



(51) International Patent Classification:

G05D 23/30 (2006.01) F24F 11/02 (2006.01)
G05D 23/24 (2006.01)

(21) International Application Number:

PCT/US2012/058187

(22) International Filing Date:

30 September 2012 (30.09.2012)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

61/550,346 21 October 2011 (21.10.2011) US

(71) Applicant: NEST LABS, INC. [US/US]; 900 Hansen Way, Palo Alto, California 94304 (US).

(72) Inventors: MATSUOKA, Yoky; 900 Hansen Way, Palo Alto, California 94304 (US). LEE, Eric, A.; 900 Hansen Way, Palo Alto, California 94304 (US). HALES, Steven, A.; 900 Hansen Way, Palo Alto, California 94304 (US). STEFANSKI, Mark, D.; 900 Hansen Way, Palo Alto,

California 94304 (US). SHARAN, Rangoli; 900 Hansen Way, Palo Alto, California 94304 (US).

(74) Agent: BERGSTROM, Robert, W.; Olympic Patent Works PLLC, P.O. Box 4277, Seattle, Washington 98194-0277 (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, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, 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, 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, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ,

[Continued on next page]

(54) Title: AUTOMATED CONTROL-SCHEDULE ACQUISITION WITHIN AN INTELLIGENT CONTROLLER





FIG. 12

(57) Abstract: The current application is directed to intelligent controllers that initially aggressively learn, and then continue, in a steady-state mode, to monitor, learn, and modify one or more control schedules that specify a desired operational behavior of a device, machine, system, or organization controlled by the intelligent controller. An intelligent controller generally acquires one or more initial control schedules through schedule-creation and schedule-modification interfaces or by accessing a default control schedule stored locally or remotely in a memory or mass-storage device. The intelligent controller then proceeds to learn, over time, a desired operational behavior for the device, machine, system, or organization controlled by the intelligent controller based on immediate-control inputs, schedule-modification inputs, and previous and current control schedules, encoding the desired operational behavior in one or more control schedules and/or sub-schedules.





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, ML, MR, NE, SN, TD, TG).

Published:

- *with international search report (Art. 21(3))*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

Declarations under Rule 4.17:

- *as to the identity of the inventor (Rule 4.17(i))*

AUTOMATED CONTROL-SCHEDULE ACQUISITION WITHIN AN INTELLIGENT CONTROLLER

CROSS-REFERENCE TO RELATED APPLICATION

This application claims the benefit of US Provisional Serial No. 61/550,346 filed October 21, 2011, which is incorporated by reference herein.

TECHNICAL FIELD

The current patent application is directed to machine learning and intelligent controllers and, in particular, to intelligent controllers, and machine-learning methods incorporated within intelligent controllers, that develop and refine one or more control schedules, over time, based on received control inputs and inputs through schedule-creation and schedule-modification interfaces.

BACKGROUND

Control systems and control theory are well-developed fields of research and development that have had a profound impact on the design and development of a large number of systems and technologies, from airplanes, spacecraft, and other vehicle and transportation systems to computer systems, industrial manufacturing and operations facilities, machine tools, process machinery, and consumer devices. Control theory encompasses a large body of practical, system-control-design principles, but is also an important branch of theoretical and applied mathematics. Various different types of controllers are commonly employed in many different application domains, from simple closed-loop feedback controllers to complex, adaptive, state-space and differential-equations-based processor-controlled control system.

Many controllers are designed to output control signals to various dynamical components of a system based on a control model and sensor feedback from the system. Many systems are designed to exhibit a predetermined behavior or mode of operation, and the control components of the system are therefore designed, by traditional design and optimization techniques, to ensure that the predetermined system behavior transpires under

- 2 -

normal operational conditions. A more difficult control problem involves design and implementation of controllers that can produce desired system operational behaviors that are specified following controller design and implementation. Theoreticians, researchers, and developers of many different types of controllers and automated systems continue to seek approaches to controller design to produce controllers with the flexibility and intelligence to control systems to produce a wide variety of different operational behaviors, including operational behaviors specified after controller design and manufacture.

SUMMARY

The current application is directed to intelligent controllers that initially aggressively learn, and then continue, in a steady-state mode, to monitor, learn, and modify one or more control schedules that specify a desired operational behavior of a device, machine, system, or organization controlled by the intelligent controller. An intelligent controller generally acquires one or more initial control schedules through schedule-creation and schedule-modification interfaces or by accessing a default control schedule stored locally or remotely in a memory or mass-storage device. The intelligent controller then proceeds to learn, over time, a desired operational behavior for the device, machine, system, or organization controlled by the intelligent controller based on immediate-control inputs, schedule-modification inputs, and previous and current control schedules, encoding the desired operational behavior in one or more control schedules and/or sub-schedules. Control-schedule learning is carried out in at least two different phases, the first of which favors frequent immediate-control inputs and the remaining of which generally tend to minimize the frequency of immediate-control inputs. Learning occurs following monitoring periods and learning results may be propagated from a new provisional control schedule or sub-schedule associated with a completed monitoring period to one or more related control schedules or sub-schedules associated with related time periods.

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 illustrates a smart-home environment.

- 3 -

Figure 2 illustrates integration of smart-home devices with remote devices and systems.

Figure 3 illustrates information processing within the environment of intercommunicating entities illustrated in Figure 2.

Figure 4 illustrates a general class of intelligent controllers to which the current application is directed.

Figure 5 illustrates additional internal features of an intelligent controller.

Figure 6 illustrates a generalized computer architecture that represents an example of the type of computing machinery that may be included in an intelligent controller, server computer, and other processor-based intelligent devices and systems.

Figure 7 illustrates features and characteristics of an intelligent controller of the general class of intelligent controllers to which the current application is directed.

Figure 8 illustrates a typical control environment within which an intelligent controller operates.

Figure 9 illustrates the general characteristics of sensor output.

Figures 10A-D illustrate information processed and generated by an intelligent controller during control operations.

Figures 11A-E provide a transition-state-diagram-based illustration of intelligent-controller operation.

Figure 12 provides a state-transition diagram that illustrates automated control-schedule learning.

Figure 13 illustrates time frames associated with an example control schedule that includes shorter-time-frame sub-schedules.

Figures 14A-C show three different types of control schedules.

Figures 15A-G show representations of immediate-control inputs that may be received and executed by an intelligent controller, and then recorded and overlaid onto control schedules, such as those discussed above with reference to Figures 14A-C, as part of automated control-schedule learning.

- 4 -

Figures 16A-E illustrate one aspect of the method by which a new control schedule is synthesized from an existing control schedule and recorded schedule changes and immediate-control inputs.

Figures 17A-E illustrate one approach to resolving schedule clusters.

Figures 18A-B illustrate the effect of a prospective schedule change entered by a user during a monitoring period.

Figures 19A-B illustrate the effect of a retrospective schedule change entered by a user during a monitoring period.

Figures 20A-C illustrate overlay of recorded data onto an existing control schedule, following completion of a monitoring period, followed by clustering and resolution of clusters.

Figures 21A-B illustrate the setpoint-spreading operation.

Figures 22A-B illustrate schedule propagation.

Figures 23A-C illustrate new-provisional-schedule propagation using P -value vs. t control-schedule plots.

Figures 24A-I illustrate a number of example rules used to simplify a pre-existing control schedule overlaid with propagated setpoints as part of the process of generating a new provisional schedule.

Figures 25A-M illustrate an example implementation of an intelligent controller that incorporates the above-described automated-control-schedule-learning method.

Figure 26 illustrates three different week-based control schedules corresponding to three different control modes for operation of an intelligent controller.

Figure 27 illustrates a state-transition diagram for an intelligent controller that operates according to seven different control schedules.

Figures 28A-C illustrate one type of control-schedule transition that may be carried out by an intelligent controller.

Figures 29-30 illustrate types of considerations that may be made by an intelligent controller during steady-state-learning phases.

Figure 31A illustrates a perspective view of an intelligent thermostat.

- 5 -

Figures 31B-31C illustrate the intelligent thermostat being controlled by a user.

Figure 32 illustrates an exploded perspective view of the intelligent thermostat and an HVAC-coupling wall dock.

Figures 33A-33B illustrate exploded front and rear perspective views of the intelligent thermostat.

Figures 34A-34B illustrate exploded front and rear perspective views, respectively, of the head unit.

Figures 35A-35B illustrate exploded front and rear perspective views, respectively, of the head unit frontal assembly.

Figures 36A-36B illustrate exploded front and rear perspective views, respectively, of the backplate unit.

Figure 37 shows a perspective view of a partially assembled head unit.

Figure 38 illustrates the head unit circuit board.

Figure 39 illustrates a rear view of the backplate circuit board.

Figures 40A-D illustrate steps for achieving initial learning.

Figures 41A-M illustrate a progression of conceptual views of a thermostat control schedule.

Figures 42A-42B illustrate steps for steady-state learning.

DETAILED DESCRIPTION

The current application is directed to a general class of intelligent controllers that includes many different specific types of intelligent controllers that can be applied to, and incorporated within, many different types of devices, machines, systems, and organizations. Intelligent controllers control the operation of devices, machines, systems, and organizations. The general class of intelligent controllers to which the current application is directed include automated learning components that allow the intelligent controllers to learn desired operational behaviors of the devices, machines, systems, and organizations which they control and incorporate the learned information into control schedules. The subject matter of this patent specification relates to the subject matter of the

- 6 -

following commonly assigned applications, each of which is incorporated by reference herein: US Ser. No. 13/269,501 filed October 7, 2011; and U.S. Ser. No. 13/317,423 filed October 17, 2011.

The current application discloses, in addition to methods and implementations for automated control-schedule learning, a specific example of an intelligent thermostat controller, or intelligent thermostat, and a specific control-schedule-learning method for the intelligent thermostat that serves as a detailed example of the automated control-schedule-learning methods employed by the general class of intelligent controllers to which the current application is directed. The intelligent thermostat is an example of a smart-home device.

The detailed description includes three subsections: (1) an overview of the smart-home environment; (2) automated control-schedule learning; and (3) automated control-schedule learning in the context of an intelligent thermostat. The first subsection provides a description of one area technology that offers many opportunities for of application and incorporation of automated-control-schedule-learning methods. The second subsection provides a detailed description of automated control-schedule learning, including a first, general implementation. The third subsection provides a specific example of an automated-control-schedule-learning method incorporated within an intelligent thermostat.

Overview of the Smart-Home Environment

Figure 1 illustrates a smart-home environment. The smart-home environment 100 includes a number of intelligent, multi-sensing, network-connected devices. These smart-home devices intercommunicate and are integrated together within the smart-home environment. The smart-home devices may also communicate with cloud-based smart-home control and/or data-processing systems in order to distribute control functionality, to access higher-capacity and more reliable computational facilities, and to integrate a particular smart home into a larger, multi-home or geographical smart-home-device-based aggregation.

- 7 -

The smart-home devices may include one more intelligent thermostats 102, one or more intelligent hazard-detection units 104, one or more intelligent entryway-interface devices 106, smart switches, including smart wall-like switches 108, smart utilities interfaces and other services interfaces, such as smart wall-plug interfaces 110, and a wide variety of intelligent, multi-sensing, network-connected appliances 112, including refrigerators, televisions, washers, dryers, lights, audio systems, intercom systems, mechanical actuators, wall air conditioners, pool-heating units, irrigation systems, and many other types of intelligent appliances and systems.

In general, smart-home devices include one or more different types of sensors, one or more controllers and/or actuators, and one or more communications interfaces that connect the smart-home devices to other smart-home devices, routers, bridges, and hubs within a local smart-home environment, various different types of local computer systems, and to the Internet, through which a smart-home device may communicate with cloud-computing servers and other remote computing systems. Data communications are generally carried out using any of a large variety of different types of communications media and protocols, including wireless protocols, such as Wi-Fi, ZigBee, 6LoWPAN, various types of wired protocols, including CAT6 Ethernet, HomePlug, and other such wired protocols, and various other types of communications protocols and technologies. Smart-home devices may themselves operate as intermediate communications devices, such as repeaters, for other smart-home devices. The smart-home environment may additionally include a variety of different types of legacy appliances and devices 140 and 142 which lack communications interfaces and processor-based controllers.

Figure 2 illustrates integration of smart-home devices with remote devices and systems. Smart-home devices within a smart-home environment 200 can communicate through the Internet 202 via 3G/4G wireless communications 204, through a hubbed network 206, or by other communications interfaces and protocols. Many different types of smart-home-related data, and data derived from smart-home data 208, can be stored in, and retrieved from, a remote system 210, including a cloud-based remote system. The remote system may include various types of statistics, inference, and indexing engines 212 for data

- 8 -

processing and derivation of additional information and rules related to the smart-home environment. The stored data can be exposed, via one or more communications media and protocols, in part or in whole, to various remote systems and organizations, including charities 214, governments 216, academic institutions 218, businesses 220, and utilities 222. In general, the remote data-processing system 210 is managed or operated by an organization or vendor related to smart-home devices or contracted for remote data-processing and other services by a homeowner, landlord, dweller, or other smart-home-associated user. The data may also be further processed by additional commercial-entity data-processing systems 213 on behalf of the smart-homeowner or manager and/or the commercial entity or vendor which operates the remote data-processing system 210. Thus, external entities may collect, process, and expose information collected by smart-home devices within a smart-home environment, may process the information to produce various types of derived results which may be communicated to, and shared with, other remote entities, and may participate in monitoring and control of smart-home devices within the smart-home environment as well as monitoring and control of the smart-home environment. Of course, in many cases, export of information from within the smart-home environment to remote entities may be strictly controlled and constrained, using encryption, access rights, authentication, and other well-known techniques, to ensure that information deemed confidential by the smart-home manager and/or by the remote data-processing system is not intentionally or unintentionally made available to additional external computing facilities, entities, organizations, and individuals.

Figure 3 illustrates information processing within the environment of intercommunicating entities illustrated in Figure 2. The various processing engines 212 within the external data-processing system 210 can process data with respect to a variety of different goals, including provision of managed services 302, various types of advertizing and communications 304, social-networking exchanges and other electronic social communications 306, and for various types of monitoring and rule-generation activities 308. The various processing engines 212 communicate directly or indirectly with smart-home devices 310-313, each of which may have data-consumer ("DC"), data-source ("DS"), services-consumer ("SC"), and services-source ("SS") characteristics. In addition, the

- 9 -

processing engines may access various other types of external information 316, including information obtained through the Internet, various remote information sources, and even remote sensor, audio, and video feeds and sources.

Automated Schedule Learning

Figure 4 illustrates a general class of intelligent controllers to which the current application is directed. The intelligent controller 402 controls a device, machine, system, or organization 404 via any of various different types of output control signals and receives information about the controlled entity and the environment from sensor output received by the intelligent controller from sensors embedded within the controlled entity 404, the intelligent controller 402, or in the environment of the intelligent controller and/or controlled entity. In Figure 4, the intelligent controller is shown connected to the controlled entity 404 via a wire or fiber-based communications medium 406. However, the intelligent controller may be interconnected with the controlled entity by alternative types of communications media and communications protocols, including wireless communications. In many cases, the intelligent controller and controlled entity may be implemented and packaged together as a single system that includes both the intelligent controller and a machine, device, system, or organization controlled by the intelligent controller. The controlled entity may include multiple devices, machines, system, or organizations and the intelligent controller may itself be distributed among multiple components and discrete devices and systems. In addition to outputting control signals to controlled entities and receiving sensor input, the intelligent controller also provides a user interface 410-413 through which a human user or remote entity, including a user-operated processing device or a remote automated control system, can input immediate-control inputs to the intelligent controller as well as create and modify the various types of control schedules. In Figure 4, the intelligent controller provides a graphical-display component 410 that displays a control schedule 416 and includes a number of input components 411-413 that provide a user interface for input of immediate-control directives to the intelligent controller for controlling the controlled entity or entities and input of scheduling-interface commands that

- 10 -

control display of one or more control schedules, creation of control schedules, and modification of control schedules.

To summarize, the general class of intelligent controllers to which the current is directed receive sensor input, output control signals to one or more controlled entities, and provide a user interface that allows users to input immediate-control command inputs to the intelligent controller for translation by the intelligent controller into output control signals as well as to create and modify one or more control schedules that specify desired controlled-entity operational behavior over one or more time periods. These basic functionalities and features of the general class of intelligent controllers provide a basis upon which automated control-schedule learning, to which the current application is directed, can be implemented.

Figure 5 illustrates additional internal features of an intelligent controller. An intelligent controller is generally implemented using one or more processors 502, electronic memory 504-507, and various types of microcontrollers 510-512, including a microcontroller 512 and transceiver 514 that together implement a communications port that allows the intelligent controller to exchange data and commands with one or more entities controlled by the intelligent controller, with other intelligent controllers, and with various remote computing facilities, including cloud-computing facilities through cloud-computing servers. Often, an intelligent controller includes multiple different communications ports and interfaces for communicating by various different protocols through different types of communications media. It is common for intelligent controllers, for example, to use wireless communications to communicate with other wireless-enabled intelligent controllers within an environment and with mobile-communications carriers as well as any of various wired communications protocols and media. In certain cases, an intelligent controller may use only a single type of communications protocol, particularly when packaged together with the controlled entities as a single system. Electronic memories within an intelligent controller may include both volatile and non-volatile memories, with low-latency, high-speed volatile memories facilitating execution of control routines by the one or more processors and slower, non-volatile memories storing control

- 11 -

routines and data that need to survive power-on/power-off cycles. Certain types of intelligent controllers may additionally include mass-storage devices.

Figure 6 illustrates a generalized computer architecture that represents an example of the type of computing machinery that may be included in an intelligent controller, server computer, and other processor-based intelligent devices and systems. The computing machinery includes one or multiple central processing units ("CPUs") 602-605, one or more electronic memories 608 interconnected with the CPUs by a CPU/memory-subsystem bus 610 or multiple busses, a first bridge 612 that interconnects the CPU/memory-subsystem bus 610 with additional busses 614 and 616 and/or other types of high-speed interconnection media, including multiple, high-speed serial interconnects. These busses and/or serial interconnections, in turn, connect the CPUs and memory with specialized processors, such as a graphics processor 618, and with one or more additional bridges 620, which are interconnected with high-speed serial links or with multiple controllers 622-627, such as controller 627, that provide access to various different types of mass-storage devices 628, electronic displays, input devices, and other such components, subcomponents, and computational resources.

Figure 7 illustrates features and characteristics of an intelligent controller of the general class of intelligent controllers to which the current application is directed. An intelligent controller includes controller logic 702 generally implemented as electronic circuitry and processor-based computational components controlled by computer instructions stored in physical data-storage components, including various types of electronic memory and/or mass-storage devices. It should be noted, at the onset, that computer instructions stored in physical data-storage devices and executed within processors comprise the control components of a wide variety of modern devices, machines, and systems, and are as tangible, physical, and real as any other component of a device, machine, or system. Occasionally, statements are encountered that suggest that computer-instruction-implemented control logic is "merely software" or something abstract and less tangible than physical machine components. Those familiar with modern science and technology understand that this is not the case. Computer instructions executed by processors must be physical entities stored in physical devices. Otherwise, the processors

- 12 -

would not be able to access and execute the instructions. The term "software" can be applied to a symbolic representation of a program or routine, such as a printout or displayed list of programming-language statements, but such symbolic representations of computer programs are not executed by processors. Instead, processors fetch and execute computer instructions stored in physical states within physical data-storage devices.

The controller logic accesses and uses a variety of different types of stored information and inputs in order to generate output control signals 704 that control the operational behavior of one or more controlled entities. The information used by the controller logic may include one or more stored control schedules 706, received output from one or more sensors 708-710, immediate control inputs received through an immediate-control interface 712, and data, commands, and other information received from remote data-processing systems, including cloud-based data-processing systems 713. In addition to generating control output 704, the controller logic provides an interface 714 that allows users to create and modify control schedules and may also output data and information to remote entities, other intelligent controllers, and to users through an information-output interface.

Figure 8 illustrates a typical control environment within which an intelligent controller operates. As discussed above, an intelligent controller 802 receives control inputs from users or other entities 804 and uses the control inputs, along with stored control schedules and other information, to generate output control signals 805 that control operation of one or more controlled entities 808. Operation of the controlled entities may alter an environment within which sensors 810-812 are embedded. The sensors return sensor output, or feedback, to the intelligent controller 802. Based on this feedback, the intelligent controller modifies the output control signals in order to achieve a specified goal or goals for controlled-system operation. In essence, an intelligent controller modifies the output control signals according to two different feedback loops. The first, most direct feedback loop includes output from sensors that the controller can use to determine subsequent output control signals or control-output modification in order to achieve the desired goal for controlled-system operation. In many cases, a second feedback loop involves environmental or other feedback 816 to users which, in turn, elicits subsequent

- 13 -

user control and scheduling inputs to the intelligent controller 802. In other words, users can either be viewed as another type of sensor that outputs immediate-control directives and control-schedule changes, rather than raw sensor output, or can be viewed as a component of a higher-level feedback loop.

There are many different types of sensors and sensor output. In general, sensor output is directly or indirectly related to some type of parameter, machine state, organization state, computational state, or physical environmental parameter. Figure 9 illustrates the general characteristics of sensor output. As shown in a first plot 902 in Figure 9, a sensor may output a signal, represented by curve 904, over time, with the signal directly or indirectly related to a parameter P , plotted with respect to the vertical axis 906. The sensor may output a signal continuously or at intervals, with the time of output plotted with respect to the horizontal axis 908. In certain cases, sensor output may be related to two or more parameters. For example, in plot 910, a sensor outputs values directly or indirectly related to two different parameters P_1 and P_2 , plotted with respect to axes 912 and 914, respectively, over time, plotted with respect to vertical axis 916. In the following discussion, for simplicity of illustration and discussion, it is assumed that sensors produce output directly or indirectly related to a single parameter, as in plot 902 in Figure 9. In the following discussion, the sensor output is assumed to be a set of parameter values for a parameter P . The parameter may be related to environmental conditions, such as temperature, ambient light level, sound level, and other such characteristics. However, the parameter may also be the position or positions of machine components, the data states of memory-storage address in data-storage devices, the current drawn from a power supply, the flow rate of a gas or fluid, the pressure of a gas or fluid, and many other types of parameters that comprise useful information for control purposes.

Figures 10A-D illustrate information processed and generated by an intelligent controller during control operations. All the figures show plots, similar to plot 902 in Figure 9, in which values of a parameter or another set of control-related values are plotted with respect to a vertical axis and time is plotted with respect to a horizontal axis. Figure 10A shows an idealized specification for the results of controlled-entity operation. The vertical axis 1002 in Figure 10A represents a specified parameter value, P_s . For

- 14 -

example, in the case of an intelligent thermostat, the specified parameter value may be temperature. For an irrigation system, by contrast, the specified parameter value may be flow rate. Figure 10A is the plot of a continuous curve 1004 that represents desired parameter values, over time, that an intelligent controller is directed to achieve through control of one or more devices, machines, or systems. The specification indicates that the parameter value is desired to be initially low 1006, then rise to a relatively high value 1008, then subside to an intermediate value 1010, and then again rise to a higher value 1012. A control specification can be visually displayed to a user, as one example, as a control schedule.

Figure 10B shows an alternate view, or an encoded-data view, of a control schedule corresponding to the control specification illustrated in Figure 10A. The control schedule includes indications of a parameter-value increase 1016 corresponding to edge 1018 in Figure 10A, a parameter-value decrease 1020 corresponding to edge 1022 in Figure 10A, and a parameter-value increase 1024 corresponding to edge 1016 in Figure 10A. The directional arrows plotted in Figure 10B can be considered to be setpoints, or indications of desired parameter changes at particular points in time within some period of time.

The control schedules learned by an intelligent controller represent a significant component of the results of automated learning. The learned control schedules may be encoded in various different ways and stored in electronic memories or mass-storage devices within the intelligent controller, within the system controlled by the intelligent controller, or within remote data-storage facilities, including cloud-computing-based data-storage facilities. In many cases, the learned control schedules may be encoded and stored in multiple locations, including control schedules distributed among internal intelligent-controller memory and remote data-storage facilities. A setpoint change may be stored as a record with multiple fields, including fields that indicate whether the setpoint change is a system-generated setpoint or a user-generated setpoint, whether the setpoint change is an immediate-control-input setpoint change or a scheduled setpoint change, the time and date of creation of the setpoint change, the time and date of the last edit of the setpoint change, and other such fields. In addition, a setpoint may be associated with two or more parameter values. As one example, a range setpoint may indicate a range of

- 15 -

parameter values within which the intelligent controller should maintain a controlled environment. Setpoint changes are often referred to as "setpoints."

Figure 10C illustrates the control output by an intelligent controller that might result from the control schedule illustrated in Figure 10B. In this figure, the magnitude of an output control signal is plotted with respect to the vertical axis 1026. For example, the control output may be a voltage signal output by an intelligent thermostat to a heating unit, with a high-voltage signal indicating that the heating unit should be currently operating and a low-voltage output indicating that the heating system should not be operating. Edge 1028 in Figure 10C corresponds to setpoint 1016 in Figure 10B. The width of the positive control output 1030 may be related to the length, or magnitude, of the desired parameter-value change, indicated by the length of setpoint arrow 1016. When the desired parameter value is obtained, the intelligent controller discontinues output of a high-voltage signal, as represented by edge 1032. Similar positive output control signals 1034 and 1036 are elicited by setpoints 1020 and 1024 in Figure 10B.

Finally, Figure 10D illustrates the observed parameter changes, as indicated by sensor output, resulting from control, by the intelligent controller, of one or more controlled entities. In Figure 10D, the sensor output, directly or indirectly related to the parameter P , is plotted with respect to the vertical axis 1040. The observed parameter value is represented by a smooth, continuous curve 1042. Although this continuous curve can be seen to be related to the initial specification curve, plotted in Figure 10A, the observed curve does not exactly match that specification curve. First, it may take a finite period of time 1044 for the controlled entity to achieve the parameter-valued change represented by setpoint 1016 in the control schedule plotted in Figure 10B. Also, once the parameter value is obtained, and the controlled entity directed to discontinue operation, the parameter value may begin to fall 1046, resulting in a feedback-initiated control output to resume operation of the controlled entity in order to maintain the desired parameter value. Thus, the desired high-level constant parameter value 1008 in Figure 10A may, in actuality, end up as a time-varying curve 1048 that does not exactly correspond to the control specification 1004. The first level of feedback, discussed above with reference to Figure 8, is used by the intelligent controller to control one or more control entities so that the observed parameter value, over

- 16 -

time, as illustrated in Figure 10D, matches the specified time behavior of the parameter in Figure 10A as closely as possible. The second level feedback control loop, discussed above with reference to Figure 8, may involve alteration of the specification, illustrated in Figure 10A, by a user, over time, either by changes to stored control schedules or by input of immediate-control directives, in order to generate a modified specification that produces a parameter-value/time curve reflective of a user's desired operational results.

There are many types of controlled entities and associated controllers. In certain cases, control output may include both an indication of whether the controlled entity should be currently operational as well as an indication of a level, throughput, or output of operation when the controlled entity is operational. In other cases, the control out may be simply a binary activation/deactivation signal. For simplicity of illustration and discussion, the latter type of control output is assumed in the following discussion.

Figures 11A-E provide a transition-state-diagram-based illustration of intelligent-controller operation. In these diagrams, the disk-shaped elements, or nodes, represent intelligent-controller states and the curved arrows interconnecting the nodes represent state transitions. Figure 11A shows one possible state-transition diagram for an intelligent controller. There are four main states 1102-1105. These states include: (1) a quiescent state 1102, in which feedback from sensors indicate that no controller outputs are currently needed and in which the one or more controlled entities are currently inactive or in maintenance mode; (2) an awakening state 1103, in which sensor data indicates that an output control may be needed to return one or more parameters to within a desired range, but the one or more controlled entities have not yet been activated by output control signals; (3) an active state 1104, in which the sensor data continue to indicate that observed parameters are outside desired ranges and in which the one or more controlled entities have been activated by control output and are operating to return the observed parameters to the specified ranges; and (4) an incipient quiescent state 1105, in which operation of the one or more controlled entities has returned the observed parameter to specified ranges but feedback from the sensors has not yet caused the intelligent controller to issue output control signals to the one or more controlled entities to deactivate the one or more controlled entities. In general, state transitions flow in a clockwise direction, with the

- 17 -

intelligent controller normally occupying the quiescent state 1102, but periodically awakening, in step 1103, due to feedback indications in order to activate the one or more controlled entities, in state 1104, to return observed parameters back to specified ranges. Once the observed parameters have returned to specified ranges, in step 1105, the intelligent controller issues deactivation output control signals to the one or more controlled entities, returning to the quiescent state 1102.

Each of the main-cycle states 1102-1105 is associated with two additional states: (1) a schedule-change state 1106-1109 and a control-change state 1110-1113. These states are replicated so that each main-cycle state is associated with its own pair of schedule-change and control-change states. This is because, in general, schedule-change and control-change states are transient states, from which the controller state returns either to the original main-cycle state from which the schedule-change or control-change state was reached by a previous transition or to a next main-cycle state in the above-described cycle. Furthermore, the schedule-change and control-change states are a type of parallel, asynchronously operating state associated with the main-cycle states. A schedule-change state represents interaction between the intelligent controller and a user or other remote entity carrying out control-schedule-creation, control-schedule-modification, or control-schedule-management operations through a displayed-schedule interface. The control-change states represent interaction of a user or other remote entity to the intelligent controller in which the user or other remote entity inputs immediate-control commands to the intelligent controller for translation into output control signals to the one or more controlled entities.

Figure 11B is the same state-transition diagram shown in Figure 11A, with the addition of circled, alphanumeric labels, such as circled, alphanumeric label 1116, associated with each transition. Figure 11C provides a key for these transition labels. Figures 11B-C thus together provide a detailed illustration of both the states and state transitions that together represent intelligent-controller operation.

To illustrate the level of detail contained in Figures 11B-C, consider the state transitions 1118-1120 associated with states 1102 and 1106. As can be determined from the table provided in Figure 11C, the transition 1118 from state 1102 to state 1106 involves

- 18 -

a control-schedule change made by either a user, a remote entry, or by the intelligent controller itself to one or more control schedules stored within, or accessible to, the intelligent controller. In general, following the schedule change, operation transitions back to state 1102 via transition 1119. However, in the relatively unlikely event that the schedule change has resulted in sensor data that was previously within specified ranges now falling outside newly specified ranges, the state transitions instead, via transition 1120, to the awakening state 1103.

Automated control-schedule learning by the intelligent controller, in fact, occurs largely as a result of intelligent-controller operation within the schedule-change and control-change states. Immediate-control inputs from users and other remote entities, resulting in transitions to the control-change states 1110-1113, provide information from which the intelligent controller learns, over time, how to control the one or more controlled entities in order to satisfy the desires and expectations of one or more users or remote entities. The learning process is encoded, by the intelligent controller, in control-schedule changes made by the intelligent controller while operating in the schedule-change states 1106-1109. These changes are based on recorded immediate-control inputs, recorded control-schedule changes, and current and historical control-schedule information. Additional sources of information for learning may include recorded output control signals and sensor inputs as well as various types of information gleaned from external sources, including sources accessible through the Internet. In addition to the previously described states, there is also an initial state or states 1130 that represent a first-power-on state or state following a reset of the intelligent controller. Generally, a boot operation followed by an initial-configuration operation or operations leads from the one or more initial states 1130, via transitions 1132 and 1134, to one of either the quiescent state 1102 or the awakening state 1103.

Figures 11D-E illustrate, using additional shading of the states in the state-transition diagram shown in Figure 11A, two modes of automated control-schedule learning carried out by an intelligent controller to which the current application is directed. The first mode, illustrated in Figure 11D, is a steady-state mode. The steady-state mode seeks optimal or near-optimal control with minimal immediate-control input. While learning

- 19 -

continues in the steady-state mode, the learning is implemented to respond relatively slowly and conservatively to immediate-control input, sensor input, and input from external information sources with the presumption that steady-state learning is primarily tailored to small-grain refinement of control operation and tracking of relatively slow changes in desired control regimes over time. In steady-state learning and general intelligent-controller operation, the most desirable state is the quiescent state 1102, shown crosshatched in Figure 11D to indicate this state as the goal, or most desired state, of steady-state operation. Light shading is used to indicate that the other main-cycle states 1103-1105 have neutral or slightly favored status in the steady-state mode of operation. Clearly, these states are needed for intermediate or continuous operation of controlled entities in order to maintain one or more parameters within specified ranges, and to track scheduled changes in those specified ranges. However, these states are slightly disfavored in that, in general, a minimal number, or minimal cumulative duration, of activation and deactivation cycles of the one or more controlled entities often leads to most optimal control regimes, and minimizing the cumulative time of activation of the one or more controlled entities often leads to optimizing the control regime with respect to energy and/or resource usage. In the steady-state mode of operation, the schedule-change and control-change states 1110-1113 are highly disfavored, because the intent of automated control-schedule learning is for the intelligent controller to, over time, devise one or more control schedules that accurately reflect a user's or other remote entity's desired operational behavior. While, at times, these states may be temporarily frequently inhabited as a result of changes in desired operational behavior, changes in environmental conditions, or changes in the controlled entities, a general goal of automated control-schedule learning is to minimize the frequency of both schedule changes and immediate-control inputs. Minimizing the frequency of immediate-control inputs is particularly desirable in many optimization schemes.

