

(19) **DANMARK**

(10) **DK/EP 3804268 T3**



(12) **Oversættelse af
europæisk patentskrift**

Patent- og
Varemærkestyrelsen

-
- (51) Int.Cl.: **H 04 L 9/40 (2022.01)** **F 03 D 7/00 (2006.01)** **F 03 D 17/00 (2016.01)**
F 03 D 80/00 (2016.01) **G 06 F 21/55 (2013.01)**
- (45) Oversættelsen bekendtgjort den: **2023-11-20**
- (80) Dato for Den Europæiske Patentmyndigheds bekendtgørelse om meddelelse af patentet: **2023-08-23**
- (86) Europæisk ansøgning nr.: **19731829.8**
- (86) Europæisk indleveringsdag: **2019-05-23**
- (87) Den europæiske ansøgnings publiceringsdag: **2021-04-14**
- (86) International ansøgning nr.: **US2019033664**
- (87) Internationalt publikationsnr.: **WO2019226853**
- (30) Prioritet: **2018-05-24 US 201815988515**
- (84) Designerede stater: **AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR**
- (73) Patenthaver: **General Electric Renovables España, S.L., Calle Roc Boronat 78, , 08005 Barcelona, Spanien**
- (72) Opfinder: **ABBASZADEH, Masoud, c/o General Electric Company One Research Circle, Niskayuna, New York 12309, USA**
- (74) Fuldmægtig i Danmark: **COPA COPENHAGEN PATENTS K/S, Rosenørns Allé 1, 2. sal, 1970 Frederiksberg C, Danmark**
- (54) Benævnelse: **SYSTEM OG FREMGANGSMÅDE TIL ANOMALI- OG CYBERTRUSSELEDETEKTION I EN VINDMØLLE**
- (56) Fremdragne publikationer:
EP-A1- 3 239 884
KR-A- 20140 109 132

DESCRIPTION

BACKGROUND

[0001] Industrial control systems that operate wind turbines are increasingly connected to the Internet. As a result, these control systems are vulnerable to threats, such as cyber-attacks (e.g., associated with a computer virus, malicious software, *etc.*) that could disrupt electric power generation and distribution, damage turbines, *etc.* Current methods primarily consider attack detection in Information Technology ("IT," such as, computers that store, retrieve, transmit, manipulate data) and Operation Technology ("OT," such as direct monitoring devices and communication bus interfaces). Cyber-attacks can still penetrate through these protection layers and reach the physical "domain" as seen in 2010 with the Stuxnet attack. Such attacks can diminish the performance of a control system and may cause total shut down or even catastrophic damage to a wind turbine. Currently, no methods are available to automatically detect, during a cyber-incident, attacks at the domain layer where sensors, controllers, and actuators are located. In some cases, multiple attacks may occur simultaneously (e.g., more than one actuator, sensor, or parameter inside control system devices might be altered maliciously by an unauthorized party at the same time). Note that some subtle consequences of cyber-attacks, such as stealthy attacks occurring at the domain layer, might not be readily detectable (e.g., when only one monitoring node, such as a sensor node, is used in a detection algorithm). It may also be important to determine when a monitoring node is experiencing a fault (as opposed to a malicious attack) and, in some cases, exactly what type of fault is occurring. Existing approaches to protect wind turbines, such as failure and diagnostics technologies, may not adequately address these problems - especially when multiple, simultaneous attacks and/faults occur since such multiple faults/failure diagnostic technologies are not designed for detecting stealthy attacks in an automatic manner.

EP3239884A1 describes a domain level threat detection for industrial asset control system;

KR20140109132A describes a warning generation method through monitoring of power curve.

[0002] It would therefore be desirable to protect a wind turbine from cyber-attacks in an automatic and accurate manner even when attacks percolate through the IT and OT layers and directly harm control systems.

SUMMARY

[0003] The invention is defined by the appended claims.

[0004] According to some embodiments, a plurality of monitoring nodes each generate a

series of current monitoring node values over time that represent a current operation of a wind turbine. An abnormality detection computer platform may receive the series of current monitoring node values and generate a set of current feature vectors. The abnormality detection computer platform may also access an abnormality detection model having a plurality of decision boundaries created using wind information (e.g., wind speed and/or acceleration) along with at least one of a set of normal feature vectors and a set of abnormal feature vectors. The abnormality detection computer platform may then select one of the decision boundaries based on current wind information associated with the wind turbine and execute the abnormality detection model and transmit an abnormality alert signal based on the set of current feature vectors and the selected decision boundary

[0005] Some embodiments comprise: means for receiving a series of current monitoring node values from a plurality of monitoring nodes each generating a series of current monitoring node values over time that represent a current operation of a wind turbine; means for generating a set of current feature vectors based on the series of current monitoring node values; means for accessing, by an abnormality detection computer platform, an abnormality detection model having a plurality of decision boundaries created using wind information along with at least one of a set of normal feature vectors and a set of abnormal feature vectors; means for selecting one of the decision boundaries based on current wind information associated with the wind turbine; and means for executing the abnormality detection model and transmitting an abnormality alert signal based on the set of current feature vectors and the selected decision boundary

[0006] Some technical advantages of some embodiments disclosed herein are improved systems and methods to protect one or more wind turbines from cyber-attacks and faults in an automatic and accurate manner.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007]

FIG. 1 is a high-level block diagram of a wind turbine abnormality model creation system that may be provided in accordance with some embodiments.

FIG. 2 is a wind turbine model creation method according to some embodiments.

FIG. 3 is as wind turbine abnormality alert method according to some embodiments.

FIG. 4 illustrates an off-line process in accordance with some embodiments.

FIG. 5 illustrates an example set of monitoring nodes for a wind turbine according to some embodiments.

FIG. 6 is an example associated with a wind turbine engine in accordance with some embodiments.

FIG. 7 is an example of a global threat protection system in accordance with some embodiments when multiple wind turbines (or monitoring nodes) are involved in a system.

FIG. 8 illustrates three dimensions of wind turbine sensor outputs in accordance with some embodiments.

FIG. 9 is a wind turbine abnormality alert system according to some embodiments.

FIGS. 10 through 12 illustrate boundaries and locations of feature vectors for various wind turbine parameters in accordance with some embodiments.

FIG. 13 illustrates a block diagram view of a wind turbine abnormality detection system in accordance with some embodiments.

FIG. 14 is a high-level block diagram of an off-line training process to protect a wind turbine according to some embodiments.

FIG. 15 is a high-level block diagram of an on-line operation process to protect a wind turbine in accordance with some embodiments.

FIG. 16 is a block diagram of a more detailed off-line training process to protect a wind turbine according to some embodiments.

FIG. 17 is a block diagram of a more detailed on-line operation process to protect a wind turbine in accordance with some embodiments.

FIG. 18 is a wind turbine asset protection boundary decision creation display in accordance with some embodiments.

FIG. 19 is a wind turbine model apparatus according to some embodiments.

FIG. 20 is a portion of a tabular decision boundary database in accordance with some embodiments.

FIG. 21 is one example of evaluation results for a wind turbine protection system according to some embodiments.

DETAILED DESCRIPTION

[0008] In the following detailed description, numerous specific details are set forth in order to provide a thorough understanding of embodiments. However, it will be understood by those of ordinary skill in the art that the embodiments may be practiced without these specific details. In other instances, well-known methods, procedures, components and circuits have not been described in detail so as not to obscure the embodiments.

[0009] One or more specific embodiments of the present invention will be described below. In an effort to provide a concise description of these embodiments, all features of an actual implementation may not be described in the specification. It should be appreciated that in the development of any such actual implementation, as in any engineering or design project, numerous implementation-specific decisions must be made to achieve the developers' specific goals, such as compliance with system-related and business-related constraints, which may vary from one implementation to another. Moreover, it should be appreciated that such a development effort might be complex and time consuming, but would nevertheless be a routine undertaking of design, fabrication, and manufacture for those of ordinary skill having the benefit of this disclosure.

[0010] Some embodiments described herein may provide a system and method for reliable cyber-threat detection for a wind turbine. Note that wind farms are usually in remote geographical locations and critical failure-prone components (such as pitch systems) may be located at hard-to-reach locations in the turbine (and, as a result, the cost of maintenance associated with failures that are not detected quickly could be substantial). To avoid such a result, given wind information (*e.g.*, wind speed and/or acceleration) along with a plurality of monitoring nodes values over time (*e.g.*, sensor, actuator, and/or controller nodes), a set of abnormality detection decision boundaries may be built in feature space. The term "feature" may refer to, for example, mathematical characterizations of data. Examples of features as applied to data might include the maximum and minimum, mean, standard deviation, variance, settling time, Fast Fourier Transform ("FFT") spectral components, linear and non-linear principal components, independent components, sparse coding, deep learning, *etc.* The type and number of features for each monitoring node might be optimized using domain-knowledge and/or a feature discovery process. The features may be, for example, calculated over a sliding window with consecutive samples of specified duration from time series data. The length of the window and the duration of overlap for each batch may be determined from domain knowledge and an inspection of the data or using batch processing. The features may be computed at the local level (associated with each monitoring node) and/or the global level (associated with all the monitoring nodes, *i.e.*, an entire wind turbine or a farm of wind turbines). The time-domain values of the nodes or their extracted features may be, according to some embodiments, normalized for better numerical conditioning.

[0011] FIG. 1 is a high-level architecture of a system 100 in accordance with some embodiments. The system 100 may include monitoring node sensors 110 MN_i through MN_N , a wind monitoring node sensor 112 W , a "normal space" data source 120, and an "abnormal space" data source 130. The normal space data source 120 might store, for each of the plurality of monitoring nodes 110, a series of normal values over time that represent normal operation of a wind turbine (*e.g.*, generated by a model or collected from actual sensor data as illustrated by the dashed line in FIG. 1). The abnormal space data source 130 might store, for each of the monitoring nodes 110, a series of attacked values that represent an abnormal operation of the wind turbine (*e.g.*, when the system is experiencing one or more cyber-attacks and/or faults).

[0012] Information from the normal space data source 120 and the abnormal space data source 130 may be provided to an abnormality detection model creation computer 160 that uses this data to create a set of decision boundaries (that is, boundaries that separates normal behavior from abnormal behavior in various wind conditions). The decision boundary may then be used by an abnormality detection computer 150 executing an abnormality detection model 155. The abnormality detection model 155 may, for example, monitor streams of data from the monitoring nodes 110, 112 comprising data from sensor nodes, actuator nodes, and/or any other critical monitoring nodes (e.g., sensor nodes MN_i through MN_N and W) and automatically output an abnormality alert signal to one or more remote monitoring devices 170 when appropriate (e.g., for display to an operator). As used herein, the term "automatically" may refer to, for example, actions that can be performed with little or no human intervention. According to some embodiments, information about detected abnormalities may be transmitted back to a wind turbine control system.

[0013] As used herein, devices, including those associated with the system 100 and any other device described herein, may exchange information via any communication network which may be one or more of a Local Area Network ("LAN"), a Metropolitan Area Network ("MAN"), a Wide Area Network ("WAN"), a proprietary network, a Public Switched Telephone Network ("PSTN"), a Wireless Application Protocol ("WAP") network, a Bluetooth network, a wireless LAN network, and/or an Internet Protocol ("IP") network such as the Internet, an intranet, or an extranet. Note that any devices described herein may communicate via one or more such communication networks.

[0014] The abnormality detection model creation computer 160 may store information into and/or retrieve information from various data stores, such as the normal space data source 120 and/or the abnormal space data source 130. The various data sources may be locally stored or reside remote from the abnormality detection model creation computer 160. Although a single abnormality detection model creation computer 160 is shown in FIG. 1, any number of such devices may be included. Moreover, various devices described herein might be combined according to embodiments of the present invention. For example, in some embodiments, the abnormality detection model creation computer 160 and one or more data sources 120, 130 might comprise a single apparatus. The abnormality detection model creation computer 160 functions may be performed by a constellation of networked apparatuses, in a distributed processing or cloud-based architecture.

[0015] A user may access the system 100 via one of the monitoring devices 170 (e.g., a Personal Computer ("PC"), tablet, or smartphone) to view information about and/or manage abnormality information in accordance with any of the embodiments described herein. In some cases, an interactive graphical display interface may let a user define and/or adjust certain parameters (e.g., abnormality detection trigger levels) and/or provide or receive automatically generated recommendations or results from the abnormality detection model creation computer 160 and/or abnormality detection computer 150.

[0016] For example, FIG. 2 illustrates a model creation method that might be performed by

some or all of the elements of the system 100 described with respect to FIG. 1. The flow charts described herein do not imply a fixed order to the steps, and embodiments of the present invention may be practiced in any order that is practicable. Note that any of the methods described herein may be performed by hardware, software, or any combination of these approaches. For example, a computer-readable storage medium may store thereon instructions that when executed by a machine result in performance according to any of the embodiments described herein.

[0017] At S210, the system may retrieve, for each of a plurality of monitoring nodes, a series of normal values over time that represent normal operation of a wind turbine and a set of normal feature vectors may be generated. Similarly, at S220 the system may retrieve, for each of the plurality of monitoring nodes, a series of abnormal values over time that represent an abnormal operation of the wind turbine and a set of abnormal feature vectors may be generated. The series of normal and/or abnormal values might be obtained, for example, by running Design of Experiments ("DoE") on a wind turbine. At S230, a set of decision boundaries may be automatically calculated for an abnormality detection model based on wind information (e.g., wind speed and/or acceleration), the set of normal feature vectors, and the set of abnormal feature vectors. According to some embodiments, each decision boundary might be associated with a line, a hyperplane, and/or a non-linear boundary separating normal space from abnormal space. Moreover, a decision boundary might comprise a multi-class decision boundary separating normal space, attacked space, and degraded operation space (e.g., associated with a wind turbine fault). In addition, note that the abnormality detection model might be associated with feature mapping functions and/or feature parameters.

