

5

METHODS FOR DETERMINING THE STATE OF HEALTH OF AN INDUSTRIAL PROCESS

10 FIELD OF INVENTION

The invention relates to the monitoring of industrial processes for current or impending anomalies.

15 BACKGROUND

In the dynamic industrial processes, complexity is increasingly growing as gradually more and various process units are included where the process health of each of these units contributes to the overall process health. Different statistical analysis methods
20 have been widely used in industrial plants to investigate the different data signals as independent from each other. However, it is important to consider the hierarchical relations and causality directions in taking a step forward in detecting overall health issues.

25 EP 3 379 357 B1 discloses to group technical status data from a process into signal groups according to, inter alia, the physical location of the respective signals within the industrial process system, or a correlation analysis. Signals from each group are then analyzed by a machine-learning model dedicated to this group.

30 OBJECTIVE OF THE INVENTION

It is the objective of the present invention to further improve the accuracy with which the state of health of an industrial process as a whole may be determined.

- 2 -

This objective is achieved by the methods according to the independent claims. Further advantageous embodiments are detailed in the respective dependent claims.

DISCLOSURE OF THE INVENTION

5

The invention provides a method for determining the state of health of an industrial process. This process is executed by at least one industrial plant that comprises an arrangement of entities. These entities may have any level of granularity. For example, an entity may correspond to a piece of equipment, such as a pump, a tank, a reactor vessel, a pipe or a valve. But an entity may also, for example, correspond to a sub-division of the process (like "fabrication of substance A", "fabrication of substance B" and "manufacture of product C from substances A and B"), or to a process module of a modular plant. The state of each entity is characterized by a set of entity state variables.

15

In the course of the method, values, and/or time series of values, of the entity state variables are obtained. For each entity, these values, and/or time series of values, are provided to a model corresponding to the respective entity, so that a prediction of the state of health of the respective entity is obtained. This model may, for example, predict at least one further variable of the entity from which the state of health of the entity may then be inferred. The model may also, for example, directly predict the state of health of the entity. The state of health may be in the form of a numerical score (for example, on a scale between 0 and 1, where 1 is best), but it may also be in the form of one of several available discrete classes. The classes may, for example, represent discrete grades (like from A down to F), or a binary classification as "healthy" or "unhealthy".

25

Models corresponding to entities may, for example, be machine learning models. From a set of training examples, these models learn the unknown propagation from the inputs (i.e., the state variables) to the outputs (i.e., the further variable or the state of health). Alternatively or in combination, for entities for which there is at least some a priori knowledge about the propagation from the inputs to the outputs, a simulation model may be used. Depending on the accuracy and level of detail with which such a simulation model mimics the behavior of the entity, the model may be expensive to compute. Therefore, a surrogate approximation of the simulation model may be used in place of the actual simulation model. The surrogate approximation may, for example,

35

- 3 -

be a machine learning model that learns the behavior of the simulation model. Different types of models may be used for different entities.

5 Based at least in part on the layout of the industrial plant, propagation paths for anomalies between the entities that are involved in the process are determined.

Based at least in part on these propagation paths, importances of the states of health of the individual entities for the overall state of health of the process are determined. Based at least in part on these importances, the individual states of health are
10 aggregated to obtain the overall states of health of the process.

The inventors have found that decomposing the determining of the state of health into these method steps provides, at each of these steps, the opportunity to consider a priori knowledge about the entities, or the interaction between these entities in the
15 process as a whole. Such a priori knowledge may manifest itself in the models of the entities, in the propagation paths for anomalies, in the importances of the individual entity states of health for the overall state of health of the process, and in the mode by which the individual states of health are aggregated. The multiple opportunities to consider such knowledge at the different steps allow for a wider range of the available
20 knowledge to be utilized. Therefore, given a concrete heap of available a priori knowledge, the accuracy with which the overall state of health may be determined is improved.