Figure 11E, in contrast to Figure 11D, illustrates an aggressive-learning mode in which the intelligent controller generally operates for a short period of time following transitions within the one or more initial states 730 to the main-cycle states 1102-1103. During the aggressive-learning mode, in contrast to steady-state operational mode

- 20 -

shown in Figure 11D, the quiescent state 1102 is least favored and the schedule-change and control-change states 1106-1113 are most favored, with states 1103-1105 having neutral desirability. In the aggressive-learning mode or phase of operation, the intelligent controller seeks frequent immediate-control inputs and schedule changes in order to quickly and aggressively acquire one or more initial control schedules. As discussed below, by using relatively rapid immediate-control-input relaxation strategies, the intelligent controller, while operating in aggressive-learning mode, seeks to compel a user or other remote entity to provide immediate-control inputs at relatively short intervals in order to quickly determine the overall shape and contour of an initial control schedule. Following completion of the initial aggressive learning and generation of adequate initial control schedules, relative desirability of the various states reverts to those illustrated in Figure 11D as the intelligent controller begins to refine control schedules and track longer-term changes in control specifications, the environment, the control system, and other such factors. Thus, the automated control-schedule-learning methods and intelligent controllers incorporating these methods to which the current application is directed feature an initial aggressive-learning mode that is followed, after a relatively short period of time, by a long-term, steady-state learning mode.

Figures 12 provide a state-transition diagram that illustrates automated control-schedule learning. Automated learning occurs during normal controller operation, illustrated in Figures 11A-C, and thus the state-transition diagram shown in Figure 12 describes operation behaviors of an intelligent controller that occur in parallel with the intelligent-controller operation described in Figures 11A-C. Following one or more initial states 1202, corresponding to the initial states 1130 in Figure 11B, the intelligent controller enters an initial-configuration learning state 1204 in which the intelligent controller attempts to create one or more initial control schedules based on one or more of default control schedules stored within the intelligent controller or accessible to the intelligent controller, an initial-schedule-creation dialog with a user or other remote entity through a schedule-creation interface, by a combination of these two approaches, or by additional approaches. The initial-configuration learning mode 1204 occurs in parallel with transitions 1132 and 1134 in Figure 11B. During the initial-learning mode, learning from

- 21 -

manually entered setpoint changes does not occur, as it has been found that users often make many such changes inadvertently, as they manipulate interface features to explore the controller's features and functionalities.

Following initial configuration, the intelligent controller transitions next to the aggressive-learning mode 1206, discussed above with reference to Figure 11E. The aggressive-learning mode 1206 is a learning-mode state which encompasses most or all states except for state 1130 of the states in Figure 11B. In other words, the aggressive-learning-mode state 1206 is a learning-mode state parallel to the general operational states discussed in Figures 11A-E. As discussed above, during aggressive learning, the intelligent controller attempts to create one or more control schedules that are at least minimally adequate to specify operational behavior of the intelligent controller and the entities which it controls based on frequent input from users or other remote entities. Once aggressive learning is completed, the intelligent controller transitions forward through a number of steady-state learning phases 1208-1210. Each transition downward, in the state-transition diagram shown in Figure 12, through the series of steady-state learning-phase states 1208-1210, is accompanied by changes in learning-mode parameters that result in generally slower, more conservative approaches to automated control-schedule learning as the one or more control schedules developed by the intelligent controller in previous learning states become increasingly accurate and reflective of user desires and specifications. The determination of whether or not aggressive learning is completed may be made based on a period of time, a number of information-processing cycles carried out by the intelligent controller, by determining whether the complexity of the current control schedule or schedules is sufficient to provide a basis for slower, steady-state learning, and/or on other considerations, rules, and thresholds. It should be noted that, in certain implementations, there may be multiple aggressive-learning states.

Figure 13 illustrates time frames associated with an example control schedule that includes shorter-time-frame sub-schedules. The control schedule is graphically represented as a plot with the horizontal axis 1302 representing time. The vertical axis 1303 generally represents one or more parameter values. As discussed further, below, a control schedule specifies desired parameter values as a function of time. The

- 22 -

control schedule may be a discrete set of values or a continuous curve. The specified parameter values are either directly or indirectly related to observable characteristics in environment, system, device, machine, or organization that can be measured by, or inferred from measurements obtained from, any of various types of sensors. In general, sensor output serves as at least one level of feedback control by which an intelligent controller adjusts the operational behavior of a device, machine, system, or organization in order to bring observed parameter values in line with the parameter values specified in a control schedule. The control schedule used as an example in the following discussion is incremented in hours, along the horizontal axis, and covers a time span of one week. The control schedule includes seven sub-schedules 1304-1310 that correspond to days. As discussed further below, in an example intelligent controller, automated control-schedule learning takes place at daily intervals, with a goal of producing a robust weekly control schedule that can be applied cyclically, week after week, over relatively long periods of time. As also discussed below, an intelligent controller may learn even longer-period control schedules, such as yearly control schedules, with monthly, weekly, daily, and even hourly sub-schedules organized hierarchically below the yearly control schedule. In certain cases, an intelligent controller may generate and maintain shorter-time-frame control schedules, including hourly control schedules, minute-based control schedules, or even control schedules incremented in milliseconds and microseconds. Control schedules are, like the stored computer instructions that together compose control routines, tangible, physical components of control systems. Control schedules are stored as physical states in physical storage media. Like control routines and programs, control schedules are necessarily tangible, physical control-system components that can be accessed and used by processor-based control logic and control systems.

Figures 14A-C show three different types of control schedules. In Figure 14A, the control schedule is a continuous curve 1402 representing a parameter value, plotted with respect to the vertical axis 1404, as a function of time, plotted with respect to the horizontal axis 1406. The continuous curve comprises only horizontal and vertical sections. Horizontal sections represent periods of time at which the parameter is desired to remain constant and vertical sections represent desired changes in the parameter value at

- 23 -

particular points in time. This is a simple type of control schedule and is used, below, in various examples of automated control-schedule learning. However, automated control-schedule-learning methods can also learn more complex types of schedules. For example, Figure 14B shows a control schedule that includes not only horizontal and vertical segments, but arbitrarily angled straight-line segments. Thus, a change in the parameter value may be specified, by such a control schedule, to occur at a given rate, rather than specified to occur instantaneously, as in the simple control schedule shown in Figure 14A. Automated-control-schedule-learning methods may also accommodate smooth-continuous-curve-based control schedules, such as that shown in Figure 14C. In general, the characterization and data encoding of smooth, continuous-curve-based control schedules, such as that shown in Figure 14C, is more complex and includes a greater amount of stored data than the simpler control schedules shown in Figures 14B and 14A.

In the following discussion, it is generally assumed that a parameter value tends to relax towards lower values in the absence of system operation, such as when the parameter value is temperature and the controlled system is a heating unit. However, in other cases, the parameter value may relax toward higher values in the absence of system operation, such as when the parameter value is temperature and the controlled system is an air conditioner. The direction of relaxation often corresponds to the direction of lower resource or expenditure by the system. In still other cases, the direction of relaxation may depend on the environment or other external conditions, such as when the parameter value is temperature and the controlled system is an HVAC system including both heating and cooling functionality.

Turning to the control schedule shown in Figure 14A, the continuous-curve-represented control schedule 1402 may be alternatively encoded as discrete setpoints corresponding to vertical segments, or edges, in the continuous curve. A continuous-curve control schedule is generally used, in the following discussion, to represent a stored control schedule either created by a user or remote entity via a schedule-creation interface provided by the intelligent controller or created by the intelligent controller based on already-existing control schedules, recorded immediate-control inputs, and/or recorded sensor data, or a combination of these types of information.

- 24 -

Immediate-control inputs are also graphically represented in parameter-value versus time plots. Figures 15A-G show representations of immediate-control inputs that may be received and executed by an intelligent controller, and then recorded and overlaid onto control schedules, such as those discussed above with reference to Figures 14A-C, as part of automated control-schedule learning. An immediate-control input is represented graphically by a vertical line segment that ends in a small filled or shaded disk. Figure 15A shows representations of two immediate-control inputs 1502 and 1504. An immediate-control input is essentially equivalent to an edge in a control schedule, such as that shown in Figure 14A, that is input to an intelligent controller by a user or remote entity with the expectation that the input control will be immediately carried out by the intelligent controller, overriding any current control schedule specifying intelligent-controller operation. An immediate-control input is therefore a real-time setpoint input through a control-input interface to the intelligent controller.

Because an immediate-control input alters the current control schedule, an immediate-control input is generally associated with a subsequent, temporary control schedule, shown in Figure 15A as dashed horizontal and vertical lines that form a temporary-control-schedule parameter vs. time curve extending forward in time from the immediate-control input. Temporary control schedules 1506 and 1508 are associated with immediate-control inputs 1502 and 1504, respectively, in Figure 15A.

Figure 15B illustrates an example of immediate-control input and associated temporary control schedule. The immediate-control input 1510 is essentially an input setpoint that overrides the current control schedule and directs the intelligent controller to control one or more controlled entities in order to achieve a parameter value equal to the vertical coordinate of the filled disk 1512 in the representation of the immediate-control input. Following the immediate-control input, a temporary constant-temperature control-schedule interval 1514 extends for a period of time following the immediate-control input, and the immediate-control input is then relaxed by a subsequent immediate-control-input endpoint, or subsequent setpoint 1516. The length of time for which the immediate-control input is maintained, in interval 1514, is a parameter of automated control-schedule learning. The direction and magnitude of the subsequent immediate-control-input endpoint setpoint

- 25 -

1516 represents one or more additional automated-control-schedule-learning parameters. Please note that an automated-control-schedule-learning parameter is an adjustable parameter that controls operation of automated control-schedule learning, and is different from the one or more parameter values plotted with respect to time that comprise control schedules. The parameter values plotted with respect to the vertical axis in the example control schedules to which the current discussion refers are related directly or indirectly to observables, including environmental conditions, machines states, and the like.

Figure 15C shows an existing control schedule on which an immediate-control input is superimposed. The existing control schedule called for an increase in the parameter value P , represented by edge 1520, at 7:00 a.m. (1522 in Figure 15C). The immediate-control input 1524 specifies an earlier parameter-value change of somewhat less magnitude. Figures 15D-G illustrate various subsequent temporary control schedules that may obtain, depending on various different implementations of intelligent-controller logic and/or current values of automated-control-schedule-learning parameter values. In Figures 15D-G, the temporary control schedule associated with an immediate-control input is shown with dashed line segments and that portion of the existing control schedule overridden by the immediate-control input is shown by dotted line segments. In one approach, shown in Figure 15D, the desired parameter value indicated by the immediate-control input 1524 is maintained for a fixed period of time 1526 after which the temporary control schedule relaxes, as represented by edge 1528, to the parameter value that was specified by the control schedule at the point in time that the immediate-control input is carried out. This parameter value is maintained 1530 until the next scheduled setpoint, which corresponds to edge 1532 in Figure 15C, at which point the intelligent controller resumes control according to the control schedule.

In an alternative approach shown in Figure 15E, the parameter value specified by the immediate-control input 1524 is maintained 1532 until a next scheduled setpoint is reached, in this case the setpoint corresponding to edge 1520 in the control schedule shown in Figure 15C. At the next setpoint, the intelligent controller resumes control according to the existing control schedule. This approach is often desirable,

- 26 -

because users often expect a manually entered setpoint to remain in force until a next scheduled setpoint change.

In a different approach, shown in Figure 15F, the parameter value specified by the immediate-control input 1524 is maintained by the intelligent controller for a fixed period of time 1534, following which the parameter value that would have been specified by the existing control schedule at that point in time is resumed 1536.

In the approach shown in Figure 15G, the parameter value specified by the immediate-control input 1524 is maintained 1538 until a setpoint with opposite direction from the immediate-control input is reached, at which the existing control schedule is resumed 1540. In still alternative approaches, the immediate-control input may be relaxed further, to a lowest-reasonable level, in order to attempt to optimize system operation with respect to resource and/or energy expenditure. In these approaches, generally used during aggressive learning, a user is compelled to positively select parameter values greater than, or less than, a parameter value associated with a minimal or low rate of energy or resource usage.

In one example implementation of automated control-schedule learning, an intelligent controller monitors immediate-control inputs and schedule changes over the course of a monitoring period, generally coinciding with the time span of a control schedule or sub-schedule, while controlling one or more entities according to an existing control schedule except as overridden by immediate-control inputs and input schedule changes. At the end of the monitoring period, the recorded data is superimposed over the existing control schedule and a new provisional schedule is generated by combining features of the existing control schedule and schedule changes and immediate-control inputs. Following various types of resolution, the new provisional schedule is promoted to the existing control schedule for future time intervals for which the existing control schedule is intended to control system operation.

Figures 16A-E illustrate one aspect of the method by which a new control schedule is synthesized from an existing control schedule and recorded schedule changes and immediate-control inputs. Figure 16A shows the existing control schedule for a monitoring period. Figure 16B shows a number of recorded immediate-control inputs

- 27 -

superimposed over the control schedule following the monitoring period. As illustrated in Figure 16B, there are six immediate-control inputs 1602-1607. In a clustering technique, clusters of existing-control-schedule setpoints and immediate-control inputs are detected. One approach to cluster detection is to determine all time intervals greater than a threshold length during which neither existing-control-schedule setpoints nor immediate-control inputs are present, as shown in Figure 16C. The horizontal, double-headed arrows below the plot, such as double-headed arrow 1610, represent the intervals of greater than the threshold length during which neither existing-control-schedule setpoints nor immediate-control inputs are present in the superposition of the immediate-control inputs onto the existing control schedule. Those portions of the time axis not overlapping by these intervals are then considered to be clusters of existing-control-schedule setpoints and immediate-control inputs, as shown in Figure 16D. A first cluster 1612 encompasses existing-control-schedule setpoints 1614-1616 and immediate-control inputs 1602 and 1603. A second cluster 1620 encompasses immediate-control inputs 1604 and 1605. A third cluster 1622 encompasses only existing-control-schedule setpoint 1624. A fourth cluster 1626 encompasses immediate-control inputs 1606 and 1607 as well as the existing-control-schedule setpoint 1628. In one cluster-processing method, each cluster is reduced to zero, one, or two setpoints in a new provisional schedule generated from the recorded immediate-control inputs and existing control schedule. Figure 16E shows an exemplary new provisional schedule 1630 obtained by resolution of the four clusters identified in Figure 16D.

Cluster processing is intended to simplify the new provisional schedule by coalescing the various existing-control-schedule setpoints and immediate-control inputs within a cluster to zero, one, or two new-control-schedule setpoints that reflect an apparent intent on behalf of a user or remote entity with respect to the existing control schedule and the immediate-control inputs. It would be possible, by contrast, to generate the new provisional schedule as the sum of the existing-control-schedule setpoints and immediate-control inputs. However, that approach would often lead to a ragged, highly variable, and fine-grained control schedule that generally does not reflect the ultimate desires of users or other remote entities and which often constitutes a parameter-value vs. time curve that

- 28 -

cannot be achieved by intelligent control. As one example, in an intelligent thermostat, two setpoints 15 minutes apart specifying temperatures that differ by ten degrees may not be achievable by an HVAC system controlled by an intelligent controller. It may be the case, for example, that under certain environmental conditions, the HVAC system is capable of raising the internal temperature of a residence by a maximum of only five degrees per hour. Furthermore, simple control schedules can lead to a more diverse set of optimization strategies that can be employed by an intelligent controller to control one or more entities to produce parameter values, or P values, over time, consistent with the control schedule. An intelligent controller can then optimize the control in view of further constraints, such as minimizing energy usage or resource utilization.

There are many possible approaches to resolving a cluster of existing-control-schedule setpoints and immediate-control inputs into one or two new provisional schedule setpoints. Figures 17A-E illustrate one approach to resolving schedule clusters. In each of Figures 17A-E, three plots are shown. The first plot shows recorded immediate-control inputs superimposed over an existing control schedule. The second plot reduces the different types of setpoints to a single generic type of equivalent setpoints, and the final plot shows resolution of the setpoints into zero, one, or two new provisional schedule setpoints.

Figure 17A shows a cluster 1702 that exhibits an obvious increasing P -value trend, as can be seen when the existing-control-schedule setpoints and immediate-control inputs are plotted together as a single type of setpoint, or event, with directional and magnitude indications with respect to actual control produced from the existing-control-schedule setpoints and immediate-control inputs 704 within an intelligent controller. In this case, four out of the six setpoints 706-709 resulted in an increase in specified P value, with only a single setpoint 710 resulting in a slight decrease in P value and one setpoint 712 produced no change in P value. In this and similar cases, all of the setpoints are replaced by a single setpoint specifying an increase in P value, which can be legitimately inferred as the intent expressed both in the existing control schedule and in the immediate-control inputs. In this case, the single setpoint 716 that replaces the cluster of setpoints 704 is placed at the time of the first setpoint in the cluster and specifies a new P value equal to the highest P value specified by any setpoint in the cluster.

- 29 -

The cluster illustrated in Figure 17B contains five setpoints 718-722. Two of these setpoints specify a decrease in P value, two specify an increase in P value, and one had no effect. As a result, there is no clear P -value-change intent demonstrated by the collection of setpoints, and therefore the new provisional schedule 724 contains no setpoints over the cluster interval, with the P value maintained at the initial P value of the existing control schedule within the cluster interval.

Figure 17C shows a cluster exhibiting a clear downward trend, analogous to the upward trend exhibited by the clustered setpoints shown in Figure 17A. In this case, the four cluster setpoints are replaced by a single new provisional schedule setpoint 726 at a point in time corresponding to the first setpoint in the cluster and specifying a decrease in P value equivalent to the lowest P value specified by any of the setpoints in the cluster.

In Figure 17D, the cluster includes three setpoints 730-732. The setpoint corresponding to the existing-control-schedule setpoint 730 and a subsequent immediate-control setpoint 731 indicate a clear intent to raise the P value at the beginning of the cluster interval and the final setpoint 732 indicates a clear intent to lower the P value at the end of the cluster interval. In this case, the three setpoints are replaced by two setpoints 734 and 736 in the new provisional schedule that mirror the intent inferred from the three setpoints in the cluster. Figure 17E shows a similar situation in which three setpoints in the cluster are replaced by two new-provisional-schedule setpoints 738 and 740, in this case representing a temporary lowering and then subsequent raising of the P value as opposed to the temporary raising and subsequent lowering of the P value in the new provisional schedule in Figure 17B.

There are many different computational methods that can recognize the trends of clustered setpoints discussed with reference to Figures 17A-E. These trends provide an example of various types of trends that may be computationally recognized. Different methods and strategies for cluster resolution are possible, including averaging, curve fitting, and other techniques. In all cases, the goal of cluster resolution is to resolve multiple setpoints into a simplest possible set of setpoints that reflect a user's intent, as judged from the existing control schedule and the immediate-control inputs.

- 30 -

Figures 18A-B illustrate the effect of a prospective schedule change entered by a user during a monitoring period. In Figures 18A-B, and in subsequent figures, a schedule-change input by a user is represented by a vertical line 1802 ending in a small filled disk 1804 indicating a specified P value. The setpoint is placed with respect to the horizontal axis at a time at which the setpoint is scheduled to be carried out. A short vertical line segment 1806 represents the point in time that the schedule change was made by a user or remote entity, and a horizontal line segment 1808 connects the time of entry with the time for execution of the setpoint represented by vertical line segments 1806 and 1802, respectively. In the case shown in Figure 18A, a user altered the existing control schedule, at 7:00 a.m. 1810, to include setpoint 1802 at 11:00 a.m. In cases such as those shown in Figure 18A, where the schedule change is prospective and where the intelligent controller can control one or more entities according to the changed control schedule within the same monitoring period, the intelligent controller simply changes the control schedule, as indicated in Figure 18B, to reflect the schedule change. In one automated-control-schedule-learning method, therefore, prospective schedule changes are not recorded. Instead, the existing control schedule is altered to reflect a user's or remote entity's desired schedule change.

Figures 19A-B illustrate the effect of a retrospective schedule change entered by a user during a monitoring period. In the case shown in Figure 19A, a user input three changes to the existing control schedule at 6:00 p.m. 1902, including deleting an existing setpoint 1904 and adding two new setpoints 1906 and 1908. All of these schedule changes would impact only a future monitoring period controlled by the modified control schedule, since the time at which they were entered is later than the time at which the changes in P value are scheduled to occur. For these types of schedule changes, the intelligent controller records the schedule changes in a fashion similar to the recording of immediate-control inputs, including indications of the fact that this type of setpoint represents a schedule change made by a user through a schedule-modification interface rather than an immediate-control input.

Figure 19B shows a new provisional schedule that incorporates the schedule changes shown in Figure 19A. In general, schedule changes are given relatively large

- 31 -

deference by the currently described automated-control-schedule-learning method. Because a user has taken the time and trouble to make schedule changes through a schedule-change interface, it is assumed that the schedule changes are strongly reflective of the user's desires and intentions. As a result, as shown in Figure 19B, the deletion of existing setpoint 1904 and the addition of the two new setpoints 1906 and 1908 are entered into the existing control schedule to produce the new provisional schedule 1910. Edge 1912 corresponds to the schedule change represented by setpoint 1906 in Figure 19A and edge 1914 corresponds to the schedule change represented by setpoint 1908 in Figure 19A. In summary, either for prospective schedule changes or retrospective schedule changes made during a monitoring period, the schedule changes are given great deference during learning-based preparation of a new provisional schedule that incorporates both the existing control schedule and recorded immediate-control inputs and schedule changes made during the monitoring period.

Figures 20A-C illustrate overlay of recorded data onto an existing control schedule, following completion of a monitoring period, followed by clustering and resolution of clusters. As shown in Figure 20A, a user has input six immediate-control inputs 2004-2009 and two retrospective schedule changes 2010 and 2012 during the monitoring period, which are overlain or superimposed on the existing control schedule 2002,. As shown in Figure 20B, clustering produces four clusters 2014-2017. Figure 20C shows the new provisional schedule obtained by resolution of the clusters. Cluster 2014, with three existing-control-schedule setpoints and two immediate-control setpoints, is resolved to new-provisional-schedule setpoints 2020 and 2022. Cluster 2 (2015 in Figure 20B), containing two immediate-control setpoints and two retrospective-schedule setpoints, is resolved to setpoints 2024 and 2026. Cluster 3 (2016 in Figure 20B) is resolved to the existing-control-schedule setpoint 2028, and cluster 4 (2017 in Figure 20B), containing two immediate-control setpoints and an existing-control-schedule setpoint, is resolved to setpoint 2030. In preparation for a subsequent schedule-propagation step, each of the new-provisional-schedule setpoints is labeled with an indication of whether or not the setpoint parameter value is derived from an immediate-control setpoint or from either an existing-control-schedule setpoint or retrospective schedule-change setpoint. The latter two

- 32 -

categories are considered identical, and setpoints of those categories are labeled with the character "s" in Figure 20C, while the setpoints with temperatures derived from immediate-control setpoints, 2020 and 2022, are labeled "i." As discussed further, below, only setpoints labeled "i" are propagated to additional, related sub-schedules of a higher-level control schedule.

An additional step that may follow clustering and cluster resolution and precede new-provisional-schedule propagation, in certain implementations, involves spreading apart setpoints derived from immediate-control setpoints in the new provisional schedule. Figures 21A-B illustrate the setpoint-spreading operation. Figure 21A shows a new provisional schedule with setpoints labeled, as discussed above with reference to Figure 20C, with either "s" or "i" in order to indicate the class of setpoints from which the setpoints were derived. In this new provisional schedule 2102, two setpoints labeled "i" 2104 and 2106 are separated by a time interval 2108 of length less than a threshold time interval for separation purposes. The spreading operation detects pairs of "i" labeled setpoints that are separated, in time, by less than the threshold time interval and moves the latter setpoint of the pair forward, in time, so that the pair of setpoints are separated by at least a predetermined fixed-length time interval 2110 in Figure 21B. In a slightly more complex spreading operation, in the case that the latter setpoint of the pair would be moved closer than the threshold time to a subsequent setpoint, the latter setpoint may be moved to a point in time halfway between the first setpoint of the pair and the subsequent setpoint. The intent of the spreading operation is to ensure adequate separation between setpoints for schedule simplicity and in order to produce a control schedule that can be realized under intelligent-controller control of a system.

A next operation carried out by the currently discussed automated-control-schedule-learning method is propagation of a new provisional sub-schedule, created, as discussed above, following a monitoring period, to related sub-schedules in a higher-level control schedule. Schedule propagation is illustrated in Figures 22A-B. Figure 22A shows a higher-level control schedule 2202 that spans a week in time and that includes daily sub-schedules, such as the Saturday sub-schedule 2204. In Figure 22A, the Monday sub-schedule 2206 has recently been replaced by a new provisional Monday sub-schedule

- 33 -

following the end of a monitoring period, indicated in Figure 22A by crosshatching oppositely slanted from the crosshatching of the sub-schedules corresponding to other days of the week. As shown in Figure 22B, the schedule-propagation technique used in the currently discussed automated-control-schedule-learning method involves propagating the new provisional Monday sub-schedule 2206 to other, related sub-schedules 2208-2211 in the higher-level control schedule 2202. In this case, weekday sub-schedules are considered to be related to one another, as are weekend sub-schedules, but weekend sub-schedules are not considered to be related to weekday sub-schedules. Sub-schedule propagation involves overlaying the "i"-labeled setpoints in the new provisional schedule 2206 over related existing control schedules, in this case sub-schedules 2208-2211, and then resolving the setpoint-overlaid existing control schedules to produce new provisional sub-schedules for the related sub-schedules. In Figure 22B, overlaying of "i"-labeled setpoints from new provisional sub-schedule 2206 onto the related sub-schedules 2208-2211 is indicated by bi-directional crosshatching. Following resolution of these overlaid setpoints and existing sub-schedules, the entire higher-level control schedule 2202 is then considered to be the current existing control schedule for the intelligent controller. In other words, following resolution, the new provisional sub-schedules are promoted to existing sub-schedules. In certain cases, the sub-schedule propagation rules may change, over time. As one example, propagation may occur to all days, initially, of a weekly schedule but may then more selectively propagate weekday sub-schedules to weekdays and week-end-day sub-schedules to week-end-days. Other such rules may be employed for propagation of sub-schedules.

As discussed above, there can be multiple hierarchical layers of control schedules and sub-schedules maintained by an intelligent controller, as well as multiple sets of hierarchically related control schedules. In these cases, schedule propagation may involve relatively more complex propagation rules for determining to which sub-schedules a newly created provisional sub-schedule should be propagated. Although propagation is shown, in Figure 22B, in the forward direction in time, propagation of a new provisional schedule or new provisional sub-schedule may be carried out in either a forward or reverse direction with respect to time. In general, new-provisional-schedule propagation is

governed by rules or by tables listing those control schedules and sub-schedules considered to be related to each control schedule and/or sub-schedule.

Figures 23A-C illustrate new-provisional-schedule propagation using P -value vs. t control-schedule plots. Figure 23A shows an existing control schedule 2302 to which the "i"-labeled setpoints in a new provisional schedule are propagated. Figure 23B shows the propagated setpoints with "i" labels overlaid onto the control schedule shown in Figure 23A. Two setpoints 2304 and 2306 are overlaid onto the existing control schedule 2302. The existing control schedule includes four existing setpoints 2308-2311. The second of the propagated setpoints 2306 lowers the parameter value to a level 2312 greater than the corresponding parameter-value level 2314 of the existing control schedule 2302, and therefore overrides the existing control schedule up to existing setpoint 2310. In this simple case, no further adjustments are made, and the propagated setpoints are incorporated in the existing control schedule to produce a new provisional schedule 2316 shown in Figure 23C. When setpoints have been propagated to all related control schedules or sub-schedules, and new provisional schedules and sub-schedules are generated for them, the propagation step terminates, and all of the new provisional schedules and sub-schedules are together considered to be a new existing higher-level control schedule for the intelligent controller.

Following propagation and overlaying of "i"-labeled setpoints onto a new provisional schedule to a related sub-schedule or control schedule, as shown in Figure 23B, numerous rules may be applied to the overlying setpoints and existing control schedule in order to simplify and to make realizable the new provisional schedule generated from the propagated setpoints and existing control schedule. Figures 24A-I illustrate a number of example rules used to simplify a existing control schedule overlaid with propagated setpoints as part of the process of generating a new provisional schedule. Each of Figures 24A-I include two P -value vs. t plots, the first showing a propagated setpoint overlying a existing control schedule and the second showing resolution of the propagated setpoint to generate a portion of a new provisional schedule obtained by resolving a propagated setpoint.

- 35 -

The first, left-hand P -value vs. t plot 2402 in Figure 24A shows a propagated setpoint 2404 overlying an existing control schedule 2405. Figure 24A also illustrates terminology used in describing many of the example rules used to resolve propagated setpoints with existing control schedules. In Figure 24A a first existing setpoint, pe_1 2406, precedes the propagated setpoint 2404 in time by a length of time a 2407 and a second existing setpoint of the existing control schedule, pe_2 2408, follows the propagated setpoint 2404 in time by a length of time b 2409. The P -value difference between the first existing-control-schedule setpoint 2406 and the propagated setpoint 2404 is referred to as " ΔP " 2410. The right-hand P -value vs. t plot 2412 shown in Figure 24A illustrates a first propagated-setpoint-resolution rule. As shown in this Figure, when ΔP is less than a threshold ΔP and b is less than a threshold Δt , then the propagated setpoint is deleted. Thus, resolution of the propagated setpoint with the existing control schedule, by rule 1, removes the propagated setpoint, as shown in the right-hand side of Figure 24A.

Figures 24B-I illustrate an additional example of propagated-setpoint-resolution rules in similar fashion to the illustration of the first propagated-setpoint-resolution rule in Figure 24A. Figure 24B illustrates a rule that, when b is less than a threshold Δt and when the first rule illustrated in Figure 24A does not apply, then the new propagated setpoint 2414 is moved ahead in time by a value Δt_2 2416 from existing setpoint pe_1 and existing setpoint pe_2 is deleted.

Figure 24C illustrates a third rule applied when neither of the first two rules are applicable to a propagated setpoint. If a is less than a threshold value Δt , then the propagated setpoint is moved back in time by a predetermined value Δt_3 from pe_2 and the existing setpoint pe_1 is deleted.

Figure 24D illustrates a fourth row applicable when none of the first three rules can be applied to a propagated setpoint. In this case, the P value of the propagated setpoint becomes the P value for the existing setpoint pe_1 and the propagated setpoint is deleted.

When none of the first four rules, described above with reference to Figures 24A-D, are applicable, then additional rules may be tried in order to resolve a propagated setpoint with an existing control schedule. Figure 24E illustrates a fifth rule. When b is

- 36 -

less than a threshold Δt and ΔP is less than a threshold Δp , then, as shown in Figure 24E, the propagated setpoint 2424 is deleted. In other words, a propagated setpoint too close to an existing control schedule setpoint is not incorporated into the new provisional control schedule. The existing setpoints may also be reconsidered, during propagated-setpoint resolution. For example, as shown in Figure 24F, when a second existing setpoint pe_2 that occurs after a first existing setpoint pe_1 results in a change in the parameter value ΔP less than a threshold ΔP , then the second existing setpoint pe_2 may be removed. Such proximal existing setpoints may arise due to the deference given to schedule changes following previous monitoring periods. Similarly, as shown in Figure 24G, when a propagated setpoint follows an existing setpoint, and the change in the parameter value ΔP produced by the propagated setpoint is less than a threshold ΔP value, then the propagated setpoint is deleted. As shown in Figure 24H, two existing setpoints that are separated by less than a threshold Δt value may be resolved into a single setpoint coincident with the first of the two existing setpoints. Finally, in similar fashion, a propagated setpoint that is too close, in time, to an existing setpoint may be deleted.

