of machines transporting rigid objects. In: Bradley, L. and Trave-Massuyes, L (eds.). *22nd International Workshop on Qualitative Reasoning QR08*, Boulder, Co, 2008

[Struss et al. 08b] Struss, P., Kather, A., Schneider, D., Voigt, T.: Qualitative modeling for diagnosis of machines transporting rigid objects. In: Grastien, A. and Stumptner, M. (eds.), *19th International Workshop on Principles of Diagnosis* (DX08), Blue Moutains, Australia

[Struss 08] P. Struss: Model-based Problem Solving. In: van Harmelen, F., Lifschitz, V., and Porter, B. (eds.). *Handbook of Knowledge Representation*, Elsevier, 2008

[Williams 86] Williams, B.: Doing Time: Putting qualitative reasoning on firmer ground." In: *Proceedings of the National Conference on Artificial Intelligence*, Philadelphia, PA, 1986, pp. 105-113