[0018] The decision boundaries can then be used to detect cyber-attacks and/or faults associated with a wind turbine. For example, FIG. 3 is an abnormality alert method according to some embodiments. At S310, the system may receive, from a plurality of monitoring nodes, a series of current values over time that represent a current operation of a wind turbine. At S320, an abnormality detection platform computer may then generate, based on the received series of current values, a set of current feature vectors.

[0019] At S330, the system may determine current wind information. For example, the system may determine a current wind speed, a current wind acceleration, or other type of wind information (e.g., identifying a wind gust). At S340, the system selects an appropriate decision boundary based on the current wind information. At S350, an abnormality detection model may be executed to transmit an abnormality alert signal based on the set of current feature vectors and the selected decision boundary when appropriate (e.g., when a cyber-attack or fault is detected). According to some embodiments, one or more response actions may be performed when an abnormality alert signal is transmitted. For example, the system might automatically shut down all or a portion of the wind turbine (e.g., to let the detected potential cyber-attack or fault be further investigated). As other examples, one or more parameters might be automatically modified, a software application might be automatically triggered to capture data and/or isolate possible causes, etc.

[0020] Some embodiments described herein may take advantage of the physics of a wind turbine control system by learning *a priori* from tuned high-fidelity equipment models and/or actual "on the job" data to detect single or multiple simultaneous adversarial threats to the system or naturally occurring faults. Moreover, according to some embodiments, all monitoring node data may be converted to features using advanced feature-based methods, and the real-time operation of the wind turbine control system may be monitoring in substantially real-time. Abnormalities may be detected by classifying the monitored data as being "normal" or disrupted (or degraded). Each decision boundary may be constructed using dynamic models and may help enable early detection of vulnerabilities (and potentially avert catastrophic failures) allowing an operator to restore the wind turbine control system to normal operation in a timely fashion.

[0021] Note that an appropriate set of multi-dimensional feature vectors, which may be extracted automatically (*e.g.*, via an algorithm) and/or be manually input, might comprise a good predictor of measured data in a low dimensional vector space. According to some embodiments, appropriate decision boundaries may be constructed in a multi-dimensional space using a data set which is obtained via scientific principles associated with DoE techniques. Moreover, multiple algorithmic methods (*e.g.*, support vector machines or machine learning techniques) may be used to generate decision boundaries. Since boundaries may be driven by measured data (or data generated from high-fidelity models), defined boundary margins may help to create a threat zone in a multi-dimensional feature space. Moreover, the margins may be dynamic in nature and adapted based on a transient or steady state model of the equipment and/or be obtained while operating the system as in self-learning systems from incoming data stream. According to some embodiments, a training method may be used for supervised learning to teach decision boundaries. This type of supervised learning may take into account on operator's knowledge about system operation (*e.g.*, the differences between normal and abnormal operation).

[0022] FIG. 4 illustrates an off-line boundary creation process 400 in accordance with some embodiments. Information about threats, spoofing, attack vectors, vulnerabilities, faults, *etc.* 410 may be provided to models 420 and/or a training and evaluation database 450 created using DoE techniques. The models 420 may, for example, simulate data 430 from monitoring nodes to be used to compute features that are assembled into a feature vector 440 to be stored in the training and evaluation database 450. The data in the training and evaluation database 450 may then be used to compute decision boundaries 460 to distinguish between normal operation and abnormal operation. According to some embodiments, the process 400 may include a prioritization of monitoring nodes and anticipated attack or fault vectors to form one or more data sets to develop decision boundaries. Attack vectors are abnormal values at critical inputs where malicious attacks can be created at the domain level that will make the system go into threatened/abnormal space. In addition, the models 420 may comprise high-fidelity models that can be used to create a data set (*e.g.*, a set that describes abnormal space as "levels of threat conditions in the system versus quantities from the monitoring nodes"). The data 430 from the monitoring nodes might be, for example, quantities that are captured for a length of from 60 to 80 seconds from sensor nodes, actuator nodes, and/or controller nodes

(and a similar data set may be obtained for "levels of normal operating conditions in the system versus quantities from the abnormal nodes"). This process will result in data sets for "abnormal space" and "normal space." The 60 to 80 seconds long quantities may be used to compute features 440 using feature engineering to create feature vectors. These feature vectors along with wind information can then be used to obtain a set of decision boundaries that separate the data sets for abnormal space and normal space (used to detect an anomaly such as a cyber-attack or fault).

[0023] Since faults or attacks might be multi-prong (e.g., multiple attacks might happen at once), DoE experiments may be designed to capture the abnormal space (e.g., using full factorial, Taguchi screening, central composite, and/or Box-Behnken). When models are not available, these DoE methods can also be used to collect data from real-world wind turbines. Experiments may run, for example, using different combinations of simultaneous attacks, faults, and/or wind conditions. Similar experiments may be run to create a data set for the normal operating space. According to some embodiments, the system may detect "degraded" or faulty operation as opposed to a threat or attack. Such decisions may require the use of a data set for a degraded and/or faulty operating space.

[0024] In addition to wind information, systems may analyze and/or monitor various types of wind turbine information. For example, FIG. 5 is an example set 500 of monitoring nodes for a wind turbine according to some embodiments. The set 500 might include sensors 502 that provide information about, for example, an azimuth angle, generator speed, blade pitch angles, generator power, etc. The set 500 may further include actuators 504 including commands, generator torque, yaw angles, etc. In addition, the set 500 may include controllers 506 that are selected based on controller parameters.

[0025] Note that many different types of features may be utilized in accordance with any of the embodiments described herein, including principal components (weights constructed with natural basis sets) and statistical features (e.g., mean, variance, skewness, kurtosis, maximum, minimum values of time series signals, location of maximum and minimum values, independent components, etc.). Other examples include deep learning features (e.g., generated by mining experimental and/or historical data sets) and frequency domain features (e.g., associated with coefficients of Fourier or wavelet transforms). Embodiments may also be associated with time series analysis features, such as cross-correlations, auto-correlations, orders of the autoregressive, moving average model, parameters of the model, derivatives and integrals of signals, rise time, settling time, neural networks, etc. Still other examples include logical features (with semantic abstractions such as "yes" and "no"), geographic/position locations, and interaction features (mathematical combinations of signals from multiple threat nodes and specific locations). Embodiments may incorporate any number of features, with more features allowing the approach to become more accurate as the system learns more about the physical process, faults, and/or threats. According to some embodiments, dissimilar values from monitoring nodes may be normalized to unit-less space, which may allow for a simple way to compare outputs and strength of outputs.

[0026] FIG. 6 is an example 600 associated with a wind turbine in accordance with some embodiments. In particular, the example includes a controller and actuator portion 610 subject to actuator and controller attacks, a wind turbine portion 620 subject to state attacks, and sensors 930 subject to sensor attacks. By way of examples only, the sensors 630 might comprise physical and/or virtual sensors associated with wind speed, temperatures, blade parameters, power levels, *etc.* The actuators might be associated with, for example, generators. By monitoring the information in the wind turbine 620, an anomaly detection platform may be able to detect faults and/or cyber-attacks (*e.g.*, using feature vectors and a decision boundary) that could potentially cause a large amount of damage.

[0027] FIG. 7 is an example of a global wind turbine protection system 700 in accordance with some embodiments when multiple wind turbines are involved in a system. In particular, the system 700 includes three wind turbines (A, B, and C) and batches of values 710 from monitoring nodes are collected for each generated over a period of time (*e.g.*, 60 to 80 seconds). According to some embodiments, the batches of values 710 from monitoring nodes overlap in time. The values 710 from monitoring nodes may, for example, be stored in a matrix 720 arranged by time ($t_1, t_2, \text{etc.}$) and by type of monitoring node ($S_i, S_5, \text{etc.}$). Feature engineering components 730 may use information in each matrix 720 to create a feature vector 740 for each of the three wind turbines (*e.g.*, the feature vector 740 for wind turbine C might include $F_{Sci}, F_{Sc2}, \text{etc.}$). The three feature vectors 740 may then be combined into a single global feature vector 750 for the system 700. Interaction features 760 may be applied (*e.g.*, associated with $A*B*C, A+B+C, \text{etc.}$) and an abnormality detection engine 770 may compare the result with a decision boundary (selected based on wind conditions) and output an abnormality alert signal when appropriate.

[0028] FIG. 8 illustrates 800 three dimensions of wind turbine monitoring node outputs in accordance with some embodiments. In particular, a graph 810 plots monitoring node outputs ("+" in three dimensions, such as dimensions associated with Principal Component Features ("PCF"): $w_1, w_2, \text{and } w_3$). Moreover, the graph 810 includes an indication of a first normal operating space decision boundary 820 associated with a first wind condition. Although a single contiguous boundary 820 is illustrated in FIG. 8, embodiments might be associated with multiple regions. For example, the graph 810 includes an indication of a second normal operating space decision boundary 830 associated with a second wind condition. Note that PCF information may be represented as weights in reduced dimensions. For example, data from each monitoring node may be converted to low dimensional features (*e.g.*, weights). According to some embodiments, monitoring node data is normalized as follows:

$$S_{normalized}(k) = \frac{S_{nominal}(k) - S_{original}(k)}{\bar{S}_{nominal}}$$

where S stands for a threat node quantity at " k " instant of time. Moreover, output may then be expressed as a weighted linear combination of basis functions as follows:

$$S = S_0 + \sum_{j=1}^N w_j \psi_j$$

where S_0 is the average threat node output with all threats, w_j is the j^{th} weight, and ψ_j is the j^{th}

basis vector. According to some embodiments, natural basis vectors are obtained using a covariance of the threat nodes' data matrix. Once the basis vectors are known, weight may be found using the following equation (assuming that the basis sets are orthogonal):

$$w_j = (S - S_0)^T \psi_j$$

Note that weights may be an example of features used in a feature vector.

[0029] Thus, embodiments may enable the passive detection of indications of wind turbine abnormal operations using real-time signals from monitoring nodes. Moreover, the detection framework may allow for the development of tools that facilitate proliferation of the invention to various systems (e.g., different turbine models) in multiple geolocations. According to some embodiments, distributed detection systems enabled by this technology (across multiple types of equipment and systems) will allow for the collection of coordinated data to help detect faults and multi-prong attacks. Note that the feature-based approaches described herein may allow for extended feature vectors and/or incorporate new features into existing vectors as new learnings and alternate sources of data become available. As a result, embodiments may detect a relatively wide range of faults and/or cyber-threats (e.g., stealth, replay, covert, injection attacks, etc.) as the systems learn more about their characteristics. Embodiments may also reduce false positive rates as systems incorporate useful key new features and remove ones that are redundant or less important. Note that the detection systems described herein may provide early warning to wind turbine operators so that an attack may be thwarted (or the effects of the attack may be blunted), reducing damage to equipment.

[0030] According to some embodiments, a system may further localize an origin of an abnormal condition to a particular monitoring node. For example, the localizing may be performed in accordance with a time at which a decision boundary associated with one monitoring node was crossed as compared to a time at which a decision boundary associated with another monitoring node was crossed. According to some embodiments, an indication of the particular monitoring node might be included in an abnormality alert signal.

[0031] Some embodiments of the algorithm may utilize feature-based learning techniques based on high-fidelity physics models and/or machine operation data (which would allow the algorithm to be deployed on any system) to establish a high dimensional decision boundary. As a result, detection may occur with more precision using multiple signals, making the detection more accurate with less false positives. Moreover, embodiments may detect multiple faults or attacks on control signals, and rationalize where the root cause fault or attack originated. For example, the algorithm may decide if a signal is anomalous because of a previous signal attack, or if it is instead independently under attack. This may be accomplished, for example, by monitoring the evolution of the features as well as by accounting for time delays between attacks.

[0032] A cyber-attack detection and localization algorithm may process a real-time wind turbine signal data stream and then compute features (multiple identifiers) which can then be compared to the signal-specific decision boundary. A block diagram of a system 900 utilizing a signal-specific wind turbine abnormality detection and localization algorithm according to some

embodiments is provided in FIG. 9. In particular, a wind turbine 932 provides information to sensors 934 which helps controllers with electronics and processors 936 adjust actuators 938. An abnormality detection system 960 may include one or more high-fidelity physics based models 942 associated with the wind turbine 932 to create normal data 910 and/or abnormal data 920. The normal data 910 and abnormal data 920 may be accessed by a feature discovery component 944 and processed by decision boundary algorithms 946 along with associated wind information while off-line (e.g., not necessarily while the wind turbine 932 is operating). The decision boundary algorithms 946 may generate an abnormality detection model including decision boundaries for various monitoring nodes and/or wind conditions. Each decision boundary may separate two data sets in a high dimensional space which is constructed by running a binary classification algorithm, such as a support vector machine using the normal data 910 and abnormal data 920 for each monitoring node signal (e.g., from the sensors 934, controllers 936, and/or the actuators 938).