In certain implementations, a significant distinction is made between user-entered setpoint changes and automatically generated setpoint changes. The former setpoint changes are referred to as "anchor setpoints," and are not overridden by learning. In many cases, users expect that the setpoints which they manually enter should not be changed. Additional rules, heuristics, and consideration can be used to differentiate setpoint changes for various levels of automated adjustment during both aggressive and steady-state learning. It should also be noted that setpoints associated with two parameter values that indicate a parameter-value range may be treated in different ways during comparison operations used in pattern matching and other automated learning calculations and determinations. For example, a range setpoint change may need to match another range setpoint change in both parameters to be deemed to be equivalent or identical.

Next, an example implementation of an intelligent controller that incorporates the above-described automated-control-schedule-learning method is provided. Figures 25A-M illustrate an example implementation of an intelligent controller that incorporates the above-described automated-control-schedule-learning method. At the

- 37 -

onset, it should be noted that the following implementation is but one of many different possible implementations that can be obtained by varying any of many different design and implementation parameters, including modular organization, control structures, data structures, programming language, hardware components, firmware, and other such design and implementation parameters. Many different types of control schedules may be used by different types of intelligent controllers applied to different control domains. Automated-control-schedule-learning methods incorporated into intelligent-controller logic may significantly vary depending on the types and numbers of control schedules that specify intelligent-controller operation. The time periods spanned by various different types of control schedules and the granularity, in time, of control schedules may vary widely depending on the control tasks for which particular controllers are designed.

Figure 25A shows the highest-level intelligent-controller control logic. This high-level control logic comprises an event-handling loop in which various types of control-related events are handled by the intelligent controller. In Figure 25A, four specific types of control-related events are handled, but, in general, the event-handling loop may handle many additional types of control-related events that occur at lower levels within the intelligent-controller logic. Examples include communications events, in which the intelligent controller receives or transmits data to remote entities, such as remote smart-home devices and cloud-computing servers. Other types of control-related events include control-related events related to system activation and deactivation according to observed parameters and control schedules, various types of alarms and timers that may be triggered by sensor data falling outside of control-schedule-specified ranges for detection and unusual or rare events that require specialized handling. Rather than attempt to describe all the various different types of control-related events that may be handled by an intelligent controller, Figure 25A illustrates handling of four example control-related events.

In step 2502, the intelligent controller waits for a next control-related event to occur. When a control-related event occurs, control flows to step 2504, and the intelligent controller determines whether an immediate-control input has been input by a user or remote entity through the immediate-control-input interface. When an immediate-control input has been input by a user or other remote entity, as determined in step 2504,

- 38 -

the intelligent controller carries out the immediate control input, in step 2505, generally by changing internally stored specified ranges for parameter values and, when needed, activating one or more controlled entities, and then the immediate-control input is recorded in memory, in step 2506. When an additional setpoint or other schedule feature needs to be added to terminate the immediate-control input, as determined in step 2507, then the additional setpoint or other schedule feature is added to the control schedule, in step 2508. Examples of such added setpoints are discussed above with reference to Figures 15A-G. When the control-related event that triggered exit from step 2502 is a timer event indicating that the current time is that of a scheduled setpoint or scheduled control, as determined in step 2509, then the intelligent controller carries out the scheduled controller setpoint in step 2510. When the scheduled control carried out in step 2510 is a temporary scheduled control added in step 2508 to terminate an immediate-control input, as determined in step 2511, then the temporary scheduled control is deleted in step 2512. When the control-related event that triggered exit from step 2502 is a change made by a user or remote entity to the control schedule via the control-schedule-change interface, as determined in step 2513, then, when the schedule change is prospective, as determined in step 1514, the schedule change is made by the intelligent controller to the existing control schedule in step 2515, as discussed above with reference to Figures 18A-B. Otherwise, the schedule change is retrospective, and is recorded by the intelligent controller into memory in step 2516 for later use in varying a new provisional schedule at the termination of the current monitoring period.

When the control-related event that triggered exit from 2502 is a timer event associated with the end of the current monitoring period, as determined in step 2517, then a monitoring-period routine is called, in step 2518, to process recorded immediate-control inputs and schedule changes, as discussed above with reference to Figures 15A-24F. When additional control-related events occur after exit from step 2502, which are generally queued to an occurred event queue, as determined in step 2519, control flows back to step 2504 for handling a next queued event. Otherwise, control flows back in step 2502 where the intelligent controller waits for a next control-related event.

- 39 -

Figure 25B provides a control-flow diagram for the routine "monitoring period" called in step 2518 in Figure 25A. In step 2522, the intelligent controller accesses a state variable that stores an indication of the current learning mode. When the current learning mode is an aggressive-learning mode, as determined in step 2523, the routine "aggressive monitoring period" is called in step 2524. Otherwise, the routine "steady-state monitoring period" is called, in step 2525. While this control-flow diagram is simple, it clearly shows the feature of automated-control-schedule-learning discussed above with reference to Figures 11D-E and Figure 12. Automated-control-schedule learning is bifurcated into an initial, aggressive-learning period followed by a subsequent steady-state learning period.

Figure 25C provides a control-flow diagram for the routine "aggressive monitoring period" called in step 2524 of Figure 25B. This routine is called at the end of each monitoring period. In the example discussed above, a monitoring period terminates at the end of each daily control schedule, immediately after 12:00 p.m. However, monitoring periods, in alternative implementations, may occur at a variety of other different time intervals and may even occur variably, depending on other characteristics and parameters. Monitoring periods are generally the smallest-granularity time periods corresponding to control schedules or sub-schedules, as discussed above.

In step 2527, the intelligent controller combines all recorded immediate-control inputs with the existing control schedule, as discussed above with reference to Figures 16B and 20A. In step 2528, the routine "cluster" is called in order to partition the recorded immediate-control inputs and schedule changes and existing-control-schedule setpoints to clusters, as discussed above with reference to Figures 16C-D and Figure 20B. In step 2529, the intelligent controller calls the routine "simplify clusters" to resolve the various setpoints within each cluster, as discussed above with reference to Figures 16A-20C. In step 2530, the intelligent controller calls the routine "generate new schedule" to generate a new provisional schedule following cluster resolution, as discussed above with reference to Figures 20C and 21A-B. In step 2531, the intelligent controller calls the routine "propagateNewSchedule" discussed above with reference to Figures 22A-24I, in order to propagate features of the provisional schedule generated in step 2530 to related

- 40 -

sub-schedules and control schedules of the intelligent controller's control schedule. In step 2532, the intelligent controller determines whether or not the currently completed monitoring period is the final monitoring period in the aggressive-learning mode. When the recently completed monitoring period is the final monitoring period in the aggressive-monitoring learning mode, as determined in step 2532, then, in step 2533, the intelligent controller sets various state variables that control the current learning mode to indicate that the intelligent controller is now operating in the steady-state learning mode and, in step 2534, sets various learning parameters to parameter values compatible with phase I of steady-state learning.

Many different learning parameters may be used in different implementations of automated control-schedule learning. In the currently discussed implementation, learning parameters may include the amount of time that immediate-control inputs are carried out before termination by the intelligent controller, the magnitudes of the various threshold Δt and threshold ΔP values used in cluster resolution and resolution of propagated setpoints with respect to existing control schedules. Finally, in step 2535, the recorded immediate-control inputs and schedule changes, as well as clustering information and other temporary information derived and stored during creation of a new provisional schedule and propagation of the provisional schedule are deleted and the learning logic is reinitialized to begin a subsequent monitoring period.

Figure 25D provides a control-flow diagram for the routine "cluster" called in step 2528 of Figure 25C. In step 2537, the local variable Δt_{int} is set to a learning-mode and learning-phase-dependent value. Then, in the *while*-loop of steps 2538-2542, the routine "interval cluster" is repeatedly called in order to generate clusters within the existing control schedule until one or more clustering criteria are satisfied, as determined in step 2540. Prior to satisfaction of the clustering criteria, the value of Δt_{int} is incremented, in step 2542, prior to each next call to the routine "interval cluster" in step 2539, in order to alter the next clustering for satisfying the clustering criteria. The variable Δt_{int} corresponds to the minimum length of time between setpoints that results in the setpoints being classified as belonging to two different clusters, as discussed above with reference to Figure 16C, or

- 41 -

the time period 1610 the interval between two clusters. Decreasing Δt_{int} generally produces additional clusters.

Various different types of clustering criteria may be used by an intelligent controller. In general, it is desirable to generate a sufficient number of clusters to produce adequate control-schedule simplification, but too many clusters result in additional control-schedule complexity. The clustering criteria are designed, therefore, to choose a Δt_{int} sufficient to produce a desirable level of clustering that leads to a desirable level of control-schedule simplification. The *while*-loop continues while the value of Δt_{int} remains within an acceptable range of values. When the clustering criteria fails to be satisfied by repeated calls to the routine "intervalCluster" in the *while*-loop of steps 2538-2542, then, in step 2543, one or more alternative clustering methods may be employed to generate clusters, when needed for control-schedule simplification. Alternative methods may involve selecting clusters based on local maximum and minimum parameter values indicated in the control schedule or, when all else fails, by selecting, as cluster boundaries, a number of the longest setpoint-free time intervals within the setpoints generated in step 2537.

Figure 25E provides a control-flow diagram for the routine "interval cluster" called in step 2539 of Figure 25D. In step 2545, the intelligent controller determines whether or not a setpoint coincides with the beginning time of the control schedule corresponding to the monitoring period. When a setpoint does coincide with the beginning of the time of the control schedule, as determined in step 2545, then the local variable "startCluster" is set to the start time of the control schedule and the local variable "numCluster" is set to 1, in step 2546. Otherwise, the local variable "numCluster" is set to 0 in step 2547. In step 2548, the local variable "lastSP" is set to the start time of the control schedule and the local variable "curT" is set to "lastSP" plus a time increment Δt_{inc} in step 2548. The local variable "curT" is an indication of the current time point in the control schedule being considered, the local variable "numCluster" is an indication of the number of setpoints in a next cluster that is being created, the local variable "startCluster" is an indication of the point in time of the first setpoint in the cluster, and the local variable "lastSP" is an indication of the time of the last detected setpoint in the control schedule. Next, in the *while*-loop of steps 2549-2559, the control schedule corresponding to the

- 42 -

monitoring period is traversed, from start to finish, in order to generate a sequence of clusters from the control schedule. In step 2550, a local variable Δt is set to the length of the time interval between the last detected setpoint and the current point in time that is being considered. When there is a setpoint that coincides with the current point in time, as determined in step 2551, then a routine "nextSP" is called, in step 2552, to consider and process the setpoint. Otherwise, when Δt is greater than Δt_{int} , as determined in step 2553, then, when a cluster is being processed, as determined in step 2554, the cluster is closed and stored, in step 2555, and the local variable "numCluster" is reinitialized to begin processing of a next cluster. The local variable "curt" is incremented, in step 2556, and the *while*-loop continues to iterate when *curT* is less than or equal to the time at which the control schedule ends, as determined in step 2557. When the *while*-loop ends, and when a cluster was being created, as determined in step 2558, then that cluster is closed and stored in step 2559.

Figure 25F provides a control-flow diagram for the routine "next SP" called in step 2552 of Figure 25E. In step 2560, the intelligent controller determines whether or not a cluster was being created at the time of the routine call. When a cluster was being created, and when Δt is less than Δt_{int} , as determined in step 2561, then the current setpoint is added to the cluster in step 2562. Otherwise, the currently considered cluster is closed and stored, in step 2563. When a cluster was not being created, then the currently detected setpoint becomes the first setpoint in a new cluster, in step 2564.

Figure 25G provides a control-flow diagram for the routine "simplify clusters" called in step 2529 of Figure 25C. This routine is a simple *for*-loop, comprising steps 2566-2568 in which each cluster, determined by the routine "cluster" called in step 2528 of Figure 25C, is simplified, as discussed above with reference to Figures 16A-21D. The cluster is simplified by a call to the routine "simplify" in step 2567.

Figure 25H is a control-flow diagram for the routine "simplify" called in step 2567 of Figure 25G. In step 2570, the intelligent controller determines whether or not the currently considered cluster contains any schedule-change setpoints. When the currently considered cluster contains schedule-change setpoints, then any immediate-control setpoints are removed, in step 2572. When the cluster contains only a single schedule-change setpoint, as determined in step 2573, then that single schedule-change setpoint is

- 43 -

left to represent the entire cluster, in step 2574. Otherwise, the multiple schedule changes are resolved into zero, one, or two setpoints to represent the cluster as discussed above with reference to Figures 17A-E in step 2575. The zero, one, or two setpoints are then entered into the existing control schedule in step 2576. When the cluster does not contain any schedule-change setpoints, as determined in step 2570, and when the setpoints in the cluster can be replaced by a single setpoint, as determined in step 2577, as discussed above with reference to Figures 17A and 17C, then the setpoints of the cluster are replaced with a single setpoint, in step 2578, as discussed above with reference to Figures 17A and 17C. Note that, as discussed above with reference to Figures 20A-C, the setpoints are associated with labels "s" and "i" to indicate whether they are derived from scheduled setpoints or from immediate-control setpoints. Similarly, when the setpoints of the cluster can be replaced by two setpoints, as determined in step 2579, then the cluster is replaced by the two setpoints with appropriate labels, as discussed above with reference to Figures 17D-E, in step 2580. Otherwise, the condition described with reference to Figure 717B has occurred, in which case all of the remaining setpoints are deleted from the cluster in step 2581.

Figure 25I provides a control-flow diagram for the routine "generate new schedule" called in step 2530 of Figure 25C. When the new provisional schedule includes two or more immediate-control setpoints, as determined in step 2583, then the routine "spread" is called in step 2584. This routine spreads "i"-labeled setpoints, as discussed above to Figures 21A-B. The control schedule is then stored as a new current control schedule for the time period in step 2585 with the indications of whether the setpoints are derived from immediate-control setpoints or schedule setpoints retained for a subsequent propagation step in step 2586.

Figure 25J provides a control-flow diagram for the routine "spread," called in step 2584 in Figure 25I. In step 2587, the local variable "first" is set to the first immediate-control setpoint in the provisional schedule. In step 2588, the variable "second" is set to the second immediate-control setpoint in the provisional schedule. Then, in the *while*-loop of steps 2589-2599, the provisional schedule is traversed in order to detect pairs of immediate-control setpoints that are closer together, in time, than a threshold length of

- 44 -

time Δt_j . The second setpoint is moved, in time, in steps 2592-2596, either by a fixed time interval Δt_s or to a point halfway between the previous setpoint and the next setpoint, in order to spread the immediate-control setpoints apart.

Figure 25K provides a control-flow diagram for the routine "propagate new schedule" called in step 2531 of Figure 25C. This routine propagates a provisional schedule created in step 2530 in Figure 25C to related sub-schedules, as discussed above with reference to Figures 22A-B. In step 2599a, the intelligent controller determines the additional sub-schedules or schedules to which the provisional schedule generated in step 2530 should be propagated. Then, in the *for*-loop of steps 2599b-2599e, the retained immediate-control setpoints, retained in step 2586 in Figure 25I, are propagated to a next related control schedule and those setpoints, along with existing-control-schedule setpoints in the next control schedule, are resolved by a call to the routine "resolve additional schedule," in step 2599d.

Figure 25L provides a control-flow diagram for the routine "resolve additional schedule," called in step 2599d of Figure 25K. In step 2599f, the intelligent controller accesses a stored set of schedule-resolution rules, such as those discussed above with reference to Figures 24A-I, and sets the local variable *j* to the number of schedule-resolution rules to be applied. Again, in the nested *for*-loops of steps 2599g-2599n, the rules are applied to each immediate-control setpoint in the set of setpoints generated in step 2599c of Figure 25K. The rules are applied in sequence to each immediate-control setpoint until either the setpoint is deleted, as determined in step 2599j, or until the rule is successfully applied to simplify the schedule, in step 2599k. Once all the propagated setpoints have been resolved in the nested *for*-loops of steps 2599g-2599n, then the schedule is stored as a new provisional schedule, in step 2599o.

Figure 25M provides a control-flow diagram for the routine "steady-state monitoring" called in step 2525 of Figure 25B. This routine is similar to the routine "aggressive monitoring period" shown in Figure 25C and called in step 2524 of Figure 25B. Many of the steps are, in fact, nearly identical, and will not be again described, in the interest of brevity. However, step 2599q is an additional step not present in the routine "aggressiveMonitoringPeriod." In this step, the immediate-control setpoints and schedule-

- 45 -

change setpoints overlaid on the existing-control-schedule setpoints are used to search a database of recent historical control schedules in order to determine whether or not the set of setpoints is more closely related to another control schedule to which the intelligent controller is to be targeted or shifted. When the control-schedule shift is indicated by this search, as determined in step 2599h, then the shift is carried out in step 2599i, and the stored immediate-controls and schedule changes are combined with a sub-schedule of the target schedule to which the intelligent controller is shifted, in step 2599t, prior to carrying out generation of the new provisional schedule. The historical-search routine, called in step 2599q, may also filter the recorded immediate-control setpoints and schedule-change setpoints recorded during the monitoring period with respect to one or more control schedules or sub-schedules corresponding to the monitoring period. This is part of a more conservative learning approach, as opposed to the aggressive learning approach used in the aggressive-learning mode, that seeks to only conservatively alter a control schedule based on inputs recorded during a monitoring period. Thus, while the procedure carried out at the end of a monitoring period are similar both for the aggressive-learning mode and the steady-state learning mode, schedule changes are carried out in a more conservative fashion during steady-state learning, and the schedule changes become increasingly conservative with each successive phase of steady-state learning. With extensive recent and historical control-schedule information at hand, the intelligent controller can make intelligent and increasingly accurate predictions of whether immediate-control inputs and schedule changes that occurred during the monitoring period reflect the user's desire for long-term changes to the control schedule or, instead, reflect temporary control changes related to temporally local events and conditions.

As mentioned above, an intelligent controller may employ multiple different control schedules that are applicable over different periods of time. For example, in the case of a residential HVAC thermostat controller, an intelligent controller may use a variety of different control schedules applicable to different seasons during the year; perhaps a different control schedule for winter, summer, spring, and fall. Other types of intelligent controllers may use a number of control schedules for various different periods of control that span minutes and hours to months, years, and even greater periods of time.

- 46 -

Figure 26 illustrates three different week-based control schedules corresponding to three different control modes for operation of an intelligent controller. Each of the three control schedules 2602-2604 is a different week-based control schedule that controls intelligent-controller operation for a period of time until operational control is shifted, in step 2599s of Figure 25M, to another of the control schedules. Figure 27 illustrates a state-transition diagram for an intelligent controller that operates according to seven different control schedules. The modes of operation controlled by the particular control schedules are shown as disks, such as disk 2702, and the transitions between the modes of operation are shown as curved arrows, such as curved arrow 2704. In the case shown in Figure 27, the state-transition diagram expresses a deterministic, higher-level control schedule for the intelligent controller comprising seven different operational modes, each controlled by a particular control schedule. Each of these particular control schedules may, in turn, be composed of additional hierarchical levels of sub-schedules. The automated-learning methods to which the current application is directed can accommodate automated learning of multiple control schedules and sub-schedules, regardless of their hierarchical organization. Monitoring periods generally encompass the shortest-time, smallest-grain sub-schedules in a hierarchy, and transitions between sub-schedules and higher-level control schedules are controlled by higher-level control schedules, such as the transition-state-diagram-expressed higher-level control schedule illustrated in Figure 27, by the sequential ordering of sub-schedules within a larger control schedule, such as the daily sub-schedules within a weekly control schedule discussed with reference to Figure 13, or according to many other control-schedule organizations and schedule-shift criteria.

Figures 28A-C illustrate one type of control-schedule transition that may be carried out by an intelligent controller. Figure 28A shows the existing control schedule according to which the intelligent controller is currently operating. Figure 28B shows recorded immediate-control inputs over a recently completed monitoring period superimposed over the control schedule shown in Figure 28A. These immediate-control inputs 2802-2805 appear to represent a significant departure from the existing control schedule 2800. In step 2599q of Figure 25M, an intelligent controller may consider various alternative control schedules or historical control schedules, including control schedule

- 47 -

2810, shown in Figure 28C, that may be alternate control schedules for the recently completed monitoring period. As it turns out, resolution of the immediate-control inputs with the existing control schedule would produce a control schedule very close to control schedule 2810 shown in Figure 28C. This then provides a strong indication to the intelligent controller that the recorded immediate-control inputs may suggest a need to shift control to control schedule 2810 rather than to modify the existing control schedule and continue using the modified control schedule. Although this is one type of schedule-change transition that may occur in step 2599s in Figure 25M, other schedule-change shifts may be controlled by knowledge of the current date, day of the week, and perhaps knowledge of various environmental parameters that together specify use of multiple control schedules to be used to control intelligent-control operations.

Figures 29-30 illustrate types of considerations that may be made by an intelligent controller during steady-state-learning phases. In Figure 29, the plot of a new provisional schedule 2902 is shown, along with similar plots of 15 recent or historical control schedules or provisional schedules applicable to the same time period 2904-2918. Visual comparison of the new provisional schedule 2902 to the recent and historical provisional schedules 2904-2918 immediately reveals that the new provisional schedule represents a rather radical change in the control regime. During steady-state learning, such radical changes may not be propagated or used to replace existing control schedules, but may instead be recorded and used for propagation or replacement purposes only when the accumulated record of recent and historical provisional schedules provide better support for considering the provisional schedule as an indication of future user intent. For example, as shown in Figure 30, were the new provisional schedule compared to a record of recent and/or historical control schedules 3002-3016, the intelligent controller would be far more likely to use new provisional schedule 2902 for replacement or propagation purposes.

Automated Schedule Learning in the Context of an Intelligent Thermostat

An implementation of automated control-schedule learning is included in a next-described intelligent thermostat. The intelligent thermostat is provided with a

- 48 -

selectively layered functionality that exposes unsophisticated users to a simple user interface, but provides advanced users with an ability to access and manipulate many different energy-saving and energy tracking capabilities. Even for the case of unsophisticated users who are only exposed to the simple user interface, the intelligent thermostat provides advanced energy-saving functionality that runs in the background. The intelligent thermostat uses multi-sensor technology to learn the heating and cooling environment in which the intelligent thermostat is located and to optimize energy-saving settings.

The intelligent thermostat also learns about the users, beginning with a setup dialog in which the user answers a few simple questions, and then continuing, over time, using multi-sensor technology to detect user occupancy patterns and to track the way the user controls the temperature using schedule changes and immediate-control inputs. On an ongoing basis, the intelligent thermostat processes the learned and sensed information, automatically adjusting environmental control settings to optimize energy usage while, at the same time, maintaining the temperature within the environment at desirable levels, according to the learned occupancy patterns and comfort preferences of one or more users.

Advantageously, the selectively layered functionality of the intelligent thermostat allows for effective operation in a variety of different technological circumstances within home and business environments. For simple environments having no wireless home network or Internet connectivity, the intelligent thermostat operates effectively in a standalone mode, learning and adapting to an environment based on multi-sensor technology and user input. However, for environments that have home network or Internet connectivity, the intelligent thermostat operates effectively in a network-connected mode to offer additional capabilities.

When the intelligent thermostat is connected to the Internet via a home network, such as through IEEE 802.11 (Wi-Fi) connectivity, the intelligent thermostat may: (1) provide real-time or aggregated home energy performance data to a utility company, intelligent thermostat data service provider, intelligent thermostats in other homes, or other data destinations; (2) receive real-time or aggregated home energy performance data from a utility company, intelligent thermostat data service provider, intelligent thermostats in other

- 49 -

homes, or other data sources; (3) receive new energy control instructions and/or other upgrades from one or more intelligent thermostat data service providers or other sources; (4) receive current and forecasted weather information for inclusion in energy-saving control algorithm processing; (5) receive user control commands from the user's computer, network-connected television, smart phone, and/or other stationary or portable data communication appliance; (6) provide an interactive user interface to a user through a digital appliance; (7) receive control commands and information from an external energy management advisor, such as a subscription-based service aimed at leveraging collected information from multiple sources to generate energy-saving control commands and/or profiles for their subscribers; (8) receive control commands and information from an external energy management authority, such as a utility company to which limited authority has been voluntarily given to control the intelligent thermostat in exchange for rebates or other cost incentives; (9) provide alarms, alerts, or other information to a user on a digital appliance based on intelligent thermostat-sensed HVAC-related events; (10) provide alarms, alerts, or other information to the user on a digital appliance based on intelligent thermostat-sensed non-HVAC related events; and (11) provide a variety of other useful functions enabled by network connectivity.

Figure 31A illustrates a perspective view of an intelligent thermostat. The intelligent thermostat 3100 has a sleek, elegant appearance. The intelligent thermostat 3100 comprises a circular main body 3108 with a diameter of about 8 cm and that has a visually pleasing outer finish, such as a satin nickel or chrome finish. A cap-like structure comprising a rotatable outer ring 3106, a sensor ring 3104, and a circular display monitor 3102 is separated from the main body 3108 by a small peripheral gap 3110. The outer ring 3106 may have an outer finish identical to that of the main body 3108, while the sensor ring 3104 and circular display monitor 3102 may have a common circular glass (or plastic) outer covering that is gently arced in an outward direction and that provides a sleek yet solid and durable-looking overall appearance. The sensor ring 3104 contains any of a wide variety of sensors, including infrared sensors, visible-light sensors, and acoustic sensors. The glass or plastic that covers the sensor ring 3104 may be smoked or mirrored such that the sensors themselves are not visible to the user. An air venting functionality may be provided, via the

- 50 -

peripheral gap 3110, which allows the ambient air to be sensed by the internal sensors without the need for gills or grill-like vents.

Figures 31B-31C illustrate the intelligent thermostat being controlled by a user. The intelligent thermostat 3100 is controlled by two types of user input: (1) a rotation of the outer ring 3106 (Figure 31B); and (2) an inward push on the outer ring 3106 (Figure 31C) until an audible and/or tactile “click” occurs. The inward push may cause the outer ring 3106 to move forward, while in another implementation, the entire cap-like structure, including both the outer ring 3106 and the glass covering of the sensor ring 3104 and circular display monitor 3102, move inwardly together when pushed. The sensor ring 3104, the circular display monitor 3102, and the common glass covering do not rotate with outer ring 3106 in one implementation.

By rotation of the outer ring 3106, or ring rotation, and inward pushing of the outer ring 3106, or inward click, the intelligent thermostat 3100 can receive all necessary information from the user for basic setup and operation. The outer ring 3106 is mechanically mounted in a manner that provides a smooth yet viscous feel to the user, for further promoting an overall feeling of elegance while also reducing spurious or unwanted rotational inputs. The intelligent thermostat 3100 recognizes three fundamental user inputs: (1) ring rotate left, (2) ring rotate right, and (3) inward click. In other implementations, more complex fundamental user inputs can be recognized, such as double-click or triple-click inward presses and speed-sensitive or acceleration-sensitive rotational inputs.

Figure 32 illustrates an exploded perspective view of the intelligent thermostat and an HVAC-coupling wall dock. The HVAC-coupling wall dock 3206 has the functionality as a very simple, elemental, standalone thermostat when the intelligent thermostat 3100 is removed, the elemental thermostat including a standard temperature readout/setting dial 3208 and a simple COOL-OFF-HEAT switch 3209. This can prove useful for a variety of situations, such as when the intelligent thermostat 3100 needs to be removed for service or repair for an extended period of time. In one implementation, the elemental thermostat components 3208 and 3209 are entirely mechanical in nature, so that no electrical power is needed to trip the control relays. In other implementations, simple electronic controls, such as electrical up/down buttons and/or an LCD readout, are

- 51 -

provided. In other implementations, a subset of the advanced functionalities of the intelligent thermostat 3100 can be provided, such as elemental network access to allow remote control that provides a brain-stem functionality while the intelligent thermostat is temporarily removed.

Figures 33A-33B illustrate exploded front and rear perspective views of the intelligent thermostat. Figures 33A-33B show the intelligent thermostat 3302 with respect to its two main components: (1) the head unit 3204; and (2) the back plate 3306. In the drawings shown, the *z* direction is outward from the wall, the *y* direction is the head-to-toe direction relative to a walk-up user, and the *x* direction is the user's left-to-right direction.

Figures 34A-34B illustrate exploded front and rear perspective views, respectively, of the head unit. Head unit 3304 includes a head unit frame 3310, the outer ring 3311, a head unit frontal assembly 3312, a front lens 3313, and a front grille 3314. Electrical components on the head unit frontal assembly 3312 can connect to electrical components on the backplate 3306 via ribbon cables and/or other plug type electrical connectors.

Figures 35A-35B illustrate exploded front and rear perspective views, respectively, of the head unit frontal assembly. Head unit frontal assembly 3312 comprises a head unit circuit board 3316, a head unit front plate 3318, and an LCD module 3322. The components of the front side of head unit circuit board 3316 are hidden behind an RF shield in Figure 35A. A rechargeable Lithium-Ion battery 3325 is located on the back of the head unit circuit board 3316, which, in one implementation, has a nominal voltage of 3.7 volts and a nominal capacity of 560 mAh. To extend battery life, the battery 3325 is normally not charged beyond 450 mAh by the thermostat battery charging circuitry. Moreover, although the battery 3325 is rated to be capable of being charged to 4.2 volts, the intelligent thermostat battery-charging circuitry normally does not charge the intelligent thermostat beyond 3.95 volts. Also shown in Figure 35B is an optical finger navigation module 3324 that is configured and positioned to sense rotation of the outer ring 3311. The module 3324 uses methods analogous to the operation of optical computer mice to sense the movement of a texturable surface on a facing periphery of the outer ring 3311. Notably, the module 3324 is one of the very few sensors that are controlled by the relatively power-intensive

- 52 -

head unit microprocessor rather than the relatively low-power backplate microprocessor. This is achievable without excessive power drain because the head unit microprocessor is already awake when a user is manually turning the dial, avoiding excessive wake-up power drain. Advantageously, very fast response can also be provided by the head unit microprocessor. Also shown in Figure 35A is a Fresnel lens 3320 that operates in conjunction with a PIR motion sensor disposed thereunderneath.

Figures 36A-36B illustrate exploded front and rear perspective views, respectively, of the backplate unit. Backplate unit 3306 comprises a backplate rear plate 3330, a backplate circuit board 3332, and a backplate cover 3339. Figure 36A shows the HVAC wire connectors 3334 that include integrated wire-insertion-sensing circuitry and two relatively large capacitors 3336 that are used by the power stealing circuitry that is mounted on the back side of the backplate circuit board 3332.

Figure 37 shows a perspective view of a partially assembled head unit. In certain implementations, placement of grille member 3314 over the Fresnel lens 3320 and an associated PIR motion sensor 3344 conceals and protects these PIR sensing elements, while horizontal slots in the grille member 3314 allow the PIR motion sensing hardware, despite being concealed, to detect the lateral motion of occupants in a room or area. A temperature sensor 3340 uses a pair of thermal sensors to more accurately measure ambient temperature. A first or upper thermal sensor 3341 associated with temperature sensor 3340 gathers temperature data closer to the area outside or on the exterior of the thermostat while a second or lower thermal sensor 3342 collects temperature data more closely associated with the interior of the housing. In one implementation, each of the temperature sensors 3341 and 3342 comprises a Texas Instruments TMP112 digital temperature sensor chip, while the PIR motion sensor 3344 comprises PerkinElmer DigiPyro PYD 1998 dual-element pyrodetector.

To more accurately determine the ambient temperature, the temperature taken from the lower thermal sensor 3342 is considered in conjunction with the temperatures measured by the upper thermal sensor 3341 and when determining the effective ambient temperature. This configuration can be used to compensate for the effects of internal heat produced in the thermostat by the microprocessor(s) and/or other electronic

- 53 -

components, obviating or minimizing temperature measurement errors that might otherwise be suffered. In some implementations, the accuracy of the ambient temperature measurement may be further enhanced by thermally coupling upper thermal sensor 3341 of temperature sensor 3340 to grille member 3314 as the upper thermal sensor 3341 better reflects the ambient temperature than lower thermal sensor 3342.

