



- (51) **International Patent Classification:**
A47C 27/08 (2006.01) A61B 5/00 (2006.01)
- (21) **International Application Number:**
PCT/US2018/067606
- (22) **International Filing Date:**
27 December 2018 (27.12.2018)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
62/611,172 28 December 2017 (28.12.2017) US
- (71) **Applicant: SLEEP NUMBER CORPORATION**
[US/US]; 1001 Third Avenue South, Minneapolis, MN 55404 (US).
- (72) **Inventors; and**
(71) **Applicants: SAYADI, Omid** [IR/US]; 7024 Via Anacapa, San Jose, CA 95139 (US). **DEMIRLI, Ramazan** [US/US]; 175 West St. James Street, #1409, San Jose, CA 95110 (US). **YOUNG, Steven, Jay** [US/US]; 105 Milani Court, Los Gatos, CA 95030 (US).
- (74) **Agent: NELSON, Stuart, A.**; Fish & Richardson P.C., P.O. Box 1022, Minneapolis, MN 55440-1022 (US).

- (81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) **Designated States** (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

(54) **Title:** BED HAVING SLEEP STAGE DETECTING FEATURE




FIG. 1

(57) **Abstract:** A bed system (100) comprises a first bed (112) having a first mattress, a first pressure sensor (146), and a first controller (124) in data communication with the first pressure sensor (146), the first controller (124) configured to receive pressure readings and to transmit the first pressure readings to a remote server (1204). The bed system (100) further includes a second bed that includes a second mattress having an inflatable chamber. The bed system (100) further includes a second pressure sensor in fluid communication with the second mattress inflatable chamber to sense pressure applied to the second mattress of the inflatable chamber. The bed system further includes a second controller in data communication with the second pressure sensor, the controller configured to receive the one or more sleep-state classifiers and to run the received sleep-state classifiers in order to collect one or more sleep-state votes. The second controller is further configured to determine a sleep-state and operate the bed system according to the determined sleep-state.



Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

- *with international search report (Art. 21(3))*

BED HAVING SLEEP STAGE DETECTING FEATURE

[0001] The present document relates to a bed with sensors used for sleep stage detection.

CROSS-REFERENCE TO RELATED APPLICATIONS

[0002] This application claims priority to U.S. Application Serial No. 62/611,172, filed on December 28, 2017. The disclosure of the prior application is considered part of the disclosure of this application, and is incorporated in its entirety into this application.

BACKGROUND

[0003] In general, a bed is a piece of furniture used as a location to sleep or relax. Many modern beds include a soft mattress on a bed frame. The mattress may include springs, foam material, and/or an air chamber to support the weight of one or more occupants.

SUMMARY

[0004] In one aspect, a bed system includes a first bed that includes a first mattress. The system further includes a first pressure sensor in communication with the mattress to sense pressure applied to the first mattress. The system further includes a first controller in data communication with the first pressure sensor, the first controller configured to: receive, from the first pressure sensor, first pressure readings indicative of the sensed pressure of the inflatable chamber. The first controller is further configured to transmit the first pressure readings to a remote server such that the remote server is able to generate one or more sleep-state classifiers that, when run by a controller on incoming pressure readings, provides a sleep-state vote. The system further includes a second bed that includes a second mattress. The system further includes a second pressure sensor in communication with the second mattress to sense pressure applied to the second mattress. The system further includes a second controller in data communication with the second pressure sensor, the controller configured to: receive the one or more sleep-state classifiers. The second controller is further configured to run the received sleep-state classifiers

on second pressure readings in order to collect one or more sleep-state votes from the running sleep-state classifiers. The second controller is further configured to determine, from the one or more sleep-state votes, a sleep-state of a user on the second bed. The second controller is further configured to responsive to the determined sleep-state, operate the bed system according to the determined sleep-state. Other systems, devices, methods, and computer-readable media may be used.

[0005] Implementations can include any, all, or none of the following features. Operating the bed system according to the determined sleep-state includes one of the list including turning on a light, turning off a light, turning on a warming feature, changing bed temperature, changing HVAC temperature, changing firmness of the inflatable chamber, and articulating a foundation of the bed system. The bed system including the remote server. The remote server is physically remote from the first controller and the second controller; and wherein the remote server is in data communication with the first controller and the second controller. The remote server is configured to: generate training data from the first pressure data; generate, from the training data, the one or more sleep-state classifiers; and send, to the second controller, the one or more sleep-state classifiers. Generating, from the training data, the one or more sleep-state classifiers includes generating a feature set from the training data; mapping the training data to a kernel space; training a classifier with the feature set so that, based on the training data in kernel space, the classifier is able to classify unseen data. Training a classifier includes unsupervised training. The unsupervised training includes at least one of the group including k-means clustering, mixture modeling, hierarchical clustering, self-organizing mapping, and hidden Markov modelling. Training a classifier includes supervised training. The supervised training includes providing the remote server with a set of annotations for the training data. The annotations for the training data are provided by a human. The annotations for the training data are provided programmatically. Generating the one or more presence classifiers includes training a deep learning model on the training data; Training the deep learning model on the training data includes generating an initial neural network configured to receive pressure data and generate presence votes. The presence vote includes a presence classification and a confidence value. Generating the one or more presence classifiers includes determining a loss value for the initial neural network; and iteratively refining, beginning with the initial neural network, to a

final neural network having a lower loss value than the initial neural network. The iterative refining is performed with a gradient descent process until a lower loss value cannot be found with the gradient descent process. A particular sleep-state classifier is used for multiple users in multiple beds. The sleep-state classifiers are personalized for a single user such that the sleep-state classifiers are generated from training data of the single user's use of the bed system and the sleep-state classifiers are used to detect sleep-state of the single user on the second bed. A second set of sleep-state classifiers are personalized for a second user such that the second set of sleep-state classifiers are generated from training data of the second user's use of the bed system and the second set of sleep-state classifiers are used to detect sleep-state of the second user on the second bed. Determining, from the one or more sleep-state votes, a sleep-state of a user on the second bed is personalized for a single user such that votes from different classifiers are weighed based on the classifiers historical accuracy for that user. The first bed and the second bed are separate beds. The first bed and the second bed are the same beds. To run the received sleep-state classifiers on second pressure readings in order to collect one or more sleep-state votes from the running sleep-state classifiers, the second controller is configured to run the received sleep-state classifiers on a plurality of sleep-state classifiers in order to collect one or more sleep-state votes from the running sleep-state classifiers. The system of claim 1, wherein: the first mattress includes a first inflatable chamber; the first pressure sensor is in fluid communication with the first inflatable chamber to sense the pressure applied to the first mattress; the second mattress includes a second inflatable chamber; and the second pressure sensor is in fluid communication with the second inflatable chamber to sense the pressure applied to the second mattress.

[0006] In one aspect, a bed system includes a first bed includes a first mattress. The system further includes means for supporting the first mattress. The system further includes a first pressure sensor in fluid communication with the first mattress inflatable chamber to sense pressure applied to the first mattress of the first inflatable chamber. The system further includes a first controller in data communication with the first pressure sensor, the first controller configured to: receive, from the first pressure sensor, first pressure readings indicative of the sensed pressure of the inflatable chamber. The system further includes transmit the first pressure readings to a remote server such that the remote server is able to generate

one or more sleep-state classifiers that, when run by a controller on incoming pressure readings, provides a sleep-state vote. The system further includes a second bed includes a second mattress having an inflatable chamber. The system further includes a second pressure sensor in fluid communication with the second mattress inflatable chamber to sense pressure applied to the second mattress of the inflatable chamber. The system further includes a second controller in data communication with the second pressure sensor, the controller configured to: receive the one or more sleep-state classifiers. The system further includes run the received sleep-state classifiers on second pressure readings in order to collect one or more sleep-state votes from the running sleep-state classifiers. The system further includes determine, from the one or more sleep-state votes, a sleep-state of a user on the second bed. The system further includes responsive to the determined sleep-state, operate the bed system according to the determined sleep-state a bed system includes a first bed includes a first mattress. The system further includes means for supporting the first mattress. The system further includes a first pressure sensor in fluid communication with the first mattress inflatable chamber to sense pressure applied to the first mattress of the first inflatable chamber. The system further includes a first controller in data communication with the first pressure sensor, the first controller configured to: receive, from the first pressure sensor, first pressure readings indicative of the sensed pressure of the inflatable chamber. The system further includes transmit the first pressure readings to a remote server such that the remote server is able to generate one or more sleep-state classifiers that, when run by a controller on incoming pressure readings, provides a sleep-state vote. The system further includes a second bed includes a second mattress having an inflatable chamber. The system further includes a second pressure sensor in fluid communication with the second mattress inflatable chamber to sense pressure applied to the second mattress of the inflatable chamber. The system further includes a second controller in data communication with the second pressure sensor, the controller configured to: receive the one or more sleep-state classifiers. The system further includes run the received sleep-state classifiers on second pressure readings in order to collect one or more sleep-state votes from the running sleep-state classifiers. The system further includes determine, from the one or more sleep-state votes, a sleep-state of a user on the second bed. The system

further includes responsive to the determined sleep-state, operate the bed system according to the determined sleep-state

[0007] Implementations can include any, all, or none of the following features.

[0008] The technology described here may be used to provide a number of potential advantages. Sleep stage sensing related to a bed may be improved by the use of machine learning techniques. For example, sleep stage sensing may be made faster and/or more accurate. Noisy and complex sensor data may be quickly and efficiently converted into accurate presence information. By utilizing user-specific training data, sleep stage categorizers may be tailored to specific users and more accurately detect presence of the particular user in a bed.

[0009] Other features, aspects and potential advantages will be apparent from the accompanying description and figures.

DESCRIPTION OF DRAWINGS

[0010] FIG. 1 shows an example air bed system.

[0011] FIG. 2 is a block diagram of an example of various components of an air bed system.

[0012] FIG. 3 shows an example environment including a bed in communication with devices located in and around a home.

[0013] FIGs. 4A and 4B are block diagrams of example data processing systems that can be associated with a bed.

[0014] FIGs. 5 and 6 are block diagrams of examples of motherboards that can be used in a data processing system that can be associated with a bed.

[0015] FIG. 7 is a block diagram of an example of a daughterboard that can be used in a data processing system that can be associated with a bed.

[0016] FIG. 8 is a block diagram of an example of a motherboard with no daughterboard that can be used in a data processing system that can be associated with a bed.

[0017] FIG. 9 is a block diagram of an example of a sensory array that can be used in a data processing system that can be associated with a bed.

[0018] FIG. 10 is a block diagram of an example of a control array that can be used in a data processing system that can be associated with a bed

- [0019] FIG. 11 is a block diagram of an example of a computing device that can be used in a data processing system that can be associated with a bed.
- [0020] FIGs. 12-16 are block diagrams of example cloud services that can be used in a data processing system that can be associated with a bed.
- [0021] FIG. 17 is a block diagram of an example of using a data processing system that can be associated with a bed to automate peripherals around the bed.
- [0022] FIG. 18 is a schematic diagram that shows an example of a computing device and a mobile computing device.
- [0023] FIGs. 19A and 19B are swimlane diagrams of example processes for training and using machine-learning classifiers to determine user sleep state.
- [0024] FIG. 20 is a flowchart of an example process for training classifiers on pressure signals.
- [0025] FIG. 21 is a flowchart of an example process to encourage restful sleep.
- [0026] FIG. 22 shows an example system for generating new classifiers.
- [0027] FIG. 23 shows an example system for generating new classifiers.
- [0028] Like reference symbols in the various drawings indicate like elements.

DETAILED DESCRIPTION

- [0029] A bed that detects sleep stages of a user may use machine-learning techniques to identify a sleep state of a user or users on the bed. For example, an airbed may collect pressure signals for a particular user over a period of time. These pressure signals may be used to train one or more personalized categorizers that are each able to categorize a live pressure signal into a sleep state (e.g., asleep or awake, in Rapid Eye Movement (REM), in Non Rapid Eye Movement (NREM), deep, light, N1, N2, N3). One of these categorizers, or a group of these categorizers, can then be used by the bed on live pressure readings to determine the sleep state of the user. Based on this sleep state, the bed or another device may be actuated or driven (e.g., alter the firmness of the mattress or the articulation of the bedframe).
- [0030] In some cases, the computing power available to a controller on a bed is too limited to train categorizers, but is powerful enough execute a classifier on a stream of live pressure data. In such cases, pressure data from a bed may be collected by a controller and sent to a cloud service. The cloud service can then train one or

more classifiers based on the pressure data from that bed and optionally other beds. The classifiers can then be transferred to a bed for use, either the same bed, different beds, or both the same bed and different beds.

[0031] Example Airbed Hardware

[0032] FIG. 1 shows an example air bed system 100 that includes a bed 112. The bed 112 includes at least one air chamber 114 surrounded by a resilient border 116 and encapsulated by bed ticking 118. The resilient border 116 can comprise any suitable material, such as foam.

[0033] As illustrated in FIG. 1, the bed 112 can be a two chamber design having first and second fluid chambers, such as a first air chamber 114A and a second air chamber 114B. In alternative embodiments, the bed 112 can include chambers for use with fluids other than air that are suitable for the application. In some embodiments, such as single beds or kids' beds, the bed 112 can include a single air chamber 114A or 114B or multiple air chambers 114A and 114B. First and second air chambers 114A and 114B can be in fluid communication with a pump 120. The pump 120 can be in electrical communication with a remote control 122 via control box 124. The control box 124 can include a wired or wireless communications interface for communicating with one or more devices, including the remote control 122. The control box 124 can be configured to operate the pump 120 to cause increases and decreases in the fluid pressure of the first and second air chambers 114A and 114B based upon commands input by a user using the remote control 122. In some implementations, the control box 124 is integrated into a housing of the pump 120.

[0034] The remote control 122 can include a display 126, an output selecting mechanism 128, a pressure increase button 129, and a pressure decrease button 130. The output selecting mechanism 128 can allow the user to switch air flow generated by the pump 120 between the first and second air chambers 114A and 114B, thus enabling control of multiple air chambers with a single remote control 122 and a single pump 120. For example, the output selecting mechanism 128 can be a physical control (e.g., switch or button) or an input control displayed on display 126. Alternatively, separate remote control units can be provided for each air chamber and can each include the ability to control multiple air chambers. Pressure increase and decrease buttons 129 and 130 can allow a user to increase or decrease the pressure, respectively, in the air chamber selected with the output selecting mechanism 128.

Adjusting the pressure within the selected air chamber can cause a corresponding adjustment to the firmness of the respective air chamber. In some embodiments, the remote control 122 can be omitted or modified as appropriate for an application. For example, in some embodiments the bed 112 can be controlled by a computer, tablet, smart phone, or other device in wired or wireless communication with the bed 112.

[0035] FIG. 2 is a block diagram of an example of various components of an air bed system. For example, these components can be used in the example air bed system 100. As shown in FIG. 2, the control box 124 can include a power supply 134, a processor 136, a memory 137, a switching mechanism 138, and an analog to digital (A/D) converter 140. The switching mechanism 138 can be, for example, a relay or a solid state switch. In some implementations, the switching mechanism 138 can be located in the pump 120 rather than the control box 124.

[0036] The pump 120 and the remote control 122 are in two-way communication with the control box 124. The pump 120 includes a motor 142, a pump manifold 143, a relief valve 144, a first control valve 145A, a second control valve 145B, and a pressure transducer 146. The pump 120 is fluidly connected with the first air chamber 114A and the second air chamber 114B via a first tube 148A and a second tube 148B, respectively. The first and second control valves 145A and 145B can be controlled by switching mechanism 138, and are operable to regulate the flow of fluid between the pump 120 and first and second air chambers 114A and 114B, respectively.

[0037] In some implementations, the pump 120 and the control box 124 can be provided and packaged as a single unit. In some alternative implementations, the pump 120 and the control box 124 can be provided as physically separate units. In some implementations, the control box 124, the pump 120, or both are integrated within or otherwise contained within a bed frame or bed support structure that supports the bed 112. In some implementations, the control box 124, the pump 120, or both are located outside of a bed frame or bed support structure (as shown in the example in FIG. 1).

[0038] The example air bed system 100 depicted in FIG. 2 includes the two air chambers 114A and 114B and the single pump 120. However, other implementations can include an air bed system having two or more air chambers and one or more pumps incorporated into the air bed system to control the air chambers. For example,

a separate pump can be associated with each air chamber of the air bed system or a pump can be associated with multiple chambers of the air bed system. Separate pumps can allow each air chamber to be inflated or deflated independently and simultaneously. Furthermore, additional pressure transducers can also be incorporated into the air bed system such that, for example, a separate pressure transducer can be associated with each air chamber.

[0039] In use, the processor 136 can, for example, send a decrease pressure command to one of air chambers 114A or 114B, and the switching mechanism 138 can be used to convert the low voltage command signals sent by the processor 136 to higher operating voltages sufficient to operate the relief valve 144 of the pump 120 and open the control valve 145A or 145B. Opening the relief valve 144 can allow air to escape from the air chamber 114A or 114B through the respective air tube 148A or 148B. During deflation, the pressure transducer 146 can send pressure readings to the processor 136 via the A/D converter 140. The A/D converter 140 can receive analog information from pressure transducer 146 and can convert the analog information to digital information useable by the processor 136. The processor 136 can send the digital signal to the remote control 122 to update the display 126 in order to convey the pressure information to the user.

[0040] As another example, the processor 136 can send an increase pressure command. The pump motor 142 can be energized in response to the increase pressure command and send air to the designated one of the air chambers 114A or 114B through the air tube 148A or 148B via electronically operating the corresponding valve 145A or 145B. While air is being delivered to the designated air chamber 114A or 114B in order to increase the firmness of the chamber, the pressure transducer 146 can sense pressure within the pump manifold 143. Again, the pressure transducer 146 can send pressure readings to the processor 136 via the A/D converter 140. The processor 136 can use the information received from the A/D converter 140 to determine the difference between the actual pressure in air chamber 114A or 114B and the desired pressure. The processor 136 can send the digital signal to the remote control 122 to update display 126 in order to convey the pressure information to the user.

[0041] Generally speaking, during an inflation or deflation process, the pressure sensed within the pump manifold 143 can provide an approximation of the

pressure within the respective air chamber that is in fluid communication with the pump manifold 143. An example method of obtaining a pump manifold pressure reading that is substantially equivalent to the actual pressure within an air chamber includes turning off pump 120, allowing the pressure within the air chamber 114A or 114B and the pump manifold 143 to equalize, and then sensing the pressure within the pump manifold 143 with the pressure transducer 146. Thus, providing a sufficient amount of time to allow the pressures within the pump manifold 143 and chamber 114A or 114B to equalize can result in pressure readings that are accurate approximations of the actual pressure within air chamber 114A or 114B. In some implementations, the pressure of the air chambers 114A and/or 114B can be continuously monitored using multiple pressure sensors (not shown).

[0042] In some implementations, information collected by the pressure transducer 146 can be analyzed to determine various states of a person lying on the bed 112. For example, the processor 136 can use information collected by the pressure transducer 146 to determine a heart rate or a respiration rate for a person lying in the bed 112. For example, a user can be lying on a side of the bed 112 that includes the chamber 114A. The pressure transducer 146 can monitor fluctuations in pressure of the chamber 114A and this information can be used to determine the user's heart rate and/or respiration rate. As another example, additional processing can be performed using the collected data to determine a sleep state of the person (e.g., awake, light sleep, deep sleep). For example, the processor 136 can determine when a person falls asleep and, while asleep, the various sleep states of the person.

[0043] Additional information associated with a user of the air bed system 100 that can be determined using information collected by the pressure transducer 146 includes motion of the user, presence of the user on a surface of the bed 112, weight of the user, heart arrhythmia of the user, and apnea. Taking user presence detection for example, the pressure transducer 146 can be used to detect the user's presence on the bed 112, e.g., via a gross pressure change determination and/or via one or more of a respiration rate signal, heart rate signal, and/or other biometric signals. For example, a simple pressure detection process can identify an increase in pressure as an indication that the user is present on the bed 112. As another example, the processor 136 can determine that the user is present on the bed 112 if the detected pressure increases above a specified threshold (so as to indicate that a person or other object

above a certain weight is positioned on the bed 112). As yet another example, the processor 136 can identify an increase in pressure in combination with detected slight, rhythmic fluctuations in pressure as corresponding to the user being present on the bed 112. The presence of rhythmic fluctuations can be identified as being caused by respiration or heart rhythm (or both) of the user. The detection of respiration or a heartbeat can distinguish between the user being present on the bed and another object (e.g., a suit case) being placed upon the bed.

[0044] In some implementations, fluctuations in pressure can be measured at the pump 120. For example, one or more pressure sensors can be located within one or more internal cavities of the pump 120 to detect fluctuations in pressure within the pump 120. The fluctuations in pressure detected at the pump 120 can indicate fluctuations in pressure in one or both of the chambers 114A and 114B. One or more sensors located at the pump 120 can be in fluid communication with the one or both of the chambers 114A and 114B, and the sensors can be operative to determine pressure within the chambers 114A and 114B. The control box 124 can be configured to determine at least one vital sign (e.g., heart rate, respiratory rate) based on the pressure within the chamber 114A or the chamber 114B.

[0045] In some implementations, the control box 124 can analyze a pressure signal detected by one or more pressure sensors to determine a heart rate, respiration rate, and/or other vital signs of a user lying or sitting on the chamber 114A or the chamber 114B. More specifically, when a user lies on the bed 112 positioned over the chamber 114A, each of the user's heart beats, breaths, and other movements can create a force on the bed 112 that is transmitted to the chamber 114A. As a result of the force input to the chamber 114A from the user's movement, a wave can propagate through the chamber 114A and into the pump 120. A pressure sensor located at the pump 120 can detect the wave, and thus the pressure signal output by the sensor can indicate a heart rate, respiratory rate, or other information regarding the user.

[0046] With regard to sleep state, air bed system 100 can determine a user's sleep state by using various biometric signals such as heart rate, respiration, and/or movement of the user. While the user is sleeping, the processor 136 can receive one or more of the user's biometric signals (e.g., heart rate, respiration, and motion) and determine the user's present sleep state based on the received biometric signals. In some implementations, signals indicating fluctuations in pressure in one or both of the

chambers 114A and 114B can be amplified and/or filtered to allow for more precise detection of heart rate and respiratory rate.

[0047] The control box 124 can perform a pattern recognition algorithm or other calculation based on the amplified and filtered pressure signal to determine the user's heart rate and respiratory rate. For example, the algorithm or calculation can be based on assumptions that a heart rate portion of the signal has a frequency in the range of 0.5-4.0 Hz and that a respiration rate portion of the signal has a frequency in the range of less than 1 Hz. The control box 124 can also be configured to determine other characteristics of a user based on the received pressure signal, such as blood pressure, tossing and turning movements, rolling movements, limb movements, weight, the presence or lack of presence of a user, and/or the identity of the user. Techniques for monitoring a user's sleep using heart rate information, respiration rate information, and other user information are disclosed in U.S. Patent Application Publication No. 20100170043 to Steven J. Young et al., titled "APPARATUS FOR MONITORING VITAL SIGNS," the entire contents of which is incorporated herein by reference.

[0048] For example, the pressure transducer 146 can be used to monitor the air pressure in the chambers 114A and 114B of the bed 112. If the user on the bed 112 is not moving, the air pressure changes in the air chamber 114A or 114B can be relatively minimal, and can be attributable to respiration and/or heartbeat. When the user on the bed 112 is moving, however, the air pressure in the mattress can fluctuate by a much larger amount. Thus, the pressure signals generated by the pressure transducer 146 and received by the processor 136 can be filtered and indicated as corresponding to motion, heartbeat, or respiration.

[0049] In some implementations, rather than performing the data analysis in the control box 124 with the processor 136, a digital signal processor (DSP) can be provided to analyze the data collected by the pressure transducer 146. Alternatively, the data collected by the pressure transducer 146 could be sent to a cloud-based computing system for remote analysis.

[0050] In some implementations, the example air bed system 100 further includes a temperature controller configured to increase, decrease, or maintain the temperature of a bed, for example for the comfort of the user. For example, a pad can be placed on top of or be part of the bed 112, or can be placed on top of or be part of

one or both of the chambers 114A and 114B. Air can be pushed through the pad and vented to cool off a user of the bed. Conversely, the pad can include a heating element that can be used to keep the user warm. In some implementations, the temperature controller can receive temperature readings from the pad. In some implementations, separate pads are used for the different sides of the bed 112 (e.g., corresponding to the locations of the chambers 114A and 114B) to provide for differing temperature control for the different sides of the bed.

[0051] In some implementations, the user of the air bed system 100 can use an input device, such as the remote control 122, to input a desired temperature for the surface of the bed 112 (or for a portion of the surface of the bed 112). The desired temperature can be encapsulated in a command data structure that includes the desired temperature as well as identifies the temperature controller as the desired component to be controlled. The command data structure can then be transmitted via Bluetooth or another suitable communication protocol to the processor 136. In various examples, the command data structure is encrypted before being transmitted. The temperature controller can then configure its elements to increase or decrease the temperature of the pad depending on the temperature input into remote control 122 by the user.

[0052] In some implementations, data can be transmitted from a component back to the processor 136 or to one or more display devices, such as the display 126. For example, the current temperature as determined by a sensor element of temperature controller, the pressure of the bed, the current position of the foundation or other information can be transmitted to control box 124. The control box 124 can then transmit the received information to remote control 122 where it can be displayed to the user (e.g., on the display 126).

[0053] In some implementations, the example air bed system 100 further includes an adjustable foundation and an articulation controller configured to adjust the position of a bed (e.g., the bed 112) by adjusting the adjustable foundation that supports the bed. For example, the articulation controller can adjust the bed 112 from a flat position to a position in which a head portion of a mattress of the bed is inclined upward (e.g., to facilitate a user sitting up in bed and/or watching television). In some implementations, the bed 112 includes multiple separately articulable sections. For example, portions of the bed corresponding to the locations of the chambers 114A

and 114B can be articulated independently from each other, to allow one person positioned on the bed 112 surface to rest in a first position (e.g., a flat position) while a second person rests in a second position (e.g., an reclining position with the head raised at an angle from the waist). In some implementations, separate positions can be set for two different beds (e.g., two twin beds placed next to each other). The foundation of the bed 112 can include more than one zone that can be independently adjusted. The articulation controller can also be configured to provide different levels of massage to one or more users on the bed 112.

