



(19) **United States**

(12) **Patent Application Publication**  
**Kohn et al.**

(10) **Pub. No.: US 2017/0271984 A1**

(43) **Pub. Date: Sep. 21, 2017**

(54) **USING BATTERY DC CHARACTERISTICS TO CONTROL POWER OUTPUT**

**Publication Classification**

(71) Applicant: **Atigeo Corp.**, Bellevue, WA (US)

(51) **Int. Cl.**  
*H02M 3/156* (2006.01)  
*G05B 15/02* (2006.01)  
*H02J 7/00* (2006.01)

(72) Inventors: **Wolf Kohn**, Seattle, WA (US); **Vishnu Vettrivel**, Bothell, WA (US); **Jonathan Cross**, Bellevue, WA (US); **Pengbo Zhang**, Seattle, WA (US); **Michael Luis Sandoval**, Bellevue, WA (US); **Brian Schaper**, Renton, WA (US); **Neel Master**, Santa Monica, CA (US); **Brandon Weiss**, Woodbridge, VA (US); **David Kettler**, Seattle, WA (US)

(52) **U.S. Cl.**  
CPC ..... *H02M 3/156* (2013.01); *H02J 7/0068* (2013.01); *G05B 15/02* (2013.01)

(57) **ABSTRACT**

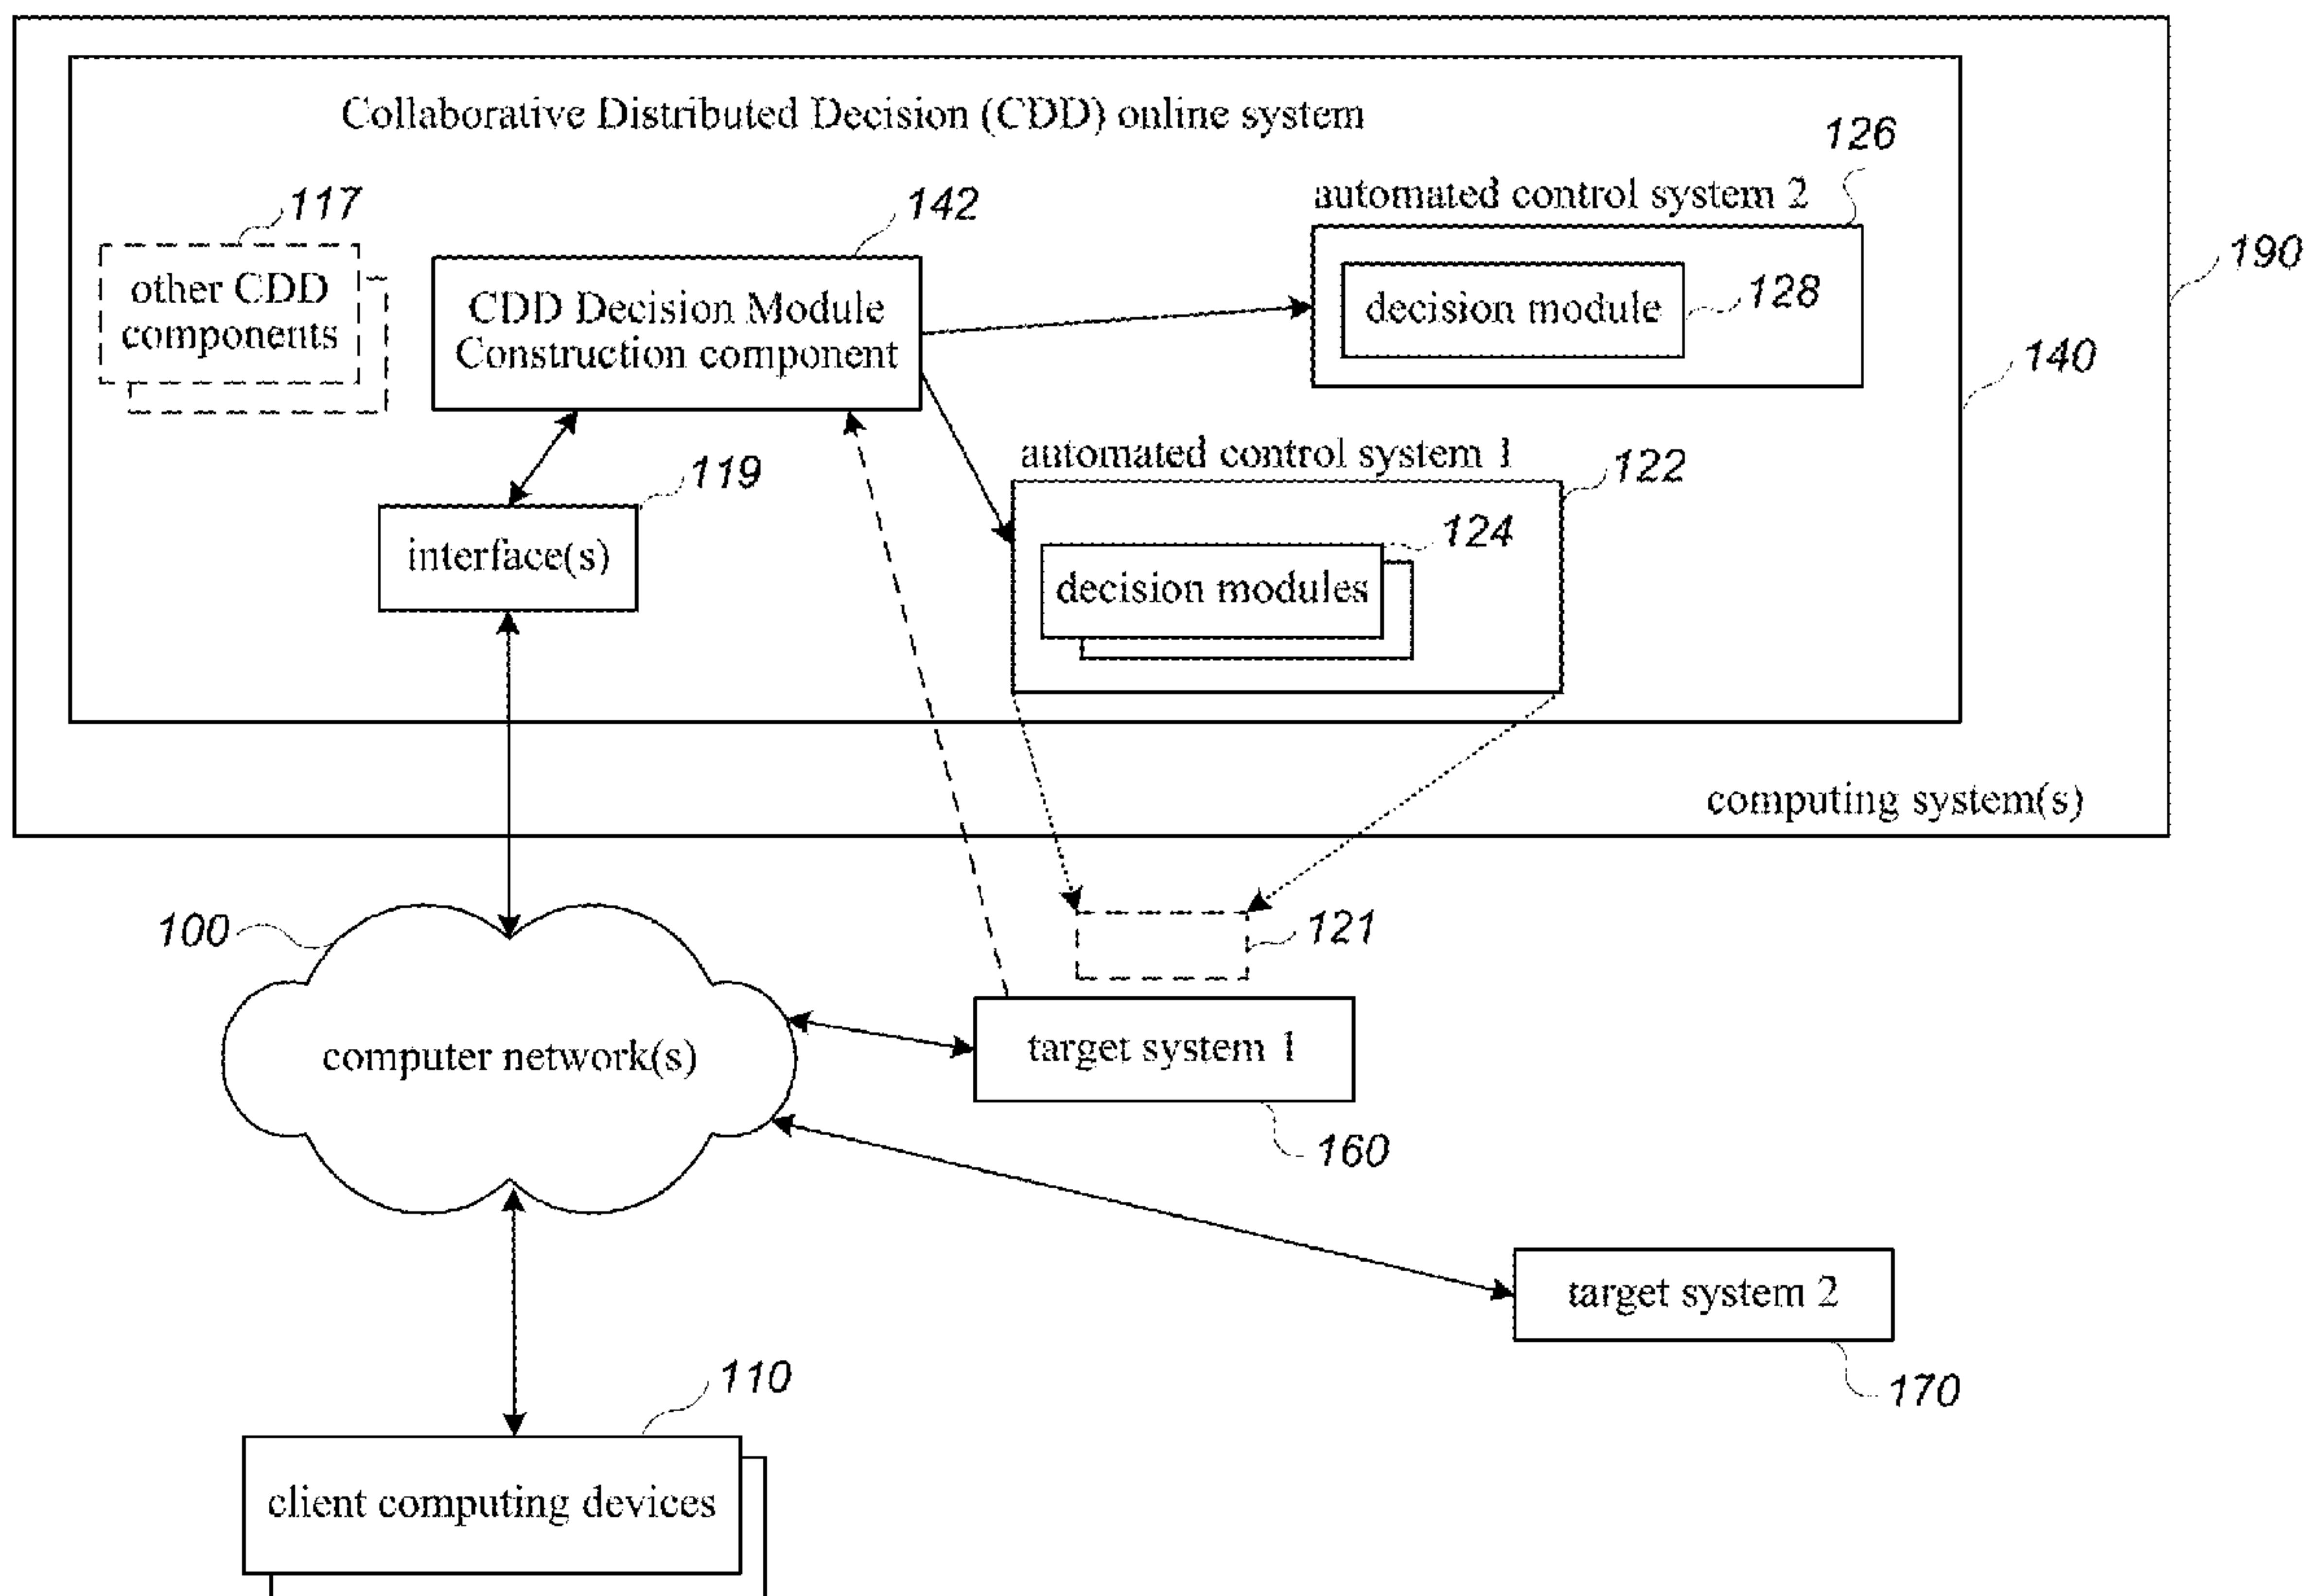
Techniques are described for implementing automated control systems to control operations of specified physical target systems, such as with one or more batteries used to store and provide electrical power. Characteristics of each battery's state may be used to perform automated control of DC power from the battery, such as in a real-time manner and to optimize long-term operation of the battery. In some situations, multiple batteries are controlled by using multiple control systems each associated with one of the batteries, and with overall control being coordinated in a distributed manner using interactions between the multiple control systems. A system that includes one or more batteries to be controlled may further include additional components in some situations, such as one or more electrical sources and/or one or more electrical loads, with one non-exclusive example of a type of such system being one or more home electrical power systems.

(21) Appl. No.: **15/096,091**

(22) Filed: **Apr. 11, 2016**

**Related U.S. Application Data**

(60) Provisional application No. 62/304,034, filed on Mar. 4, 2016.



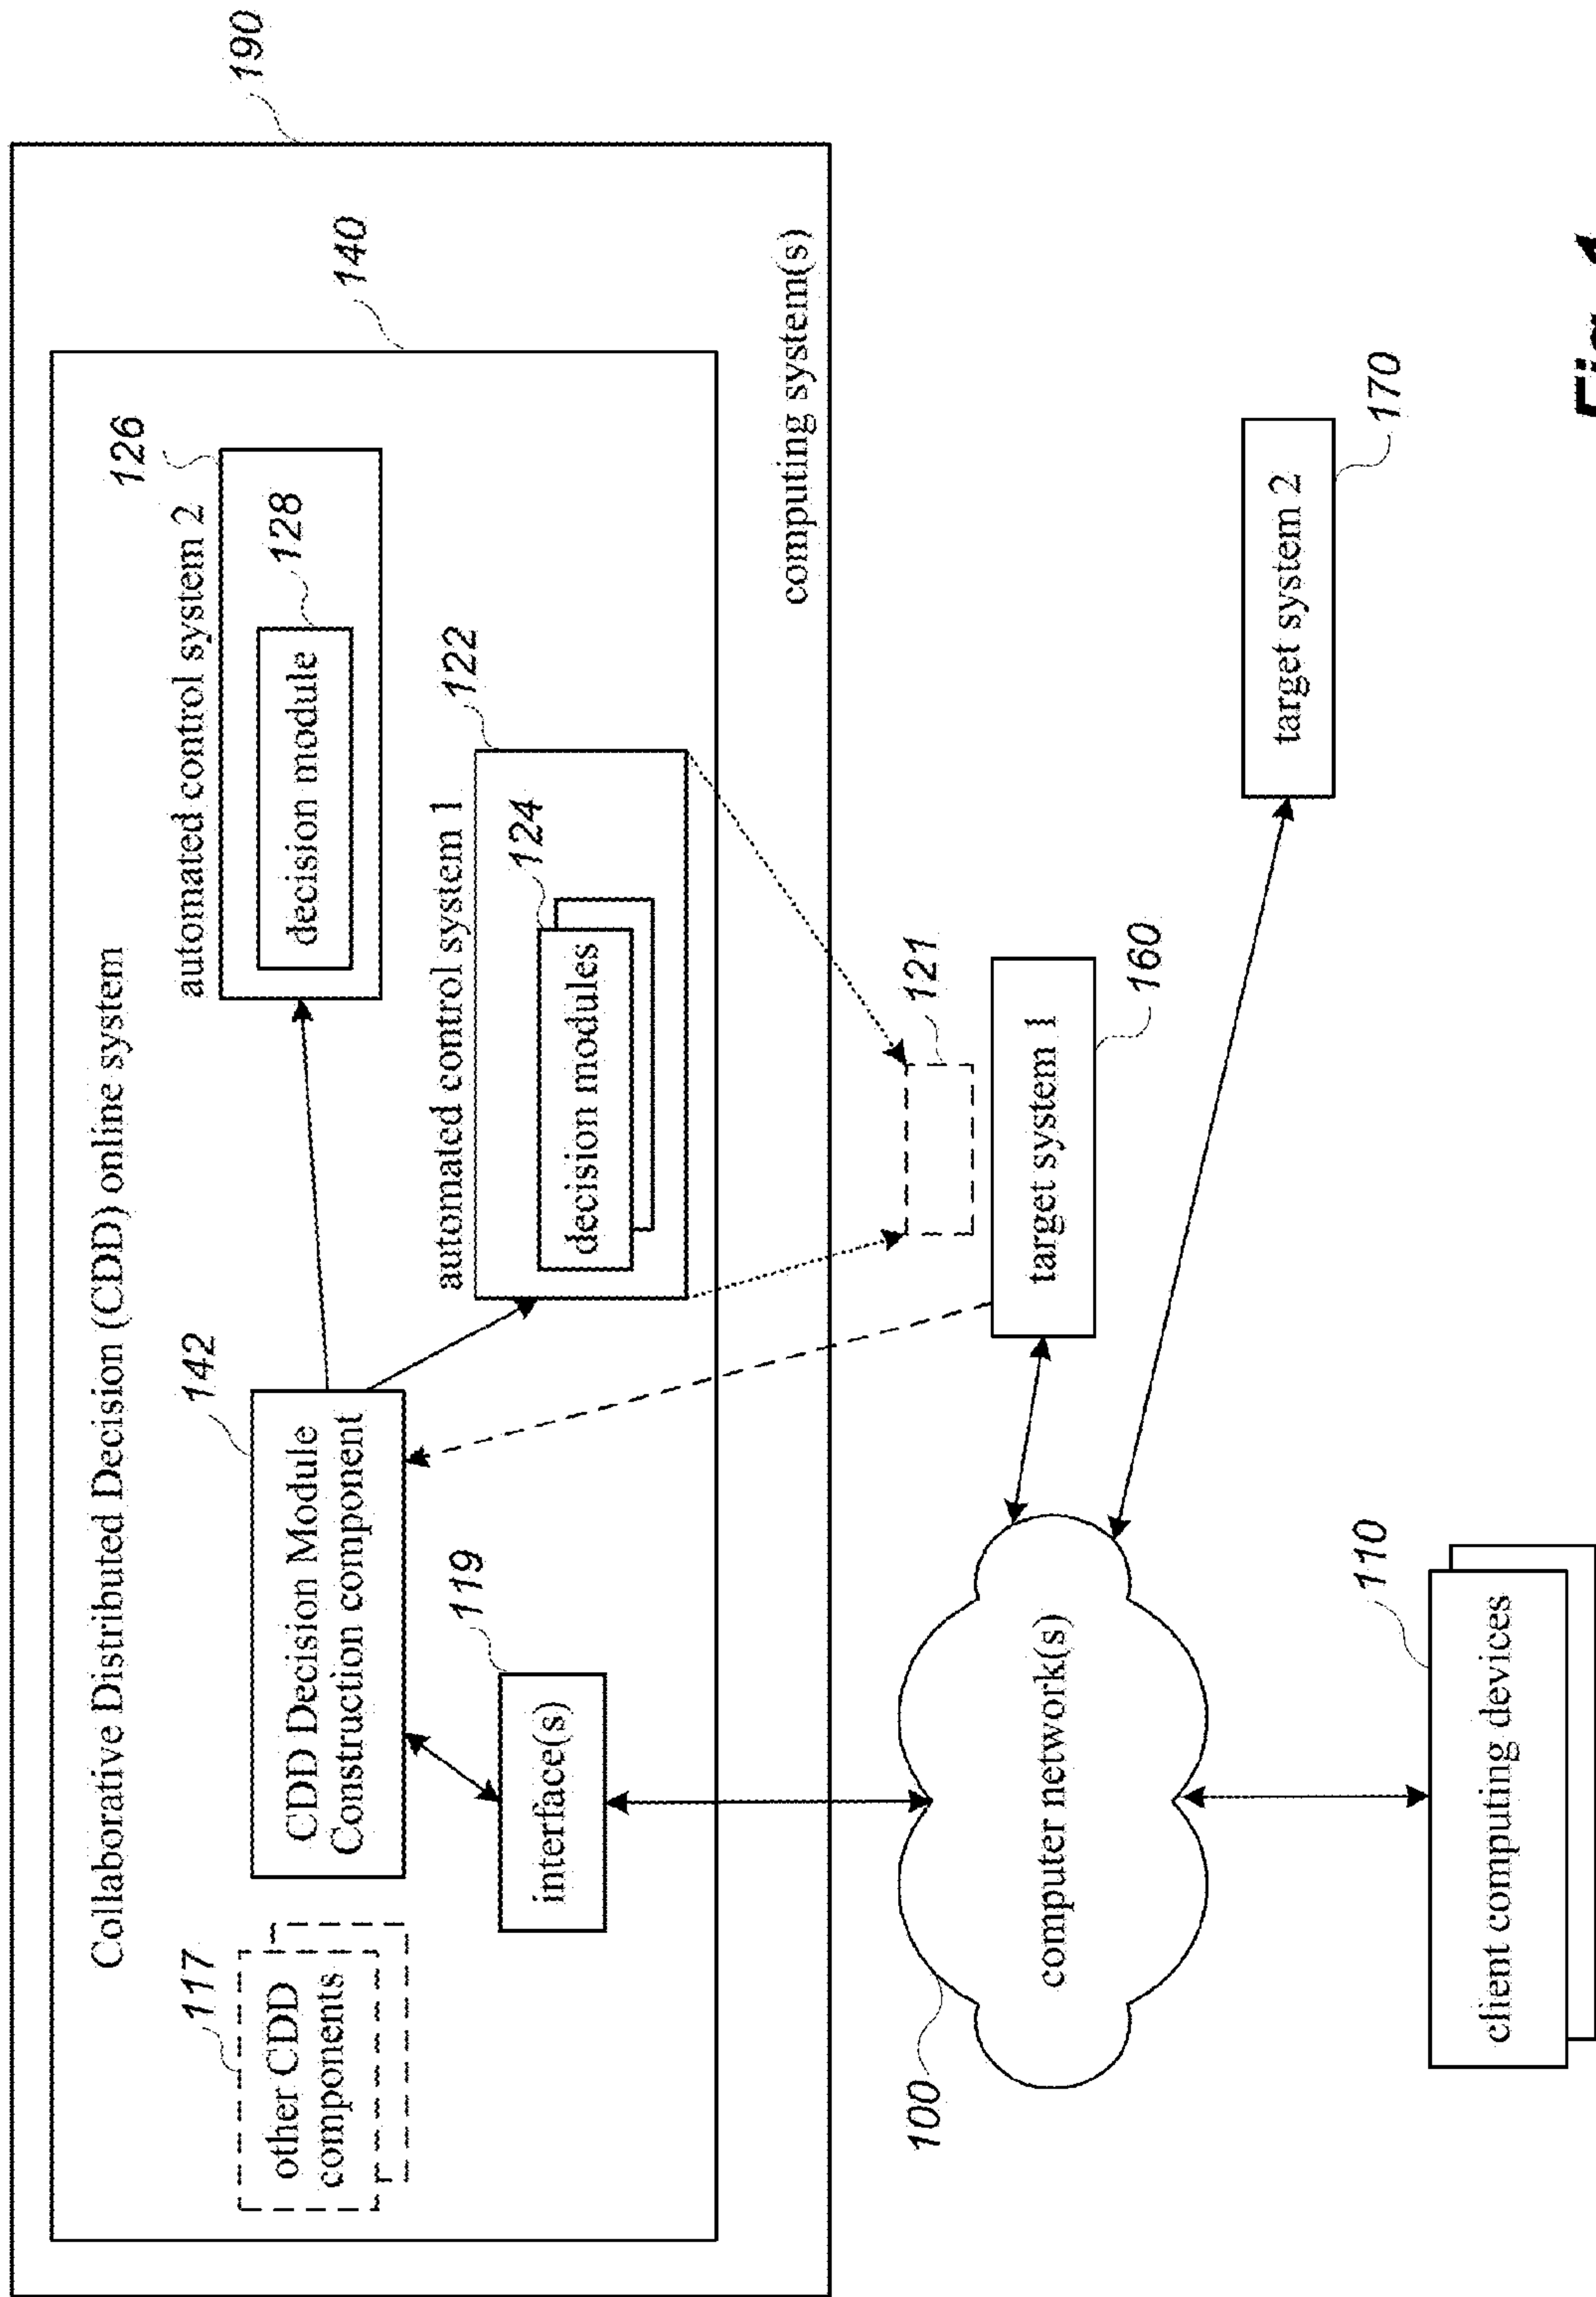
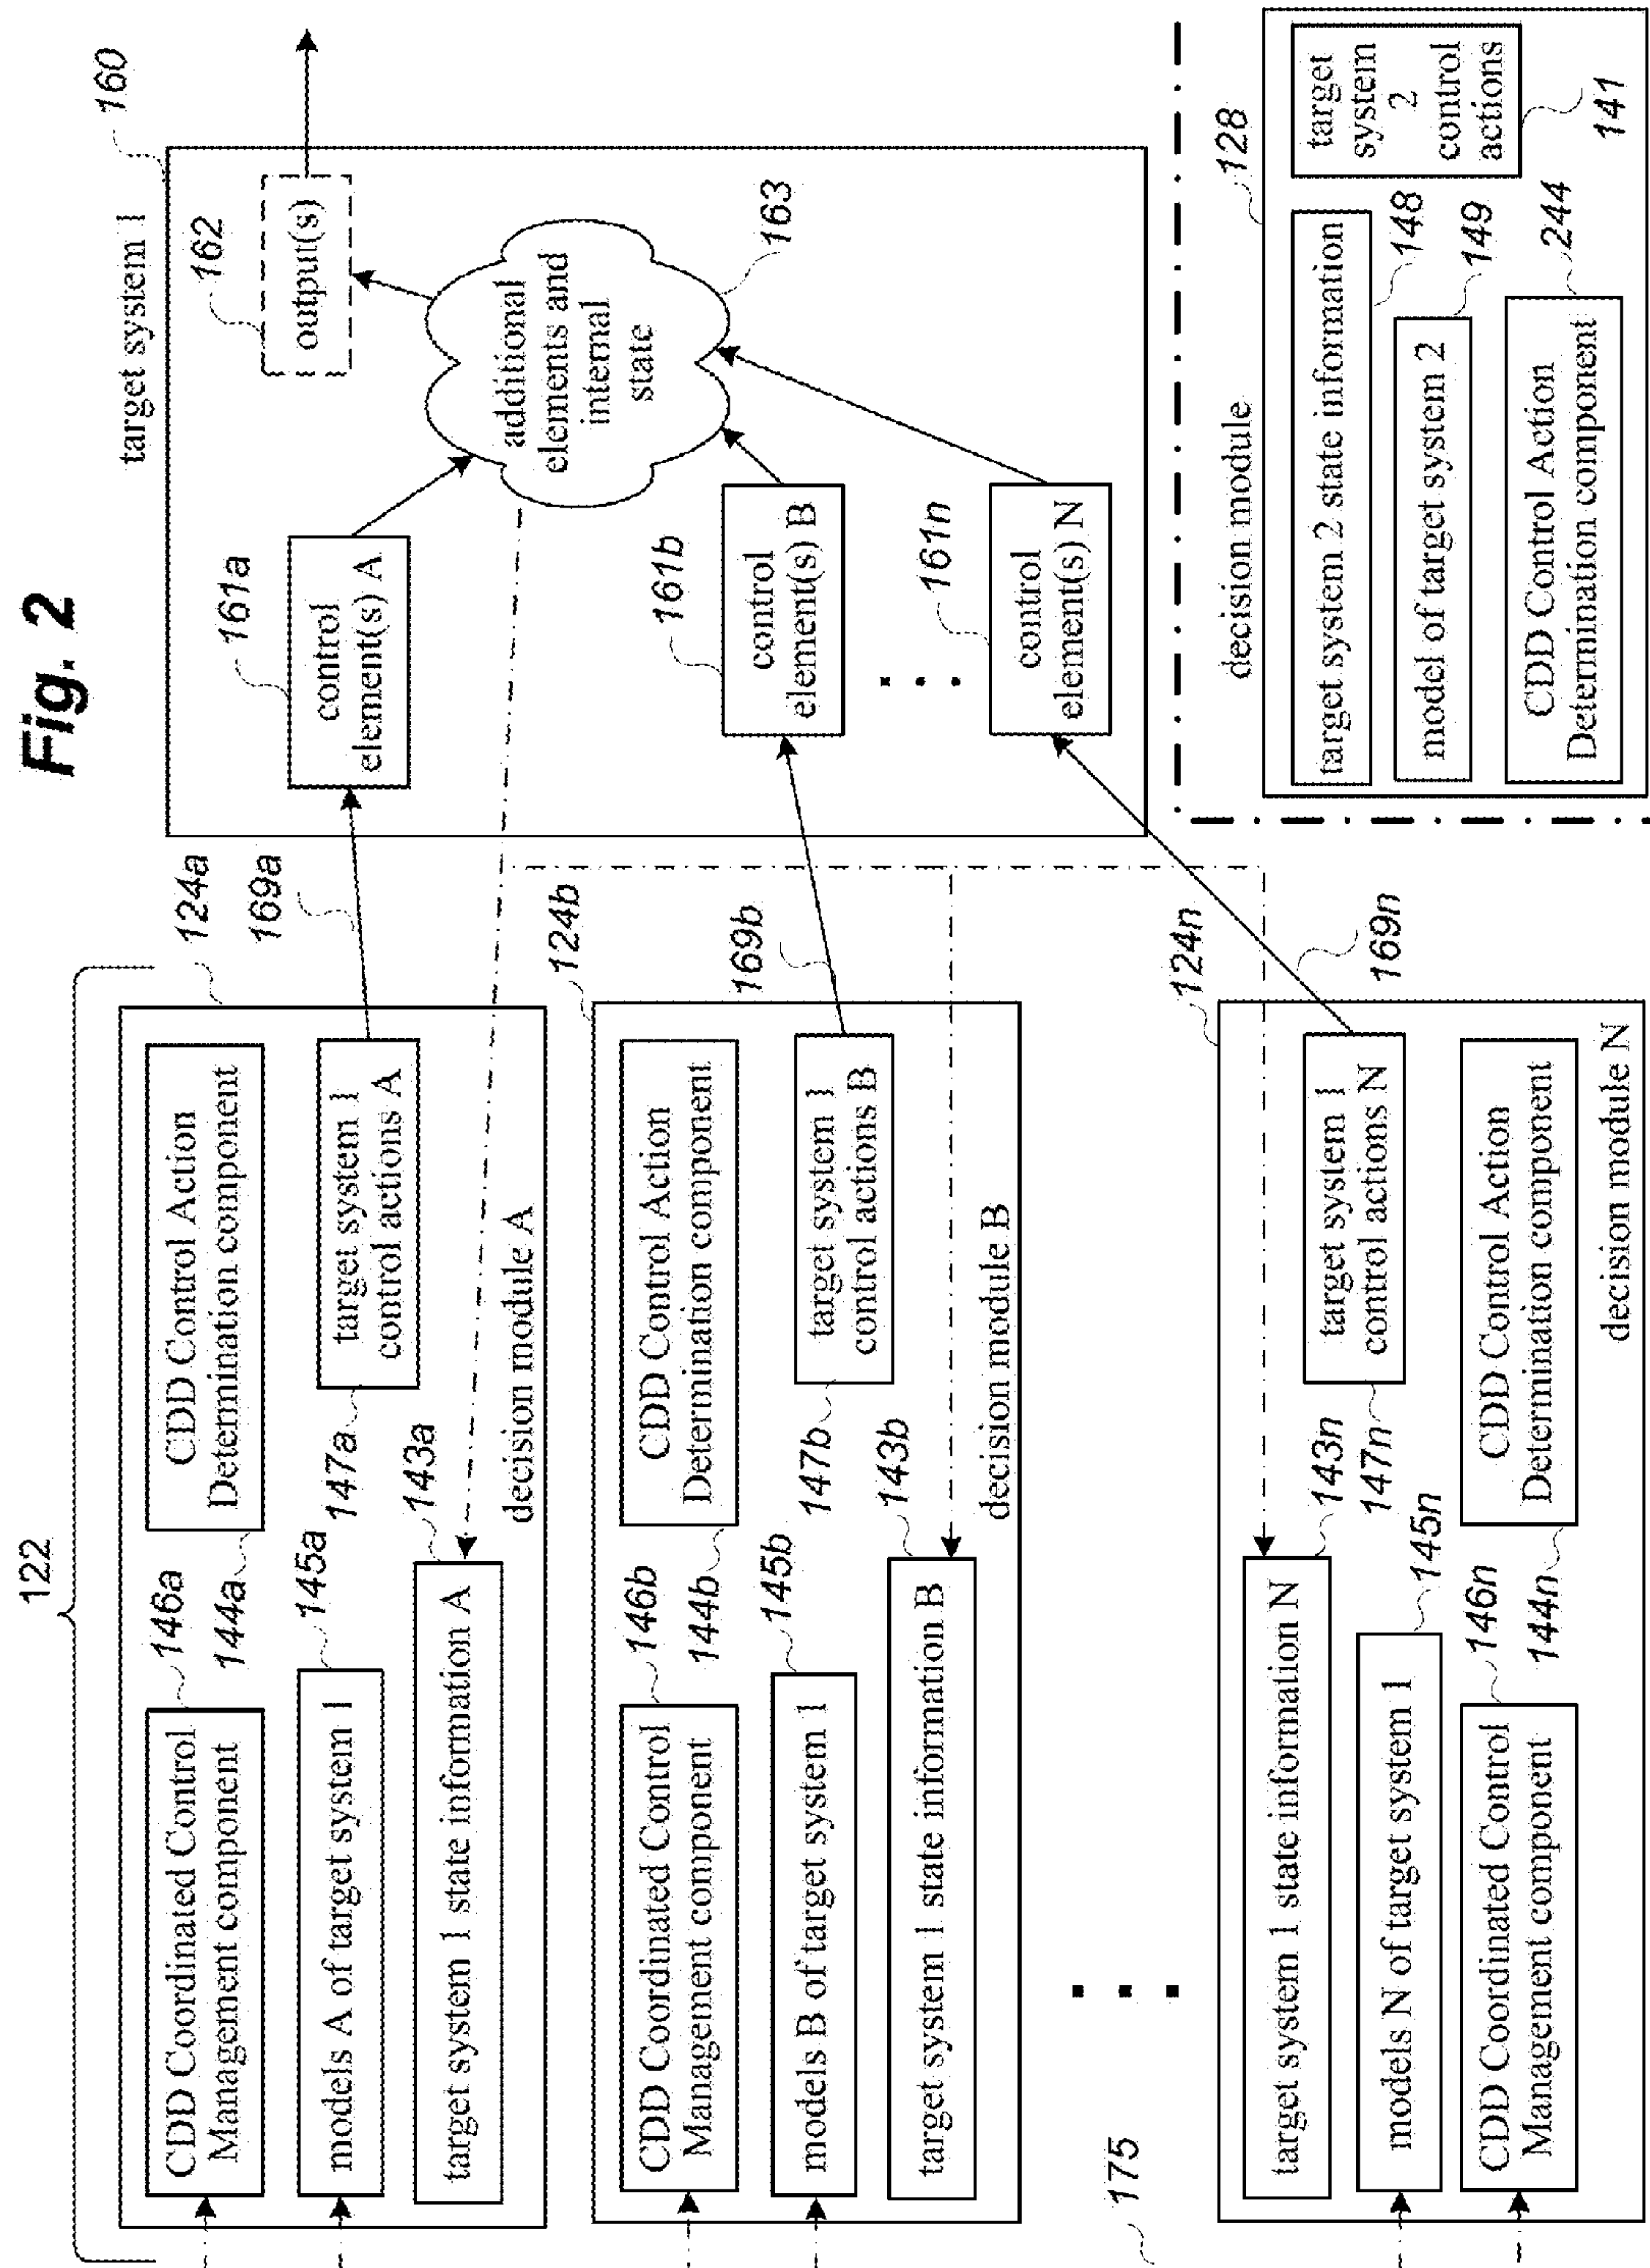