Figure 38 illustrates the head unit circuit board. The head unit circuit board 3316 comprises a head unit microprocessor 3802 (such as a Texas Instruments AM3703 chip) and an associated oscillator 3804, along with DDR SDRAM memory 3806, and mass NAND storage 3808. A Wi-Fi module 3810, such as a Murata Wireless Solutions LBWA19XSLZ module, which is based on the Texas Instruments WL1270 chipset supporting the 802.11 b/g/n WLAN standard, is provided in a separate compartment of RF shielding 3834 for Wi-Fi capability. Wi-Fi module 3810 is associated with supporting circuitry 3812 including an oscillator 3814. A ZigBee module 3816, which can be, for example, a C2530F256 module from Texas Instruments, is provided, also in a separately shielded RF compartment, for ZigBee capability. The ZigBee module 3816 is associated with supporting circuitry 3818, including an oscillator 3819 and a low-noise amplifier 3820. Display backlight voltage conversion circuitry 3822, piezoelectric driving circuitry 3824, and power management circuitry 3826 are additionally provided. A proximity sensor and an ambient light sensor (PROX/ALS), more particularly a Silicon Labs SI1142 Proximity/Ambient Light Sensor with an I2C Interface, is provided on a flex circuit 3828 that attaches to the back of the head unit circuit board by a flex circuit connector 3830. Battery-charging-supervision-disconnect circuitry 3832 and spring/RF antennas 3836 are additionally provided. A temperature sensor 3838 and a PIR motion sensor 3840 are additionally provided.

Figure 39 illustrates a rear view of the backplate circuit board. The backplate circuit board 3332 comprises a backplate processor/microcontroller 3902, such as a Texas Instruments MSP430F System-on-Chip Microcontroller that includes an on-board memory 3903. The backplate circuit board 3332 further comprises power-supply circuitry 3904, which includes power-stealing circuitry, and switch circuitry 3906 for each HVAC respective HVAC function. For each such function, the switch circuitry 3906 includes an

- 54 -

isolation transformer 3908 and a back-to-back NFET package 3910. The use of FETs in the switching circuitry allows for active power stealing, i.e., taking power during the HVAC ON cycle, by briefly diverting power from the HVAC relay circuit to the reservoir capacitors for a very small interval, such as 100 micro-seconds. This time is small enough not to trip the HVAC relay into the OFF state but is sufficient to charge up the reservoir capacitors. The use of FETs allows for this fast switching time (100 micro-seconds), which would be difficult to achieve using relays (which stay on for tens of milliseconds). Also, such relays would readily degrade with fast switching, and they would also make audible noise. In contrast, the FETS operate with essentially no audible noise. A combined temperature/humidity sensor module 3912, such as a Sensirion SHT21 module, is additionally provided. The backplate microcontroller 3902 performs polling of the various sensors, sensing for mechanical wire insertion at installation, alerting the head unit regarding current vs. setpoint temperature conditions and actuating the switches accordingly, and other functions such as looking for appropriate signal on the inserted wire at installation.

Next, an implementation of the above-described automated-control-schedule-learning methods for the above-described intelligent thermostat is provided. Figures 40A-D illustrate steps for achieving initial learning. Figures 41A-M illustrate a progression of conceptual views of a thermostat schedule. The progression of conceptual views of a thermostat schedule occurs as processing is performed according to selected steps of Figures 40A-40D, for an example one-day monitoring period during an initial aggressive-learning period. For one implementation, the steps of Figures 40A-40D are carried out by the head unit microprocessor of thermostat 3302, with or without Internet connectivity. In other implementations, one or more of the steps of Figures 40A-40D can be carried out by a cloud server to which the thermostat 3302 has network connectivity. While the example presented in Figures 41A-41M is for a heating schedule scenario, the described method is likewise applicable for cooling-schedule learning, and can be readily extended to HVAC schedules containing mixtures of heating setpoints, cooling setpoints, and/or range setpoints. While the examples of Figures 40A-41M are presented in the particular context of establishing a weekly schedule, which represents one particularly

- 55 -

appropriate time basis for HVAC schedule establishment and execution, in other implementations a bi-weekly HVAC schedule, a semi-weekly HVAC schedule, a monthly HVAC schedule, a bi-monthly HVAC schedule, a seasonal HVAC schedule, and other types of schedules may be established. While the examples of Figures 40A-41M are presented and/or discussed in terms of a typical residential installation, this is for the purpose of clarity of explanation. The methods are applicable to a wide variety of other types of enclosures, such as retail stores, business offices, industrial settings, and so forth. In the discussion that follows, the time of a particular user action or setpoint entry are generally expressed as both the day and the time of day of that action or entry, while the phrase "time of day" is generally used to express a particular time of day.

The initial learning process represents an "aggressive learning" approach in which the goal is to quickly establish an at least roughly appropriate HVAC schedule for a user or users based on a very brief period of automated observation and tracking of user behavior. Once the initial learning process is established, the thermostat 3302 then switches over to steady-state learning, which is directed to perceiving and adapting to longer-term repeated behaviors of the user or users. In most cases, the initial learning process is begun, in step 4002, in response to a new installation and startup of the thermostat 3302 in a residence or other controlled environment, often following a user-friendly setup interview. Initial learning can also be invoked by other events, such as a factory reset of the intelligent thermostat 3302 or an explicit request of a user who may wish for the thermostat 3302 to repeat the aggressive-learning phase.

In step 4004, a default beginning schedule is accessed. For one implementation, the beginning schedule is simply a single setpoint that takes effect at 8 AM each day and that includes a single setpoint temperature. This single setpoint temperature is dictated by a user response that is provided near the end of the setup interview or upon invocation of initial learning, where the user is asked whether to start learning a heating schedule or a cooling schedule. When the user chooses heating, the initial single setpoint temperature is set to 68° F, or some other appropriate heating setpoint temperature, and when the user chooses cooling, the initial single setpoint temperature is set to 80° F, or some other appropriate cooling setpoint temperature. In other implementations, the default

- 56 -

beginning schedule can be one of a plurality of predetermined template schedules that is selected directly or indirectly by the user at the initial setup interview. Figure 41A illustrates an example of a default beginning schedule having heating setpoints labeled "a" through "g".

In step 4006, a new monitoring period is begun. The selection of a one-day monitoring period has been found to provide good results in the case of control-schedule acquisition in an intelligent thermostat. However, other monitoring periods can be used, including multi-day blocks of time, sub-day blocks of time, other suitable periods, and can alternatively be variable, random, or continuous. For example, when performed on a continuous basis, any user setpoint change or scheduled setpoint input can be used as a trigger for processing that information in conjunction with the present schedule to produce a next version, iteration, or refinement of the schedule. For one implementation, in which the thermostat 3302 is a power-stealing thermostat having a rechargeable battery, the period of one day has been found to provide a suitable balance between the freshness of the schedule revisions and the need to maintain a modest computing load on the head unit microprocessor to preserve battery power.

In step 4008, throughout the day, the intelligent thermostat 3302 receives and stores both immediate-control and schedule-change inputs. Figure 41B shows a representation of a plurality of immediate-control and schedule-change user setpoint entries that were made on a typical day of initial learning, which happens to be a Tuesday in the currently described example. In the following discussion and in the accompanying drawings, including Figures 41A-41M, a preceding superscript "N" identifies a schedule-change, or non-real-time ("NRT"), setpoint entry and a preceding superscript "R" identifies an immediate-control, or real-time ("RT") setpoint entry. An encircled number represents a pre-existing scheduled setpoint. For each NRT setpoint, a succeeding subscript that identifies the entry time of that NRT setpoint is also provided. No such subscript is needed for RT setpoints, since their horizontal position on the schedule is indicative of both their effective time and their entry time. Thus, in the example shown in Figure 41B, at 7:30 AM a user made an RT setpoint entry "i" having a temperature value of 76° F, at 7:40 AM a user made another RT setpoint entry "j" having a temperature value of 72° F, at 9:30 AM a

- 57 -

user made another RT setpoint entry "l" having a temperature value of 72° F, at 11:30 AM a user made another RT setpoint entry "m" having a temperature value of 76° F, and so on. On Tuesday, at 10 AM, a user created, through a scheduling interface, an NRT setpoint entry "n" that is to take effect on Tuesdays at 12:00 PM and created an NRT setpoint entry "w" that is to take effect on Tuesdays at 9:00 PM. Subsequently, on Tuesday at 4:00 PM, a user created an NRT setpoint entry "h" that is to take effect on Mondays at 9:15 PM and created an NRT setpoint entry "k" that was to take effect on Tuesdays at 9:15 AM. Finally, on Tuesday at 8 PM, a user created an NRT setpoint entry "s" that is to take effect on Tuesdays at 6:00 PM.

Referring now to step 4010, throughout the 24-hour monitoring period, the intelligent thermostat controls the HVAC system according to whatever current version of the control schedule is in effect as well as whatever RT setpoint entries are made by the user and whatever NRT setpoint entries have been made that are causally applicable. The effect of an RT setpoint entry on the current setpoint temperature is maintained until the next pre-existing setpoint is encountered, until a causally applicable NRT setpoint is encountered, or until a subsequent RT setpoint entry is made. Thus, with reference to Figures 41A-41B, on Tuesday morning, at 6:45PM, the current operating setpoint of the thermostat changes to 73° F due to pre-existing setpoint "b," then, at 7:30 AM, the current operating setpoint changes to 76° F due to RT setpoint entry "i," then, at 7:45 AM, the current operating setpoint changes to 72° F due to RT setpoint entry "j," then, at 8:15 AM, the current operating setpoint changes to 65° F due to pre-existing setpoint entry "c," then, at 9:30 AM, the current operating setpoint changes to 72° F due to RT setpoint entry "l," then, at 11:30 AM, the current operating setpoint changes to 76° F due to RT setpoint entry "m," then at 12:00PM the current operating setpoint changes to 71° F due to NRT setpoint entry "n," then, at 12:15 PM, the current operating setpoint changes to 78° F, due to RT setpoint entry "o," and so forth. At 9:15 AM, there is no change in the current setpoint due to NRT setpoint entry "k" because it did not yet exist. By contrast, the NRT setpoint entry "n" is causally applicable because it was entered by the user at 10AM that day and took effect at its designated effective time of 12:00PM.

- 58 -

According to one optional alternative embodiment, step 4010 can be carried out so that an RT setpoint entry is only effective for a maximum of 2 hours, or other relatively brief interval, as the operating setpoint temperature, with the operating setpoint temperature returning to whatever temperature would be specified by the pre-existing setpoints on the current schedule or by any causally applicable NRT setpoint entries. This optional alternative embodiment is designed to encourage the user to make more RT setpoint entries during the initial learning period so that the learning process can be achieved more quickly. As an additional optional alternative, the initial schedule, in step 4004, is assigned with relatively low-energy setpoints, as, for example, relatively low-temperature setpoints in winter, such as 62° F, which generally produces a lower-energy control schedule. As yet another alternative, during the first few days, instead of reverting to pre-existing setpoints after 2 hours, the operating setpoint instead reverts to a lowest-energy pre-existing setpoint in the schedule.

Referring now to step 4012, at the end of the monitoring period, the stored RT and NRT setpoints are processed with respect to one another and the current schedule to generate a modified version, iteration, or refinement of the schedule, the particular steps for which are shown in Figure 40B. This processing can be carried out, for example, at 11:50 PM of the learning day, or at some other time near or around midnight. When it is determined that the initial learning is not yet complete, in step 4014, the modified version of the schedule is used for another day of initial learning, in steps 4006-4010, is yet again modified in step 4012, and the process continues until initial learning is complete. When initial learning is complete, steady-state learning begins in step 4016, described below with respect to Figures 32A-33.

For some implementations, the decision, in step 4014, regarding whether or not the initial control-schedule learning is complete is based on both the passage of time and whether there has been a sufficient amount of user behavior to record and process. For one implementation, the initial learning is considered to be complete only when two days of initial learning have passed and there have been ten separate one-hour intervals in which a user has entered an RT or NRT setpoint. Any of a variety of different criteria can be used to determine whether there has been sufficient user interaction to conclude initial learning.

- 59 -

Figure 40B illustrates steps for processing stored RT and NRT setpoints that correspond generally to step 4012 of Figure 40A. In step 4030, setpoint entries having nearby effective times are grouped into clusters, as illustrated in Figure 41C. In one implementation, any set of two or more setpoint entries for which the effective time of each member is separated by less than 30 minutes from that of at least one other member constitutes a single cluster.

In step 4032, each cluster of setpoint entries is processed to generate a single new setpoint that represents the entire cluster in terms of effective time and temperature value. This process is directed to simplifying the schedule while, at the same time, best capturing the true intent of the user by virtue of the user's setpoint-entry behavior. While a variety of different approaches, including averaging of temperature values and effective times of cluster members, can be used, one method for carrying out step 4032, described in more detail in Figure 40C, takes into account the NRT vs. RT status of each setpoint entry, the effective time of each setpoint entry, and the entry time of each setpoint entry.

Referring now to Figure 40C, which corresponds to step 4032 of Figure 40B, a determination is made, in step 4060, whether there are any NRT setpoint entries in the cluster having an entry time that is later than the earliest effective time in the cluster. When this is the case, then, in step 4064, the cluster is replaced by a single representative setpoint with both the effective time and the temperature value of the latest-entered NRT setpoint entry. This approach provides deference to the wishes of the user who has taken the time to specifically enter a desired setpoint temperature for that time. When, in step 4060, there are no such NRT setpoint entries, then, in step 4062, the cluster is replaced by a single representative setpoint with an effective time of the earliest effective cluster member and a setpoint temperature equal to that of the cluster member having the latest entry time. This approach provides deference to the wishes of the user as expressed in the immediate-control inputs and existing setpoints.

Referring again to Figure 40B, in step 4034, the new representative setpoint that determined in step 4032 is tagged with an "RT" or "NRT" label based on the type of setpoint entry from which the setpoint's temperature value was assigned. Thus, in accordance with the logic of Figure 40C, were an NRT setpoint to have the latest-occurring

- 60 -

time of entry for the cluster, the new setpoint would be tagged as "NRT." Were an RT setpoint to have the latest-occurring time of entry, the new setpoint would be tagged as "RT." In steps 4036-4038, any singular setpoint entries that are not clustered with other setpoint entries are simply carried through as new setpoints to the next phase of processing, in step 4040.

Referring to Figures 41C-41D, it can be seen that, for the "ij" cluster, which has only RT setpoint entries, the single representative setpoint "ij" is assigned to have the earlier effective time of RT setpoint entry "i" while having the temperature value of the later-entered RT setpoint entry "j," representing an application of step 4062 of Figure 40C, and that new setpoint "ij" is assigned an "RT" label in step 4034. It can further be seen that, for the "kl" cluster, which has an NRT setpoint "k" with an entry time later than the earliest effective time in that cluster, the single representative setpoint "kl" is assigned to have both the effective time and temperature value of the NRT setpoint entry "k," representing an application of step 4064 of Figure 40C, and that new setpoint "kl" is assigned an "NRT" label in step 4034. For the "mno" cluster, which has an NRT setpoint "n" but with an entry time earlier than the earliest effective time in that cluster, the single representative setpoint "mno" is assigned to have the earliest effective time of RT setpoint entry "m" while having the temperature value of the latest-entered setpoint entry "o," again representing an application of step 4062 of Figure 40C, and that new setpoint "mno" is assigned an "RT" label in step 4034. The remaining results shown in Figure 41D, all of which are also considered to be new setpoints at this stage, also follow from the methods of Figures 40B-40C.

Referring again to Figure 40B, step 4040 is next carried out after steps 4034 and 4038 and applied to the new setpoints as a group, which are shown in Figure 41D. In step 4040, any new setpoint having an effective time that is 31-60 minutes later than that of any other new setpoint is moved, in time, to have a new effective time that is 60 minutes later than that of the other new setpoint. This is shown in Figure 41E with respect to the new setpoint "q," the effective time of which is being moved to 5:00 PM so that it is 60 minutes away from the 4:00 PM effective time of the new setpoint "p." In one implementation, this process is only performed a single time based on an instantaneous

snapshot of the schedule at the beginning of step 4040. In other words, there is no iterative cascading effect with respect to these new setpoint separations. Accordingly, while step 4040 results in a time distribution of new setpoint effective times that are generally separated by at least one hour, some new setpoints having effective times separated by less than one hour may remain. These minor variances have been found to be tolerable, and often preferable to deleterious effects resulting from cascading the operation to achieve absolute one-hour separations. Furthermore, these one-hour separations can be successfully completed later in the algorithm, after processing against the pre-existing schedule setpoints. Other separation intervals may be used in alternative implementations.

Referring to step 4042 of Figure 40B, consistent with the aggressive purposes associated with initial learning, the new setpoints that have now been established for the current learning day are next replicated across other days of the week that may be expected to have similar setpoints, when those new setpoints have been tagged as “RT” setpoints. Preferably, new setpoints tagged as “NRT” are not replicated, since it is likely that the user who created the underlying NRT setpoint entry has already created similar desired NRT setpoint entries. For some implementations that have been found to be well suited for the creation of a weekly schedule, a predetermined set of replication rules is applied. These replication rules depend on which day of the week the initial learning process was first started. The replication rules are optimized to take into account the practical schedules of a large population of expected users, for which weekends are often differently structured than weekdays, while, at the same time, promoting aggressive initial-schedule establishment. For one implementation, the replication rules set forth in Table 1 are applicable.

TABLE 1

If the First Initial Learning Day was . . .	And the Current Learning Day is	Then Replicate New Setpoints Onto . . .
Any Day Mon-Thu	Any Day Mon-Fri	All Days Mon-Fri
	Sat or Sun	Sat and Sun
Friday	Fri	All 7 Days
	Sat or Sun	Sat and Sun
	Any Day Mon-Thu	All Days Mon-Fri

Saturday	Sat or Sun	Sat and Sun
	Any Day Mon-Fri	All Days Mon-Fri
Sunday	Sun	All 7 Days
	Mon or Tue	All 7 Days
	Any Day Wed-Fri	All Days Mon-Fri
	Sat	Sat and Sun

Figure 41F illustrates effects of the replication of the RT-tagged new setpoints of Figure 41E, from a Tuesday monitoring period, onto the displayed portions of the neighboring days Monday and Wednesday. Thus, for example, the RT-tagged new setpoint “x,” having an effective time of 11:00 PM, is replicated as new setpoint “x2” on Monday, and all other weekdays, and the RT-tagged new setpoint “ij,” having an effective time of 7:30 AM, is replicated as new setpoint “ij2” on Wednesday and all other weekdays. As per the rules of Table 1, all of the other RT-tagged new setpoints, including “mno,” “p,” “q,” and “u,” are also replicated across all other weekdays. Neither of the NRT-tagged new setpoints “kl” or “rst” is replicated. The NRT user setpoint entry “h,” which was entered on Tuesday by a user who desired it to be effective on Mondays, is not replicated.

Referring now to step 4044 of Figure 40B, the new setpoints and replicated new setpoints are overlaid onto the current schedule of pre-existing setpoints, as illustrated in Figure 41G, which shows the pre-existing setpoints encircled and the new setpoints not encircled. In many of the subsequent steps, the RT-tagged and NRT-tagged new setpoints are treated the same, and, when so, the “RT” and “NRT” labels are not used in describing such steps. In step 4046, a mutual-filtering and/or time-shifting of the new and pre-existing setpoints is carried out according to predetermined filtering rules that are designed to optimally or near optimally capture the pattern information and preference information, while also simplifying overall schedule complexity. While a variety of different approaches can be used, one method for carrying out the objective of step 4046 is described, in greater detail, in Figure 40D. Finally, in step 4048, the results of step 4046 become the newest version of the current schedule that is either further modified by another initial learning day or that is used as the starting schedule in the steady-state learning process.

- 63 -

Referring to Figure 40D, which sets forth one method for carrying out the processing of step 4046 of Figure 40C, a first type of any new setpoint having an effective time that is less than one hour later than that of a first pre-existing setpoint and less than one hour earlier than that of a second pre-existing setpoint is identified in step 4080. Examples of such new setpoints of the first type are circled in dotted lines in Figure 41G. The steps of Figure 40D are carried out for the entire weeklong schedule, even though only a portion of that schedule is shown in Figure 41G, for explanatory purposes. In step 4081, any new setpoints of the first type are deleted when they have effective times less than one hour earlier than the immediately subsequent pre-existing setpoint and when they have a temperature value that is not more than one degree F away from that of the immediately preceding pre-existing setpoint. For purposes of step 4081 and other steps in which a nearness or similarity evaluation between the temperature values of two setpoints is undertaken, the comparison of the setpoint values is carried out with respect to rounded versions of their respective temperature values, the rounding being to the nearest one degree F or to the nearest 0.5 degree C, even though the temperature values of the setpoints may be maintained to a precision of 0.2° F or 0.1° C for other operational purposes. When using rounding, for example, two setpoint temperatures of 77.6 ° F and 79.4 ° F are considered as 1 degree F apart when each is first rounded to the nearest degree F, and therefore not greater than 1 degree F apart. Likewise, two setpoint temperatures of 20.8° C and 21.7° C will be considered as 0.5 degree C apart when each is first rounded to the nearest 0.5 degree C, and therefore not greater than 0.5 degree C apart. When applied to the example scenario at Figure 41G, new setpoint “ij” falls within the purview of the rule in step 4081, and that new setpoint “ij” is thus deleted, as shown in Figure 41H.

Subsequent to the deletion of any new setpoints of the first type in step 4081, any new setpoint of the first type that has an effective time that is within 30 minutes of the immediately subsequent pre-existing setpoint is identified in step 4082. When such first-type setpoints are identified, they are moved, later in time, to one hour later than the immediately preceding pre-existing setpoint, and the immediately subsequent pre-existing setpoint is deleted. When applied to the example scenario at Figure 41G, new setpoint “ij2” falls within the purview of the rule in step 4082 and new setpoint “ij2” is therefore moved,

- 64 -

later in time, to one hour from the earlier pre-existing setpoint “f,” with the subsequent pre-existing setpoint “g” deleted, as shown in Figure 41H. Subsequently, in step 4084, any new setpoint of the first type that has an effective time that is within 30 minutes of the immediately preceding pre-existing setpoint there is identified. When such a first-type setpoint is identified, the setpoint is moved, earlier in time, to one hour earlier than the immediately subsequent pre-existing setpoint and the immediately preceding pre-existing setpoint is deleted. In step 4086, for each remaining new setpoint of the first type that is not subject to the purview of steps 4082 or 4084, the setpoint temperature of the immediately preceding pre-existing setpoint is changed to that of the new setpoint and that new setpoint is deleted.

In step 4087, any RT-tagged new setpoint that is within one hour of an immediately subsequent pre-existing setpoint and that has a temperature value not greater than one degree F different from an immediately preceding pre-existing setpoint is identified and deleted. In step 4088, for each new setpoint, any pre-existing setpoint that is within one hour of that new setpoint is deleted. Thus, for example, Figure 41I shows a pre-existing setpoint “a” that is less than one hour away from the new setpoint “x2,” and so the pre-existing setpoint “a” is deleted, in Figure 41J. Likewise, the pre-existing setpoint “d” is less than one hour away from the new setpoint “q,” and so the pre-existing setpoint “d” is deleted, in Figure 41J.

In step 4090, starting from the earliest effective setpoint time in the schedule and moving later in time to the latest effective setpoint time, a setpoint is deleted when the setpoint has a temperature value that differs by not more than 1 degree F or 0.5 degree C from that of the immediately preceding setpoint. As discussed above, anchor setpoints, in many implementations, are not deleted or adjusted as a result of automatic schedule learning. For example, Figure 41K shows the setpoints “mno” and “x” that are each not more than one degree F from immediately preceding setpoints, and so setpoints “mno” and “x” are deleted, in Figure 41L. Finally, in step 4092, when there are any remaining pairs of setpoints, new or pre-existing, having effective times that are less than one hour apart, the later effective setpoint of each pair is deleted. The surviving setpoints are then established as members of the current schedule, as indicated in Figure 41M, all of which are labeled

- 65 -

“pre-existing setpoints” for subsequent iterations of the initial learning process of Figure 40A or, when that process is complete, for subsequent application of steady-state learning, described below. Of course, the various time intervals for invoking the above-discussed clustering, resolving, filtering, and shifting operations may vary, in alternative implementations.

Figures 42A-42B illustrate steps for steady-state learning. Many of the same concepts and teachings described above for the initial learning process are applicable to steady-state learning, including the tracking of real-time user setpoint entries and non-real time user setpoint entries, clustering, resolving, replicating, overlaying, and final filtering and shifting.

Certain differences arise between initial and steady state learning, in that, for the steady-state learning process, there is an attention to the detection of historical patterns in the setpoint entries, an increased selectivity in the target days across which the detected setpoint patterns are replicated, and other differences. Referring to Figure 42A, the steady state learning process begins in step 4202, which can correspond to the completion of the initial learning process (Figure 40A, step 4016), and which can optionally correspond to a resumption of steady-state learning after a user-requested pause in learning. In step 4204, a suitable version of the current schedule is accessed. When the steady-state learning is being invoked immediately following initial learning, often be the case for a new intelligent-thermostat installation, the control schedule is generally the current schedule at the completion of initial learning.

However, a previously established schedule may be accessed in step 4204, in certain implementations. A plurality of different schedules that were previously built up by the intelligent thermostat 3302 over a similar period in the preceding year can be stored in the thermostat 3302, or, alternatively, in a cloud server to which it has a network connection. For example, there may be a “January” schedule that was built up over the preceding January and then stored to memory on January 31. When step 4204 is being carried out on January 1 of the following year, the previously stored “January” schedule can be accessed. In certain implementations, the intelligent thermostat 3302 may establish and store schedules that are applicable for any of a variety of time periods and then later access

- 66 -

those schedules, in step 4204, for use as the next current schedule. Similar storage and recall methods are applicable for the historical RT/NRT setpoint entry databases that are discussed further below.

In step 4206, a new day of steady-state learning is begun. In step 4208, throughout the day, the intelligent thermostat receives and tracks both real-time and non-real time user setpoint entries. In step 4210, throughout the day, the intelligent thermostat proceeds to control an HVAC system according to the current version of the schedule, whatever RT setpoint entries are made by the user, and whatever NRT setpoint entries have been made that are causally applicable.

According to one optional alternative embodiment, step 4210 can be carried out so that any RT setpoint entry is effective only for a maximum of 4 hours, after which the operating setpoint temperature is returned to whatever temperature is specified by the pre-existing setpoints in the current schedule and/or whatever temperature is specified by any causally applicable NRT setpoint entries. As another alternative, instead of reverting to any pre-existing setpoints after 4 hours, the operating setpoint instead reverts to a relatively low energy value, such as a lowest pre-existing setpoint in the schedule. This low-energy bias operation can be initiated according to a user-settable mode of operation.

At the end of the steady-state learning day, such as at or around midnight, processing steps 4212-4216 are carried out. In step 4212, a historical database of RT and NRT user setpoint entries, which may extend back at least two weeks, is accessed. In step 4214, the day's tracked RT/NRT setpoint entries are processed in conjunction with the historical database of RT/NRT setpoint entries and the pre-existing setpoints in the current schedule to generate a modified version of the current schedule, using steps that are described further below with respect to Figure 42B. In step 4216, the day's tracked RT/NRT setpoint entries are then added to the historical database for subsequent use in the next iteration of the method. Notably, in step 4218, whether there should be a substitution of the current schedule to something that is more appropriate and/or preferable is determined, such as for a change of season, a change of month, or another such change. When a schedule change is determined to be appropriate, a suitable schedule is accessed in step 4204, before the next iteration. Otherwise, the next iteration is begun in step 4206

using the most recently computed schedule. In certain implementations, step 4218 is carried out based on direct user instruction, remote instruction from an automated program running on an associated cloud server, remote instruction from a utility company, automatically based on the present date and/or current/forecasted weather trends, or based on a combination of one or more of the above criteria or other criteria.

Referring to Figure 42B, which corresponds to step 4214 of Figure 42A, steps similar to those of steps 4030-4040 of Figure 40B are carried out in order to cluster, resolve, tag, and adjust the day's tracked RT/NRT setpoint entries and historical RT/NRT setpoint entries. In step 4232, all RT-tagged setpoints appearing in the results of step 4232 are identified as pattern-candidate setpoints. In step 4234, the current day's pattern-candidate setpoints are compared to historical pattern-candidate setpoints to detect patterns, such as day-wise or week-wise patterns, of similar effective times and similar setpoint temperatures. In step 4236, for any such patterns detected in step 4234 that include a current-day pattern-candidate setpoint, the current-day pattern-candidate setpoint is replicated across all other days in the schedule for which such pattern may be expected to be applicable. As an example, Table 2 illustrates one particularly useful set of pattern-matching rules and associated setpoint replication rules.

TABLE 2

If Today Was . . .	And the Detected Match is With . . .	Then Replicate The Matched Support Onto . . .
Tue	Yesterday	All Days Mon-Fri
	Last Tuesday	Tuesdays Only
Wed	Yesterday	All Days Mon-Fri
	Last Wednesday	Wednesdays Only
Thu	Yesterday	All Days Mon-Fri
	Last Thursday	Thursdays Only
Fri	Yesterday	All Days Mon-Fri
	Last Friday	Fridays Only
Sat	Yesterday	All 7 Days of Week
	Last Saturday	Saturdays Only
Sun	Yesterday	Saturdays and Sundays
	Last Sunday	Sundays Only
Mon	Yesterday	All 7 Days of Week
	Last Monday	Mondays Only

- 68 -

For one implementation, in carrying out step 4236, the replicated setpoints are assigned the same effective time of day, and the same temperature value, as the particular current day pattern-candidate setpoint for which a pattern is detected. In other implementations, the replicated setpoints can be assigned the effective time of day of the historical pattern-candidate setpoint that was involved in the match and/or the temperature value of that historical pattern-candidate setpoint. In still other implementations, the replicated setpoints can be assigned the average effective time of day of the current and historical pattern-candidate setpoints that were matched and/or the average temperature value of the current and historical pattern-candidate setpoints that were matched.

In step 4238, the resulting replicated schedule of new setpoints is overlaid onto the current schedule of pre-existing setpoints. Also, in step 4238, any NRT-tagged setpoints resulting from step 4230 are overlaid onto the current schedule of pre-existing setpoints. In step 4240, the overlaid new and pre-existing setpoints are then mutually filtered and/or shifted in effective time using methods similar to those discussed above for step 4046 of Figure 40B. The results are then established, in step 4242, as the newest version of the current schedule.

Although the present invention has been described in terms of particular examples, it is not intended that the invention be limited to these examples. Modifications within the spirit of the invention will be apparent to those skilled in the art. For example, as discussed above, automated control-schedule learning may be employed in a wide variety of different types of intelligent controllers in order to learn one or more schedules that may span period of time from milliseconds to years. Intelligent-controller logic may include logic-circuit implementations, firmware, and computer-instruction-based routine and program implementations, all of which may vary depending on the selected values of a wide variety of different implementation and design parameters, including programming language, modular organization, hardware platform, data structures, control structures, and many other such design and implementation parameters. As discussed above, the steady-state learning mode follows aggressive learning may include multiple different phases, with the intelligent controller generally becoming increasingly conservative, with regard to

- 69 -

schedule modification, with later phases. Automated-control-schedule learning may be carried out within an individual intelligent controller, may be carried out in distributed fashion among multiple controllers, may be carried out in distributed fashion among one or more intelligent controllers and remote computing facilities, and may be carried out primarily in remote computing facilities interconnected with intelligent controllers.

It is appreciated that the previous description of the disclosed examples is provided to enable any person skilled in the art to make or use the present disclosure. Various modifications to these examples will be readily apparent to those skilled in the art, and the generic principles defined herein may be applied to other examples without departing from the spirit or scope of the disclosure. Thus, the present disclosure is not intended to be limited to the examples shown herein but is to be accorded the widest scope consistent with the principles and novel features disclosed herein.

- 70 -

CLAIMS

1. An intelligent controller comprising:
 - a processor;
 - a memory;
 - a control schedule stored in the memory;
 - a schedule interface;
 - a control interface; andinstructions stored within the memory that, when executed by the processor,
 - receive immediate-control inputs through the control interface during a monitoring period and record the received immediate-control inputs in memory;
 - receive schedule changes through the schedule interface during the monitoring period and record received schedule changes of at least one type in the memory;
 - generate an updated monitoring-period schedule, after the monitoring period, based on the recorded immediate-control inputs, recorded schedule changes, and the control schedule;
 - substitute the updated monitoring-period schedule for a portion of the control schedule corresponding to the monitoring period; and
 - propagate the updated monitoring-period schedule to additional time periods within the control schedule.