[0054] Example of a Bed in a Bedroom Environment

[0055] FIG. 3 shows an example environment 300 including a bed 302 in communication with devices located in and around a home. In the example shown, the bed 302 includes pump 304 for controlling air pressure within two air chambers 306a and 306b (as described above with respect to the air chambers 114A-114B). The pump 304 additionally includes circuitry for controlling inflation and deflation functionality performed by the pump 304. The circuitry is further programmed to detect fluctuations in air pressure of the air chambers 306a-b and used the detected fluctuations in air pressure to identify bed presence of a user 308, sleep state of the user 308, movement of the user 308, and biometric signals of the user 308 such as heart rate and respiration rate. In the example shown, the pump 304 is located within a support structure of the bed 302 and the control circuitry 334 for controlling the pump 304 is integrated with the pump 304. In some implementations, the control circuitry 334 is physically separate from the pump 304 and is in wireless or wired communication with the pump 304. In some implementations, the pump 304 and/or control circuitry 334 are located outside of the bed 302. In some implementations, various control functions can be performed by systems located in different physical locations. For example, circuitry for controlling actions of the pump 304 can be located within a pump casing of the pump 304 while control circuitry 334 for performing other functions associated with the bed 302 can be located in another portion of the bed 302, or external to the bed 302. As another example, control circuitry 334 located within the pump 304 can communicate with control circuitry 334 at a remote location through a LAN or WAN (e.g., the internet). As yet another example, the control circuitry 334 can be included in the control box 124 of FIGs. 1 and 2.

[0056] In some implementations, one or more devices other than, or in addition to, the pump 304 and control circuitry 334 can be utilized to identify user bed presence, sleep state, movement, and biometric signals. For example, the bed 302 can include a second pump in addition to the pump 304, with each of the two pumps connected to a respective one of the air chambers 306a-b. For example, the pump 304 can be in fluid communication with the air chamber 306b to control inflation and deflation of the air chamber 306b as well as detect user signals for a user located over the air chamber 306b such as bed presence, sleep state, movement, and biometric signals while the second pump is in fluid communication with the air chamber 306a to control inflation and deflation of the air chamber 306a as well as detect user signals for a user located over the air chamber 306a.

[0057] As another example, the bed 302 can include one or more pressure sensitive pads or surface portions that are operable to detect movement, including user presence, user motion, respiration, and heart rate. For example, a first pressure sensitive pad can be incorporated into a surface of the bed 302 over a left portion of the bed 302, where a first user would normally be located during sleep, and a second pressure sensitive pad can be incorporated into the surface of the bed 302 over a right portion of the bed 302, where a second user would normally be located during sleep. The movement detected by the one or more pressure sensitive pads or surface portions can be used by control circuitry 334 to identify user sleep state, bed presence, or biometric signals.

[0058] In some implementations, information detected by the bed (e.g., motion information) is processed by control circuitry 334 (e.g., control circuitry 334 integrated with the pump 304) and provided to one or more user devices such as a user device 310 for presentation to the user 308 or to other users. In the example depicted in FIG. 3, the user device 310 is a tablet device; however, in some implementations, the user device 310 can be a personal computer, a smart phone, a smart television (e.g., a television 312), or other user device capable of wired or wireless communication with the control circuitry 334. The user device 310 can be in communication with control circuitry 334 of the bed 302 through a network or through direct point-to-point communication. For example, the control circuitry 334 can be connected to a LAN (e.g., through a Wi-Fi router) and communicate with the user device 310 through the LAN. As another example, the control circuitry 334 and

the user device 310 can both connect to the Internet and communicate through the Internet. For example, the control circuitry 334 can connect to the Internet through a WiFi router and the user device 310 can connect to the Internet through communication with a cellular communication system. As another example, the control circuitry 334 can communicate directly with the user device 310 through a wireless communication protocol such as Bluetooth. As yet another example, the control circuitry 334 can communicate with the user device 310 through a wireless communication protocol such as ZigBee, Z-Wave, infrared, or another wireless communication protocol suitable for the application. As another example, the control circuitry 334 can communicate with the user device 310 through a wired connection such as, for example, a USB connector, serial/RS232, or another wired connection suitable for the application.

[0059] The user device 310 can display a variety of information and statistics related to sleep, or user 308's interaction with the bed 302. For example, a user interface displayed by the user device 310 can present information including amount of sleep for the user 308 over a period of time (e.g., a single evening, a week, a month, etc.) amount of deep sleep, ratio of deep sleep to restless sleep, time lapse between the user 308 getting into bed and the user 308 falling asleep, total amount of time spent in the bed 302 for a given period of time, heart rate for the user 308 over a period of time, respiration rate for the user 308 over a period of time, or other information related to user interaction with the bed 302 by the user 308 or one or more other users of the bed 302. In some implementations, information for multiple users can be presented on the user device 310, for example information for a first user positioned over the air chamber 306a can be presented along with information for a second user positioned over the air chamber 306b. In some implementations, the information presented on the user device 310 can vary according to the age of the user 308. For example, the information presented on the user device 310 can evolve with the age of the user 308 such that different information is presented on the user device 310 as the user 308 ages as a child or an adult.

[0060] The user device 310 can also be used as an interface for the control circuitry 334 of the bed 302 to allow the user 308 to enter information. The information entered by the user 308 can be used by the control circuitry 334 to provide better information to the user or to various control signals for controlling

functions of the bed 302 or other devices. For example, the user can enter information such as weight, height, and age and the control circuitry 334 can use this information to provide the user 308 with a comparison of the user's tracked sleep information to sleep information of other people having similar weights, heights, and/or ages as the user 308. As another example, the user 308 can use the user device 310 as an interface for controlling air pressure of the air chambers 306a and 306b, for controlling various recline or incline positions of the bed 302, for controlling temperature of one or more surface temperature control devices of the bed 302, or for allowing the control circuitry 334 to generate control signals for other devices (as described in greater detail below).

[0061] In some implementations, control circuitry 334 of the bed 302 (e.g., control circuitry 334 integrated into the pump 304) can communicate with other devices or systems in addition to or instead of the user device 310. For example, the control circuitry 334 can communicate with the television 312, a lighting system 314, a thermostat 316, a security system 318, or other house hold devices such as an oven 322, a coffee maker 324, a lamp 326, and a nightlight 328. Other examples of devices and/or systems that the control circuitry 334 can communicate with include a system for controlling window blinds 330, one or more devices for detecting or controlling the states of one or more doors 332 (such as detecting if a door is open, detecting if a door is locked, or automatically locking a door), and a system for controlling a garage door 320 (e.g., control circuitry 334 integrated with a garage door opener for identifying an open or closed state of the garage door 320 and for causing the garage door opener to open or close the garage door 320). Communications between the control circuitry 334 of the bed 302 and other devices can occur through a network (e.g., a LAN or the Internet) or as point-to-point communication (e.g., using Bluetooth, radio communication, or a wired connection). In some implementations, control circuitry 334 of different beds 302 can communicate with different sets of devices. For example, a kid bed may not communicate with and/or control the same devices as an adult bed. In some embodiments, the bed 302 can evolve with the age of the user such that the control circuitry 334 of the bed 302 communicates with different devices as a function of age of the user.

[0062] The control circuitry 334 can receive information and inputs from other devices/systems and use the received information and inputs to control actions

of the bed 302 or other devices. For example, the control circuitry 334 can receive information from the thermostat 316 indicating a current environmental temperature for a house or room in which the bed 302 is located. The control circuitry 334 can use the received information (along with other information) to determine if a temperature of all or a portion of the surface of the bed 302 should be raised or lowered. The control circuitry 334 can then cause a heating or cooling mechanism of the bed 302 to raise or lower the temperature of the surface of the bed 302. For example, the user 308 can indicate a desired sleeping temperature of 74 degrees while a second user of the bed 302 indicates a desired sleeping temperature of 72 degrees. The thermostat 316 can indicate to the control circuitry 334 that the current temperature of the bedroom is 72 degrees. The control circuitry 334 can identify that the user 308 has indicated a desired sleeping temperature of 74 degrees, and send control signals to a heating pad located on the user 308's side of the bed to raise the temperature of the portion of the surface of the bed 302 where the user 308 is located to raise the temperature of the user 308's sleeping surface to the desired temperature.

[0063] The control circuitry 334 can also generate control signals controlling other devices and propagate the control signals to the other devices. In some implementations, the control signals are generated based on information collected by the control circuitry 334, including information related to user interaction with the bed 302 by the user 308 and/or one or more other users. In some implementations, information collected from one or more other devices other than the bed 302 are used when generating the control signals. For example, information relating to environmental occurrences (e.g., environmental temperature, environmental noise level, and environmental light level), time of day, time of year, day of the week, or other information can be used when generating control signals for various devices in communication with the control circuitry 334 of the bed 302. For example, information on the time of day can be combined with information relating to movement and bed presence of the user 308 to generate control signals for the lighting system 314. In some implementations, rather than or in addition to providing control signals for one or more other devices, the control circuitry 334 can provide collected information (e.g., information related to user movement, bed presence, sleep state, or biometric signals for the user 308) to one or more other devices to allow the one or more other devices to utilize the collected information when generating control

signals. For example, control circuitry 334 of the bed 302 can provide information relating to user interactions with the bed 302 by the user 308 to a central controller (not shown) that can use the provided information to generate control signals for various devices, including the bed 302.

[0064] Still referring to FIG. 3, the control circuitry 334 of the bed 302 can generate control signals for controlling actions of other devices, and transmit the control signals to the other devices in response to information collected by the control circuitry 334, including bed presence of the user 308, sleep state of the user 308, and other factors. For example, control circuitry 334 integrated with the pump 304 can detect a feature of a mattress of the bed 302, such as an increase in pressure in the air chamber 306b, and use this detected increase in air pressure to determine that the user 308 is present on the bed 302. In some implementations, the control circuitry 334 can identify a heart rate or respiratory rate for the user 308 to identify that the increase in pressure is due to a person sitting, laying, or otherwise resting on the bed 302 rather than an inanimate object (such as a suitcase) having been placed on the bed 302. In some implementations, the information indicating user bed presence is combined with other information to identify a current or future likely state for the user 308. For example, a detected user bed presence at 11:00am can indicate that the user is sitting on the bed (e.g., to tie her shoes, or to read a book) and does not intend to go to sleep, while a detected user bed presence at 10:00pm can indicate that the user 308 is in bed for the evening and is intending to fall asleep soon. As another example, if the control circuitry 334 detects that the user 308 has left the bed 302 at 6:30am (e.g., indicating that the user 308 has woken up for the day), and then later detects user bed presence of the user 308 at 7:30am, the control circuitry 334 can use this information that the newly detected user bed presence is likely temporary (e.g., while the user 308 ties her shoes before heading to work) rather than an indication that the user 308 is intending to stay on the bed 302 for an extended period.

[0065] In some implementations, the control circuitry 334 is able to use collected information (including information related to user interaction with the bed 302 by the user 308, as well as environmental information, time information, and input received from the user) to identify use patterns for the user 308. For example, the control circuitry 334 can use information indicating bed presence and sleep states for the user 308 collected over a period of time to identify a sleep pattern for the user.

For example, the control circuitry 334 can identify that the user 308 generally goes to bed between 9:30pm and 10:00pm, generally falls asleep between 10:00pm and 11:00pm, and generally wakes up between 6:30am and 6:45am based on information indicating user presence and biometrics for the user 308 collected over a week. The control circuitry 334 can use identified patterns for a user to better process and identify user interactions with the bed 302 by the user 308.

[0066] For example, given the above example user bed presence, sleep, and wake patterns for the user 308, if the user 308 is detected as being on the bed at 3:00pm, the control circuitry 334 can determine that the user's presence on the bed is only temporary, and use this determination to generate different control signals than would be generated if the control circuitry 334 determined that the user 308 was in bed for the evening. As another example, if the control circuitry 334 detects that the user 308 has gotten out of bed at 3:00am, the control circuitry 334 can use identified patterns for the user 308 to determine that the user has only gotten up temporarily (for example, to use the rest room, or get a glass of water) and is not up for the day. By contrast, if the control circuitry 334 identifies that the user 308 has gotten out of the bed 302 at 6:40am, the control circuitry 334 can determine that the user is up for the day and generate a different set of control signals than those that would be generated if it were determined that the user 308 were only getting out of bed temporarily (as would be the case when the user 308 gets out of the bed 302 at 3:00am). For other users 308, getting out of the bed 302 at 3:00am can be the normal wake-up time, which the control circuitry 334 can learn and respond to accordingly.

[0067] As described above, the control circuitry 334 for the bed 302 can generate control signals for control functions of various other devices. The control signals can be generated, at least in part, based on detected interactions by the user 308 with the bed 302, as well as other information including time, date, temperature, etc. For example, the control circuitry 334 can communicate with the television 312, receive information from the television 312, and generate control signals for controlling functions of the television 312. For example, the control circuitry 334 can receive an indication from the television 312 that the television 312 is currently on. If the television 312 is located in a different room from the bed 302, the control circuitry 334 can generate a control signal to turn the television 312 off upon making a determination that the user 308 has gone to bed for the evening. For example, if bed

presence of the user 308 on the bed 302 is detected during a particular time range (e.g., between 8:00pm and 7:00am) and persists for longer than a threshold period of time (e.g., 10 minutes) the control circuitry 334 can use this information to determine that the user 308 is in bed for the evening. If the television 312 is on (as indicated by communications received by the control circuitry 334 of the bed 302 from the television 312) the control circuitry 334 can generate a control signal to turn the television 312 off. The control signals can then be transmitted to the television (e.g., through a directed communication link between the television 312 and the control circuitry 334 or through a network). As another example, rather than turning off the television 312 in response to detection of user bed presence, the control circuitry 334 can generate a control signal that causes the volume of the television 312 to be lowered by a pre-specified amount.

[0068] As another example, upon detecting that the user 308 has left the bed 302 during a specified time range (e.g., between 6:00am and 8:00am) the control circuitry 334 can generate control signals to cause the television 312 to turn on and tune to a pre-specified channel (e.g., the user 308 has indicated a preference for watching the morning news upon getting out of bed in the morning). The control circuitry 334 can generate the control signal and transmit the signal to the television 312 to cause the television 312 to turn on and tune to the desired station (which could be stored at the control circuitry 334, the television 312, or another location). As another example, upon detecting that the user 308 has gotten up for the day, the control circuitry 334 can generate and transmit control signals to cause the television 312 to turn on and begin playing a previously recorded program from a digital video recorder (DVR) in communication with the television 312.

[0069] As another example, if the television 312 is in the same room as the bed 302, the control circuitry 334 does not cause the television 312 to turn off in response to detection of user bed presence. Rather, the control circuitry 334 can generate and transmit control signals to cause the television 312 to turn off in response to determining that the user 308 is asleep. For example, the control circuitry 334 can monitor biometric signals of the user 308 (e.g., motion, heart rate, respiration rate) to determine that the user 308 has fallen asleep. Upon detecting that the user 308 is sleeping, the control circuitry 334 generates and transmits a control signal to turn the television 312 off. As another example, the control circuitry 334 can generate

the control signal to turn off the television 312 after a threshold period of time after the user 308 has fallen asleep (e.g., 10 minutes after the user has fallen asleep). As another example, the control circuitry 334 generates control signals to lower the volume of the television 312 after determining that the user 308 is asleep. As yet another example, the control circuitry 334 generates and transmits a control signal to cause the television to gradually lower in volume over a period of time and then turn off in response to determining that the user 308 is asleep.

[0070] In some implementations, the control circuitry 334 can similarly interact with other media devices, such as computers, tablets, smart phones, stereo systems, etc. For example, upon detecting that the user 308 is asleep, the control circuitry 334 can generate and transmit a control signal to the user device 310 to cause the user device 310 to turn off, or turn down the volume on a video or audio file being played by the user device 310.

[0071] The control circuitry 334 can additionally communicate with the lighting system 314, receive information from the lighting system 314, and generate control signals for controlling functions of the lighting system 314. For example, upon detecting user bed presence on the bed 302 during a certain time frame (e.g., between 8:00pm and 7:00am) that lasts for longer than a threshold period of time (e.g., 10 minutes) the control circuitry 334 of the bed 302 can determine that the user 308 is in bed for the evening. In response to this determination, the control circuitry 334 can generate control signals to cause lights in one or more rooms other than the room in which the bed 302 is located to switch off. The control signals can then be transmitted to the lighting system 314 and executed by the lighting system 314 to cause the lights in the indicated rooms to shut off. For example, the control circuitry 334 can generate and transmit control signals to turn off lights in all common rooms, but not in other bedrooms. As another example, the control signals generated by the control circuitry 334 can indicate that lights in all rooms other than the room in which the bed 302 is located are to be turned off, while one or more lights located outside of the house containing the bed 302 are to be turned on, in response to determining that the user 308 is in bed for the evening. Additionally, the control circuitry 334 can generate and transmit control signals to cause the nightlight 328 to turn on in response to determining user 308 bed presence or whether the user 308 is asleep. As another example, the control circuitry 334 can generate first control signals for turning off a

first set of lights (e.g., lights in common rooms) in response to detecting user bed presence, and second control signals for turning off a second set of lights (e.g., lights in the room in which the bed 302 is located) in response to detecting that the user 308 is asleep.

[0072] In some implementations, in response to determining that the user 308 is in bed for the evening, the control circuitry 334 of the bed 302 can generate control signals to cause the lighting system 314 to implement a sunset lighting scheme in the room in which the bed 302 is located. A sunset lighting scheme can include, for example, dimming the lights (either gradually over time, or all at once) in combination with changing the color of the light in the bedroom environment, such as adding an amber hue to the lighting in the bedroom. The sunset lighting scheme can help to put the user 308 to sleep when the control circuitry 334 has determined that the user 308 is in bed for the evening.

[0073] The control circuitry 334 can also be configured to implement a sunrise lighting scheme when the user 308 wakes up in the morning. The control circuitry 334 can determine that the user 308 is awake for the day, for example, by detecting that the user 308 has gotten off of the bed 302 (i.e., is no longer present on the bed 302) during a specified time frame (e.g., between 6:00am and 8:00am). As another example, the control circuitry 334 can monitor movement, heart rate, respiratory rate, or other biometric signals of the user 308 to determine that the user 308 is awake even though the user 308 has not gotten out of bed. If the control circuitry 334 detects that the user is awake during a specified time frame, the control circuitry 334 can determine that the user 308 is awake for the day. The specified time frame can be, for example, based on previously recorded user bed presence information collected over a period of time (e.g., two weeks) that indicates that the user 308 usually wakes up for the day between 6:30am and 7:30am. In response to the control circuitry 334 determining that the user 308 is awake, the control circuitry 334 can generate control signals to cause the lighting system 314 to implement the sunrise lighting scheme in the bedroom in which the bed 302 is located. The sunrise lighting scheme can include, for example, turning on lights (e.g., the lamp 326, or other lights in the bedroom). The sunrise lighting scheme can further include gradually increasing the level of light in the room where the bed 302 is located (or in one or more other rooms). The sunrise lighting scheme can also include only turning

on lights of specified colors. For example, the sunrise lighting scheme can include lighting the bedroom with blue light to gently assist the user 308 in waking up and becoming active.

[0074] In some implementations, the control circuitry 334 can generate different control signals for controlling actions of one or more components, such as the lighting system 314, depending on a time of day that user interactions with the bed 302 are detected. For example, the control circuitry 334 can use historical user interaction information for interactions between the user 308 and the bed 302 to determine that the user 308 usually falls asleep between 10:00pm and 11:00pm and usually wakes up between 6:30am and 7:30am on weekdays. The control circuitry 334 can use this information to generate a first set of control signals for controlling the lighting system 314 if the user 308 is detected as getting out of bed at 3:00am and to generate a second set of control signals for controlling the lighting system 314 if the user 308 is detected as getting out of bed after 6:30am. For example, if the user 308 gets out of bed prior to 6:30am, the control circuitry 334 can turn on lights that guide the user 308's route to a restroom. As another example, if the user 308 gets out of bed prior to 6:30am, the control circuitry 334 can turn on lights that guide the user 308's route to the kitchen (which can include, for example, turning on the nightlight 328, turning on under bed lighting, or turning on the lamp 326).

[0075] As another example, if the user 308 gets out of bed after 6:30am, the control circuitry 334 can generate control signals to cause the lighting system 314 to initiate a sunrise lighting scheme, or to turn on one or more lights in the bedroom and/or other rooms. In some implementations, if the user 308 is detected as getting out of bed prior to a specified morning rise time for the user 308, the control circuitry 334 causes the lighting system 314 to turn on lights that are dimmer than lights that are turned on by the lighting system 314 if the user 308 is detected as getting out of bed after the specified morning rise time. Causing the lighting system 314 to only turn on dim lights when the user 308 gets out of bed during the night (i.e., prior to normal rise time for the user 308) can prevent other occupants of the house from being woken by the lights while still allowing the user 308 to see in order to reach the restroom, kitchen, or another destination within the house.

[0076] The historical user interaction information for interactions between the user 308 and the bed 302 can be used to identify user sleep and awake time frames.

For example, user bed presence times and sleep times can be determined for a set period of time (e.g., two weeks, a month, etc.). The control circuitry 334 can then identify a typical time range or time frame in which the user 308 goes to bed, a typical time frame for when the user 308 falls asleep, and a typical time frame for when the user 308 wakes up (and in some cases, different time frames for when the user 308 wakes up and when the user 308 actually gets out of bed). In some implementations, buffer time can be added to these time frames. For example, if the user is identified as typically going to bed between 10:00pm and 10:30pm, a buffer of a half hour in each direction can be added to the time frame such that any detection of the user getting onto the bed between 9:30pm and 11:00pm is interpreted as the user 308 going to bed for the evening. As another example, detection of bed presence of the user 308 starting from a half hour before the earliest typical time that the user 308 goes to bed extending until the typical wake up time (e.g., 6:30 am) for the user can be interpreted as the user going to bed for the evening. For example, if the user typically goes to bed between 10:00pm and 10:30pm, if the user's bed presence is sensed at 12:30am one night, that can be interpreted as the user getting into bed for the evening even though this is outside of the user's typical time frame for going to bed because it has occurred prior to the user's normal wake up time. In some implementations, different time frames are identified for different times of the year (e.g., earlier bed time during winter vs. summer) or at different times of the week (e.g., user wakes up earlier on weekdays than on weekends).

[0077] The control circuitry 334 can distinguish between the user 308 going to bed for an extended period (such as for the night) as opposed to being present on the bed 302 for a shorter period (such as for a nap) by sensing duration of presence of the user 308. In some examples, the control circuitry 334 can distinguish between the user 308 going to bed for an extended period (such as for the night) as opposed to going to bed for a shorter period (such as for a nap) by sensing duration of sleep of the user 308. For example, the control circuitry 334 can set a time threshold whereby if the user 308 is sensed on the bed 302 for longer than the threshold, the user 308 is considered to have gone to bed for the night. In some examples, the threshold can be about 2 hours, whereby if the user 308 is sensed on the bed 302 for greater than 2 hours, the control circuitry 334 registers that as an extended sleep event. In other examples, the threshold can be greater than or less than two hours.

[0078] The control circuitry 334 can detect repeated extended sleep events to determine a typical bed time range of the user 308 automatically, without requiring the user 308 to enter a bed time range. This can allow the control circuitry 334 to accurately estimate when the user 308 is likely to go to bed for an extended sleep event, regardless of whether the user 308 typically goes to bed using a traditional sleep schedule or a non-traditional sleep schedule. The control circuitry 334 can then use knowledge of the bed time range of the user 308 to control one or more components (including components of the bed 302 and/or non-bed peripherals) differently based on sensing bed presence during the bed time range or outside of the bed time range.

[0079] In some examples, the control circuitry 334 can automatically determine the bed time range of the user 308 without requiring user inputs. In some examples, the control circuitry 334 can determine the bed time range of the user 308 automatically and in combination with user inputs. In some examples, the control circuitry 334 can set the bed time range directly according to user inputs. In some examples, the control circuitry 334 can associate different bed times with different days of the week. In each of these examples, the control circuitry 334 can control one or more components (such as the lighting system 314, the thermostat 316, the security system 318, the oven 322, the coffee maker 324, the lamp 326, and the nightlight 328), as a function of sensed bed presence and the bed time range.

[0080] The control circuitry 334 can additionally communicate with the thermostat 316, receive information from the thermostat 316, and generate control signals for controlling functions of the thermostat 316. For example, the user 308 can indicate user preferences for different temperatures at different times, depending on the sleep state or bed presence of the user 308. For example, the user 308 may prefer an environmental temperature of 72 degrees when out of bed, 70 degrees when in bed but awake, and 68 degrees when sleeping. The control circuitry 334 of the bed 302 can detect bed presence of the user 308 in the evening and determine that the user 308 is in bed for the night. In response to this determination, the control circuitry 334 can generate control signals to cause the thermostat to change the temperature to 70 degrees. The control circuitry 334 can then transmit the control signals to the thermostat 316. Upon detecting that the user 308 is in bed during the bed time range or asleep, the control circuitry 334 can generate and transmit control signals to cause

the thermostat 316 to change the temperature to 68. The next morning, upon determining that the user is awake for the day (e.g., the user 308 gets out of bed after 6:30am) the control circuitry 334 can generate and transmit control circuitry 334 to cause the thermostat to change the temperature to 72 degrees.

[0081] In some implementations, the control circuitry 334 can similarly generate control signals to cause one or more heating or cooling elements on the surface of the bed 302 to change temperature at various times, either in response to user interaction with the bed 302 or at various pre-programmed times. For example, the control circuitry 334 can activate a heating element to raise the temperature of one side of the surface of the bed 302 to 73 degrees when it is detected that the user 308 has fallen asleep. As another example, upon determining that the user 308 is up for the day, the control circuitry 334 can turn off a heating or cooling element. As yet another example, the user 308 can pre-program various times at which the temperature at the surface of the bed should be raised or lowered. For example, the user can program the bed 302 to raise the surface temperature to 76 degrees at 10:00pm, and lower the surface temperature to 68 degrees at 11:30pm.