Fig. 1

Fig. 2








Fig. 3



**Fig. 4**

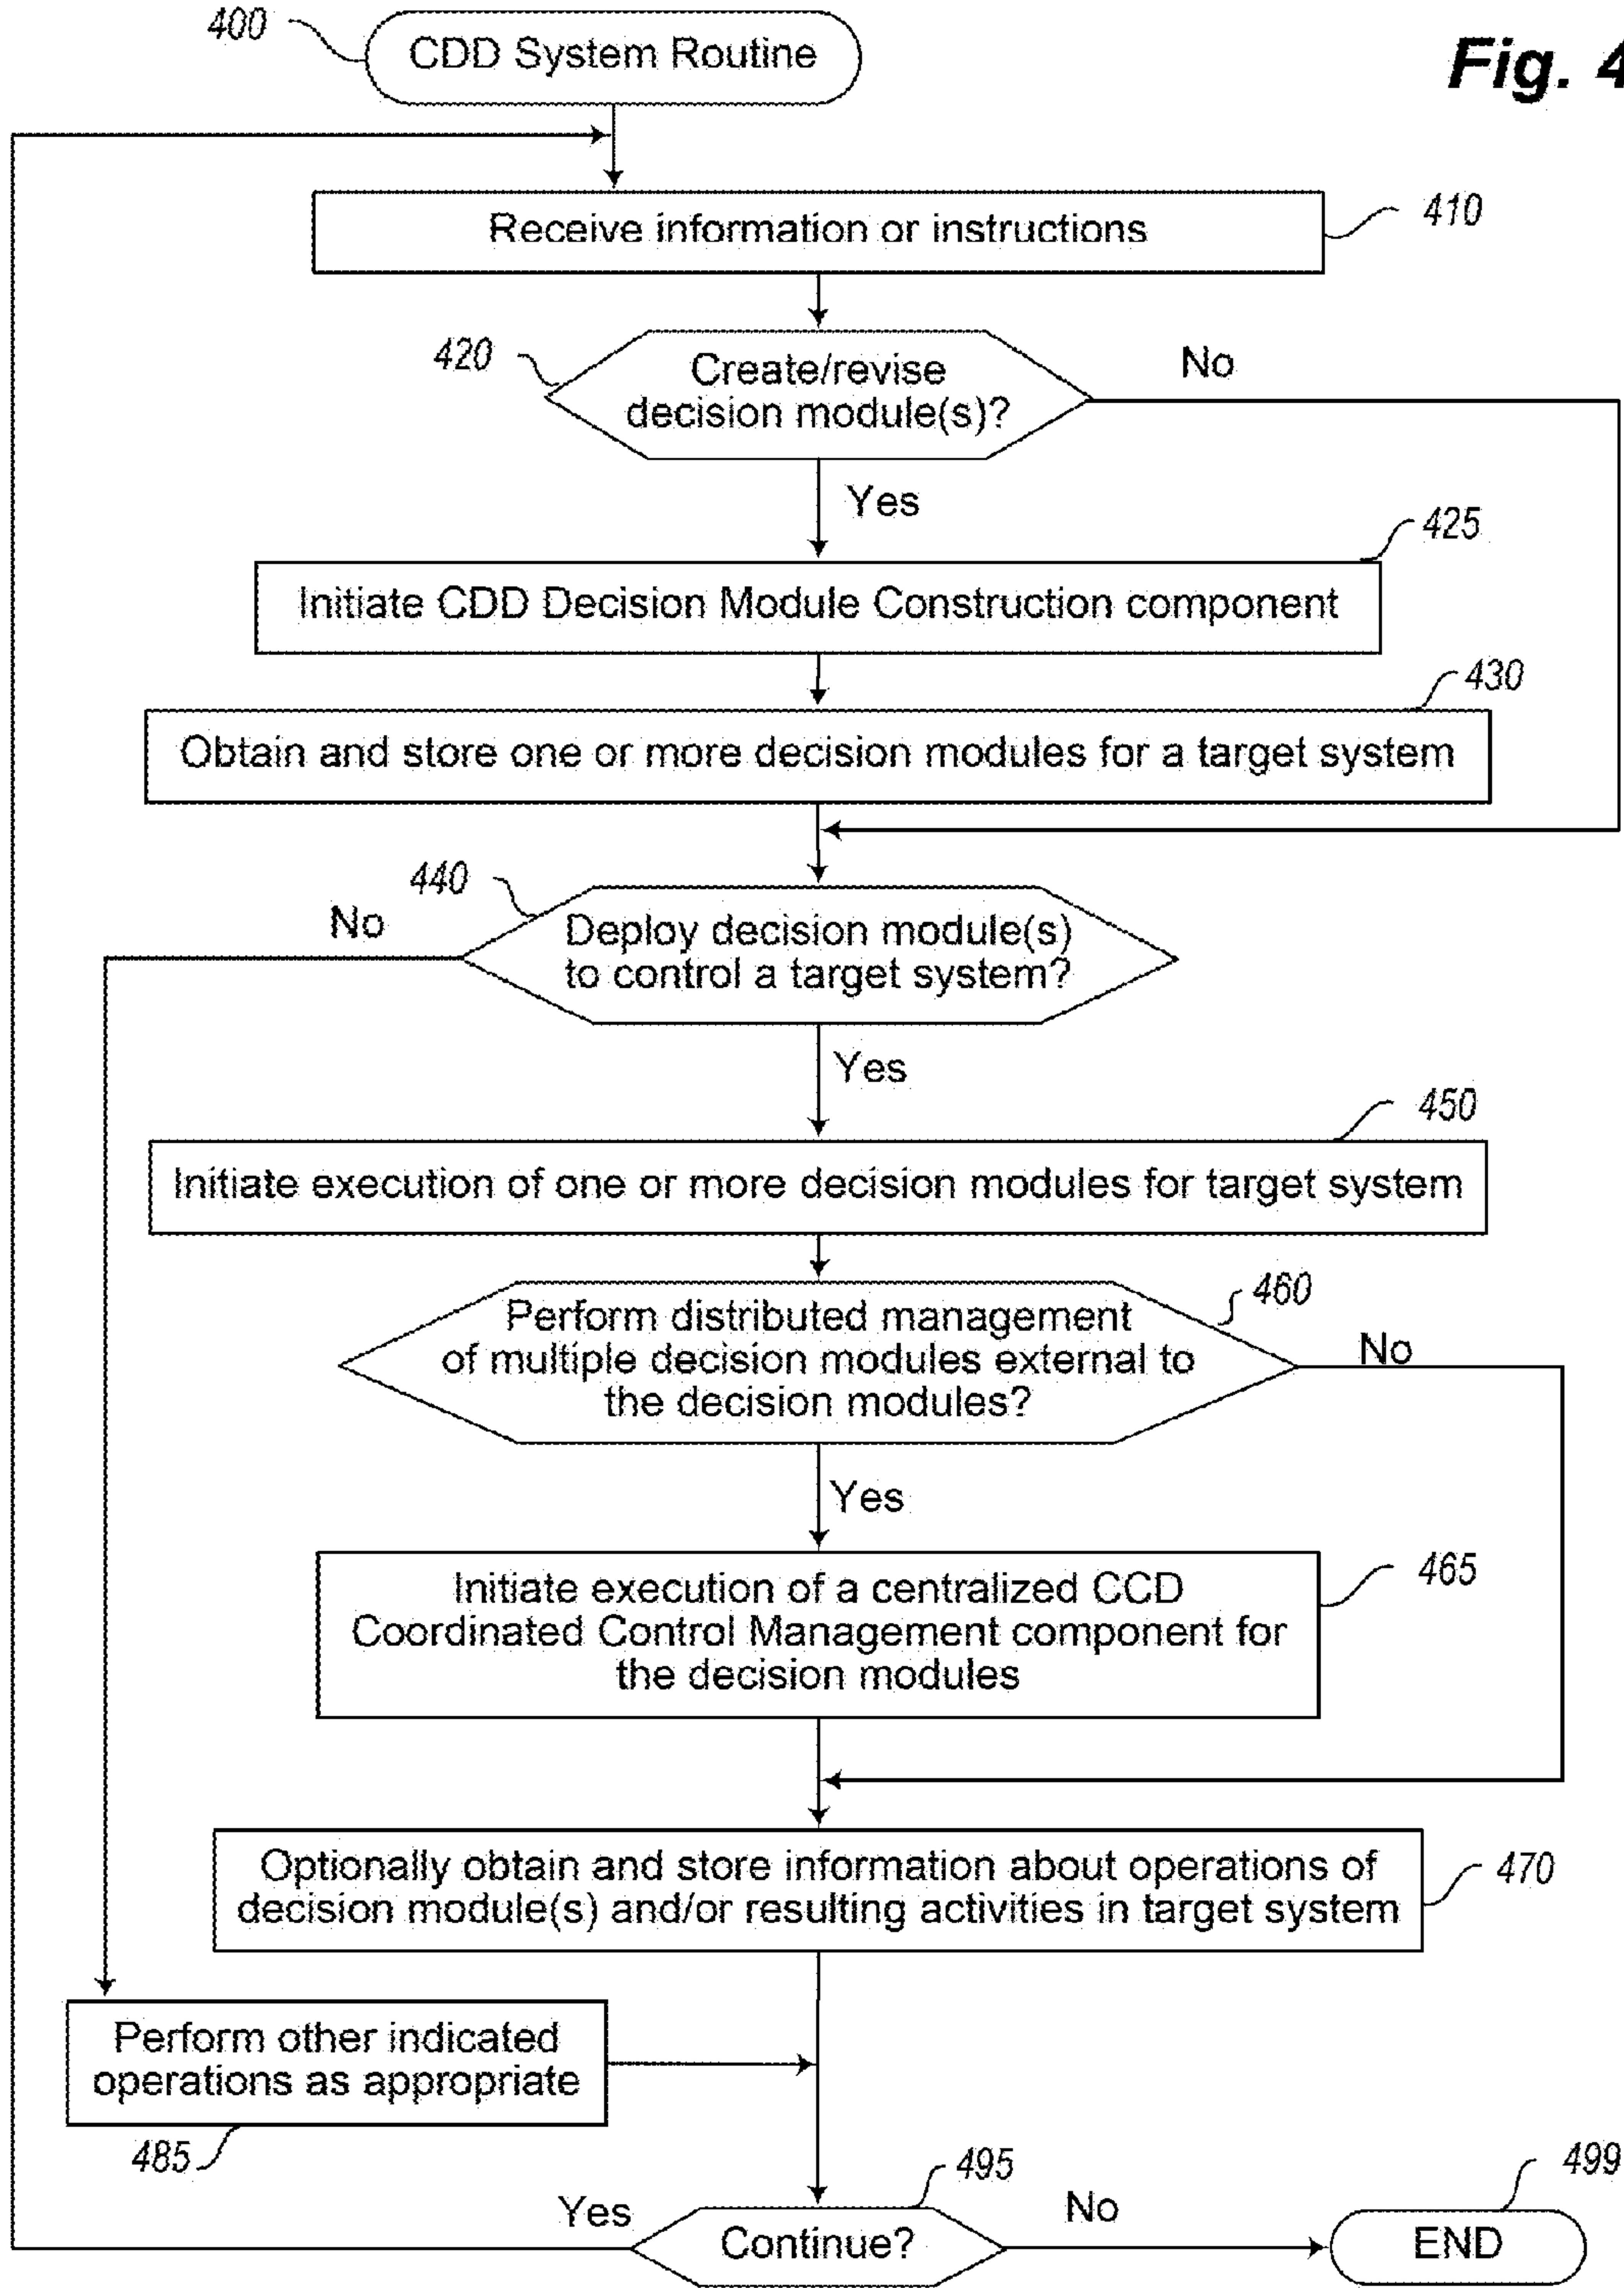


Fig. 5A

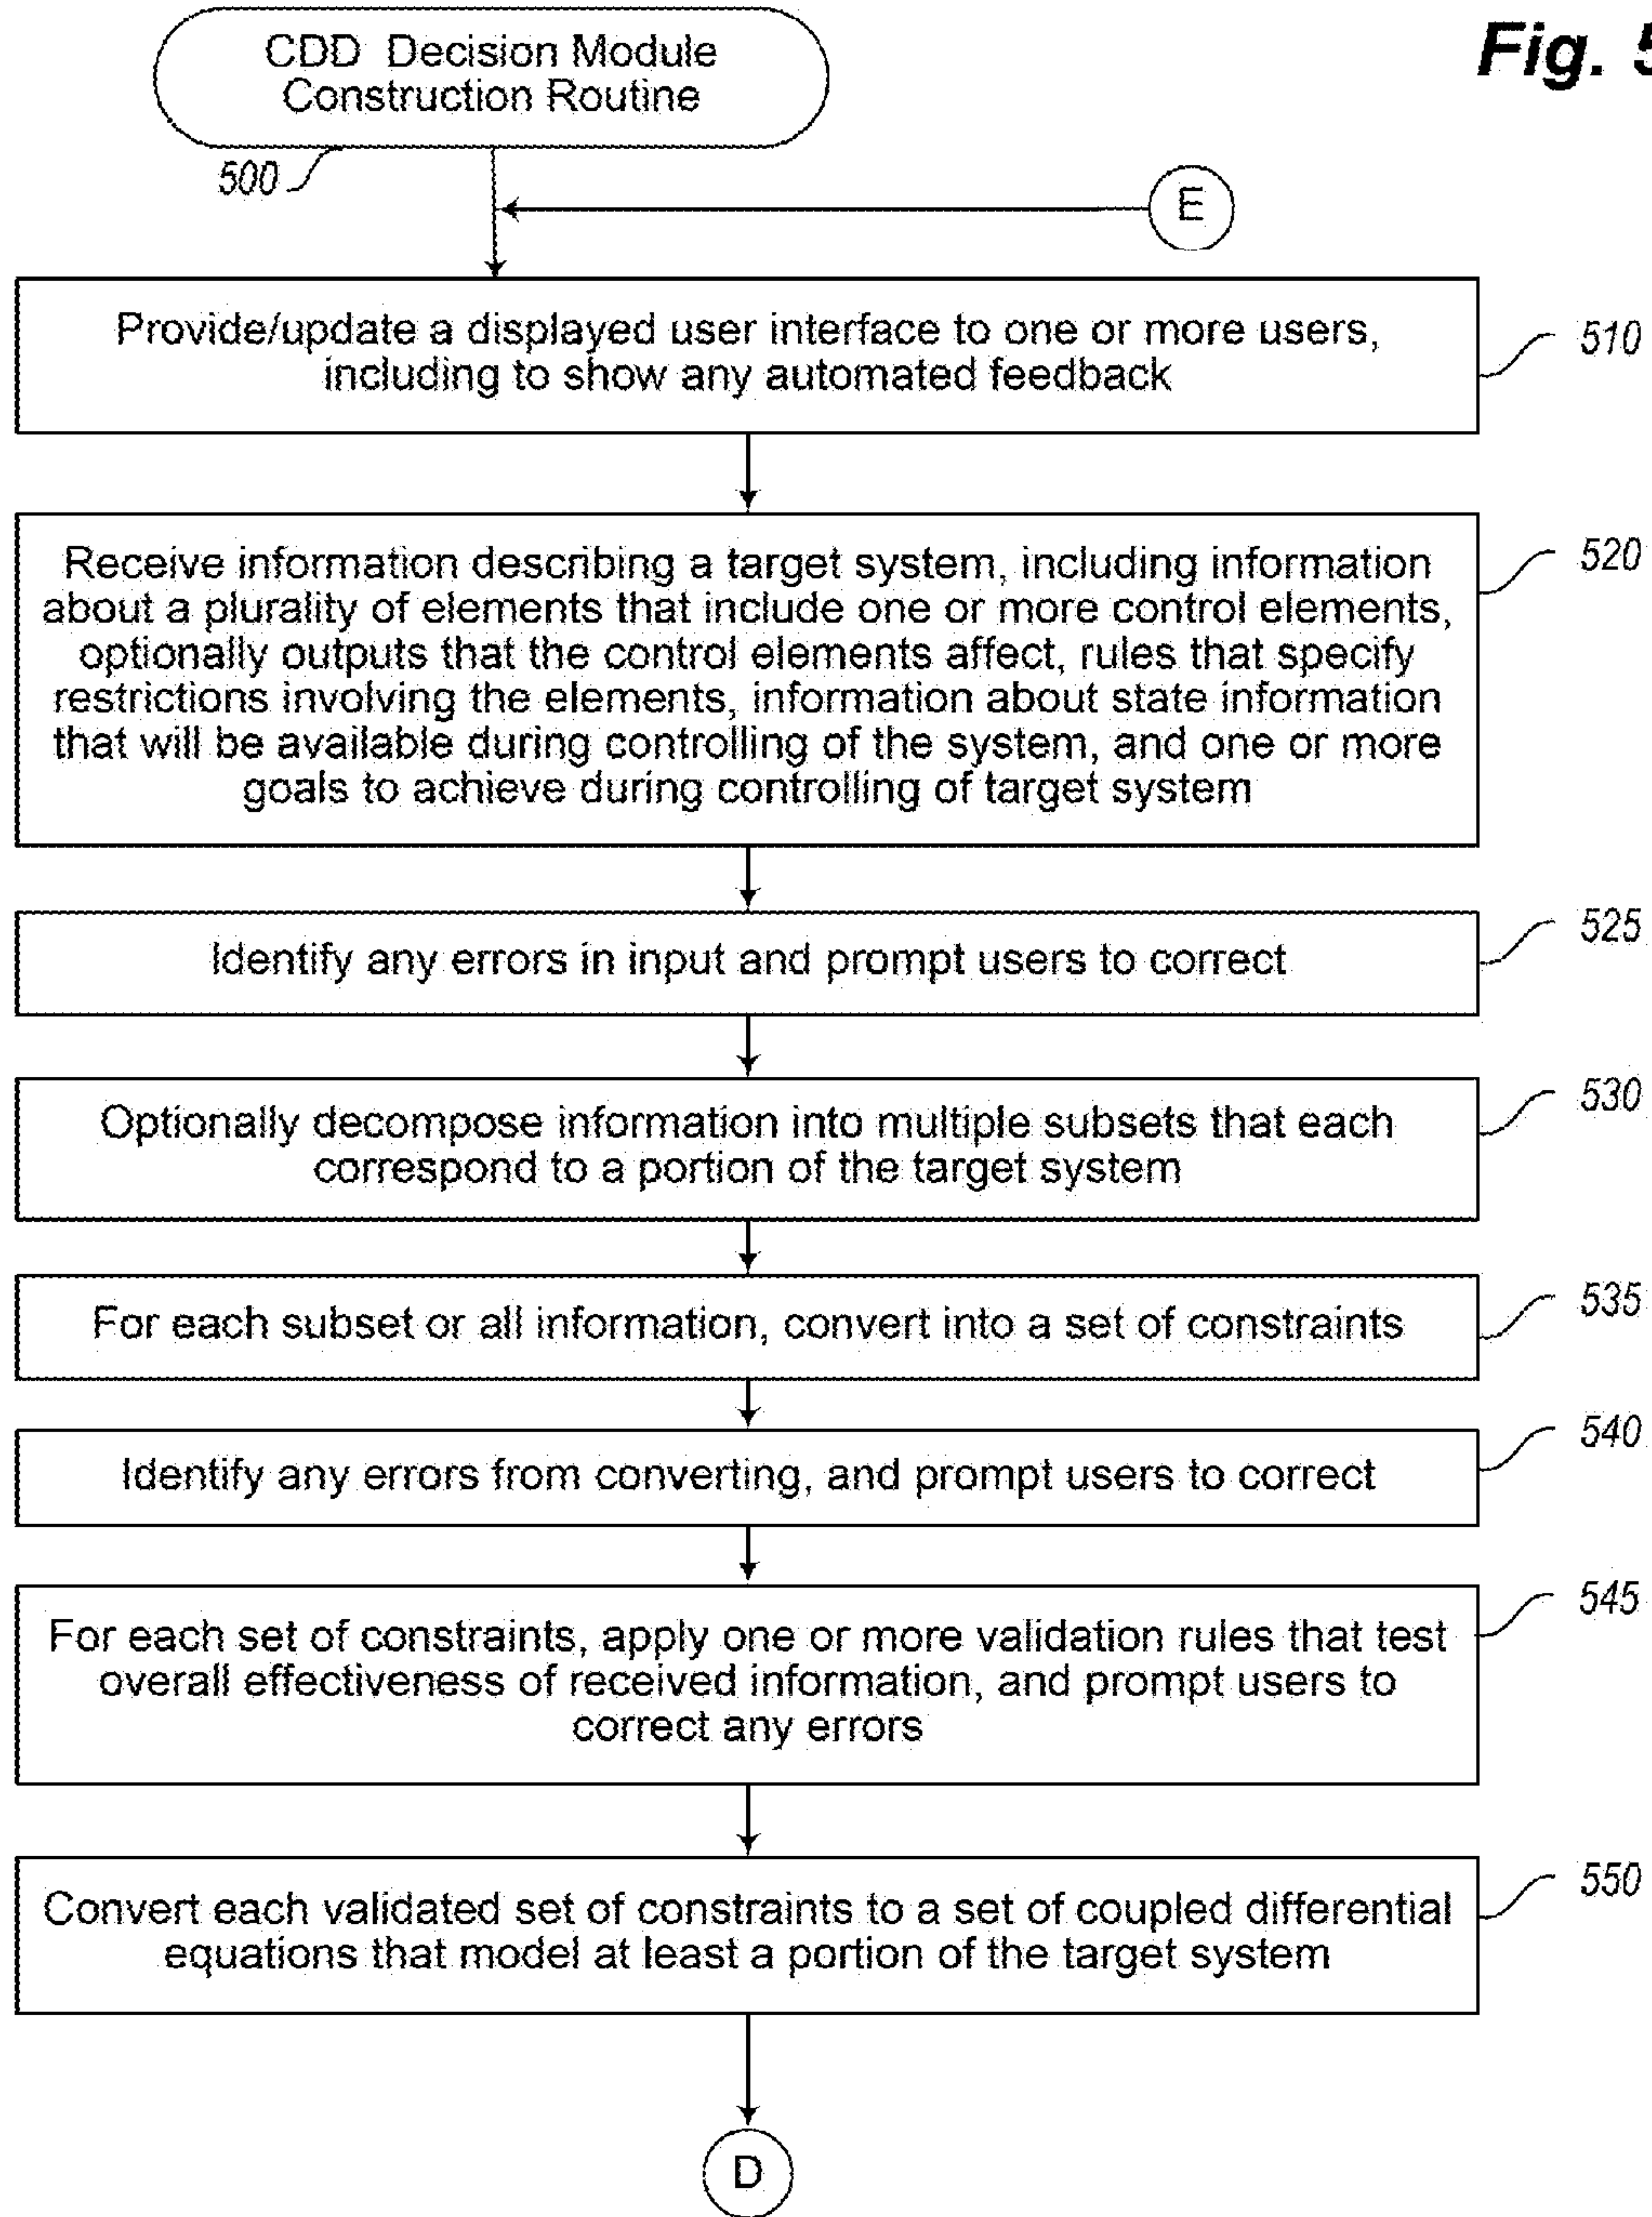
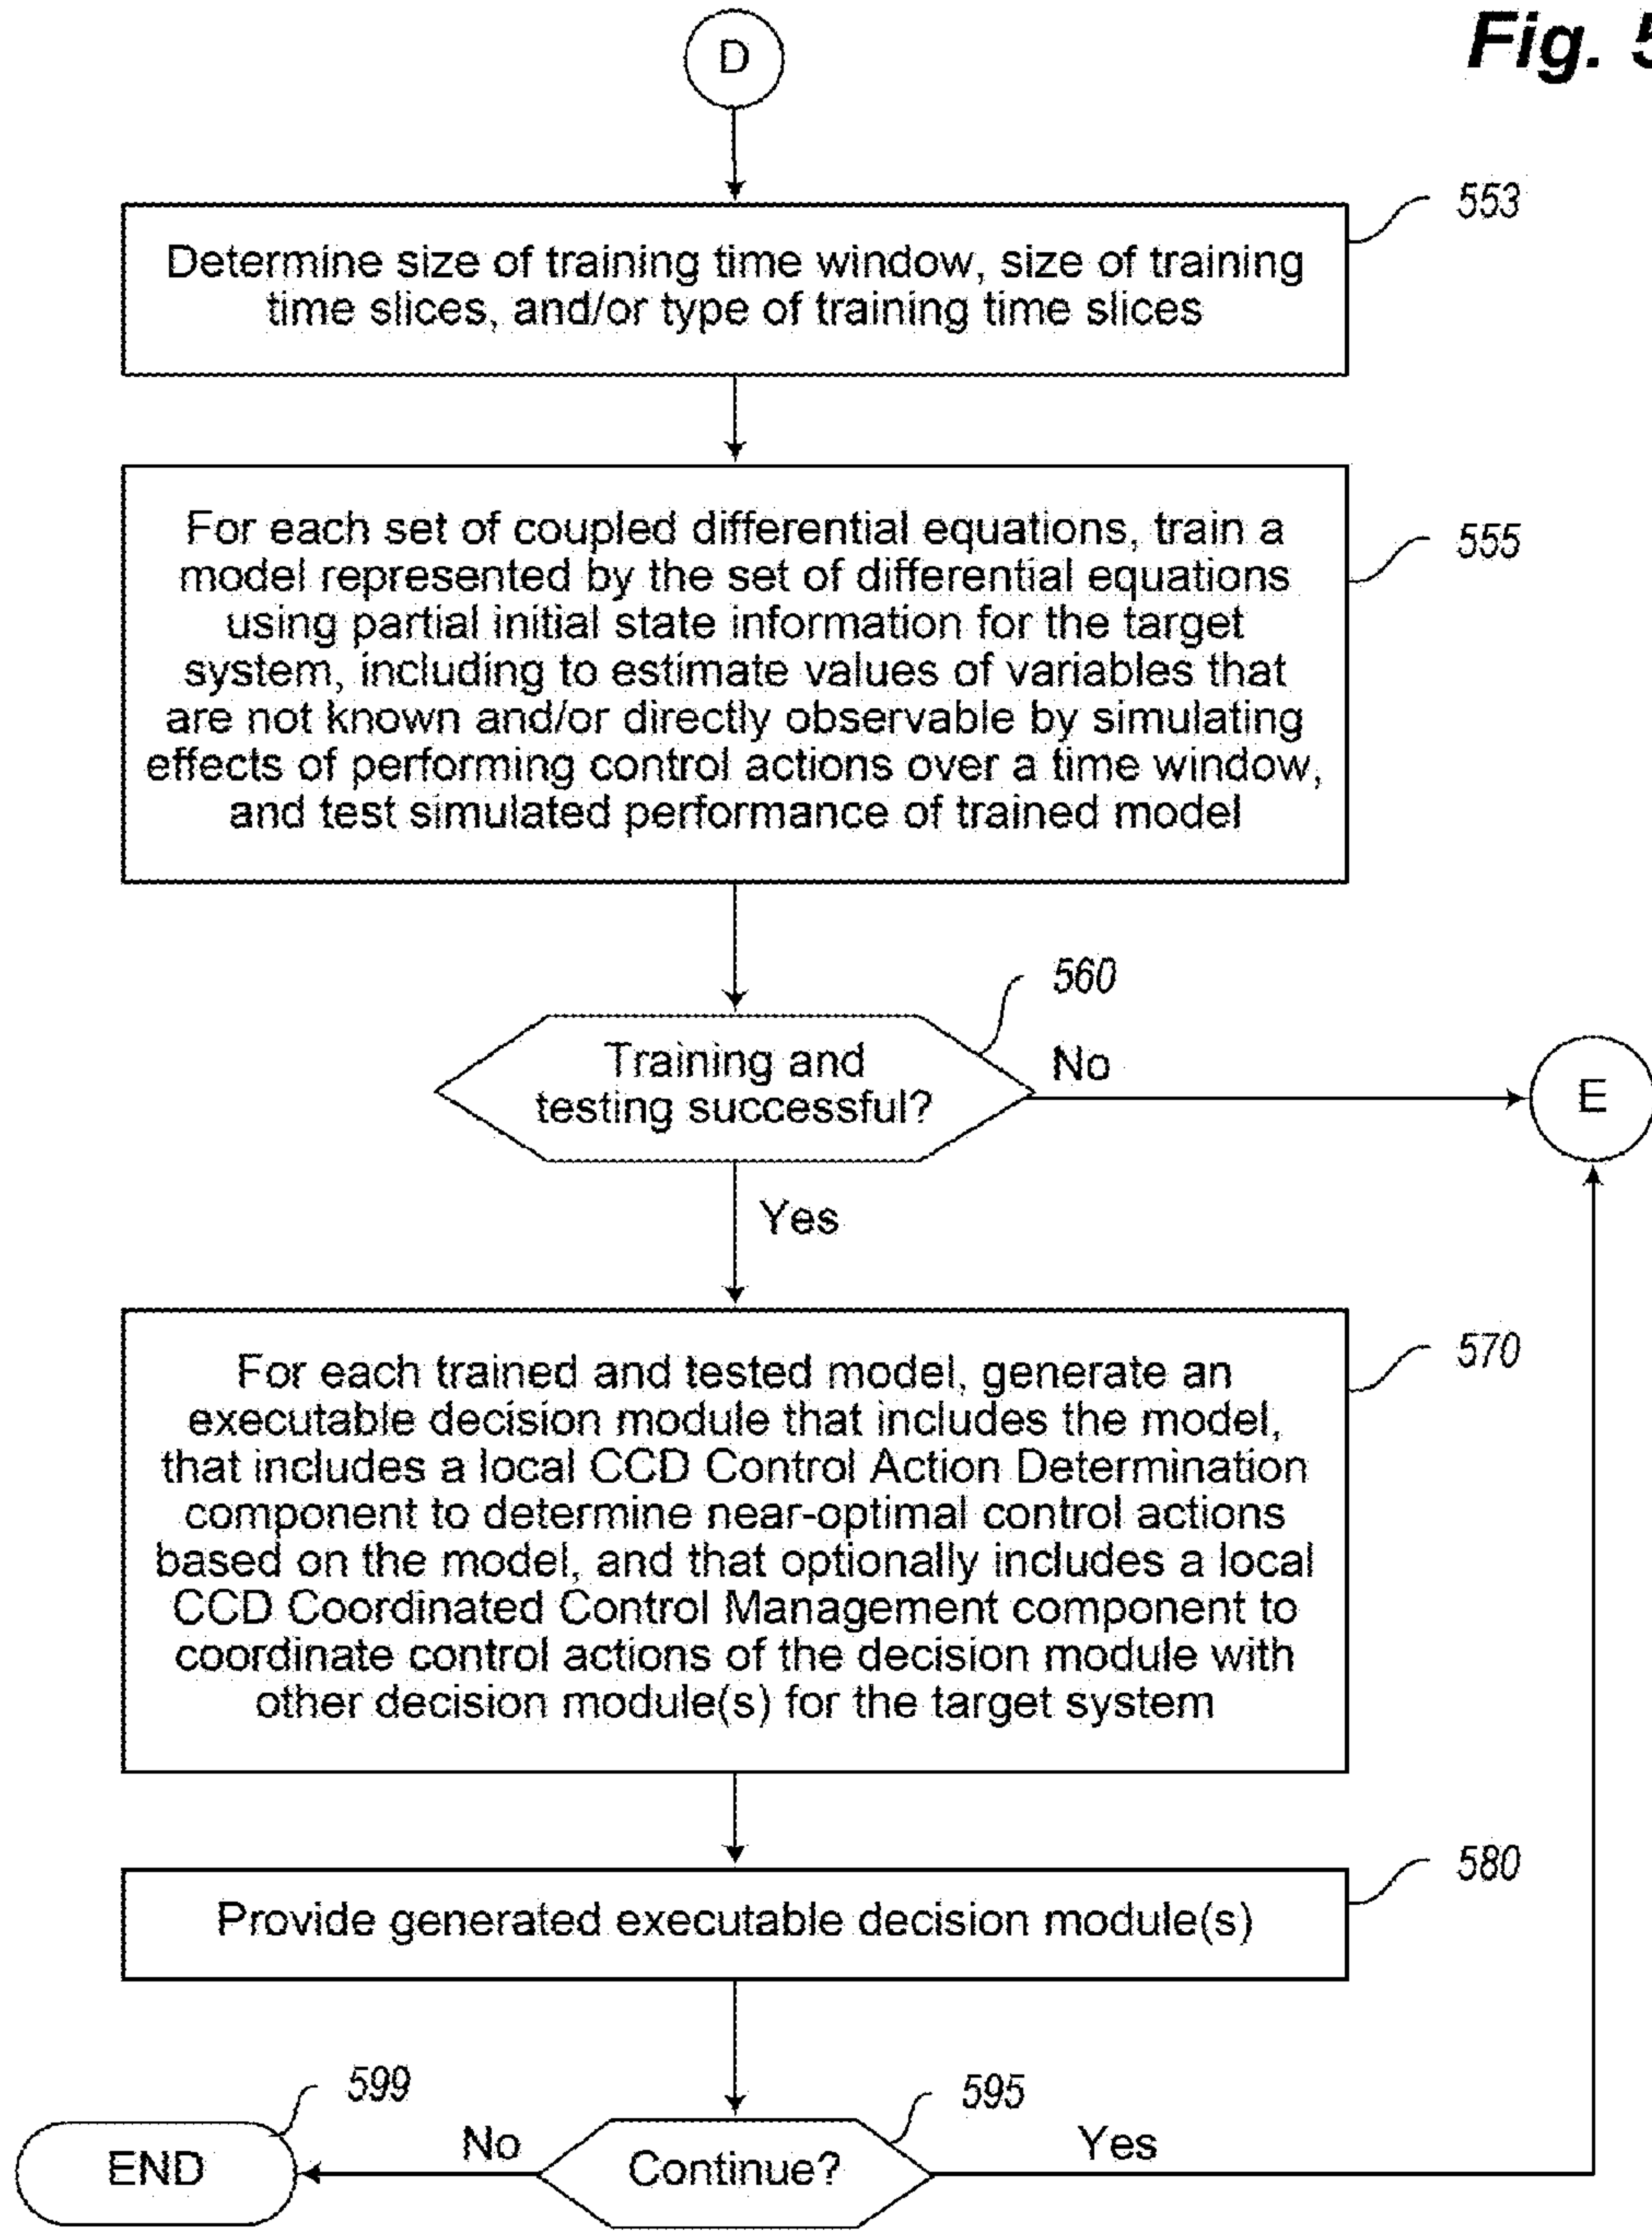