[0033] A real-time threat detection platform 950 may receive the boundaries along with streams of data from the monitoring nodes (e.g., including current wind information). The platform 950 may include a feature extraction on each monitoring node element 952 and a normalcy decision 954 with an algorithm to detect attacks and faults in individual signals using signal specific decision boundaries (selected based on current wind conditions), as well rationalize faults and attacks on multiple signals, to declare which signals are being attacked, and which became anomalous due to a previous fault or attack on the system via a localization module 956. An accommodation element 958 may generate outputs 970, such as an abnormality decision indication (e.g., a fault or threat alert signal), a controller action, and/or a list of abnormal monitoring nodes.

[0034] During real-time detection, contiguous batches of control signal data may be processed by the platform 950, normalized and the feature vector extracted. The location of the vector for each signal in high-dimensional feature space may then be compared to a corresponding decision boundary selected based on wind conditions. If it falls within the abnormal region, then abnormal operation may be declared. The algorithm may then make a decision about where the abnormality originally occurred. An abnormality may sometimes be an attack on the actuators 938 and then manifested in the sensor 934 data. Abnormality assessments might be performed in a post decision module (e.g., the localization element 956) to isolate whether the abnormality is related to the sensor, controller, or actuator (e.g., indicating which part of the monitoring node). This may be done by individually monitoring, over time, the location of the feature vector with respect to the hard decision boundary. For example, when a sensor 934 is spoofed, the attacked sensor feature vector will cross the hard decision boundary earlier than the rest of the vectors as described with respect to FIGS. 10 through 12. If a sensor 934 is declared to be anomalous, and a command to the auxiliary equipment is later determined to be anomalous, it may be determined that the original attack, such as signal spoofing, occurred on the sensor 934. Conversely, if the signal to the auxiliary equipment was determined to be anomalous first, and then later manifested in the sensor 934 feedback signal, it may be determined that the signal to the equipment was initially attacked.

[0035] According to some embodiments, it may be detected whether or not a signal is in the normal operating space (or abnormal space) through the use of localized decision boundaries and real-time computation of the specific signal features. Moreover, an algorithm may differentiate between a sensor being attacked as compared to a signal to auxiliary equipment being attacked. The control intermediary parameters and control logical(s) may also be analyzed using similar methods. Note that an algorithm may rationalize signals that become anomalous, and an attack on a signal may then be identified.

[0036] FIG. 10 illustrates 1000 boundaries and feature vectors for various wind turbine monitoring node parameters in accordance with some embodiments. In particular, for each parameter a graph includes a first axis representing value weight 1 ("w1"), a feature 1, and a second axis representing value weight 2 ("w2"), a feature 2. Values for w1 and w2 might be associated with, for example, outputs from a Principal Component Analysis ("PCA") that is performed on the input data. PCA might be one of the features that might be used by the algorithm to characterize the data, but note that other features could be leveraged.

[0037] A graph is provided for azimuth angle 1010, generator speed 1020, generator power 1030, air density 1040, blade pitch angle A 1050, and ambient temperature 1060. Each graph includes a hard boundary 1012 (solid curve), minimum boundary 1016 (dotted curve), and maximum boundary 1014 (dashed curve) and an indication associated with current feature location for each wind turbine monitoring node parameter (illustrated with an "X" on the graph). As illustrated in FIG. 10, the current monitoring node location is between the minimum and maximum boundaries (that is, the "X" is between the dotted and dashed lines). As a result, the system may determine that the operation of the wind turbine control system is normal (and no fault or threat is being detected indicating that the system is currently under attack).

[0038] FIG. 11 illustrates 1100 subsequent boundaries and feature vectors for these parameters. Consider, for example, a feature vector movement 1112 for the azimuth angle. Even though feature vector 1112 has moved, it is still within the maximum and minimum boundaries and, as a result, normal operation of that monitoring node may be determined. This is the case for the first five graphs in FIG. 11. In this example, a feature vector movement 1162 for the wind turbine ambient temperature has exceeded with maximum boundary and, as a result, abnormal operation of that monitoring node may be determined. For example, a fault or threat may exist for the ambient temperature scale factor, which is a corrective value. The result is that the feature for the intermediary monitoring node signal feature vector illustrated in FIG. 11 moves 1162 such that it is anomalous. The algorithm detects this fault or cyber-attack, and two parallel actions might be initiated. One action may be post processing of the signal to discover what was attacked, in this case if the system has been monitoring each temperature sensor, it may conclude that none of them are currently abnormal. Therefore, it may be determined that something used to calculate this feature was attacked. The other action may be to continually monitor and detect additional attacks. Such an approach may facilitate a detection of multiple signal attacks.

[0039] Given the example of FIG. 11, assume that the ambient temperature signal was

attacked. This may cause the system to respond in such a way so as to put other signals into an abnormal state. This is illustrated 1200 in FIG. 12, where the attack has already been detected and now other signals shown to be abnormal. In particular, feature movement for the azimuth angle 1212, generator speed 1222, generator power 1232, and air density 1242 have all become abnormal (joining the feature vector for the wind turbine ambient temperature 1262). Note that the feature vector for blade pitch angle A did not become abnormal. In order to decide whether or not these signals 1212, 1222, 1232, 1242 are truly currently under attack, a historical batch with pertinent feature vector information may be kept for some duration of time. Then when an attack is detected on another signal, this batch is examined, and the time at which the confirmed attack on wind turbine ambient temperature as well as several subsequent elements is analyzed.

[0040] Note that one signal rationalization might be associated with a system time delay. That is, after a wind turbine sensor is attacked there might be a period of time before the system returns to a steady state. After this delay, any signal that becomes anomalous might be due to an attack as opposed to the system responding.

[0041] The current methods for detecting abnormal conditions in monitoring nodes are limited to Fault Detection Isolation and Accommodation ("FDIA"), which itself is very limited. The cyber-attack detection and localization algorithms described herein can not only detect abnormal signals of wind turbine sensors, but can also detect signals sent to auxiliary equipment, control intermediary parameters and/or control logical(s). The algorithm can also understand multiple signal attacks. One challenge with correctly identifying a cyber-attack threat is that it may occur with multiple sensors being impacted by malware. According to some embodiments, an algorithm may identify in real-time that an attack has occurred, which sensor(s) are impacted, and declare a fault response. To achieve such a result, the detailed physical response of the system must be known to create acceptable decision boundaries. This might be accomplished, for example, by constructing data sets for normal and abnormal regions by running DoE experiments on high-fidelity models. A data set for each sensor might comprise a feature vector for given threat values (*e.g.*, temperature, blade angles, *etc.*). Full factorial, Taguchi screening, central composite and Box-Behnken are some of the known design methodologies used to create the attack space. When models are not available, these DoE methods are also used to collect data from real-world wind turbines. Experiments may be run at different combinations of simultaneous attacks. In some embodiments, the system may detect degraded/faulty operation as opposed to a cyber-attack. Such decisions might utilize a data set associated with a degraded/faulty operating space. At the end of this process, the system may create data sets such as "attack v/s normal" and "degraded v/s normal" for use while constructing decision boundaries. Further note that a decision boundary may be created for each signal using data sets in feature space. Various classification methods may be used to compute decision boundaries. For example, binary linear and non-linear supervised classifiers are examples of methods that could be used to obtain a decision boundary.

[0042] Note that multiple vector properties might be examined, and the information described with respect to FIGS. 10 through 12 may be processed to determine if the signal had been

trending in a specific direction as the attack was detected (or if it had just been moving due to noise). Had the signal been uniformly trending as the attack took place and afterward, then this signal is a response to the original attack and not an independent attack.

[0043] According to some embodiments, the system may localize or otherwise analyze an origin of the threat to a particular wind turbine monitoring node. For example, the localizing may be performed in accordance with a time at which a decision boundary associated with one monitoring node was crossed as compared to a time at which a decision boundary associated with another monitoring node was crossed. According to some embodiments, an indication of the particular monitoring node might be included in the threat alert signal.

[0044] Some embodiments described herein may take advantage of the physics of a wind turbine by learning a priori from tuned high-fidelity equipment models and/or actual "on the job" data to detect single or multiple simultaneous adversarial threats to the wind turbine. Moreover, according to some embodiments, all monitoring node data may be converted to features using advanced feature-based methods, and the real-time operation of the wind turbine may be monitoring in substantially real-time. Abnormalities may be detected by classifying the monitored data as being "normal" or disrupted (or degraded). This decision boundary may be constructed using dynamic models and may help to enable early detection of vulnerabilities (and potentially avert catastrophic failures) allowing an operator to restore the wind turbine to normal operation in a timely fashion.

[0045] Thus, some embodiments may provide an advanced anomaly detection algorithm to detect faults and cyber-attacks on, for example, key wind turbine control sensors. The algorithm may identify which signals(s) are being attacked using control signal-specific decision boundaries and may inform a wind turbine to take accommodative actions. In particular, a detection and localization algorithm might detect whether a sensor, auxiliary equipment input signal, control intermediary parameter, or control logical are in a normal or anomalous state. Some examples of wind turbine monitoring nodes that might be analyzed include: critical control sensors; control system intermediary parameters; auxiliary equipment input signals; and/or logical commands to controller.

[0046] A cyber-attack detection and localization algorithm may process a real-time wind turbine signal data stream and then compute features (multiple identifiers) which can then be compared to the sensor specific decision boundary. In some embodiments, generating features may involve simply performing an identity transform. That is, the original signal might be used as it is.

[0047] Note that a wind turbine may be associated with non-linear operations over a range of operating parameters. As a result, data variations can be substantial and determining when a fault or cyber threat is present based on operation of the control system may be difficult. FIG. 13 illustrates a block diagram view of an abnormality detection system 1300 in accordance with some embodiments. In particular, the system 1300 illustrates a wind turbine 1310 (e.g., associated with unit controllers) that transmits information about temperature, airflow, *etc.* to a

normalization function 1320. The wind turbine 1310 may also transmit information to mode processing 1330 (e.g., a reported mode of operation and current wind information) and feature processing 1340 (e.g., unit data). The normalization function 1320 might, for example, transmit a normalized monitoring node signal to feature processing 1340. Post processing 1350 may receive information from feature processing 1340 and transmit data to decision processing 1370 (which can automatically create a fault or cyber-attack warning based at least in part on data sets received from a normal/attack data sets storage unit 1360).

[0048] In this way, embodiments may provide a method and system for anomaly (a cyber-attack or a fault) detection in wind turbine that is fast and reliable at a physical layer using a plurality of monitoring nodes. The system may provide coverage for substantially all of the wind turbine operating space in both below-rated and above-rated regions as well as during gust events and/or power curtailments as will now be described. FIG. 14 is a high-level block diagram of an "off-line" training process 1400 (e.g., that might be executed before a wind turbine is operating) to protect a wind turbine according to some embodiments. After time-series of wind turbine values 1410 are received from monitoring nodes (e.g., based on a physics model or historical data), local features may be extracted for each node 1420, 1422 (for a total of P nodes). Global features may also be extracted 1430 and a dimensionality reduction process 1440 may be executed. Note that global features may include interactive features that involve two or more monitoring nodes. Also note that the global feature extraction 1430 might also utilize information directly from the time-series data 1410 (as illustrated by the dashed arrow of FIG. 14). A classification and decision boundaries computation 1450 may then be performed using wind information 1460 to create a first decision boundary 1482 (e.g., associated with a first wind condition) and a second decision boundary 1486 (e.g., associated with a second wind condition).

[0049] FIG. 15 is a high-level block diagram of an "on-line" operation process 1500 (e.g., that might be executed while a wind turbine is generating power) to protect a wind turbine in accordance with some embodiments. After time-series of current wind turbine values 1510 are received from monitoring nodes, local features may be extracted for each node 1520, 1522 (for a total of P nodes). Global features may also be extracted 1530 and a dimensionality reduction process 1540 may be executed. Note that the global feature extraction 1530 might also utilize information directly from the time-series data 1510 (as illustrated by the dashed arrow of FIG. 15). An abnormality decision 1550 may then be performed using wind information 1560 to select an appropriate decision boundary and an output may indicate if the system status is normal 1582 or abnormal 1584 (e.g., currently being cyber-attacked or experiencing a fault).

[0050] FIG. 16 is a block diagram of a more detailed off-line training process 1600 to protect a wind turbine according to some embodiments. Embodiments may receive time-series data from a collection of monitoring nodes 1610 (sensor/actuators/controller nodes) and extracts feature from the time-series data for each monitoring node 1620, 1622. The term "feature" may refer to, for example, mathematical characterizations of data. Examples of features as applied to data might include the statistical features, such as maximum and minimum, mean,

standard deviation, variance, or higher-order statistics, rise time, settling time, Fast Fourier Transform ("FFT") spectral components, linear and non-linear principal components, independent components, sparse coding, deep learning, *etc.* The structural natural frequencies of the wind turbine can also be considered as features. This may include, for example, tower natural frequencies (e.g., 1st and 2nd harmonic frequencies of the tower fore-aft and side-to-side vibrations) and blade natural frequencies (e.g., 1st and 2nd harmonic frequencies of the blade flap-wise, edgewise, and/or torsional vibrations). The natural frequencies may be computed through spectral analysis methods such as FFT or a Eigensystem Realization Algorithm ("ERA"). The type and number of features for each monitoring node might be optimized using domain-knowledge, feature engineering, or Receiver Operator Curve ("ROC") statistics. The local features for each monitoring node may be, according to some embodiments, stacked to create a global feature vector 1630. The global feature vector may also contain interactive feature involving two or more monitoring nodes, e.g. cross-correlation between two nodes as illustrated by the dashed arrow in FIG. 16. The interactive features may be selected through correlation analysis between the nodes. The rate of change of the time-domain signals or the rate of change of the features may also be used as additional features. The features may be normalized using some base values such as those in at-rated wind speed (the wind speed at which wind turbine reaches its rated power) condition. The dimension of the global feature vector can then be further reduced 1640 using any dimensionality reduction technic such as PCA. The features may be, for example, calculated over a sliding window of the signal time series. The length of the window and the duration of slide may be determined from domain knowledge and inspection of the data or using batch processing.