In particular, the method allows to separate qualitative causalities between anomalies
25 in entities from quantitative causalities. The qualitative causality resides in the topology of propagation paths between entities and contains the information which entities may, in principle, be in a chain of events that starts in one entity and escalates to a degradation, a shutdown or a failure of the process as a whole. The quantitative causality is a measure of the probability of such an escalation happening. This
30 quantitative causality may, for example, be encoded in strengths of the propagation paths, or in the importances of the individual entity states of health that are used when aggregating them to an overall state of health of the process. The quantitative causality is not necessarily dependent only on two particular entities that are connected by a propagation path. Rather, it may also depend on the layout of the plant as a whole. For
35 example, the functionality of certain entities may be realized in the plant with a certain

- 4 -

level of redundancy. One example for this is an arrangement of parallel pumps that share the work of transporting a certain amount of fluid. If one pump fails, the remaining pumps may make up at least partially for the shortfall.

5 In a particularly advantageous embodiment, the propagation paths of anomalies are determined based at least in part on a topology of material and/or energy flows between entities. Such flows are a primary reason why entities in an industrial process are dependent on each other. In particular, many chemical processes are multi-stage processes where the product of one stage is the educt that the next stage needs.

10

In a further advantageous embodiment, at least one known directional interaction path from a first entity to a second entity during normal operation of the plant may be expanded to a bidirectional propagation path for anomalies between the first entity and the second entity. This is based on the experience that an anomaly in a process may reverse a cause-and-effect relationship between different entities. For example, while 15 the flow rate of a pump is normally causal for the fill level in a downstream tank, an anomaly may cause the tank to overfill, so that it cannot accept any more liquid from the pump. This in turn causes the flow rate of the pump to come to a halt because the liquid is not compressible.

20

Therefore, in a further particularly advantageous embodiment, at least one propagation path specifically comprises a cause-effect relationship between

- a first entity state variable of a first entity and a second entity state variable of a second entity; and/or
- 25 • occurrence of a first event in a first entity and occurrence of a second event in a second entity.

One possible source for these propagation paths is a set of already existing interlocking rules between triggering events and events that happen as a consequence 30 of these triggering events. For example, if pressure in a vacuum chamber is too high, certain sensitive equipment, like a mass spectrometer or a getter pump, may be prevented from starting in order to prevent damage to this equipment.

However, propagation paths for anomalies between entities may also be determined in 35 a data-driven manner without a priori knowledge. In a further particularly advantageous

embodiment, the determining of at least one propagation path from a first entity to a second entity specifically comprises:

- computing a statistical impact of a first time series of a first entity state variable of the first entity onto a second time series of a second entity state variable of the second entity; and
- quantifying the strength of propagation of anomalies between the first and second entities based at least in part on this statistical impact.

For example, the computing of a statistical impact may comprise computing the Granger Causality, and/or the Transfer Entropy, between the first and second time series. The Granger Causality is a statistical hypothesis test for testing to which extent the first time series is useful for predicting the second time series. The Transfer Entropy focuses more on an information gain between probability distributions.

- In a further advantageous embodiment, the determining of propagation paths may comprise:
- verifying a feasibility of each propagation path from a set of candidate propagation paths; and
 - in response to a propagation path not being feasible, removing this propagation path from the set of candidate propagation paths.

In particular, when expanding a directional interaction path from a first entity to a second entity to a bidirectional interaction path between the first entity and the second entity, the reverse direction may turn out not to be feasible. For example, if material is transported from the first entity to the second entity by falling downwards, it cannot fall upwards again. Such infeasibilities of reverse interactions might even have been deliberately devised as safety features when designing the industrial plant. Therefore, removing infeasible propagation paths gives a more accurate picture of the overall state of health.

In a further advantageous embodiment, the importance of the state of health of an entity increases with the number, and/or the strengths, or propagation paths that comprise this entity. For example, a process module may produce a product that is used as educt by multiple downstream process modules. In another example, if the process comprises entities that reside on different physical plants connected by a

- 6 -

pipeline, this pipeline may be a “single point of failure” for the whole process and therefore have a very high importance in the determining of the overall state of health.