Fig. 5B





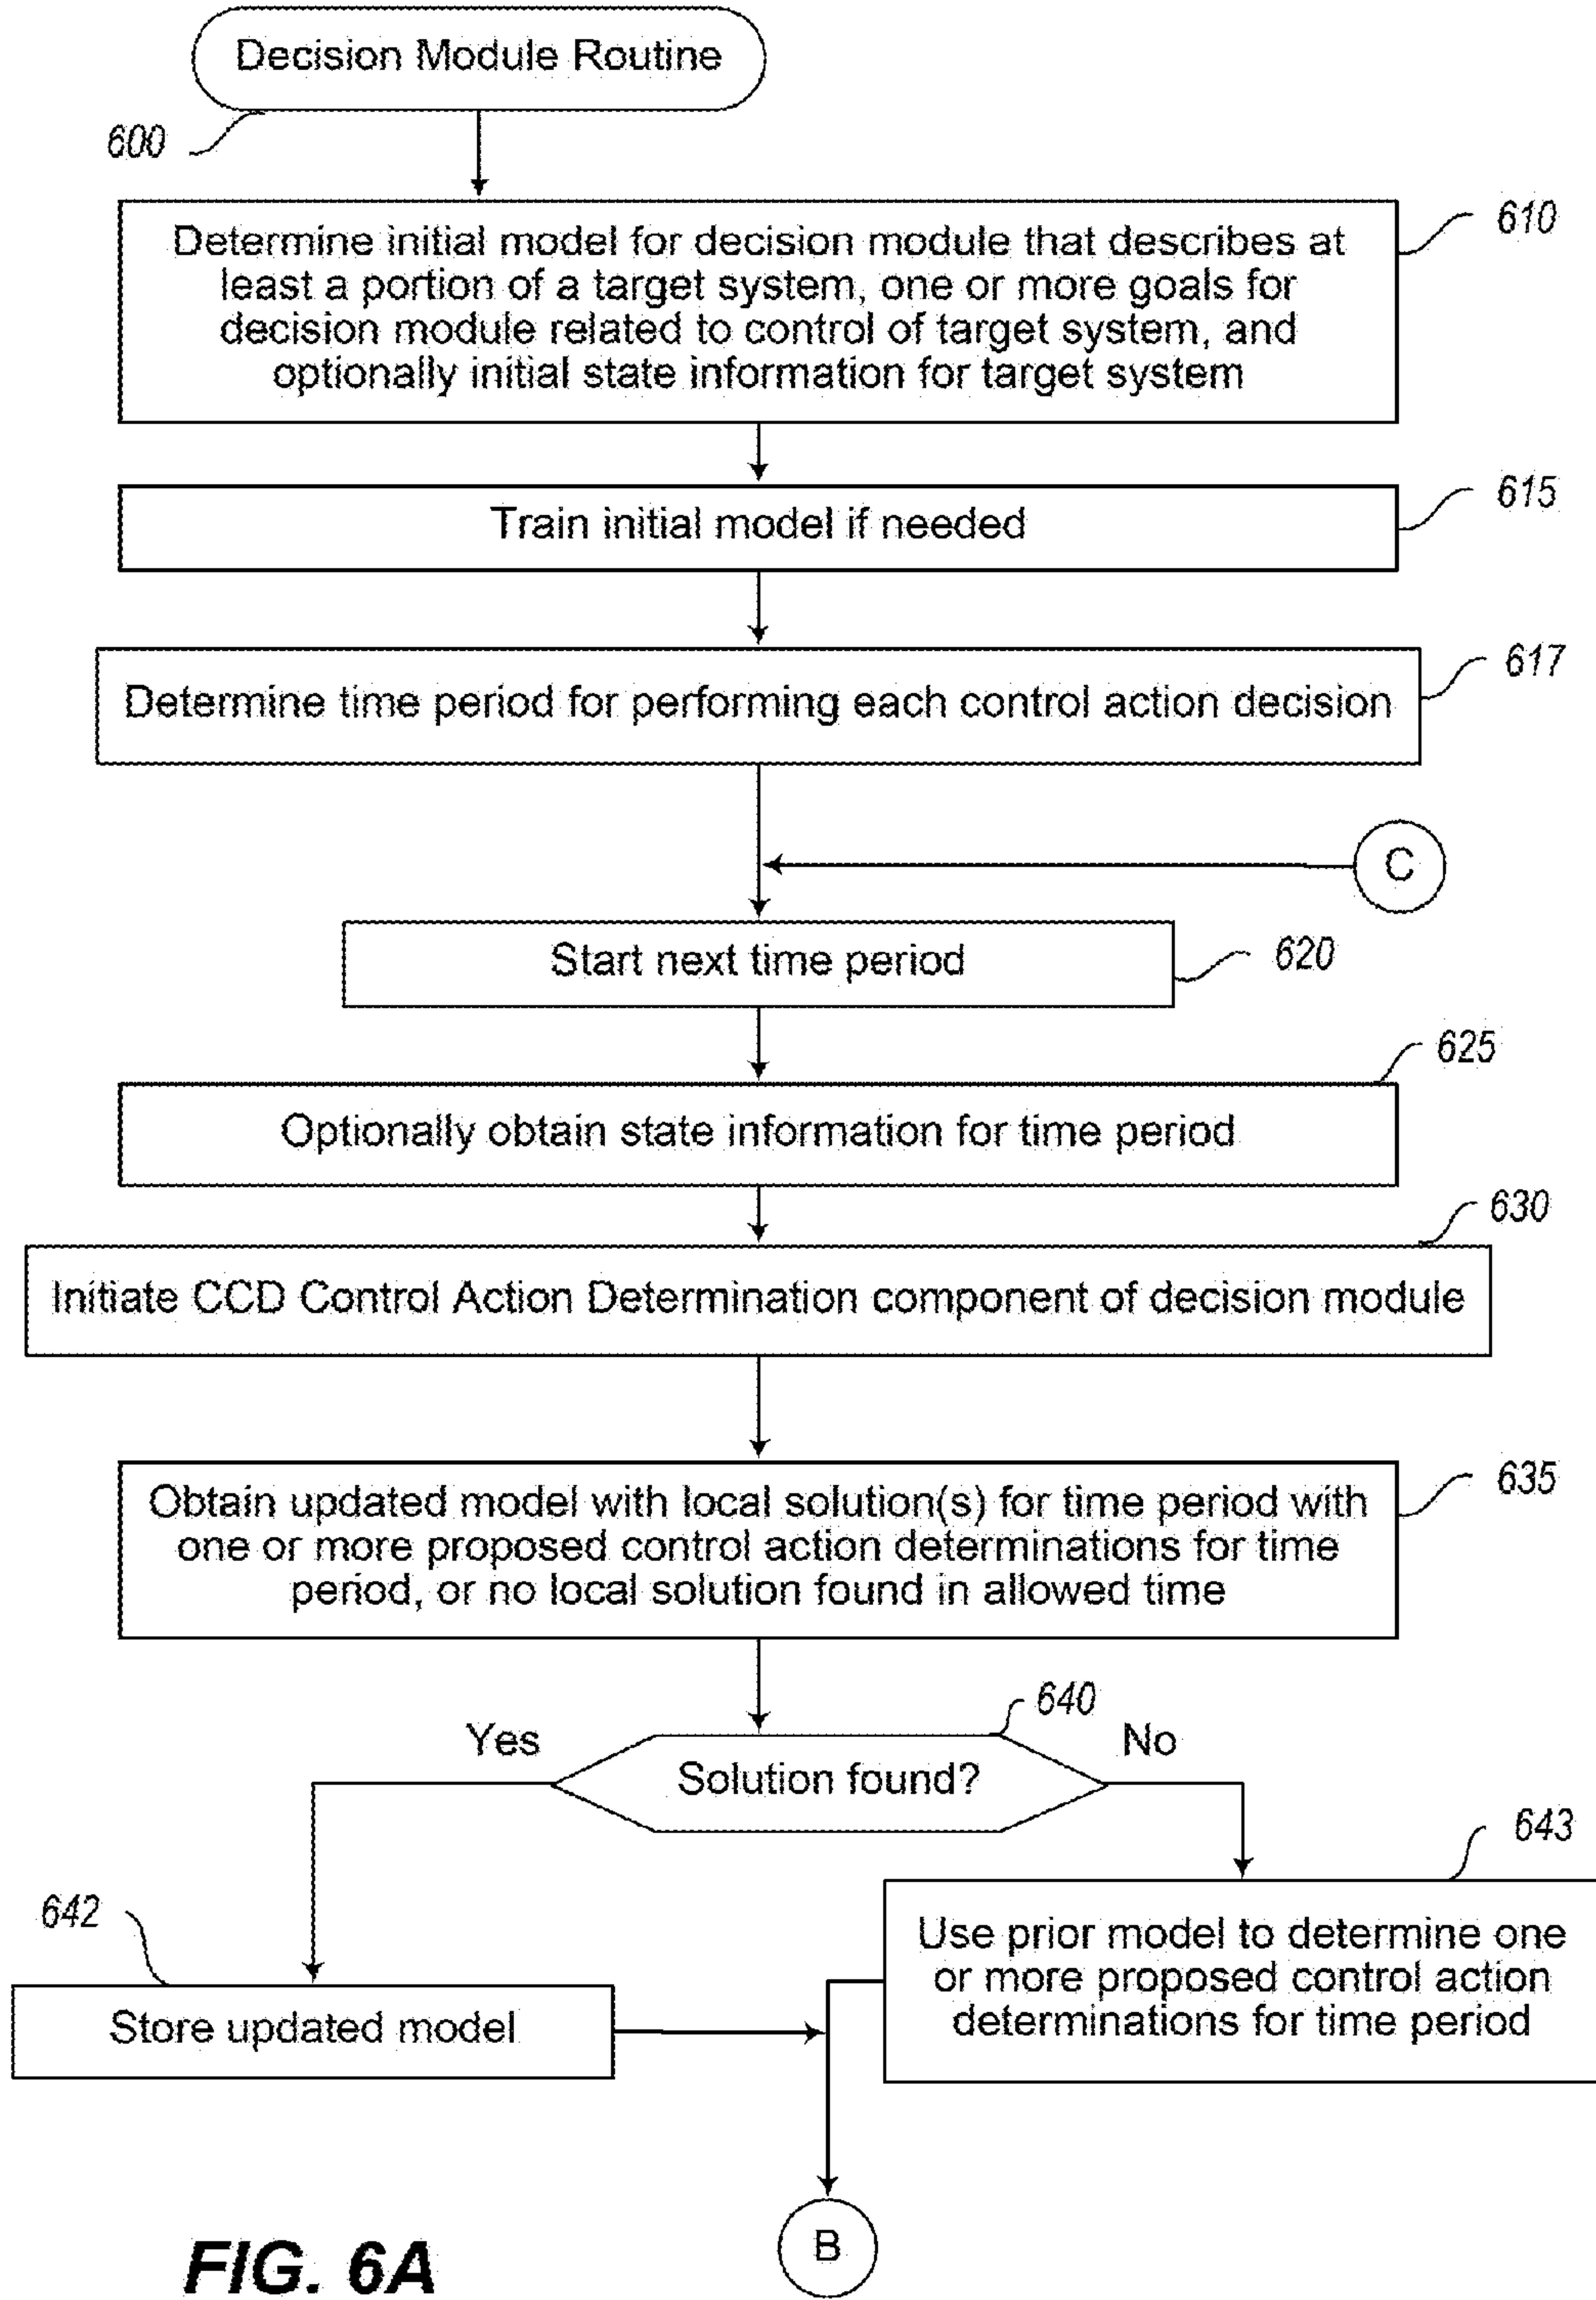
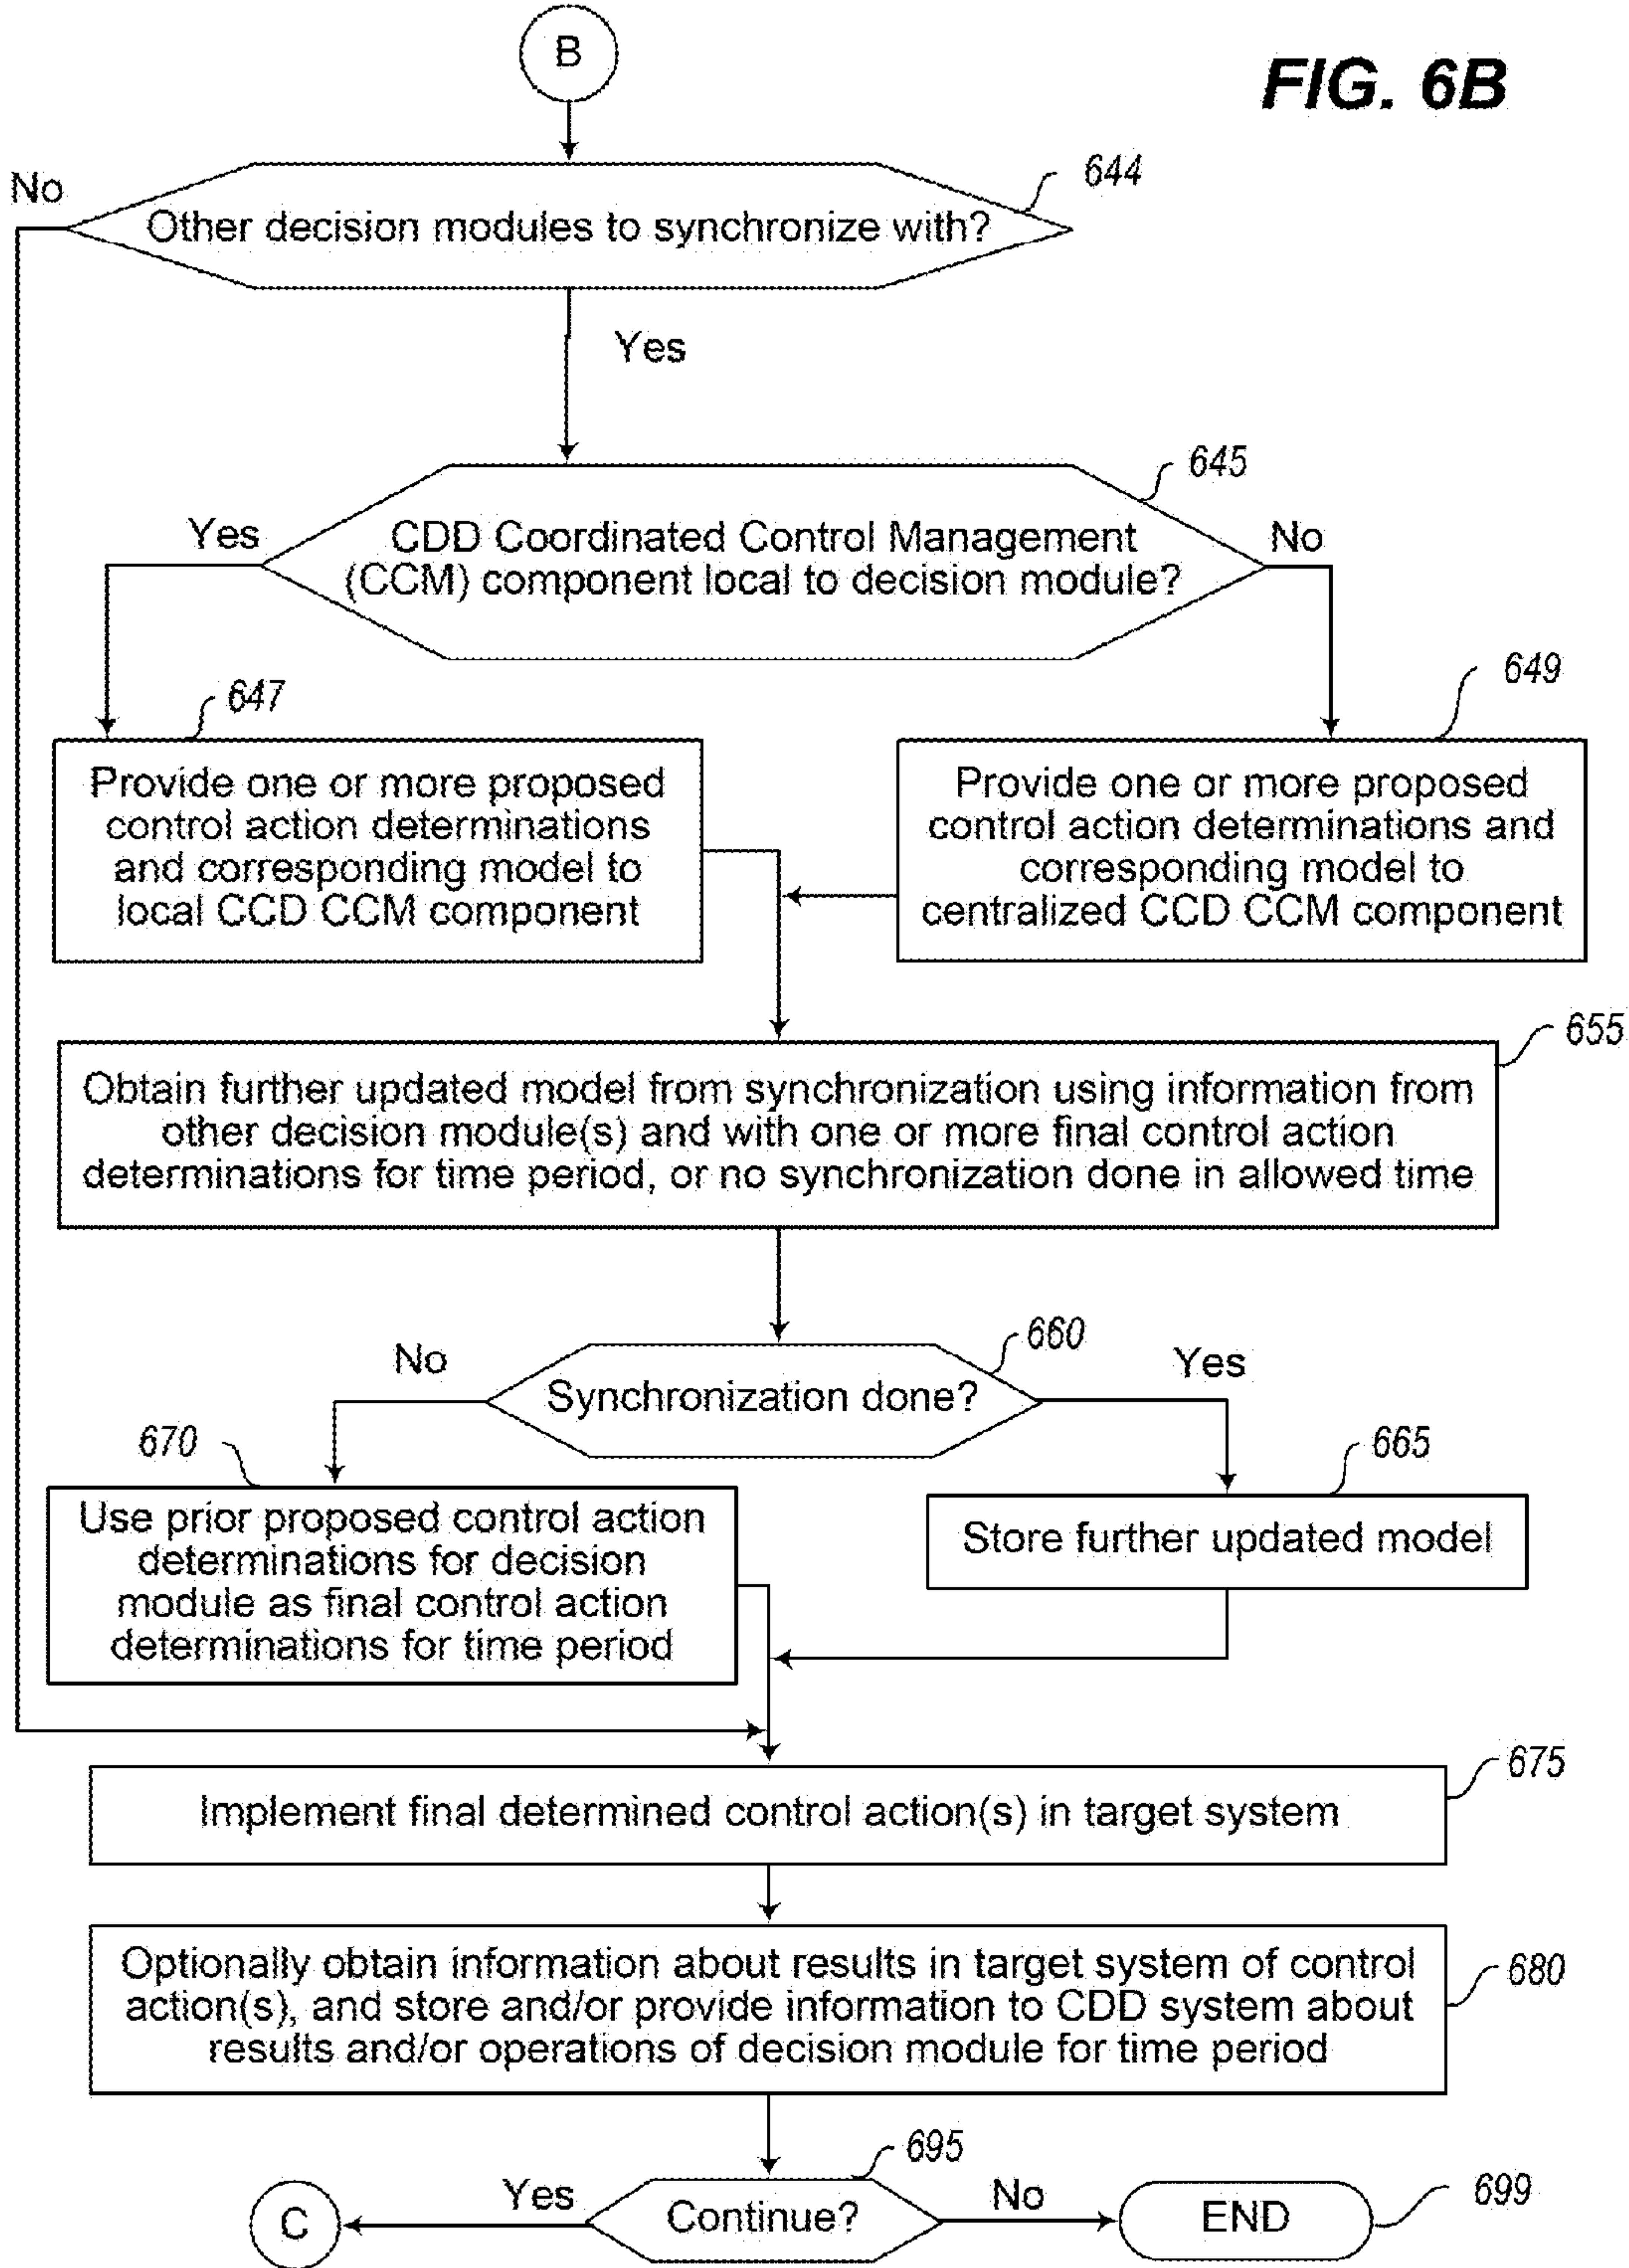