[0082] In some implementations, in response to detecting user bed presence of the user 308 and/or that the user 308 is asleep, the control circuitry 334 can cause the thermostat 316 to change the temperature in different rooms to different values. For example, in response to determining that the user 308 is in bed for the evening, the control circuitry 334 can generate and transmit control signals to cause the thermostat 316 to set the temperature in one or more bedrooms of the house to 72 degrees and set the temperature in other rooms to 67 degrees.

[0083] The control circuitry 334 can also receive temperature information from the thermostat 316 and use this temperature information to control functions of the bed 302 or other devices. For example, as discussed above, the control circuitry 334 can adjust temperatures of heating elements included in the bed 302 in response to temperature information received from the thermostat 316.

[0084] In some implementations, the control circuitry 334 can generate and transmit control signals for controlling other temperature control systems. For example, in response to determining that the user 308 is awake for the day, the control circuitry 334 can generate and transmit control signals for causing floor heating elements to activate. For example, the control circuitry 334 can cause a floor heating

system for a master bedroom to turn on in response to determining that the user 308 is awake for the day.

[0085] The control circuitry 334 can additionally communicate with the security system 318, receive information from the security system 318, and generate control signals for controlling functions of the security system 318. For example, in response to detecting that the user 308 is in bed for the evening, the control circuitry 334 can generate control signals to cause the security system to engage or disengage security functions. The control circuitry 334 can then transmit the control signals to the security system 318 to cause the security system 318 to engage. As another example, the control circuitry 334 can generate and transmit control signals to cause the security system 318 to disable in response to determining that the user 308 is awake for the day (e.g., user 308 is no longer present on the bed 302 after 6:00am). In some implementations, the control circuitry 334 can generate and transmit a first set of control signals to cause the security system 318 to engage a first set of security features in response to detecting user bed presence of the user 308, and can generate and transmit a second set of control signals to cause the security system 318 to engage a second set of security features in response to detecting that the user 308 has fallen asleep.

[0086] In some implementations, the control circuitry 334 can receive alerts from the security system 318 and indicate the alert to the user 308. For example, the control circuitry 334 can detect that the user 308 is in bed for the evening and in response, generate and transmit control signals to cause the security system 318 to engage or disengage. The security system can then detect a security breach (e.g., someone has opened the door 332 without entering the security code, or someone has opened a window when the security system 318 is engaged). The security system 318 can communicate the security breach to the control circuitry 334 of the bed 302. In response to receiving the communication from the security system 318, the control circuitry 334 can generate control signals to alert the user 308 to the security breach. For example, the control circuitry 334 can cause the bed 302 to vibrate. As another example, the control circuitry 334 can cause portions of the bed 302 to articulate (e.g., cause the head section to raise or lower) in order to wake the user 308 and alert the user to the security breach. As another example, the control circuitry 334 can generate and transmit control signals to cause the lamp 326 to flash on and off at

regular intervals to alert the user 308 to the security breach. As another example, the control circuitry 334 can alert the user 308 of one bed 302 regarding a security breach in a bedroom of another bed, such as an open window in a kid's bedroom. As another example, the control circuitry 334 can send an alert to a garage door controller (e.g., to close and lock the door). As another example, the control circuitry 334 can send an alert for the security to be disengaged.

[0087] The control circuitry 334 can additionally generate and transmit control signals for controlling the garage door 320 and receive information indicating a state of the garage door 320 (i.e., open or closed). For example, in response to determining that the user 308 is in bed for the evening, the control circuitry 334 can generate and transmit a request to a garage door opener or another device capable of sensing if the garage door 320 is open. The control circuitry 334 can request information on the current state of the garage door 320. If the control circuitry 334 receives a response (e.g., from the garage door opener) indicating that the garage door 320 is open, the control circuitry 334 can either notify the user 308 that the garage door is open, or generate a control signal to cause the garage door opener to close the garage door 320. For example, the control circuitry 334 can send a message to the user device 310 indicating that the garage door is open. As another example, the control circuitry 334 can cause the bed 302 to vibrate. As yet another example, the control circuitry 334 can generate and transmit a control signal to cause the lighting system 314 to cause one or more lights in the bedroom to flash to alert the user 308 to check the user device 310 for an alert (in this example, an alert regarding the garage door 320 being open). Alternatively, or additionally, the control circuitry 334 can generate and transmit control signals to cause the garage door opener to close the garage door 320 in response to identifying that the user 308 is in bed for the evening and that the garage door 320 is open. In some implementations, control signals can vary depend on the age of the user 308.

[0088] The control circuitry 334 can similarly send and receive communications for controlling or receiving state information associated with the door 332 or the oven 322. For example, upon detecting that the user 308 is in bed for the evening, the control circuitry 334 can generate and transmit a request to a device or system for detecting a state of the door 332. Information returned in response to the request can indicate various states for the door 332 such as open, closed but

unlocked, or closed and locked. If the door 332 is open or closed but unlocked, the control circuitry 334 can alert the user 308 to the state of the door, such as in a manner described above with reference to the garage door 320. Alternatively, or in addition to alerting the user 308, the control circuitry 334 can generate and transmit control signals to cause the door 332 to lock, or to close and lock. If the door 332 is closed and locked, the control circuitry 334 can determine that no further action is needed.

[0089] Similarly, upon detecting that the user 308 is in bed for the evening, the control circuitry 334 can generate and transmit a request to the oven 322 to request a state of the oven 322 (e.g., on or off). If the oven 322 is on, the control circuitry 334 can alert the user 308 and/or generate and transmit control signals to cause the oven 322 to turn off. If the oven is already off, the control circuitry 334 can determine that no further action is necessary. In some implementations, different alerts can be generated for different events. For example, the control circuitry 334 can cause the lamp 326 (or one or more other lights, via the lighting system 314) to flash in a first pattern if the security system 318 has detected a breach, flash in a second pattern if garage door 320 is on, flash in a third pattern if the door 332 is open, flash in a fourth pattern if the oven 322 is on, and flash in a fifth pattern if another bed has detected that a user of that bed has gotten up (e.g., that a child of the user 308 has gotten out of bed in the middle of the night as sensed by a sensor in the bed 302 of the child). Other examples of alerts that can be processed by the control circuitry 334 of the bed 302 and communicated to the user include a smoke detector detecting smoke (and communicating this detection of smoke to the control circuitry 334), a carbon monoxide tester detecting carbon monoxide, a heater malfunctioning, or an alert from any other device capable of communicating with the control circuitry 334 and detecting an occurrence that should be brought to the user 308's attention.

[0090] The control circuitry 334 can also communicate with a system or device for controlling a state of the window blinds 330. For example, in response to determining that the user 308 is in bed for the evening, the control circuitry 334 can generate and transmit control signals to cause the window blinds 330 to close. As another example, in response to determining that the user 308 is up for the day (e.g., user has gotten out of bed after 6:30am) the control circuitry 334 can generate and transmit control signals to cause the window blinds 330 to open. By contrast, if the user 308 gets out of bed prior to a normal rise time for the user 308, the control

circuitry 334 can determine that the user 308 is not awake for the day and does not generate control signals for causing the window blinds 330 to open. As yet another example, the control circuitry 334 can generate and transmit control signals that cause a first set of blinds to close in response to detecting user bed presence of the user 308 and a second set of blinds to close in response to detecting that the user 308 is asleep.

[0091] The control circuitry 334 can generate and transmit control signals for controlling functions of other household devices in response to detecting user interactions with the bed 302. For example, in response to determining that the user 308 is awake for the day, the control circuitry 334 can generate and transmit control signals to the coffee maker 324 to cause the coffee maker 324 to begin brewing coffee. As another example, the control circuitry 334 can generate and transmit control signals to the oven 322 to cause the oven to begin preheating (for users that like fresh baked bread in the morning). As another example, the control circuitry 334 can use information indicating that the user 308 is awake for the day along with information indicating that the time of year is currently winter and/or that the outside temperature is below a threshold value to generate and transmit control signals to cause a car engine block heater to turn on.

[0092] As another example, the control circuitry 334 can generate and transmit control signals to cause one or more devices to enter a sleep mode in response to detecting user bed presence of the user 308, or in response to detecting that the user 308 is asleep. For example, the control circuitry 334 can generate control signals to cause a mobile phone of the user 308 to switch into sleep mode. The control circuitry 334 can then transmit the control signals to the mobile phone. Later, upon determining that the user 308 is up for the day, the control circuitry 334 can generate and transmit control signals to cause the mobile phone to switch out of sleep mode.

[0093] In some implementations, the control circuitry 334 can communicate with one or more noise control devices. For example, upon determining that the user 308 is in bed for the evening, or that the user 308 is asleep, the control circuitry 334 can generate and transmit control signals to cause one or more noise cancellation devices to activate. The noise cancellation devices can, for example, be included as part of the bed 302 or located in the bedroom with the bed 302. As another example, upon determining that the user 308 is in bed for the evening or that the user 308 is

asleep, the control circuitry 334 can generate and transmit control signals to turn the volume on, off, up, or down, for one or more sound generating devices, such as a stereo system radio, computer, tablet, etc.

[0094] Additionally, functions of the bed 302 are controlled by the control circuitry 334 in response to user interactions with the bed 302. For example, the bed 302 can include an adjustable foundation and an articulation controller configured to adjust the position of one or more portions of the bed 302 by adjusting the adjustable foundation that supports the bed. For example, the articulation controller can adjust the bed 302 from a flat position to a position in which a head portion of a mattress of the bed 302 is inclined upward (e.g., to facilitate a user sitting up in bed and/or watching television). In some implementations, the bed 302 includes multiple separately articulable sections. For example, portions of the bed corresponding to the locations of the air chambers 306a and 306b can be articulated independently from each other, to allow one person positioned on the bed 302 surface to rest in a first position (e.g., a flat position) while a second person rests in a second position (e.g., a reclining position with the head raised at an angle from the waist). In some implementations, separate positions can be set for two different beds (e.g., two twin beds placed next to each other). The foundation of the bed 302 can include more than one zone that can be independently adjusted. The articulation controller can also be configured to provide different levels of massage to one or more users on the bed 302 or to cause the bed to vibrate to communicate alerts to the user 308 as described above.

[0095] The control circuitry 334 can adjust positions (e.g., incline and decline positions for the user 308 and/or an additional user of the bed 302) in response to user interactions with the bed 302. For example, the control circuitry 334 can cause the articulation controller to adjust the bed 302 to a first recline position for the user 308 in response to sensing user bed presence for the user 308. The control circuitry 334 can cause the articulation controller to adjust the bed 302 to a second recline position (e.g., a less reclined, or flat position) in response to determining that the user 308 is asleep. As another example, the control circuitry 334 can receive a communication from the television 312 indicating that the user 308 has turned off the television 312, and in response the control circuitry 334 can cause the articulation controller to adjust the position of the bed 302 to a preferred user sleeping position (e.g., due to the user

turning off the television 312 while the user 308 is in bed indicating that the user 308 wishes to go to sleep).

[0096] In some implementations, the control circuitry 334 can control the articulation controller so as to wake up one user of the bed 302 without waking another user of the bed 302. For example, the user 308 and a second user of the bed 302 can each set distinct wakeup times (e.g., 6:30am and 7:15am respectively). When the wakeup time for the user 308 is reached, the control circuitry 334 can cause the articulation controller to vibrate or change the position of only a side of the bed on which the user 308 is located to wake the user 308 without disturbing the second user. When the wakeup time for the second user is reached, the control circuitry 334 can cause the articulation controller to vibrate or change the position of only the side of the bed on which the second user is located. Alternatively, when the second wakeup time occurs, the control circuitry 334 can utilize other methods (such as audio alarms, or turning on the lights) to wake the second user since the user 308 is already awake and therefore will not be disturbed when the control circuitry 334 attempts to wake the second user.

[0097] Still referring to FIG. 3, the control circuitry 334 for the bed 302 can utilize information for interactions with the bed 302 by multiple users to generate control signals for controlling functions of various other devices. For example, the control circuitry 334 can wait to generate control signals for, for example, engaging the security system 318, or instructing the lighting system 314 to turn off lights in various rooms until both the user 308 and a second user are detected as being present on the bed 302. As another example, the control circuitry 334 can generate a first set of control signals to cause the lighting system 314 to turn off a first set of lights upon detecting bed presence of the user 308 and generate a second set of control signals for turning off a second set of lights in response to detecting bed presence of a second user. As another example, the control circuitry 334 can wait until it has been determined that both the user 308 and a second user are awake for the day before generating control signals to open the window blinds 330. As yet another example, in response to determining that the user 308 has left the bed and is awake for the day, but that a second user is still sleeping, the control circuitry 334 can generate and transmit a first set of control signals to cause the coffee maker 324 to begin brewing coffee, to cause the security system 318 to deactivate, to turn on the lamp 326, to turn off the

nightlight 328, to cause the thermostat 316 to raise the temperature in one or more rooms to 72 degrees, and to open blinds (e.g., the window blinds 330) in rooms other than the bedroom in which the bed 302 is located. Later, in response to detecting that the second user is no longer present on the bed (or that the second user is awake) the control circuitry 334 can generate and transmit a second set of control signals to, for example, cause the lighting system 314 to turn on one or more lights in the bedroom, to cause window blinds in the bedroom to open, and to turn on the television 312 to a pre-specified channel.

[0098] Examples of Data Processing Systems Associated with a Bed

[0099] Described here are examples of systems and components that can be used for data processing tasks that are, for example, associated with a bed. In some cases, multiple examples of a particular component or group of components are presented. Some of these examples are redundant and/or mutually exclusive alternatives. Connections between components are shown as examples to illustrate possible network configurations for allowing communication between components. Different formats of connections can be used as technically needed or desired. The connections generally indicate a logical connection that can be created with any technologically feasible format. For example, a network on a motherboard can be created with a printed circuit board, wireless data connections, and/or other types of network connections. Some logical connections are not shown for clarity. For example, connections with power supplies and/or computer readable memory may not be shown for clarity sake, as many or all elements of a particular component may need to be connected to the power supplies and/or computer readable memory.

[00100] FIG. 4A is a block diagram of an example of a data processing system 400 that can be associated with a bed system, including those described above with respect to FIGS. 1-3. This system 400 includes a pump motherboard 402 and a pump daughterboard 404. The system 400 includes a sensor array 406 that can include one or more sensors configured to sense physical phenomenon of the environment and/or bed, and to report such sensing back to the pump motherboard 402 for, for example, analysis. The system 400 also includes a controller array 408 that can include one or more controllers configured to control logic-controlled devices of the bed and/or environment. The pump motherboard 400 can be in communication with one or more computing devices 414 and one or more cloud services 410 over local networks, the

Internet 412, or otherwise as is technically appropriate. Each of these components will be described in more detail, some with multiple example configurations, below.

[00101] In this example, a pump motherboard 402 and a pump daughterboard 404 are communicably coupled. They can be conceptually described as a center or hub of the system 400, with the other components conceptually described as spokes of the system 400. In some configurations, this can mean that each of the spoke components communicates primarily or exclusively with the pump motherboard 402. For example, a sensor of the sensor array may not be configured to, or may not be able to, communicate directly with a corresponding controller. Instead, each spoke component can communicate with the motherboard 402. The sensor of the sensor array 406 can report a sensor reading to the motherboard 402, and the motherboard 402 can determine that, in response, a controller of the controller array 408 should adjust some parameters of a logic controlled device or otherwise modify a state of one or more peripheral devices. In one case, if the temperature of the bed is determined to be too hot, the pump motherboard 402 can determine that a temperature controller should cool the bed.

[00102] One advantage of a hub-and-spoke network configuration, sometimes also referred to as a star-shaped network, is a reduction in network traffic compared to, for example, a mesh network with dynamic routing. If a particular sensor generates a large, continuous stream of traffic, that traffic may only be transmitted over one spoke of the network to the motherboard 402. The motherboard 402 can, for example, marshal that data and condense it to a smaller data format for retransmission for storage in a cloud service 410. Additionally or alternatively, the motherboard 402 can generate a single, small, command message to be sent down a different spoke of the network in response to the large stream. For example, if the large stream of data is a pressure reading that is transmitted from the sensor array 406 a few times a second, the motherboard 402 can respond with a single command message to the controller array to increase the pressure in an air chamber. In this case, the single command message can be orders of magnitude smaller than the stream of pressure readings.

[00103] As another advantage, a hub-and-spoke network configuration can allow for an extensible network that can accommodate components being added, removed, failing, etc. This can allow, for example, more, fewer, or different sensors in the sensor array 406, controllers in the controller array 408, computing devices 414,

and/or cloud services 410. For example, if a particular sensor fails or is deprecated by a newer version of the sensor, the system 400 can be configured such that only the motherboard 402 needs to be updated about the replacement sensor. This can allow, for example, product differentiation where the same motherboard 402 can support an entry level product with fewer sensors and controllers, a higher value product with more sensors and controllers, and customer personalization where a customer can add their own selected components to the system 400.

[00104] Additionally, a line of air bed products can use the system 400 with different components. In an application in which every air bed in the product line includes both a central logic unit and a pump, the motherboard 402 (and optionally the daughterboard 404) can be designed to fit within a single, universal housing. Then, for each upgrade of the product in the product line, additional sensors, controllers, cloud services, etc., can be added. Design, manufacturing, and testing time can be reduced by designing all products in a product line from this base, compared to a product line in which each product has a bespoke logic control system.

[00105] Each of the components discussed above can be realized in a wide variety of technologies and configurations. Below, some examples of each component will be further discussed. In some alternatives, two or more of the components of the system 400 can be realized in a single alternative component; some components can be realized in multiple, separate components; and/or some functionality can be provided by different components.

[00106] FIG. 4B is a block diagram showing some communication paths of the data processing system 400. As previously described, the motherboard 402 and the pump daughterboard 404 may act as a hub for peripheral devices and cloud services of the system 400. In cases in which the pump daughterboard 404 communicates with cloud services or other components, communications from the pump daughterboard 404 may be routed through the pump motherboard 402. This may allow, for example, the bed to have only a single connection with the internet 412. The computing device 414 may also have a connection to the internet 412, possibly through the same gateway used by the bed and/or possibly through a different gateway (e.g., a cell service provider).

[00107] Previously, a number of cloud services 410 were described. As shown in FIG. 4B, some cloud services, such as cloud services 410d and 410e, may be

configured such that the pump motherboard 402 can communicate with the cloud service directly – that is the motherboard 402 may communicate with a cloud service 410 without having to use another cloud service 410 as an intermediary. Additionally or alternatively, some cloud services 410, for example cloud service 410f, may only be reachable by the pump motherboard 402 through an intermediary cloud service, for example cloud service 410e. While not shown here, some cloud services 410 may be reachable either directly or indirectly by the pump motherboard 402.

[00108] Additionally, some or all of the cloud services 410 may be configured to communicate with other cloud services. This communication may include the transfer of data and/or remote function calls according to any technologically appropriate format. For example, one cloud service 410 may request a copy for another cloud service's 410 data, for example, for purposes of backup, coordination, migration, or for performance of calculations or data mining. In another example, many cloud services 410 may contain data that is indexed according to specific users tracked by the user account cloud 410c and/or the bed data cloud 410a. These cloud services 410 may communicate with the user account cloud 410c and/or the bed data cloud 410a when accessing data specific to a particular user or bed.

[00109] FIG. 5 is a block diagram of an example of a motherboard 402 that can be used in a data processing system that can be associated with a bed system, including those described above with respect to FIGS. 1-3. In this example, compared to other examples described below, this motherboard 402 consists of relatively fewer parts and can be limited to provide a relatively limited feature set.

[00110] The motherboard includes a power supply 500, a processor 502, and computer memory 512. In general, the power supply includes hardware used to receive electrical power from an outside source and supply it to components of the motherboard 402. The power supply can include, for example, a battery pack and/or wall outlet adapter, an AC to DC converter, a DC to AC converter, a power conditioner, a capacitor bank, and/or one or more interfaces for providing power in the current type, voltage, etc., needed by other components of the motherboard 402.

[00111] The processor 502 is generally a device for receiving input, performing logical determinations, and providing output. The processor 502 can be a central processing unit, a microprocessor, general purpose logic circuitry, application-specific

integrated circuitry, a combination of these, and/or other hardware for performing the functionality needed.

[00112] The memory 512 is generally one or more devices for storing data. The memory 512 can include long term stable data storage (e.g., on a hard disk), short term unstable (e.g., on Random Access Memory) or any other technologically appropriate configuration.

[00113] The motherboard 402 includes a pump controller 504 and a pump motor 506. The pump controller 504 can receive commands from the processor 502 and, in response, control the function of the pump motor 506. For example, the pump controller 504 can receive, from the processor 502, a command to increase the pressure of an air chamber by 0.3 pounds per square inch (PSI). The pump controller 504, in response, engages a valve so that the pump motor 506 is configured to pump air into the selected air chamber, and can engage the pump motor 506 for a length of time that corresponds to 0.3 PSI or until a sensor indicates that pressure has been increased by 0.3 PSI. In an alternative configuration, the message can specify that the chamber should be inflated to a target PSI, and the pump controller 504 can engage the pump motor 506 until the target PSI is reached.

[00114] A valve solenoid 508 can control which air chamber a pump is connected to. In some cases, the solenoid 508 can be controlled by the processor 502 directly. In some cases, the solenoid 508 can be controlled by the pump controller 504.

[00115] A remote interface 510 of the motherboard 402 can allow the motherboard 402 to communicate with other components of a data processing system. For example, the motherboard 402 can be able to communicate with one or more daughterboards, with peripheral sensors, and/or with peripheral controllers through the remote interface 510. The remote interface 510 can provide any technologically appropriate communication interface, including but not limited to multiple communication interfaces such as WiFi, Bluetooth, and copper wired networks.

[00116] FIG. 6 is a block diagram of an example of a motherboard 402 that can be used in a data processing system that can be associated with a bed system, including those described above with respect to FIGS. 1-3. Compared to the motherboard 402 described with reference to FIG. 5, the motherboard in FIG. 6 can contain more components and provide more functionality in some applications.

[00117] In addition to the power supply 500, processor 502, pump controller 504, pump motor 506, and valve solenoid 508, this motherboard 402 is shown with a valve controller 600, a pressure sensor 602, a universal serial bus (USB) stack 604, a WiFi radio 606, a Bluetooth Low Energy (BLE) radio 608, a ZigBee radio 610, a Bluetooth radio 612 and a computer memory 512.

[00118] Similar to the way that the pump controller 504 converts commands from the processor 502 into control signals for the pump motor 506, the valve controller 600 can convert commands from the processor 502 into control signals for the valve solenoid 508. In one example, the processor 502 can issue a command to the valve controller 600 to connect the pump to a particular air chamber out of the group of air chambers in an air bed. The valve controller 600 can control the position of the valve solenoid 508 so that the pump is connected to the indicated air chamber.

[00119] The pressure sensor 602 can read pressure readings from one or more air chambers of the air bed. The pressure sensor 602 can also preform digital sensor conditioning.

[00120] The motherboard 402 can include a suite of network interfaces, including but not limited to those shown here. These network interfaces can allow the motherboard to communicate over a wired or wireless network with any number of devices, including but not limited to peripheral sensors, peripheral controllers, computing devices, and devices and services connected to the Internet 412.

[00121] FIG. 7 is a block diagram of an example of a daughterboard 404 that can be used in a data processing system that can be associated with a bed system, including those described above with respect to FIGS. 1-3. In some configurations, one or more daughterboards 404 can be connected to the motherboard 402. Some daughterboards 404 can be designed to offload particular and/or compartmentalized tasks from the motherboard 402. This can be advantageous, for example, if the particular tasks are computationally intensive, proprietary, or subject to future revisions. For example, the daughterboard 404 can be used to calculate a particular sleep data metric. This metric can be computationally intensive, and calculating the sleep metric on the daughterboard 404 can free up the resources of the motherboard 402 while the metric is being calculated. Additionally and/or alternatively, the sleep metric can be subject to future revisions. To update the system 400 with the new sleep metric, it is possible that only the daughterboard 404 that calculates that metric

need be replaced. In this case, the same motherboard 402 and other components can be used, saving the need to perform unit testing of additional components instead of just the daughterboard 404.

[00122] The daughterboard 404 is shown with a power supply 700, a processor 702, computer readable memory 704, a pressure sensor 706, and a WiFi radio 708. The processor can use the pressure sensor 706 to gather information about the pressure of the air chamber or chambers of an air bed. From this data, the processor 702 can perform an algorithm to calculate a sleep metric. In some examples, the sleep metric can be calculated from only the pressure of air chambers. In other examples, the sleep metric can be calculated from one or more other sensors. In an example in which different data is needed, the processor 702 can receive that data from an appropriate sensor or sensors. These sensors can be internal to the daughterboard 404, accessible via the WiFi radio 708, or otherwise in communication with the processor 702. Once the sleep metric is calculated, the processor 702 can report that sleep metric to, for example, the motherboard 402.

[00123] FIG. 8 is a block diagram of an example of a motherboard 800 with no daughterboard that can be used in a data processing system that can be associated with a bed system, including those described above with respect to FIGS. 1-3. In this example, the motherboard 800 can perform most, all, or more of the features described with reference to the motherboard 402 in FIG. 6 and the daughterboard 404 in FIG. 7.

[00124] FIG. 9 is a block diagram of an example of a sensory array 406 that can be used in a data processing system that can be associated with a bed system, including those described above with respect to FIGS. 1-3. In general, the sensor array 406 is a conceptual grouping of some or all the peripheral sensors that communicate with the motherboard 402 but are not native to the motherboard 402.

[00125] The peripheral sensors of the sensor array 406 can communicate with the motherboard 402 through one or more of the network interfaces of the motherboard, including but not limited to the USB stack 604, a WiFi radio 606, a Bluetooth Low Energy (BLE) radio 608, a ZigBee radio 610, and a Bluetooth radio 612, as is appropriate for the configuration of the particular sensor. For example, a sensor that outputs a reading over a USB cable can communicate through the USB stack 604.