For example, in a system with a pump and a pipe that feed in fluid to a tank, an overly high fill level of the tank may preclude the feeding of further fluid and cause back-pressure on the pump. Therefore, the pump is in a causal relationship with the high fill level of the tank. But the pressure in the pipe is not in any causal relationship with another entity. Therefore, the state of health of the pump gets a higher importance than the state of health of the pipe in the overall state of health of the process.

In a further advantageous embodiment, the aggregating of individual states of health comprises computing a weighted sum of these individual states of health, wherein the weights are determined based on the importances of the respective individual states of health. This results in an easily interpretable indicator of the overall state of health, so that a threshold value for the binary classification of the overall state as “healthy” or “unhealthy” may be easily motivated. However, alternatively or in combination to the weighted sum, the aggregation of the individual states of health may also comprise, for example, multiplicative contributions. For example, if health is measured on a scale between 0 and 1, and a certain entity is a “single point of failure” for the functioning of the process as a whole, then the contribution of this entity may be multiplicative. The overall state of health of the process can then not be better than the state of health of the “single point of failure” entity.

As discussed before, in a particularly advantageous embodiment, the industrial process may be a chemical production process that converts one or more educts into one or more products. At least one entity in the arrangement of entities may be a pipe, a pump, or a tank. These are the entities with the most and strongest propagation paths for anomalies between them.

The invention also provides a second method for determining the state of health of an industrial process. In an analogous manner to the first method, the process is executed by at least one industrial plant comprising an arrangement of entities, and the state of each such entity is characterized by a set of entity state variables. Also, values, and/or time series of values, of the entity state variables are obtained.

- 7 -

In contrast to the first method, the values, and/or the time series of values, are now provided to a model of the industrial process as a whole. This model is implemented as a machine learning model. From the outcome of the model, a prediction of the state of health of the industrial process as a whole is obtained. To this end, akin to the models
5 corresponding to individual entities that are used in the first method, the single model for the process as a whole may output a variable of the process from which the state of health may then be inferred, or it may directly predict the state of health.

The basic idea of exploiting the entity state variables to determine the overall state of health is the same that is also used in the first method. This second method, however,
10 does not presume that a priori knowledge about interactions between entities is available. For example, the single model may be directly trained in a "black box" manner on training data that comprises values and/or time series of values of the entity state variables as inputs, and states of health of the process as a whole as ground truth
15 labels. Which of the methods is more advantageous to use therefore depends on the amount of a priori information that is available in the concrete situation.

In a further advantageous embodiment, the machine learning model comprises a recurrent neural network, RNN, and/or a neural transformer network. These networks
20 are particularly well suited for processing time series of state variables as sequences.

The method may be wholly or partially computer-implemented. The invention therefore also relates to one or more computer programs with machine-readable instructions that, when executed on one or more computers and/or compute instances, cause the
25 one or more computers to perform one of the methods described above. In this context, a virtualization platform, a hardware controller, network infrastructure devices (such as switches, bridges, routers or wireless access points), as well as end devices in the network (such as sensors, actuators or other industrial field devices) that are able to execute machine readable instructions are to be regarded as computers as well.

30 The invention therefore also relates to a non-transitory storage medium, and/or to a download product, with the one or more computer programs. A download product is a product that may be sold in an online shop for immediate fulfillment by download. The invention also provides one or more computers and/or compute instances with the one

- 8 -

or more computer programs, and/or with the one or more non-transitory machine-readable storage media and/or download products.

Irrespective of which method is used to determine the state of health, the determined
5 state of health may be used to provide an early warning in case of a deteriorating
process health. This may be exploited to take action in order to keep the plant safe. For
example, in response to the determined state of health of the process meeting a
predetermined criterion (such as falling below a predetermined threshold value),
parameters that characterize the behavior of the process may be modified and/or
10 optimized such that the state of health of the process improves. Such parameters may,
for example, comprise set-point values for low-level controllers. In particular, in one
example, the one or more models that are used to predict the overall state of health of
the process may be called upon multiple times for different candidate actions to be
taken on the process in order to determine which of these candidate actions is best to
15 improve the overall state of health of the process. In this case, it is particularly
advantageous if the overall state of health can be predicted much faster than in real-
time. For example, it may then be even more advantageous to use surrogate models
as an approximation for computationally expensive simulation models.