FIG. 6A

FIG. 6B







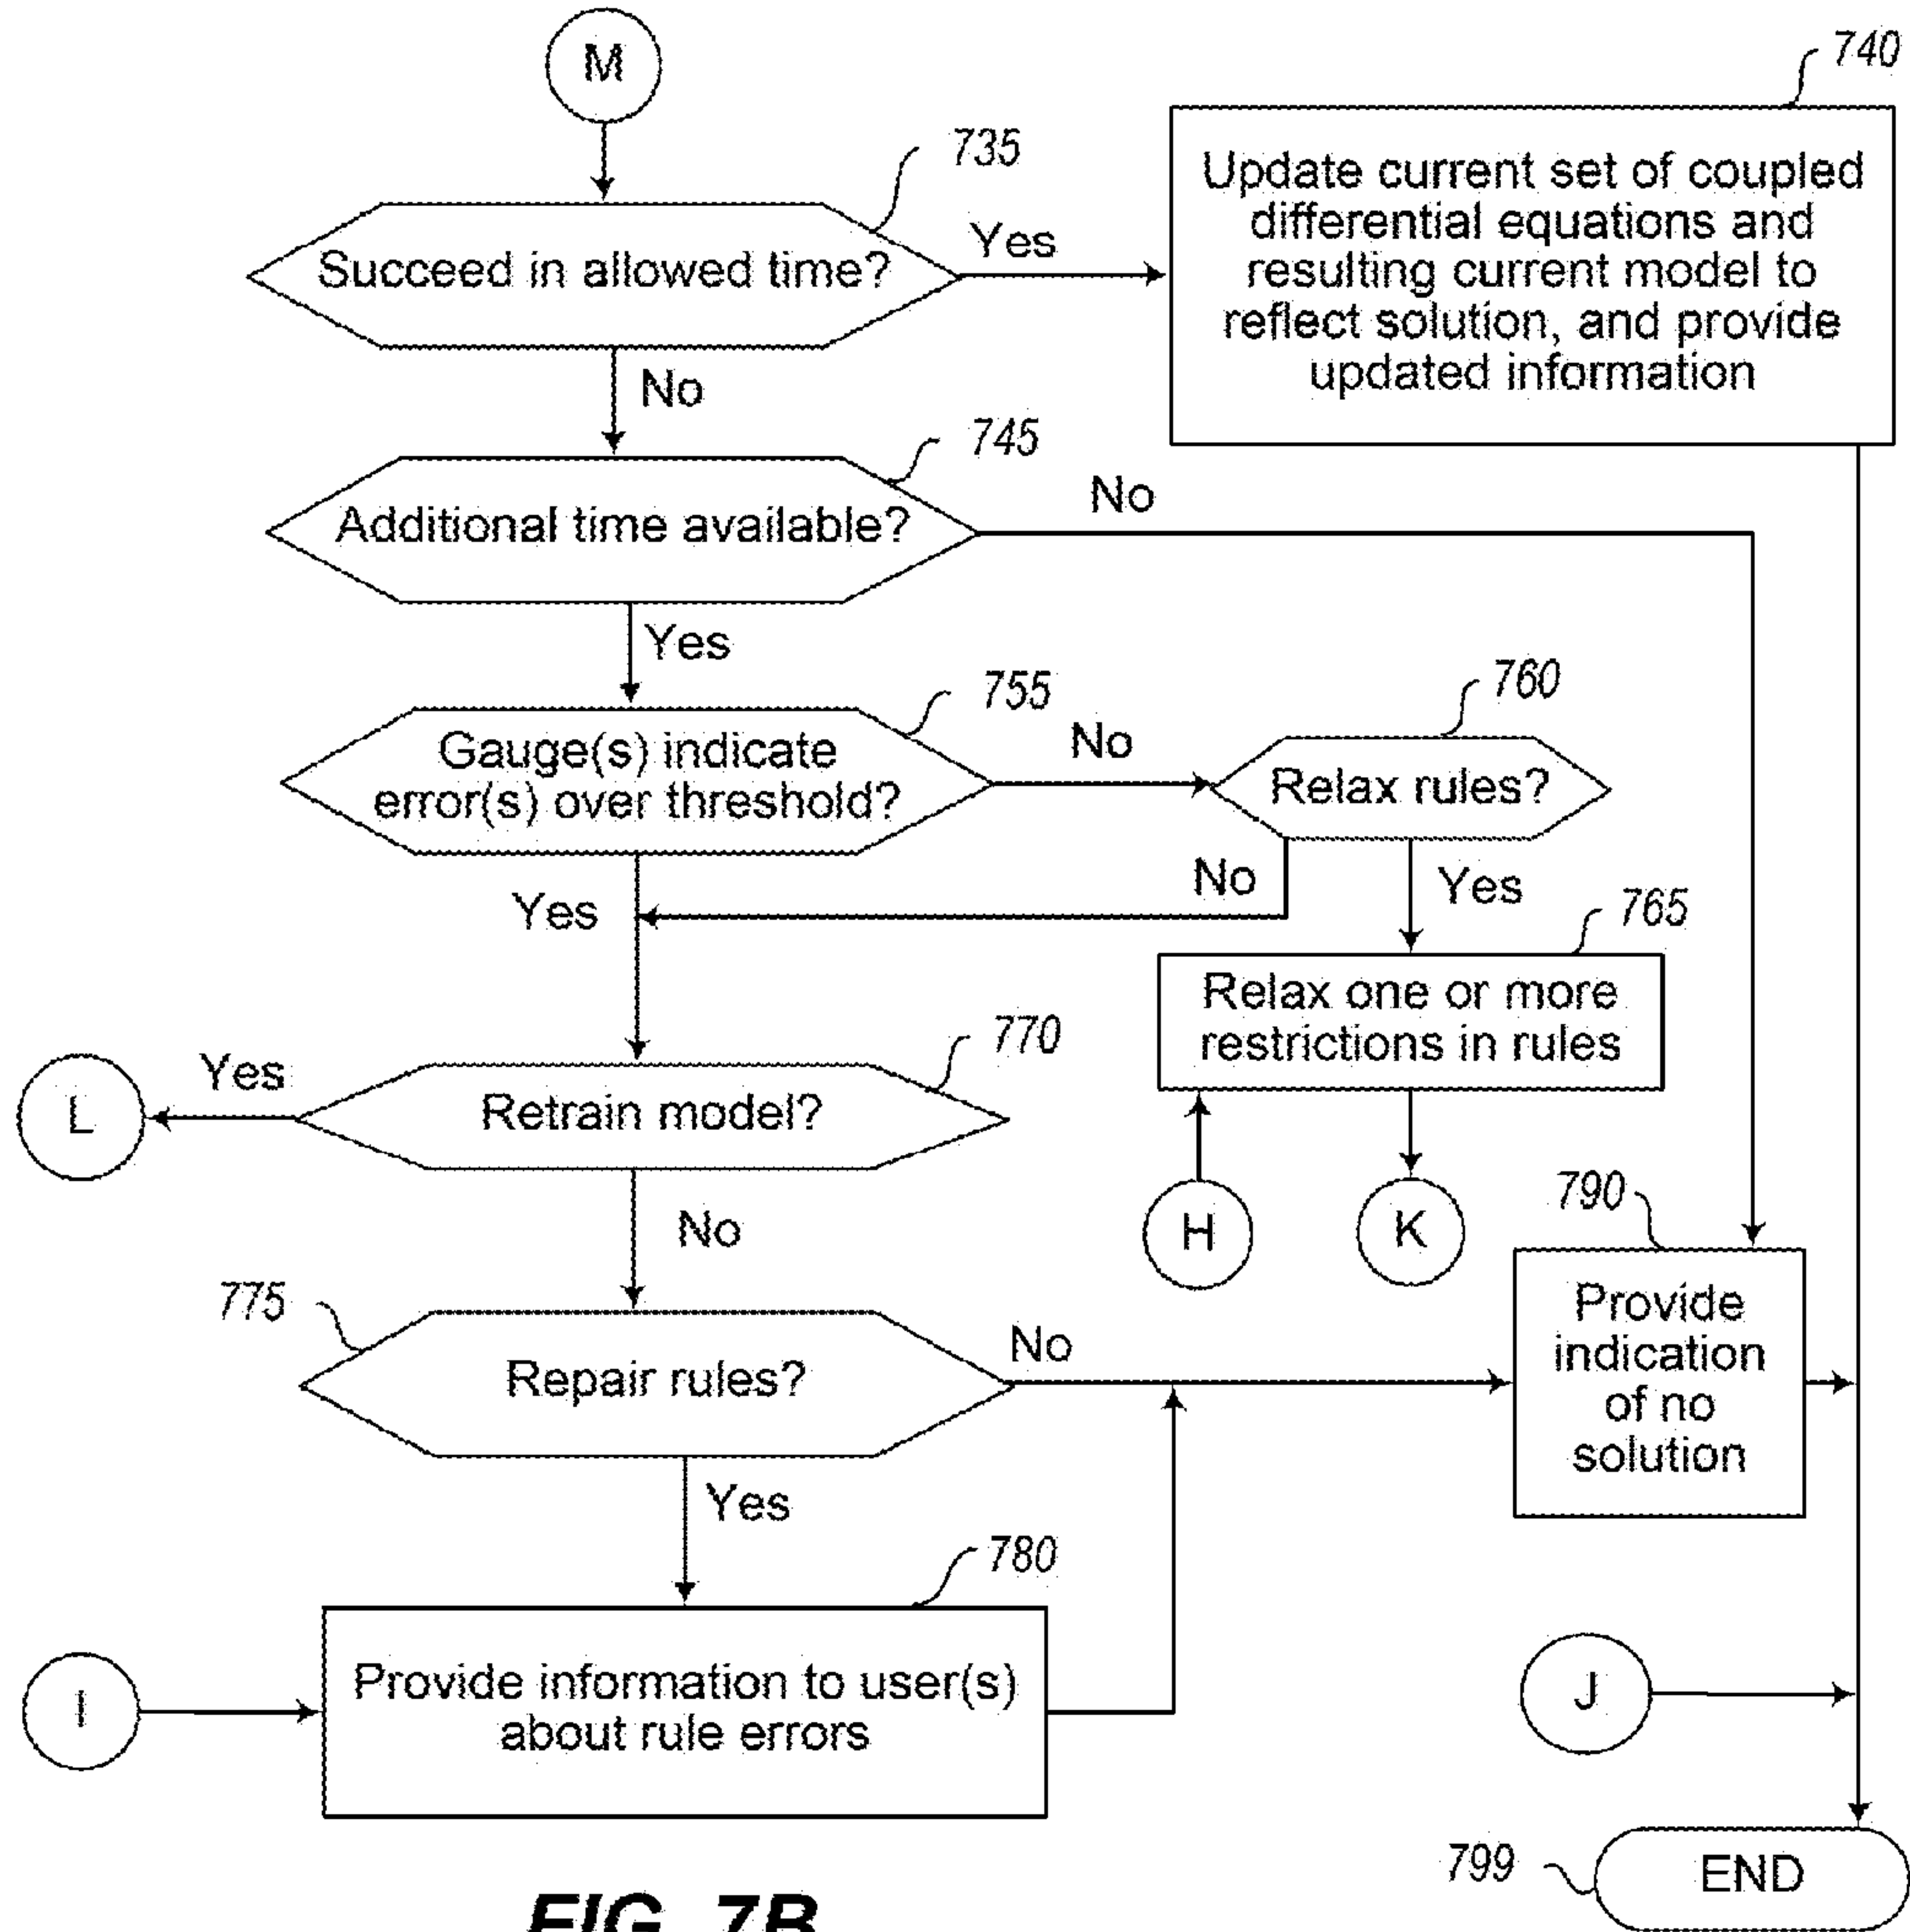


FIG. 7B

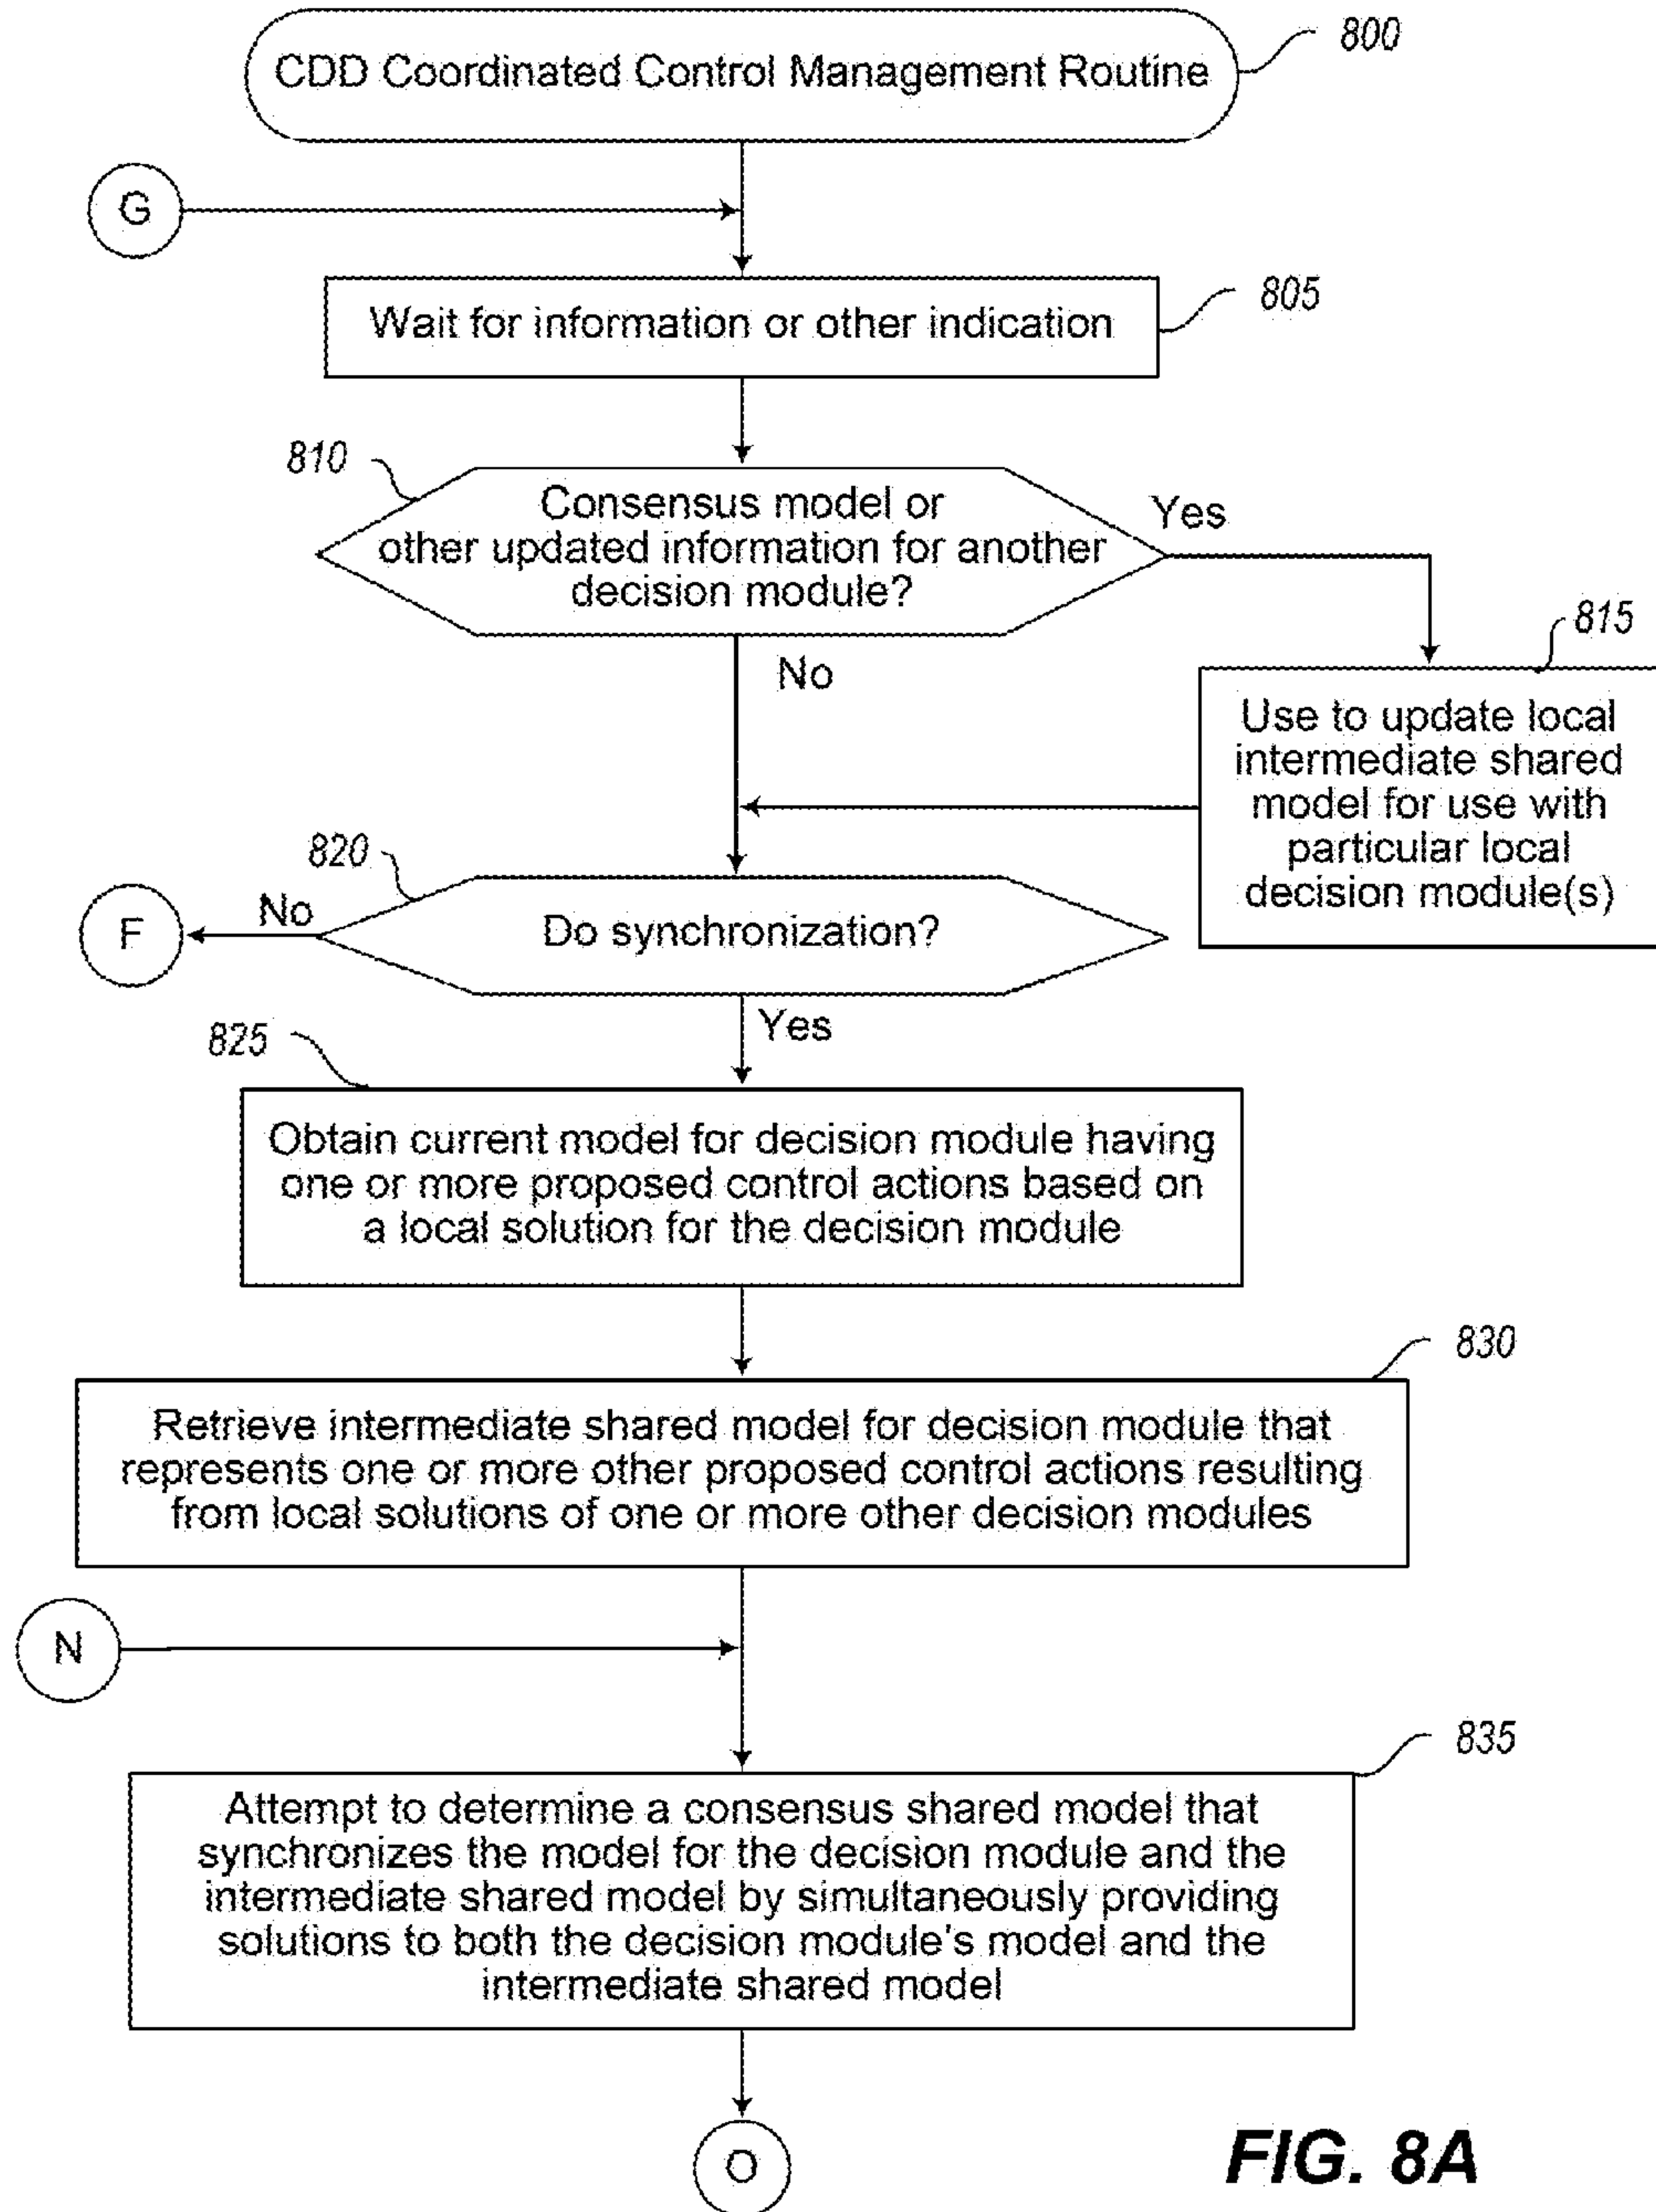


FIG. 8A

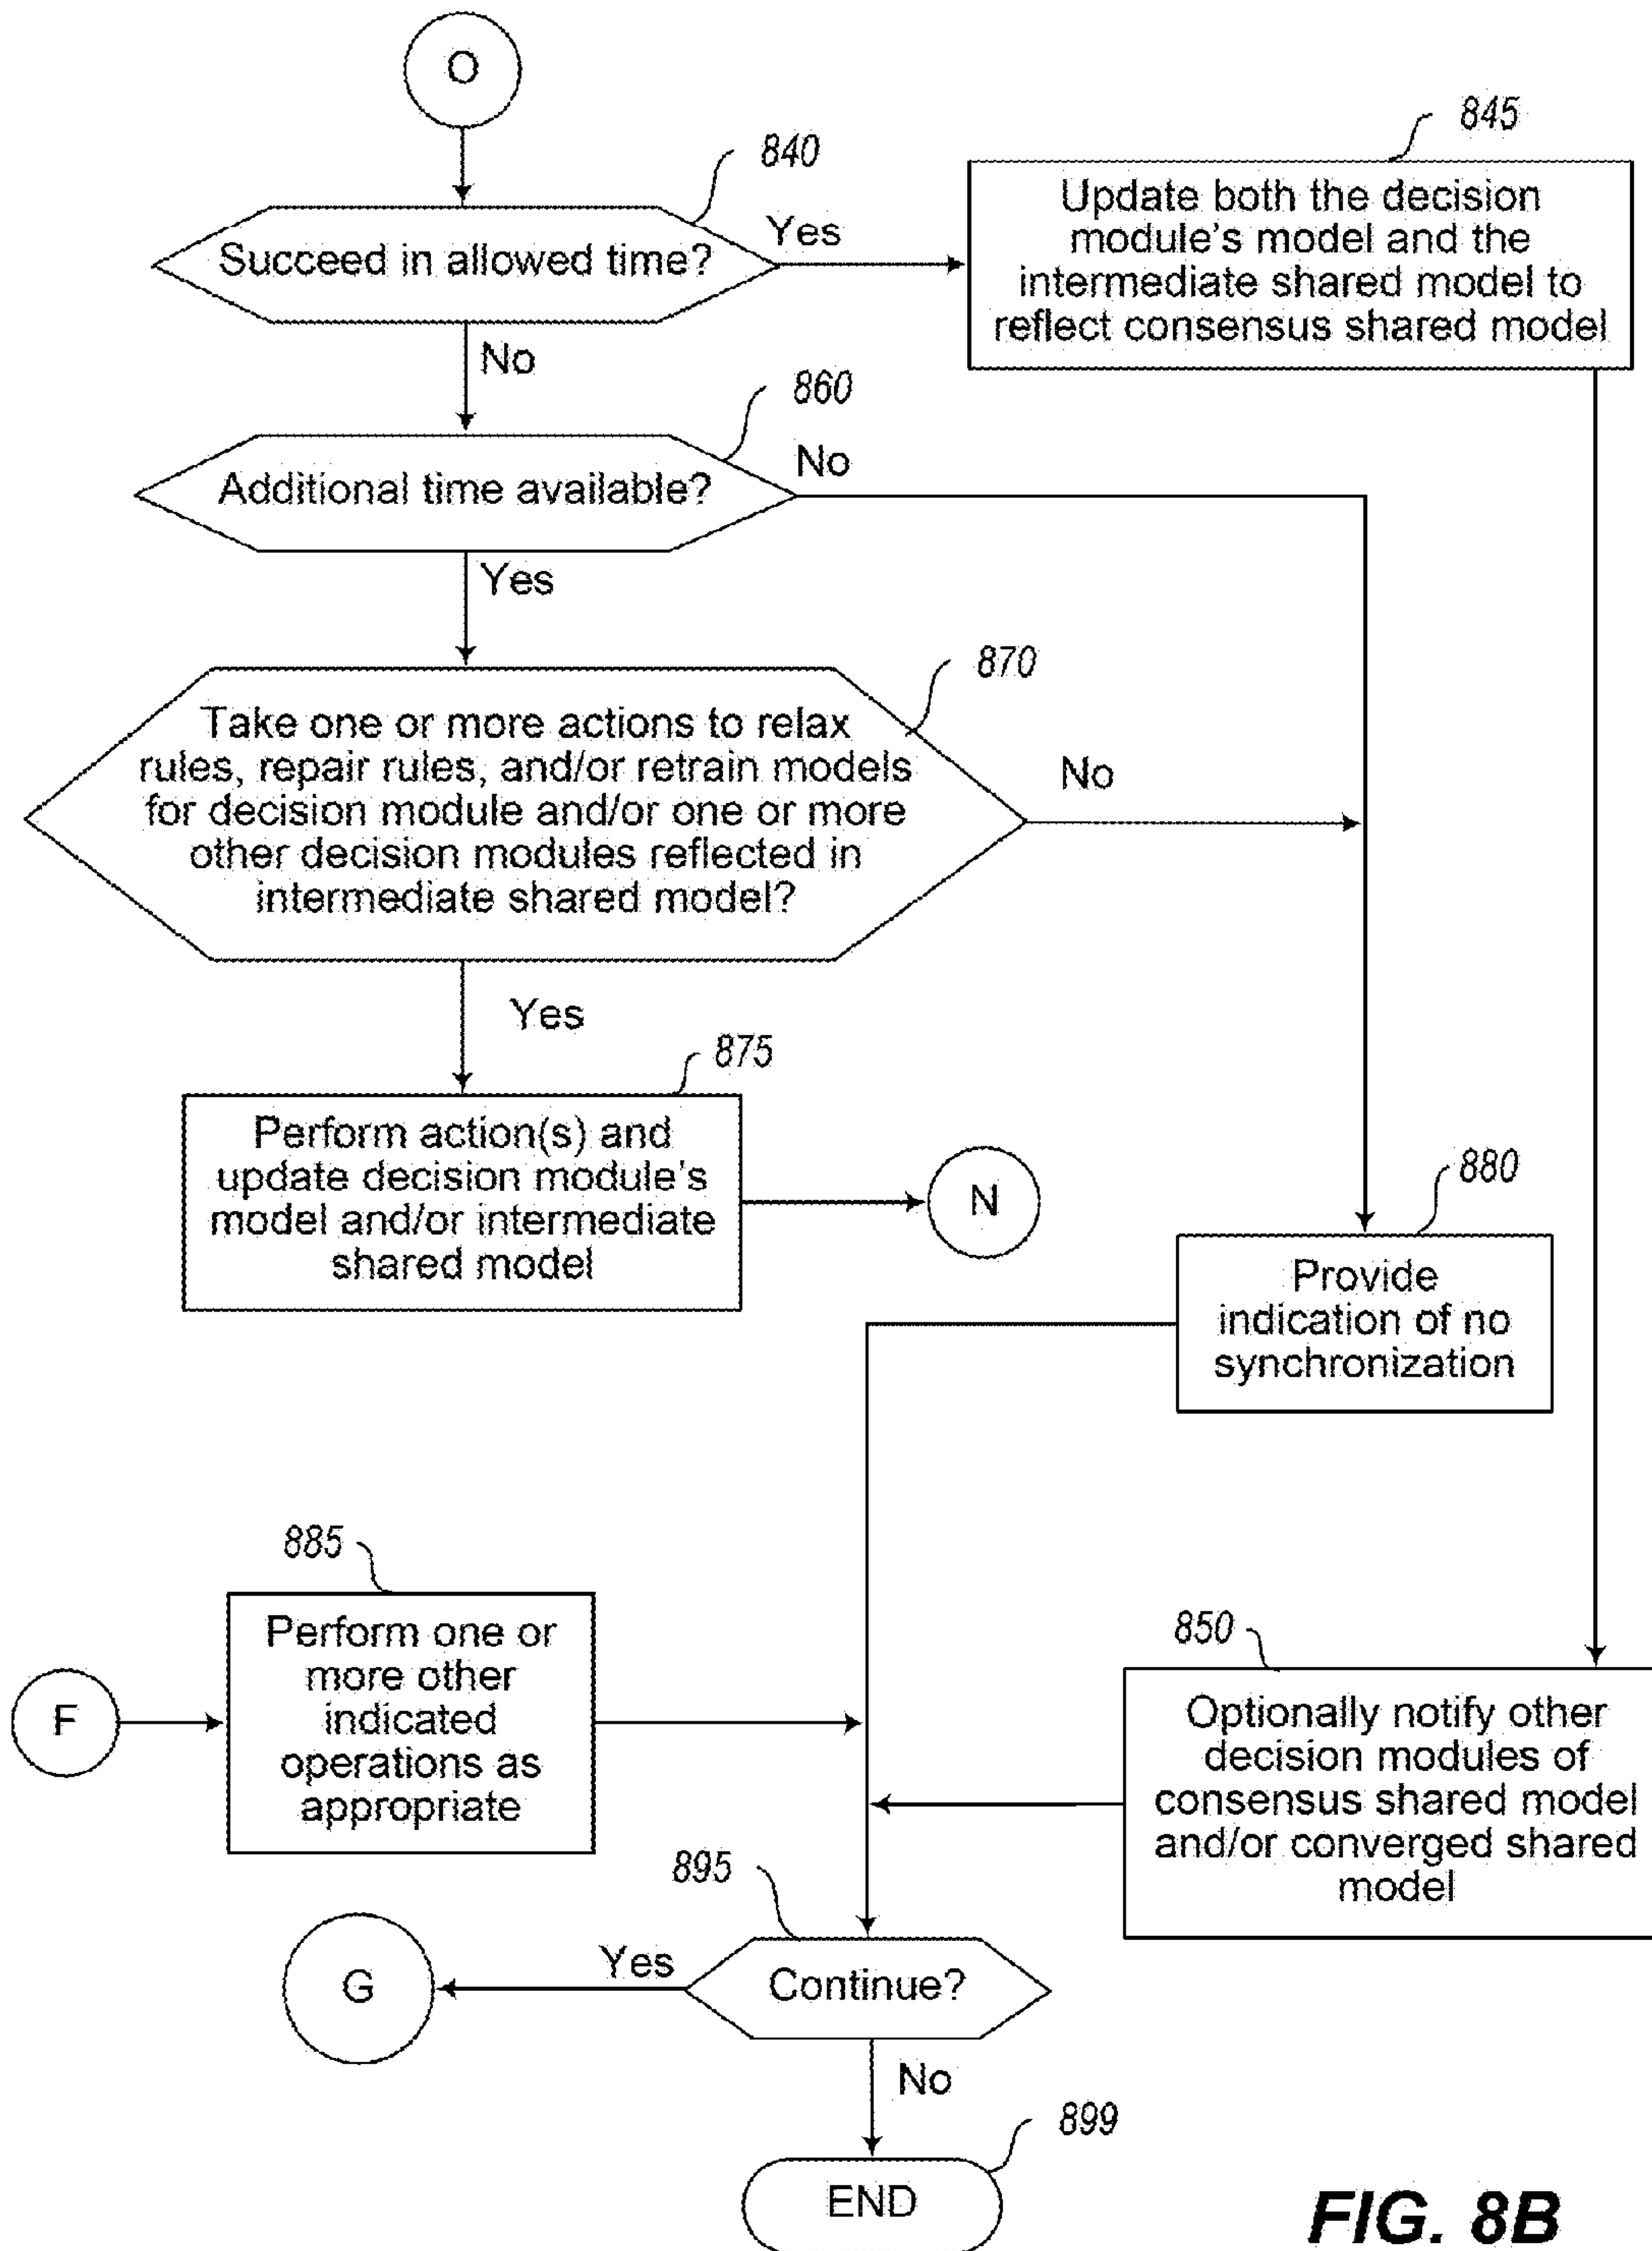