2. The intelligent controller of claim 1
 - wherein the processor includes one or more of
 - one or more local processors, and
 - one or more remote processors; and
 - wherein the memory includes one or more of
 - one or more local electronic memories,
 - one or more local mass-storage devices,
 - one or more remote electronic memories, and
 - one or more remote mass-storage devices.

- 71 -

3. The intelligent controller of claim 1 wherein control schedule stored in the memory includes one or more setpoints, each setpoint associated with a time and one or more parameter values.

4. The intelligent controller of claim 1 wherein an immediate-control input is made through the control interface to specify one or more parameter values to be implemented by one or more intelligent-controller outputs.

5. The intelligent controller of claim 1 wherein a control-schedule change is input through the schedule interface to specify a control-schedule setpoint to be implemented by one or more intelligent-controller outputs when a current time is equal to the time associated with the schedule setpoint.

6. The intelligent controller of claim 1 wherein the instructions generate an updated monitoring-period schedule, after the monitoring period, based on the recorded immediate-control inputs, recorded schedule changes, and the control schedule by:

combining the recorded immediate-control inputs, recorded schedule changes, and setpoints of the control schedule to produce a provisional schedule;

clustering the immediate-control inputs, schedule changes, and setpoints of the provisional schedule into one or more clusters;

resolving the one or more clusters within the provisional schedule; and

generating the updated monitoring-period schedule from the resolved clusters of the provisional schedule.

7. The intelligent controller of claim 6 wherein combining the recorded immediate-control inputs, recorded schedule changes, and setpoints of the control schedule to produce the provisional schedule further comprises:

representing immediate-control inputs as time-associated values of one or more parameters within a portion of the control schedule corresponding to the monitoring period;

- 72 -

representing recorded control-schedule changes as time-associated values of one or more parameters, additionally associated with input times, within the portion of the control schedule corresponding to the monitoring period; and

considering the represented immediate-control inputs, retrospective control-schedule changes, and setpoints of the provisional schedule to be events.

8. The intelligent controller of claim 7 wherein clustering the immediate-control inputs, schedule changes, and setpoints within the provisional schedule into one or more clusters further comprises:

collecting together, as a separate cluster, each group of events that are each within a first threshold time interval of another event of the separate cluster; and

collecting, as a separate cluster, each event that is not within the first threshold time interval of another event.

9. The intelligent controller of claim 8 wherein resolving the one or more clusters within the provisional schedule further includes replacing each cluster of events with one of:

no event, when no control-change trend is discerned in the cluster;

one event, when a single-event control-change trend is discerned in the cluster; and

two events, when a two-event control-change trend is discerned in the cluster.

10. The intelligent controller of claim 1 wherein the instructions propagate the updated monitoring-period schedule to additional time periods within the control schedule by:

selecting, according to one or more rules, one or more additional time periods within the control schedule related to the time period of the updated monitoring-period schedule; and

to each of the selected time periods,

copying setpoints corresponding to immediate-control inputs of the provisional monitoring-period schedule onto a portion of the control schedule corresponding to the selected time period, and

- 73 -

resolving the portion of the control schedule corresponding to the selected time period.

11. The intelligent controller of claim 10 wherein resolving the portion of the control schedule corresponding to the selected time period further comprises:

applying, to each copied setpoint, one or more rules that result in one of
deleting the copied setpoint,
deleting an existing setpoint,
moving either the copied setpoint or another setpoint in time, and
no change to the copied setpoint.

12. The intelligent controller of claim 1 further comprising:

at least three phases of control-schedule learning implemented by instructions stored in the memory, including

an initial phase,
at least one aggressive-learning phase, and
one or more steady-state learning phases.

13. The intelligent controller of claim 12

wherein, during the aggressive-learning phase, learning parameters are provided values to encourage frequent, additional immediate-control inputs; and

wherein, during the steady-state learning phase, learning parameters are provided values to discourage frequent, additional immediate-control inputs.

14. The intelligent controller of claim 1 wherein the intelligent controller controls one or more of:

an HVAC unit;
a furnace;
an air conditioner;
a heat pump;

- 74 -

an irrigation system;
a pump;
a fan;
one or more light sources;
a machine;
a device;
an organization; and
a system.

15. The intelligent controller of claim 1 wherein the setpoints of the control schedule specify one or more of:

temperature;
flow rate of a liquid or gas;
rate of energy dissipation;
pressure;
current density;
voltage;
a machine setting;
a position of a machine component;
a computational state;
a mechanical state; and
process throughput.

16. An intelligent controller comprising:

a processor;
a memory;
a schedule interface;
a control interface; and
instructions stored within the memory that, when executed by the processor,

- 75 -

acquire an initial control schedule and store the initial control schedule in memory as a current control schedule,

during an aggressive-learning phase, modify the current control schedule based on immediate-control inputs received through the control interface and schedule changes received through the schedule interface, and

during a steady-state learning phase, modify the current control schedule based on immediate-control inputs received through the control interface, schedule changes received through the schedule interface, and historical control schedules.

17. The intelligent controller of claim 16 wherein, during the aggressive-learning phase and the steady-state learning phase, the intelligent controller

receives immediate-control inputs through the control interface during a monitoring period and records the received immediate-control inputs in memory;

receives schedule changes through the schedule interface during the monitoring period and records received schedule changes of at least one type in the memory;

generates an updated monitoring-period schedule, after the monitoring period, based on the recorded immediate-control inputs, recorded schedule changes, and the current control schedule;

substitutes the updated monitoring-period schedule for a portion of the current control schedule corresponding to the monitoring period; and

propagates the updated monitoring-period schedule to additional time periods within the current control schedule.

18. The intelligent controller of claim 17 wherein the intelligent controller generates an updated monitoring-period schedule, after the monitoring period, based on the recorded immediate-control inputs, recorded schedule changes, and the current control schedule by:

combining the recorded immediate-control inputs, recorded schedule changes, and setpoints of the current control schedule to produce a provisional schedule;

- 76 -

clustering the immediate-control inputs, schedule changes, and setpoints of the provisional schedule into one or more clusters;
resolving the one or more clusters within the provisional schedule; and
generating the updated monitoring-period schedule from the resolved clusters of the provisional schedule.

19. A method that automatically learns a control schedule implemented within an intelligent controller that includes one or more processors and one or more memories, the method comprising:

acquiring an initial control schedule and storing the initial control schedule in one of the one or more memories as a current control schedule;
during each of a set of successive monitoring periods,
receiving immediate-control inputs through a control interface and recording the received immediate-control inputs in one of the one or more memories, and
receiving schedule changes through a schedule interface and recording the received schedule changes of at least one type in the memory; and
after each of the set of successive monitoring periods,
generating an updated monitoring-period schedule based on the recorded immediate-control inputs, the recorded schedule changes, and the current control schedule,
substituting the updated monitoring-period schedule for a portion of the current control schedule corresponding to the monitoring period, and
propagating the updated monitoring-period schedule to additional time periods within the current control schedule.

20. The method of claim 19 further including:

acquiring an initial control schedule and storing the initial control schedule in one of the one or more memories as a current control schedule,
during an aggressive-learning phase, modifying the current control schedule based on immediate-control inputs received through a control interface and schedule changes received through a schedule interface, and

- 77 -

during a steady-state learning phase, modifying the current control schedule based on immediate-control inputs received through the control interface, schedule changes received through the schedule interface, and historical control schedules.




FIG. 1

2/82




FIG. 2

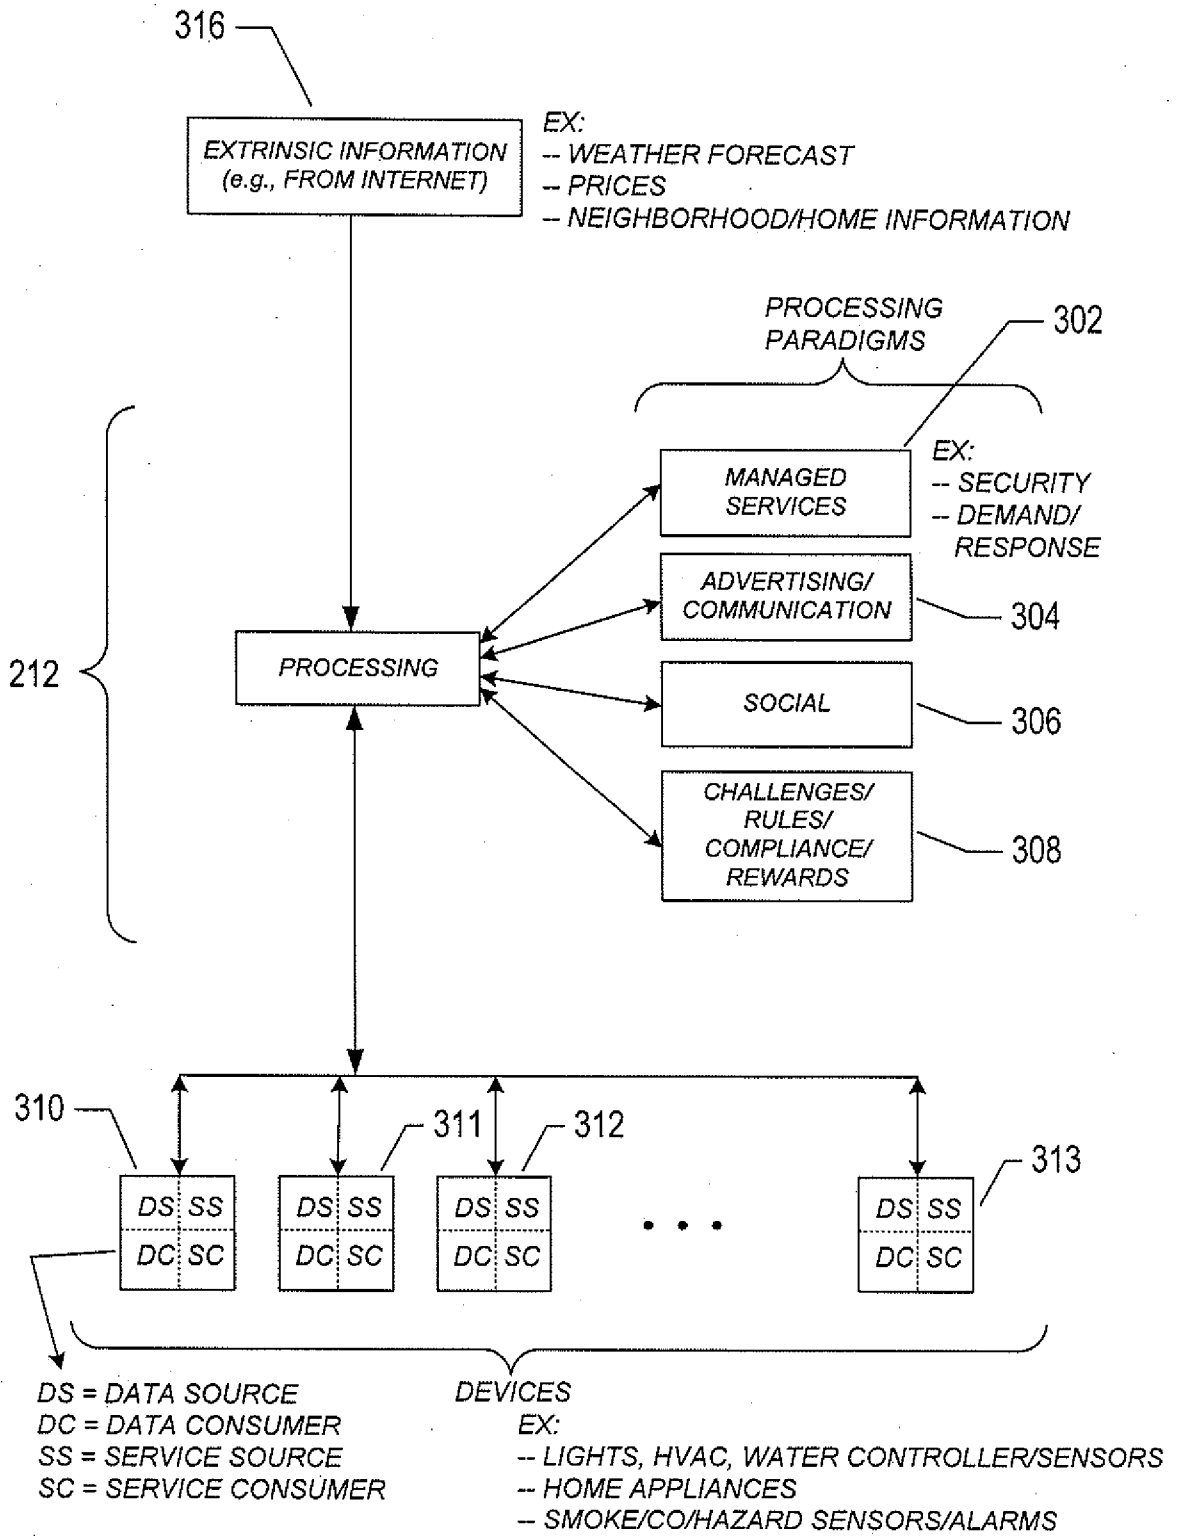


FIG. 3

4/82




FIG. 4




FIG. 5

6/82




FIG. 6

7/82




FIG. 7

8/82




FIG. 8




FIG. 9

10/82




FIG. 10A




FIG. 10B




FIG. 10C




FIG. 10D

11/82

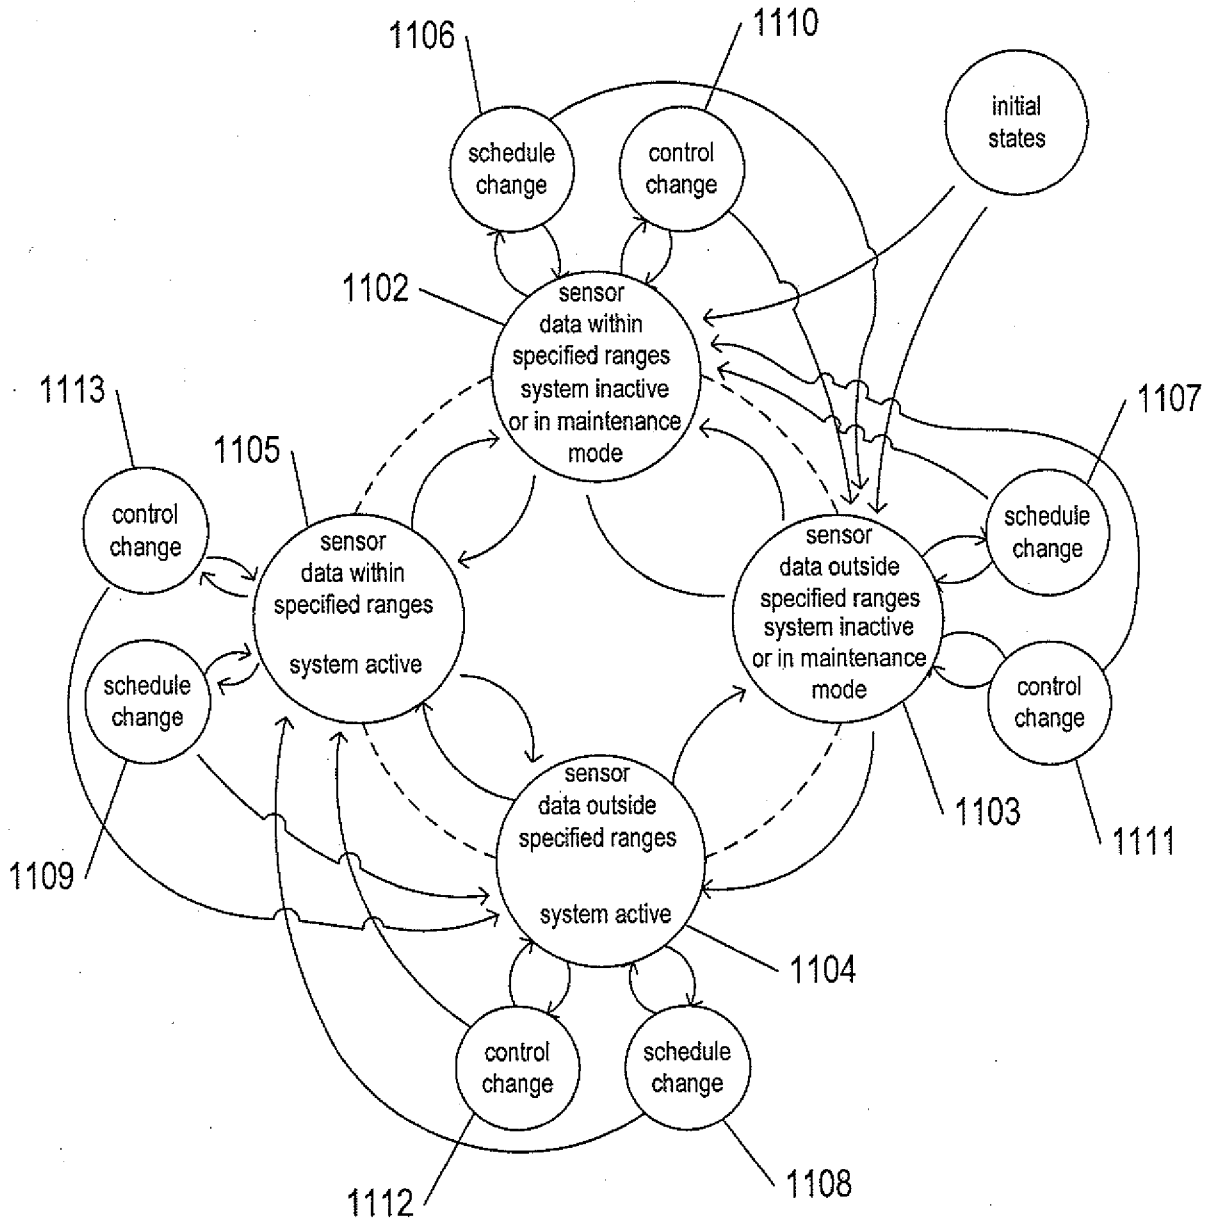


FIG. 11A

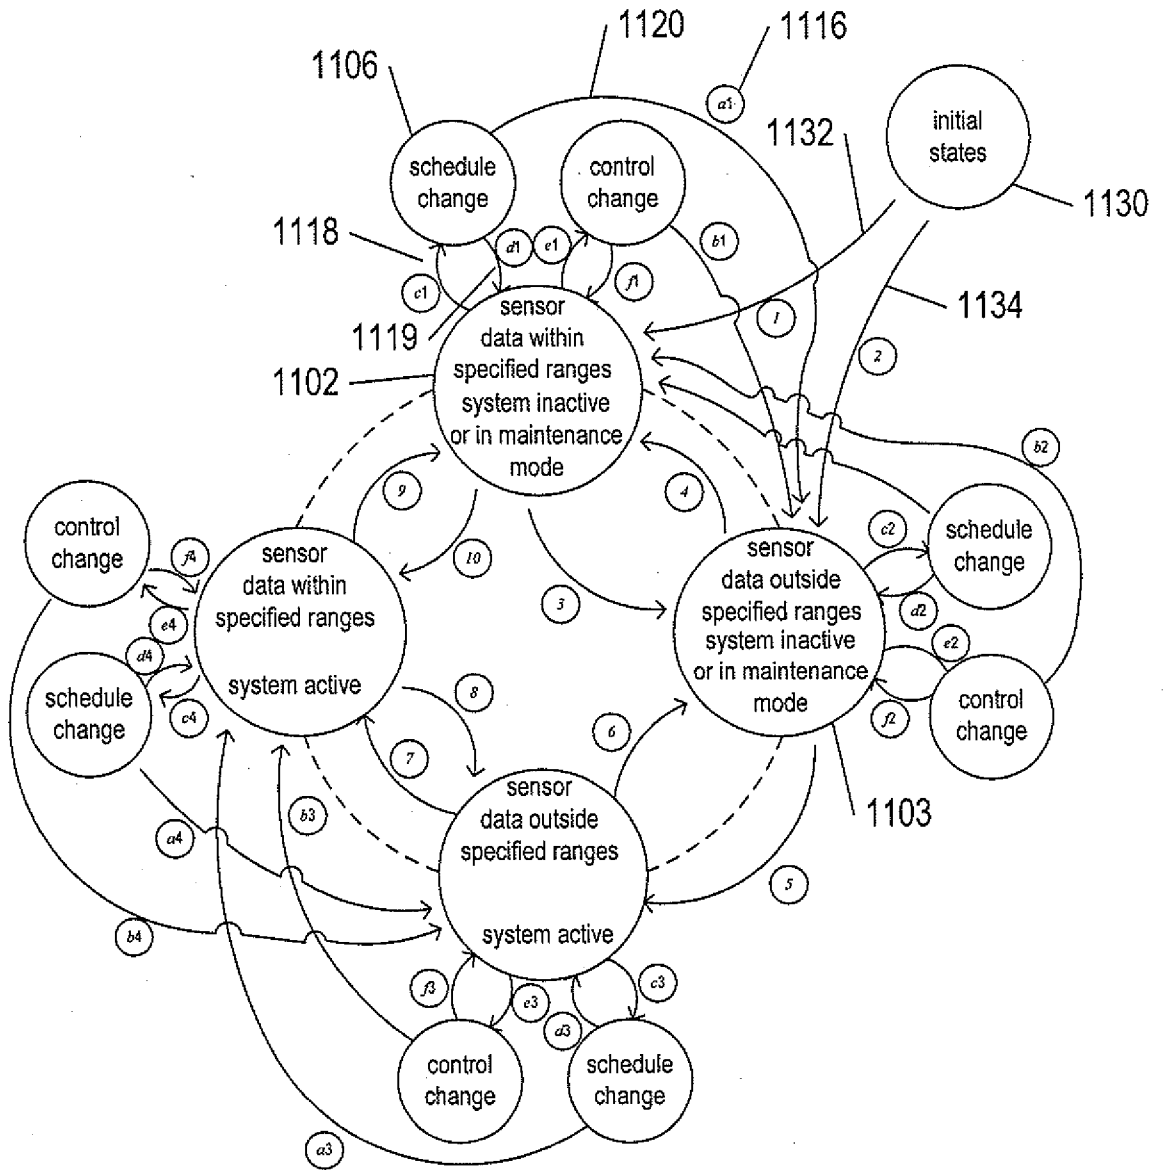


FIG. 11B

key	explanation
<i>a1, a4</i>	schedule change results in sensor data, formerly within specified ranges, to now be outside of the specified ranges
<i>a2, a3</i>	schedule change returns sensor data to within specified ranges
<i>b1, b4</i>	analogous to <i>a1, a4</i> as a result of immediate control
<i>b2, b3</i>	analogous to <i>a2, a3</i> as a result of immediate control
<i>c1, c2, c3, c4</i>	schedule changed by user or by system
<i>d1, d2, d3, d4</i>	return to former state following schedule change
<i>e1, e2, e3, e4</i>	user initiated immediate control
<i>f1, f2, f3, f4</i>	return to former state following schedule change
1,2	power-on or reset followed by initial configuration
3	change in environment monitored by sensors
4	change in environment monitored by sensors
5	controller outputs system-activation control
6	system failure or manual de-activation
7	system changes environment so that sensor data returns to within specified ranges
8	environmental changes despite system activation
9	controller detects return of sensor data to specified ranges and outputs control to deactivate system
10	manual activation of system or system failure