[00126] Some of the peripheral sensors 900 of the sensor array 406 can be bed mounted 900. These sensors can be, for example, embedded into the structure of a bed and sold with the bed, or later affixed to the structure of the bed. Other peripheral sensors 902 and 904 can be in communication with the motherboard 402, but optionally not mounted to the bed. In some cases, some or all of the bed mounted sensors 900 and/or peripheral sensors 902 and 904 can share networking hardware, including a conduit that contains wires from each sensor, a multi-wire cable or plug that, when affixed to the motherboard 402, connect all of the associated sensors with the motherboard 402. In some embodiments, one, some, or all of sensors 902, 904, 906, 908, and 910 can sense one or more features of a mattress, such as pressure, temperature, light, sound, and/or one or more other features of the mattress. In some embodiments, one, some, or all of sensors 902, 904, 906, 908, and 910 can sense one or more features external to the mattress. In some embodiments, pressure sensor 902 can sense pressure of the mattress while some or all of sensors 902, 904, 906, 908, and 910 can sense one or more features of the mattress and/or external to the mattress.

[00127] FIG. 10 is a block diagram of an example of a controller array 408 that can be used in a data processing system that can be associated with a bed system, including those described above with respect to FIGS. 1-3. In general, the controller array 408 is a conceptual grouping of some or all peripheral controllers that communicate with the motherboard 402 but are not native to the motherboard 402.

[00128] The peripheral controllers of the controller array 408 can communicate with the motherboard 402 through one or more of the network interfaces of the motherboard, including but not limited to the USB stack 604, a WiFi radio 606, a Bluetooth Low Energy (BLE) radio 608, a ZigBee radio 610, and a Bluetooth radio 612, as is appropriate for the configuration of the particular sensor. For example, a controller that receives a command over a USB cable can communicate through the USB stack 604.

[00129] Some of the controllers of the controller array 408 can be bed mounted 1000. These controllers can be, for example, embedded into the structure of a bed and sold with the bed, or later affixed to the structure of the bed. Other peripheral controllers 1002 and 1004 can be in communication with the motherboard 402, but optionally not mounted to the bed. In some cases, some or all of the bed mounted controllers 1000 and/or peripheral controllers 1002 and 1004 can share networking

hardware, including a conduit that contains wires for each controller, a multi-wire cable or plug that, when affixed to the motherboard 402, connects all of the associated controllers with the motherboard 402.

[00130] FIG. 11 is a block diagram of an example of a computing device 412 that can be used in a data processing system that can be associated with a bed system, including those described above with respect to FIGS. 1-3. The computing device 412 can include, for example, computing devices used by a user of a bed. Example computing devices 412 include, but are not limited to, mobile computing devices (e.g., mobile phones, tablet computers, laptops) and desktop computers.

[00131] The computing device 412 includes a power supply 1100, a processor 1102, and computer readable memory 1104. User input and output can be transmitted by, for example, speakers 1106, a touchscreen 1108, or other not shown components such as a pointing device or keyboard. The computing device 412 can run one or more applications 1110. These applications can include, for example, application to allow the user to interact with the system 400. These applications can allow a user to view information about the bed (e.g., sensor readings, sleep metrics), or configure the behavior of the system 400 (e.g., set a desired firmness to the bed, set desired behavior for peripheral devices). In some cases, the computing device 412 can be used in addition to, or to replace, the remote control 122 described previously.

[00132] FIG. 12 is a block diagram of an example bed data cloud service 410a that can be used in a data processing system that can be associated with a bed system, including those described above with respect to FIGS. 1-3. In this example, the bed data cloud service 410a is configured to collect sensor data and sleep data from a particular bed, and to match the sensor and sleep data with one or more users that use the bed when the sensor and sleep data was generated.

[00133] The bed data cloud service 410a is shown with a network interface 1200, a communication manager 1202, server hardware 1204, and server system software 1206. In addition, the bed data cloud service 410a is shown with a user identification module 1208, a device management 1210 module, a sensor data module 1210, and an advanced sleep data module 1214.

[00134] The network interface 1200 generally includes hardware and low level software used to allow one or more hardware devices to communicate over networks. For example the network interface 1200 can include network cards, routers, modems,

and other hardware needed to allow the components of the bed data cloud service 410a to communicate with each other and other destinations over, for example, the Internet 412. The communication manager 1202 generally comprises hardware and software that operate above the network interface 1200. This includes software to initiate, maintain, and tear down network communications used by the bed data cloud service 410a. This includes, for example, TCP/IP, SSL or TLS, Torrent, and other communication sessions over local or wide area networks. The communication manager 1202 can also provide load balancing and other services to other elements of the bed data cloud service 410a.

[00135] The server hardware 1204 generally includes the physical processing devices used to instantiate and maintain bed data cloud service 410a. This hardware includes, but is not limited to processors (e.g., central processing units, ASICs, graphical processors), and computer readable memory (e.g., random access memory, stable hard disks, tape backup). One or more servers can be configured into clusters, multi-computer, or datacenters that can be geographically separate or connected.

[00136] The server system software 1206 generally includes software that runs on the server hardware 1204 to provide operating environments to applications and services. The server system software 1206 can include operating systems running on real servers, virtual machines instantiated on real servers to create many virtual servers, server level operations such as data migration, redundancy, and backup.

[00137] The user identification 1208 can include, or reference, data related to users of beds with associated data processing systems. For example, the users can include customers, owners, or other users registered with the bed data cloud service 410a or another service. Each user can have, for example, a unique identifier, user credentials, contact information, billing information, demographic information, or any other technologically appropriate information.

[00138] The device manager 1210 can include, or reference, data related to beds or other products associated with data processing systems. For example, the beds can include products sold or registered with a system associated with the bed data cloud service 410a. Each bed can have, for example, a unique identifier, model and/or serial number, sales information, geographic information, delivery information, a listing of associated sensors and control peripherals, etc. Additionally, an index or indexes stored by the bed data cloud service 410a can identify users that are

associated with beds. For example, this index can record sales of a bed to a user, users that sleep in a bed, etc.

[00139] The sensor data 1212 can record raw or condensed sensor data recorded by beds with associated data processing systems. For example, a bed's data processing system can have a temperature sensor, pressure sensor, and light sensor. Readings from these sensors, either in raw form or in a format generated from the raw data (e.g. sleep metrics) of the sensors, can be communicated by the bed's data processing system to the bed data cloud service 410a for storage in the sensor data 1212. Additionally, an index or indexes stored by the bed data cloud service 410a can identify users and/or beds that are associated with the sensor data 1212.

[00140] The bed data cloud service 410a can use any of its available data to generate advanced sleep data 1214. In general, the advanced sleep data 1214 includes sleep metrics and other data generated from sensor readings. Some of these calculations can be performed in the bed data cloud service 410a instead of locally on the bed's data processing system, for example, because the calculations are computationally complex or require a large amount of memory space or processor power that is not available on the bed's data processing system. This can help allow a bed system to operate with a relatively simple controller and still be part of a system that performs relatively complex tasks and computations.

[00141] FIG. 13 is a block diagram of an example sleep data cloud service 410b that can be used in a data processing system that can be associated with a bed system, including those described above with respect to FIGS. 1-3. In this example, the sleep data cloud service 410b is configured to record data related to users' sleep experience.

[00142] The sleep data cloud service 410b is shown with a network interface 1300, a communication manager 1302, server hardware 1304, and server system software 1306. In addition, the sleep data cloud service 410b is shown with a user identification module 1308, a pressure sensor manager 1310, a pressure based sleep data module 1312, a raw pressure sensor data module 1314, and a non-pressure sleep data module 1316.

[00143] The pressure sensor manager 1310 can include, or reference, data related to the configuration and operation of pressure sensors in beds. For example,

this data can include an identifier of the types of sensors in a particular bed, their settings and calibration data, etc.

[00144] The pressure based sleep data 1312 can use raw pressure sensor data 1314 to calculate sleep metrics specifically tied to pressure sensor data. For example, user presence, movements, weight change, heart rate, and breathing rate can all be determined from raw pressure sensor data 1314. Additionally, an index or indexes stored by the sleep data cloud service 410b can identify users that are associated with pressure sensors, raw pressure sensor data, and/or pressure based sleep data.

[00145] The non-pressure sleep data 1316 can use other sources of data to calculate sleep metrics. For example, user entered preferences, light sensor readings, and sound sensor readings can all be used to track sleep data. Additionally, an index or indexes stored by the sleep data cloud service 410b can identify users that are associated with other sensors and/or non-pressure sleep data 1316.

[00146] FIG. 14 is a block diagram of an example user account cloud service 410c that can be used in a data processing system that can be associated with a bed system, including those described above with respect to FIGS. 1-3. In this example, the user account cloud service 410c is configured to record a list of users and to identify other data related to those users.

[00147] The user account cloud service 410c is shown with a network interface 1400, a communication manager 1402, server hardware 1404, and server system software 1406. In addition, the user account cloud service 410c is shown with a user identification module 1408, a purchase history module 1410, an engagement module 1412, and an application usage history module 1414.

[00148] The user identification module 1408 can include, or reference, data related to users of beds with associated data processing systems. For example, the users can include customers, owners, or other users registered with the user account cloud service 410a or another service. Each user can have, for example, a unique identifier, and user credentials, demographic information, or any other technologically appropriate information.

[00149] The purchase history module 1410 can include, or reference, data related to purchases by users. For example, the purchase data can include a sale's contact information, billing information, and salesperson information. Additionally,

an index or indexes stored by the user account cloud service 410c can identify users that are associated with a purchase.

[00150] The engagement 1412 can track user interactions with the manufacturer, vendor, and/or manager of the bed and or cloud services. This engagement data can include communications (e.g., emails, service calls), data from sales (e.g., sales receipts, configuration logs), and social network interactions.

[00151] The usage history module 1414 can contain data about user interactions with one or more applications and/or remote controls of a bed. For example, a monitoring and configuration application can be distributed to run on, for example, computing devices 412. This application can log and report user interactions for storage in the application usage history module 1414. Additionally, an index or indexes stored by the user account cloud service 410c can identify users that are associated with each log entry.

[00152] FIG. 15 is a block diagram of an example point of sale cloud service 1500 that can be used in a data processing system that can be associated with a bed system, including those described above with respect to FIGS. 1-3. In this example, the point of sale cloud service 1500 is configured to record data related to users' purchases.

[00153] The point of sale cloud service 1500 is shown with a network interface 1502, a communication manager 1504, server hardware 1506, and server system software 1508. In addition, the point of sale cloud service 1500 is shown with a user identification module 1510, a purchase history module 1512, and a setup module 1514.

[00154] The purchase history module 1512 can include, or reference, data related to purchases made by users identified in the user identification module 1510. The purchase information can include, for example, data of a sale, price, and location of sale, delivery address, and configuration options selected by the users at the time of sale. These configuration options can include selections made by the user about how they wish their newly purchased beds to be setup and can include, for example, expected sleep schedule, a listing of peripheral sensors and controllers that they have or will install, etc.

[00155] The bed setup module 1514 can include, or reference, data related to installations of beds that users' purchase. The bed setup data can include, for

example, the date and address to which a bed is delivered, the person that accepts delivery, the configuration that is applied to the bed upon delivery, the name or names of the person or people who will sleep on the bed, which side of the bed each person will use, etc.

[00156] Data recorded in the point of sale cloud service 1500 can be referenced by a user's bed system at later dates to control functionality of the bed system and/or to send control signals to peripheral components according to data recorded in the point of sale cloud service 1500. This can allow a salesperson to collect information from the user at the point of sale that later facilitates automation of the bed system. In some examples, some or all aspects of the bed system can be automated with little or no user-entered data required after the point of sale. In other examples, data recorded in the point of sale cloud service 1500 can be used in connection with a variety of additional data gathered from user-entered data.

[00157] FIG. 16 is a block diagram of an example environment cloud service 1600 that can be used in a data processing system that can be associated with a bed system, including those described above with respect to FIGS. 1-3. In this example, the environment cloud service 1600 is configured to record data related to users' home environment.

[00158] The environment cloud service 1600 is shown with a network interface 1602, a communication manager 1604, server hardware 1606, and server system software 1608. In addition, the environment cloud service 1600 is shown with a user identification module 1610, an environmental sensor module 1612, and an environmental factors module 1614.

[00159] The environmental sensors module 1612 can include a listing of sensors that users' in the user identification module 1610 have installed in their bed. These sensors include any sensors that can detect environmental variables – light sensors, noise sensors, vibration sensors, thermostats, etc. Additionally, the environmental sensors module 1612 can store historical readings or reports from those sensors.

[00160] The environmental factors module 1614 can include reports generated based on data in the environmental sensors module 1612. For example, for a user with a light sensor with data in the environment sensors module 1612, the

environmental factors module 1614 can hold a report indicating the frequency and duration of instances of increased lighting when the user is asleep.

[00161] In the examples discussed here, each cloud service 410 is shown with some of the same components. In various configurations, these same components can be partially or wholly shared between services, or they can be separate. In some configurations, each service can have separate copies of some or all of the components that are the same or different in some ways. Additionally, these components are only supplied as illustrative examples. In other examples each cloud service can have different number, types, and styles of components that are technically possible.

[00162] FIG. 17 is a block diagram of an example of using a data processing system that can be associated with a bed (such as a bed of the bed systems described herein) to automate peripherals around the bed. Shown here is a behavior analysis module 1700 that runs on the pump motherboard 402. For example, the behavior analysis module 1700 can be one or more software components stored on the computer memory 512 and executed by the processor 502. In general, the behavior analysis module 1700 can collect data from a wide variety of sources (e.g., sensors, non-sensor local sources, cloud data services) and use a behavioral algorithm 1702 to generate one or more actions to be taken (e.g., commands to send to peripheral controllers, data to send to cloud services). This can be useful, for example, in tracking user behavior and automating devices in communication with the user's bed.

[00163] The behavior analysis module 1700 can collect data from any technologically appropriate source, for example, to gather data about features of a bed, the bed's environment, and/or the bed's users. Some such sources include any of the sensors of the sensor array 406. For example, this data can provide the behavior analysis module 1700 with information about the current state of the environment around the bed. For example, the behavior analysis module 1700 can access readings from the pressure sensor 902 to determine the pressure of an air chamber in the bed. From this reading, and potentially other data, user presence in the bed can be determined. In another example, the behavior analysis module can access a light sensor 908 to detect the amount of light in the bed's environment.

[00164] Similarly, the behavior analysis module 1700 can access data from cloud services. For example, the behavior analysis module 1700 can access the bed

cloud service 410a to access historical sensor data 1212 and/or advanced sleep data 1214. Other cloud services 410, including those not previously described can be accessed by the behavior analysis module 1700. For example, the behavior analysis module 1700 can access a weather reporting service, a 3rd party data provider (e.g., traffic and news data, emergency broadcast data, user travel data), and/or a clock and calendar service.

[00165] Similarly, the behavior analysis module 1700 can access data from non-sensor sources 1704. For example, the behavior analysis module 1700 can access a local clock and calendar service (e.g., a component of the motherboard 402 or of the processor 502).

[00166] The behavior analysis module 1700 can aggregate and prepare this data for use by one or more behavioral algorithms 1702. The behavioral algorithms 1702 can be used to learn a user's behavior and/or to perform some action based on the state of the accessed data and/or the predicted user behavior. For example, the behavior algorithm 1702 can use available data (e.g., pressure sensor, non-sensor data, clock and calendar data) to create a model of when a user goes to bed every night. Later, the same or a different behavioral algorithm 1702 can be used to determine if an increase in air chamber pressure is likely to indicate a user going to bed and, if so, send some data to a third-party cloud service 410 and/or engage a peripheral controller 1002.

[00167] In the example shown, the behavioral analysis module 1700 and the behavioral algorithm 1702 are shown as components of the motherboard 402. However, other configurations are possible. For example, the same or a similar behavioral analysis module and/or behavior algorithm can be run in one or more cloud services, and the resulting output can be sent to the motherboard 402, a controller in the controller array 408, or to any other technologically appropriate recipient.

[00168] FIG. 18 shows an example of a computing device 1800 and an example of a mobile computing device that can be used to implement the techniques described here. The computing device 1800 is intended to represent various forms of digital computers, such as laptops, desktops, workstations, personal digital assistants, servers, blade servers, mainframes, and other appropriate computers. The mobile computing device is intended to represent various forms of mobile devices, such as

personal digital assistants, cellular telephones, smart-phones, and other similar computing devices. The components shown here, their connections and relationships, and their functions, are meant to be exemplary only, and are not meant to limit implementations of the inventions described and/or claimed in this document.

[00169] The computing device 1800 includes a processor 1802, a memory 1804, a storage device 1806, a high-speed interface 1808 connecting to the memory 1804 and multiple high-speed expansion ports 1810, and a low-speed interface 1812 connecting to a low-speed expansion port 1814 and the storage device 1806. Each of the processor 1802, the memory 1804, the storage device 1806, the high-speed interface 1808, the high-speed expansion ports 1810, and the low-speed interface 1812, are interconnected using various busses, and can be mounted on a common motherboard or in other manners as appropriate. The processor 1802 can process instructions for execution within the computing device 1800, including instructions stored in the memory 1804 or on the storage device 1806 to display graphical information for a GUI on an external input/output device, such as a display 1816 coupled to the high-speed interface 1808. In other implementations, multiple processors and/or multiple buses can be used, as appropriate, along with multiple memories and types of memory. Also, multiple computing devices can be connected, with each device providing portions of the necessary operations (e.g., as a server bank, a group of blade servers, or a multi-processor system).

[00170] The memory 1804 stores information within the computing device 1800. In some implementations, the memory 1804 is a volatile memory unit or units. In some implementations, the memory 1804 is a non-volatile memory unit or units. The memory 1804 can also be another form of computer-readable medium, such as a magnetic or optical disk.

[00171] The storage device 1806 is capable of providing mass storage for the computing device 1800. In some implementations, the storage device 1806 can be or contain a computer-readable medium, such as a floppy disk device, a hard disk device, an optical disk device, or a tape device, a flash memory or other similar solid state memory device, or an array of devices, including devices in a storage area network or other configurations. A computer program product can be tangibly embodied in an information carrier. The computer program product can also contain instructions that, when executed, perform one or more methods, such as those

described above. The computer program product can also be tangibly embodied in a computer- or machine-readable medium, such as the memory 1804, the storage device 1806, or memory on the processor 1802.

[00172] The high-speed interface 1808 manages bandwidth-intensive operations for the computing device 1800, while the low-speed interface 1812 manages lower bandwidth-intensive operations. Such allocation of functions is exemplary only. In some implementations, the high-speed interface 1808 is coupled to the memory 1804, the display 1816 (e.g., through a graphics processor or accelerator), and to the high-speed expansion ports 1810, which can accept various expansion cards (not shown). In the implementation, the low-speed interface 1812 is coupled to the storage device 1806 and the low-speed expansion port 1814. The low-speed expansion port 1814, which can include various communication ports (e.g., USB, Bluetooth, Ethernet, wireless Ethernet) can be coupled to one or more input/output devices, such as a keyboard, a pointing device, a scanner, or a networking device such as a switch or router, e.g., through a network adapter.

[00173] The computing device 1800 can be implemented in a number of different forms, as shown in the figure. For example, it can be implemented as a standard server 1820, or multiple times in a group of such servers. In addition, it can be implemented in a personal computer such as a laptop computer 1822. It can also be implemented as part of a rack server system 1824. Alternatively, components from the computing device 1800 can be combined with other components in a mobile device (not shown), such as a mobile computing device 1850. Each of such devices can contain one or more of the computing device 1800 and the mobile computing device 1850, and an entire system can be made up of multiple computing devices communicating with each other.

[00174] The mobile computing device 1850 includes a processor 1852, a memory 1864, an input/output device such as a display 1854, a communication interface 1866, and a transceiver 1868, among other components. The mobile computing device 1850 can also be provided with a storage device, such as a micro-drive or other device, to provide additional storage. Each of the processor 1852, the memory 1864, the display 1854, the communication interface 1866, and the transceiver 1868, are interconnected using various buses, and several of the

components can be mounted on a common motherboard or in other manners as appropriate.

[00175] The processor 1852 can execute instructions within the mobile computing device 1850, including instructions stored in the memory 1864. The processor 1852 can be implemented as a chipset of chips that include separate and multiple analog and digital processors. The processor 1852 can provide, for example, for coordination of the other components of the mobile computing device 1850, such as control of user interfaces, applications run by the mobile computing device 1850, and wireless communication by the mobile computing device 1850.

[00176] The processor 1852 can communicate with a user through a control interface 1858 and a display interface 1856 coupled to the display 1854. The display 1854 can be, for example, a TFT (Thin-Film-Transistor Liquid Crystal Display) display or an OLED (Organic Light Emitting Diode) display, or other appropriate display technology. The display interface 1856 can comprise appropriate circuitry for driving the display 1854 to present graphical and other information to a user. The control interface 1858 can receive commands from a user and convert them for submission to the processor 1852. In addition, an external interface 1862 can provide communication with the processor 1852, so as to enable near area communication of the mobile computing device 1850 with other devices. The external interface 1862 can provide, for example, for wired communication in some implementations, or for wireless communication in other implementations, and multiple interfaces can also be used.

[00177] The memory 1864 stores information within the mobile computing device 1850. The memory 1864 can be implemented as one or more of a computer-readable medium or media, a volatile memory unit or units, or a non-volatile memory unit or units. An expansion memory 1874 can also be provided and connected to the mobile computing device 1850 through an expansion interface 1872, which can include, for example, a SIMM (Single In Line Memory Module) card interface. The expansion memory 1874 can provide extra storage space for the mobile computing device 1850, or can also store applications or other information for the mobile computing device 1850. Specifically, the expansion memory 1874 can include instructions to carry out or supplement the processes described above, and can include secure information also. Thus, for example, the expansion memory 1874 can be

provide as a security module for the mobile computing device 1850, and can be programmed with instructions that permit secure use of the mobile computing device 1850. In addition, secure applications can be provided via the SIMM cards, along with additional information, such as placing identifying information on the SIMM card in a non-hackable manner.

[00178] The memory can include, for example, flash memory and/or NVRAM memory (non-volatile random access memory), as discussed below. In some implementations, a computer program product is tangibly embodied in an information carrier. The computer program product contains instructions that, when executed, perform one or more methods, such as those described above. The computer program product can be a computer- or machine-readable medium, such as the memory 1864, the expansion memory 1874, or memory on the processor 1852. In some implementations, the computer program product can be received in a propagated signal, for example, over the transceiver 1868 or the external interface 1862.

[00179] The mobile computing device 1850 can communicate wirelessly through the communication interface 1866, which can include digital signal processing circuitry where necessary. The communication interface 1866 can provide for communications under various modes or protocols, such as GSM voice calls (Global System for Mobile communications), SMS (Short Message Service), EMS (Enhanced Messaging Service), or MMS messaging (Multimedia Messaging Service), CDMA (code division multiple access), TDMA (time division multiple access), PDC (Personal Digital Cellular), WCDMA (Wideband Code Division Multiple Access), CDMA2000, or GPRS (General Packet Radio Service), among others. Such communication can occur, for example, through the transceiver 1868 using a radio-frequency. In addition, short-range communication can occur, such as using a Bluetooth, WiFi, or other such transceiver (not shown). In addition, a GPS (Global Positioning System) receiver module 1870 can provide additional navigation- and location-related wireless data to the mobile computing device 1850, which can be used as appropriate by applications running on the mobile computing device 1850.

[00180] The mobile computing device 1850 can also communicate audibly using an audio codec 1860, which can receive spoken information from a user and convert it to usable digital information. The audio codec 1860 can likewise generate audible sound for a user, such as through a speaker, e.g., in a handset of the mobile

computing device 1850. Such sound can include sound from voice telephone calls, can include recorded sound (e.g., voice messages, music files, etc.) and can also include sound generated by applications operating on the mobile computing device 1850.

[00181] The mobile computing device 1850 can be implemented in a number of different forms, as shown in the figure. For example, it can be implemented as a cellular telephone 1880. It can also be implemented as part of a smart-phone 1882, personal digital assistant, or other similar mobile device.

[00182] Various implementations of the systems and techniques described here can be realized in digital electronic circuitry, integrated circuitry, specially designed ASICs (application specific integrated circuits), computer hardware, firmware, software, and/or combinations thereof. These various implementations can include implementation in one or more computer programs that are executable and/or interpretable on a programmable system including at least one programmable processor, which can be special or general purpose, coupled to receive data and instructions from, and to transmit data and instructions to, a storage system, at least one input device, and at least one output device.

[00183] These computer programs (also known as programs, software, software applications or code) include machine instructions for a programmable processor, and can be implemented in a high-level procedural and/or object-oriented programming language, and/or in assembly/machine language. As used herein, the terms machine-readable medium and computer-readable medium refer to any computer program product, apparatus and/or device (e.g., magnetic discs, optical disks, memory, Programmable Logic Devices (PLDs)) used to provide machine instructions and/or data to a programmable processor, including a machine-readable medium that receives machine instructions as a machine-readable signal. The term machine-readable signal refers to any signal used to provide machine instructions and/or data to a programmable processor.

[00184] To provide for interaction with a user, the systems and techniques described here can be implemented on a computer having a display device (e.g., a CRT (cathode ray tube) or LCD (liquid crystal display) monitor) for displaying information to the user and a keyboard and a pointing device (e.g., a mouse or a trackball) by which the user can provide input to the computer. Other kinds of

devices can be used to provide for interaction with a user as well; for example, feedback provided to the user can be any form of sensory feedback (e.g., visual feedback, auditory feedback, or tactile feedback); and input from the user can be received in any form, including acoustic, speech, or tactile input.

[00185] The systems and techniques described here can be implemented in a computing system that includes a backend component (e.g., as a data server), or that includes a middleware component (e.g., an application server), or that includes a frontend component (e.g., a client computer having a graphical user interface or a Web browser through which a user can interact with an implementation of the systems and techniques described here), or any combination of such backend, middleware, or frontend components. The components of the system can be interconnected by any form or medium of digital data communication (e.g., a communication network). Examples of communication networks include a local area network (LAN), a wide area network (WAN), and the Internet.

[00186] The computing system can include clients and servers. A client and server are generally remote from each other and typically interact through a communication network. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other.

[00187] FIGs. 19A and 19B are swimlane diagrams of example processes 1900 and 1950 for training and using machine-learning classifiers to determine user sleep state, which can include sleep stage. For clarity, the processes 1900 and 1950 are being described with reference to a particular set of components. However, other system or systems can be used to perform the same or a similar process.