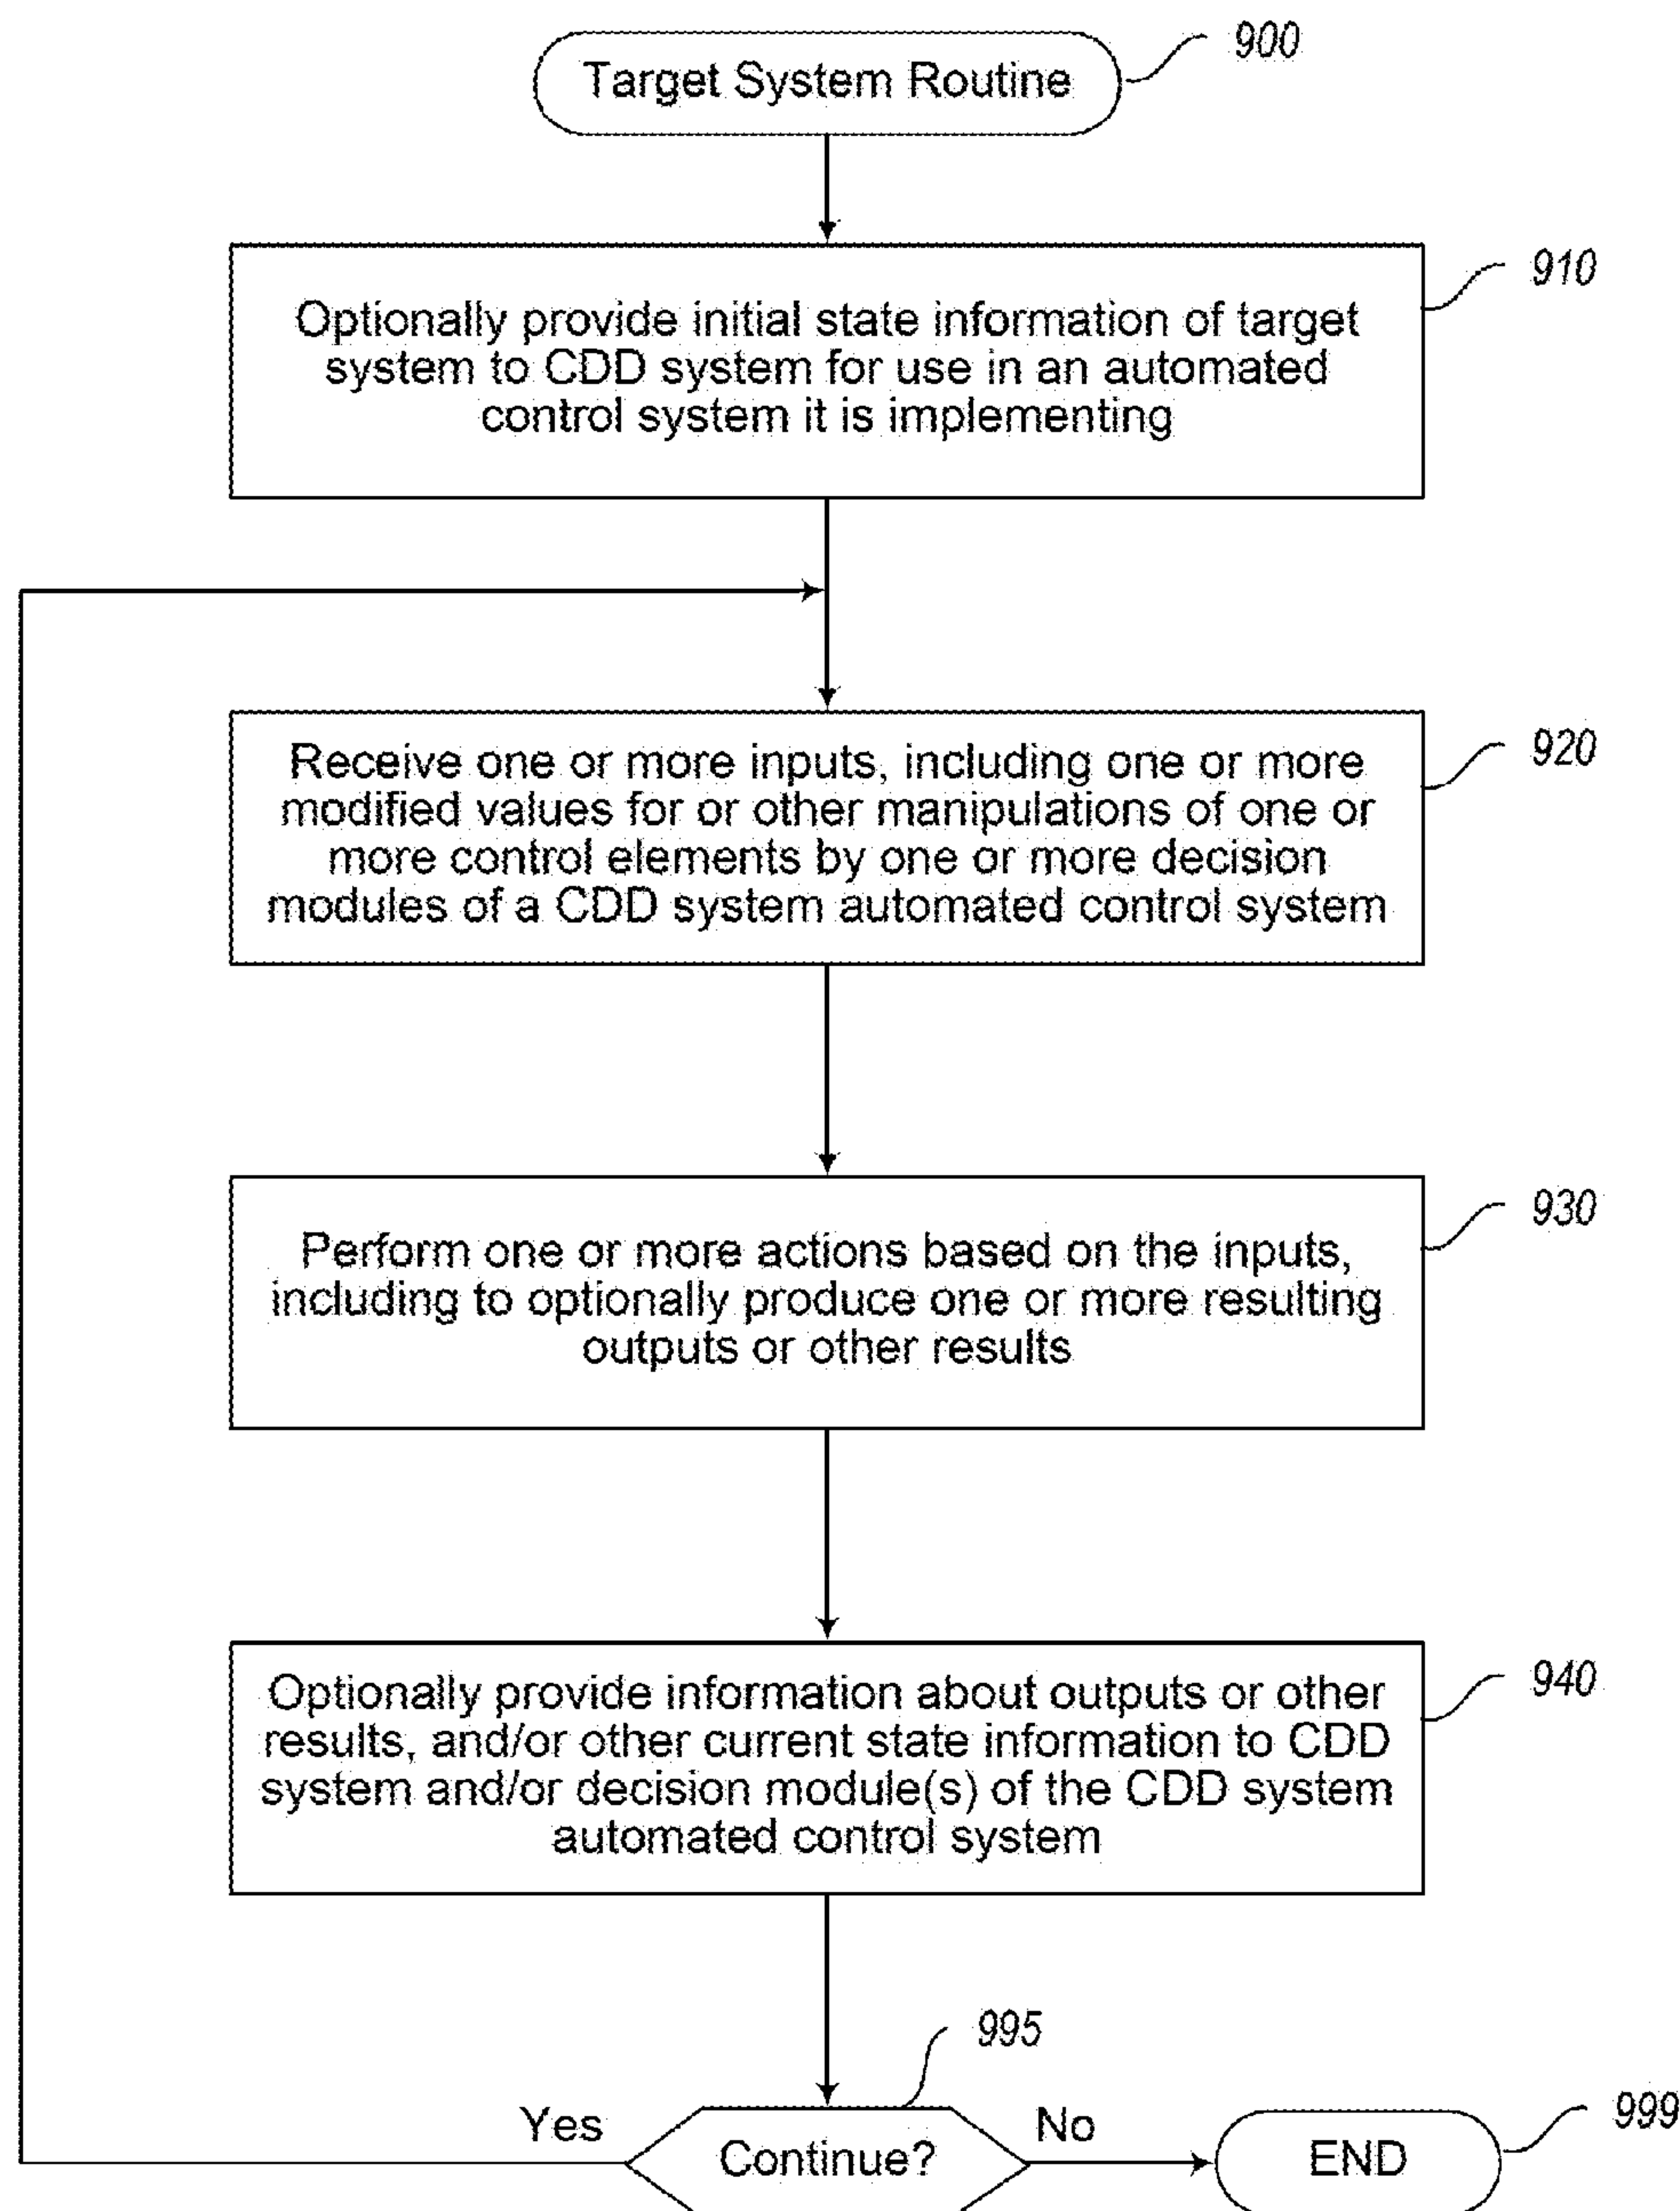


FIG. 8B



FIG. 9



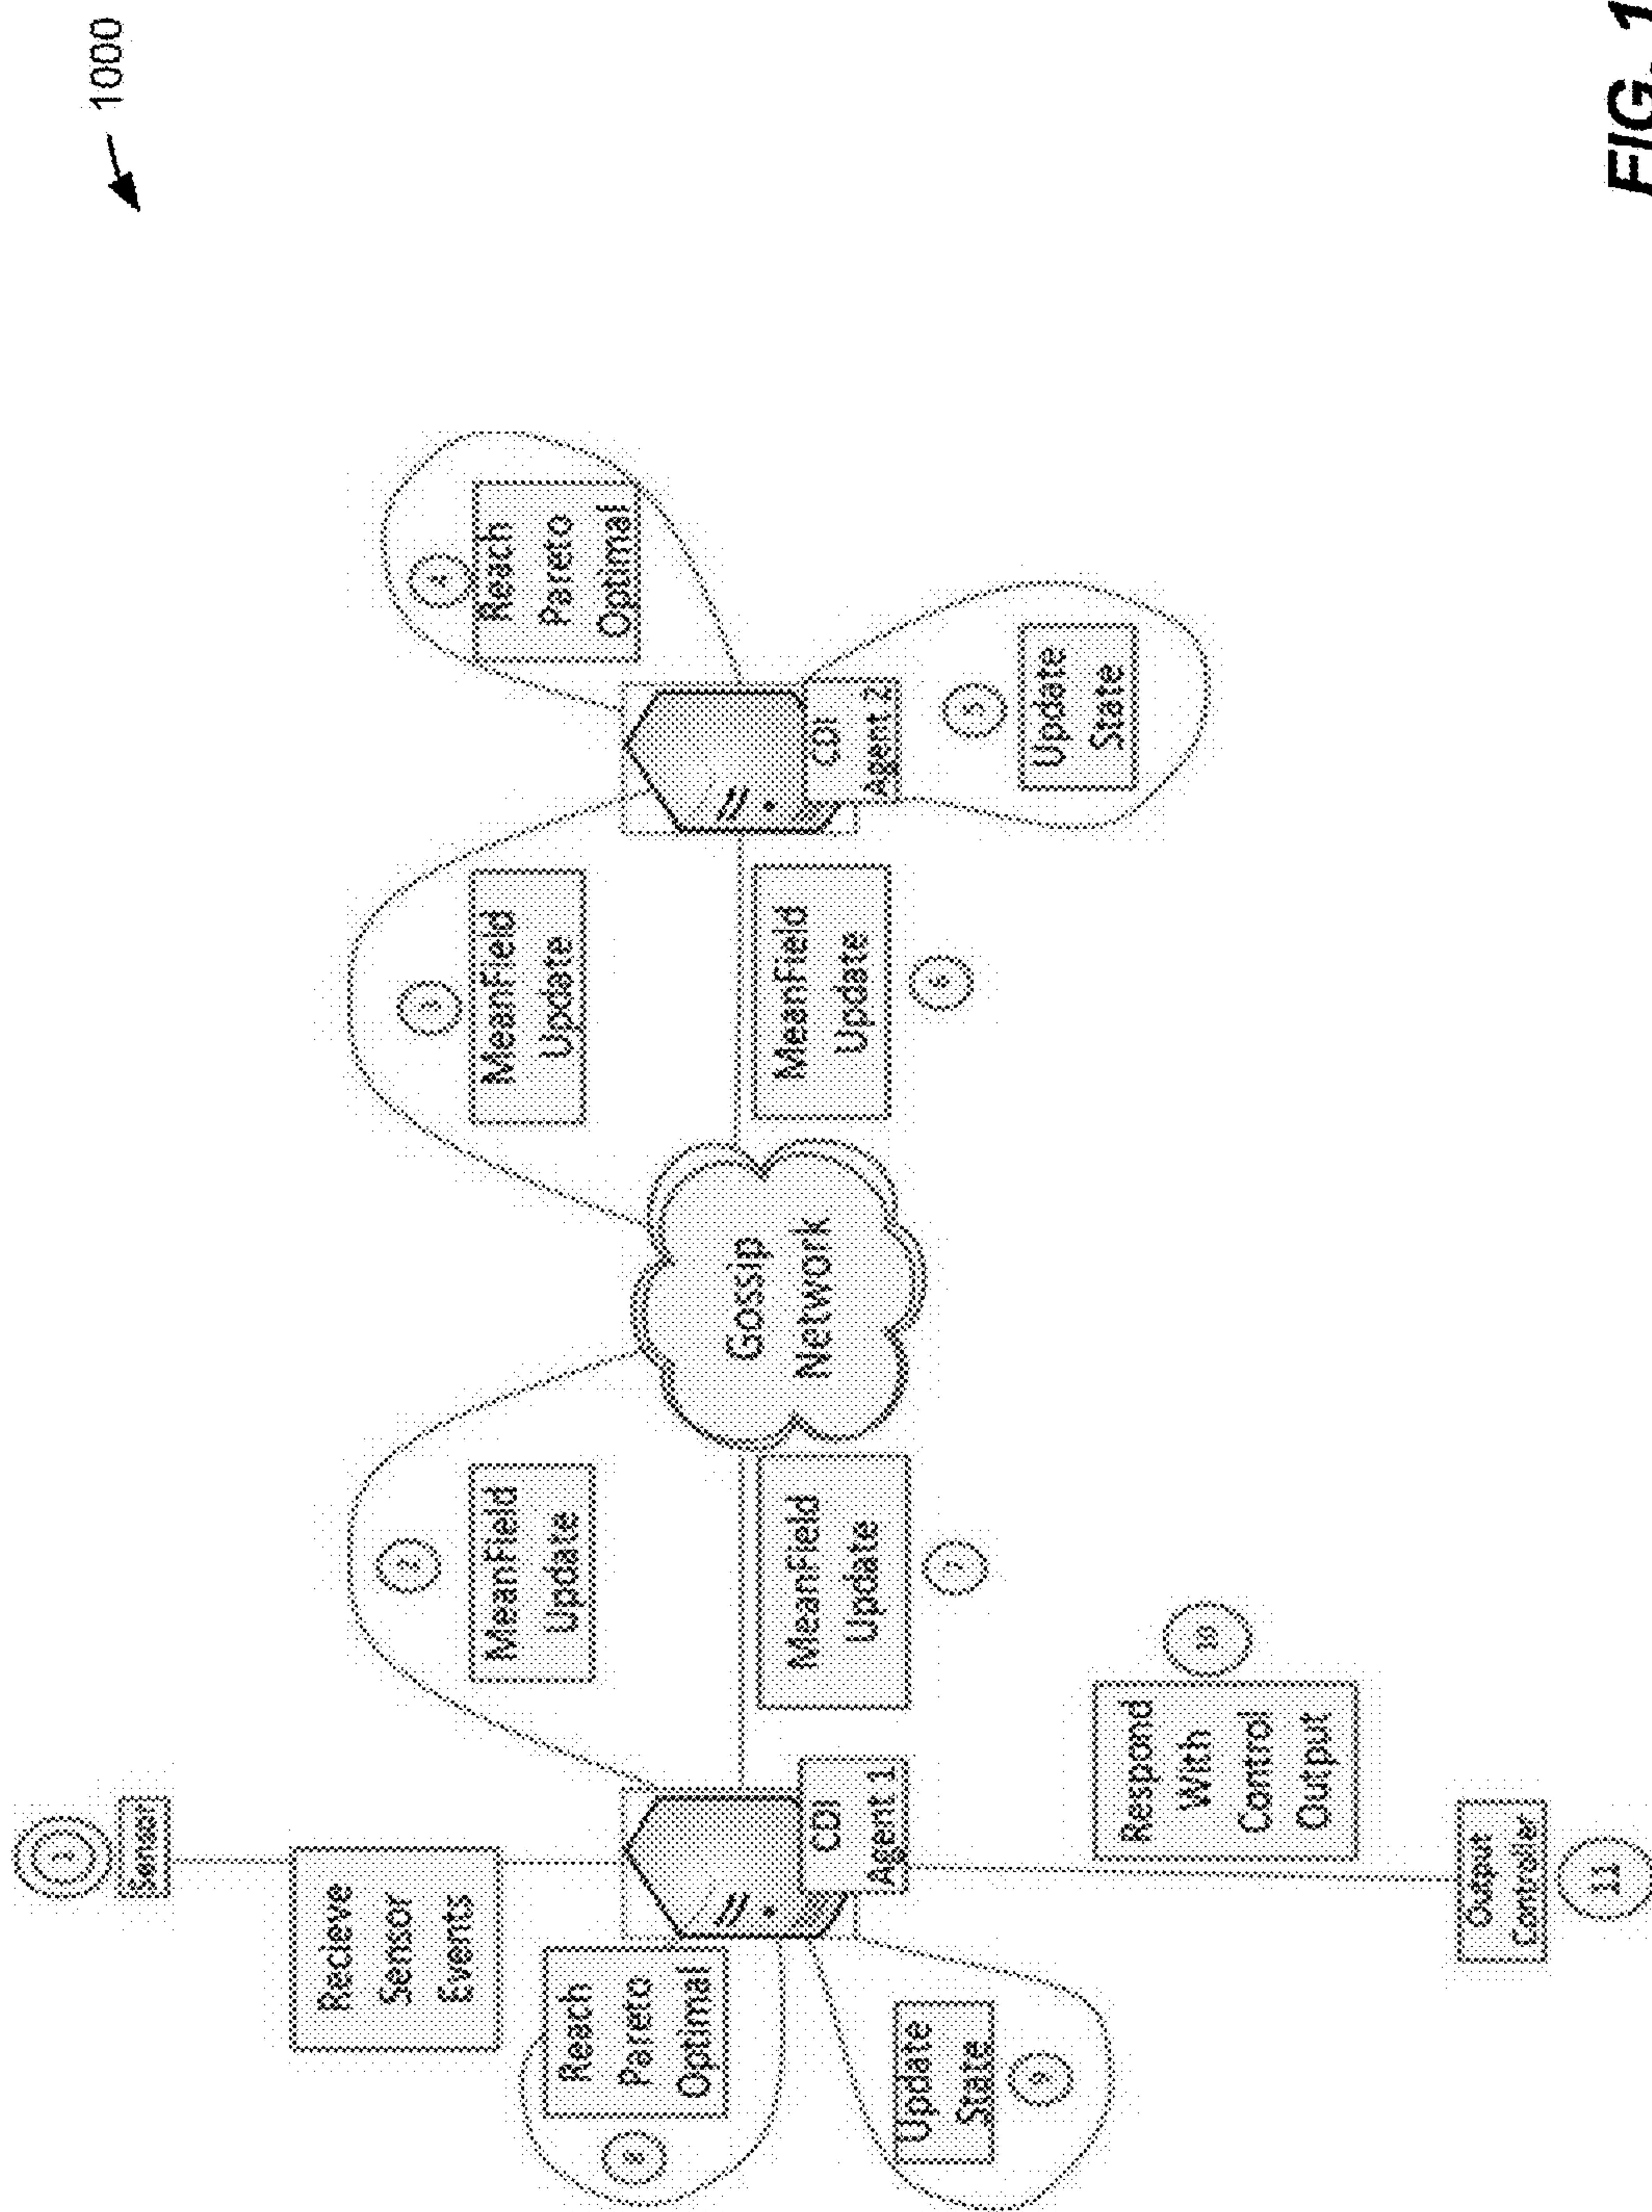
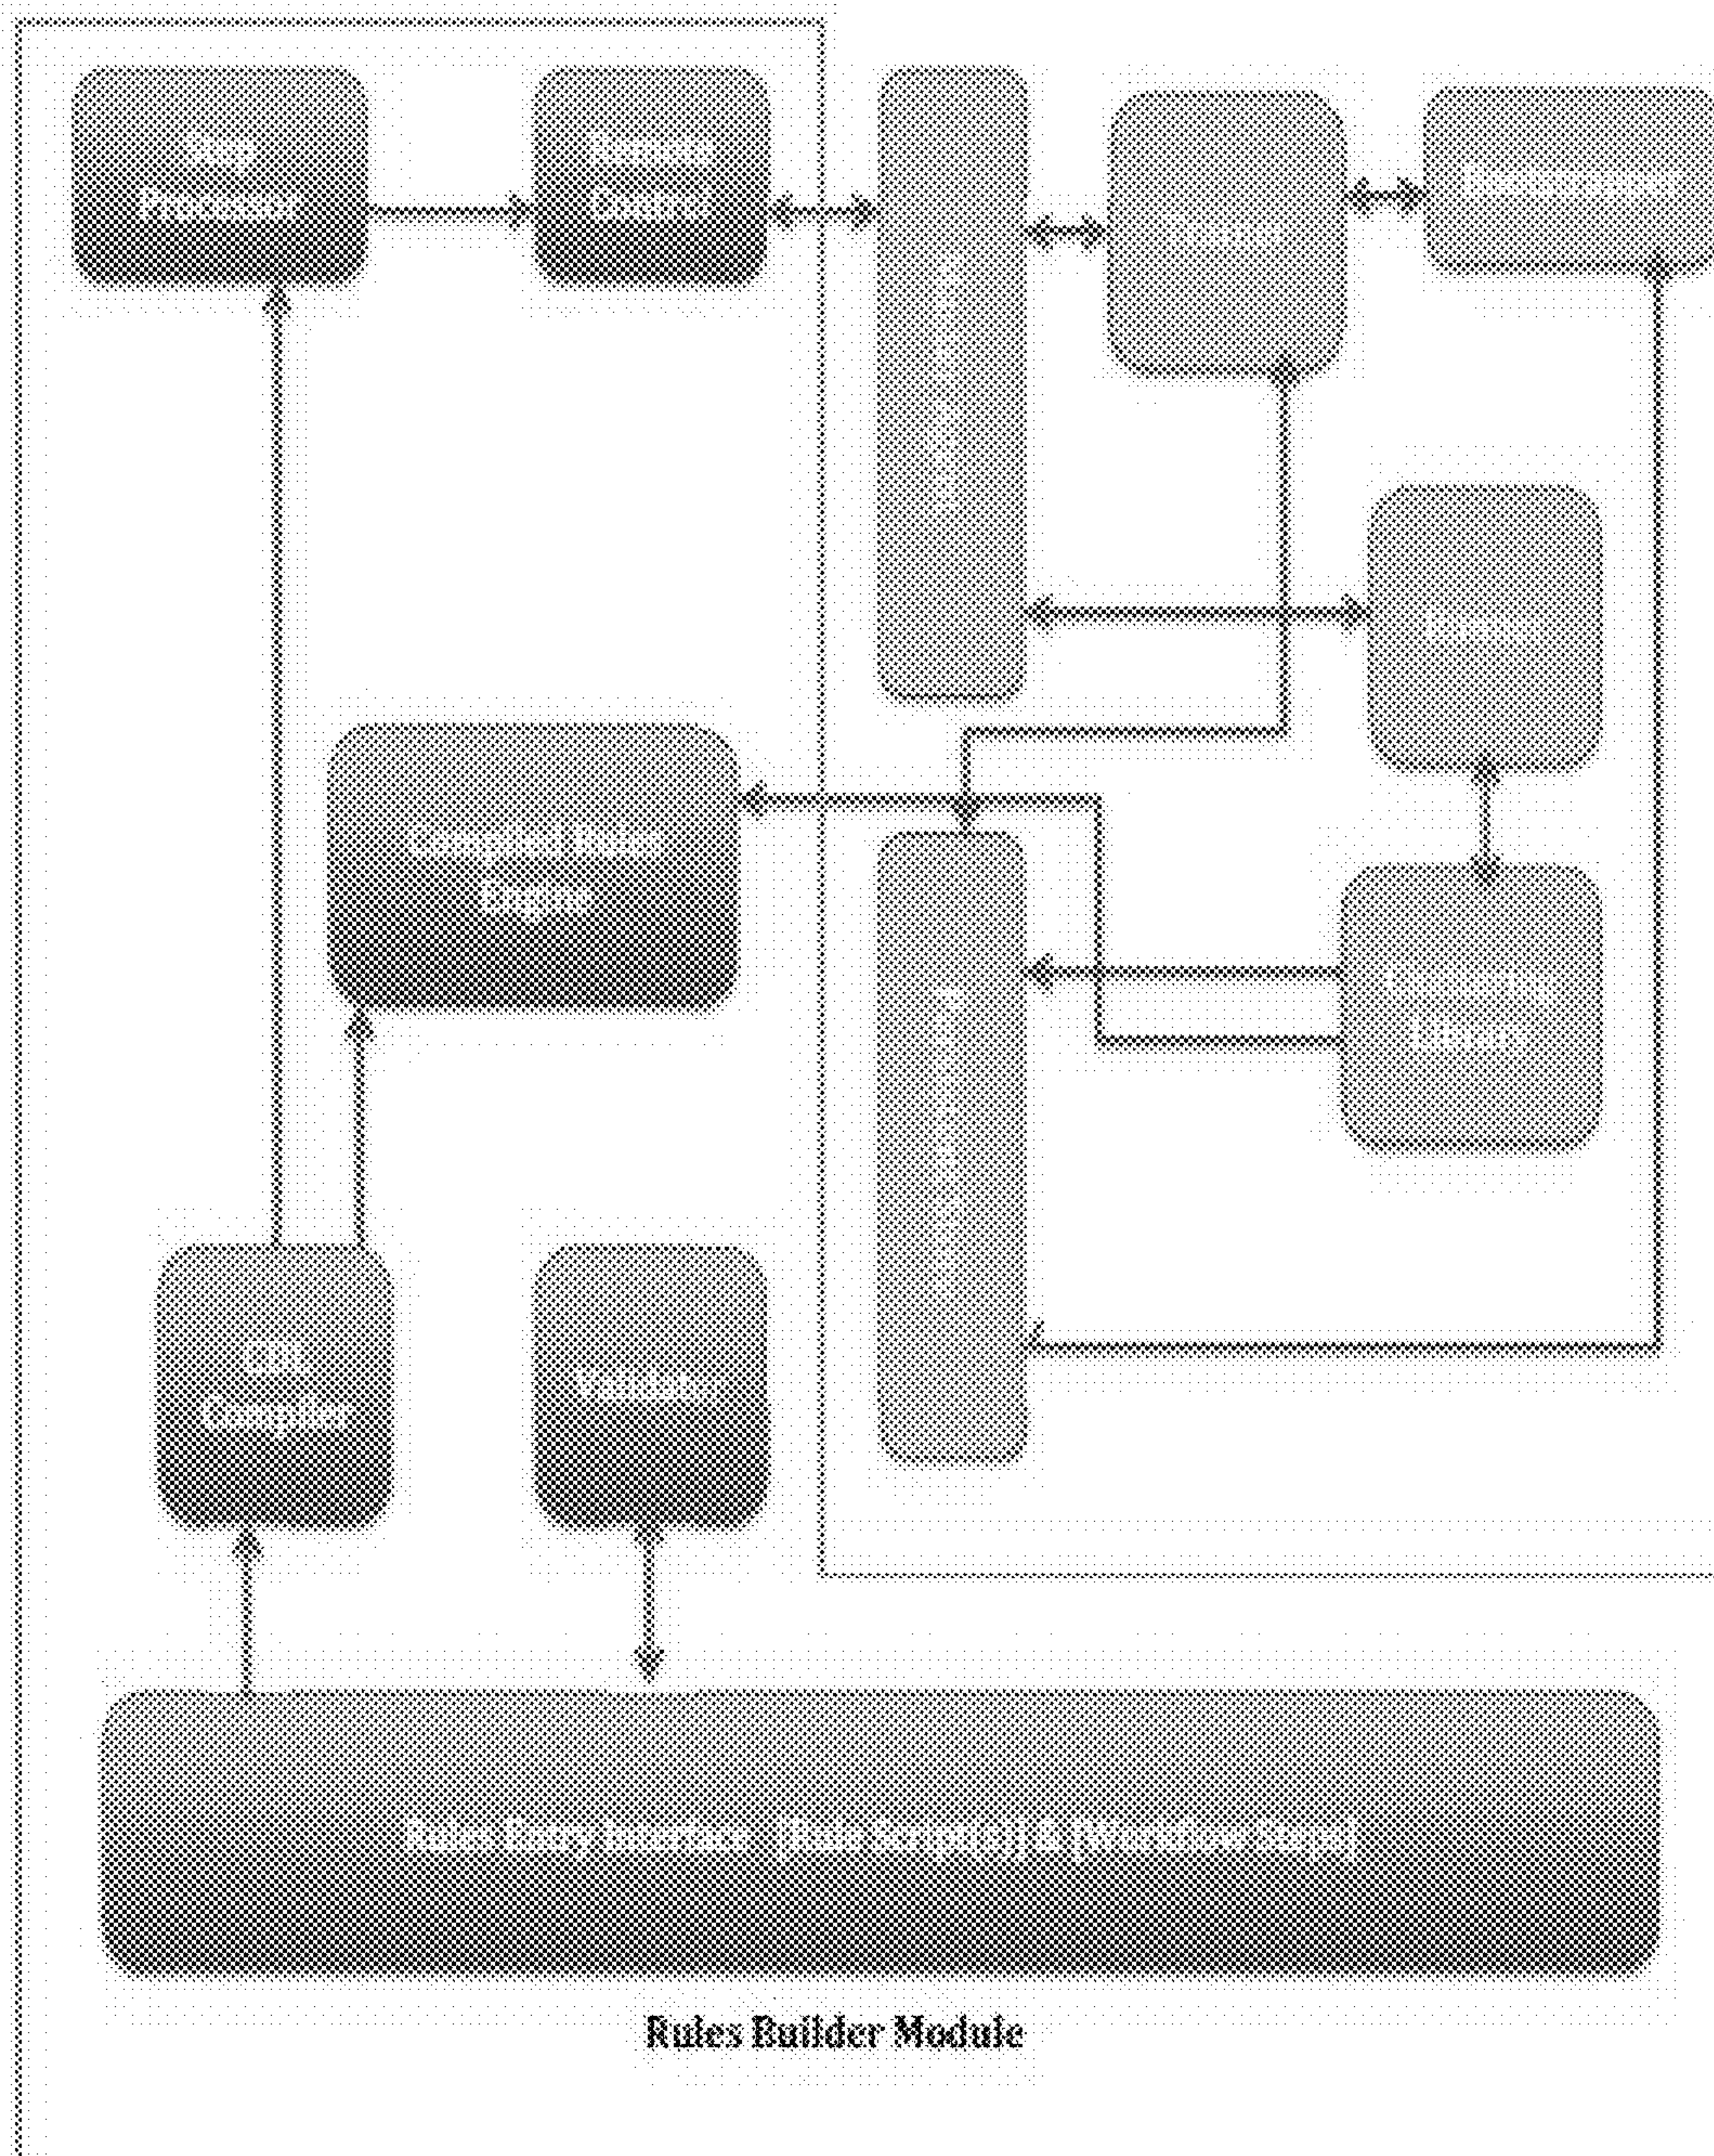