FIG. 11C

14/82

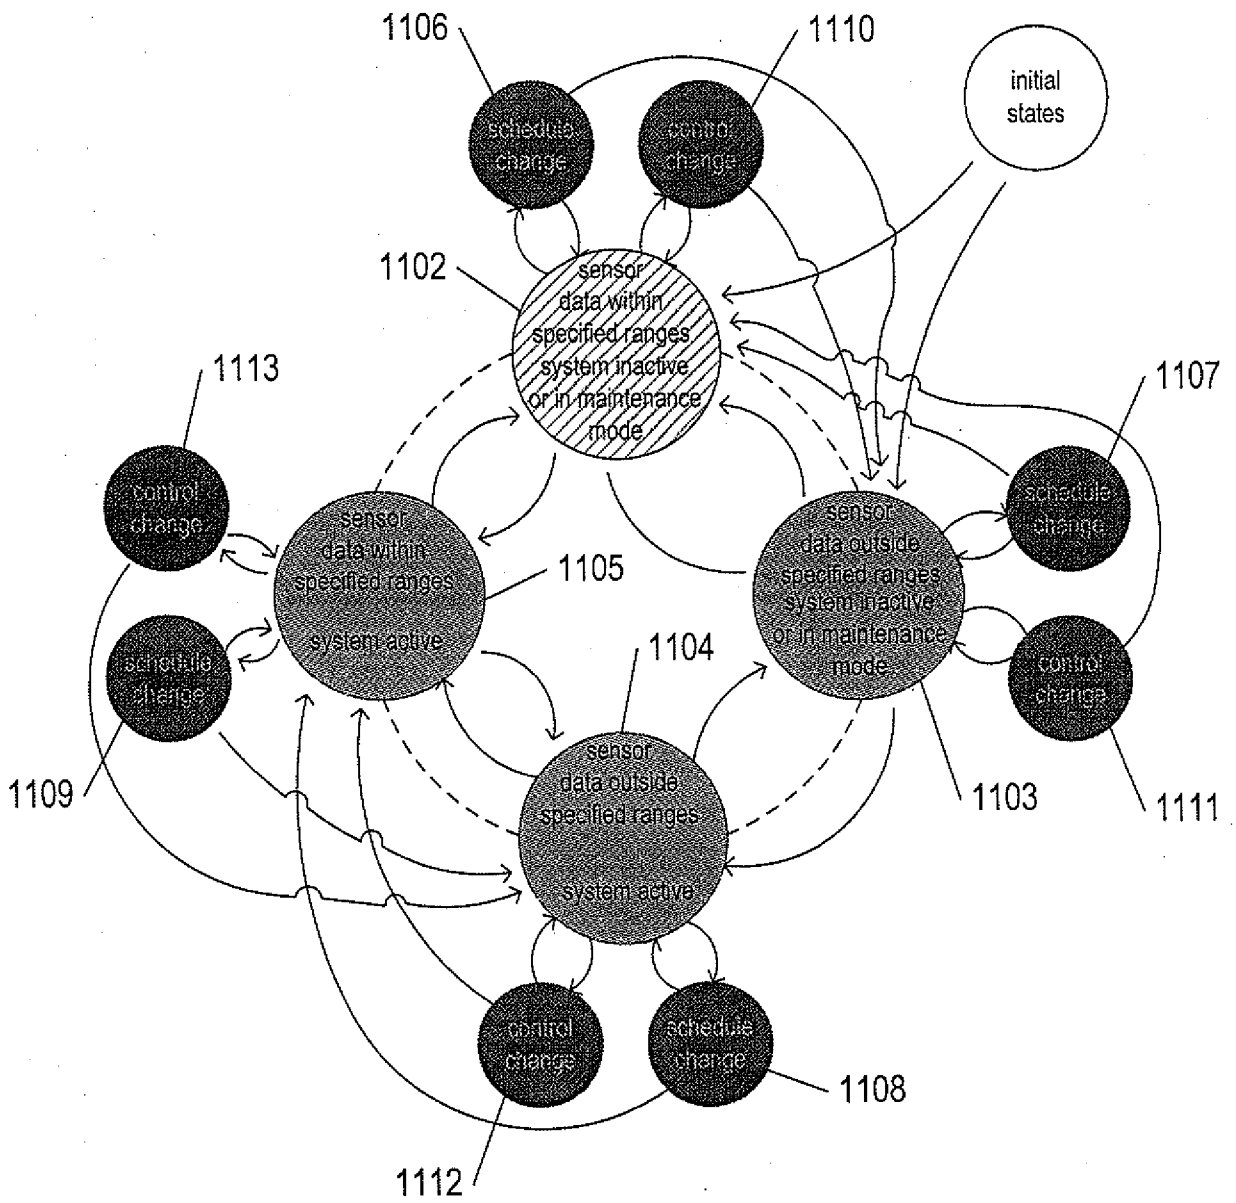


FIG. 11D

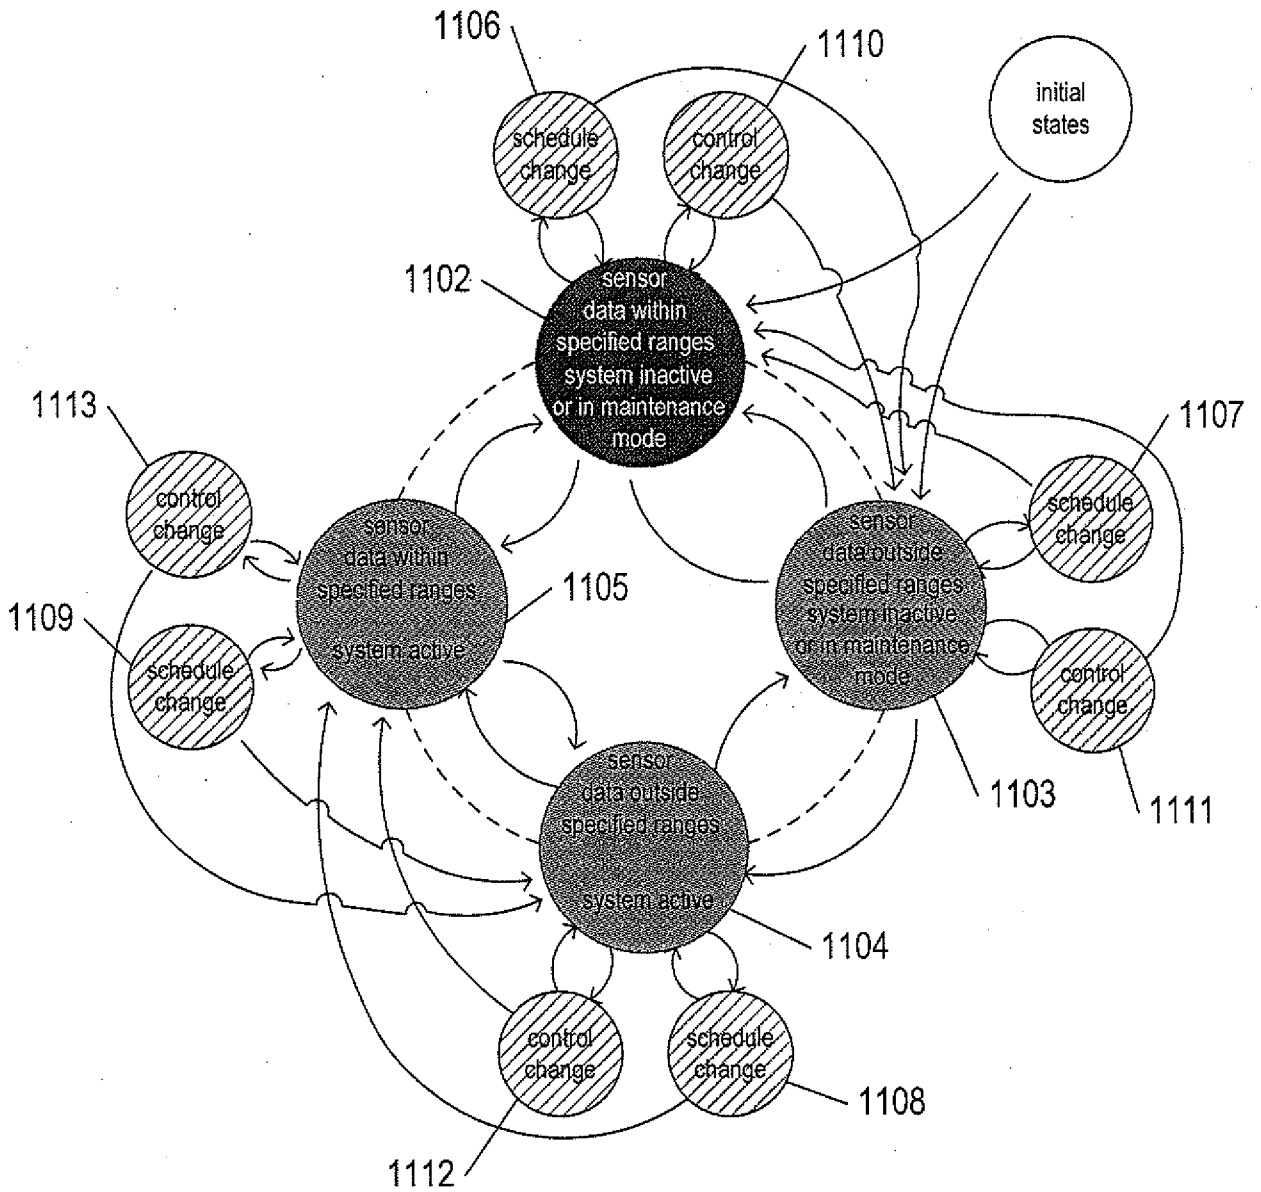


FIG. 11E

16/82




FIG. 12

17/82




FIG. 13

18/82




FIG. 14A




FIG. 14B




FIG. 14C

19/82




FIG. 15A




FIG. 15B

20/82




FIG. 15C




FIG. 15D




FIG. 15E

21/82




FIG. 15F




FIG. 15G

22/82




FIG. 16A




FIG. 16B




FIG. 16C




FIG. 16D




FIG. 16E




FIG. 17A




FIG. 17B



FIG. 17C




FIG. 17D




FIG. 17E

25/82




FIG. 18A




FIG. 18B




FIG. 19A




FIG. 19B




FIG. 20A




FIG. 20B

27/82





FIG. 20C



2108 FIG. 21A




FIG. 21B

28/82

2202




FIG. 22A

29/82




FIG. 22B

30/82




FIG. 23A




FIG. 23B




FIG. 23C

31/82




FIG. 24A





FIG. 24B


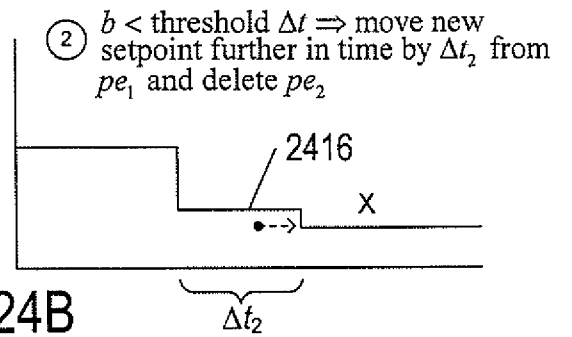


FIG. 24C


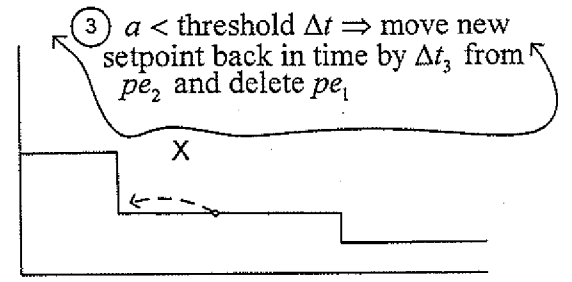


FIG. 24D



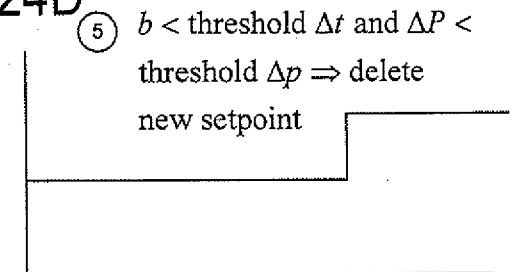


FIG. 24E



32/82




FIG. 24F




FIG. 24G




FIG. 24H




FIG. 24I

33/82

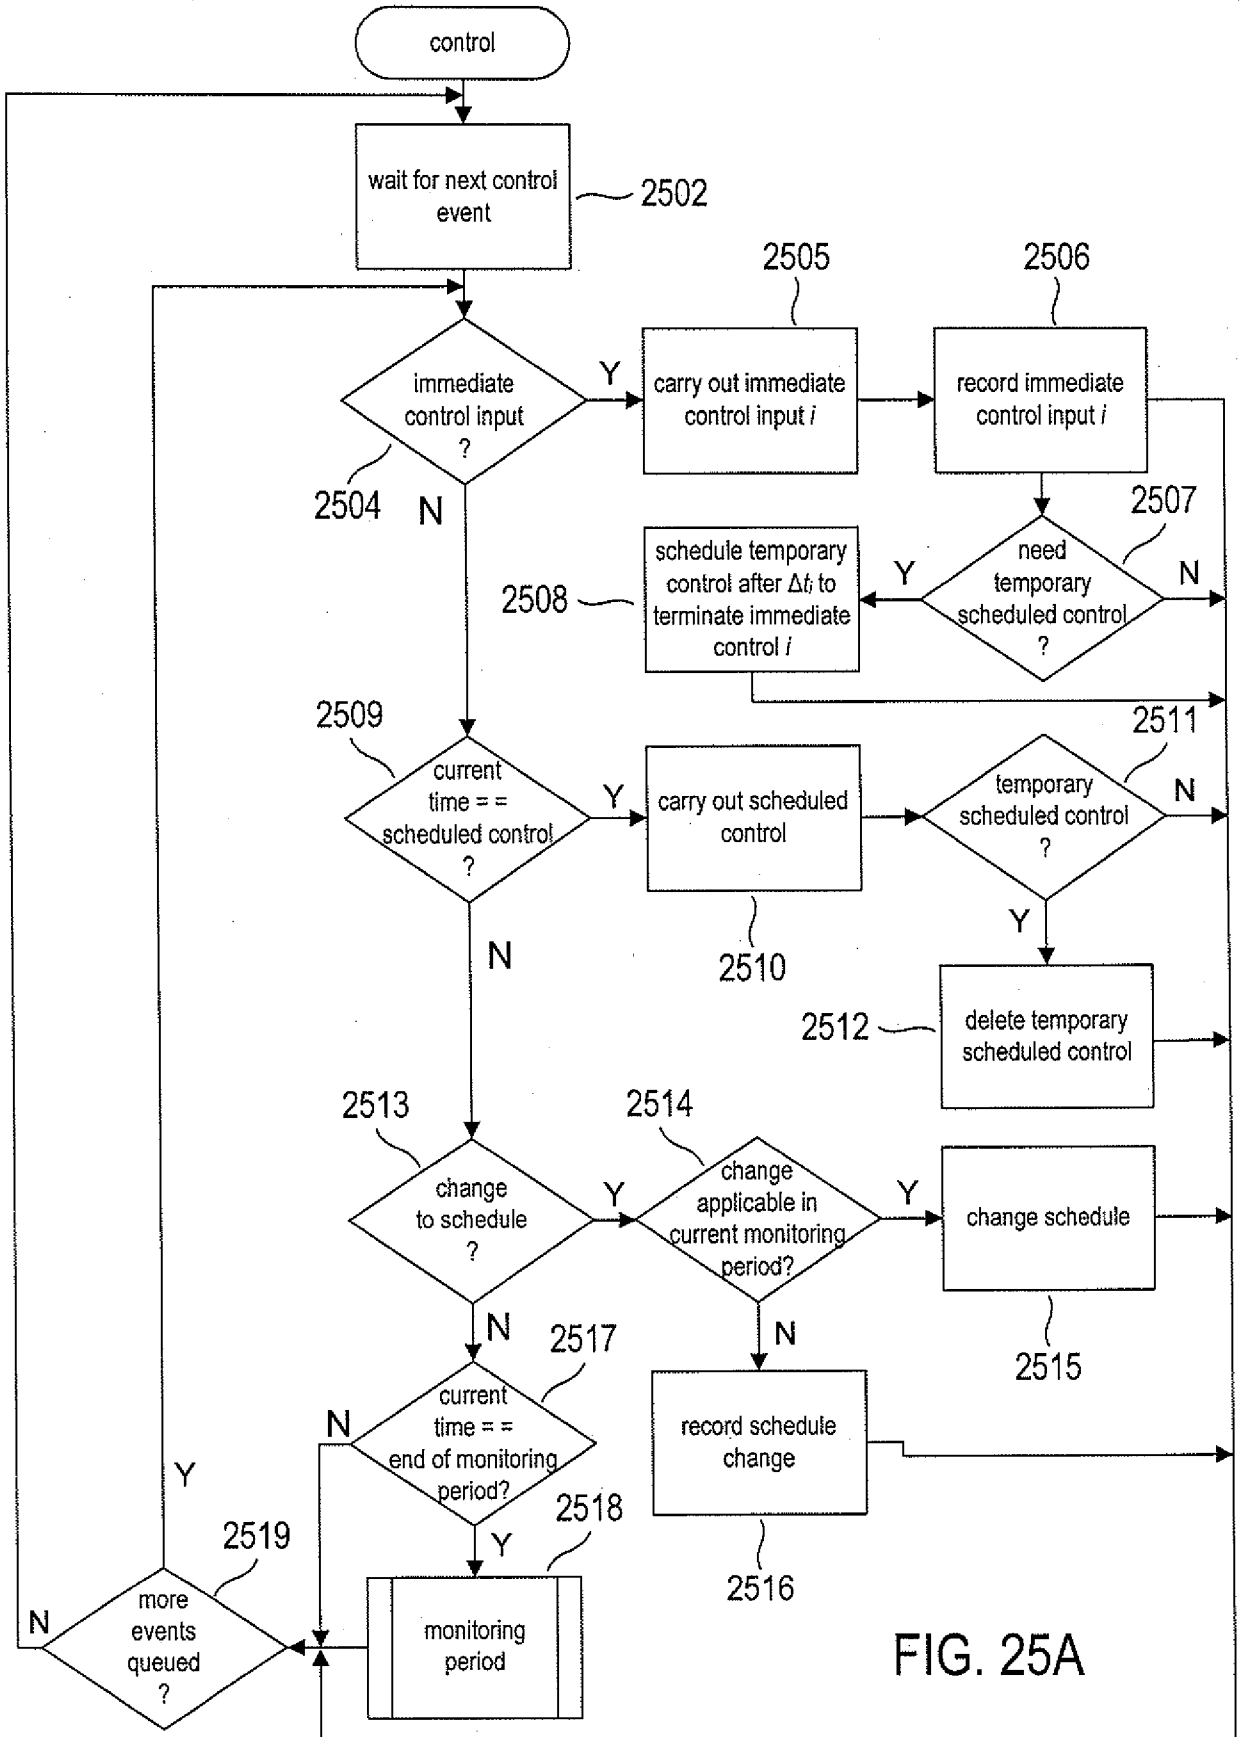


FIG. 25A

34/82

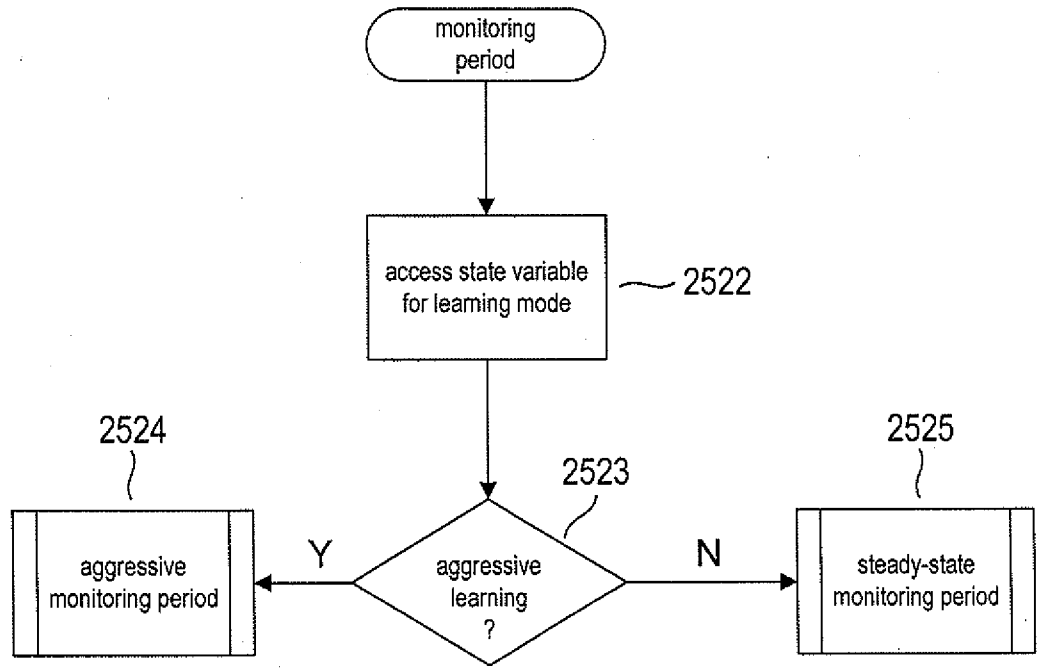


FIG. 25B

35/82

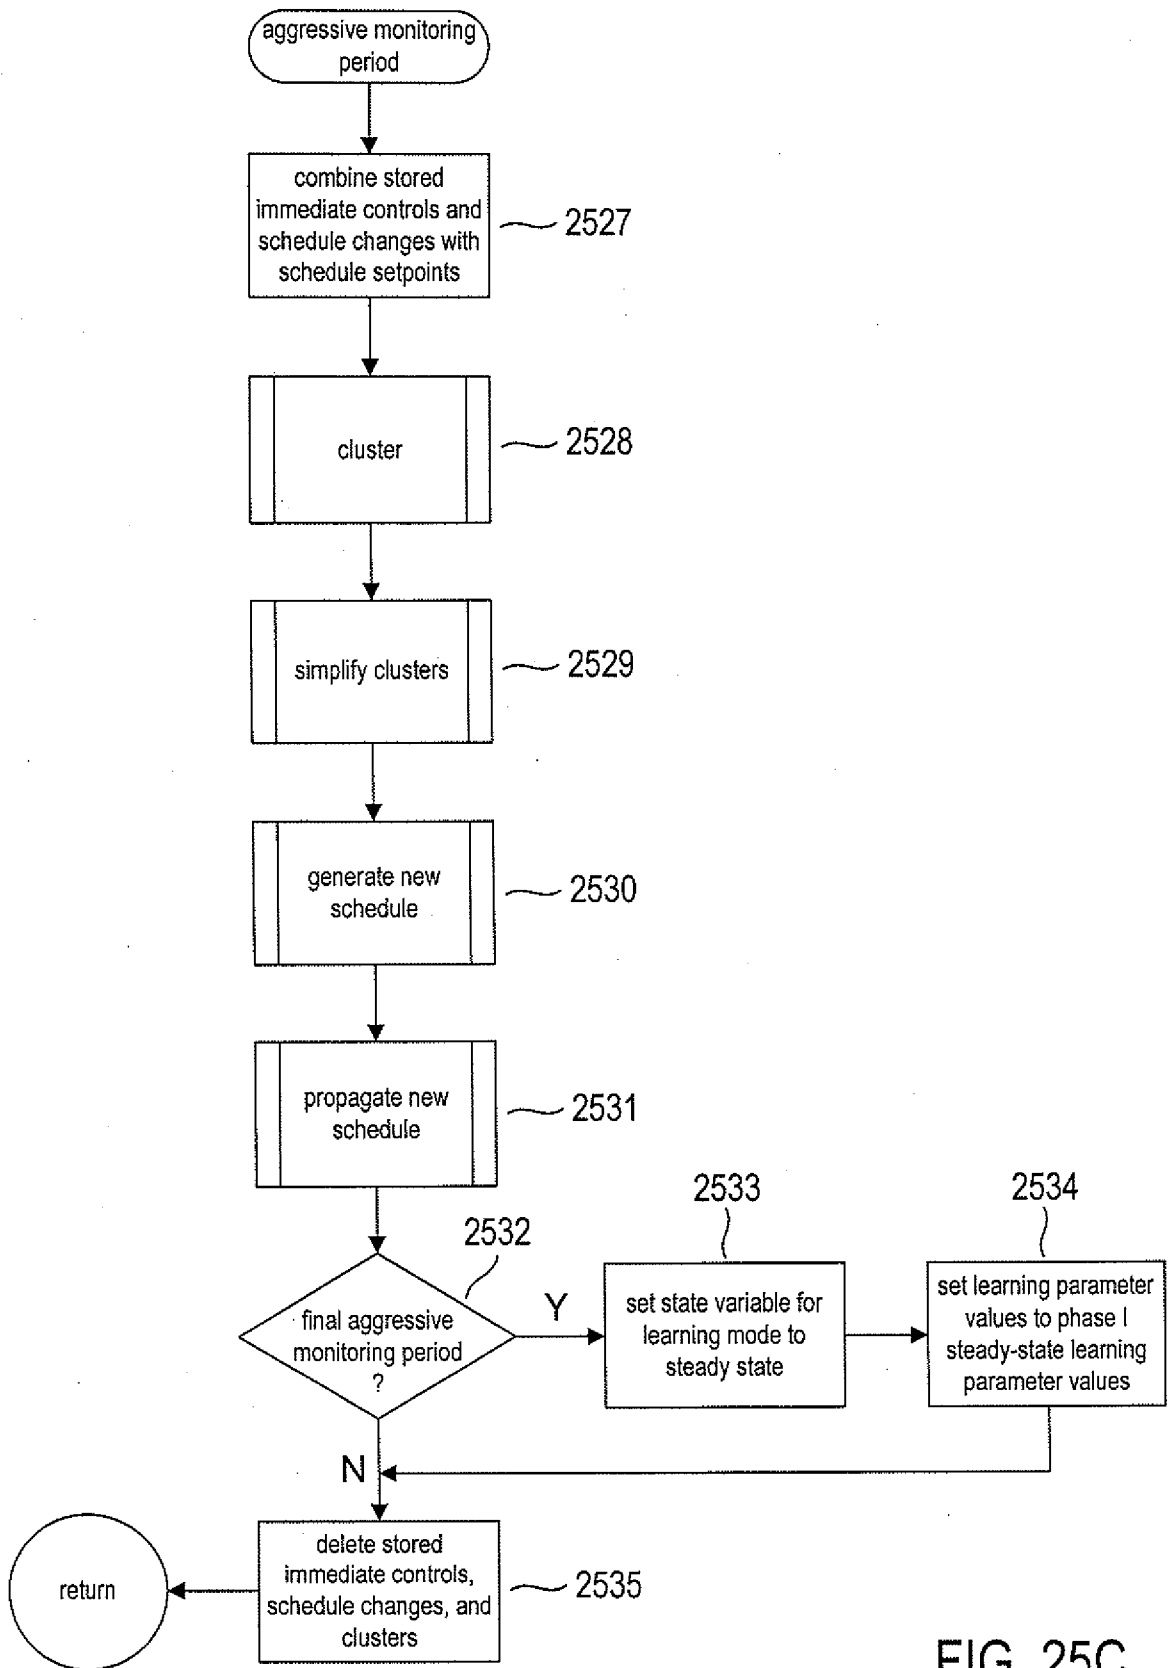


FIG. 25C

36/82




FIG. 25D

37/82




FIG. 25E

38/82

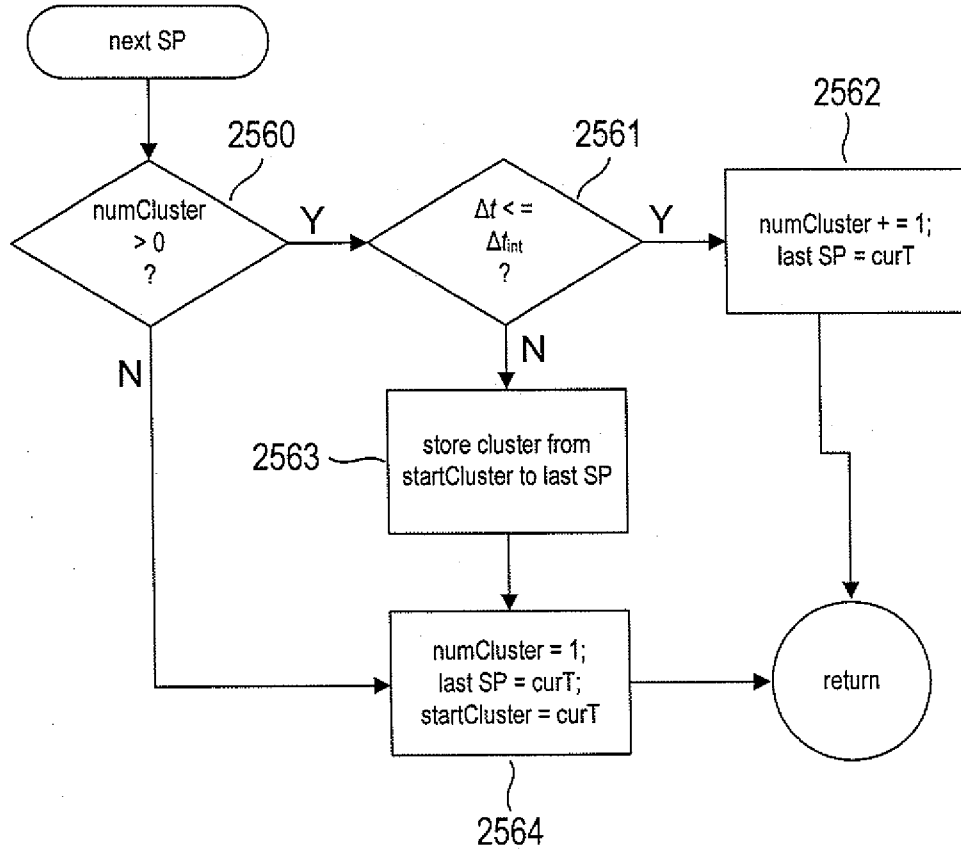


FIG. 25F

39/82




FIG. 25G

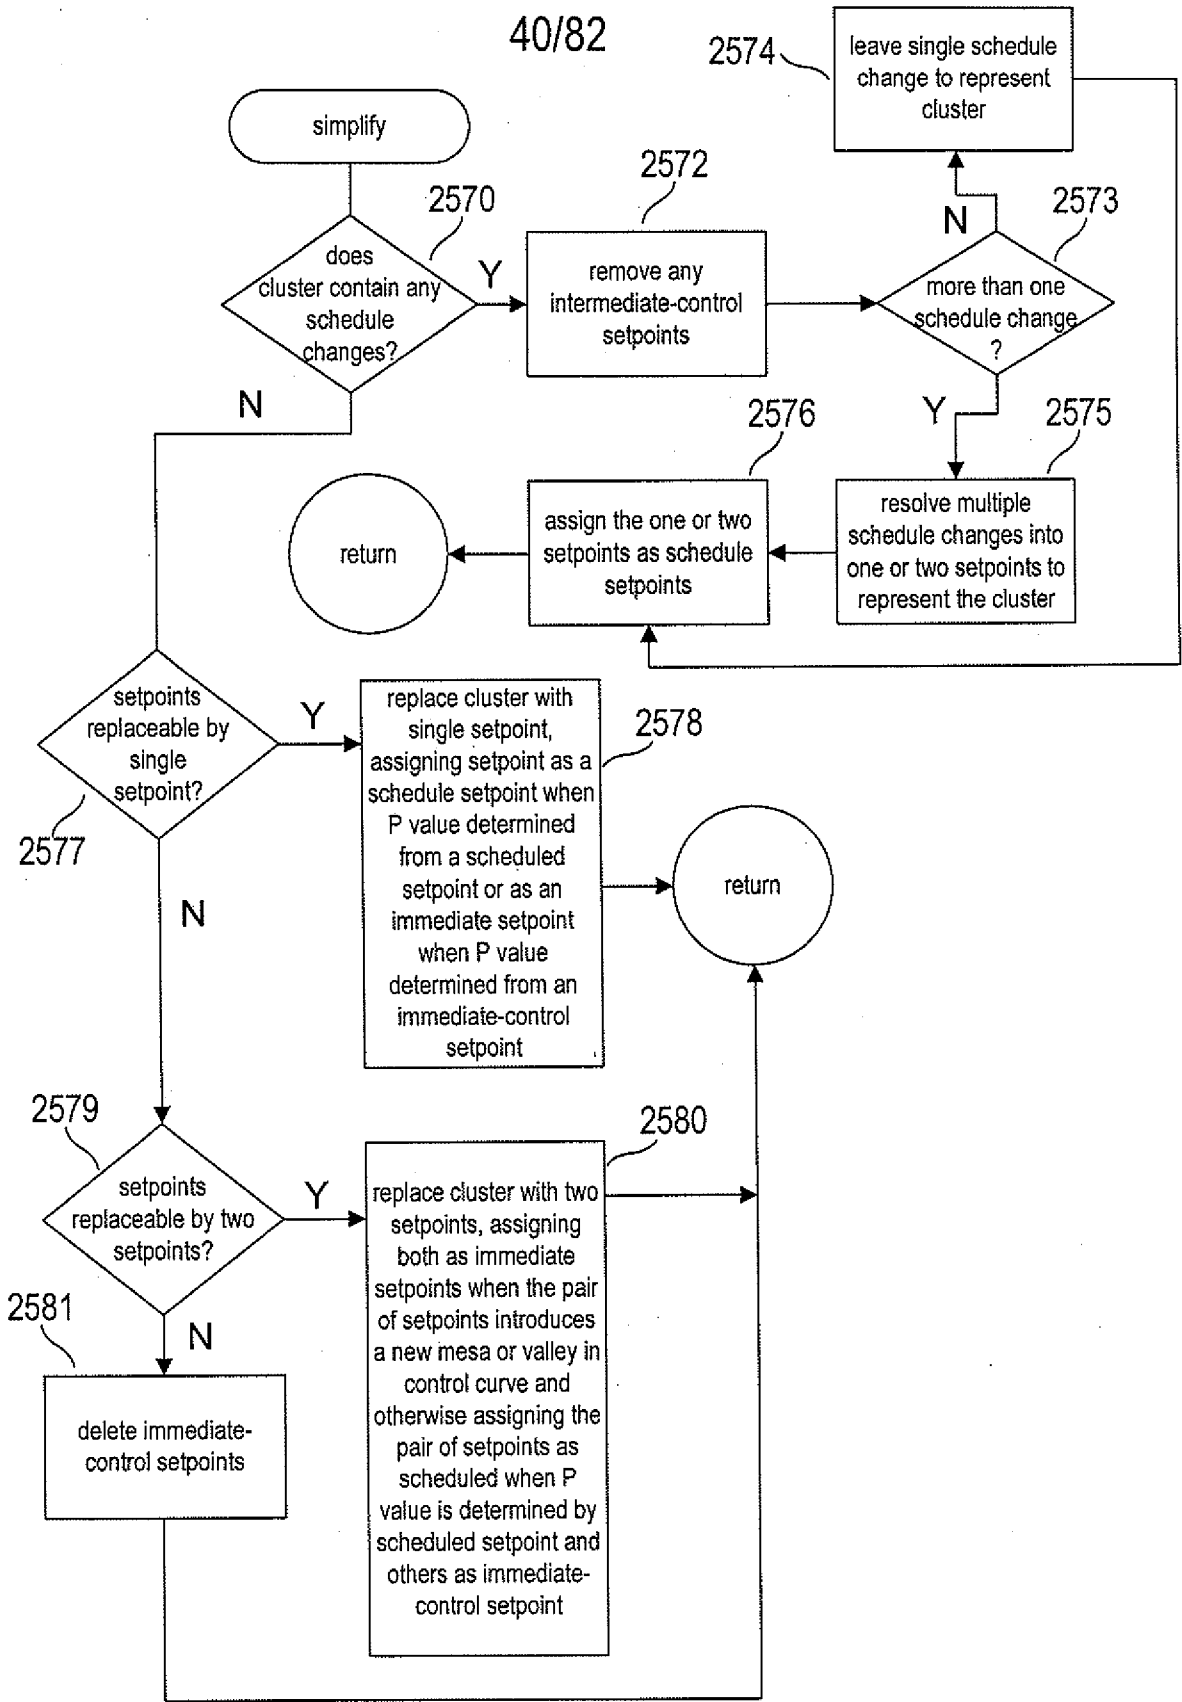


FIG. 25H

41/82

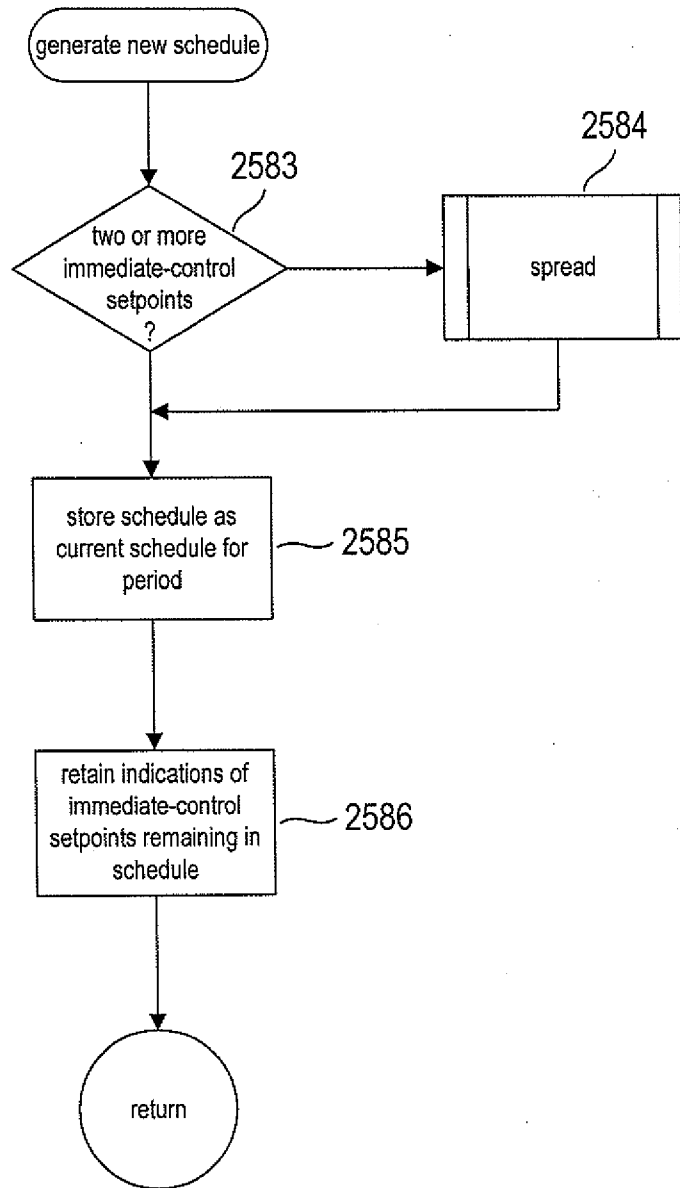


FIG. 25I

42/82

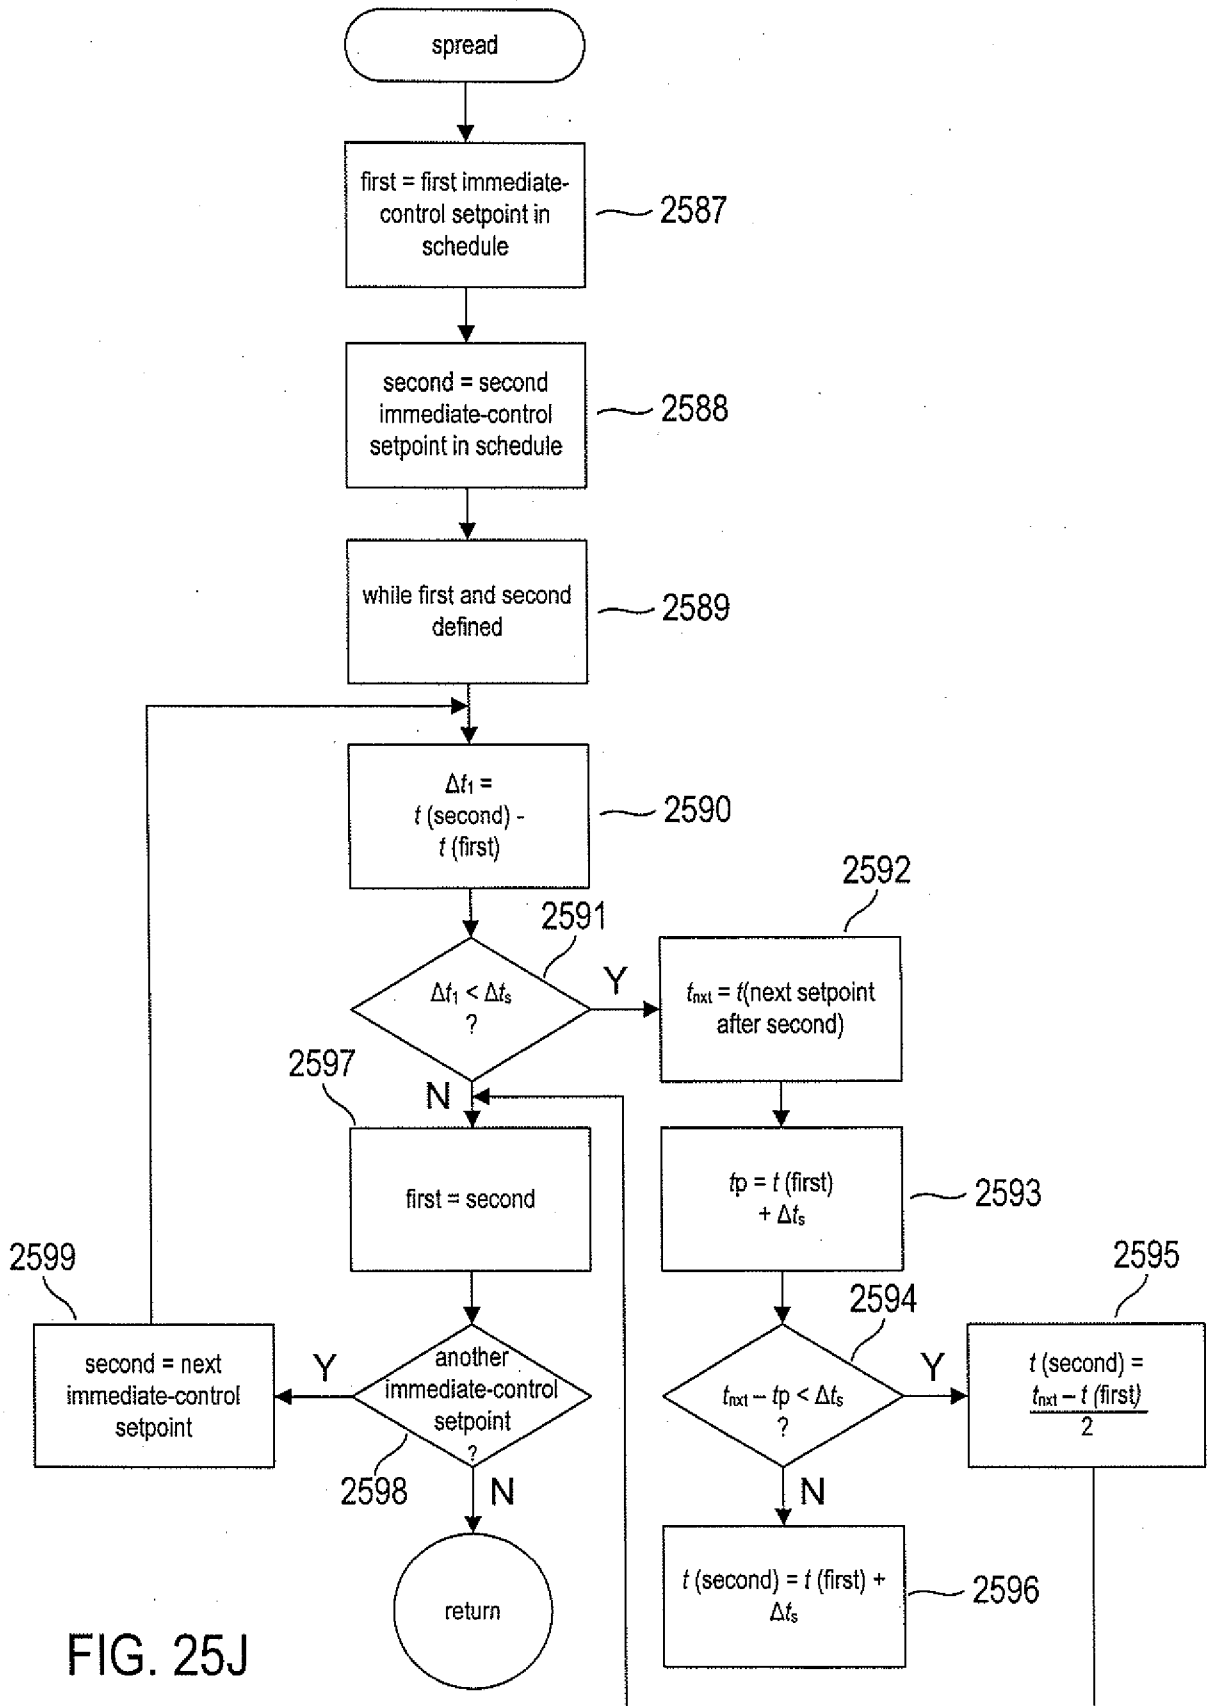


FIG. 25J

43/82

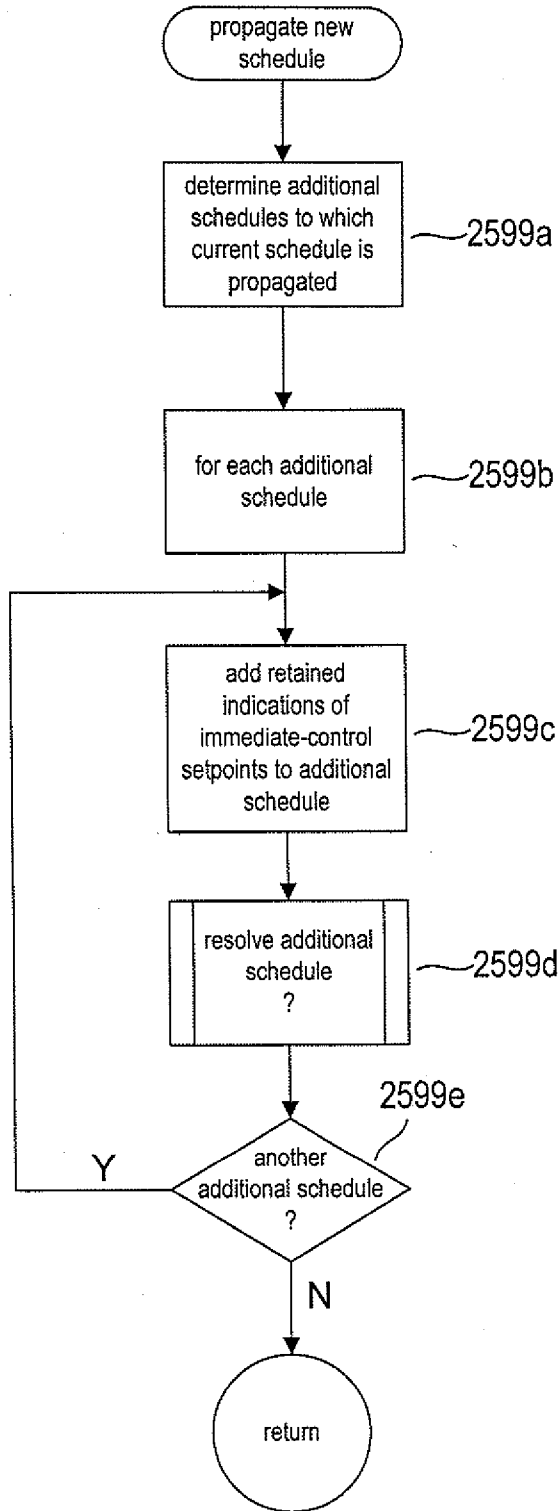


FIG. 25K

44/82

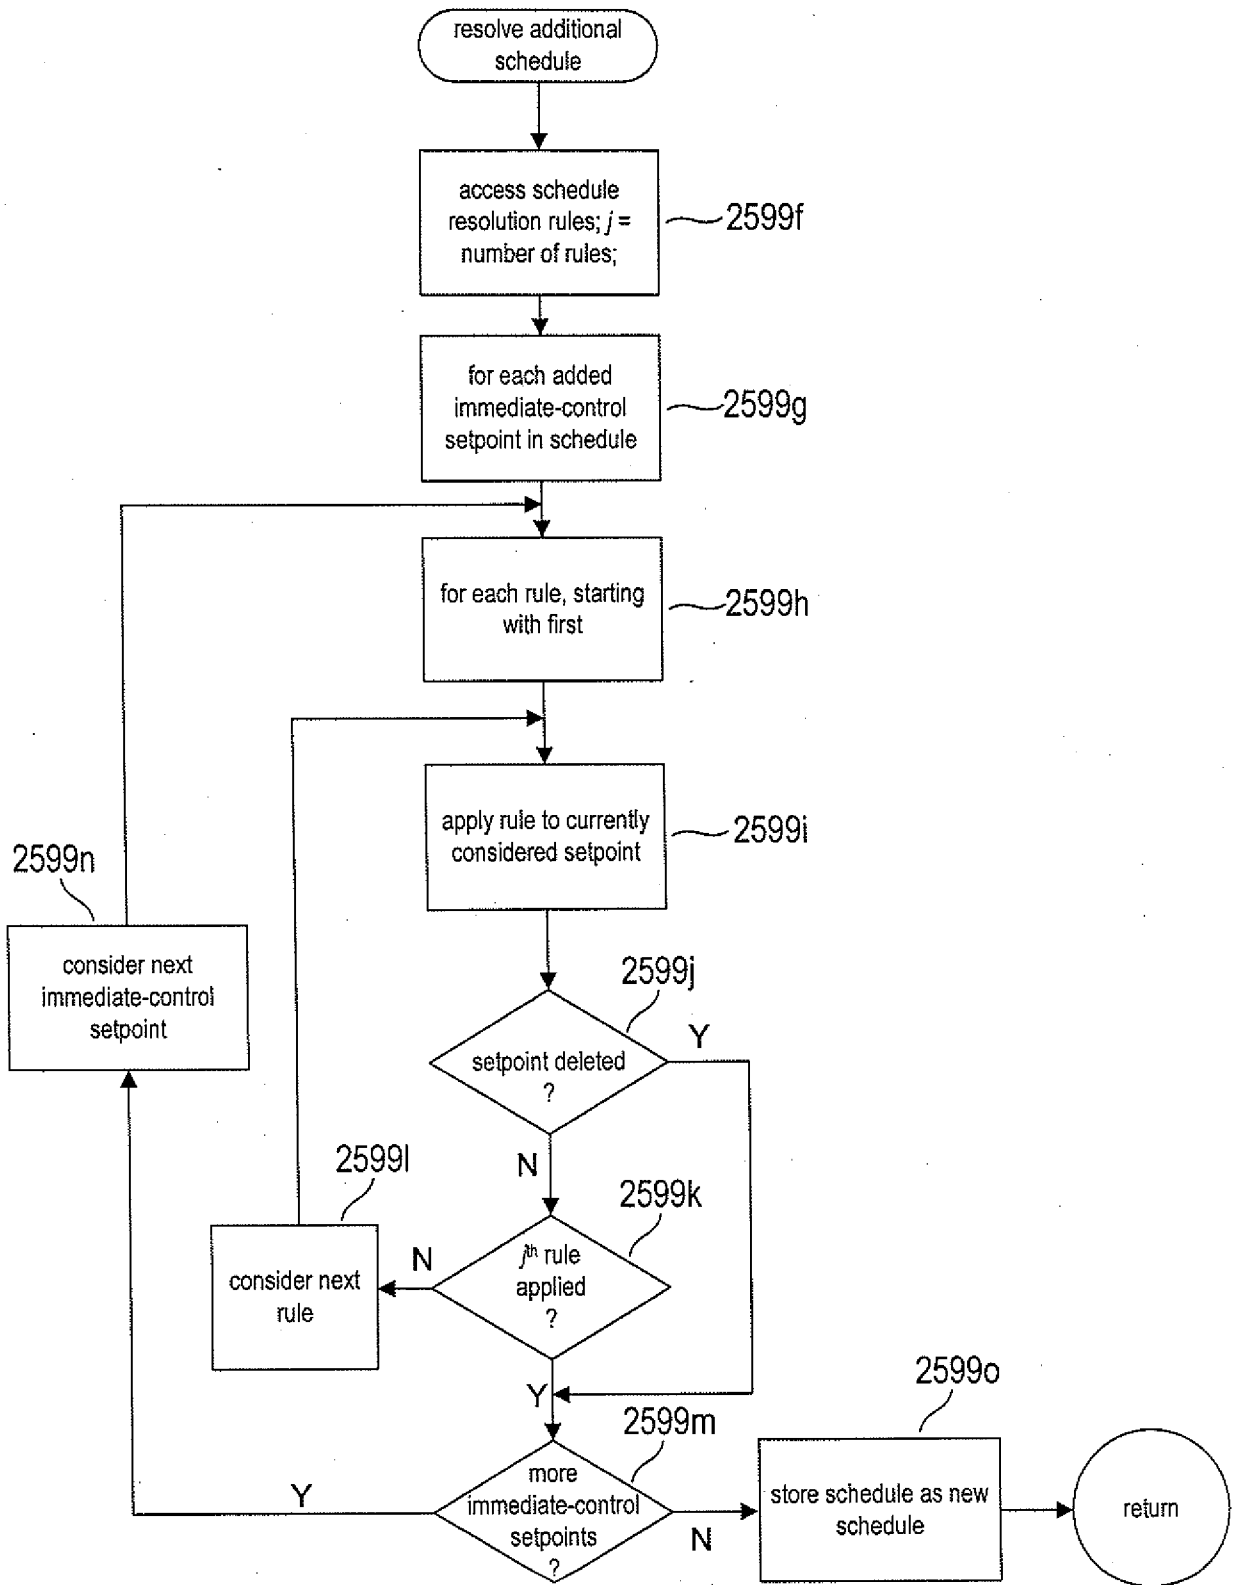


FIG. 25L

45/82

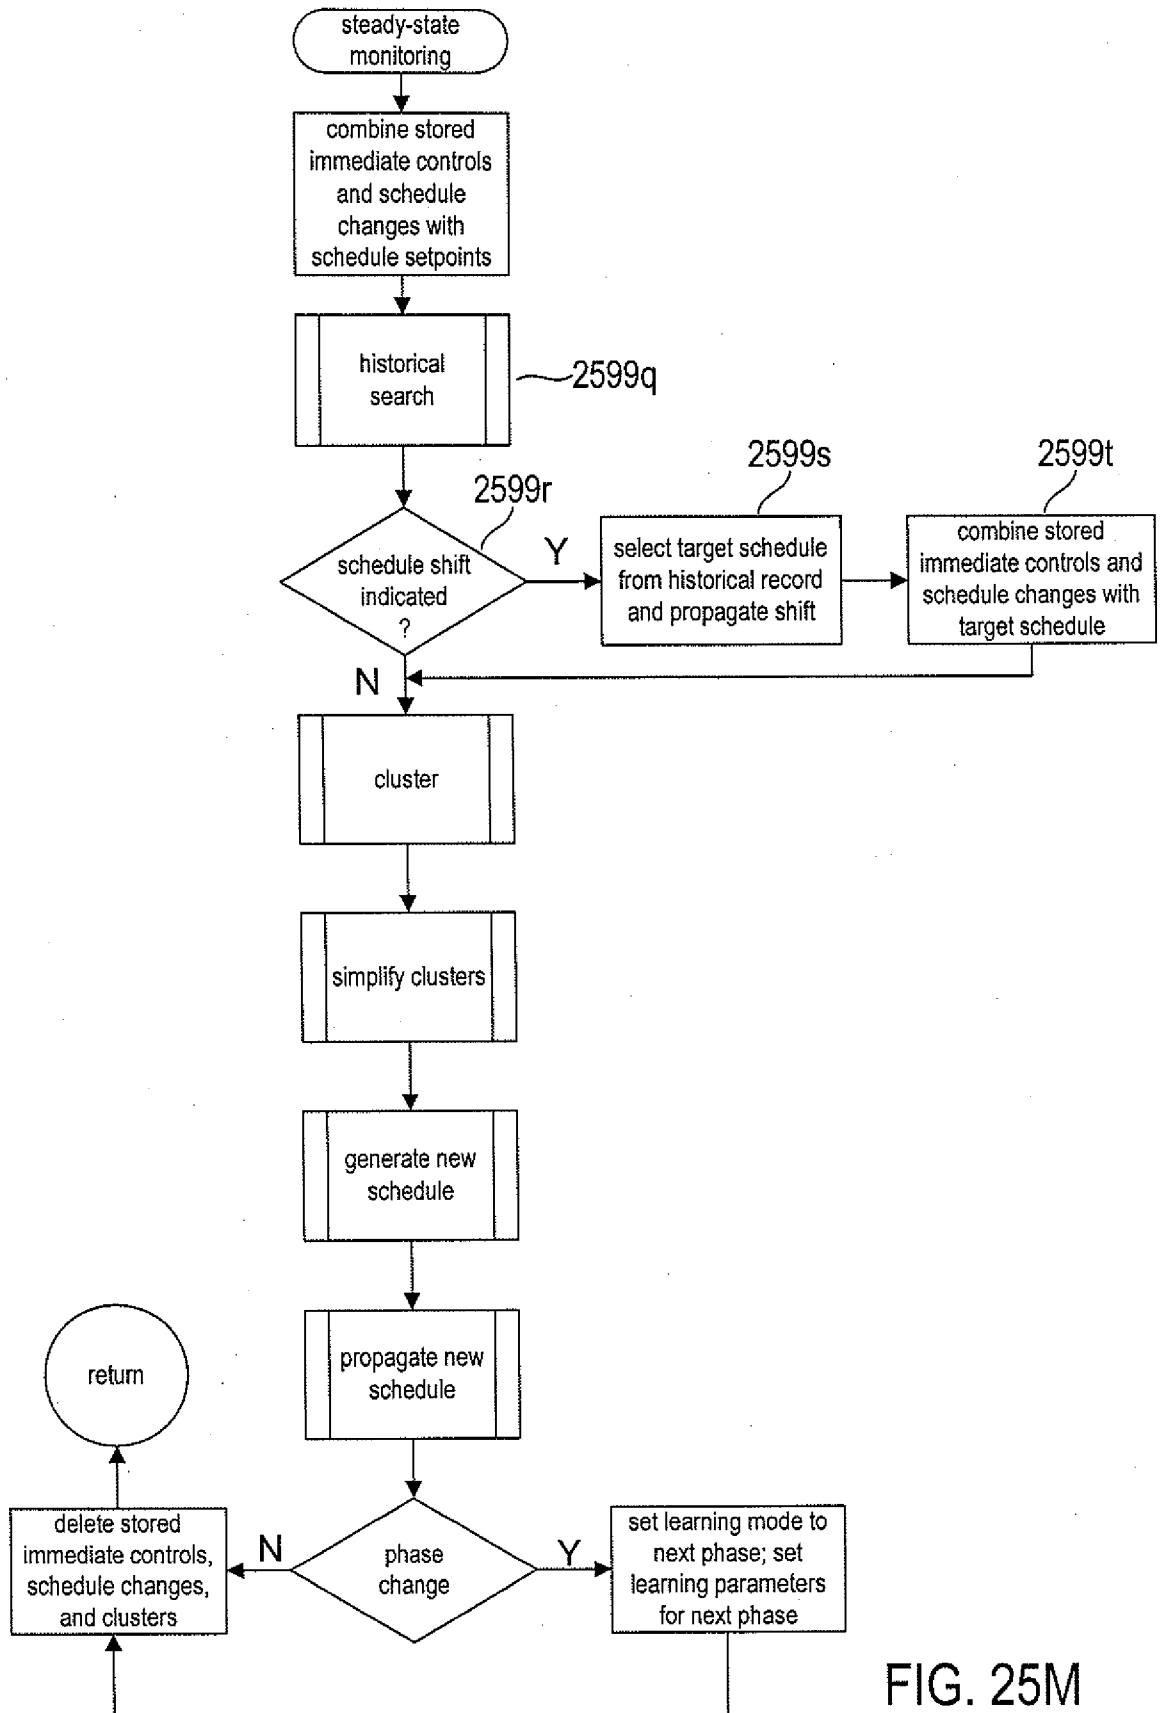


FIG. 25M




FIG. 26

47/82

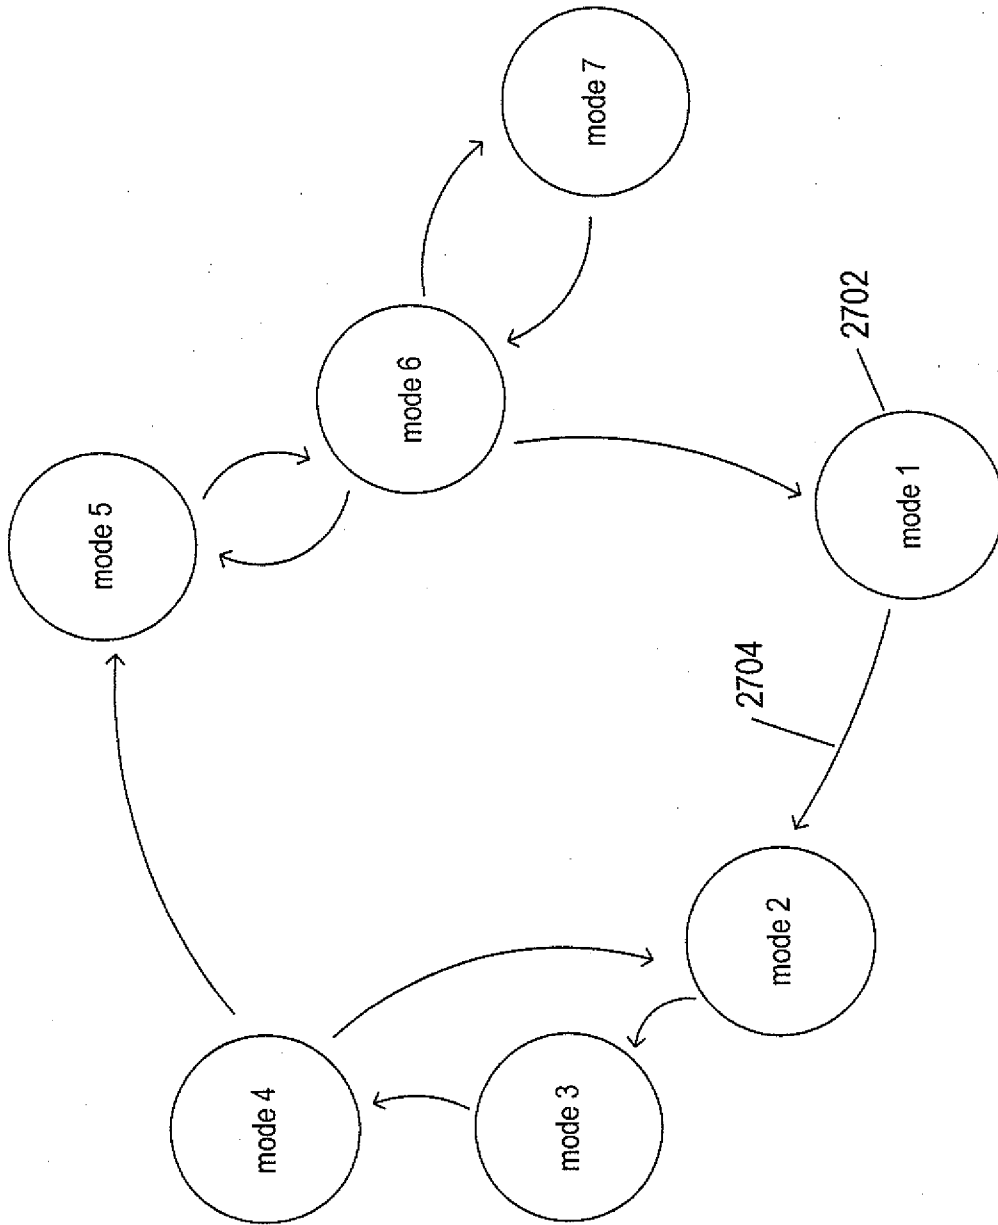


FIG. 27

48/82




FIG. 28A




FIG. 28B




FIG. 28C

49/82




FIG. 29

50/82




FIG. 30