20 DESCRIPTION OF THE FIGURES

In the following, the invention is illustrated using Figures without any intention to limit
the scope of the invention. The Figures show:

25 Figure 1: Exemplary embodiment of the method 100 for determining the state of health
of an industrial process;

Figure 2: Exemplary embodiment of the method 200 for determining the state of health
of an industrial process;

30 Figure 3: Exemplary system where the methods 100, 200 may be applied.

Figure 1 is a schematic flow chart of an embodiment of the method 100 for determining
35 the state of health 1* of an industrial process 1. The process 1 is executed by at least

one industrial plant comprising an arrangement of entities 2a-2f. The state of each such entity 2a-2f is characterized by a set of entity state variables 3a-3f.

5 In step 110, values, and/or time series of values, of the entity state variables 3a-3f are obtained.

10 In step 120, for each entity 2a-2f, the values 3a-3f, and/or time series of values 3a-3f, are provided to a model 4a-4f corresponding to the respective entity 2a-2f. In this manner, a prediction of the state of health 5a-5f of the respective entity 2a-2f is obtained.

According to block 121, the model 4a-4f corresponding to at least one entity 2a-2f may comprise

- a machine learning model; and/or
- 15 • a simulation model; and/or
- a surrogate approximation of this simulation model.

20 In step 130, based at least in part on the layout 1a of the industrial plant executing the process 1, propagation paths 6 for anomalies between said entities 2a-2f are determined.

According to block 131, such propagation paths 6 may be determined based at least in part on a topology of material and/or energy flows between entities 2a-2f.

25 According to block 132, at least one known directional interaction path from a first entity 2a-2f to a second entity 2a-2f during normal operation of the plant may be expanded to a bidirectional propagation path for anomalies between the first entity and the second entity.

30 According to block 133, at least one propagation path 6 may specifically comprise a cause-effect relationship between

- a first entity state variable 3a-3f of a first entity 2a-2f and a second entity state variable 3a-3f of a second entity 2a-2f; and/or
- 35 • occurrence of a first event in a first entity 2a-2f and occurrence of a second event in a second entity 2a-2f.

According to block 134, a statistical impact of a first time series of a first entity state variable 3a-3f of the first entity 2a-2f onto a second time series of a second entity state variable 3a-3f of the second entity 2a-2f may be computed. In particular, according to
5 block 134a, this may comprise computing the Granger Causality, and/or the Transfer Entropy, between the first and second time series. According to block 135, based at least in part on the computed statistical impact, the strength of propagation of anomalies between the first and second entities 2a-2f may be quantified.

10 According to block 136, a feasibility of each propagation path 6 from a set of candidate propagation paths 6 may be verified. According to block 137, in response to a propagation path 6 not being feasible, this propagation path 6 may be removed from the set of candidate propagation paths 6.

15 In step 140, based at least in part on the propagation paths 6, importances 7a-7f of the states of health 5a-5f of the individual entities (2a-2f) for the overall state of health (1*) of the process may be determined.

20 According to block 141, the importance 7a-7f of the state of health 5a-5f of an entity 2a-2f may increase with the number, and/or the strengths, of propagation paths 6 that comprise this entity 2a-2f.

In step 150, based at least in part on said importances 7a-7f, the individual states of health 5a-5f of the entities 2a-2f are aggregated to obtain the overall state of health 1*
25 of the process 1.

According to block 151, based on the importances 7a-7f of the respective individual states of health 5a-5f, weights for these individual states of health 5a-5f may be determined. According to block 152, using these weights, a weighted sum of these
30 individual states 5a-5f may be computed.