[0051] During a training phase, normal, attack, and fault data sets may be generated and their corresponding features may be extracted. The normal data may be collected during different operational modes of the asset and different ambient conditions. The training data may include node time-series during different wind speeds covering the above-rated and below-rated with different turbulence intensity levels, and different horizontal and vertical wind shear factors, as well as different wind directions, ambient temperatures and air densities. The training data may also include gust scenarios, such as 1-year and 50-year Extreme Operating Gust ("EOG"), Extreme Coherent Gust ("ECG"), and/or Extreme gust with Change of Direction ("ECD") occurring at different wind speeds as described in the International Electrotechnical Commission ("IEC") 61400 standard. The attack data set may comprise different attack types of different severity injected to single or multiple monitoring nodes, and the fault data set may consist of data associated with different failure modes of each monitoring node. Note that single or multiple node anomalies may be simulated for each training or evaluation case. A DoE may be used to optimize the number of simulation runs and accommodate multiple variation factors. The data sets may be generated, for example, using physics-based or data-driven models and/or collected from the field.

[0052] According to some embodiments, a wind turbine control system may contain a wind speed estimate 1660 either as a physical sensor (such as an anemometer or LIDAR) or a virtual sensor (such as a model-based observer). Note that most wind turbine control systems

have such a wind speed estimator. The wind speed estimate 1660 and its derivative (i.e., wind acceleration estimate 1670) may also be used as part of the feature vector (as illustrated by dashed arrows in FIG. 16). Note that ROC statistics may be used at the global level to optimize a selected number of principal components.

[0053] According to some embodiments, the generated data sets may be used by a classification decision boundaries computation 1650 to compute four classification decision boundaries. A first decision boundary may be computed for above-rated Normal Power Production ("NPP") 1688. This boundary 1688 may be computed using normal and abnormal data during the above-rate wind speeds and normal power production (i.e., no gust events). The above-rated region may be defined as the wind speeds from the rated wind speed, i.e. the wind speed at which the wind turbine reaches its rated power (usually between 10-12 m/s based on the type of the turbine) to the cut-off wind speed (i.e., the maximum speed at which the turbine can still operate safely, usually, between 23-25 m/s, based on the type of the wind turbine).

[0054] A second decision boundary may be computed for below-rated normal power production 1686. This boundary 1686 may be computed using normal and abnormal data during the below-rate wind speeds and normal power production (i.e., no gust events). The below-rated region is from the cut-in wind speed or the minimum wind speed at which the wind turbine can produce power (usually 3-5 m/s based on the type of the wind turbine) to the rated wind speed.

[0055] A third decision boundary may be computed for gust wind conditions 1682. This boundary 1682 may be computed using normal and abnormal data during gust events with different types and intensities. The gust may happen at any speed of the full operational range of the turbine from cut-in to cut-off speeds. Note that a single boundary may be used for gust events occurring both at below-rated and above-rated regions, since a gust starting at below-rated speed may well span into the above-rated speeds as well.

[0056] A fourth decision boundary may be computed for power curtailment 1684. This boundary 1684 may be computed using normal and abnormal data during power curtailment. Note that within certain power grid connection agreements, the output of a wind turbine may be curtailed at certain times according to power grid scheduling and planning demand (which will result in a reduction of energy production). This factor may also include the time taken for the wind turbine to become fully operational following grid curtailment. In addition, in certain jurisdictions, there may be requirements to curtail wind turbines during specific meteorological conditions to meet defined noise emission or shadow flicker criteria at nearby dwellings. Since the curtailment may happen both during below-rated and above-rated conditions, similar to the gust decision boundary 1682, a single decision boundary may be used for power curtailment 1684 in both of those regions.

[0057] Thus, a system may use separate boundaries for gust, normal power production, and curtailment to increase accuracy and reduce false alarms associated with abnormality

detection. The system may also use two separate boundaries for below-rated and above-rated conditions since the wind turbine control strategies may be substantially different in these two regions. While the below-rated control may use a constant blade pitch commands to optimize the power production, the above-rated control may use variable blade pitch commands to regulate the power and minimize structural loads. The binary classifiers can be trained using binary classification methods such as Support Vector Machines ("SVM"), an Extreme Learning Machine ("ELM") neural network, kernel-ELM (ELM with a nonlinear kernel, such as a Gaussian or wavelet kernel), decision trees, deep neural networks, *etc.* The computed classification decision boundaries may then be stored in the system to be used during real-time operation of the wind turbine.

[0058] FIG. 17 is a block diagram of a more detailed on-line operation process 1700 to protect a wind turbine in accordance with some embodiments. As before, current time-series monitoring node data 1710 may be used to extract local features 1720, 1722 and global features 1730. Dimensionality reduction 1750 may then be performed and an abnormality decision 1750 may be generated. A wind speed estimate 1760 and/or wind speed acceleration estimate 1770 may be used to determine if a gust condition currently exists 1772, such as by applying a threshold comparison to the wind acceleration estimate 1770 (and, if so, that decision boundary may be selected). If the wind turbine control system contains a gust detector, the output of that detector may be used instead. The wind speed estimate 1772 may also be used to determine if the wind turbine is in below-rated or above-rated mode 1774. This information is sent to the decision boundary selector 1750, which selects the appropriate boundary to be used for the anomaly decision. The decision boundary selector 1750 may also receive a status of power curtailment 1776 from the wind turbine control system. When the appropriate boundary is selected, the (dimensionality-reduced) global feature vector might be compared with that boundary which results in the global system status as being normal 1782 or abnormal 1784. The system status might be, for example, communicated to an operator or may be sent to automatically trigger some fault-tolerant or attack-tolerant control remedies.

[0059] Thus, embodiments may provide substantially rapid, robust, and accurate solution for anomaly detection in wind turbines even though wind may be both highly transient and hard to predict. Moreover, fast and reliable anomaly detection for wind turbines may be provided such that accurate turbine status is available both during natural faults (and hence performing as a fault detector) and cyber-attacks. This may increase wind turbine reliability and availability, provide status monitoring capability and anomaly analytics, *etc.*

[0060] By way of example, embodiments might be associated with a wind turbine in a high-fidelity simulation environment. The system might have 16 monitoring nodes (12 sensor nodes and 4 actuator nodes). A normal data set may be generated by simulating the wind turbine in a closed loop control system during various operation conditions using factors, such as wind speed, turbulence intensity level, wind direction, vertical and horizontal shear factors, *etc.*

[0061] A normal data set might include 100,000 data points (each data point containing the values of all 16 monitoring nodes) in the time space. By way of example only, 5 statistical

features may be selected for for each monitoring node. The features may be computed, according to some embodiments, over a 25 second sliding window. This will result in 4,000 normal events in the feature space. The global feature vector might be the original of size 80 which is then reduced to 25 using PCA analysis.

[0062] An abnormal training data set may be generated by simulating the wind turbine in a closed loop control system during various operation conditions, using the same operational variation factors as described above. An attack vector may be selected to be the critical sensors and actuators, such as three sensors (rotor speed, power, tower top fore-aft acceleration) and three actuators (pitch commands). The simulated attacks may be injected into the wind turbine using Pseudo-Random Binary Sequence ("PRBS") bias attacks on the attack vector. Note that this might be associated with a multi-prong attack simulation. The magnitude of each attack may be selected as a percentage of the current value of the corresponding node. The abnormal data set may consist of 15,000 data points (each data point containing the values of all 16 monitoring nodes) in the time space, which results in 600 abnormal events in the feature space.

[0063] A classification decision boundary may then be trained for using kernel-ELM with Gaussian kernel. FIG. 18 is an example of a wind turbine protection display 1800 that might be used, for example, to provide a graphical depiction 1810 to an operator and/or to provide an interactive interface allowing an operator to adjust system components as appropriate. Selection of an item on the display 1800 (e.g., via a touchscreen or computer mouse pointer 1820) may let the operator see more information about that particular item and/or adjust operation of that item (e.g., by altering how global features are extracted).

[0064] Note that the embodiments described herein may be implemented using any number of different hardware configurations. For example, FIG. 19 is a block diagram of a wind turbine protection platform 1900 that may be, for example, associated with the system 100 of FIG. 1 and/or any other system described herein. The wind turbine protection platform 1900 comprises a processor 1910, such as one or more commercially available Central Processing Units ("CPUs") in the form of one-chip microprocessors, coupled to a communication device 1920 configured to communicate via a communication network (not shown in FIG. 19). The communication device 1960 may be used to communicate, for example, with one or more remote monitoring nodes, user platforms, digital twins, etc. The wind turbine protection platform 1900 further includes an input device 1940 (e.g., a computer mouse and/or keyboard to input wind turbine and/or predictive modeling information) and/or an output device 1950 (e.g., a computer monitor to render a display, provide alerts, transmit recommendations, and/or create reports). According to some embodiments, a mobile device, monitoring physical system, and/or PC may be used to exchange information with the wind turbine protection platform 1900.

[0065] The processor 1910 also communicates with a storage device 1930. The storage device 1930 may comprise any appropriate information storage device, including combinations of magnetic storage devices (e.g., a hard disk drive), optical storage devices, mobile

telephones, and/or semiconductor memory devices. The storage device 1930 stores a program 1912 and/or wind turbine protection model 1914 for controlling the processor 1910. The processor 1910 performs instructions of the programs 1912, 1914, and thereby operates in accordance with any of the embodiments described herein. For example, the processor 1910 may receive a series of current monitoring node values and generate a set of current feature vectors. The processor 1910 may also access an abnormality detection model having a plurality of decision boundaries created using wind information (e.g., wind speed and/or acceleration) along with at least one of a set of normal feature vectors and a set of abnormal feature vectors. The processor 1910 may then select one of the decision boundaries based on current wind information associated with the wind turbine and execute the abnormality detection model and transmit an abnormality alert signal based on the set of current feature vectors and the selected decision boundary

[0066] The programs 1912, 1914 may be stored in a compressed, uncompiled and/or encrypted format. The programs 1912, 1914 may furthermore include other program elements, such as an operating system, clipboard application, a database management system, and/or device drivers used by the processor 1910 to interface with peripheral devices.

[0067] As used herein, information may be "received" by or "transmitted" to, for example: (i) the wind turbine protection platform 1900 from another device; or (ii) a software application or module within the wind turbine protection platform 1900 from another software application, module, or any other source.

[0068] In some embodiments (such as the one shown in FIG. 19), the storage device 1930 further stores a decision boundary database 2000. An example of a database that may be used in connection with the wind turbine protection platform 1900 will now be described in detail with respect to FIG. 21. Note that the database described herein is only one example, and additional and/or different information may be stored therein. Moreover, various databases might be split or combined in accordance with any of the embodiments described herein.

[0069] Referring to FIG. 20, a table is shown that represents the decision boundary database 2000 that may be stored at the wind turbine protection platform 1900 according to some embodiments. The table may include, for example, entries identifying decision boundaries associated with wind turbines to be protected. The table may also define fields 2002, 2004, 2006, 2008, 2010, 2012 for each of the entries. The fields 2002, 2004, 2006, 2008, 2010, 2012 may, according to some embodiments, specify: a decision boundary identifier 2002, a decision boundary description 2004, input requirements 2006, a boundary definition 2008, a currently active indication 2010, a status 2012, and a cause 2014. The decision boundary database 2000 may be created and updated, for example, off-line when a new physical system is monitored or modeled.

[0070] The decision boundary identifier 2002 and description 2004 may define a particular normal/abnormal boundary for a wind turbine that will be protected. The input requirements 2006 define when the boundary will be selected for use (e.g., based on current wind

information). The boundary definition 2008 might comprise, for example, an equation separating normal space from abnormal space. The currently active indication 2010 might indicate if that particular decision boundary is currently being used to detect abnormal operation of the wind turbine. If it is currently being used, the status 2012 might indicate if the wind turbine is currently "normal" or "abnormal" and (in the case of abnormal operations), why the system is operating abnormally (e.g., due to a fault or a cyber-attack on one or more monitoring nodes).

[0071] FIG. 21 is one example of evaluation results 2100 for a wind turbine protection system in accordance with the simulation described with respect to FIGS. 18 through 21. As illustrated by the predicted class 2110 and the actual class 2120 of the wind turbine, all the abnormal events were detected (that is, there were 0% false negatives) while only 1.36% of the normal cases were misclassified as abnormal (that is, there was a 1.36% false positive rate).

[0072] The following illustrates various additional embodiments of the invention. These do not constitute a definition of all possible embodiments, and those skilled in the art will understand that the present invention is applicable to many other embodiments. Further, although the following embodiments are briefly described for clarity, those skilled in the art will understand how to make any changes, if necessary, to the above-described apparatus and methods to accommodate these and other embodiments and applications.