51/82




FIG. 31A




FIG. 31B




FIG. 31C

52/82




FIG. 32

53/82




FIG. 33A

54/82




FIG. 33B

55/82




FIG. 34A

56/82




FIG. 34B

57/82




FIG. 35A

58/82




FIG. 35B

59/82




FIG. 36A

60/82




FIG. 36B

61/82




FIG. 37

62/82




FIG. 38




FIG. 39

63/82

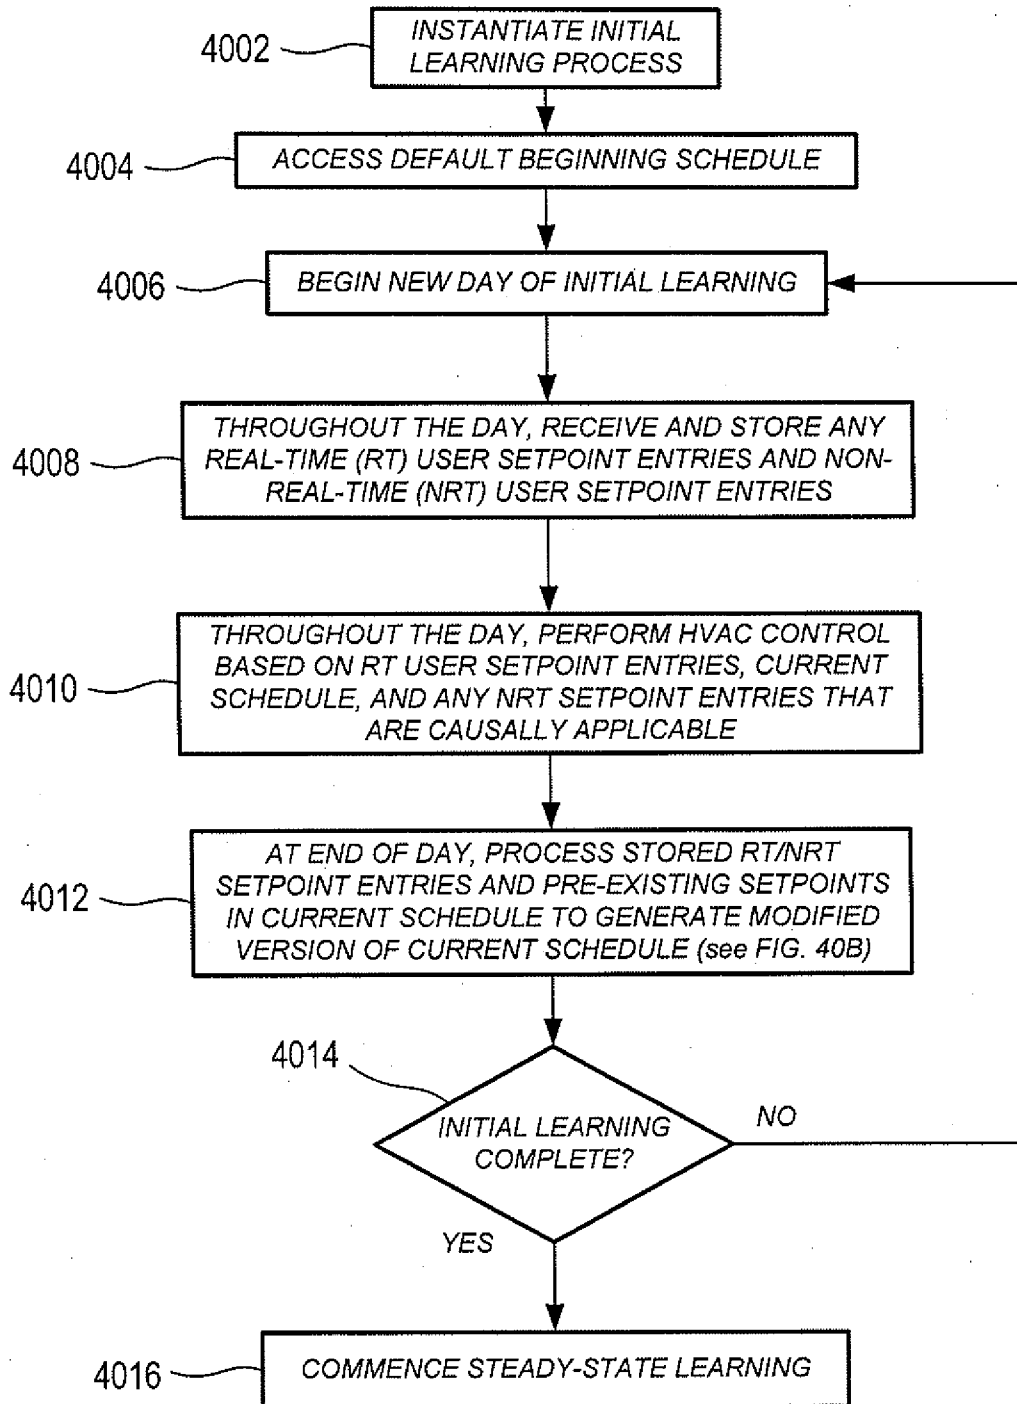


FIG. 40A

4012

64/82

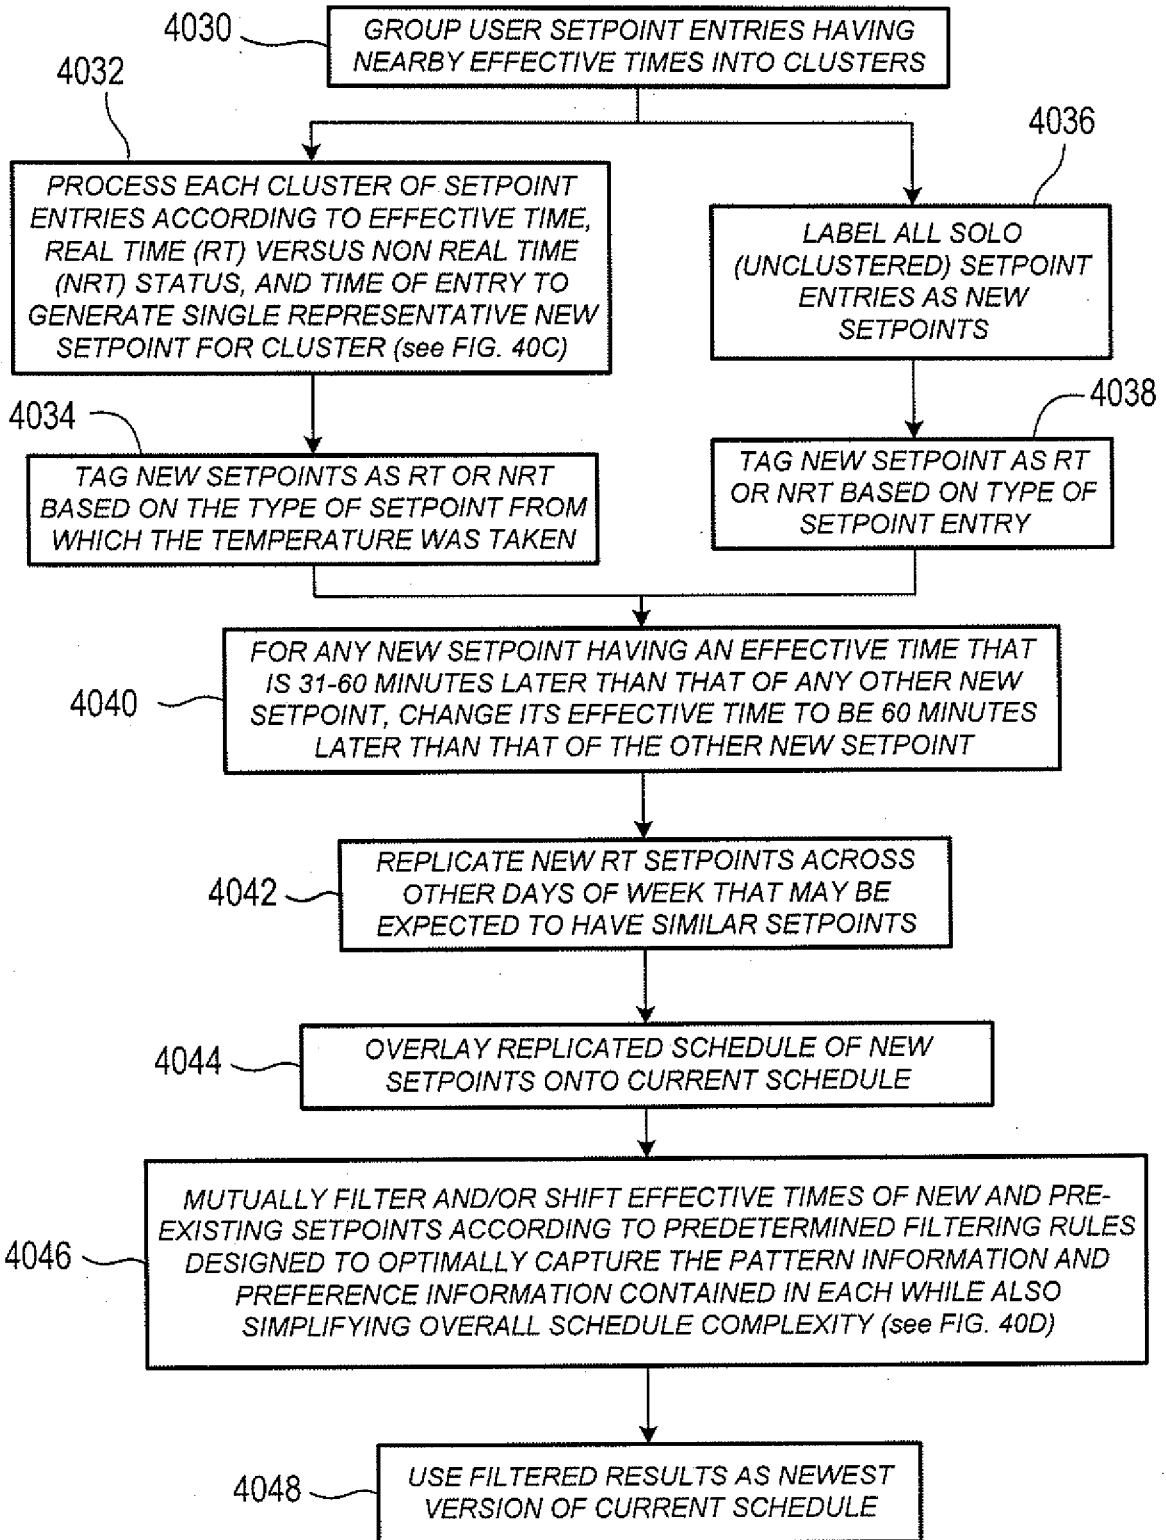


FIG. 40B

65/82

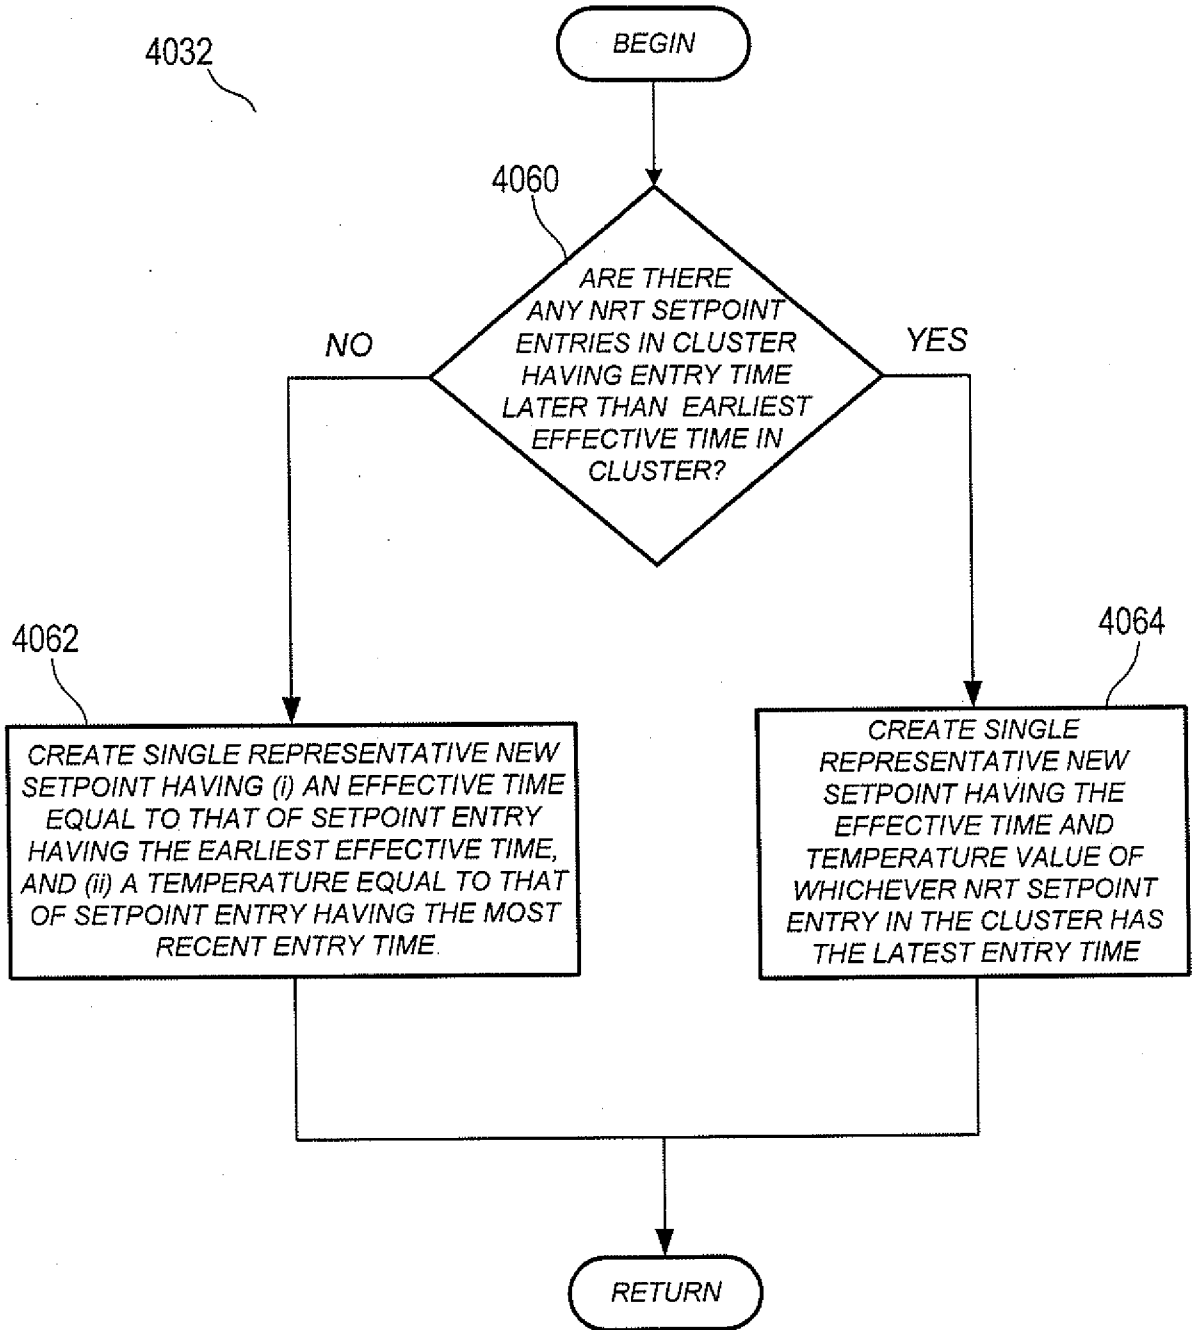


FIG. 40C

66/82

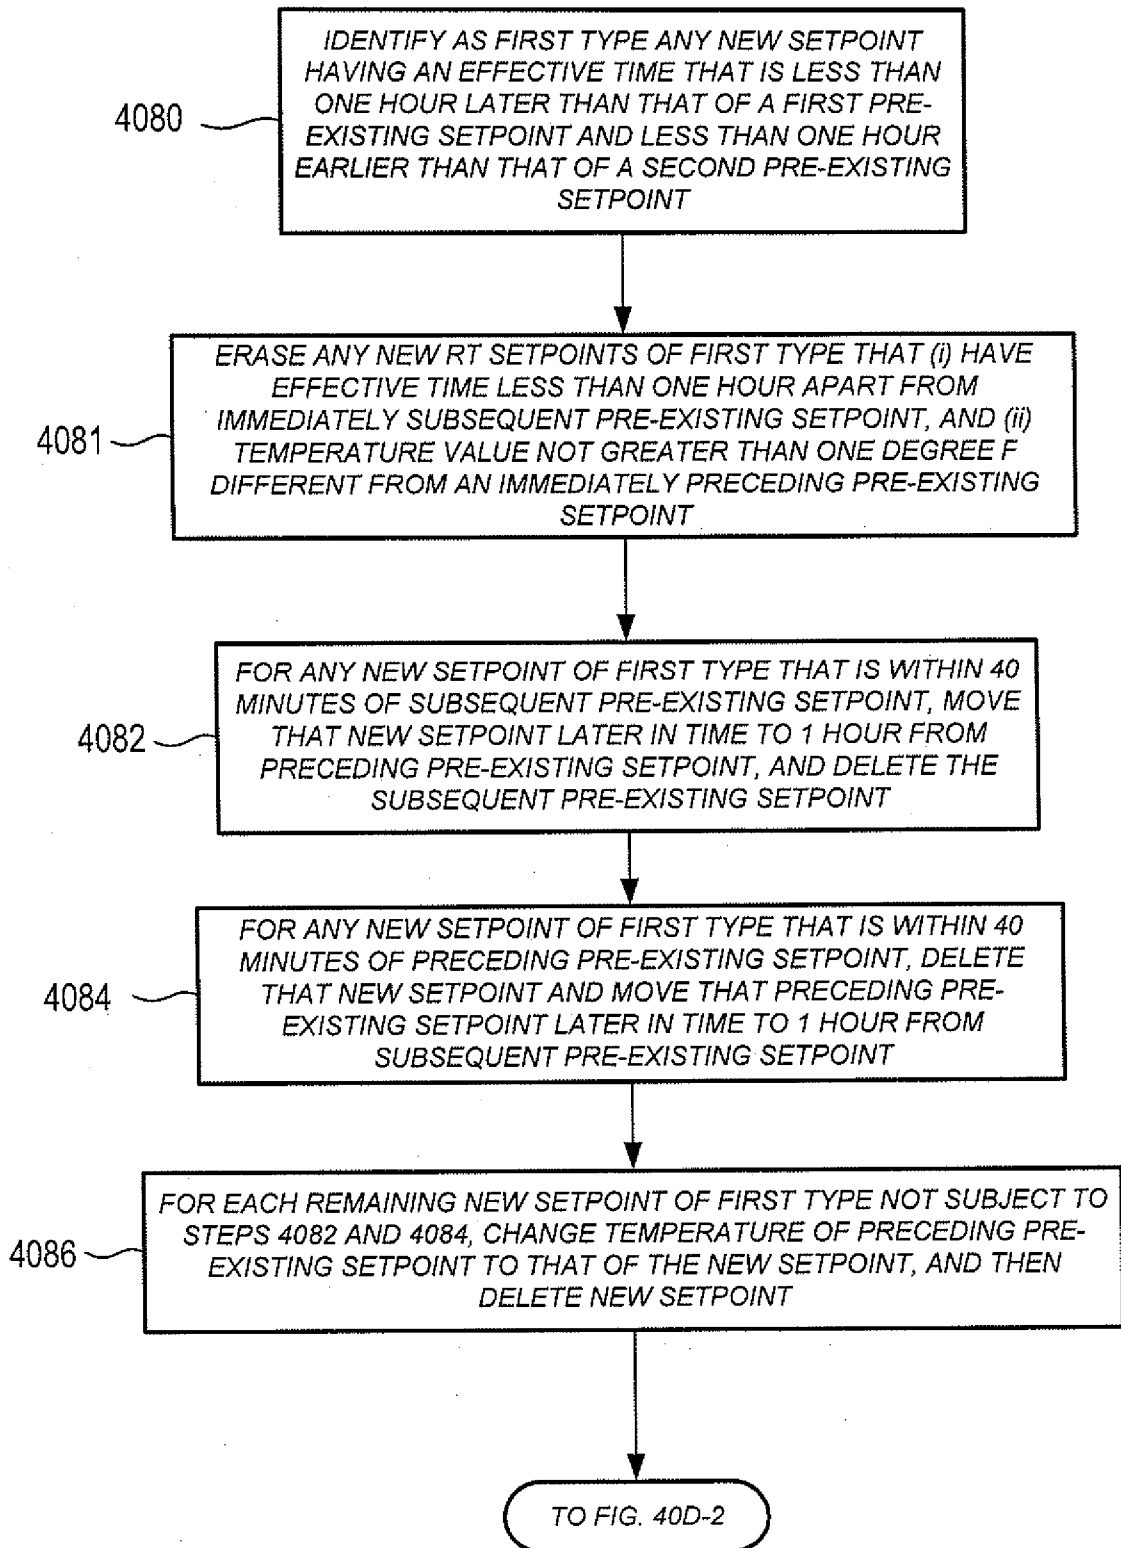


FIG. 40D-1

67/82

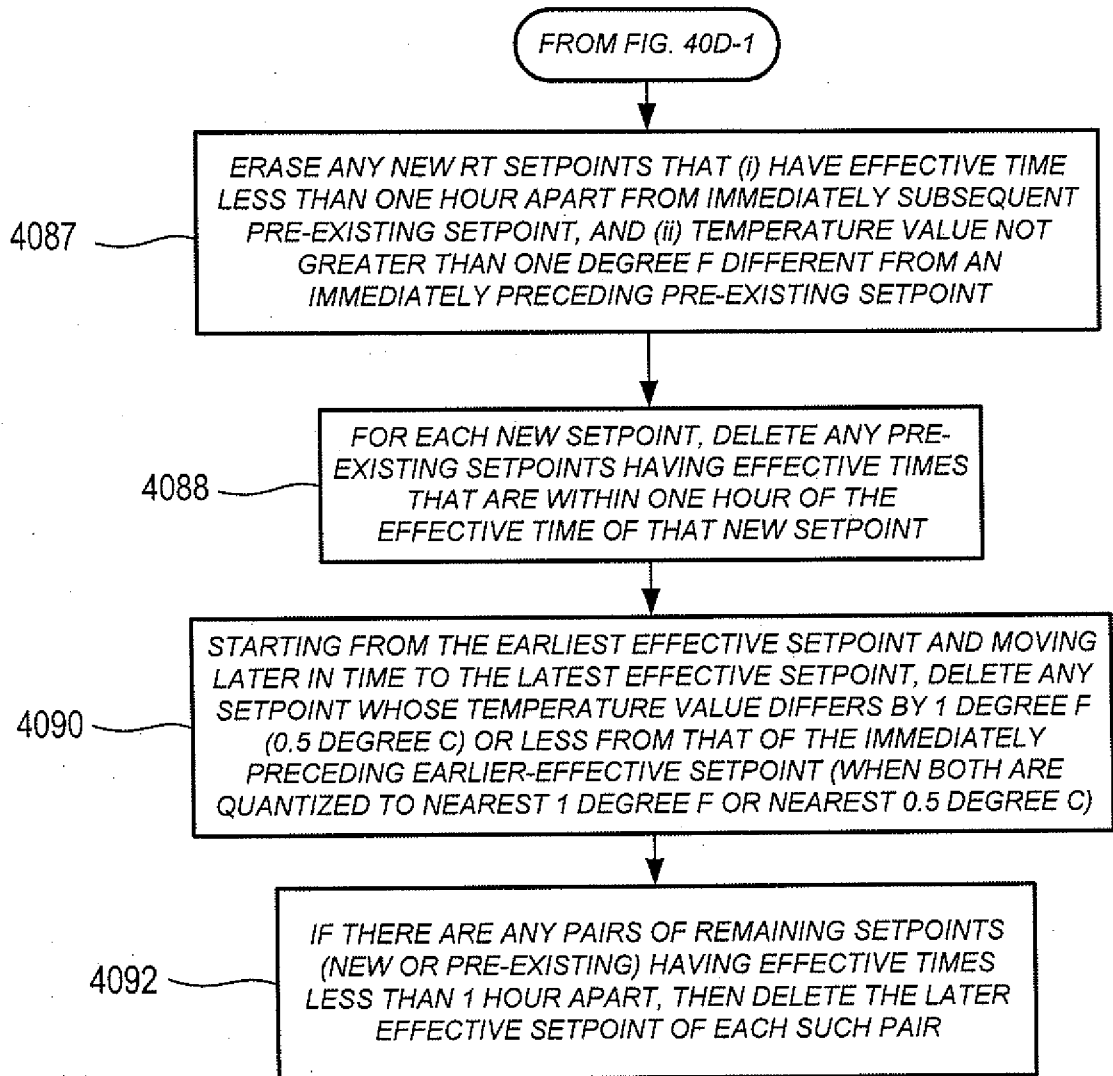


FIG. 40D-2

68/82





FIG. 41A








FIG. 41B




70/82

FIG. 41C



71/82

FIG. 41D



72/82

FIG. 41E






FIG. 41G




75/82

FIG. 41H




76/82

FIG. 411




77/82

FIG. 41J




78/82

FIG. 41K




79/82

FIG. 41L



80/82

FIG. 41M



81/82

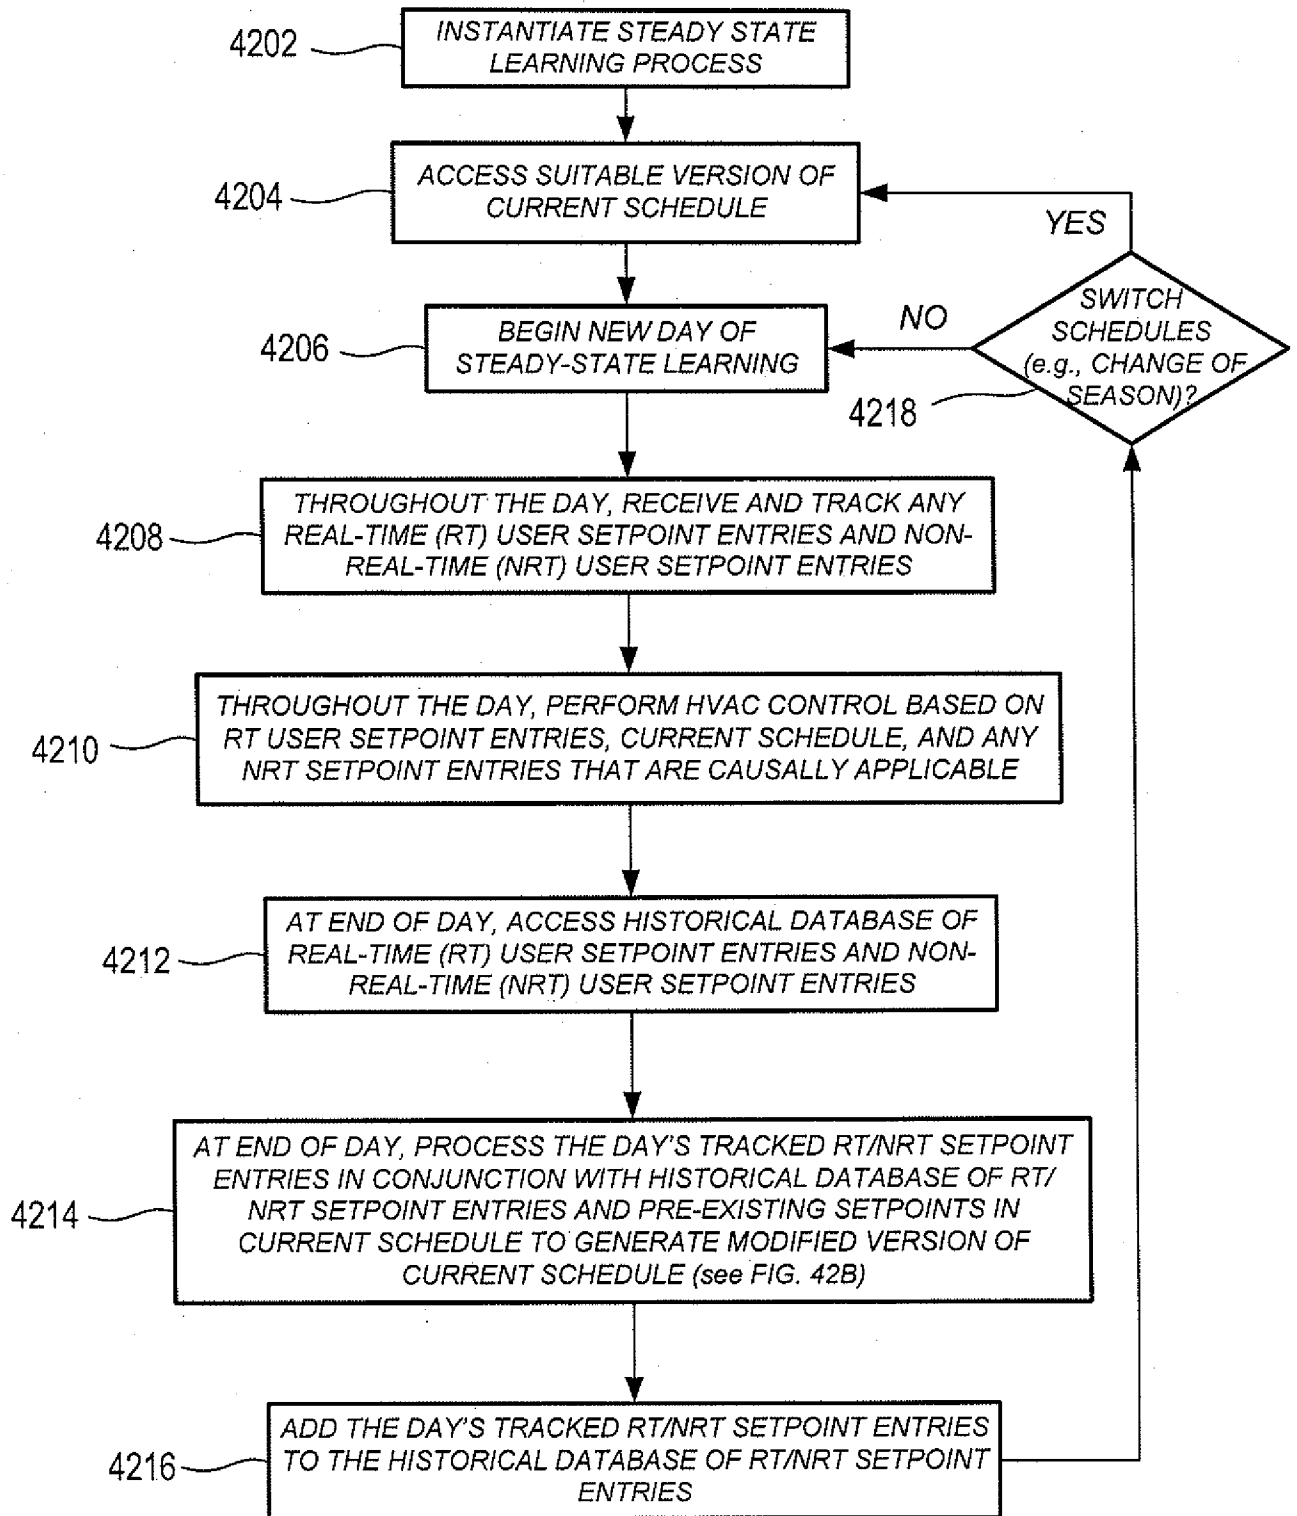


FIG. 42A

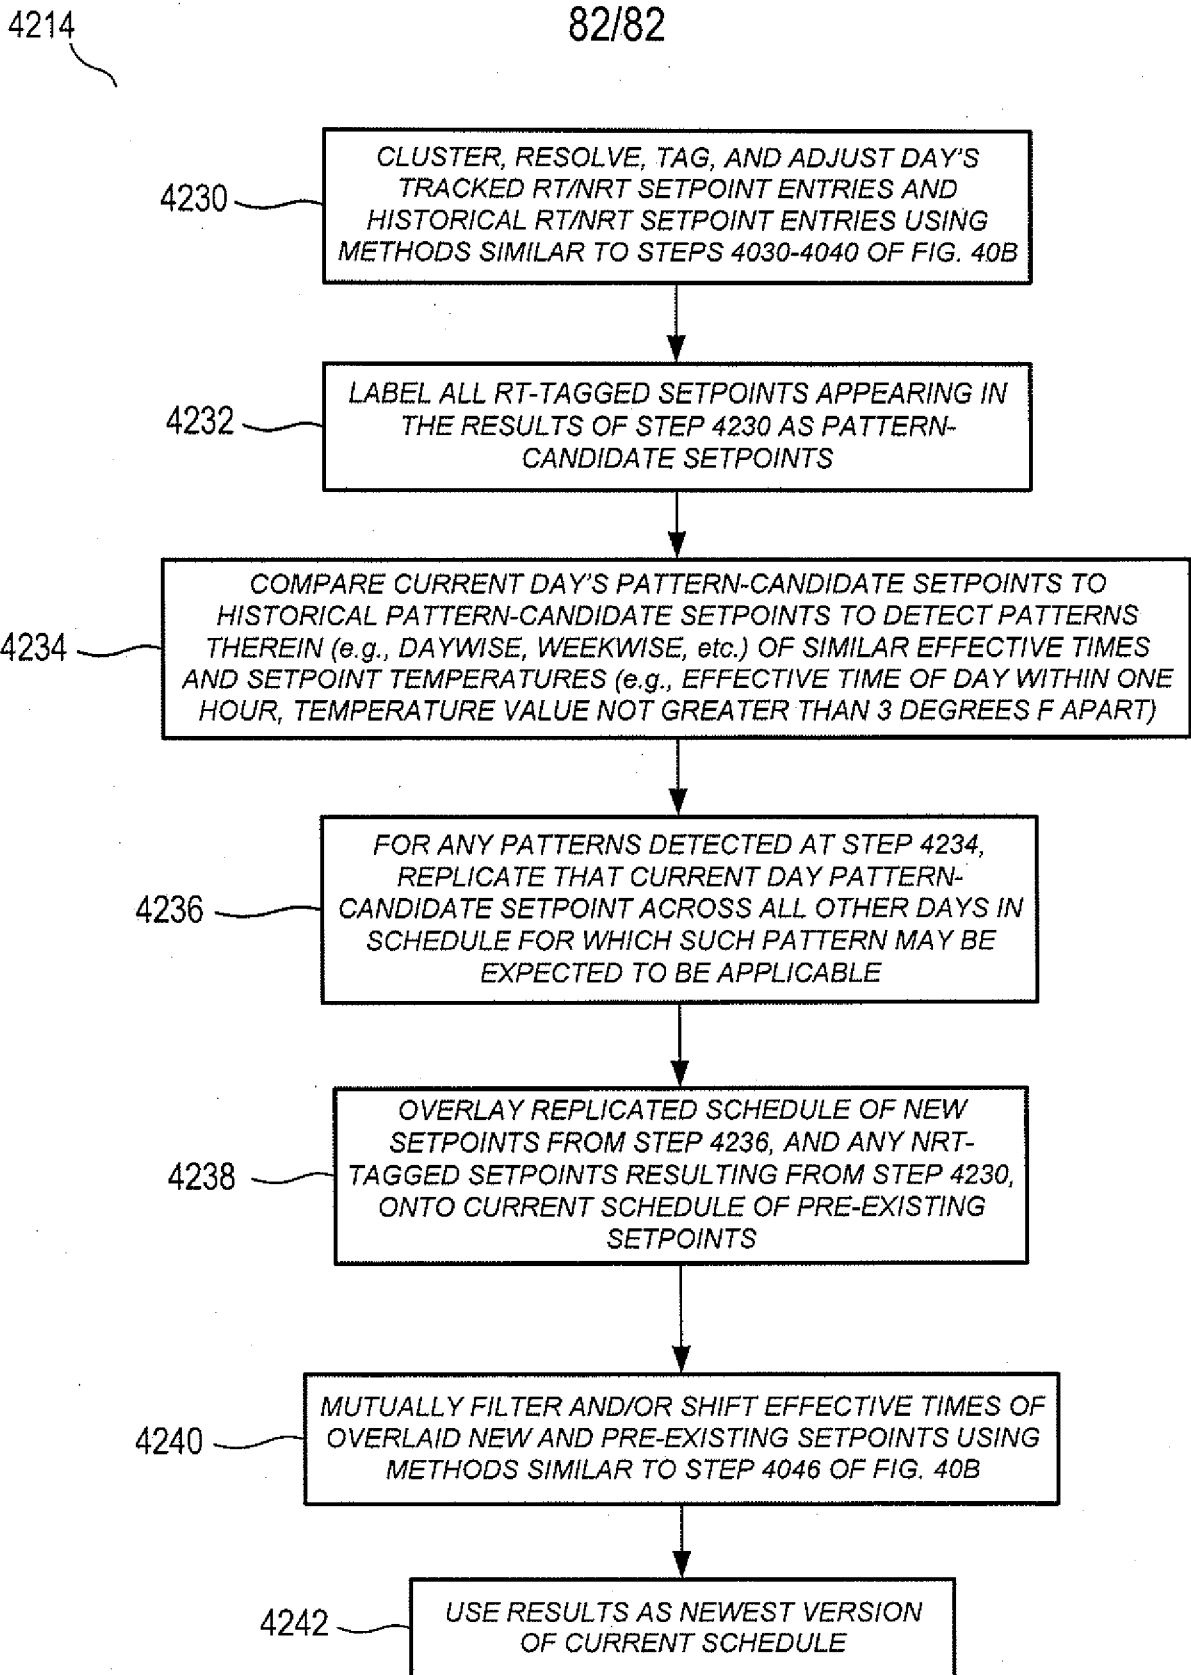


FIG. 42B

A. CLASSIFICATION OF SUBJECT MATTER**G05D 23/30(2006.01)i, G05D 23/24(2006.01)i, F24F 11/02(2006.01)i**

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)

G05D 23/30; G05B 13/02; G05B 13/04; G05B 15/02; G06F 1/26; G05D 23/19; G05B 23/02

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

eKOMPASS(KIPO internal) & Keywords: intelligent, controller, schedule, modification

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y A	US 2011-0196539 A1 (NAIR, AJAY N. et al.) 11 August 2011 See paragraphs [0019]-[0035]; claims 1-20; and figures 1-5.	1-5, 14-15, 19 6-13, 16-18, 20
Y	US 2009-0195349 A1 (FRADER-THOMPSON, SETH et al.) 06 August 2009 See paragraphs [0087]-[0162]; claims 1,3,19,23,44; and figures 1,3,6.	1-5, 14-15, 19
A	US 8036760 B2 (MEHTA, ASHISH et al.) 11 October 2011 See column 3, line 16 - column 4, line 5; claim 1; and figures 1-3.	1-20