Figure 2 is an exemplary embodiment of the method 200 for determining the state of health 1* of an industrial process 1. Like the method 100, the method 200 starts from the situation that the process 1 is executed by at least one industrial plant comprising

an arrangement of entities 2a-2f, and that the state of each such entity 2a-2f is characterized by a set of entity state variables 3a-3f.

5 In step 210, values, and/or time series of values, of the entity state variables 3a-3f are obtained.

In step 220, the values, and/or the time series of values, are provided to a model 1b of the industrial process 1. This model 1b is implemented as a machine learning model. From the outcome of the model, a prediction of the state of health 1* of the industrial
10 process is obtained. The model 1b may directly output the sought state of health 1*, or it may output a variable of the process 1 that is indicative of the state of health 1*.

According to block 221, the machine learning model 1b comprises a recurrent neural network, RNN, and/or a neural transformer network that is particularly well suited for
15 processing time series and other sequences of values.

Figure 3 shows an exemplary system where the methods 100 and/or 200 may be applied. The system comprises concrete examples for the entities 2a-2f. Fluid is pumped by a pump 2a through a valve 2b into a tank 2c. The tank 2c may also be filled
20 from another source via a valve 2d. Fluid from the tank 2c passes through a pressurized pipe 2e to another valve 2f. The pressure in the pipe 2e may be measured by means of a pressure gauge that converts the pressure p to an electrical current I.

List of reference signs:

	1	industrial process
5	1*	state of health of process 1
	1a	topology of plant executing process 1
	1b	model of process 1
	2a-2f	entities of plant executing process 1
	3a-3f	entity state variables for entities 2a-2f
10	4a-4f	models corresponding to entities 2a-2f
	5a-5f	states of health of entities 2a-2f
	6	propagation paths for anomalies between entities 2a-2f
	7a-7f	importances of states of health 5a-5f
	100	first method for determining state of health 1*
15	110	obtaining entity state variables 3a-3f
	120	providing values 3a-3f to models 4a-4f
	121	using particular types of models
	130	determining propagation paths 6 between entities 2a-2f
	131	evaluating topology of material and/or energy flows
20	132	expanding directional paths to bidirectional propagation paths 6
	133	considering particular cause-effect relationships
	134	computing statistical impact between time series
	134a	computing Granger Causality or Transfer Entropy
	135	quantifying strength of propagation
25	136	verifying feasibility of propagation path 6
	137	removing infeasible propagation path 6
	140	determining importances 7a-7f
	141	making importances 7a-7f dependent on propagation paths 6
	150	aggregating states of health 5a-5f using importances 7a-7f
30	151	computing weights based on importances 7a-7f
	152	computing weighted sum
	200	second method for determining state of health 1*
	210	obtaining entity state variables 3a-3f
	220	providing values to model 1b of process 1
35	221	using RNN or transformer network

p	pressure
I	electrical current

Claims:

1. A method (100) for determining the state of health (1*) of an industrial process (1), wherein the process (1) is executed by at least one industrial plant comprising an arrangement of entities (2a-2f) and the state of each such entity (2a-2f) is characterized by a set of entity state variables (3a-3f), comprising the steps of:
 - obtaining (110) values, and/or time series of values, of the entity state variables (3a-3f);
 - for each entity (2a-2f), providing (120) the values (3a-3f), and/or time series of values (3a-3f), to a model (4a-4f) corresponding to the respective entity (2a-2f), thereby obtaining a prediction of the state of health (5a-5f) of the respective entity (2a-2f);
 - determining (130), based at least in part on the layout (1a) of the industrial plant executing the process (1), propagation paths (6) for anomalies between said entities (2a-2f);
 - determining (140), based at least in part on said propagation paths (6), importances (7a-7f) of the states of health (5a-5f) of the individual entities (2a-2f) for the overall state of health (1*) of the process; and
 - aggregating (150), based at least in part on said importances (7a-7f), the individual states of health (5a-5f) of the entities (2a-2f) to obtain the overall state of health (1*) of the process (1).
2. The method (100) of claim 1, wherein the propagation paths (6) of anomalies are determined (131) based at least in part on a topology of material and/or energy flows between entities (2a-2f).
3. The method (100) of any one of claims 1 to 2, wherein at least one known directional interaction path from a first entity (2a-2f) to a second entity (2a-2f) during normal operation of the plant is expanded (132) to a bidirectional propagation path for anomalies between the first entity and the second entity.
4. The method (100) of any one of claims 1 to 3, wherein at least one propagation path (6) specifically comprises (133) a cause-effect relationship between
 - a first entity state variable (3a-3f) of a first entity (2a-2f) and a second entity state variable (3a-3f) of a second entity (2a-2f); and/or