[0073] Although specific hardware and data configurations have been described herein, note that any number of other configurations may be provided in accordance with embodiments of the present invention (e.g., some of the information associated with the databases described herein may be combined or stored in external systems). For example, although some embodiments are focused on wind turbines, any of the embodiments described herein could be applied to other types of wind turbines including power grids, dams, locomotives, airplanes, and autonomous vehicles (including automobiles, trucks, drones, submarines, *etc.*).

REFERENCES CITED IN THE DESCRIPTION

Cited references

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- EP3239884A1 [0001]
- KR20140109132A [0001]

KRAV

1. System (100) til at beskytte en vindmølle (620), omfattende:

- 5 en flerhed af overvågningsknodepunkter (110, 112) der hver genererer en række aktuelle overvågningsknodepunktverdier over tid, som repræsenterer en aktuel drift af vindmøllen (620); og
- en abnormitetsdetektionscomputerplatform (150) koblet til flerheden af overvågningsknodepunkter (110, 112), til at:
- 10 modtage rækken af aktuelle overvågningsknodepunktverdier og generere et sæt aktuelle funktionsvektorer (eng: feature vectors),
- tilgå en abnormitetsdetektionsmodel (155) der har en flerhed af beslutningsgrænser (eng: decision boundaries) dannet ved hjælp af vindinformation (1560) sammen
- 15 med mindst én af: et sæt normale funktionsvektorer (eng: feature vectors) og et sæt abnorme funktionsvektorer,
- vælge en af beslutningsgrænserne baseret på aktuelle vindoplysninger (1560) der er knyttet til vindmøllen (620), og
- udføre abnormitetsdetektionsmodellen (155) og transmittere et
- 20 abnormitetsalarmsignal baseret på sættet af aktuelle funktionsvektorer og den valgte beslutningsgrænse;
- hvor systemet er konfigureret til at hente, for hver af flerheden af overvågningsknodepunkter, en række abnorme verdier over tid, som repræsenterer
- 25 angrebet drift af vindmøllen, og til at generere sættet af abnorme funktionsvektorer der er et sæt af angrebsfunktionsvektorer (eng: attack feature vectors).

2. System (100) ifølge krav 1, hvor vindinformationen (1560) er forbundet med mindst én af: en vindhastighed og en vindacceleration.

30

3. System (100) ifølge et hvilket som helst af de foregående krav, hvor abnormitetsdetektionsmodellen (155) indbefatter en under-normal-effektproduktionsbeslutningsgrænse (1686) og en over-normal-effektproduktionsbeslutningsgrænse (1688), og hvor nævnte valg i det mindste delvist er baseret

35 på en aktuel vindhastighed og en tærskelværdi forbundet med at vindmøllen (620) når sin

nominelle effekt.

4. System 100 ifølge krav 3, hvor abnormitetsdetektionsmodellen 155 yderligere inkluderer en vindstødsbeslutningsgrænse 1682.

5

5. System (100) ifølge krav 4, hvor vindstødsbeslutningsgrænsen (1682) er forbundet med mindst én af: et ekstremt driftsvindstød; et ekstremt, kohærent vindstød; og et ekstremt vindstød med retningsændring.

10

6. System (100) ifølge krav 4, hvor abnormitetsdetektionsmodellen (155) yderligere indbefatter en effektbegrænsningsbeslutningsgrænse (1684).

7. System (100) ifølge et hvilket som helst af de foregående krav, hvor mindst én beslutningsgrænse er baseret på mindst én global funktion (eng: global feature) (750).

15

8. System (100) ifølge et hvilket som helst af de foregående krav, yderligere omfattende:

en normalrumsdatakilde (120), der for hver af flerheden af overvågningsknudepunkter (110, 112) lagrer en række normale overvågningsknudepunktsværdier over tid, som repræsenterer normal drift af vindmøllen (620);

20

en abnormitetsrumsdatakilde (130), der for hver af flerheden af overvågningsknudepunkter (110, 112) lagrer en række abnorme overvågningsknudepunktsværdier over tid, som repræsenterer en abnorm drift af vindmøllen (620); og

25

en computer til dannelse af en abnormitetsdetektionsmodel (155), koblet til normalrumsdatakilden og abnormitetsrumsdatakilden, til at:

modtage rækken af normale overvågningsknudepunktsværdier og generere sættet af abnorme funktionsvektorer,

30

modtage rækken af abnorme overvågningsknudepunktsværdier og generere sættet af abnorme funktionsvektorer, og

automatisk beregne og udlæse flerheden af beslutningsgrænser for abnormitetsdetektionsmodellen (155) baseret på vindinformation (1560), sættet af normale funktionsvektorer og sættet af abnorme funktionsvektorer.

35

9. System (100) ifølge krav 8, hvor mindst én af rækken af normale overvågningsknudepunktsværdier og rækken af abnorme overvågningsknudepunktsværdier er forbundet med en high-fidelity-udstyrsmodel (eng: high fidelity equipment model).
- 5 10. System (100) ifølge krav 8, hvor flerheden af beslutningsgrænser eksisterer i et multidimensionelt rum og er forbundet med mindst én af: en dynamisk model, design af eksperimentdata, maskinlæringsteknikker, en støttevektormaskine, en fuldfaktoriel proces, Taguchi-screening, en central sammensat metodologi (eng: central composite methodology), en Box-Behnken-metode, driftsforhold i den virkelige verden, et fuldfaktorielt design, et
- 10 screeningsdesign og et centralt sammensat design.

DRAWINGS


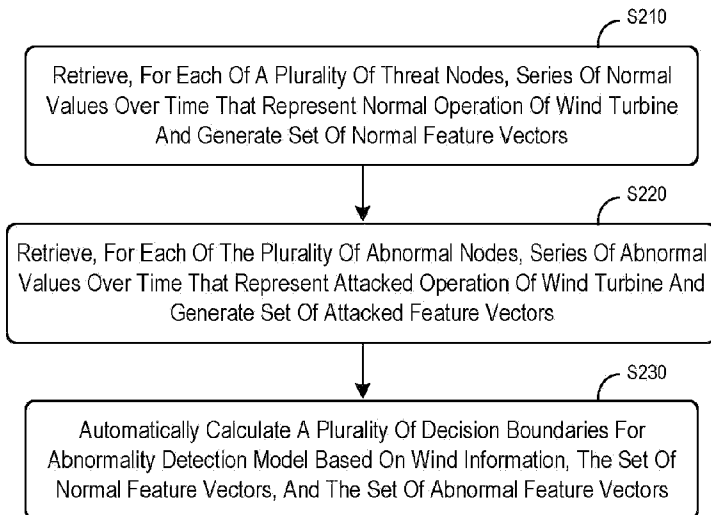
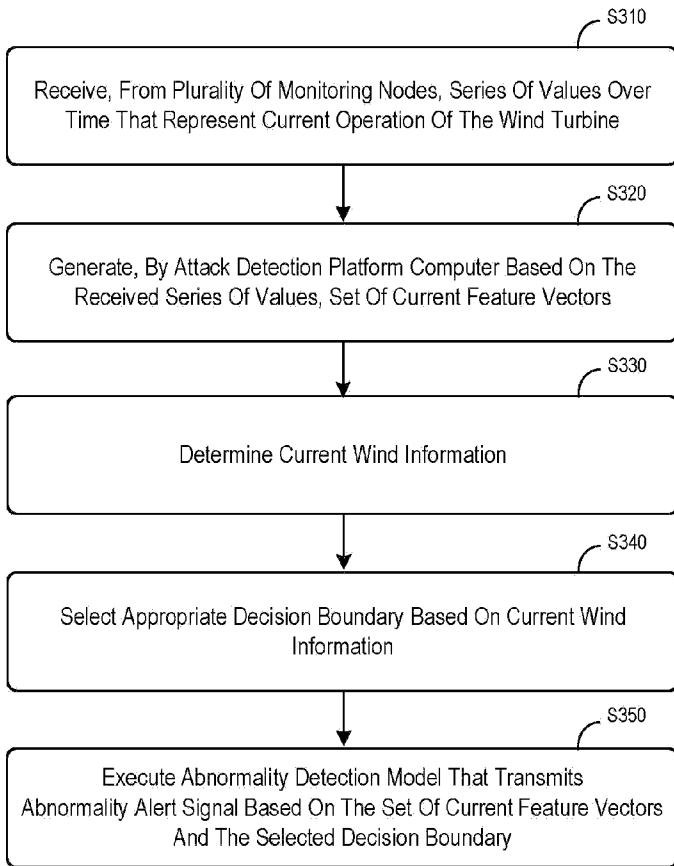


FIG. 1

**FIG. 2**

**FIG. 3**

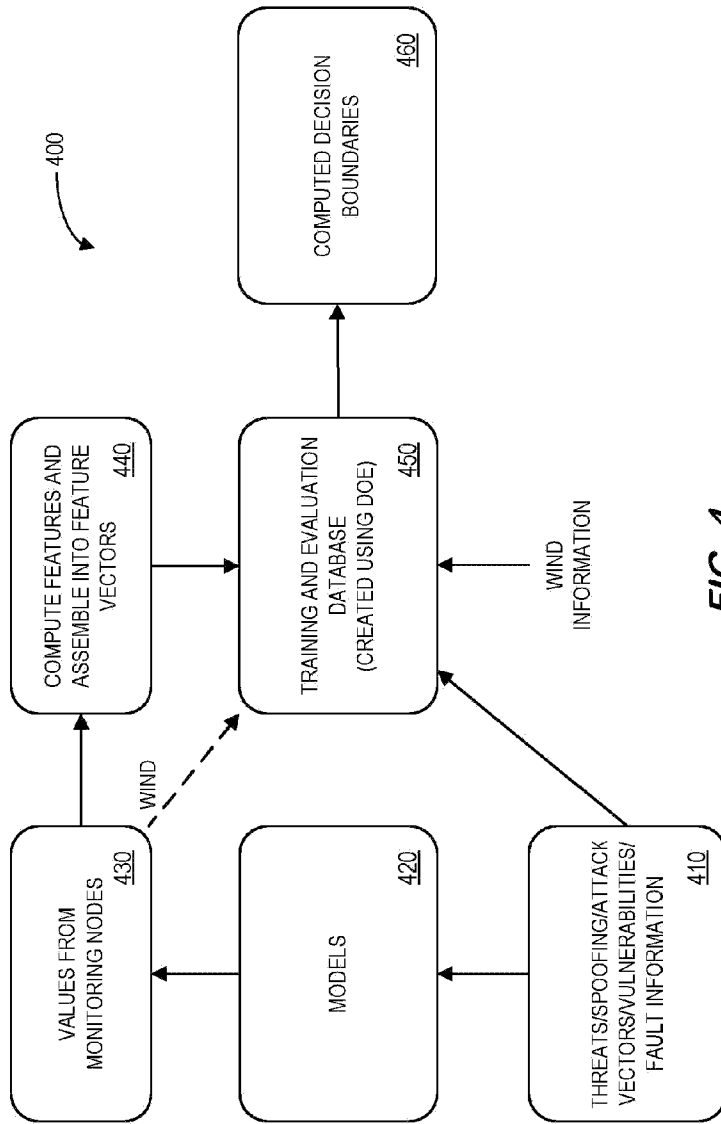


FIG. 4

500 

SENSORS 502	ACTUATORS 504	CONTROLLER 506
<ul style="list-style-type: none"> • Azimuth Angle • Generator Speed • 3 Blade Pitch Angles • Generator Power • Tower Top Fore-Aft Acceleration • Lower Top Side-to-Side Acceleration • Shaft D and Q Moments • Nacelle Yaw Angle • Wind Speed Estimate • Break Status • Air Density • Gust Status • Blade Sensors (IMU, strain gauge, etc., if any) • Ambient Temperature • Component Temps (gear box oil, pitch bearing, etc.) • Power and Speed Set-points 	<ul style="list-style-type: none"> • 3 Blade Pitch • Commands • Generator Torque • Nacelle Yaw Angle • Blade Actuators (flaps, jet, etc., if any) 	<p>Selected based on the controller parameters</p>