[00188] In the process 1900, a bed system uses the readings of a pressure sensor 1902 to learn what effect a user has on the pressure of the bed for the purposes of categorizing the user's sleep (or lack thereof). The bed system is able to use this pressure of the airbed as a signal for a decision engine that classifies the user sleep into one of a plurality of possible sleep states. These sleep states may include two states (e.g., asleep or awake), five states (e.g., awake, N1, N2, N3, and REM), three states (e.g., awake, light sleep, deep sleep), or another set of states.

[00189] In operation, the bed can determine the sleep state of the user and operate according to the sleep state. For example, the bed can determine an optimal

next sleep stage for a sleeper to ensure maximum physiological and mental recovery during sleep.

[00190] In one example, REM is the optimal sleep stage for mental recovery and deep sleep is the optimal sleep stage for physiological recovery. For example, if the sleeper is currently in NREM, and has not been in REM within the past hour, to maximize mental recovery, the best next state will be determined as REM.

[00191] In another example, if the sleeper is in deep sleep (e.g., N3) but has not been long enough in N3, the next best stage will still be N3 to keep the user in N3 longer to maximize physiological recovery.

[00192] In another example, if the sleeper is in lighter stages of sleep but has been restless and therefore not able to go to deep sleep or REM, the next optimal state is getting out of light sleep and going to REM or Deep sleep.

[00193] In order to encourage a user to stay in a current sleep stage or to enter a different sleep stage, the bed or environment can change in a way shown to encourage maintenance or change of sleep stage. For example, bed pressure and temperature can change.

[00194] A pressure sensor 1902 senses pressure 1912. For example, the pressure sensor may create a live stream of pressure readings that reflect the pressure inside of an air bladder within a bed system. This live stream of pressure readings may be provided to a bed controller 1904 in the form of analog or digital information on a substantially constant basis, thus reflecting pressure as within the air bladder due to a user (or other object) on the bed system or when the bed is empty.

[00195] The bed controller 1904 receives the pressure readings 1914. For example, the bed controller 1904 can place pressure readings in a computer memory structure such as a rolling buffer that makes the most recent *N* readings available to the bed controller. The bed controller 1904 may aggregate these pressure readings, subsample the readings, or store them all individually.

[00196] The bed controller 1904 transmits the pressure readings 1916 and a cloud reporting service 1906 receives the pressure readings 1918. For example, the bed controller 1904 can transmit all pressure readings or determine that some pressure readings – and not others – should be transmitted to the cloud reporting service 1906 that is configured to receive pressure readings and in some cases other types of data. The sensor readings sent to the cloud reporting service 1906 may be unchanged by the

bed controller 1904, aggregated (e.g., averages, maximums and minimums, etc.), or otherwise changed by the bed controller 1904.

[00197] A classifier factory 1908 generates classifiers from the pressure readings 1920. The classifier factory 1908 can train classifiers by first obtaining a large set of pre-classified pressure variation patterns. For example, one bed or many beds may report pressure data to a cloud reporting service 1906. This pressure data may be tagged, recorded, and stored for analysis in the creation of pressure classifiers to be used by the bed controller 1904 and/or other bed controllers.

[00198] The classifier factory 1908 can generate features from the pressure readings. For example, the stream of pressure signals may be broken into buffers of, for example, 1 second, 2.125 seconds, 30 seconds, or 3 seconds, to generate features in time or frequency domains. These features may be direct measure of pressure within those buffers.

[00199] In some cases, the classifier factory 1908 can generate features directly. In some cases (not shown), the bed controller 1904 can generate features and send features (as opposed to pressure, as shown) to the cloud reporting service 1906.

[00200]

[00201] For example, such features may include a maximum, minimum, or random pressure value. These features may be derived from the pressure readings within those buffers. For example, such features may include an average pressure, a standard deviation, a slope value that indicates an increase or decrease over time within that buffer, user motion, respiration measurement, cardiac measurement, and cardiorespiratory coupling measurement from the pressure variations. For example, rate/amplitude/duration of inhalation, rate/amplitude/duration of exhalation, rate of inhalation-exhalation cycle, the amplitude, width and location of fundamental frequency of breathing, the heart rate, rate of atrial depolarization, rate/amplitude/duration of atrial repolarization, rate/amplitude/duration of ventricular depolarization, rate of ventricular repolarization, etc. may be used. The values of the feature vectors may be in binary or numerical form. For each buffer, the values may be stored in a predetermined order creating a vector that is composed of a series of fields, where every vector has the same series of fields and data in those fields. Some other features may be computed from the transform domain representations of the pressure signal such as from the Fourier or Wavelet Transform coefficients.

[00202] As another example, the classifier factory can identify instances within the pressure readings where the pressure readings match a pattern or rules for a pattern. In one example, a constant or fluctuating pattern may be identified in the motion signal or at least one physiological signal (cardiac beats, breathing or cardiorespiratory coupling) derived from the pressure readings. Such patterns may be identified, and corresponding synthetic information about the pattern (e.g., timestamp, duration, rate of change, frequency of change, max change, slope of change, etc) may be synthesized from the pressure signal and/or other outside information (e.g., a real-time clock).

[00203] The classifier factory 1908 can combine or reduce the features. For example, the extracted features can be combined using principal component analysis. For a principal component analysis of the features, the classifier factory 1908 can determine a subset of all features that are discriminant of the sleep state of the user. That is, the classifier factory 1908 can sort features into those features that are useful for determining state and those features that are less useful, and the more useful features may be kept. This process may be done on a trial-and-error basis, in which random combinations of features are tested. This process may be done with the use of one or more systematic processes. For example, a linear discriminant analysis or generalized discriminant analysis may be used.

[00204] In some cases, a proper subset of features may be selected out of the set of all available features. This selection may be done once per classifier if multiple classifiers are being created. Alternatively, this selection may be done once for a plurality or all classifiers if multiple classifiers are being created.

[00205] For example, a random (or pseudorandom) number may be generated and that number of features may be removed. In some cases, a plurality of features may be aggregated into a single aggregate feature. For example, for a case in which a plurality of fluctuating physiological patterns are identified in the pressure readings, the patterns and/or synthetic data related to the patterns may be aggregated. For example, the consecutive heart rate differences may be aggregated into a mean, a standard deviation, a minimum, and/or a maximum heart rate.

[00206] The classifier factory 1908 can process the features. For example, the remaining features may then be processed to rationalize their values so that each feature is handled with a weight that corresponds to how discriminant the feature is.

If a feature is found to be highly discriminant so that is highly useful in classifying state, that feature may be given a larger weight than other features. If a second feature is found to be less discriminant than other features, that second feature can be given a lower weight.

[00207] Once mapped into kernel space, the features can be standardized to center the data points at a predetermined mean and to scale the features to have unit standard deviation. This can allow the features to all have, for example, a mean value of 0 and a standard deviation of 1. The extracted features are then converted to a vector format using the same vector format as described above.

[00208] In some cases, the remaining features can be processed by applying a kernel function to map the input data into a kernel space. A kernel space allows a high-dimensional space (e.g., the vector space populated with vectors of feature data) to be clustered such that different clusters can represent different states. The kernel function may be of any appropriate format, including linear, quadratic, polynomial, radial basis, multilayer perceptron, or custom.

[00209] In some cases, the classifier factory can use machine-learning techniques that do not create features. For example deep learning networks such as convolutional networks, deep feed forward or deep recurrent networks can be used.

[00210] The classifier factory 1908 can train the classifiers. For example, a pattern recognizer algorithm can use the vectors of extracted features and their corresponding sleep state labels as a dataset to train the classifiers with which new pressure readings can be classified. In some cases, this can include storing the classifiers with the training data for later use.

[00211] The classifier factory 1908 can transmit the classifiers 1922 and the bed controller 1904 can receive the classifiers 1924. For example, the classifier or classifiers created by the classifier factory 1908 can be transmitted to the bed controller 1904 and/or other bed controllers. In some cases, the classifiers can be transmitted on non-transitory computer readable mediums like a compact disk (CD), a Universal Serial Bus (USB) drive, or other device. The classifiers may be loaded onto the bed controller 1904 and/or other bed controllers as part of a software installation, as part of a software update, or as part of another process. In some cases, the classifier factory 1908 can transmit a message to the bed controller 1904 and/or other bed controllers, and the message can contain data defining one or more classifiers that

use a stream of pressure readings to classify the bed into one of a plurality of sleep states. In some configurations, the classifier factory 1908 can transmit the classifiers at once, either in one message or a series of messages near each other in time. In some configurations, the classifier factory 1908 can send the classifiers separated in time. For example, the classifier factory 1908 may generate and transmit classifiers. Later, with more pressure sensor data available, the classifier factory 1908 may generate an updated classifier or a new classifier unlike one already created.

[00212] The classifier may be defined in one or more data structures. For example, the classifier factory 1908 can record a classifier in an executable or interpretable files such as a software library, executable file, or object file. The classifier may be stored, used, or transmitted as a structured data object such as an extensible markup language (XML) document or a JavaScript object notation (JSON) object. In some examples, a classifier may be created in a binary or script format that the bed controller 1904 can run (e.g., execute or interpret). In some examples, a classifier may be created in a format that is not directly run, but in a format with data that allows the bed controller 1904 to construct the classifier according to the data.

[00213] The bed controller 1904 can also use the stream of pressure readings to determine sleep state 1926. For example, the bed controller 1904 can run one or more classifiers using data from the stream of pressure readings. The classifier can categorize this data into one of a plurality of states (e.g., awake, N1, N2, N3, REM). For example, the classifier may convert the data stream into a vector format described above. The classifier may then examine the vector to mathematically determine if the vector is more like training data labeled as one state or more like training data labeled as another state. Once this similarity is calculated, the categorizer can return a response indicating that state.

[00214] The bed controller 1904 can use more than one classifier. That is, the bed controller 1904 may have access to a plurality of classifiers that each function differently and/or use different training data to generate classifications. In such cases, classifier decisions can be treated as a vote and vote aggregation can be used to determine sleep state. If only one classifier is used, the vote of that classifier is the only vote and the vote is used as the sleep state. If there are multiple classifiers, the different classifiers can produce conflicting votes, and the bed controller can select a vote-winning sleep state.

[00215] Various vote-counting schemes are possible. In some cases, the bed controller 1094 can count the votes for each sleep state and the sleep state with the most votes is the determined sleep state. In some cases, the bed controller 1904 can use other vote-counting schemes. For example, votes from different classifiers may be weighed based on the classifiers historical accuracy. In such a scheme, classifiers that have been historically shown to be more accurate can be given greater weight while classifiers with lesser historical accuracy can be given less weight. This accuracy may be tracked on a population level or on a particular user level.

[00216] In some instances, votes may be cast by systems other than a machine-learning system, and those votes may be incorporated into the vote totals to impact the outcomes of the voting decision. For example, non-machine-learning pressure categorizing algorithms may cast votes based on, for example, comparisons with threshold values.

[00217] In some instances, the system may have different operational modes, and may tally votes differently depending on the mode. For example, when a bed is in the process of adjusting or when the adjustable foundation is moving or a portion of the bed is elevated, different vote strategies may be used. In some modes, some classifiers may be given greater weight or lesser weight or no weight as compared to some other modes. This may be useful, for example, when a classifier is shown to be accurate in one mode (e.g. with the bed flat) versus another mode (e.g., with the head of the bed elevated by the foundation).

[00218] In some cases, the bed controller 1904 can ensure that there is a user in bed before determining sleep state. For example, the bed controller can initially determine if the user is in the bed or if the bed is empty. If the user is determined to be in the bed, the bed controller 1904 can determine if the user is asleep in the bed and if so in what sleep state.

[00219] The bed controller 1904 selects a device operation 1928. For example, responsive to a determination that the user system is in a particular sleep state, or in response to a determination that a sleep state change has occurred, the bed controller 1904 can select a device operation to be processed. A ruleset stored in computer-readable storage, e.g. locally or on a remote machine, can identify actions that a user or another system have requested based on sleep state. For example, a user can document through a graphical user interface that they wish their front-door to lock

when they change from an awake state to an asleep state, and they wish their bed to encourage them to optimize the sleep states, for example, maximize deep, maximize REM sleep, or maximize the most restful sleep possible. That is to say, the door should lock when the user falls asleep if they have forgotten to do so manually, and the bed should encourage the most restful night of sleep possible.

[00220] In some cases, the rule set may be organized to encourage a user to stay in or transition to another sleep stage. For example, the user's environment may change based on a rule to increase or decrease bed firmness, bed temperature, Heating ventilation and air conditioning (HVAC) temperature, etc., in a way that encourages a user to transition to another sleep stage or to remain in the current sleep stage.

[00221] Based on the ruleset and the sleep stage determination, the bed controller 1904 can send messages to appropriate device controllers 1910 in order to engage the peripherals or bed-system elements called for. For example, based on the sleep determination, the bed controller 1904 can send a message to a pump to adjust the firmness of the bed-system, a message to a foot-warming controller to engage foot heaters, and a message to a light controller to adjust lighting.

[00222] A device controller 1910 can control a peripheral device 1930. For example, a light controller may initiate a script for the lighting in the room around the bed to begin dimming and shifting the color of the light toward the red end of the visible light spectrum.

[00223] The particular sleep state, including sleep stage, of a user of the bed can be used to drive one or more home automation commands that can be used to improve the sleep experience of a user. A few are described here.

[00224] Bed firmness can be adjusted depending on sleep stage. For example, when a user is in REM sleep, their body becomes atonic, that is, it lacks muscle tone and puts up very little resistance to outside pressure. However, in other sleep stages, the body can have more muscle tone and put up comparatively more resistance to outside pressure. As such, in order to provide a consistent level of comfort, body alignment, etc., the bed controller can select a firmer pressure when the user is in REM, and provide a softer pressure when the user is in other sleep stages. This has been shown in some situations to increase user comfort, maintain a steady body alignment, and provide for a more restful sleep.

[00225] Alarms to awaken the user can be adjusted based on sleep stage. For example, in order to allow for a more complete REM cycle, an alarm to awaken the user may be adjusted so that it goes off outside of a REM cycle. The user may enable a setting to allow the alarm to be moved up to, for example, fifteen minutes if the user is in REM stage when the alarm is scheduled to go off.

[00226] Temperature can be adjusted based on sleep stage. For example, a user may feel more comfortable in a warm bed, but like many users may actually sleep more restfully (e.g., with less motion) when sleeping in a cooler bed. The user may save settings that increase the temperature of the bed when in a sleep state of “awake,” but then cool the bed when the user is in a particular sleep state (e.g., asleep) or sleep stage (e.g. N1.)

[00227] In general, the process 1900 can be organized into a training time and an operating time. The training time can include actions that are generally used to create sleep stage classifiers, while the operating time can include actions that are generally used to determine a sleep state with the classifiers. Depending on the configuration of the bed system, the actions of one or both of the times may be engaged or suspended. For example, when a user newly purchases a bed, the bed may have access to no pressure readings caused by the user on the bed. When the user begins using the bed for the first few nights, the bed system can collect those pressure readings and supply them to the cloud reporting service 1906 once a critical mass of readings have been collected (e.g. a certain number of readings, a certain number of nights, a certain number of expected entry and exit events based on different tests or heuristics).

[00228] The bed system may operate in the training time to update or expand the classifiers. The bed controller 1904 may continue actions of the training time after receipt of the classifiers. For example, the bed controller 1904 may transmit pressure readings to the cloud reporting service 1906 on a regular basis, when computational resources are free, at user direction, etc. The classifier factory 1908 may generate and transmit new or updated classifiers, or may transmit messages indicating that one or more classifiers on the bed controller 1904 should be retired.

[00229] The bed controller 1904 can receive rules and setting that define how the home-automation connected to the bed-system should operate. With the

classifiers, the bed system can perform the actions of the operating time in order to cause the home-automation to perform according to the rules and settings.

[00230] The bed system can use the same pressure readings from the pressure sensor 1902 to operate in the training time and the operating time concurrently. For example, the bed system can use the stream of pressure readings to determine a sleep state and control the environment based on sleep categorizers that are currently in use. In addition, the bed system can also use the same pressure readings from the stream of pressure readings in the training time actions to improve the categorizers. In this way, a single stream of pressure readings may be used to both improve the function of the bed system and to drive automation events.

[00231] In some cases, a generic set of classifiers may be used instead of, or in conjunction with, personalized classifier. For example, when a bed is newly purchased or reset to factory settings, the bed system may operate with generic or default sleep state classifiers that are created based on population-level, not individual, pressure readings. That is, generic classifiers may be created for use in a bed system before the bed system has had an opportunity to learn about the particular pressure readings associated with a particular user. These generic classifiers may be generated using machine learning techniques, such as those described in this document, on population-level training data. These generic classifiers may additionally or alternatively be generated using non-machine learning techniques. For example, a classifier may include a threshold value (e.g., pressure, pressure change over time, heart rate change over time, changes in physiological components of the pressure readings over time), and a pressure measure over that threshold may be used to determine one sleep state while pressure readings and/or physiological values under that threshold may be used to determine another sleep state.

[00232] While a particular number, order, and arrangement of elements are described here, other alternatives are possible. For example, while the generation of classifiers 1920 is described as being performed on a classifier factory 1908, classifiers can be instead or additionally generated by the bed controller 1906, possibly without reporting pressure data to a cloud service.

[00233] In some implementations, the bed system may accommodate two users. In such a case the process 1900 can be adapted in one or more way to accommodate two users. For example, for each user, the bed system may use two sets

of classifiers (with or without some classifiers being simultaneously in both sets.) For example, one set may be used when the other side of the bed is occupied, and one set may be used when the other side of the bed is occupied. This may be useful, for example, when the presence or absence of the second user has an impact on pressure readings on the first user's side of the bed.

[00234] In some cases, the user may wish to control their home-automation environment contingent upon the sleep-state of both users. For example, a rule may specify that the front-door lock should be engaged only when both users are asleep in the bed, not just when one user is asleep in the bed. In another example, lighting may be controlled based on both user's sleep state. When one user is asleep in the bed, and the other is determined to be awake, this can indicate one user may be reading or otherwise needing to see. As such, the bed system can keep on dim lighting. When both users are present in bed and both transition to sleep, the lighting can be controlled to turn off.

[00235] This example is shown with a single bed controller 1904 providing pressures 1916 and then later receiving classifiers 1924. However, it will be understood that this system is applicable with many more beds and bed controllers. For example, pressures may be received from many bed controllers (e.g., hundreds of thousands), and training data can be synthesized from these many beds, providing data about bed use by many users. The classifiers can then be distributed to some, none, or all of those beds that provided training data. For example, some beds may receive a software updated with new classifiers. Or as another example, the new classifiers may only be included on newly manufactured beds. Or as another example, each bed may receive classifiers that are particularly tailored to the users of that particular bed.

[00236] FIG. 19B is a swimlane diagram of an example process 1950 for training and using machine-learning classifiers to determine and classify sleep stage in a bed. Unlike in the process 1900, the process 1950 includes generating classifiers 1952 with the use of deep learning styles of machine learning. In the example shown, a deep neural network (DNN) that is a computer model (as opposed to an organic brain) is being described. However, as will be understood, other types of artificial neural networks and/or other types of deep learning (e.g., convolutional neural networks, recurrent neural network, long short-term memory-LSTM, etc.) may be

used in the process 1950. Further, it will be understood that other types of machine learning can be used in the processes 1900 and 1950 in order to generate classifiers (1920 and 1952.)

[00237] In general, in the classifier generation 1952, the classifier factory 1908 receives labeled training data from the cloud reporting service. However, unlike in the process 1900, explicit features are not created as a stand-alone process. Instead, the training of the classifiers works directly on the labeled training data, not features created from the labeled training data.

[00238] The classifier factory 1908 generates classifiers from the pressure/acoustic readings 1952. For example, the classifier factory 1908 may perform artificial neural network type machine learning to generate the classifiers. The classifier factory 1908 can train classifiers by first obtaining a large set of pre-classified reading variation patterns. For example, one bed or many beds may report reading data to a cloud reporting service 1906. This reading data may be labeled, recorded, and stored for analysis in the creation of pressure classifiers to be used by the bed controller 1904 and/or other bed controllers.

[00239] The tagged data is provided to one or more DNN trainers. The DNN trainers generate an initial DNN by arranging groups of artificial neurons into layers, and then connecting the output of one layer with the input of another layer. Generally speaking, these artificial neurons are computer-operable functions that take several inputs, perform their function, and produce output. Often these functions are defined based on a two-part mathematical function – first some linear combination is performed, then a non-linear function (also called activation function) is performed. However, as will be understood, any technologically appropriate function may be used.

[00240] Neurons in one layer are all grouped, and the output of each neuron in the layer is provided as an input to neurons of the next layer. The number of connections between each layer is a function of the number of inputs of each neuron in the layer. For example, for a network in which each layer has ten neurons and each neuron has three inputs, the network would have thirty (i.e. ten time three) connections between one layer and the next. The number of layers, number of neurons per layer, and number of inputs per neuron are each parameters that the classifier factory 1908 can adjust in the process of initializing an DNN. For example,

the network may have tens of layers, each layer having hundreds of neurons, each neuron having tens of inputs. More or less complexity (numbers of layers, neurons, and/or inputs) is possible.

[00241] Each connection, from one neuron's output to the next neuron's input, is given a weight value. This weight value is initialized, for example to a random (or pseudorandom) number, or by selecting from a list of possible weights. When the output of one neuron is passed to the input of the next neuron, the value is adjusted by the weight. For example, the weight may be a number ranging from 0 to 1, and the value passed may be multiplied by the weight.

[00242] With this initial DNN generated, it is capable of receiving the training data and operating on the training data. That is, the training data, stored on disk as an ordered sequence of binary data, can be provided as input into the head of the DNN (that is, the original input neuron for the first layer of the DNN.) As will be understood, providing input the first layer of the DNN causes the DNN to execute neurons of the layers of the DNN and produce an output in the form of a second ordered sequence of binary data. Here, the second ordered sequence of binary data may then be interpreted as a classification with a confidence score – that is, the output “tells” a reader what state the DNN has classified the data into (e.g., no-sleep, light-sleep, heavy-sleep, N1, N2, REM) along with a confidence value from 0 to 1.

[00243] With the initial DNN generated, the classifier factory 1908 can refine the DNN to improve the classification results created by the DNN. In order to do so, the classifier factory 1908 can calculate a loss function and iteratively modify the DNN until the loss function for the DNN passes a test such as falling below a threshold or failing to improve over iterative refinements.

[00244] A loss function can be selected that defines how well the DNN has classified a sample of tagged training data. In the example with a confidence of values 0 to 1, a loss function may be used that assigns a loss-value of 1 for an incorrect classification, and a loss value of 1-confidence for a correct classification. In this way, an incorrect classification provides a maximum value loss, while a correct classification provides a small loss when confidence is high.

[00245] The classifier factory 1908 begins refining the DNN in order to reduce the loss value of the DNN. For example, the classifier factory 1908 can iteratively

perform the steps of i) adjusting the DNN, ii) providing training data to the DNN, and iii) calculate the loss value for the DNN.

[00246] In order to adjust the DNN, the classifier factory can select one or more optimization algorithms. In general, many of these algorithms operate by adjusting the weights of connections between neuron outputs and neuron inputs. In doing so, they adjust the actual, weighted inputs that are used by neurons of the DNN, which produces a different results for the DNN.

[00247] One of these algorithms is called a gradient descent algorithm. Gradient descent is a first-order iterative optimization algorithm for finding a minimum of the loss function. In each iteration of the gradient descent, the current weights of the connections between neurons are considered and modified in a way that reduces the loss value for the DNN by at least a small amount. To make these modifications, the classifier factory 1908 can determine the gradient of the loss function for the DNN with respect to all of the weights of the DNN. Using the gradient, new weights that would reduce the loss function by a learning rate are calculated. The gradient descent algorithm may also incorporate elements to avoid being trapped in local minima. Example elements include stochastic, batch, and mini-batch gradient descents.

[00248] Once the DNN has been adjusted, the classifier factory 1908 can generate a classifier or classifiers from the DNN. For example, the classifier factory 1908 can identify neurons with all input weights of zero and remove them, as they do not contribute to the classifications performed with the DNN.

[00249] FIG. 20 is a flowchart of an example process 2000 for training classifiers on pressure signals. The process 2000 may be used, for example, as part of the process to generate classifiers 1920.

[00250] A feature set is determined 2002. For example, raw pressure data can be separated into rolling windows of pressure data. In one example, each window represents 1 second, 1 minute, or 10 seconds of pressure data with 100 pressure readings each. A vector can be created, with the first 100 fields of the vector being used to store each of the 100 pressure readings in order. Additional fields in the vector are used to store data calculated from the 100 pressure readings. For example, a field may be used to store the difference between the first pressure reading and the last pressure reading. This difference value may be used as an approximate proxy of

the movement, with a negative difference indicating an increased movement and a positive difference indicating a decreased movement. Another field may be used to store the greatest difference between any two adjacent pressure readings, which may be indicative of the amount of noise in the pressure readings. Furthermore, some features may be created without having a clear or logical description outside of their mathematical determination. For example, a count of pressure readings with odd or even values may be stored in one field of the vector. These fields may be defined by human design, or may be generated programmatically.

[00251] Training data is mapped to kernel space 2004. For example, the vectors may be mapped into a high-dimensional space. This high dimensional space may have the same number of dimensions as the vectors have fields, or a subset of N fields of the vector may be used and the vector can be mapped to an N dimensional space. A kernel function may be found that is able to partition the space into partitions that each have one cluster of vectors in them. For example, in a 2D space, the vectors may map to one cluster around the coordinate [1,1] and another cluster around the coordinate [100, 100]. A decision boundary $y=100-x$ would thus partition the space so that one cluster is generally above the line of the function and one cluster is generally below the line of the function.

[00252] Finding the kernel function may be an automated process, or it may involve human interaction. For example, a Monte Carlo process may be used to search for a kernel function in an automated process. In a human-involved process, a computer may present a human with a series of 2 dimension views of the vector and the human can create 2 dimensional functions to partition the 2 dimensional space, and the computer system can compose a higher dimensional function from these 2 dimensional functions.