FIG. 10



FIG. 11

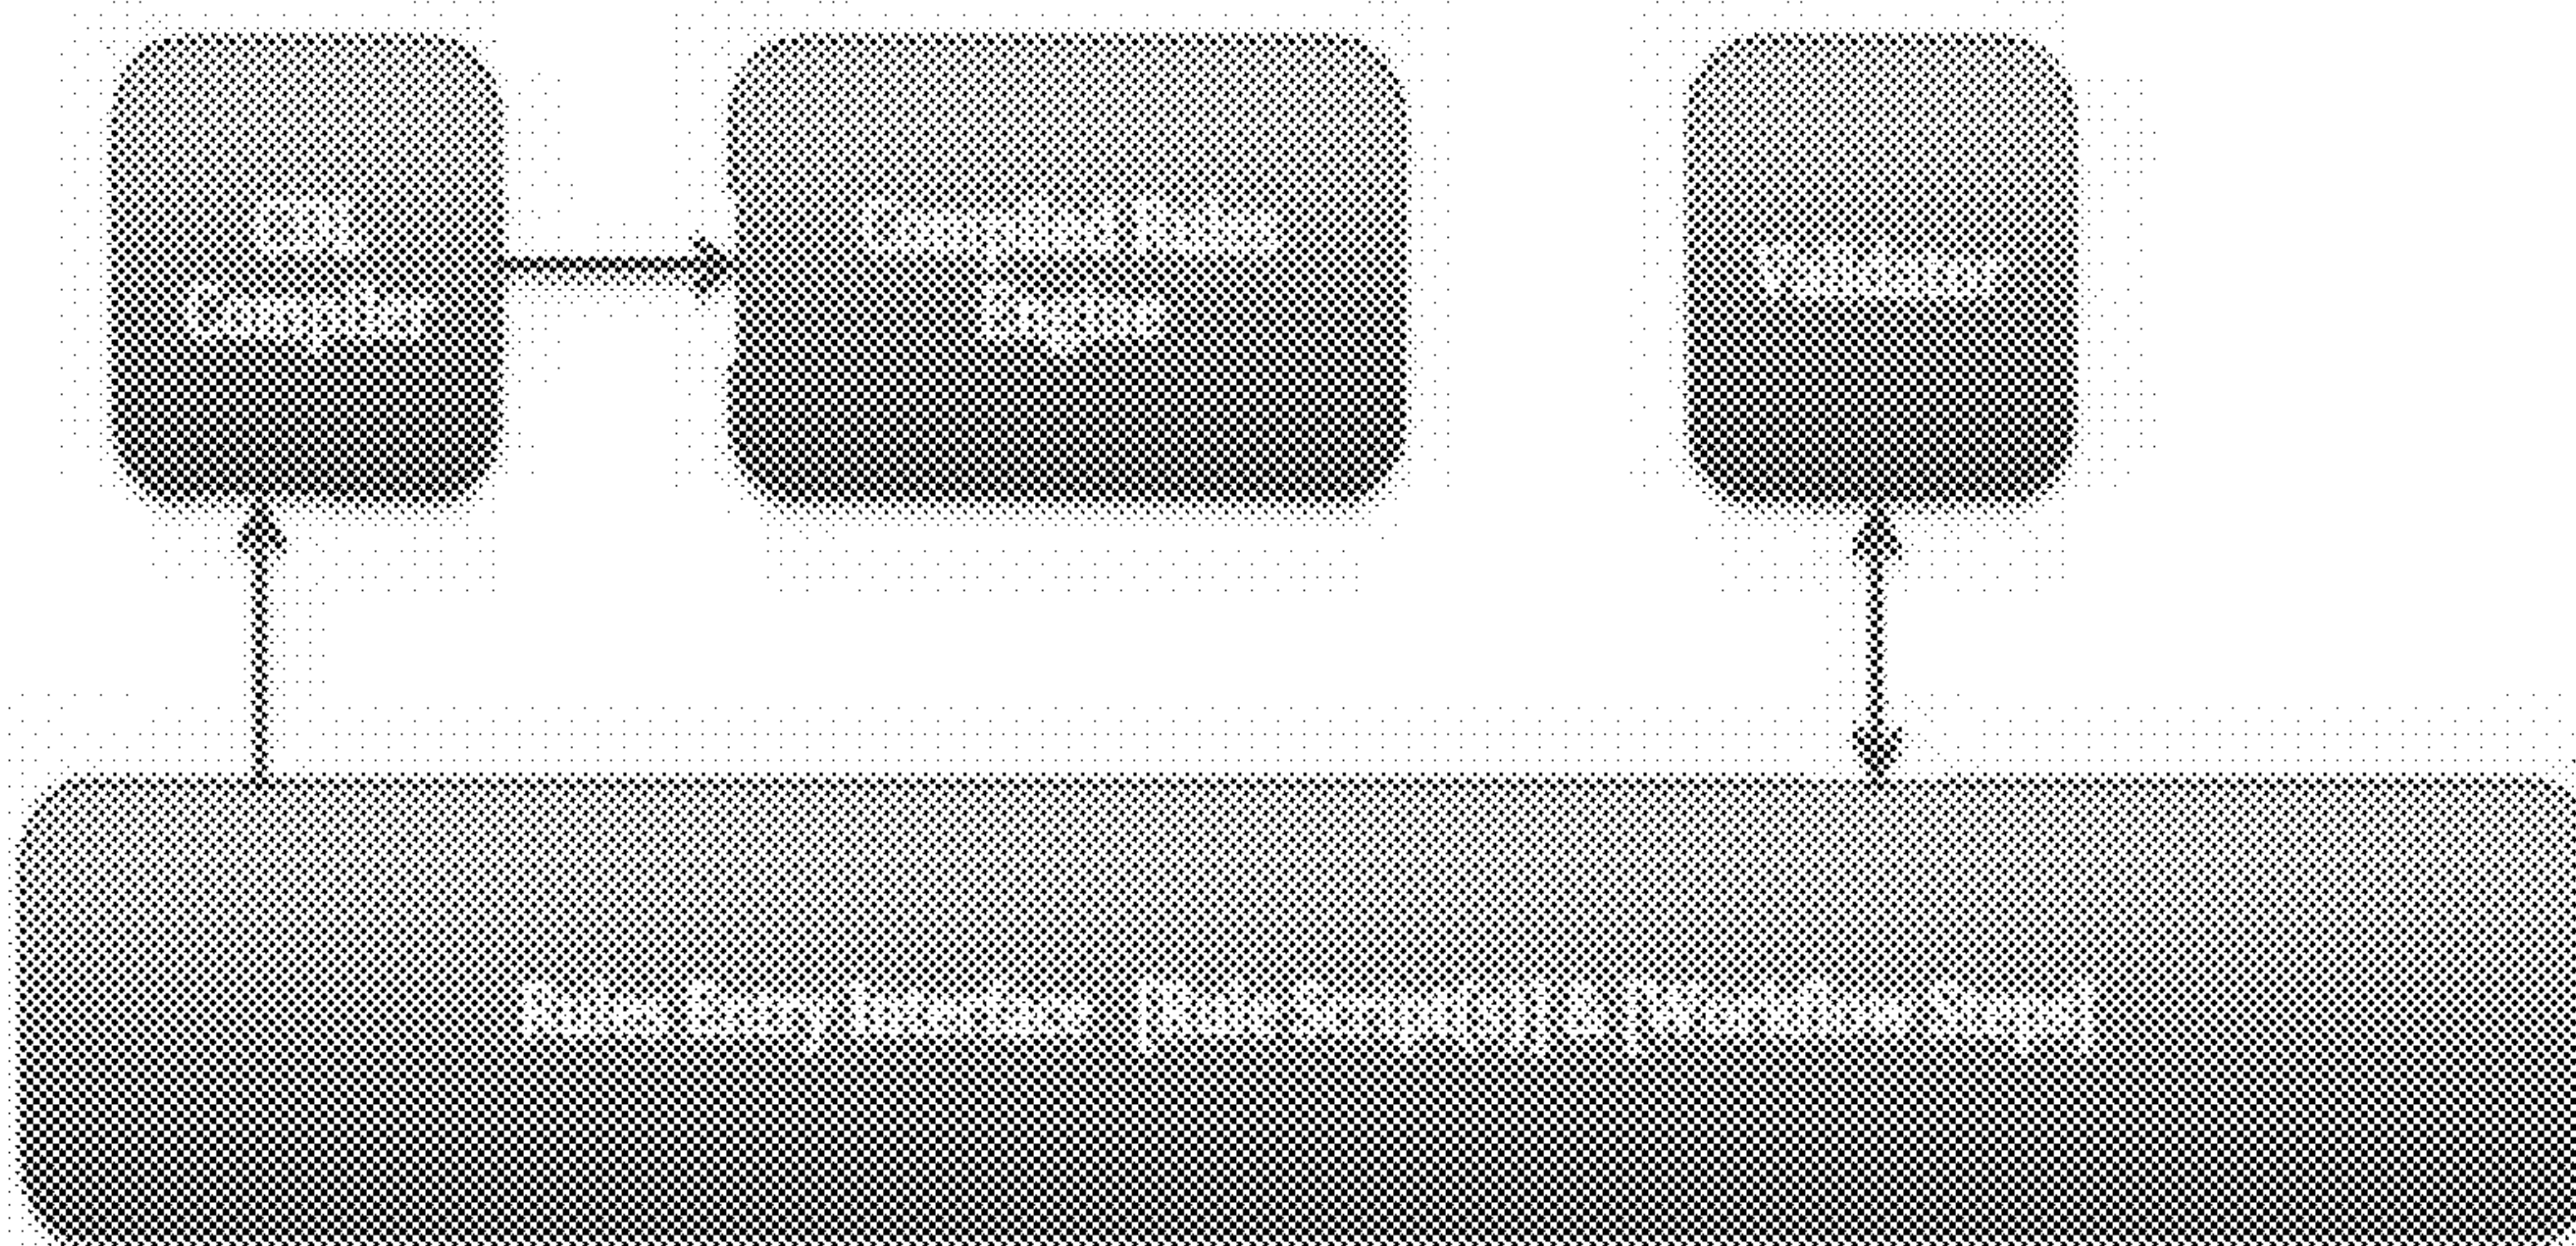
1100





**FIG. 12**

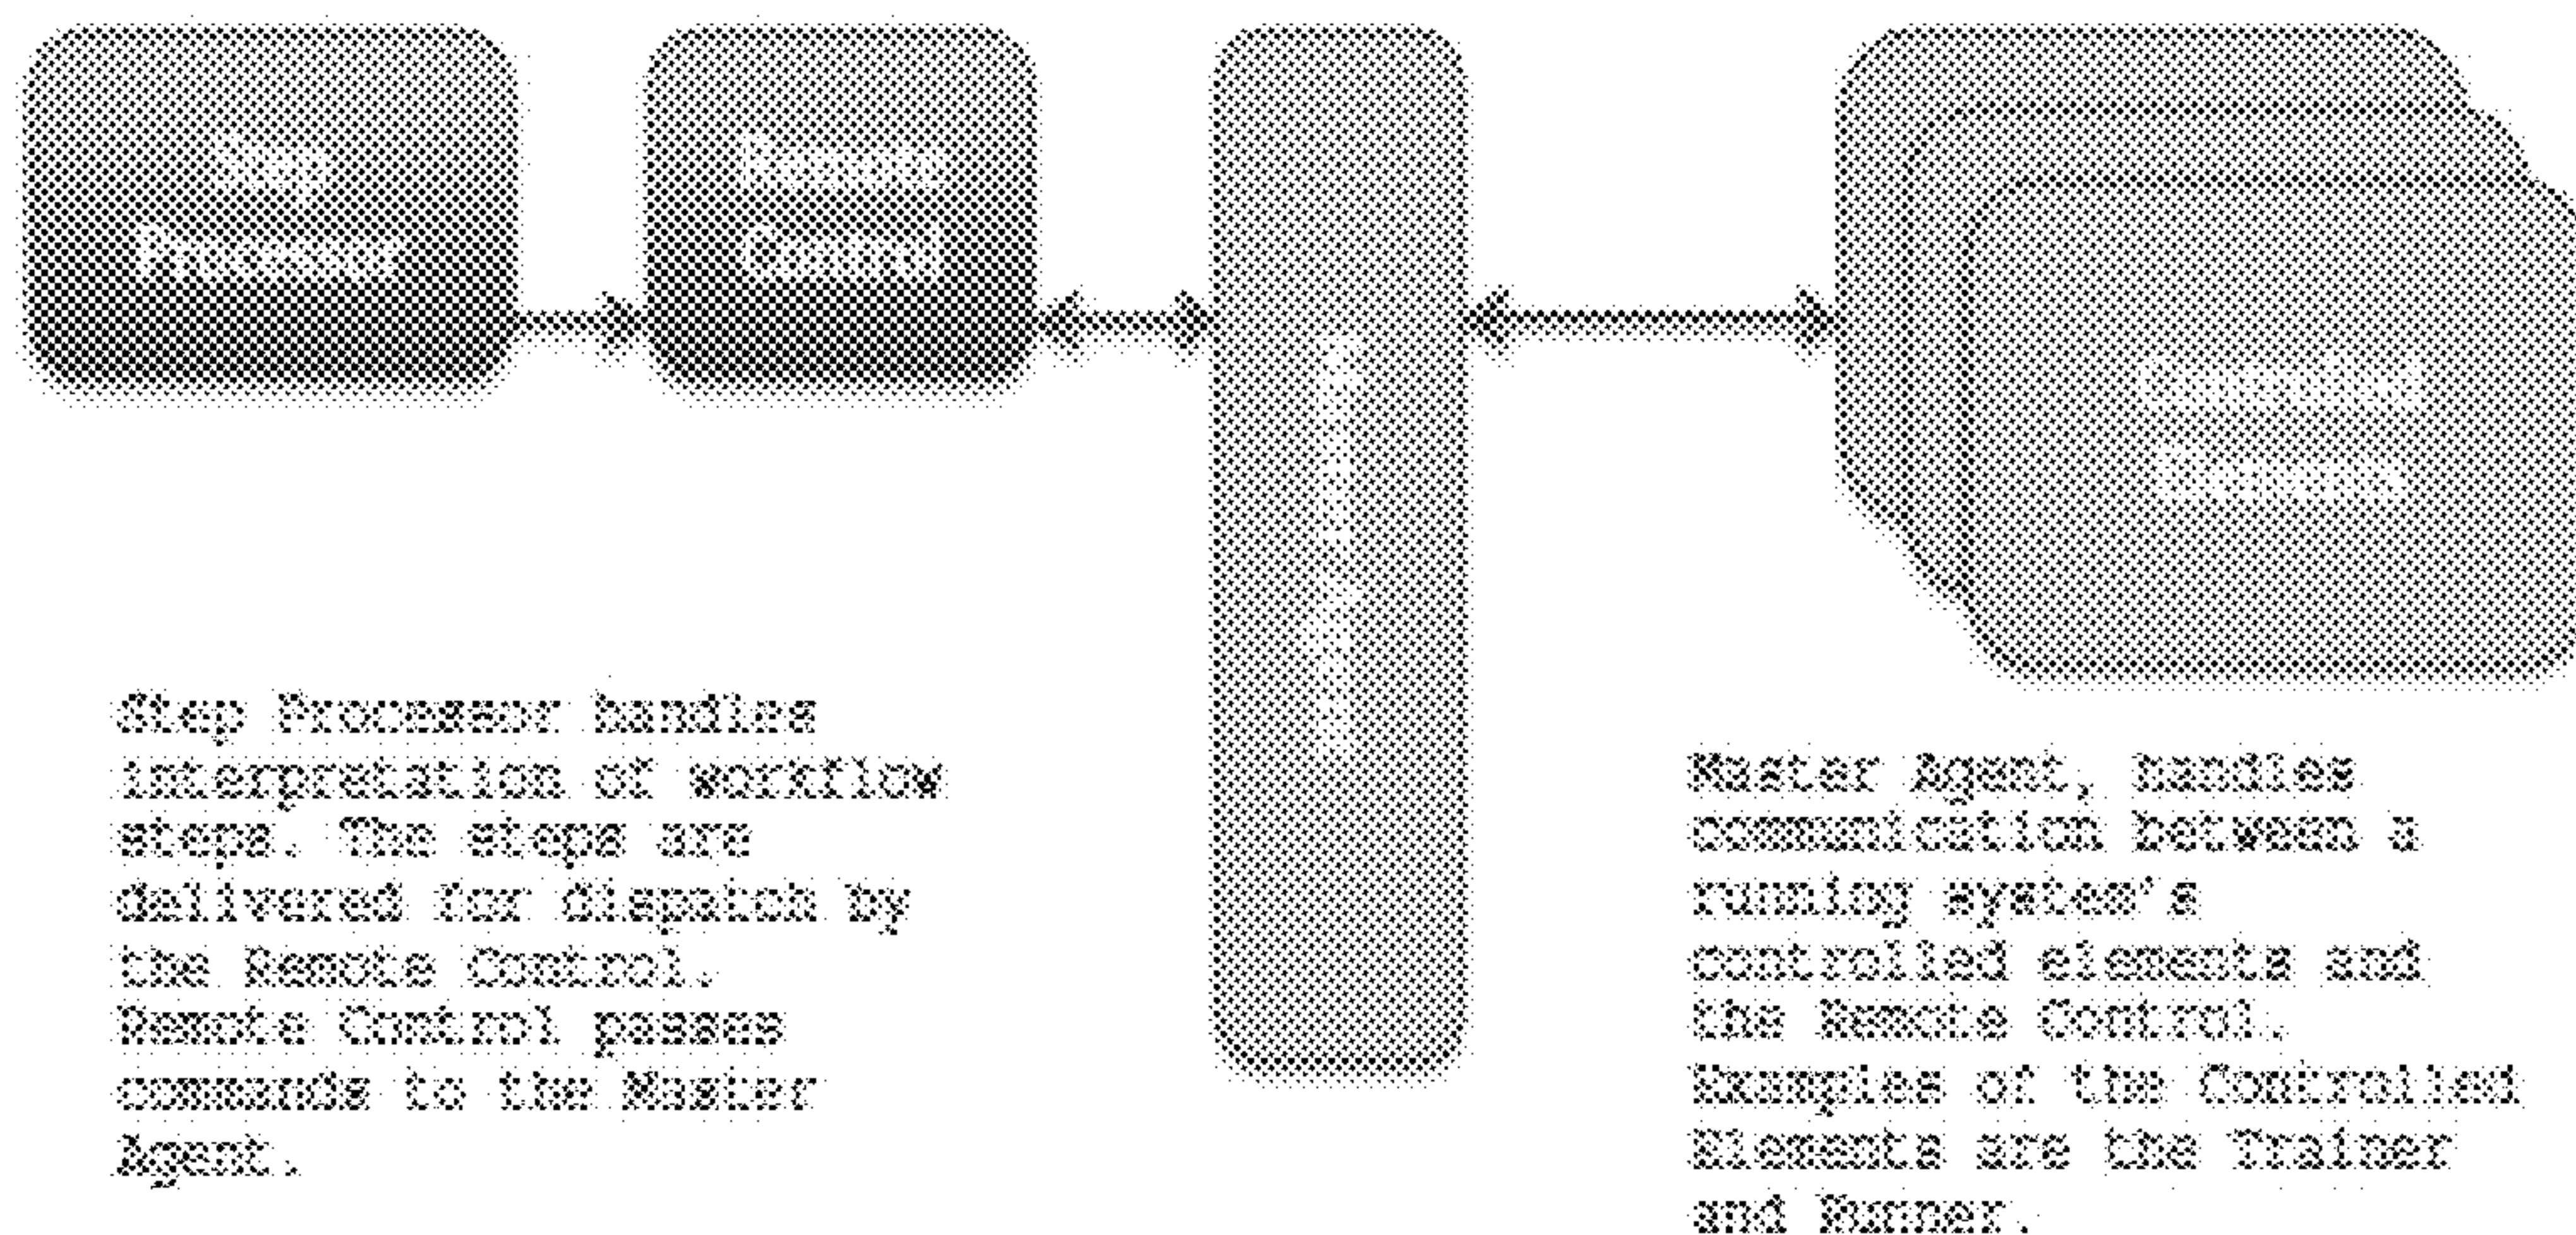
1200



**FIG. 13**

1300

Diagram C.) Workflow Step Processing with Remote Control Messaging to Master Agent