A	US 2011-0202198 A1 (VENKATAKRISHNAN, NATARAJAN et al.) 18 August 2011 See claims 1,3,8; and figure 1.	1-20
A	US 2011-0153089 A1 (TIEMANN, DIRK et al.) 23 June 2011 See paragraphs [0009]-[0018],[0038]; claims 1-3; and figures 2-3,5-7	1-20

 Further documents are listed in the continuation of Box C. 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 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

20 FEBRUARY 2013 (20.02.2013)

Date of mailing of the international search report

21 FEBRUARY 2013 (21.02.2013)

Name and mailing address of the ISA/KR



Facsimile No. 82-42-472-7140

Authorized officer

BYUN, Sung Cheal

Telephone No. 82-42-481-8262



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2012/058187

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2011-0196539 A1	11.08.2011	None	
US 2009-0195349 A1	06.08.2009	CA 2713702 A1 CA 2749628 A1 CA 2749659 A1 US 2011-0061014 A1 US 2011-0063126 A1 US 8255090 B2 WO 2009-097400 A1	06.08.2009 26.02.2012 26.02.2012 10.03.2011 17.03.2011 28.08.2012 06.08.2009
US 8036760 B2	11.10.2011	CN 101533261 A EP 2045673 A2 EP 2045673 A3 GB 2453261 A JP 2009-087350 A US 2009-0112335 A1	16.09.2009 08.04.2009 13.10.2010 01.04.2009 23.04.2009 30.04.2009
US 2011-0202198 A1	18.08.2011	CA 2752987 A1 US 2011-202189 A1 US 2011-202190 A1 US 2011-202194 A1 US 2011-202195 A1 US 2011-202196 A1 US 2011-202293 A1 US 2011-202783 A1 US 2011-202910 A1 US 8295990 B2	03.07.2012 18.08.2011 18.08.2011 18.08.2011 18.08.2011 18.08.2011 18.08.2011 18.08.2011 18.08.2011 23.10.2012
US 2011-0153089 A1	23.06.2011	CA 2725358 A1 CN 102084306 A EP 2294487 A1 EP 2294487 B1 WO 2010-000077 A1	07.01.2010 01.06.2011 16.03.2011 28.03.2012 07.01.2010