[00253] Classifiers are trained with mapped feature sets 2006. With the feature sets now clustered, the training data can be trained in order to identify which clusters are indicative of a particular state. In some cases, this may be a supervised training. In supervised training, a human can identify clusters and provide labels for each cluster. For example, each time window may be tagged by a different process to identify the bed state when the pressure readings for the time window are generated. In some cases, an explicit test may be run to generate the data. A human subject may be instructed to sleep on the bed with attached accessories that can measure signals to

determine reference sleep stages. Logs from this test session may be annotated with the different sleep states so that pressure data when a user is in light sleep are so labeled, pressure data when the user is in REM are so labeled, etc.

[00254] In some cases, other sleep stage identification processes may be used. For example, a threshold analysis may be used to produce reliable state annotations, but such an analysis may require significantly longer than a few second of pressure data. In such a case, a threshold analysis may be run over historic pressure data to label the state of the pressure data. Because this historic analysis can be run after-the-fact, it may be useful for this purpose even if it is not useful or not as useful for real-time state determination for purposes such as home automation. That is to say, an analysis that takes 3 minutes of data to make a determination may be used here even if the analysis would produce an unacceptable 3-minute lag in turning off lights.

[00255] In some cases, the training may be unsupervised training. For example, the training may be performed only with analysis of the pressure data and no outside intelligence provided. This may include unsupervised clustering of the data. Clustering techniques include, but are not limited to, k-means clustering, mixture modeling, hierarchical clustering, self-organizing mapping, and hidden Markov modelling. This may also or alternatively include unsupervised labeling of the data. For example, instead of training the data with a predetermined set of a predetermined number of states, instead the supervision may produce a number of clusters and use that number of clusters to determine the number of possible states. These states may be given a unique identifier that does not have any particular meaning (e.g., cluster1, cluster2, stateA, stateB). Then, once supervision is finished, a human can analyze the state information to determine meaningful labels for the states.

[00256] FIG. 21 is a flowchart of an example process 2100 to encourage restful sleep. The process 2100 can be used, for example, as part of the process to determine sleep state 1926.

[00257] In the process 2100, it is determined if sleep state should change. In order to make this determination, a variety of data can be used. For example, a most current sleeps state 2104 and/or previous sleep states from the time the user enters the bed 2106 can be used. In some implementations, a bed controller can log the sleep states and sleep state transitions that a user goes through as they sleep.

[00258] Motion information 2108 and physiologic information 2110 can be used. For example, a bed controller can log body motion, heart action, breathing action, body temperature, and other data about this user. This data can be compared with sleep state information in order to identify if a sleep state change for the user should occur.

[00259] Once a sleep state transition is determined to be beneficial for the user, the process 2100 can move to determining the next sleep state 2112. Determining the next sleep state can involve comparing the data 2104-2110 with a sleep state configuration 2114 in order to select one of a plurality of possible sleep states for the user to transition to. For example, the sleep state configuration 2114 can include a plurality of rules defining a next ideal sleep state for a user along with home-automation instruction (e.g., change bed pressure) designed to encourage the user to move to the determined next sleep state.

[00260] FIG. 22 shows an example system 2200 for generating new classifiers. In this example, a set of beds 2202 generates pressure readings that are used to generate classifiers that are installed on a set of beds 2208. For example, the beds 2202 can report pressure readings to a classifier server 2204. The classifier server 2204 can generate classifiers, and provide the classifiers to a software server 2206. The software server 2206 can generate a software installation or update for the beds 2208.

[00261] This type of system may be used, for example, in preparing a new model of bed or operating system for market. In this case, the new bed or operating system may not have a large user-base of bed to provide a variety of training data. Instead, pressure readings from existing beds may be used to create classifiers. These classifiers can be included in a software installation for the new beds, or in a software update. This installation can take the form of a networked installation or update, or may be provided with a physical data-storage device.

[00262] FIG. 23 shows an example system 2300 for generating new classifiers. In this example, a set of beds 2302 generates pressure readings that are used to generate classifiers that are installed on the set of beds 22302. For example, the beds 2302 can report pressure readings to a classifier server 2304. The classifier server 2304 can generate classifiers, and provide the classifiers to a software server 2306.

The software server 2306 can generate a software installation or update for the beds 2302.

[00263] This type of system may be used, for example, to update the beds 2302. For example, the system 2300 may periodically generate new classifiers that are designed to be of higher accuracy than existing classifiers on the beds 2302. This accuracy increase may be a result of having more data available for training, improved techniques for generating classifiers, or from increased personalization of data or classifiers. These classifiers can be included in a software installation for the beds, or in a software update. This installation can take the form of a networked installation or update, or may be provided with a physical data-storage device.

[00264] The foregoing detailed description and some embodiments have been given for clarity of understanding only. No unnecessary limitations are to be understood therefrom. It will be apparent to those skilled in the art that many changes can be made in the embodiments described without departing from the scope of the invention. For example, a different order and type of operations may be used to generate classifiers. Additionally, a bed system may aggregate output from classifiers in different ways. Thus, the scope of the present invention should not be limited to the exact details and structures described herein, but rather by the structures described by the language of the claims, and the equivalents of those structures. Any feature or characteristic described with respect to any of the above embodiments can be incorporated individually or in combination with any other feature or characteristic, and are presented in the above order and combinations for clarity only.

WHAT IS CLAIMED IS:

1. A bed system comprising:
 - a first bed comprising:
 - a first mattress;
 - a first pressure sensor in communication with the first mattress to sense pressure applied to the first mattress; and
 - a first controller in data communication with the first pressure sensor, the first controller configured to:
 - receive, from the first pressure sensor, first pressure readings indicative of the sensed pressure of the inflatable chamber; and
 - transmit the first pressure readings to a remote server such that the remote server is able to generate one or more sleep-state classifiers that, when run by a controller on incoming pressure readings, provide a sleep-state vote;
 - a second bed comprising:
 - a second mattress;
 - a second pressure sensor in communication with the second mattress to sense pressure applied to the second mattress; and
 - a second controller in data communication with the second pressure sensor, the controller configured to:
 - receive the one or more sleep-state classifiers;
 - run the received sleep-state classifiers on second pressure readings in order to collect one or more sleep-state votes from the running sleep-state classifiers;
 - determine, from the one or more sleep-state votes, a sleep-state of a user on the second bed;
 - responsive to the determined sleep-state, operate the bed system according to the determined sleep-state.
2. The bed system of claim 1, wherein operating the bed system according to the determined sleep-state comprises one of the list comprising turning on a light, turning off a light, turning on a warming feature, changing bed temperature, changing HVAC temperature, changing firmness of the second mattress, and articulating a foundation of the bed system.

3. The bed system of any of the claims 1 to 2, the bed system further comprising the remote server.
4. The bed system of any of the claims 1 to 3, wherein the remote server is physically remote from the first controller and the second controller; and
wherein the remote server is in data communication with the first controller and the second controller.
5. The bed system of any of the claims 1 to 4, wherein the remote server is configured to:
 - generate training data from the first pressure data;
 - generate, from the training data, the one or more sleep-state classifiers; and
 - send, to the second controller, the one or more sleep-state classifiers.
6. The bed system of any of the claims 1 to 5, wherein generating, from the training data, the one or more sleep-state classifiers comprises:
 - generating a feature set from the training data;
 - mapping the training data to a kernel space;
 - training a classifier with the feature set so that, based on the training data in kernel space, the classifier is able to classify unseen data.
7. The bed system of any of the claims 1 to 6, wherein training a classifier comprises unsupervised training.
8. The bed system of any of the claims 1 to 7, wherein the unsupervised training comprises at least one of the group comprising k-means clustering, mixture modeling, hierarchical clustering, self-organizing mapping, and hidden Markov modelling.
9. The bed system of any of the claims 1 to 8, wherein training a classifier comprises supervised training.
10. The bed system of any of the claims 1 to 9, wherein the supervised training comprises providing the remote server with a set of annotations for the training data.

11. The bed system of any of the claims 1 to 10, wherein the annotations for the training data are provided by a human.
12. The bed system of any of the claims 1 to 11, wherein the annotations for the training data are provided programmatically.
13. The bed system of any of the claims 1 to 12, wherein generating the one or more presence classifiers comprises training a deep learning model on the training data.
14. The bed system of any of the claims 1 to 13, wherein training the deep learning model on the training data comprises generating an initial neural network configured to receive pressure data and generate presence votes.
15. The bed system of any of the claims 1 to 14, wherein the presence vote comprises a presence classification and a confidence value.
16. The bed system of any of the claims 1 to 15, wherein generating the one or more presence classifiers comprises:
 - determining a loss value for the initial neural network; and
 - iteratively refining, beginning with the initial neural network, to a final neural network having a lower loss value than the initial neural network.
17. The bed system of claim 16, wherein the iterative refining is performed with a gradient descent process until a lower loss value cannot be found with the gradient descent process.
18. The bed system of any of the claims 1 to 17, wherein a particular sleep-state classifier is used for multiple users in multiple beds.
19. The bed system of any of the claims 1 to 18, wherein the sleep-state classifiers are personalized for a single user such that the sleep-state classifiers are generated from training data of the single user's use of the bed system and the sleep-state classifiers are used to detect sleep-state of the single user on the second bed.

20. The bed system of any of the claims 1 to 19, wherein a second set of sleep-state classifiers are personalized for a second user such that the second set of sleep-state classifiers are generated from training data of the second user's use of the bed system and the second set of sleep-state classifiers are used to detect sleep-state of the second user on the second bed.
21. The bed system of any of the claims 1 to 20, wherein determining, from the one or more sleep-state votes, a sleep-state of a user on the second bed is personalized for a single user such that votes from different classifiers are weighed based on the classifiers historical accuracy for that user.
22. The bed system of any of the claims 1 to 21, wherein the first bed and the second bed are separate beds.
23. The bed system of any of the claims 1 to 22, wherein the first bed and the second bed are the same beds.
24. The bed system of any of the claims 1 to 23, wherein to run the received sleep-state classifiers on second pressure readings in order to collect one or more sleep-state votes from the running sleep-state classifiers, the second controller is configured to run the received sleep-state classifiers on a plurality of sleep-state classifiers in order to collect one or more sleep-state votes from the running sleep-state classifiers.
25. The bed system of any of the claims 1 to 24, wherein:
 - the first mattress comprises a first inflatable chamber;
 - the first pressure sensor is in fluid communication with the first inflatable chamber to sense the pressure applied to the first mattress;
 - the second mattress comprises a second inflatable chamber; and
 - the second pressure sensor is in fluid communication with the second inflatable chamber to sense the pressure applied to the second mattress.
26. A bed system comprising:
 - a first bed comprising:
 - a first mattress;

means for supporting the first mattress;
a first pressure sensor in communication with the first mattress to sense pressure applied to the first mattress; and
a first controller in data communication with the first pressure sensor, the first controller configured to:
receive, from the first pressure sensor, first pressure readings indicative of the sensed pressure of the inflatable chamber; and
transmit the first pressure readings to a remote server such that the remote server is able to generate one or more sleep-state classifiers that, when run by a controller on incoming pressure readings, provide a sleep-state vote;
a second bed comprising:
a second mattress;
a second pressure sensor in communication with the second mattress to sense pressure applied to the second mattress; and
a second controller in data communication with the second pressure sensor, the controller configured to:
receive the one or more sleep-state classifiers;
run the received sleep-state classifiers on second pressure readings in order to collect one or more sleep-state votes from the running sleep-state classifiers;
determine, from the one or more sleep-state votes, a sleep-state of a user on the second bed;
responsive to the determined sleep-state, operate the bed system according to the determined sleep-state.

27. A bed configured to determine sleep state of a user in a first bed based at least in part on sensed data of a second user in a second bed.