**FIG. 14**

← 1400

Diagram B: Bootstrapping, Training, Running and Persisting state information

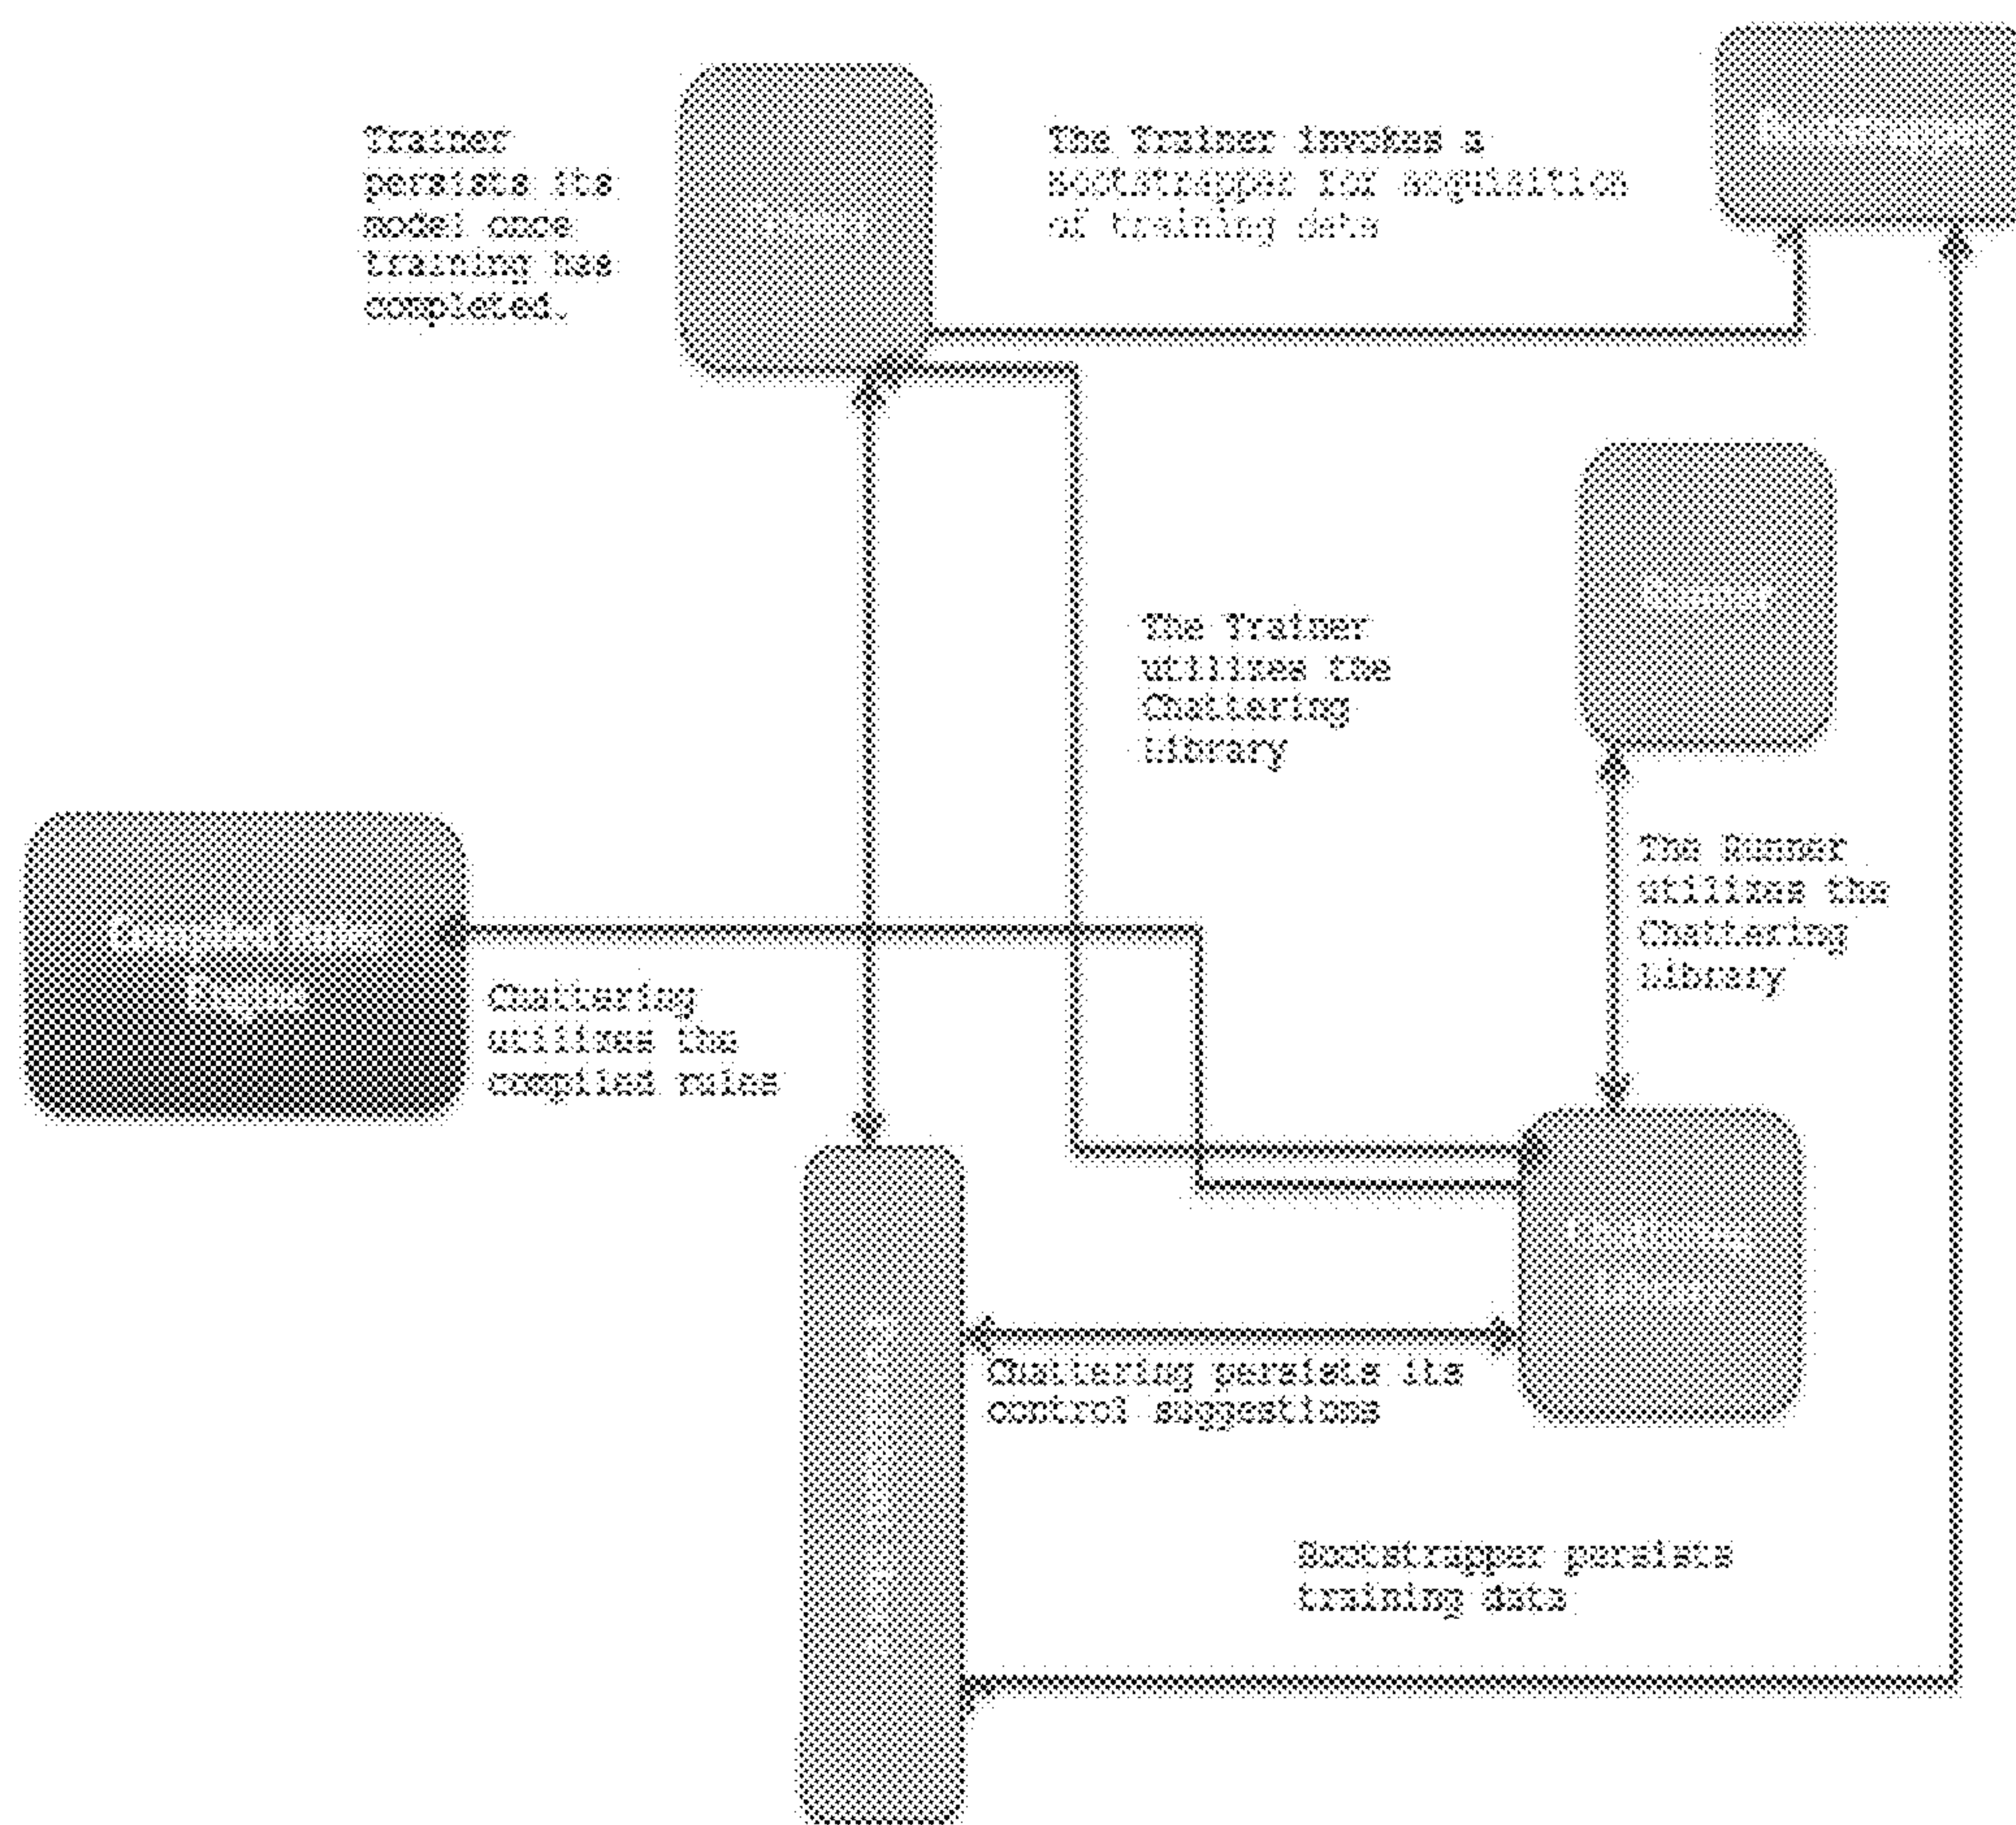




FIG. 15

1500

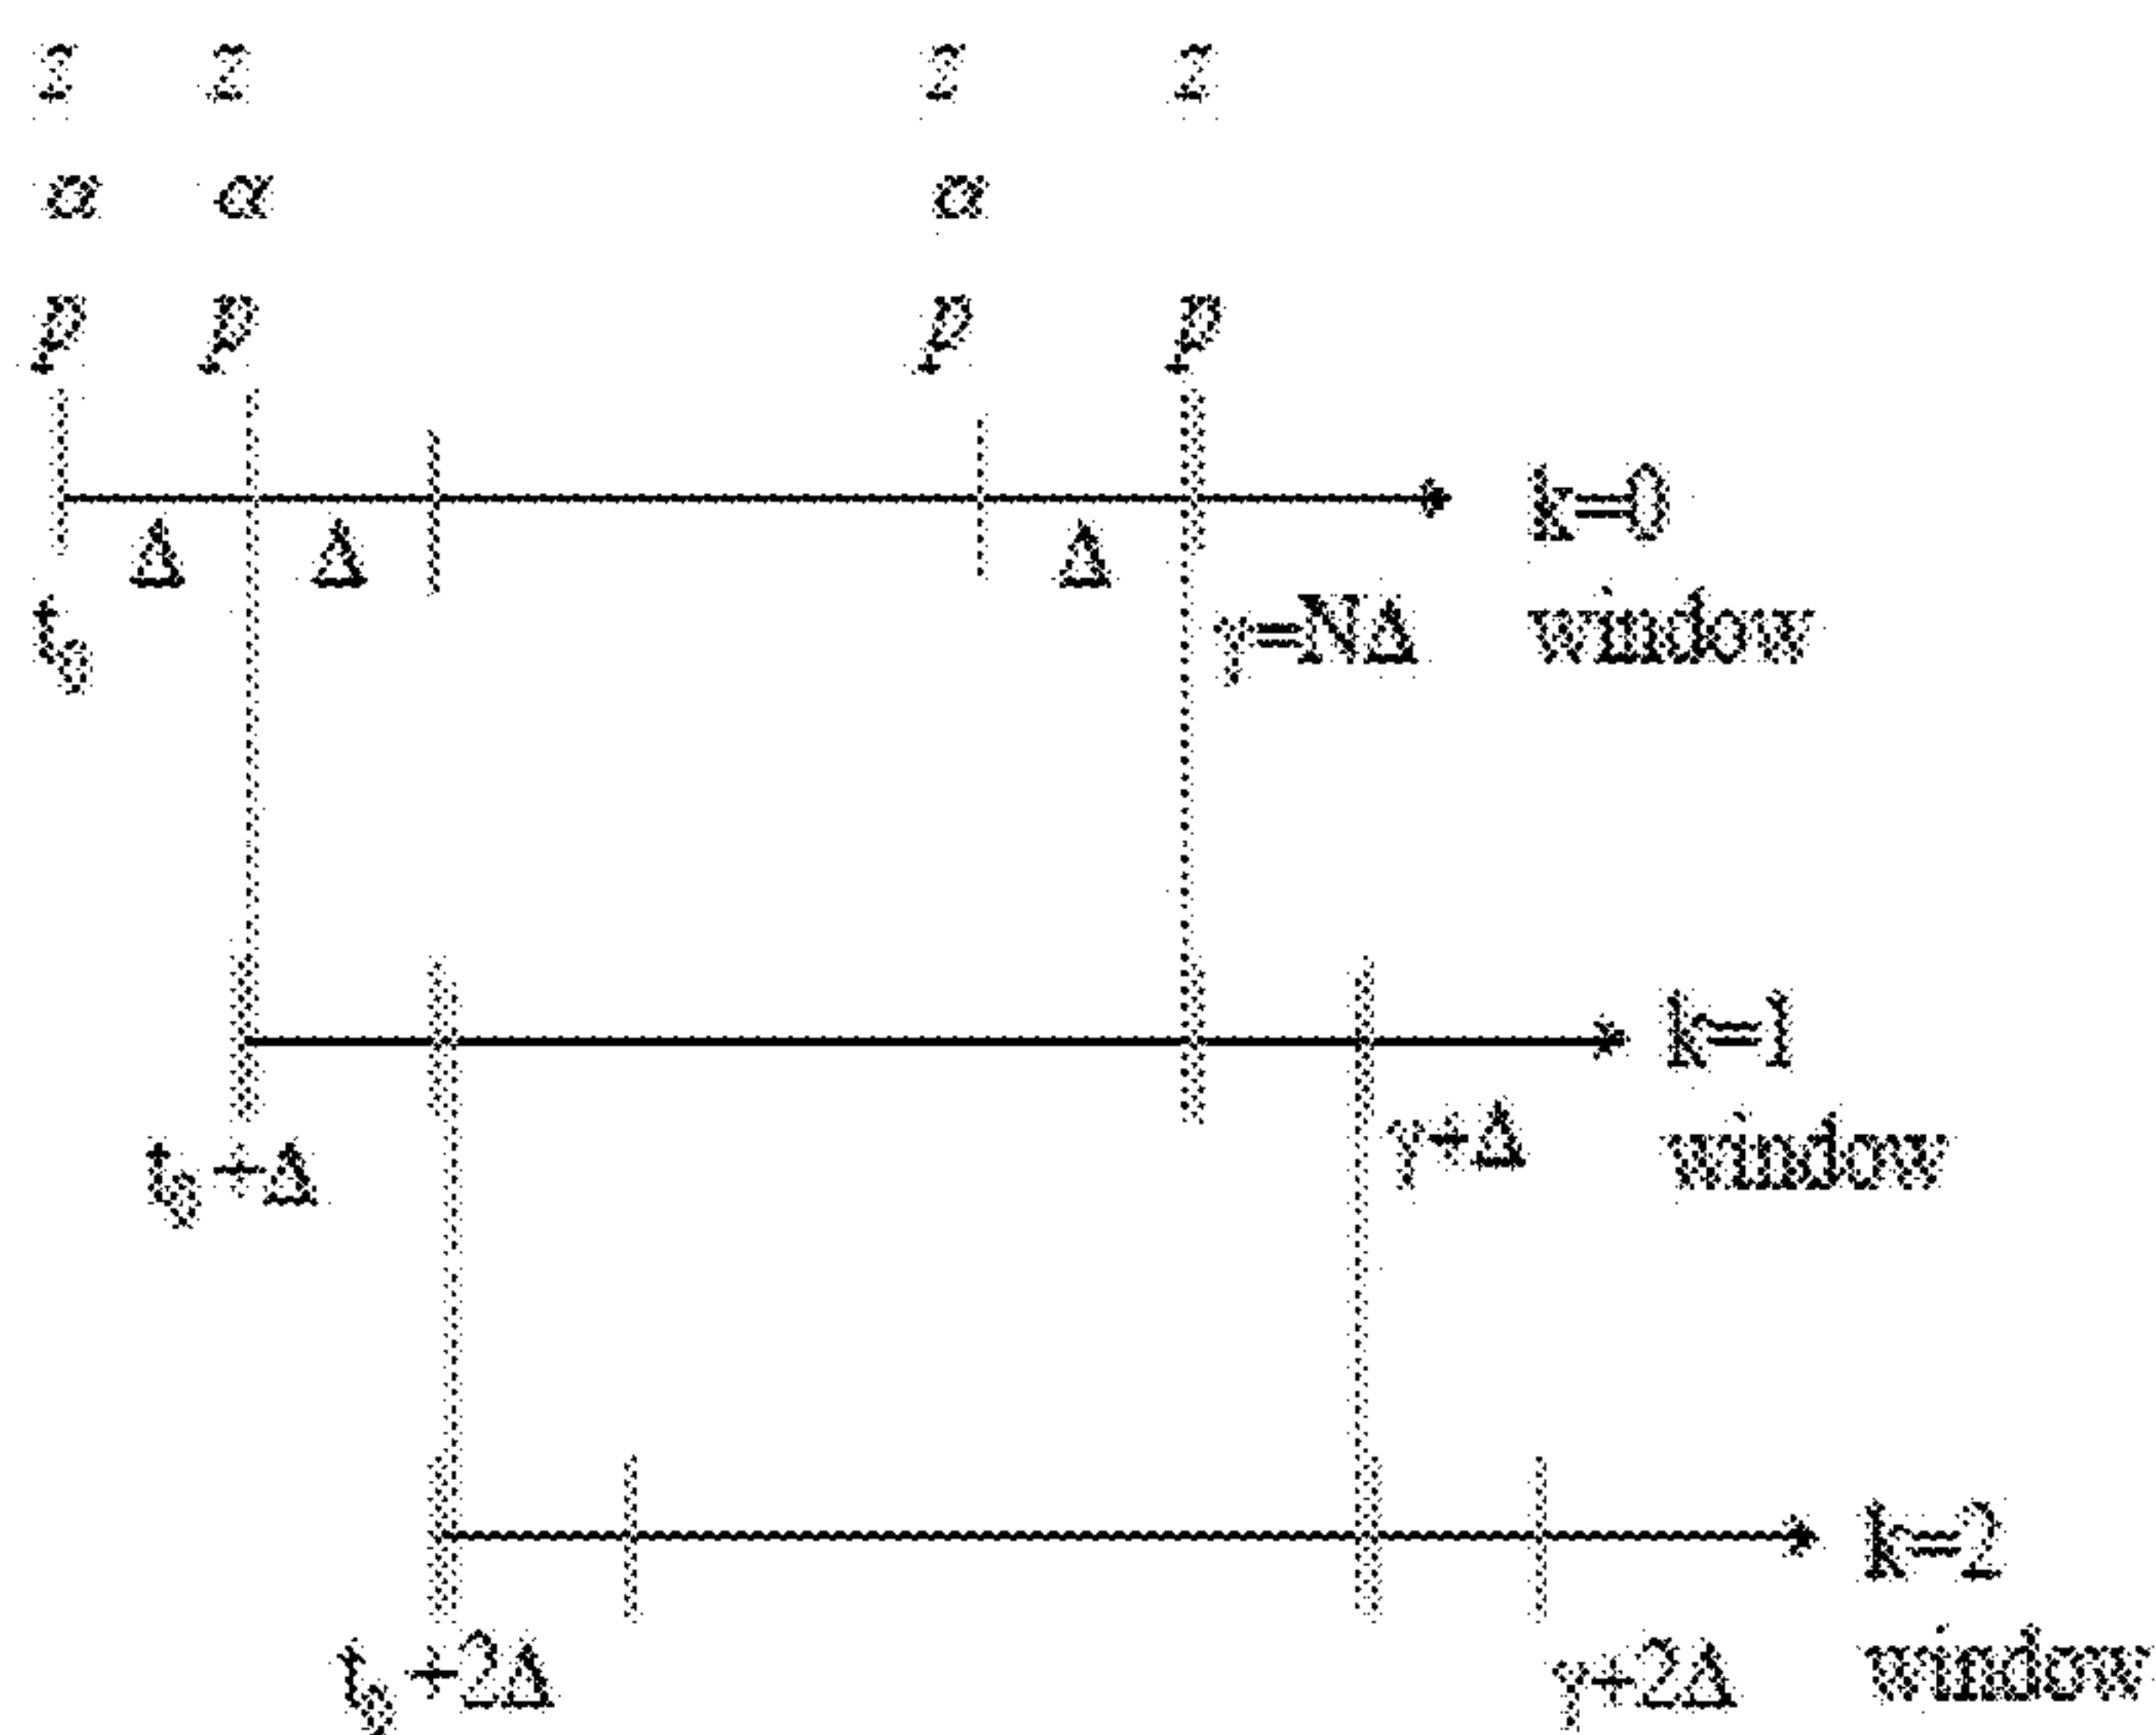


FIG. 16

1600






FIG. 17

1700

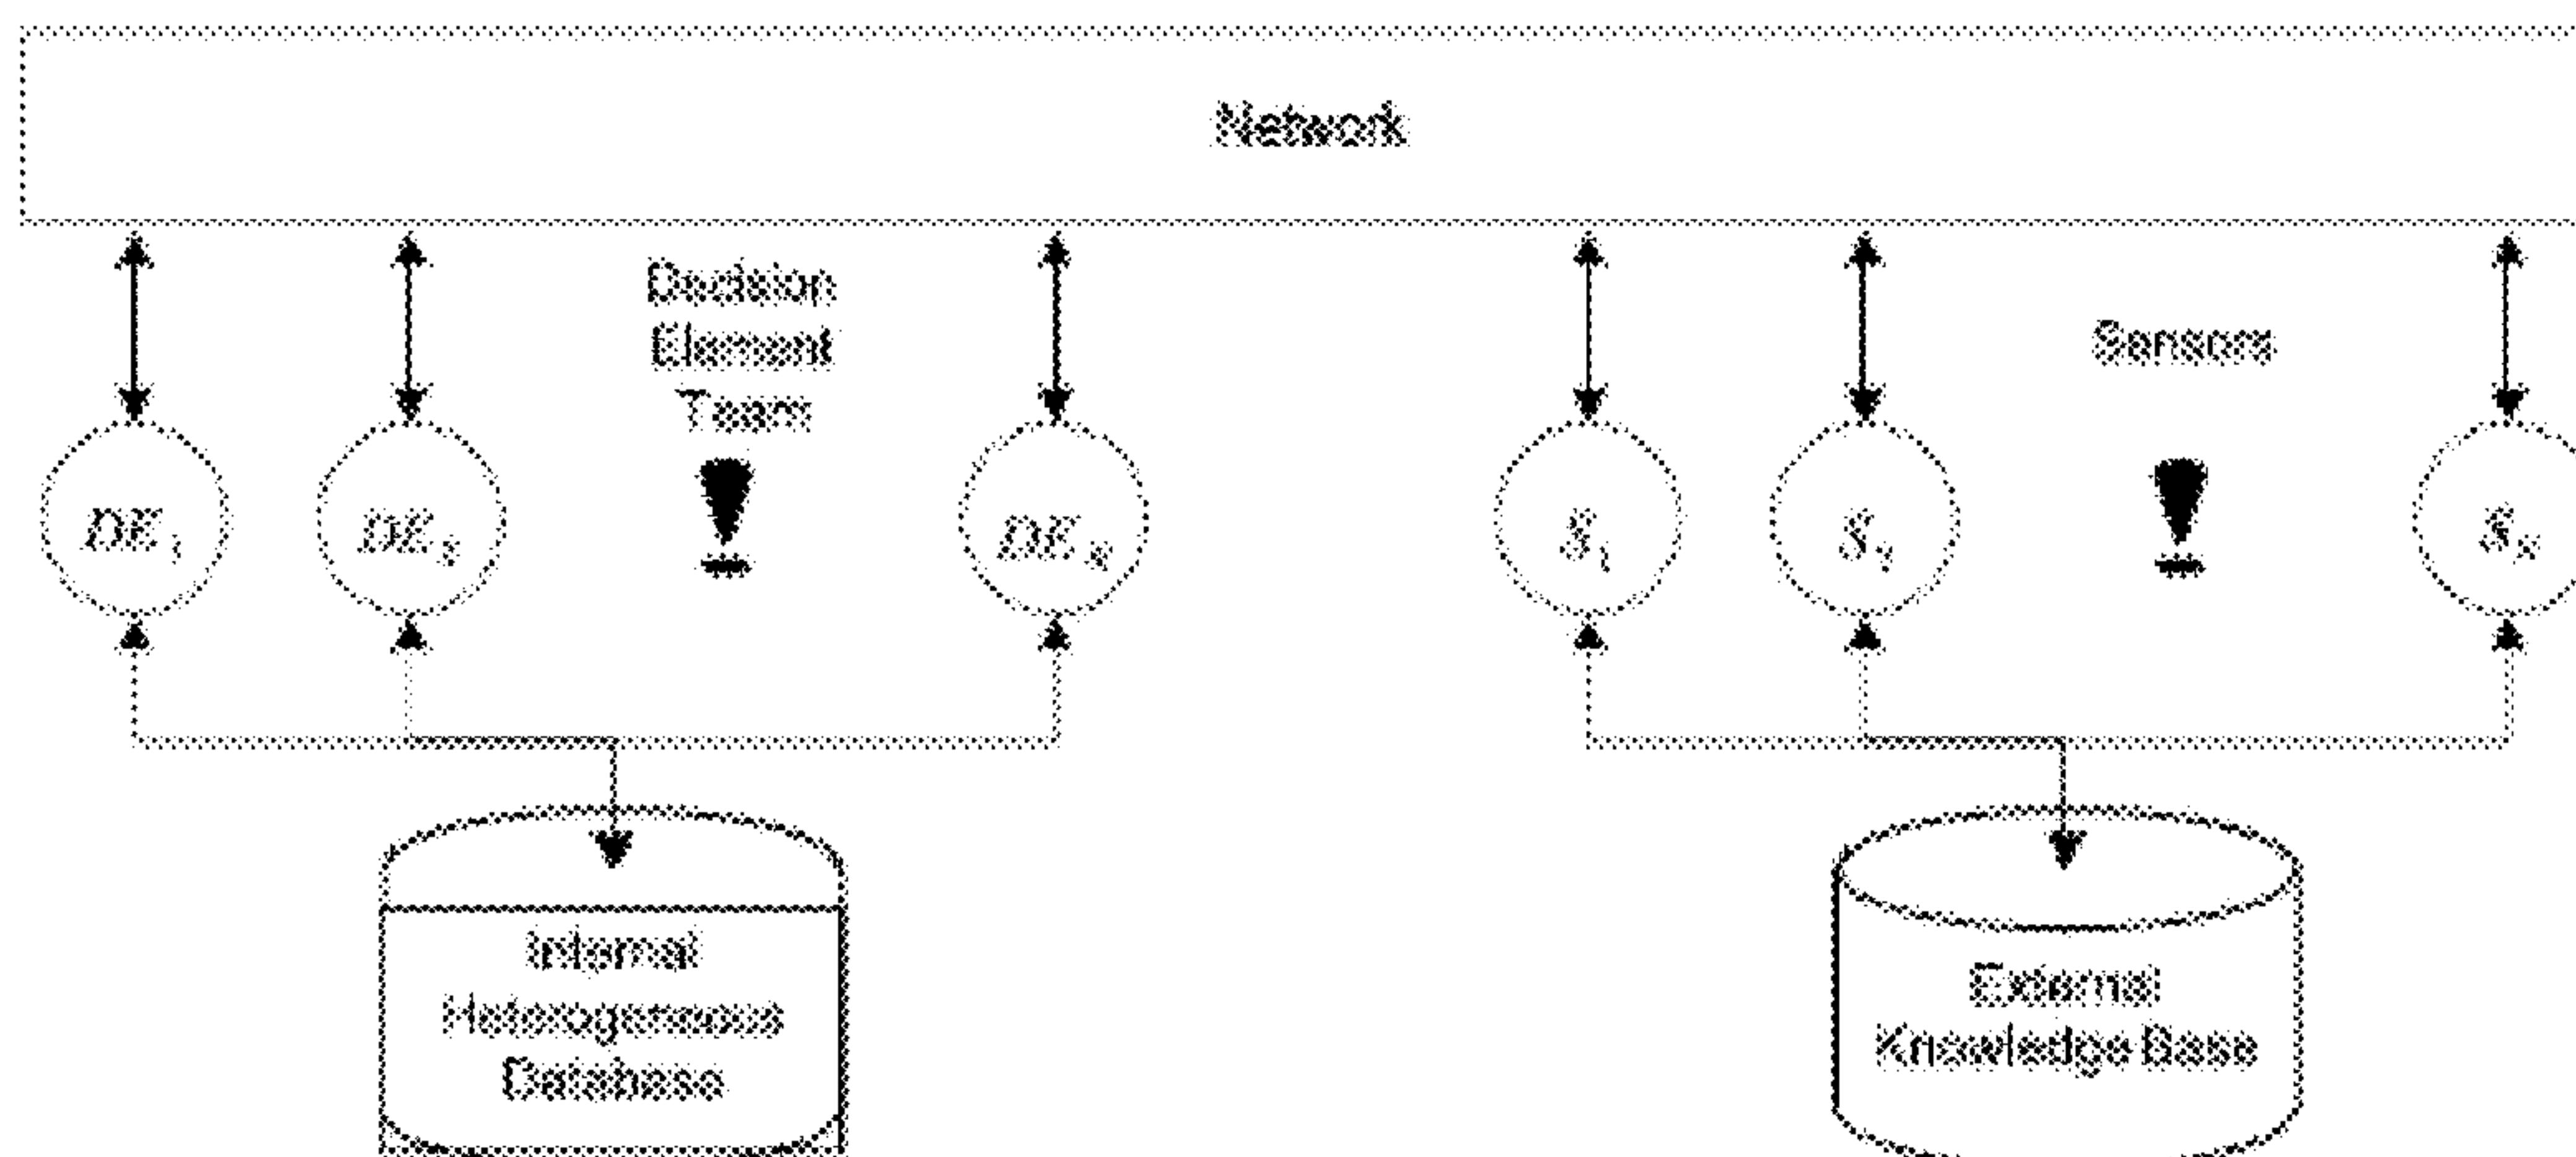