FIG. 5




FIG. 6

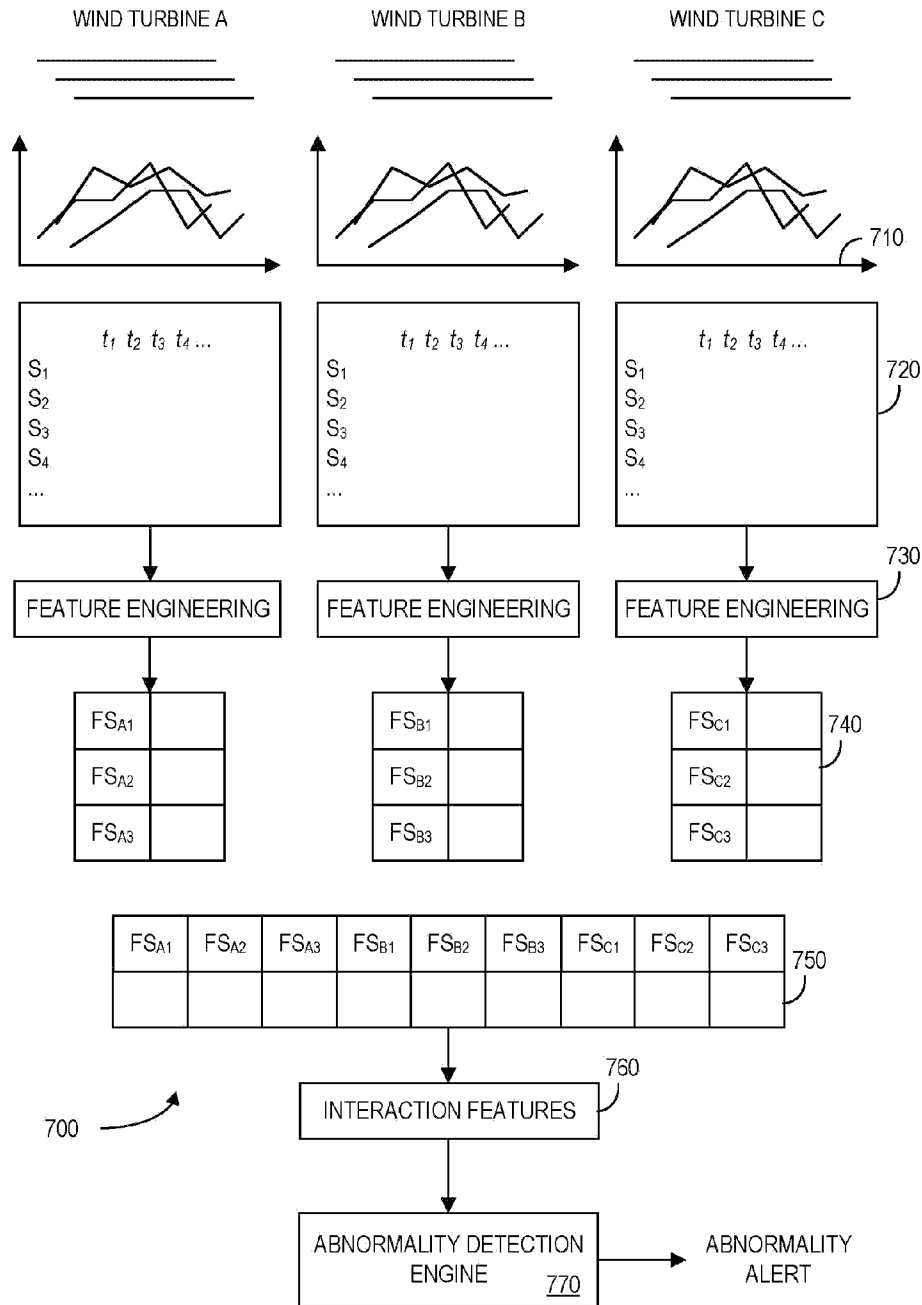


FIG. 7




FIG. 8




FIG. 9

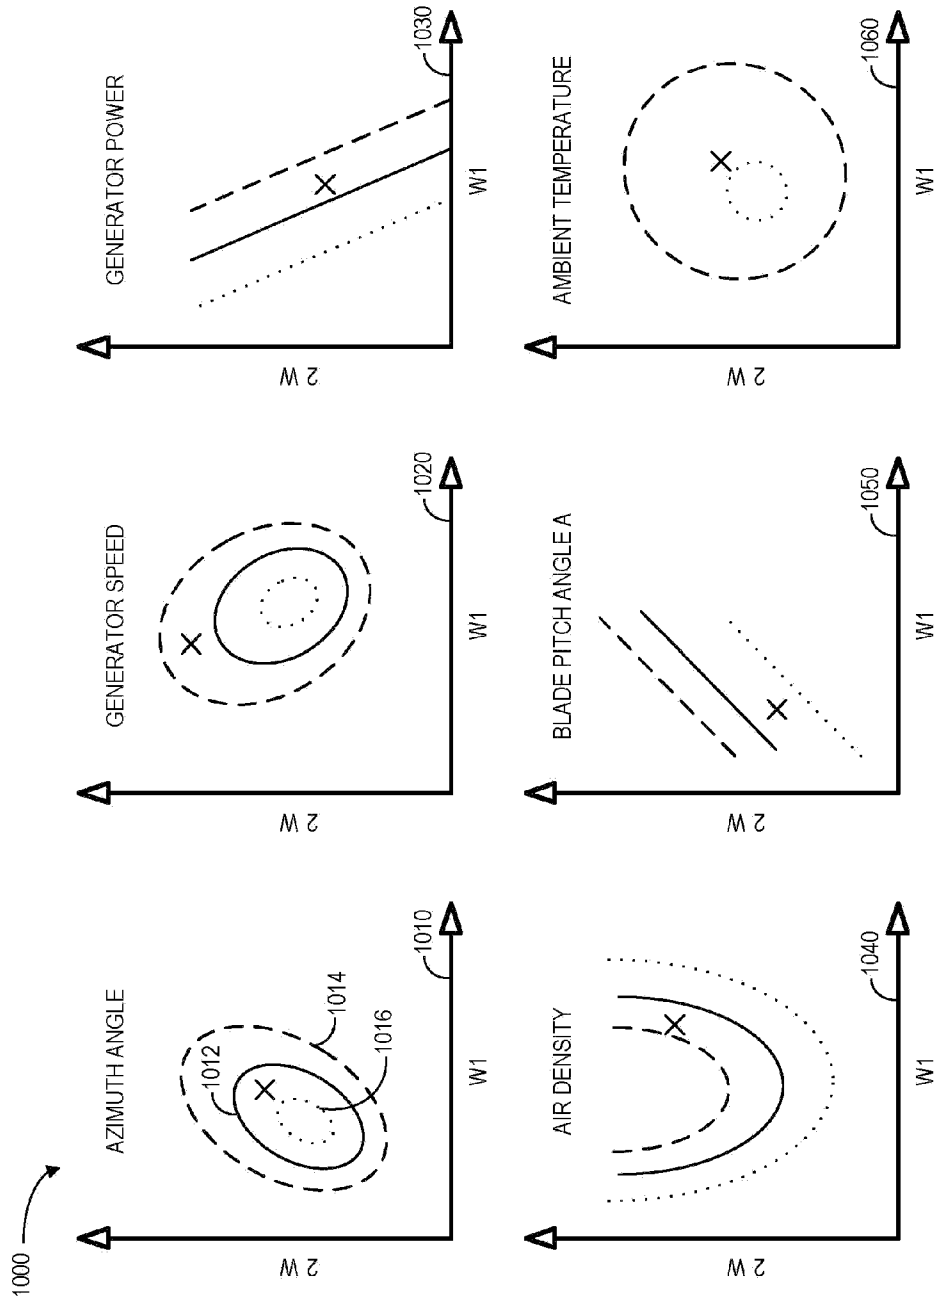


FIG. 10

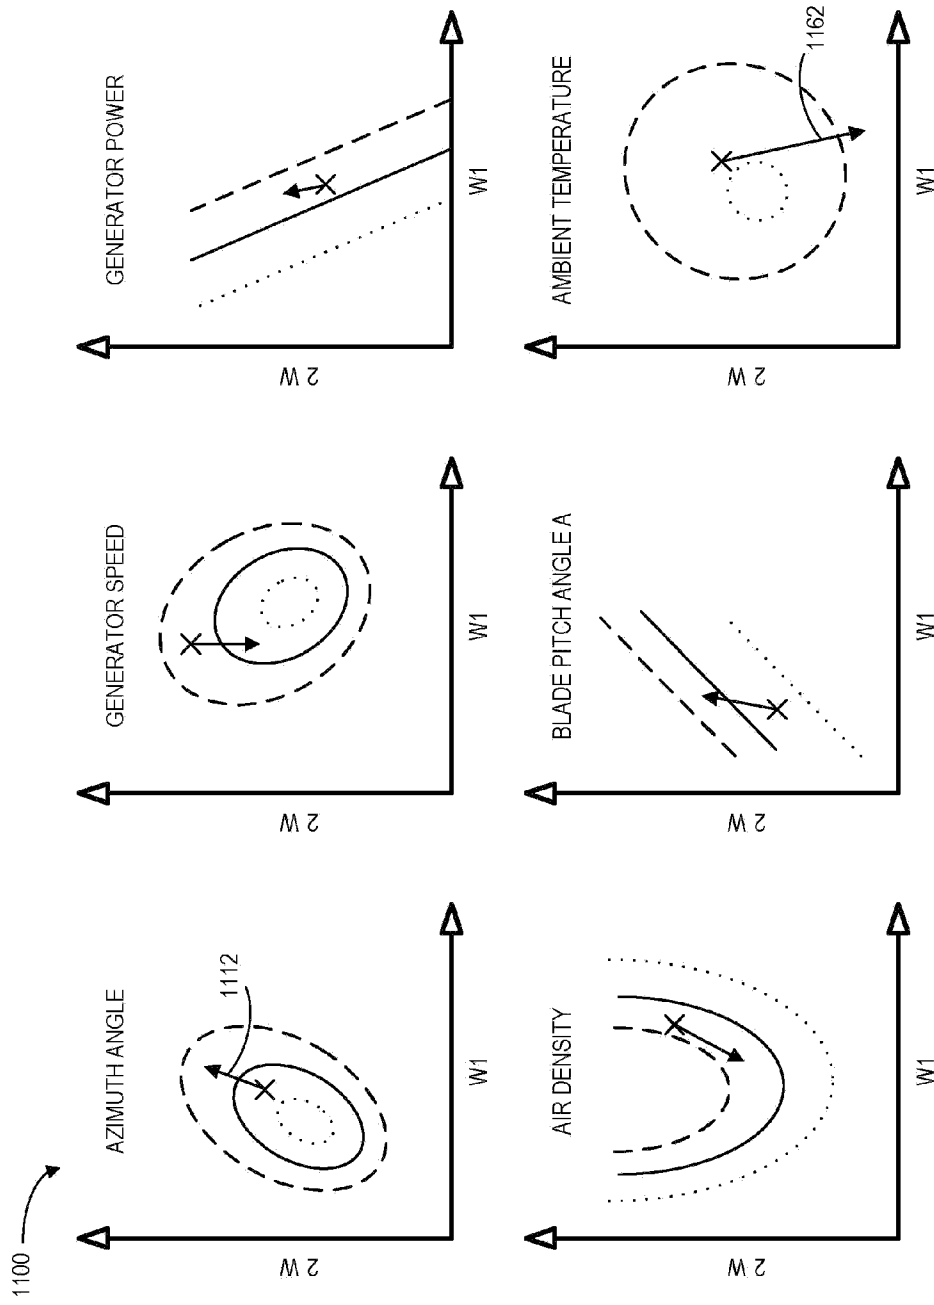


FIG. 11

1100

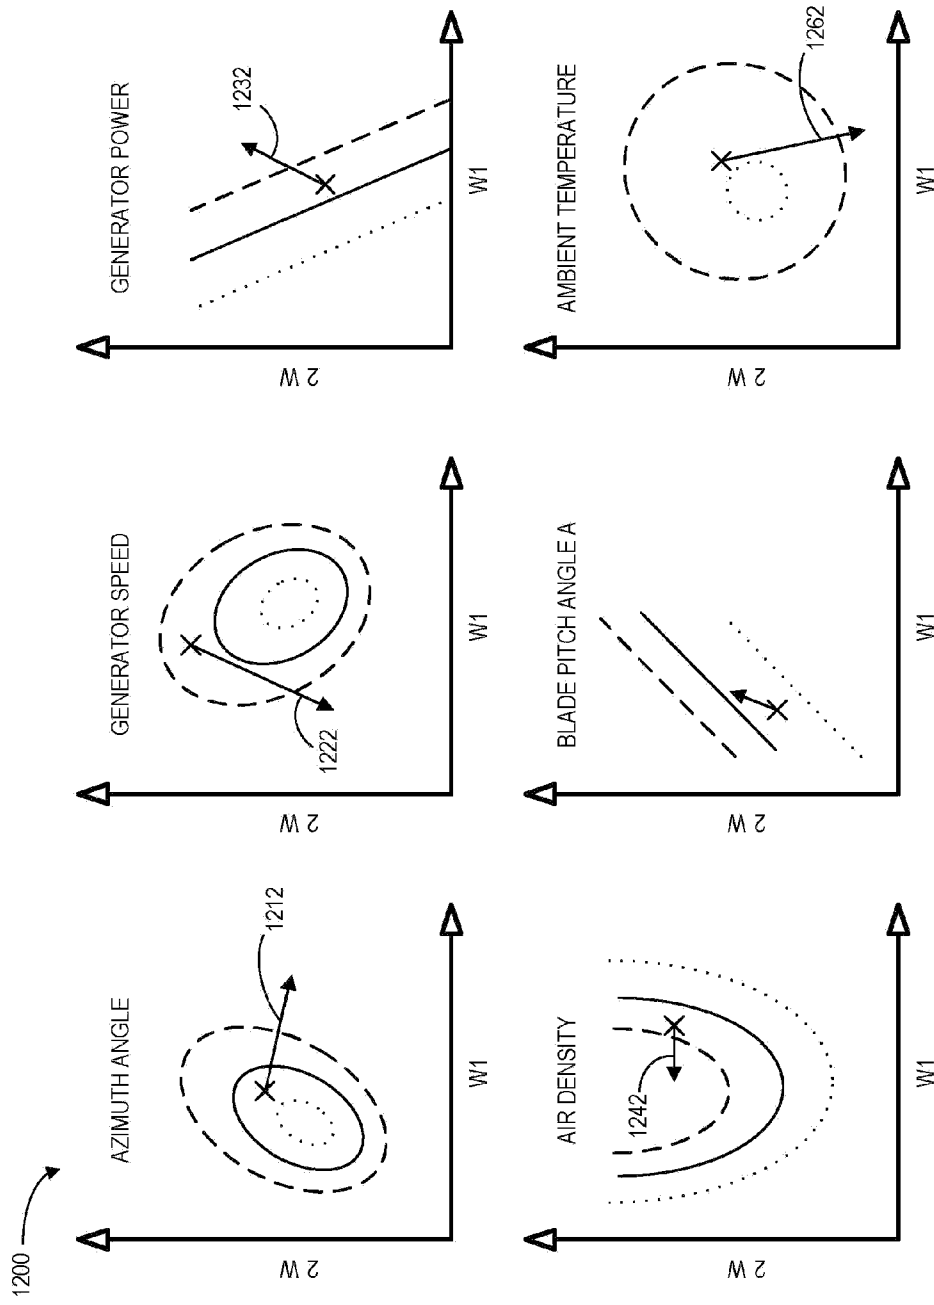


FIG. 12

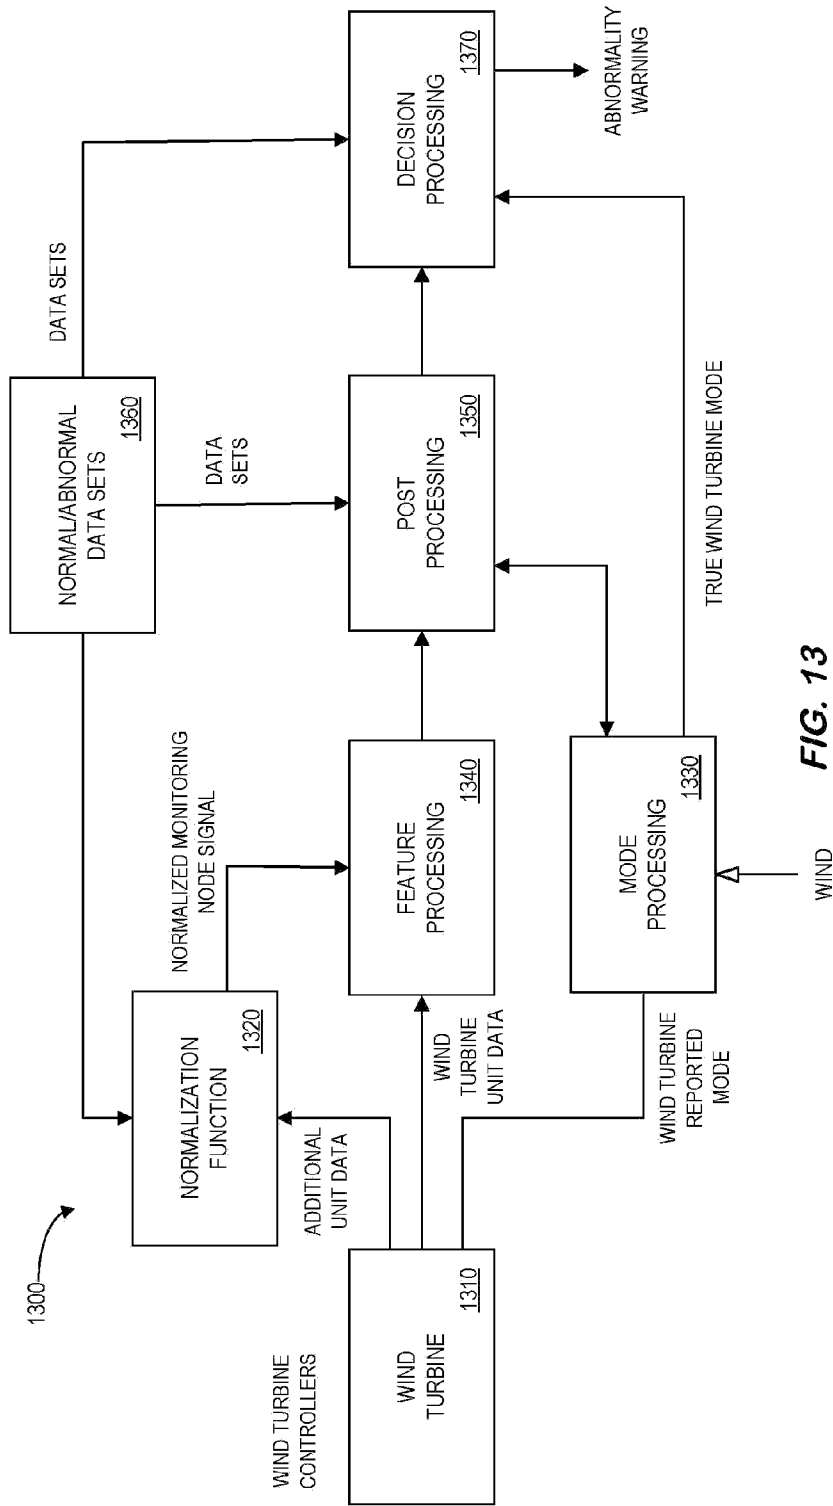


FIG. 13

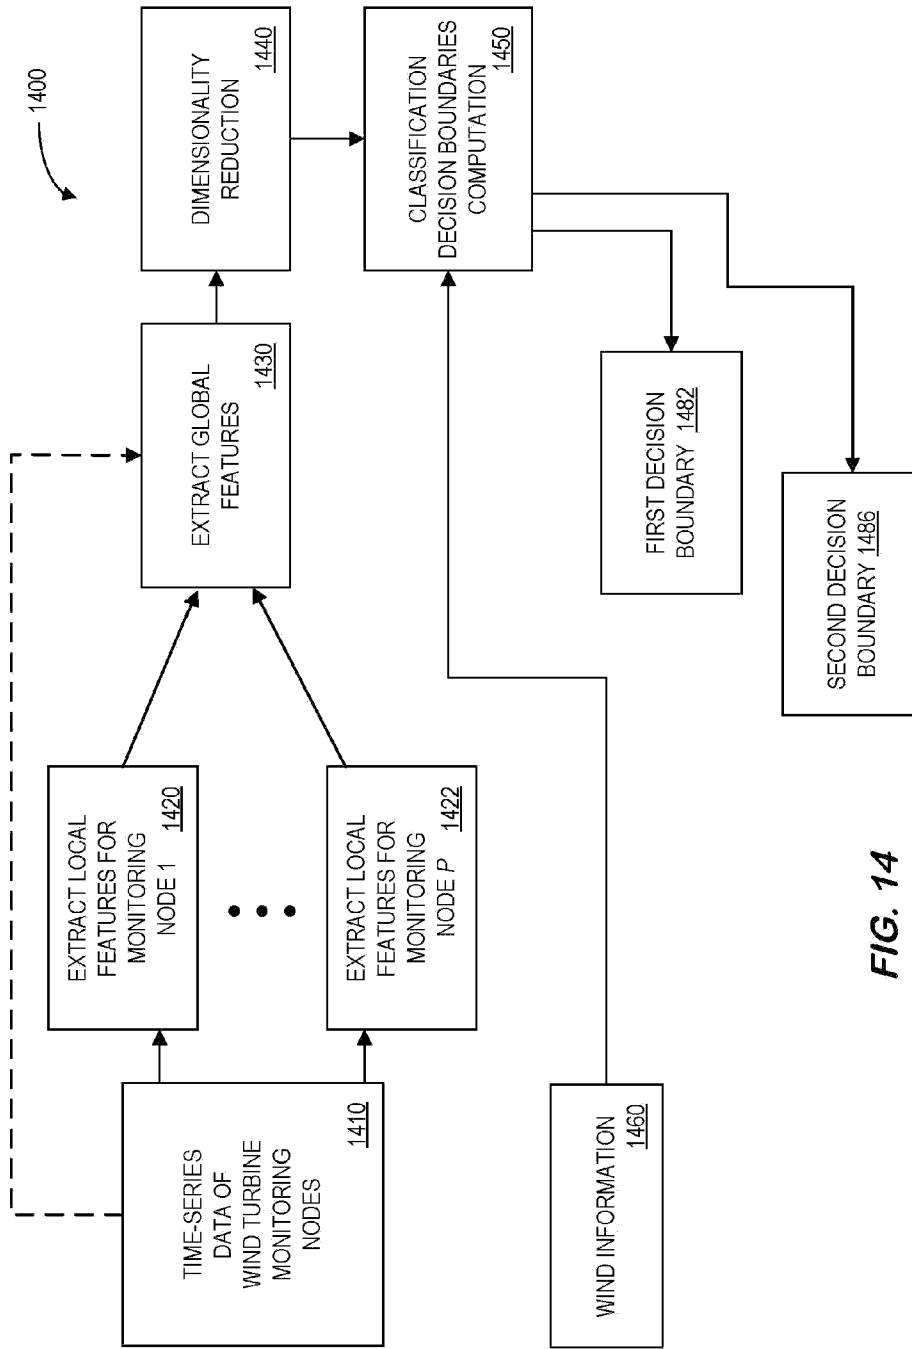


FIG. 14