FIG. 1

100 ↘

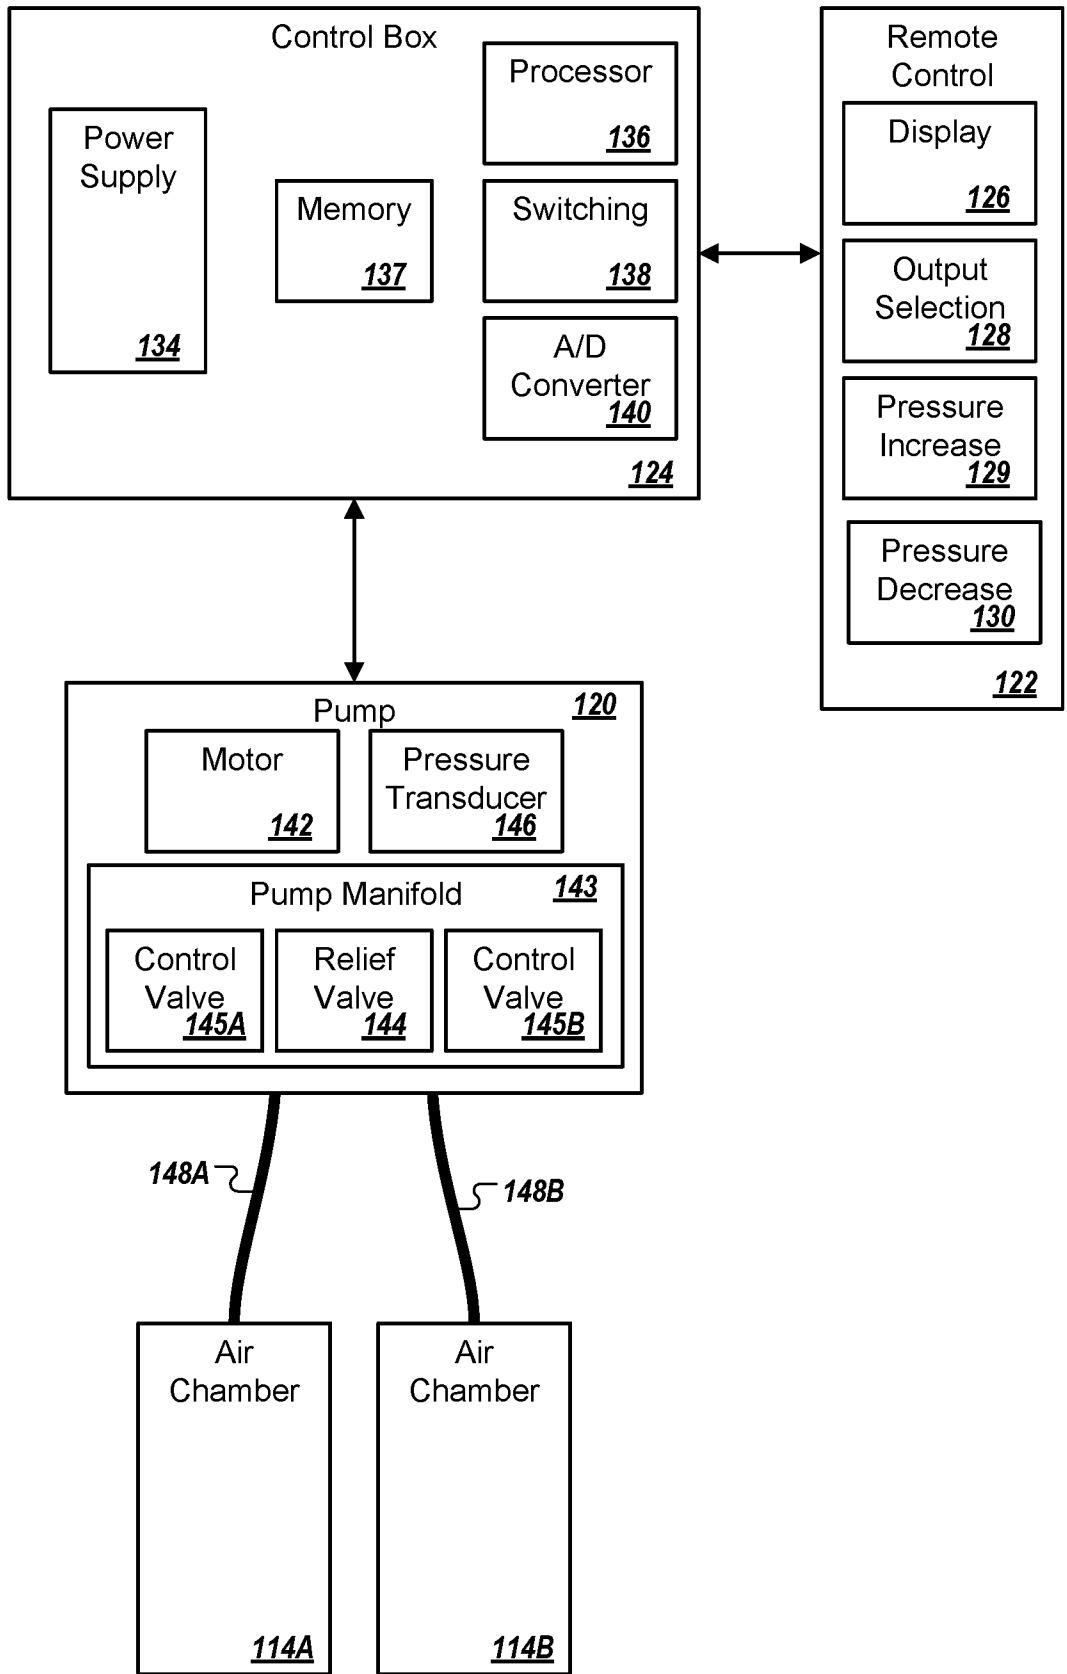


FIG. 2




FIG. 3




FIG. 4A

400 ↷




FIG. 4B

402 ↗

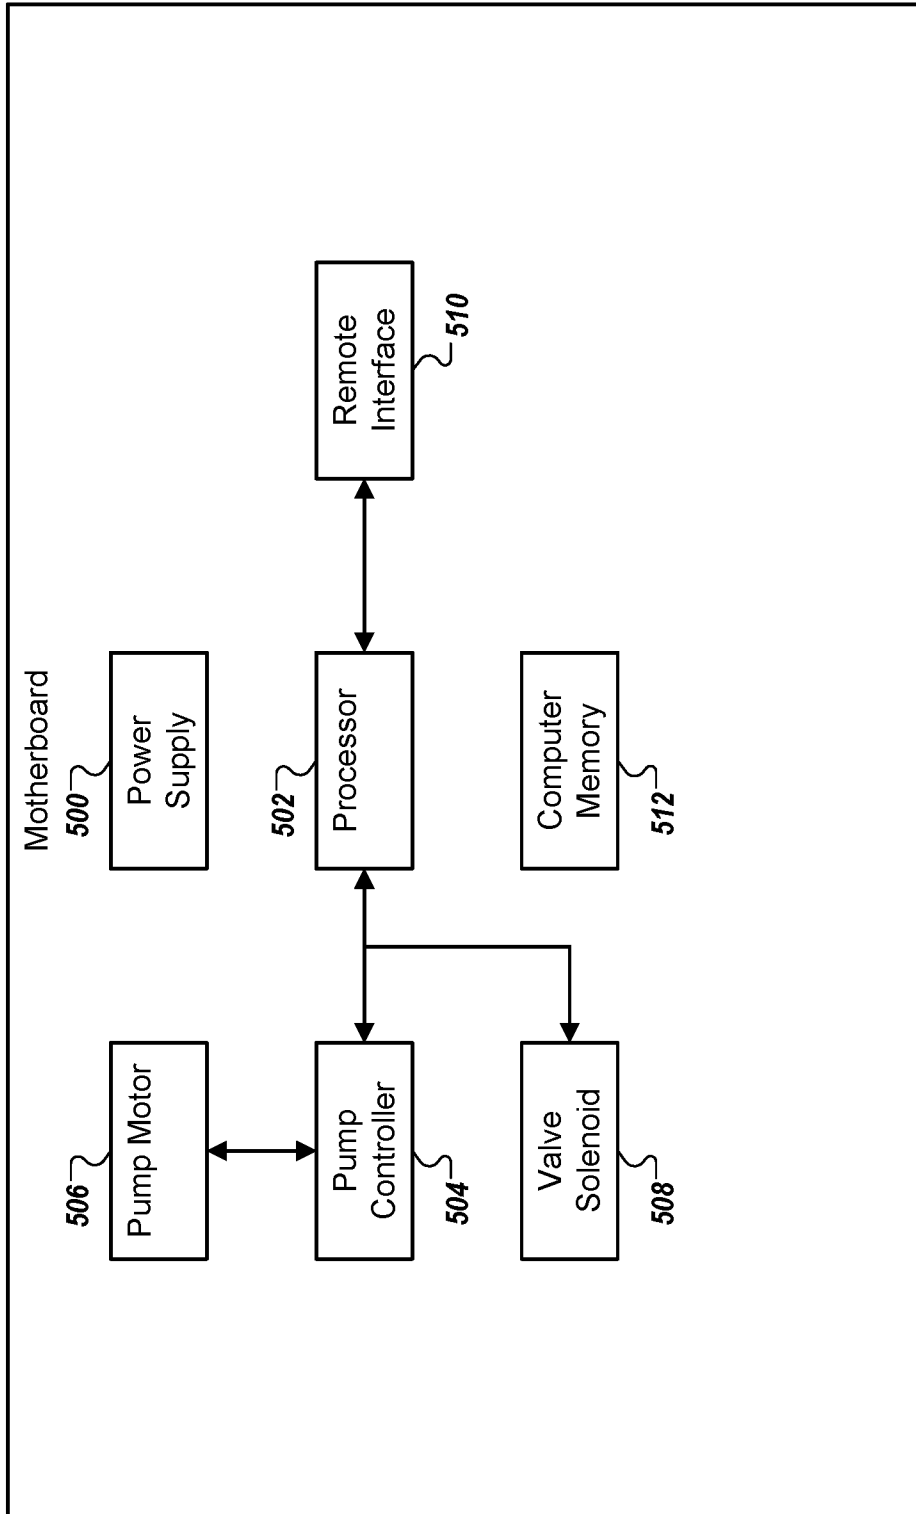


FIG. 5

402 ↗




FIG. 6

404 ↷

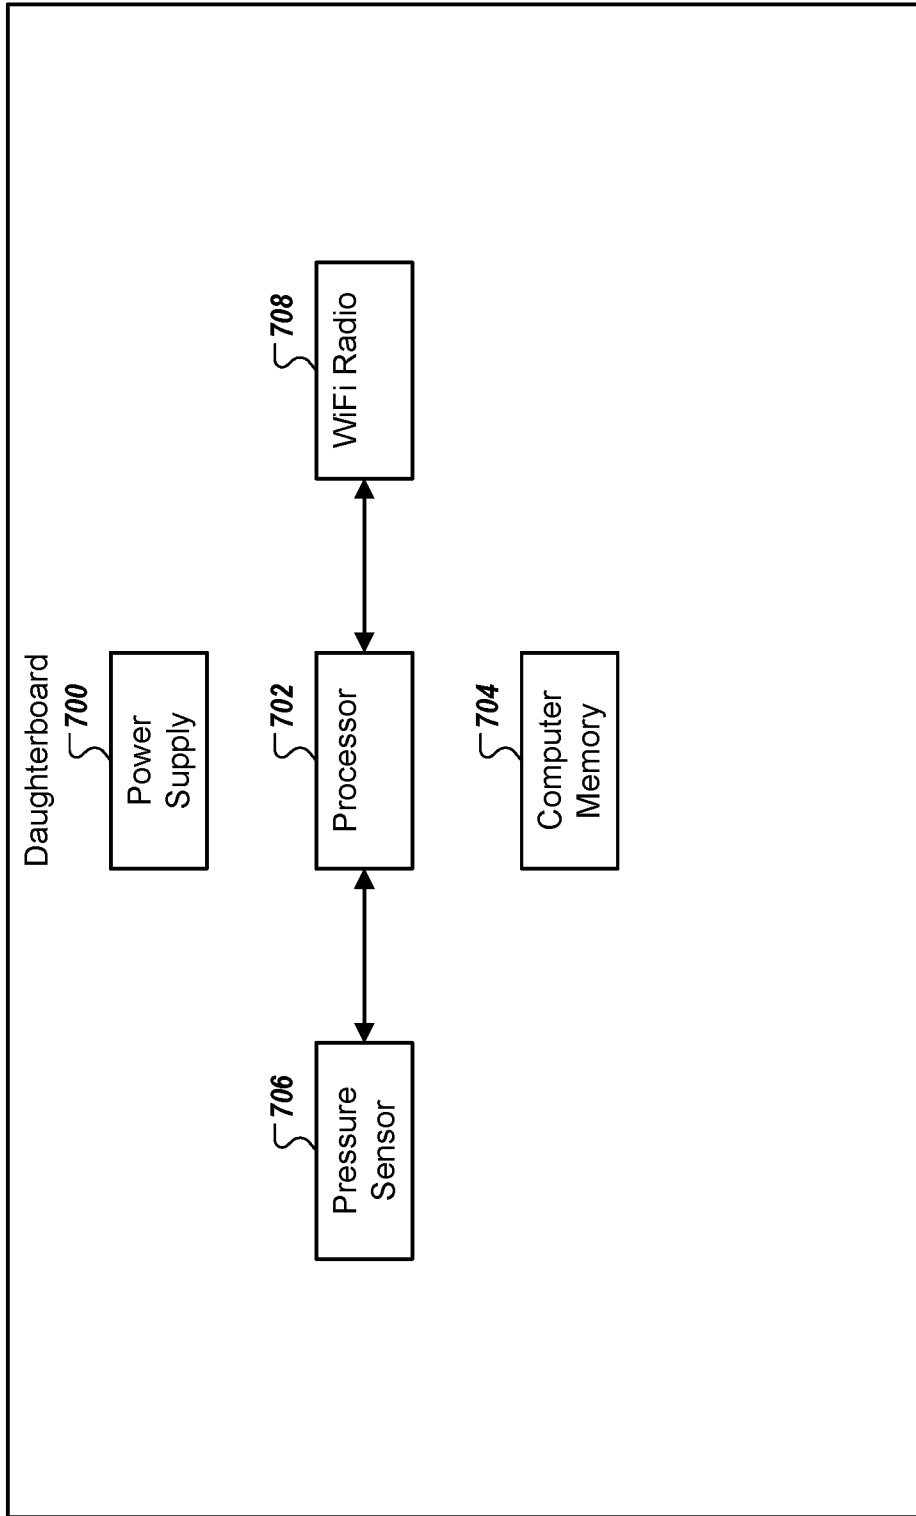


FIG. 7

800 ↗

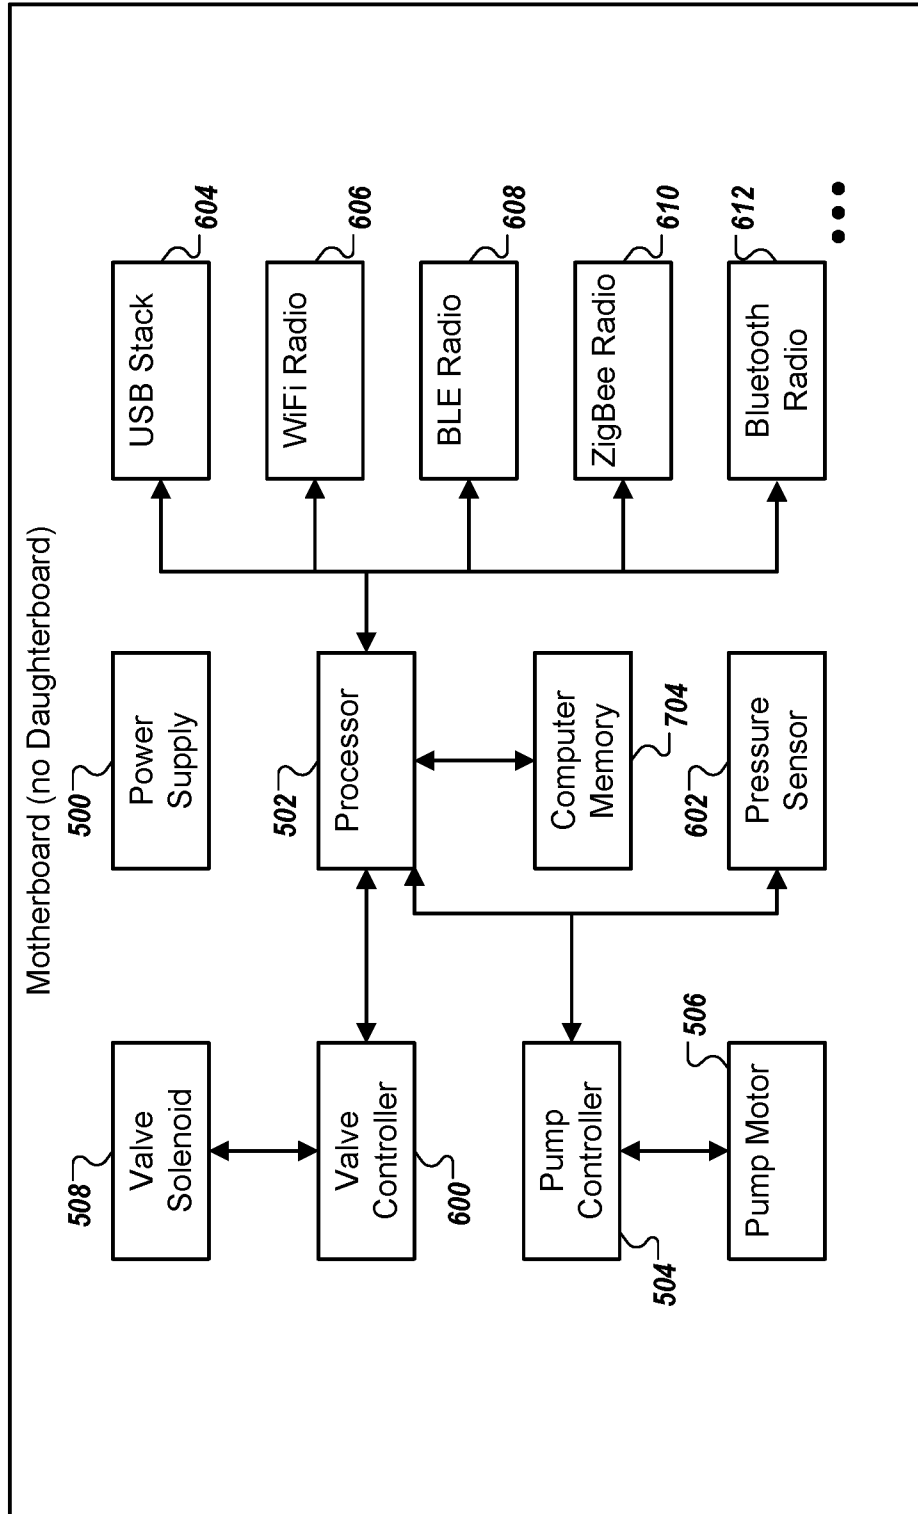


FIG. 8

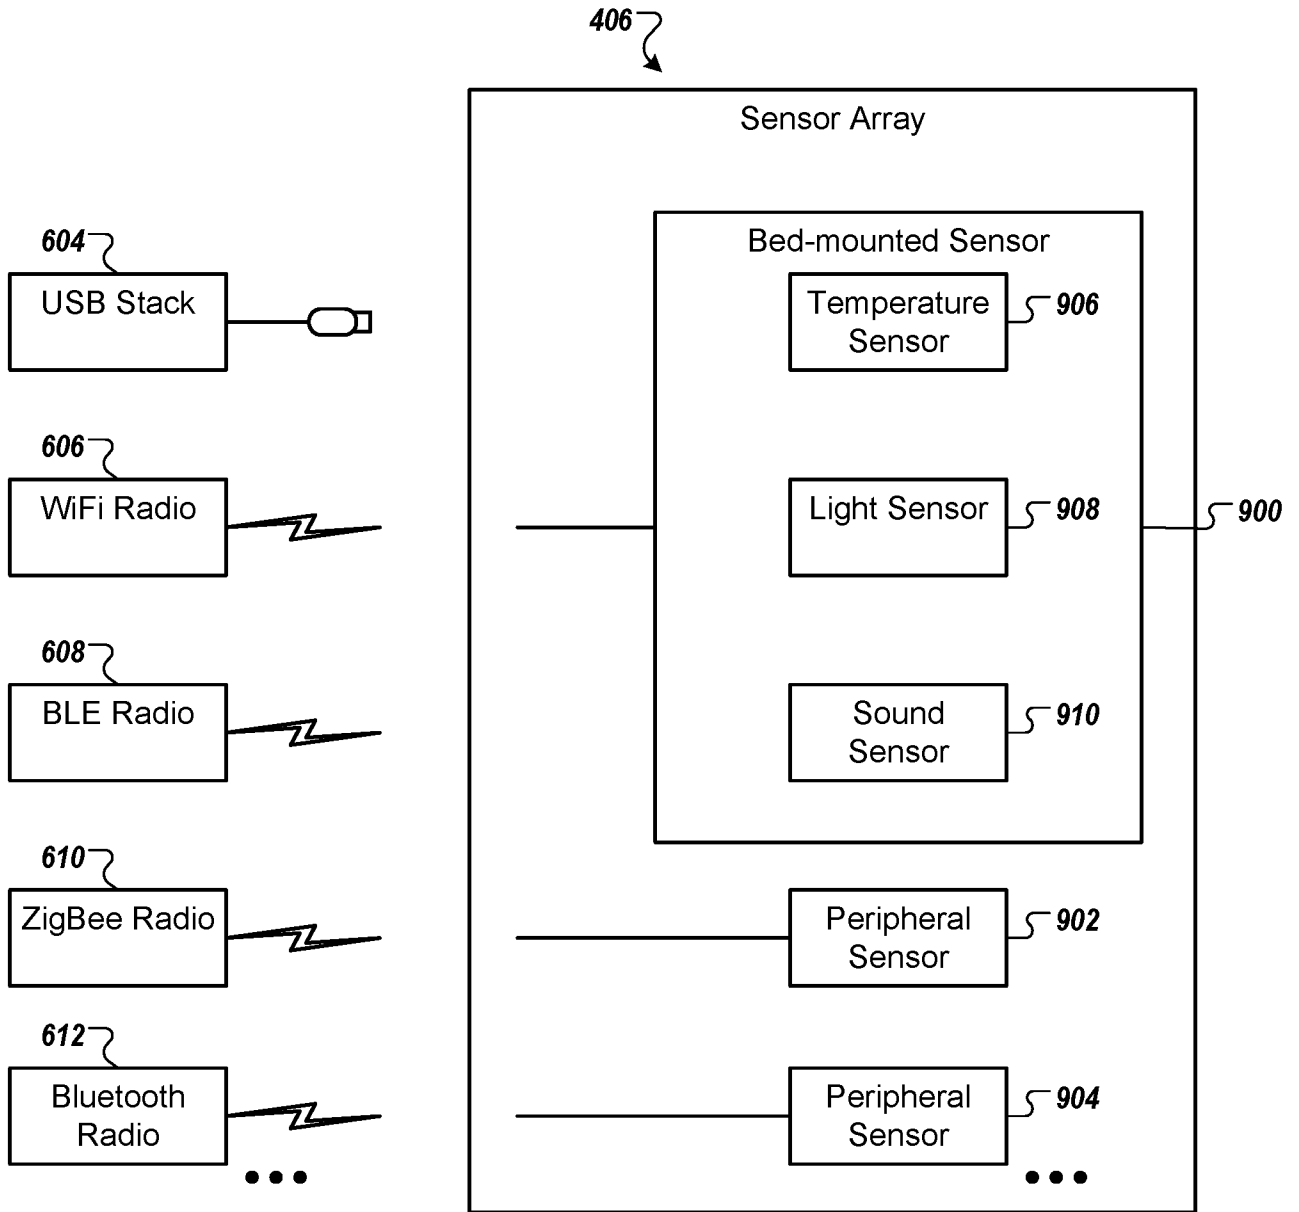


FIG. 9

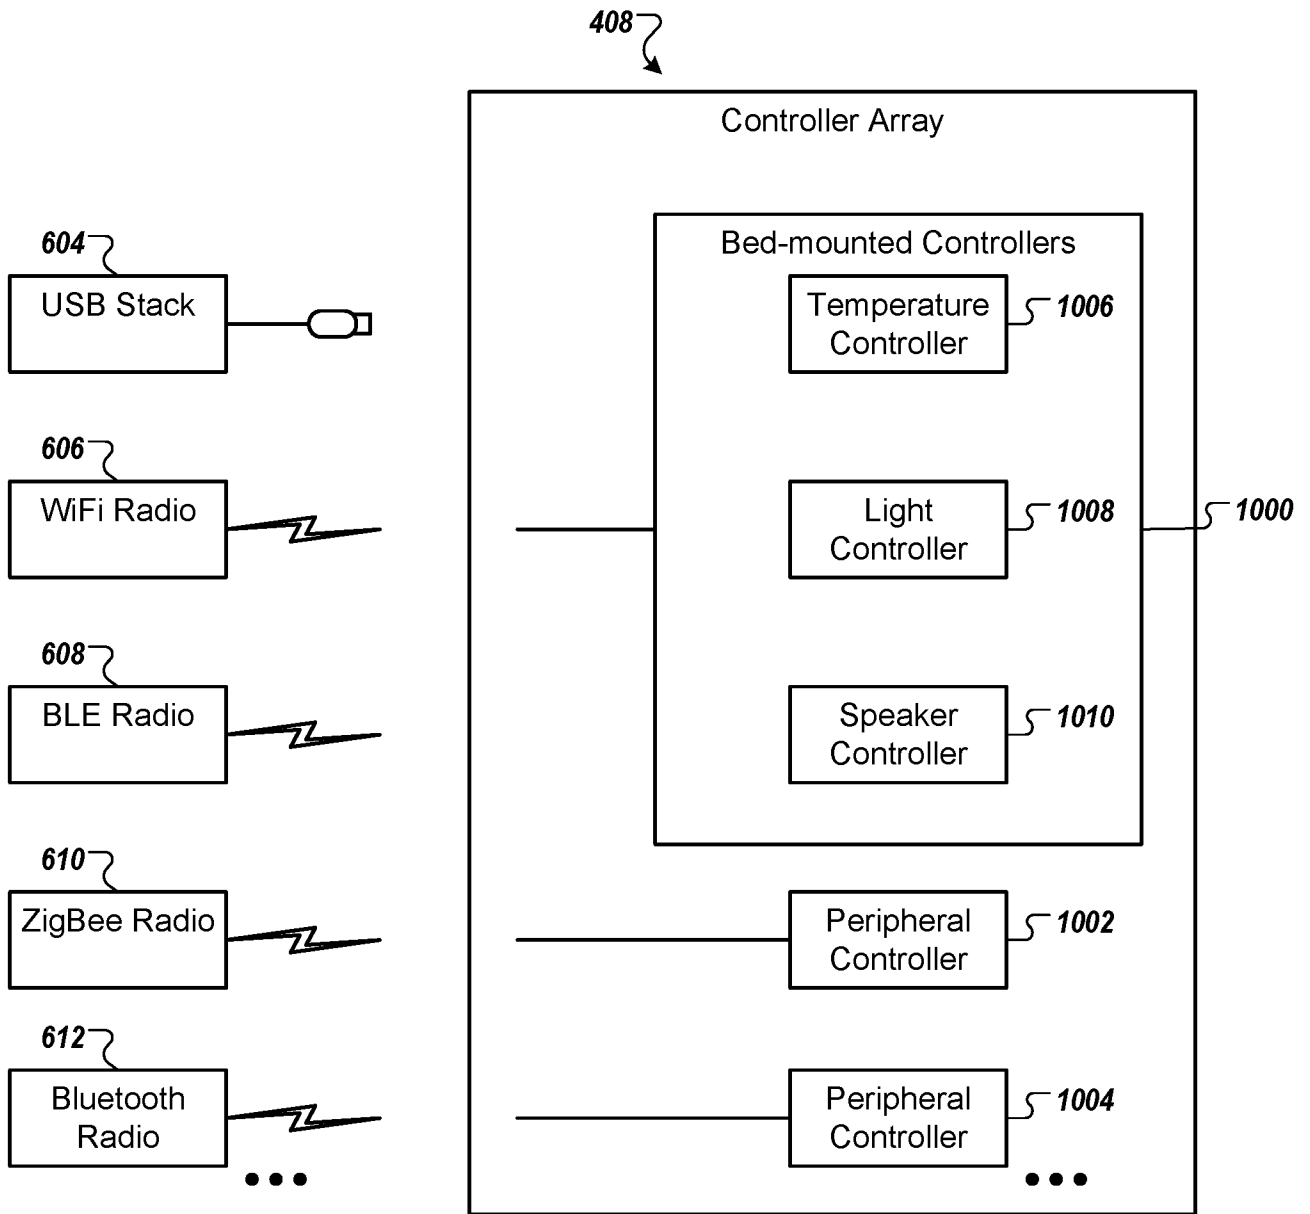


FIG. 10

412 ↗

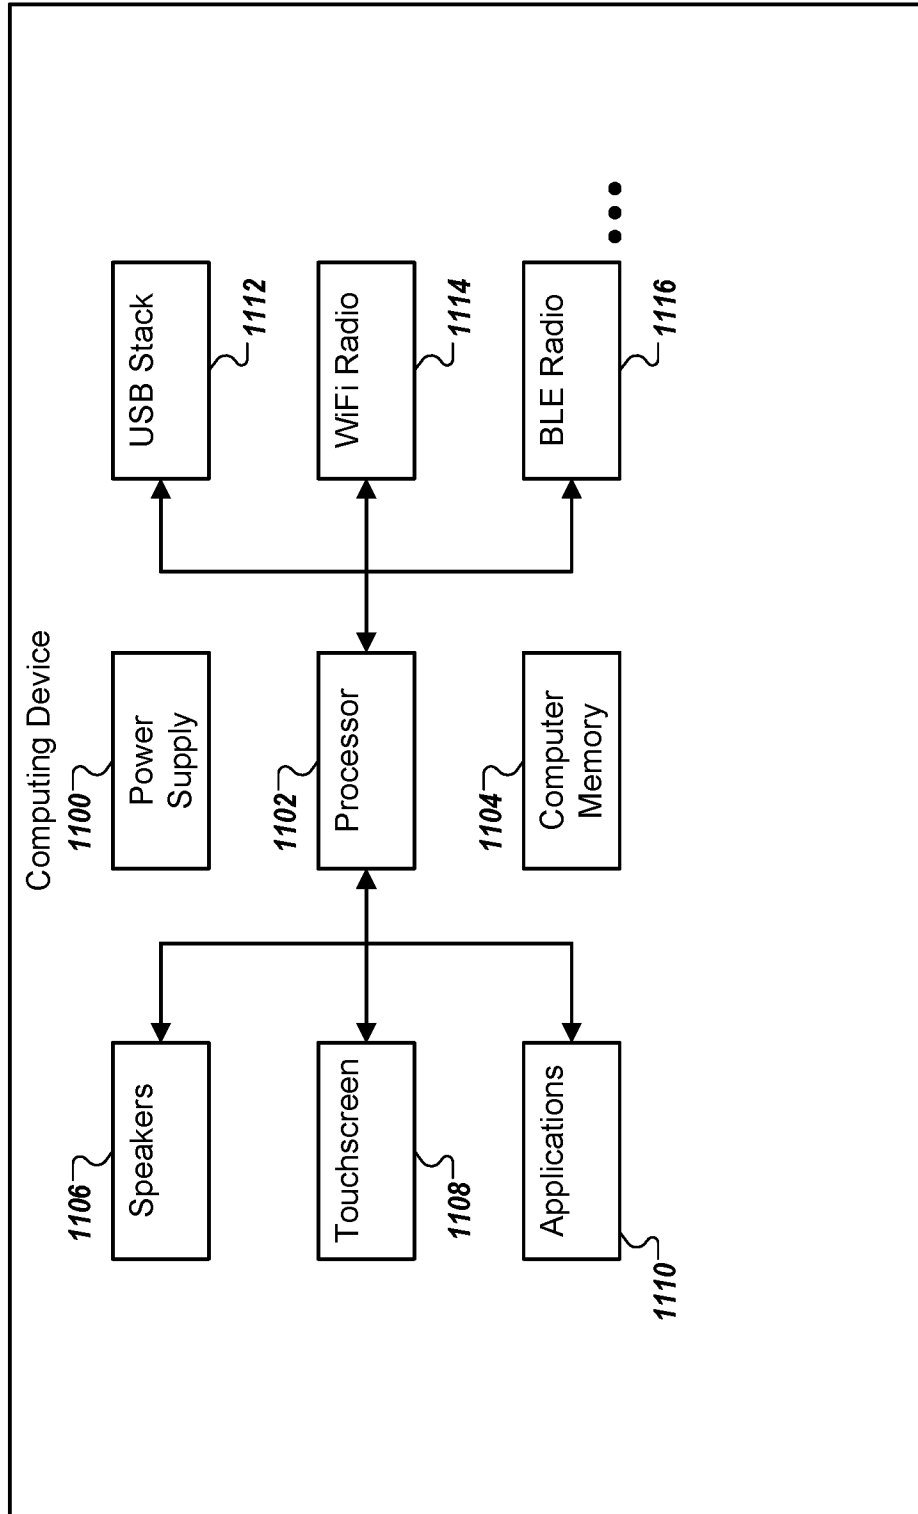


FIG. 11

410a ↷

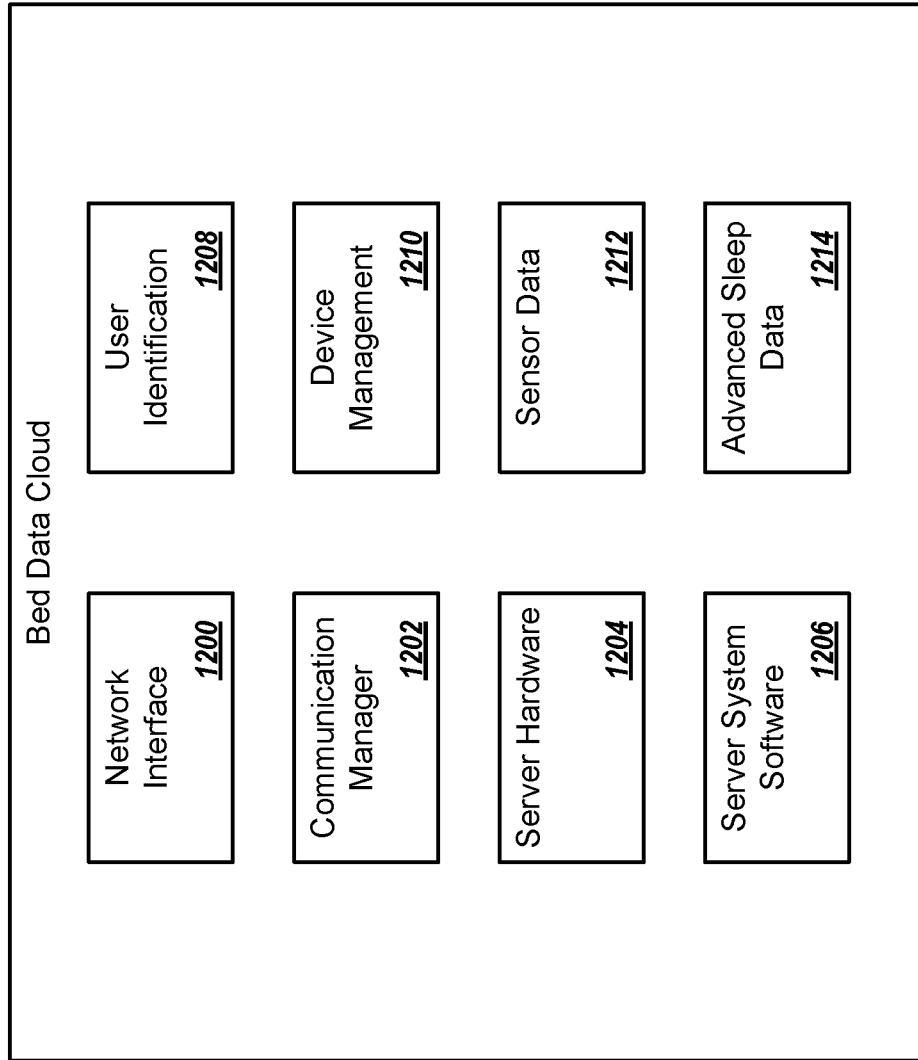


FIG. 12

410b ↷

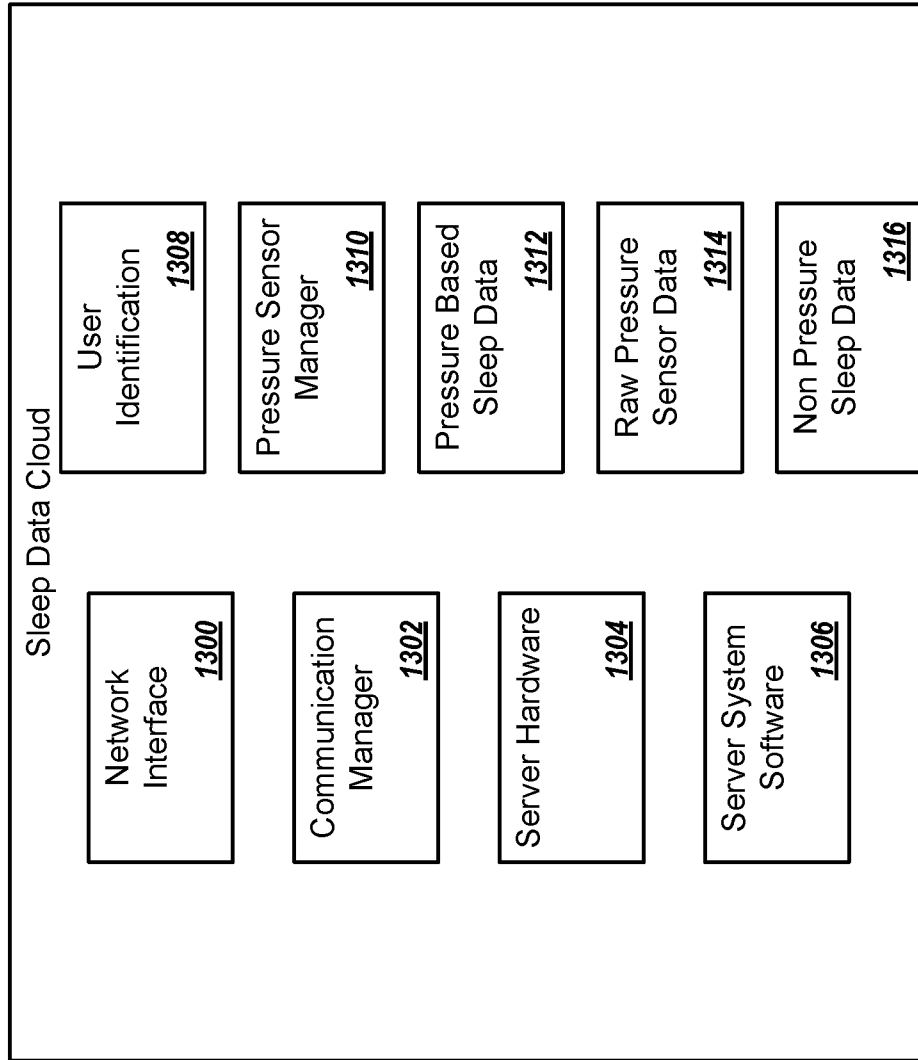


FIG. 13

410c ↷

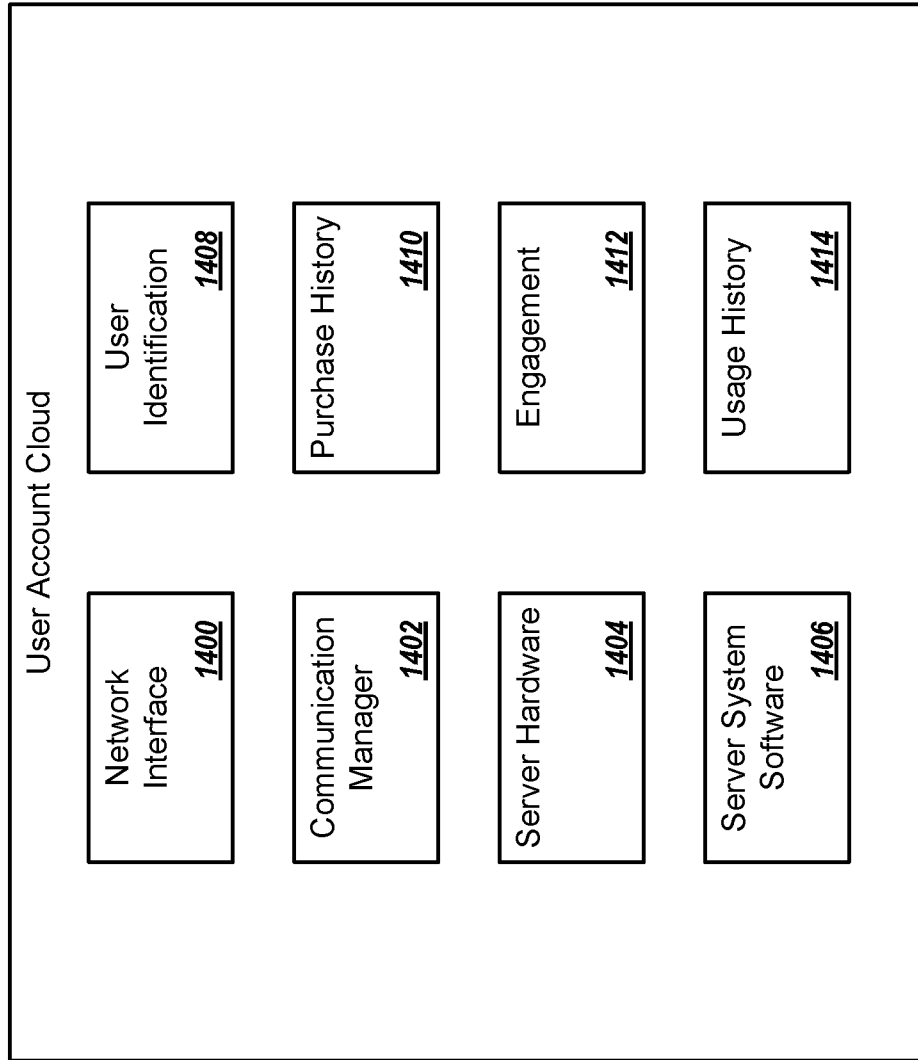


FIG. 14

1500 ↗




FIG. 15

1600 ↗




FIG. 16

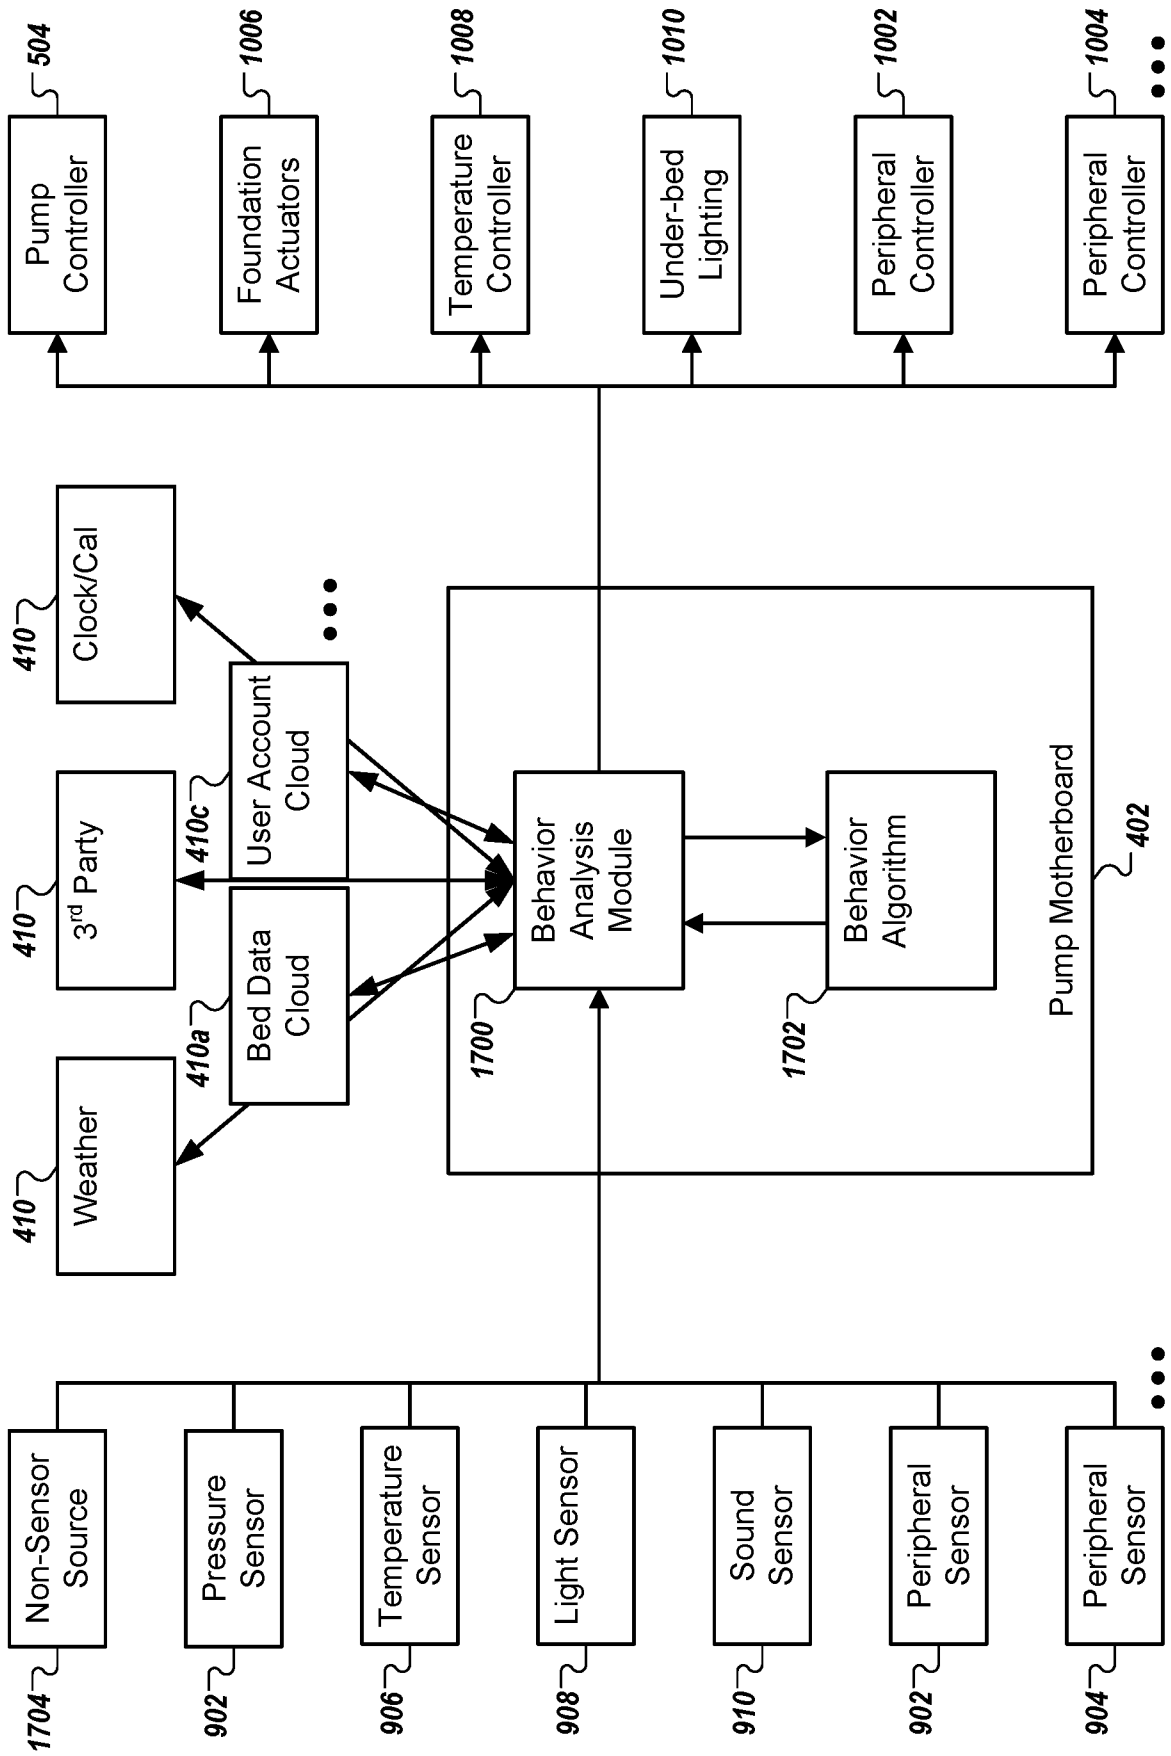


FIG. 17




FIG. 18

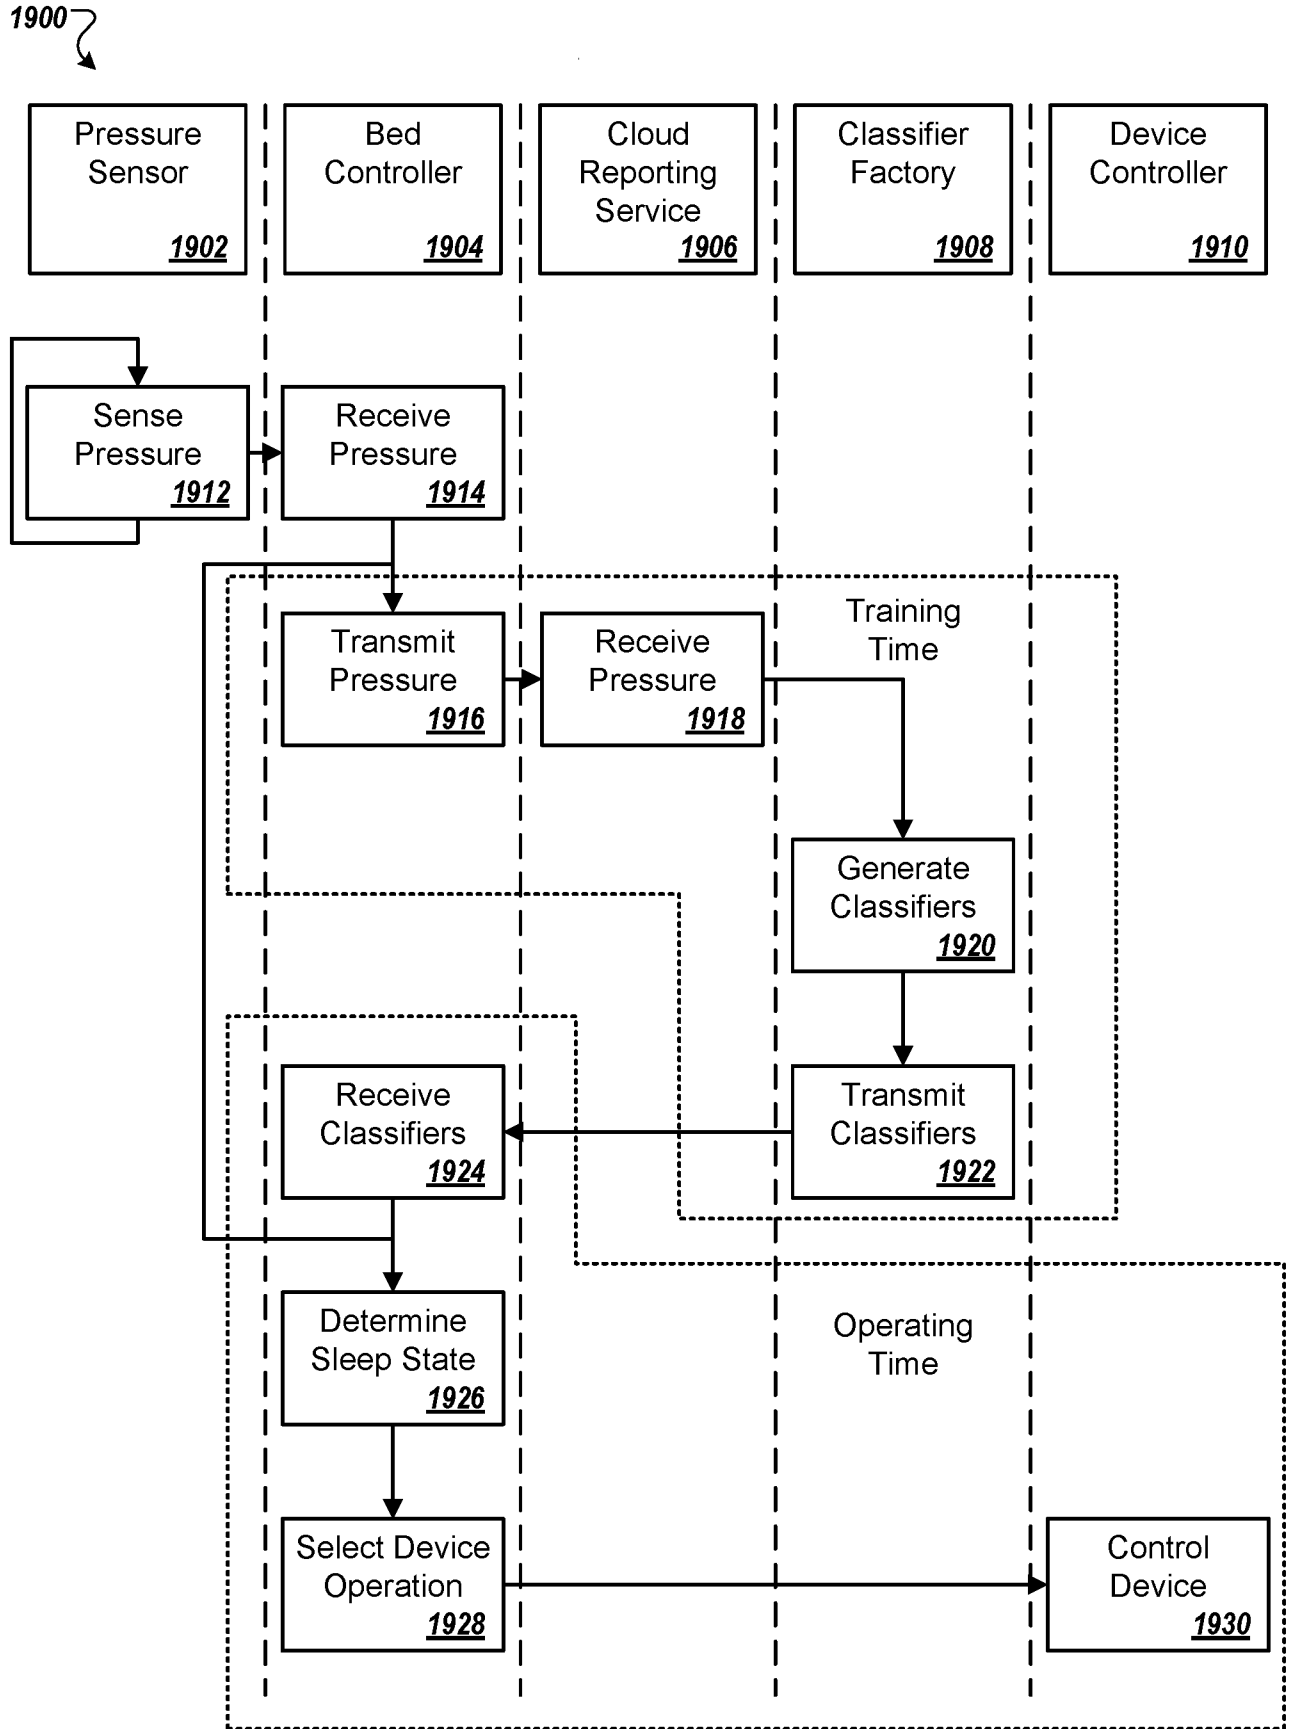


FIG. 19A




FIG. 19B

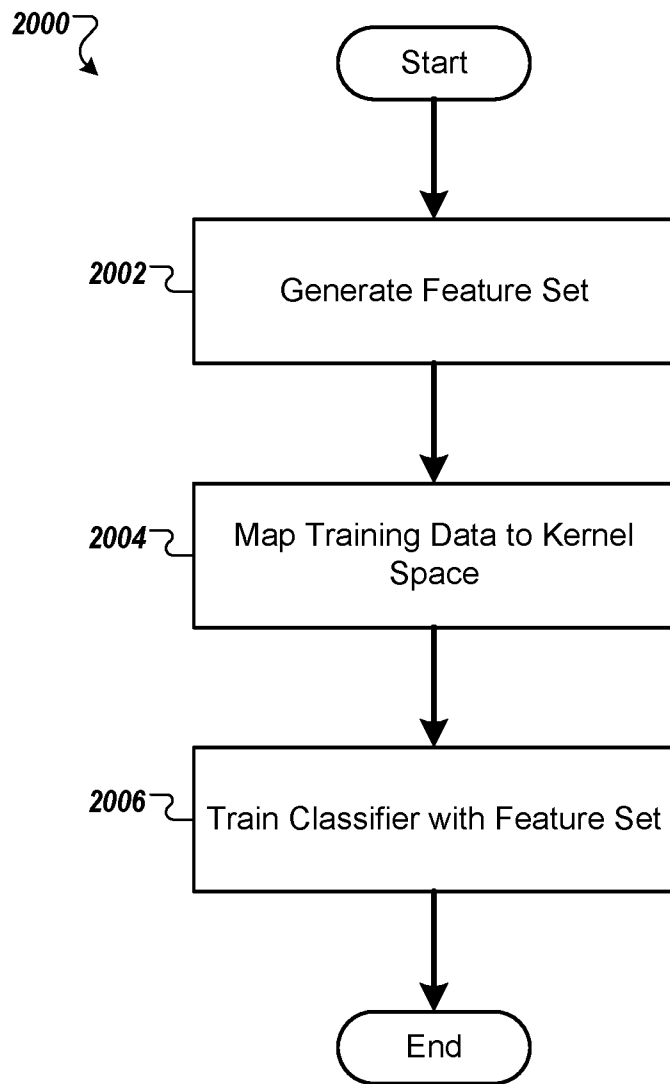


FIG. 20

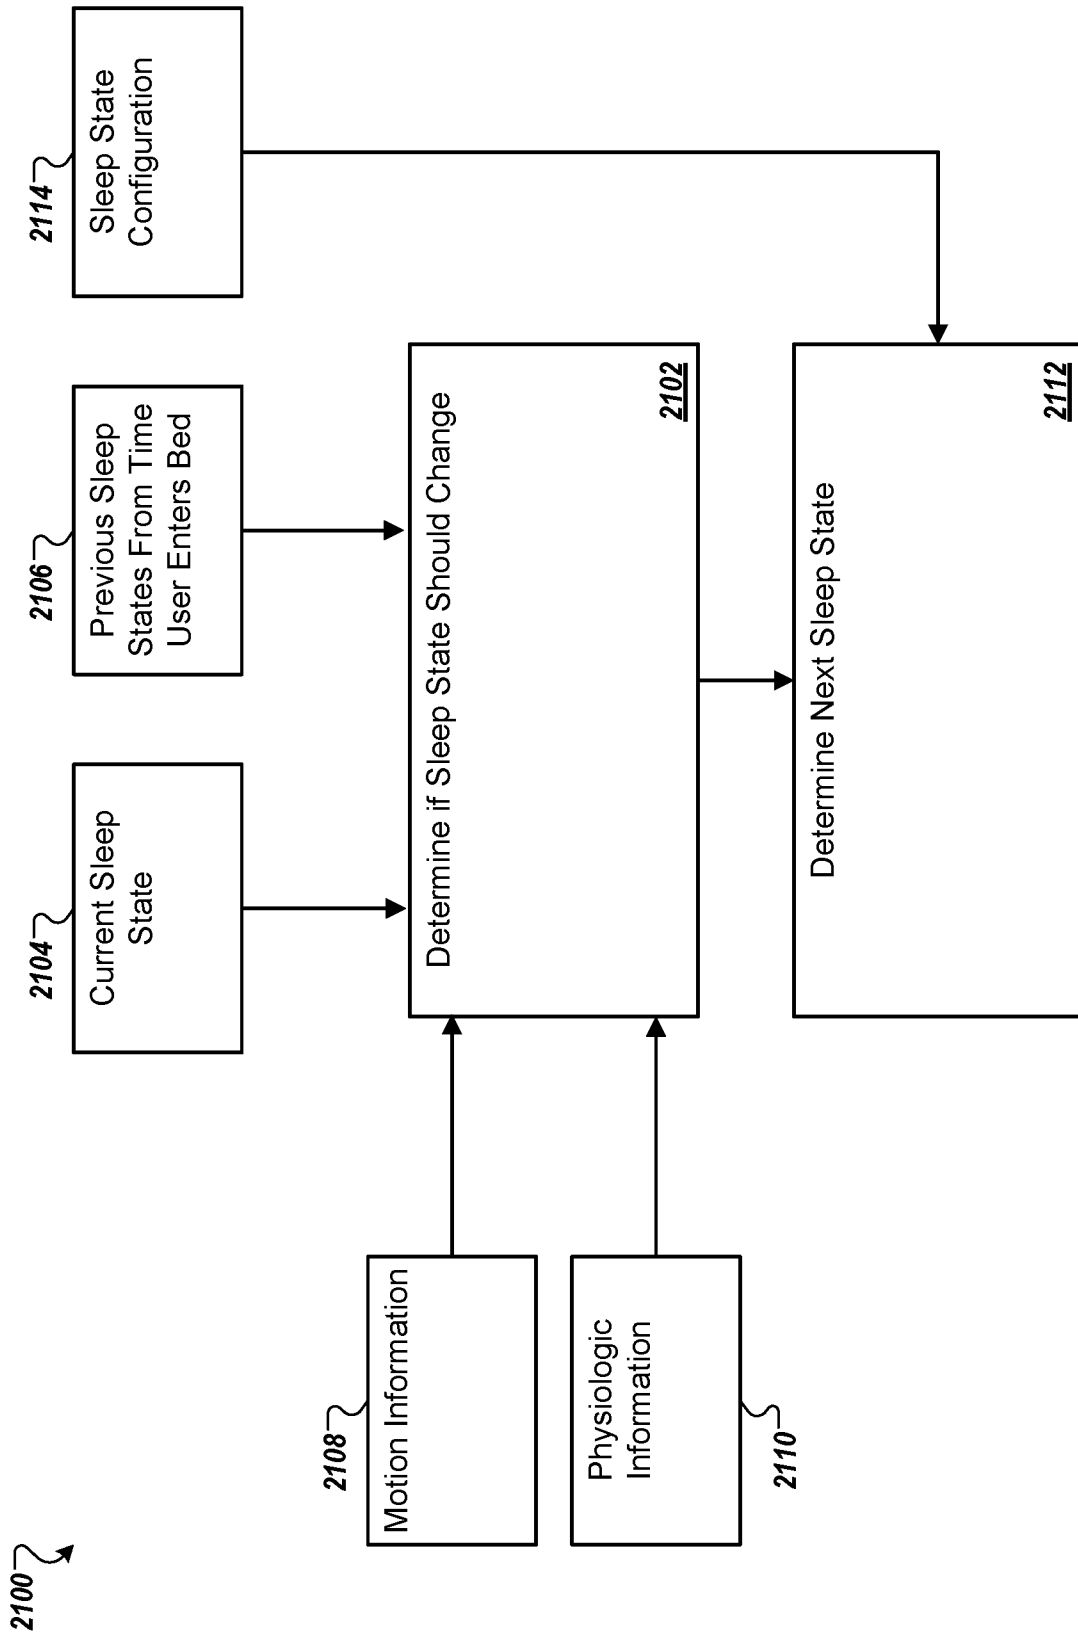


FIG. 21

2200 ↷




FIG. 22




FIG. 23

INTERNATIONAL SEARCH REPORT

International application No
PCT/US2018/067606

A. CLASSIFICATION OF SUBJECT MATTER INV. A47C27/08 A61B5/00 ADD.		
According to International Patent Classification (IPC) or to both national classification and IPC		
B. FIELDS SEARCHED Minimum documentation searched (classification system followed by classification symbols) A47C A61B A61G		
Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched		
Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) EPO-Internal, WPI Data		
C. DOCUMENTS CONSIDERED TO BE RELEVANT		
Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	WO 2016/053398 A1 (DARMA INC [US]) 7 April 2016 (2016-04-07)	27
A	paragraphs [0040], [0041], [0057], [0059], [0091], [0093]; claims; figures -----	1-26
X	WO 2016/011403 A1 (SELECT COMFORT CORP [US]) 21 January 2016 (2016-01-21)	27