- occurrence of a first event in a first entity (2a-2f) and occurrence of a second event in a second entity (2a-2f).

5. The method (100) of any one of claims 1 to 4, wherein the determining (130) of at least one propagation path (6) from a first entity (2a-2f) to a second entity (2a-2f) specifically comprises:

- computing (134) a statistical impact of a first time series of a first entity state variable (3a-3f) of the first entity (2a-2f) onto a second time series of a second entity state variable (3a-3f) of the second entity (2a-2f); and
- quantifying (135) the strength of propagation of anomalies between the first and second entities (2a-2f) based at least in part on this statistical impact.

6. The method (100) of claim 5, wherein the computing (134) of a statistical impact specifically comprises: computing (134a) the Granger Causality, and/or the Transfer Entropy, between the first and second time series.

7. The method (100) of any one of claims 1 to 6, wherein the determining (130) of propagation paths (6) specifically comprises:

- verifying (136) a feasibility of each propagation path (6) from a set of candidate propagation paths (6); and
- in response to a propagation path (6) not being feasible, removing (137) this propagation path (6) from the set of candidate propagation paths (6).

8. The method (100) of any one of claims 1 to 7, wherein the importance (7a-7f) of the state of health (5a-5f) of an entity (2a-2f) increases (141) with the number, and/or the strengths, of propagation paths (6) that comprise this entity (2a-2f).

9. The method (100) of any one of claims 1 to 8, wherein the model (4a-4f) corresponding to at least one entity (2a-2f) comprises (121):

- a machine learning model; and/or
- a simulation model; and/or
- a surrogate approximation of this simulation model.

10. The method (100) of any one of claims 1 to 9, wherein the aggregating (150) of individual states of health (5a-5f) comprises computing (152) a weighted sum of these individual states of health (5a-5f), wherein the weights are determined (151) based on the importances (7a-7f) of the respective individual states of health (5a-5f).

11. The method (100) of any one of claims 1 to 10, wherein the industrial process (1) is a chemical production process that converts one or more educts into one or more products, and at least one entity in the arrangement of entities is a pipe, a pump, or a tank.
- 5 12. A method (200) for determining the state of health (1*) of an industrial process, wherein the process is executed by at least one industrial plant comprising an arrangement of entities (2a-2f) and the state of each such entity (2a-2f) is characterized by a set of entity state variables (3a-3f), comprising the steps of:
- 10
- obtaining (210) values, and/or time series of values, of the entity state variables (3a-3f); and
 - providing (220) the values, and/or the time series of values, to a model (1b) of the industrial process (1), wherein this model (1b) is implemented as a machine learning model, thereby obtaining a prediction of the state of health
- 15 (1*) of the industrial process (1).
13. The method (200) of claim 12, wherein the machine learning model (1b) comprises (221) a recurrent neural network, RNN, and/or a neural transformer network.
14. A computer program, comprising machine-readable instructions that, when executed by one or more computers, cause the one or more computers to perform a method (100, 200) of any one of claims 1 to 13.
- 20
15. A non-transitory machine-readable storage medium, and/or a download product, with the computer program of claim 14.
16. One or more computers with the computer program
- 25 of claim 14, and/or with the non-transitory machine-readable storage medium and/or download product of claim 15.