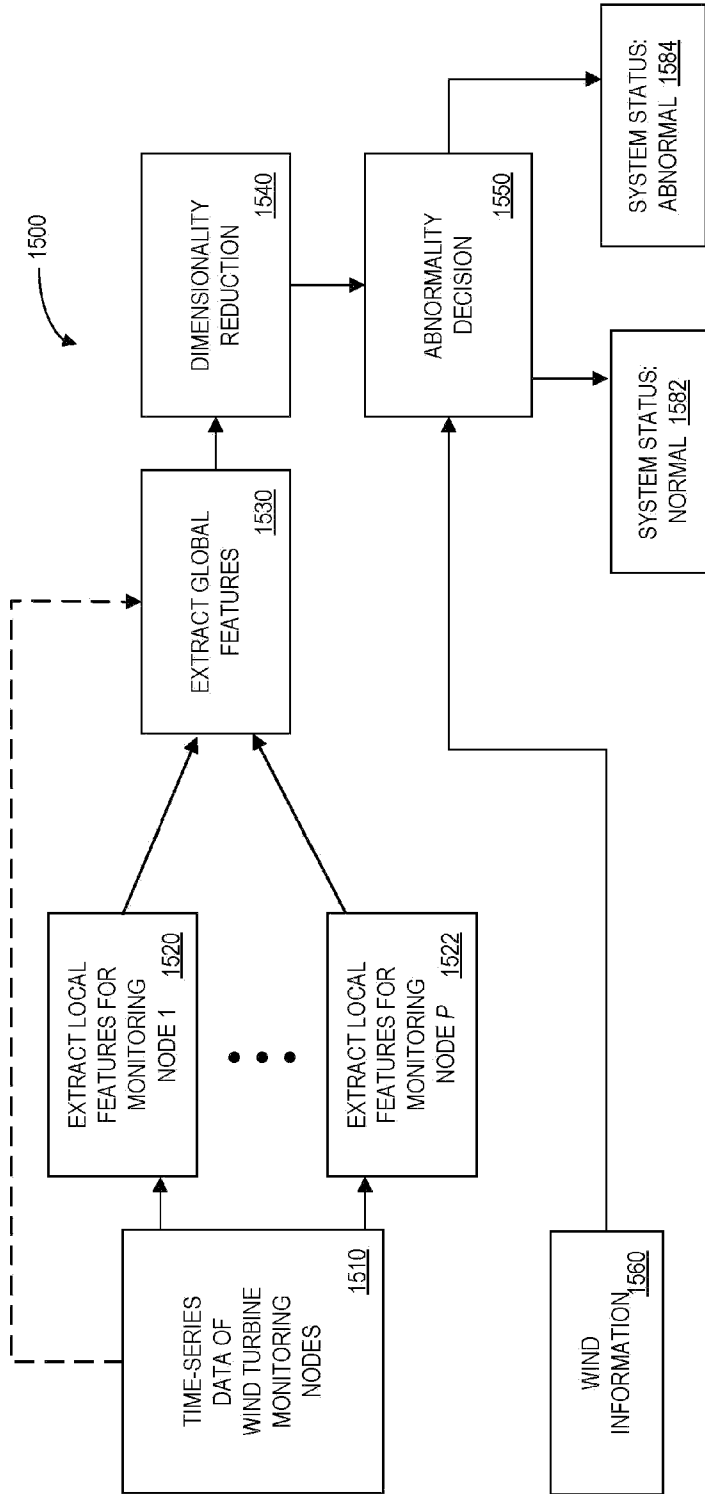


FIG. 15

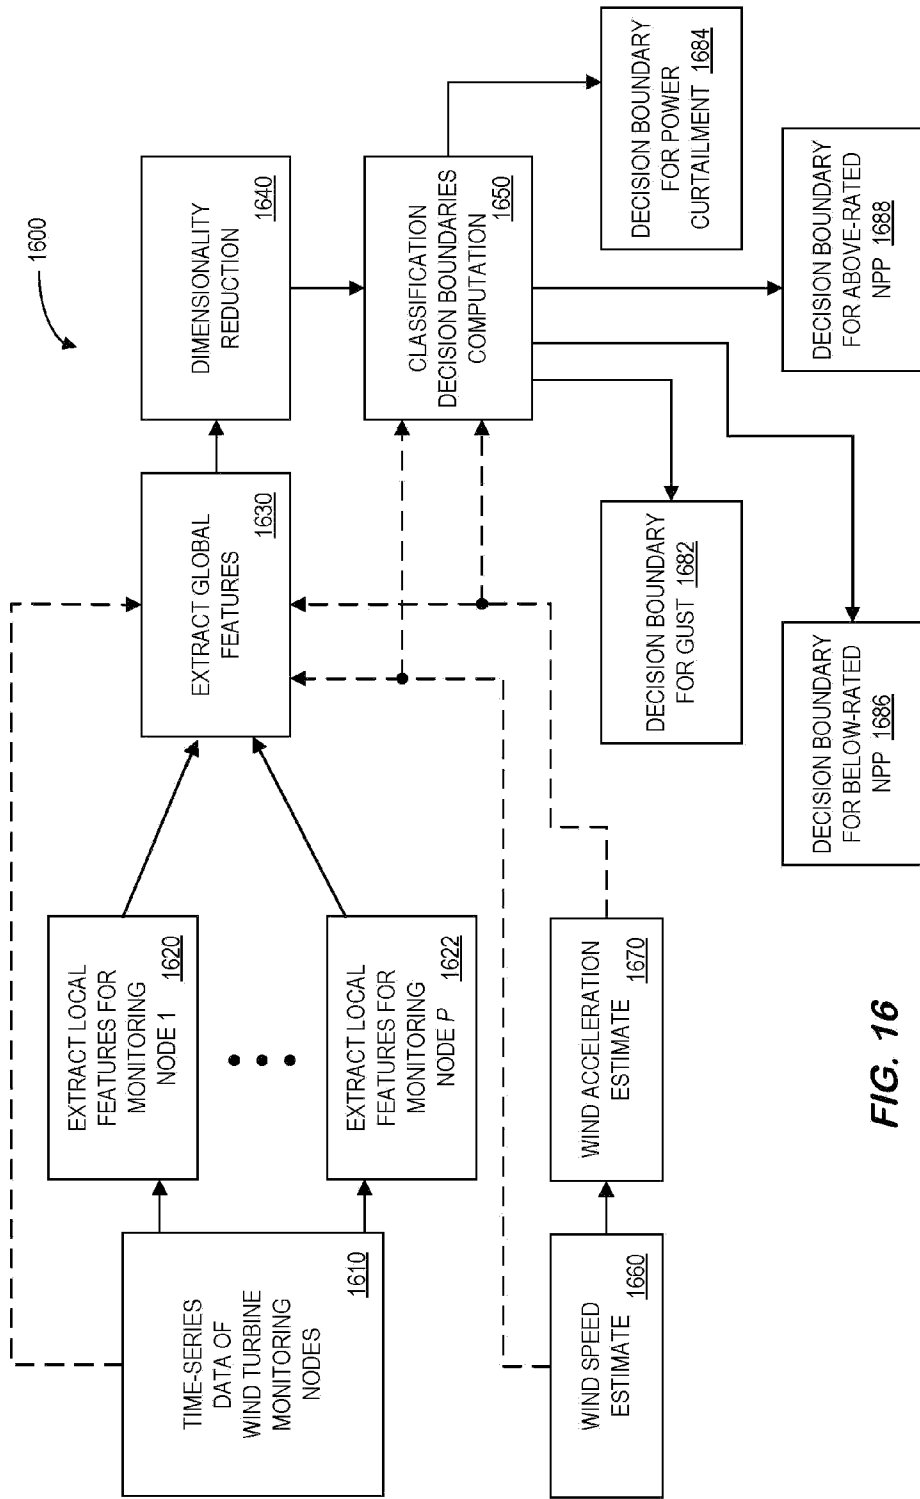


FIG. 16

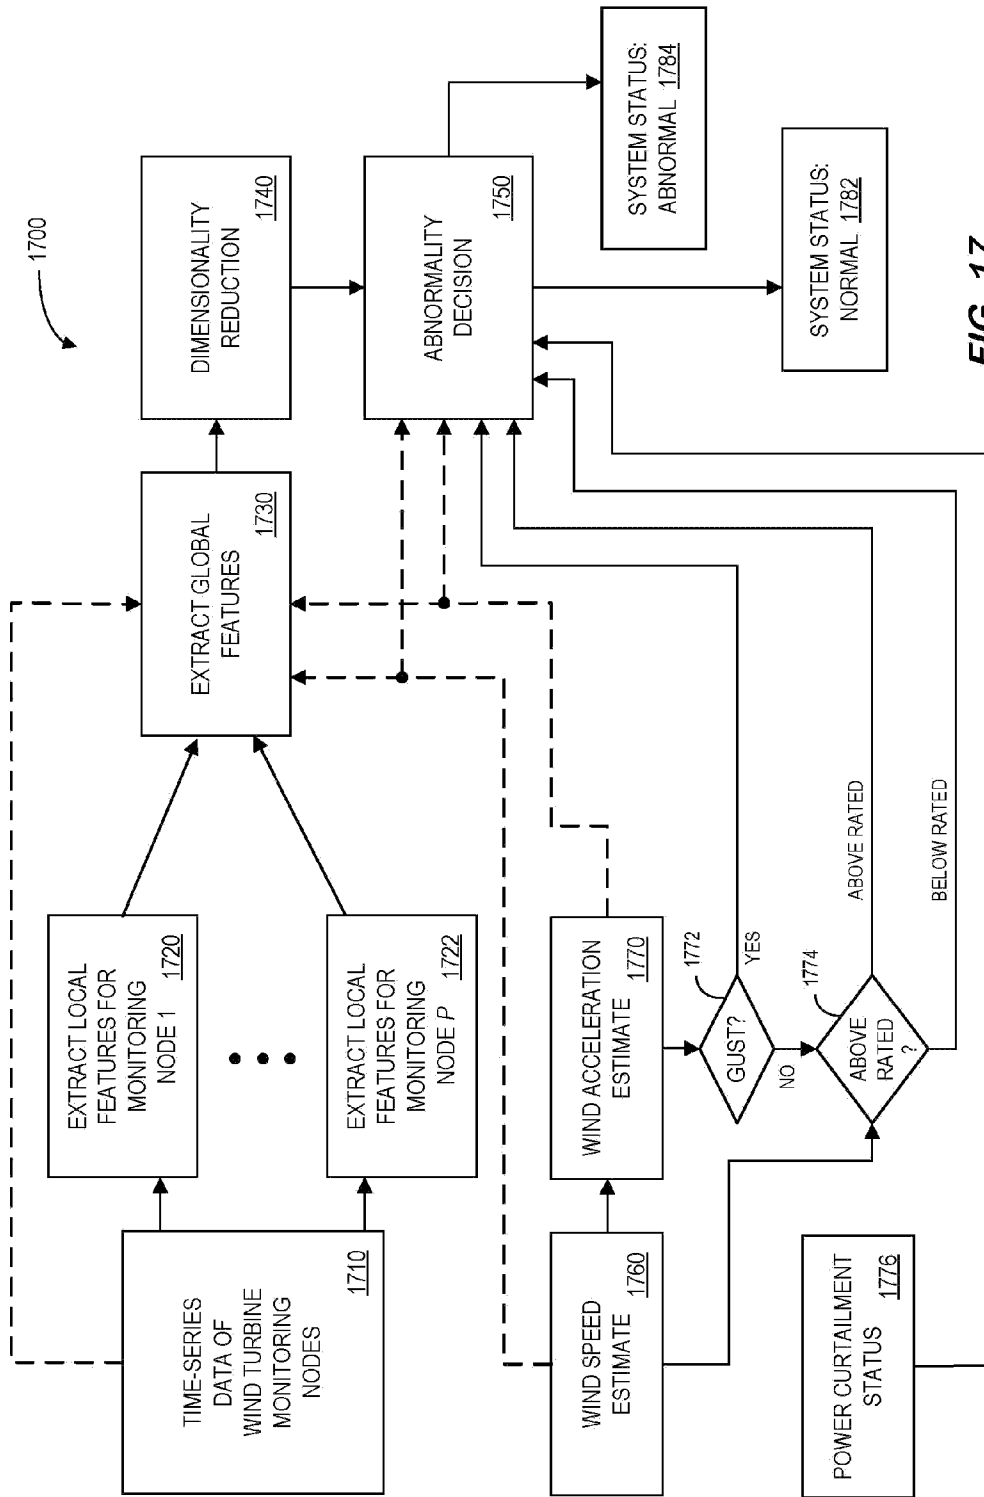


FIG. 17

1800

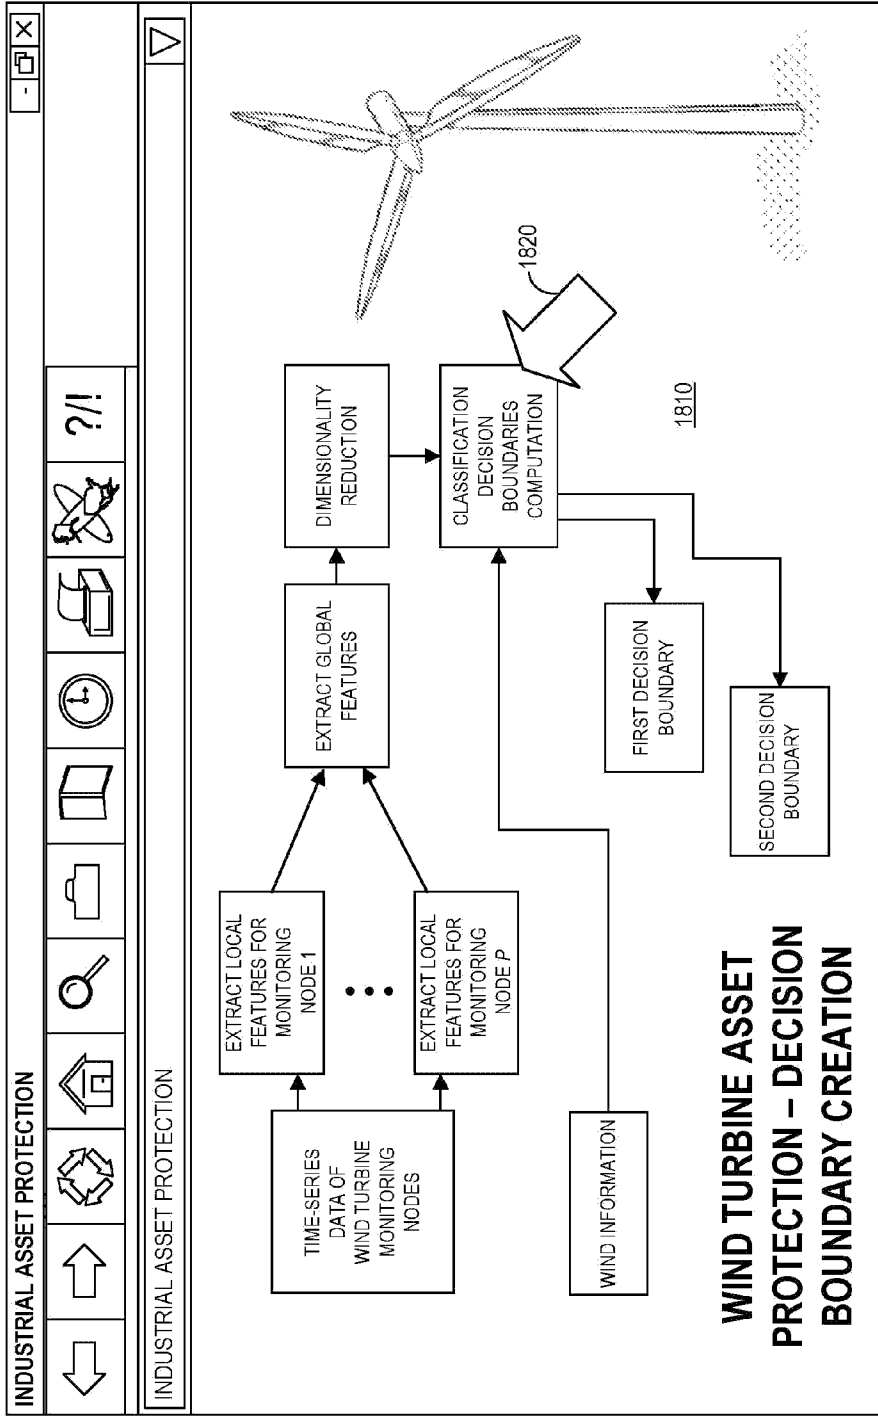


FIG. 18






FIG. 19

2000 

DECISION BOUNDARY IDENTIFIER 2002	DECISION BOUNDARY DESCRIPTION 2004	INPUT REQUIREMENTS 2006	BOUNDARY DEFINITION 2008	CURRENTLY ACTIVE? 2010	STATUS 2012	CAUSE 2014
DB_1	GUST	IR_101	EQ. 1	NO		
DB_2	BELOW-RATED NPP	IR_102	EQ. 2	YES	ABNORMAL	CYBER-ATTACK
DB_3	ABOVE-RATED NPP	IR_103	EQ. 3	NO		
DB_4	POWER CURTAILMENT	IR_104	EQ. 4	NO		

FIG. 20

2100 

STATUS		PREDICTED CLASS <u>2110</u>	
		NORMAL	ABNORMAL
ELM	%		
ACTUAL CLASS <u>2120</u>	NORMAL	98.64	1.36
	ABNORMAL	0	100

FIG. 21