A	paragraph [0031]; figures -----	1-26
<input type="checkbox"/> Further documents are listed in the continuation of Box C. <input checked="" type="checkbox"/> See patent family annex.		
* Special categories of cited documents :		
"A" document defining the general state of the art which is not considered to be of particular relevance "E" earlier application or patent but published on or after the international filing date "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) "O" document referring to an oral disclosure, use, exhibition or other means "P" document published prior to the international filing date but later than the priority date claimed		"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art "&" document member of the same patent family
Date of the actual completion of the international search <p align="center">15 February 2019</p>		Date of mailing of the international search report <p align="center">25/02/2019</p>
Name and mailing address of the ISA/ European Patent Office, P.B. 5818 Patentlaan 2 NL - 2280 HV Rijswijk Tel. (+31-70) 340-2040, Fax: (+31-70) 340-3016		Authorized officer <p align="center">Kis, Pál</p>

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No PCT/US2018/067606

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 2016053398 A1	07-04-2016	CN 106793878 A	31-05-2017
		CN 107072565 A	18-08-2017
		JP 2017534353 A	24-11-2017
		JP 2017535316 A	30-11-2017
		US 2016089031 A1	31-03-2016
		US 2016089059 A1	31-03-2016
		US 2017215735 A1	03-08-2017
		US 2017215736 A1	03-08-2017
		WO 2016053398 A1	07-04-2016
		WO 2016054294 A1	07-04-2016

WO 2016011403 A1	21-01-2016	CA 2955365 A1	21-01-2016
		CN 107072406 A	18-08-2017
		EP 3169193 A1	24-05-2017
		US 2016015184 A1	21-01-2016
		US 2019021513 A1	24-01-2019
		WO 2016011403 A1	21-01-2016