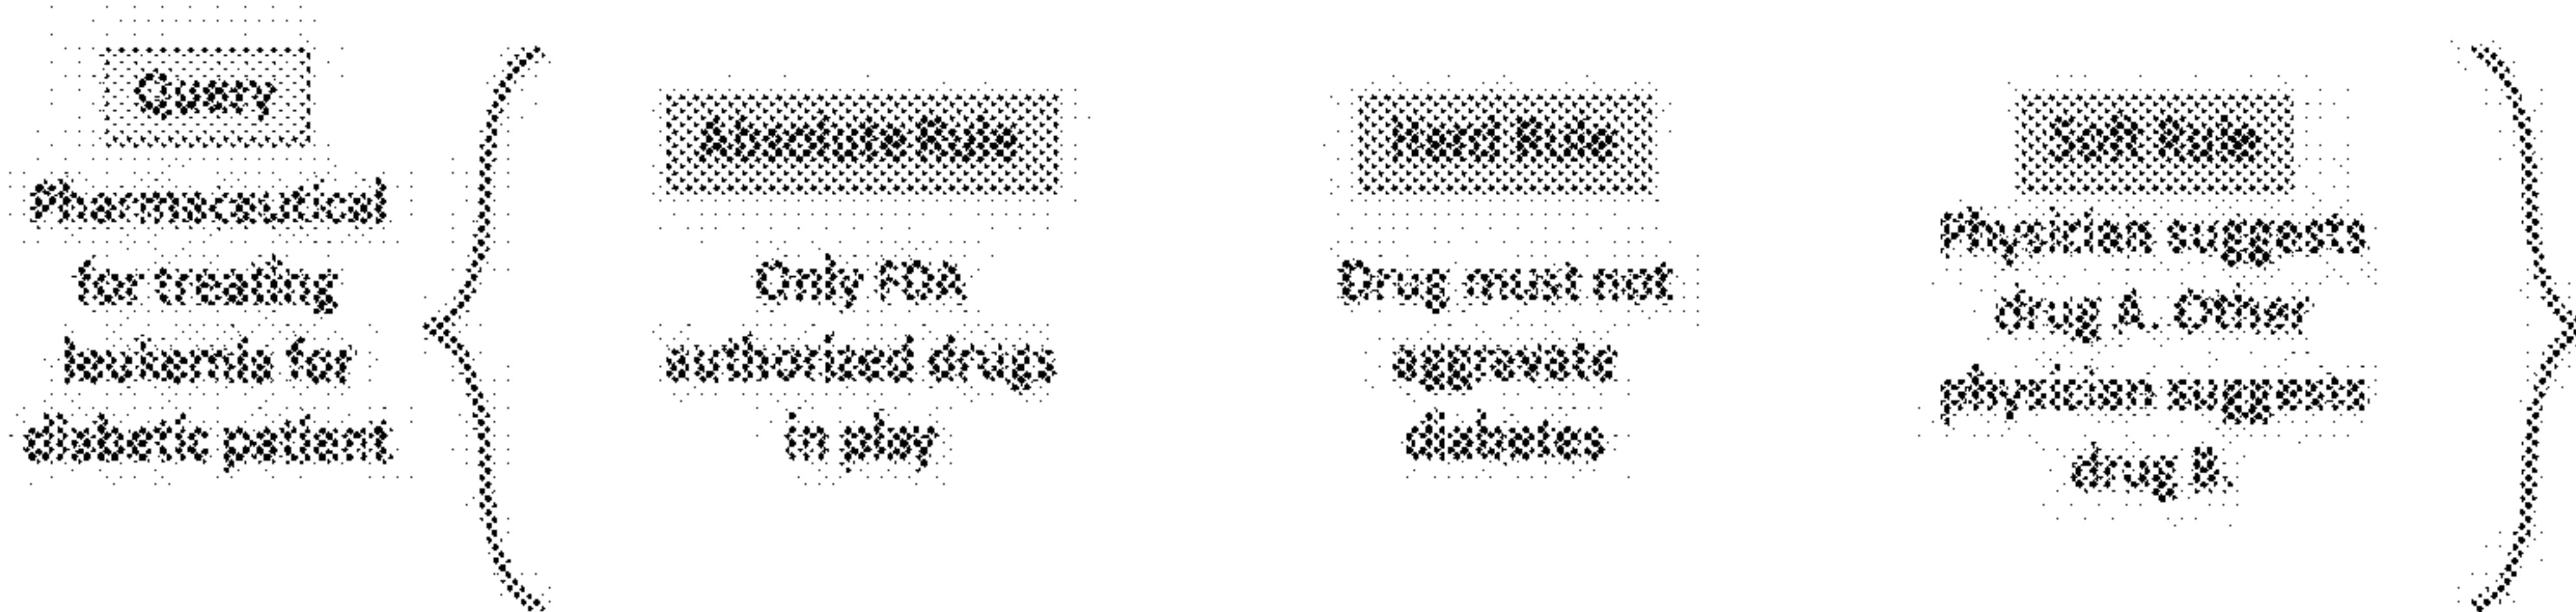


FIG. 18

1800



2000

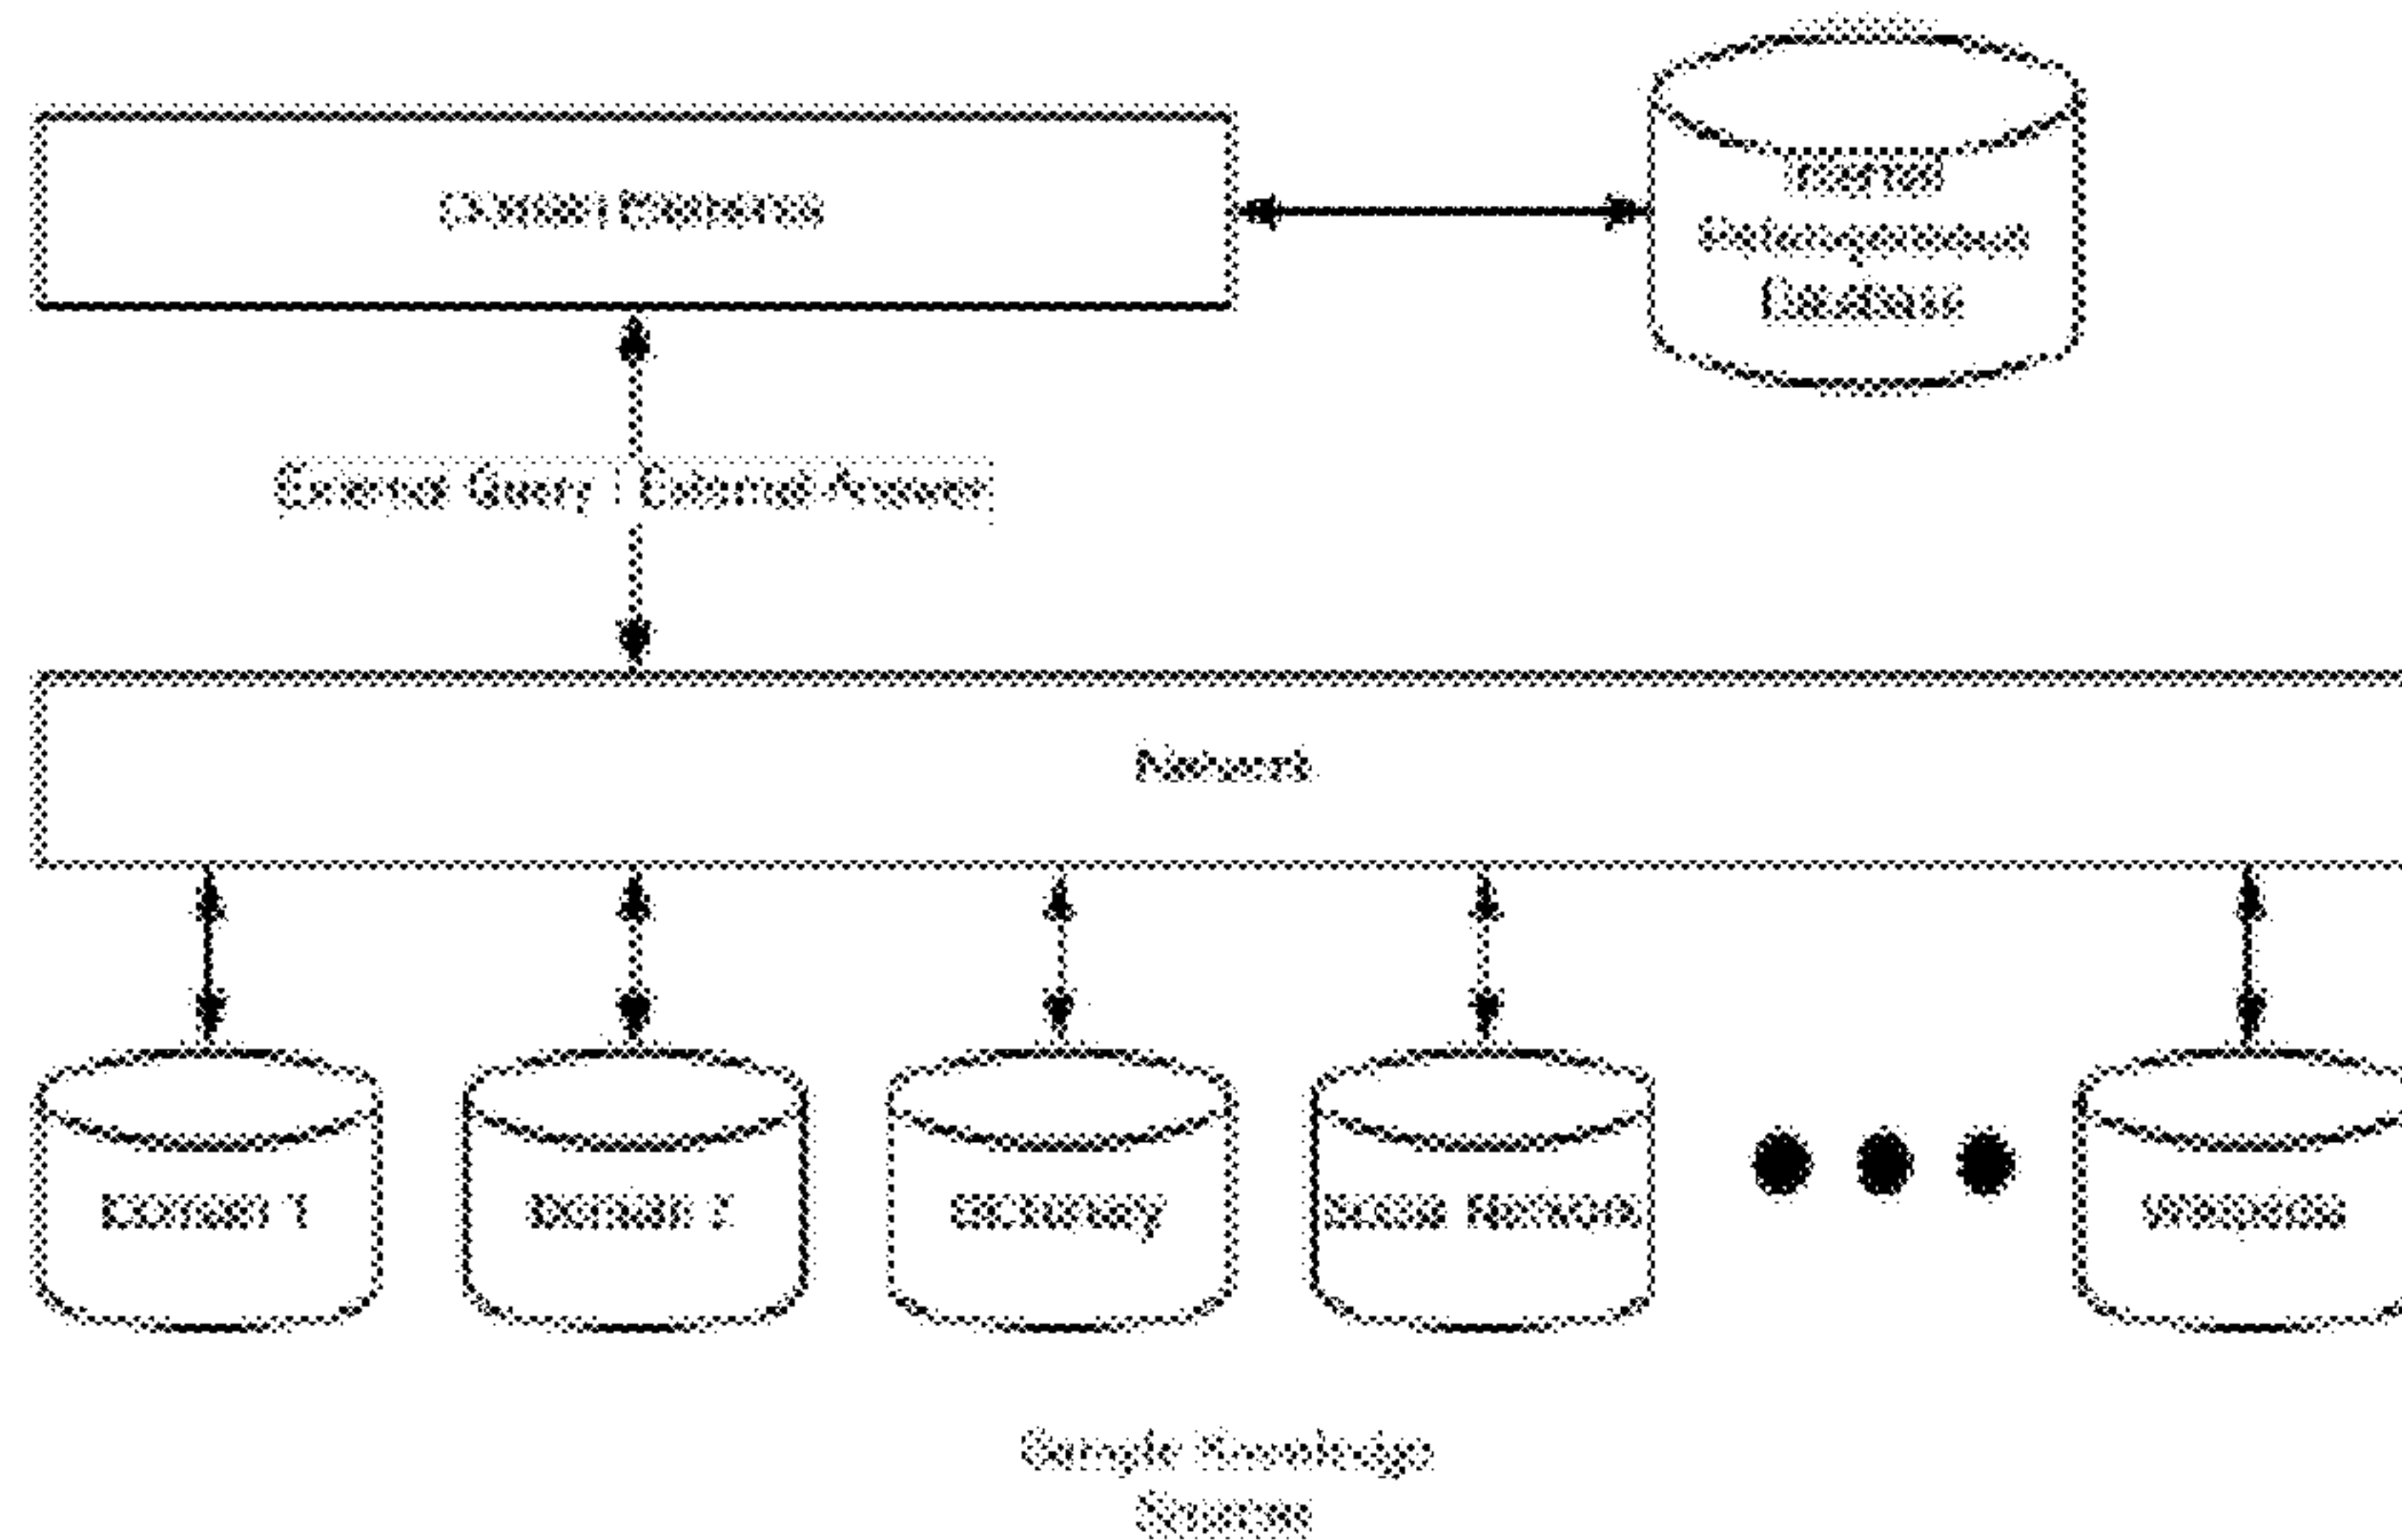


FIG. 20


FIG. 19

1900




FIG. 21

2100







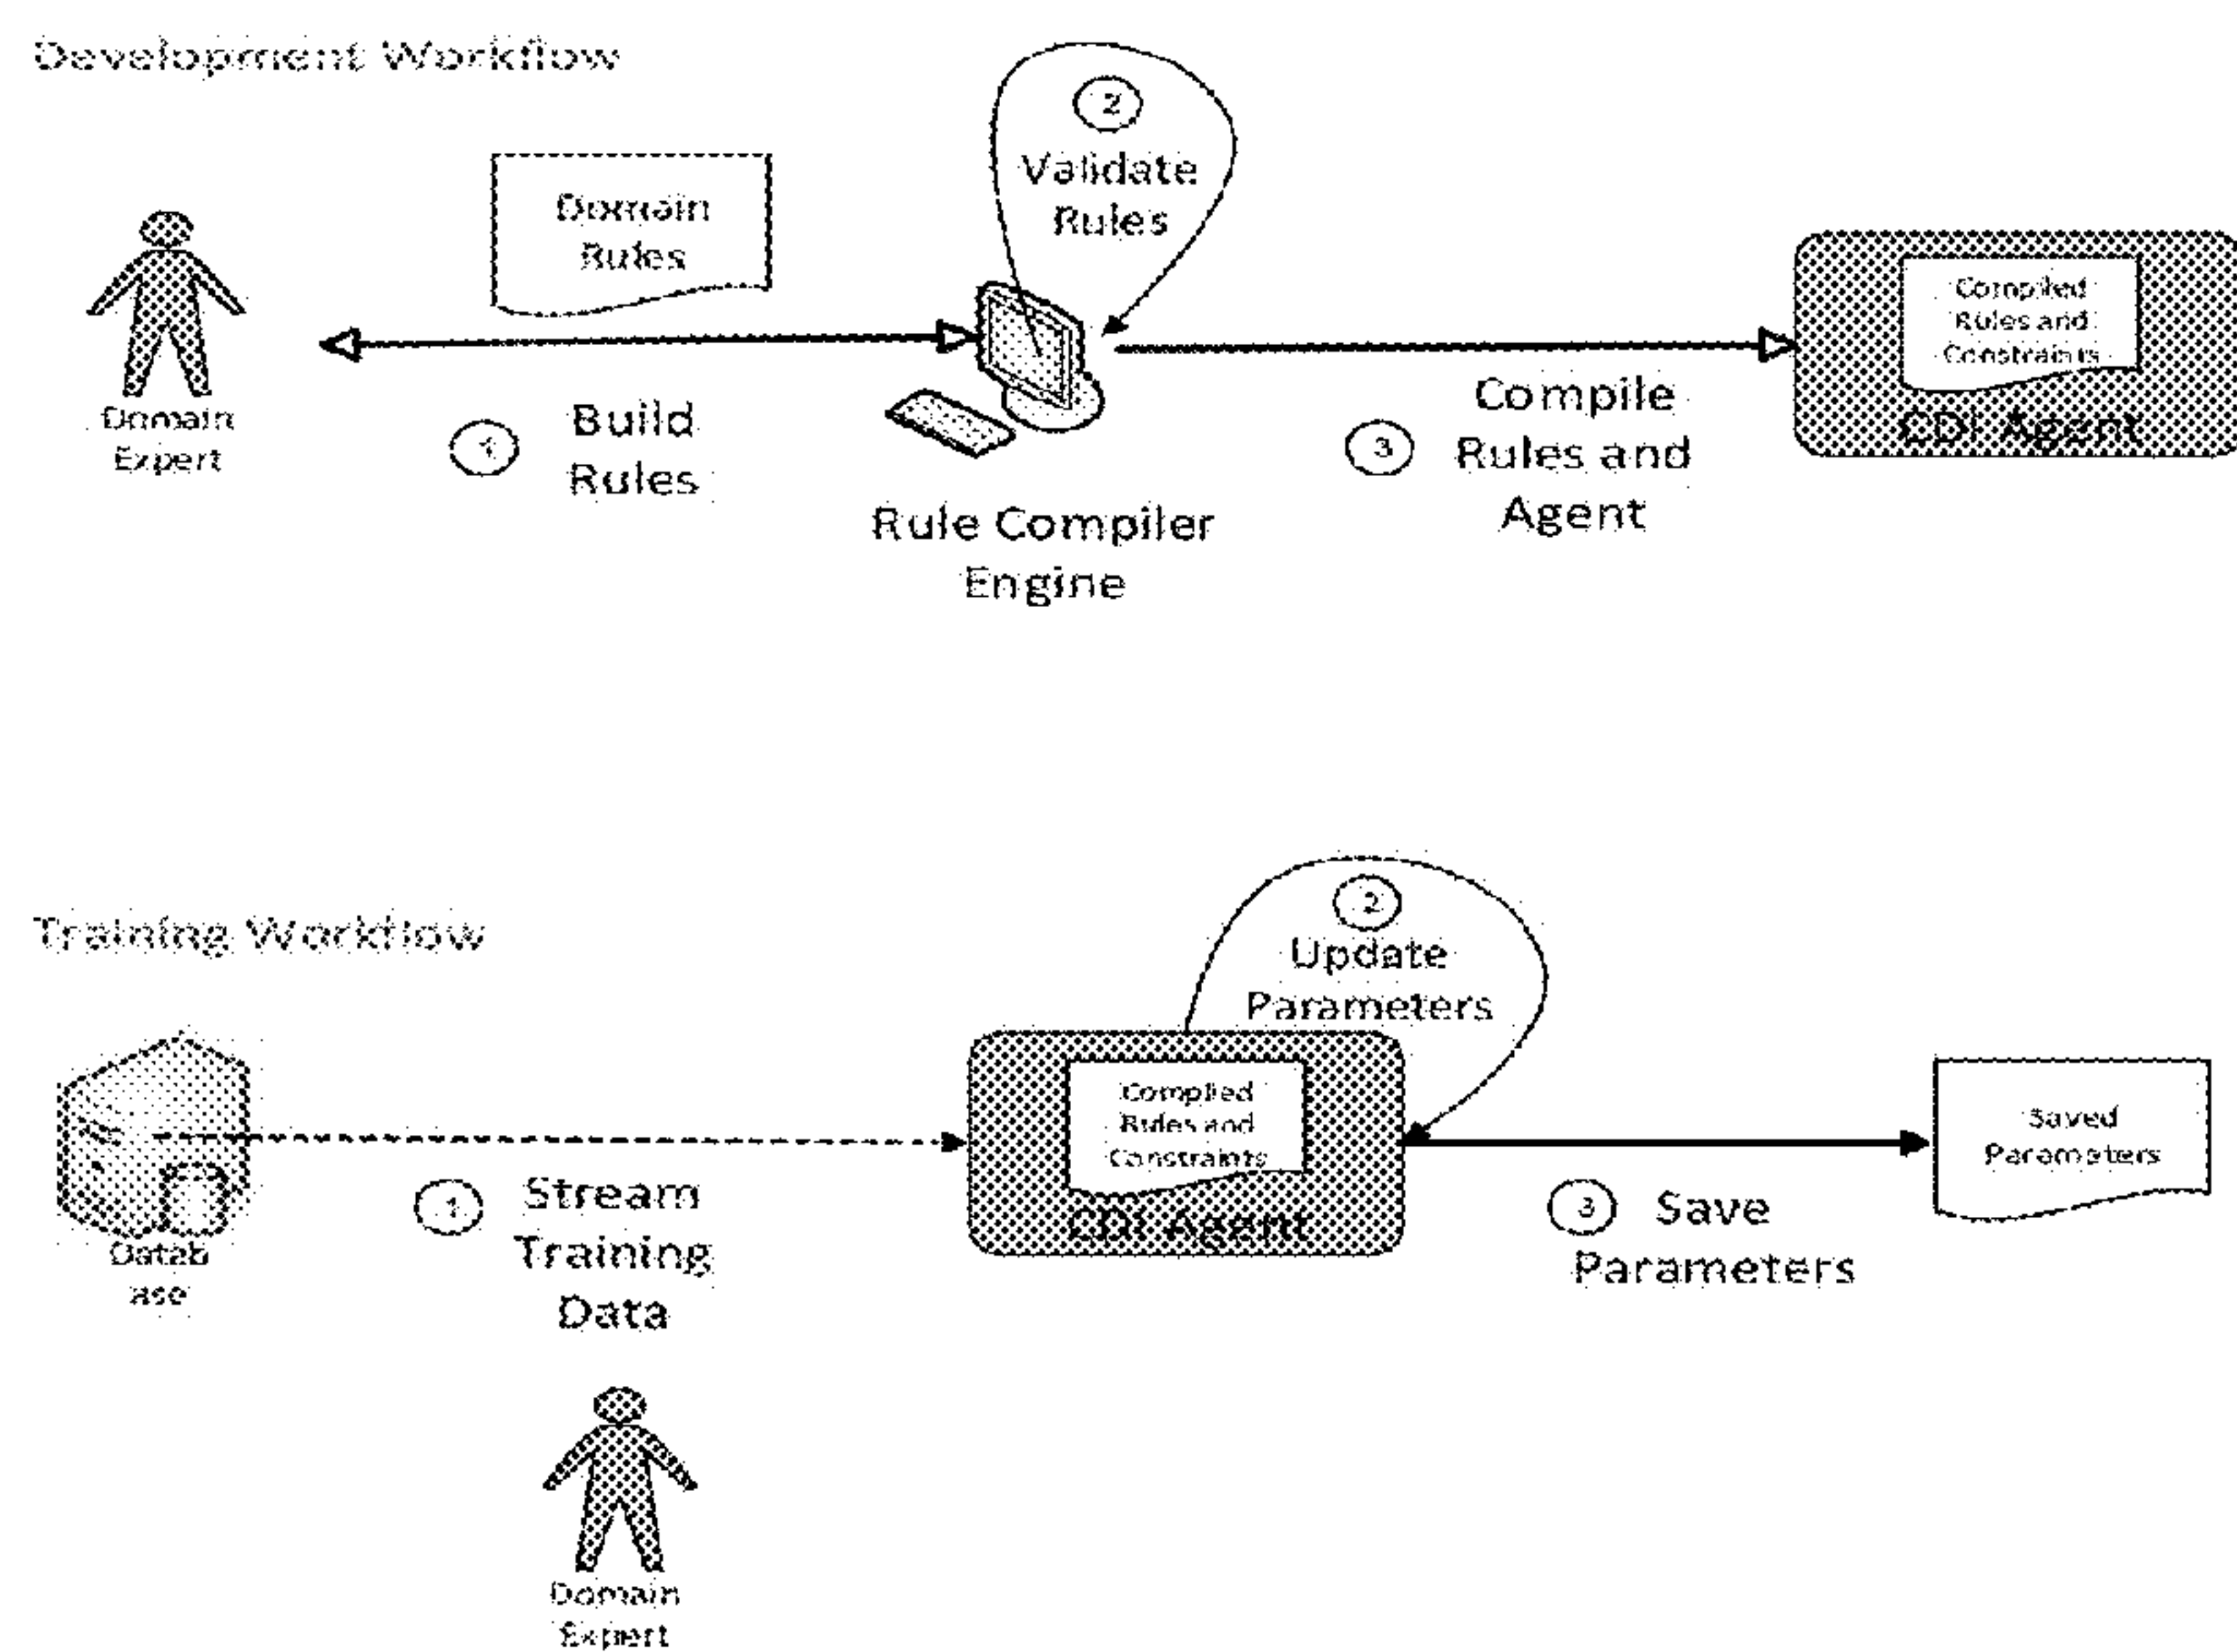
COOPERATIVE DISTRIBUTED INFERENCING (CDI) AGENT



**FIG. 23**

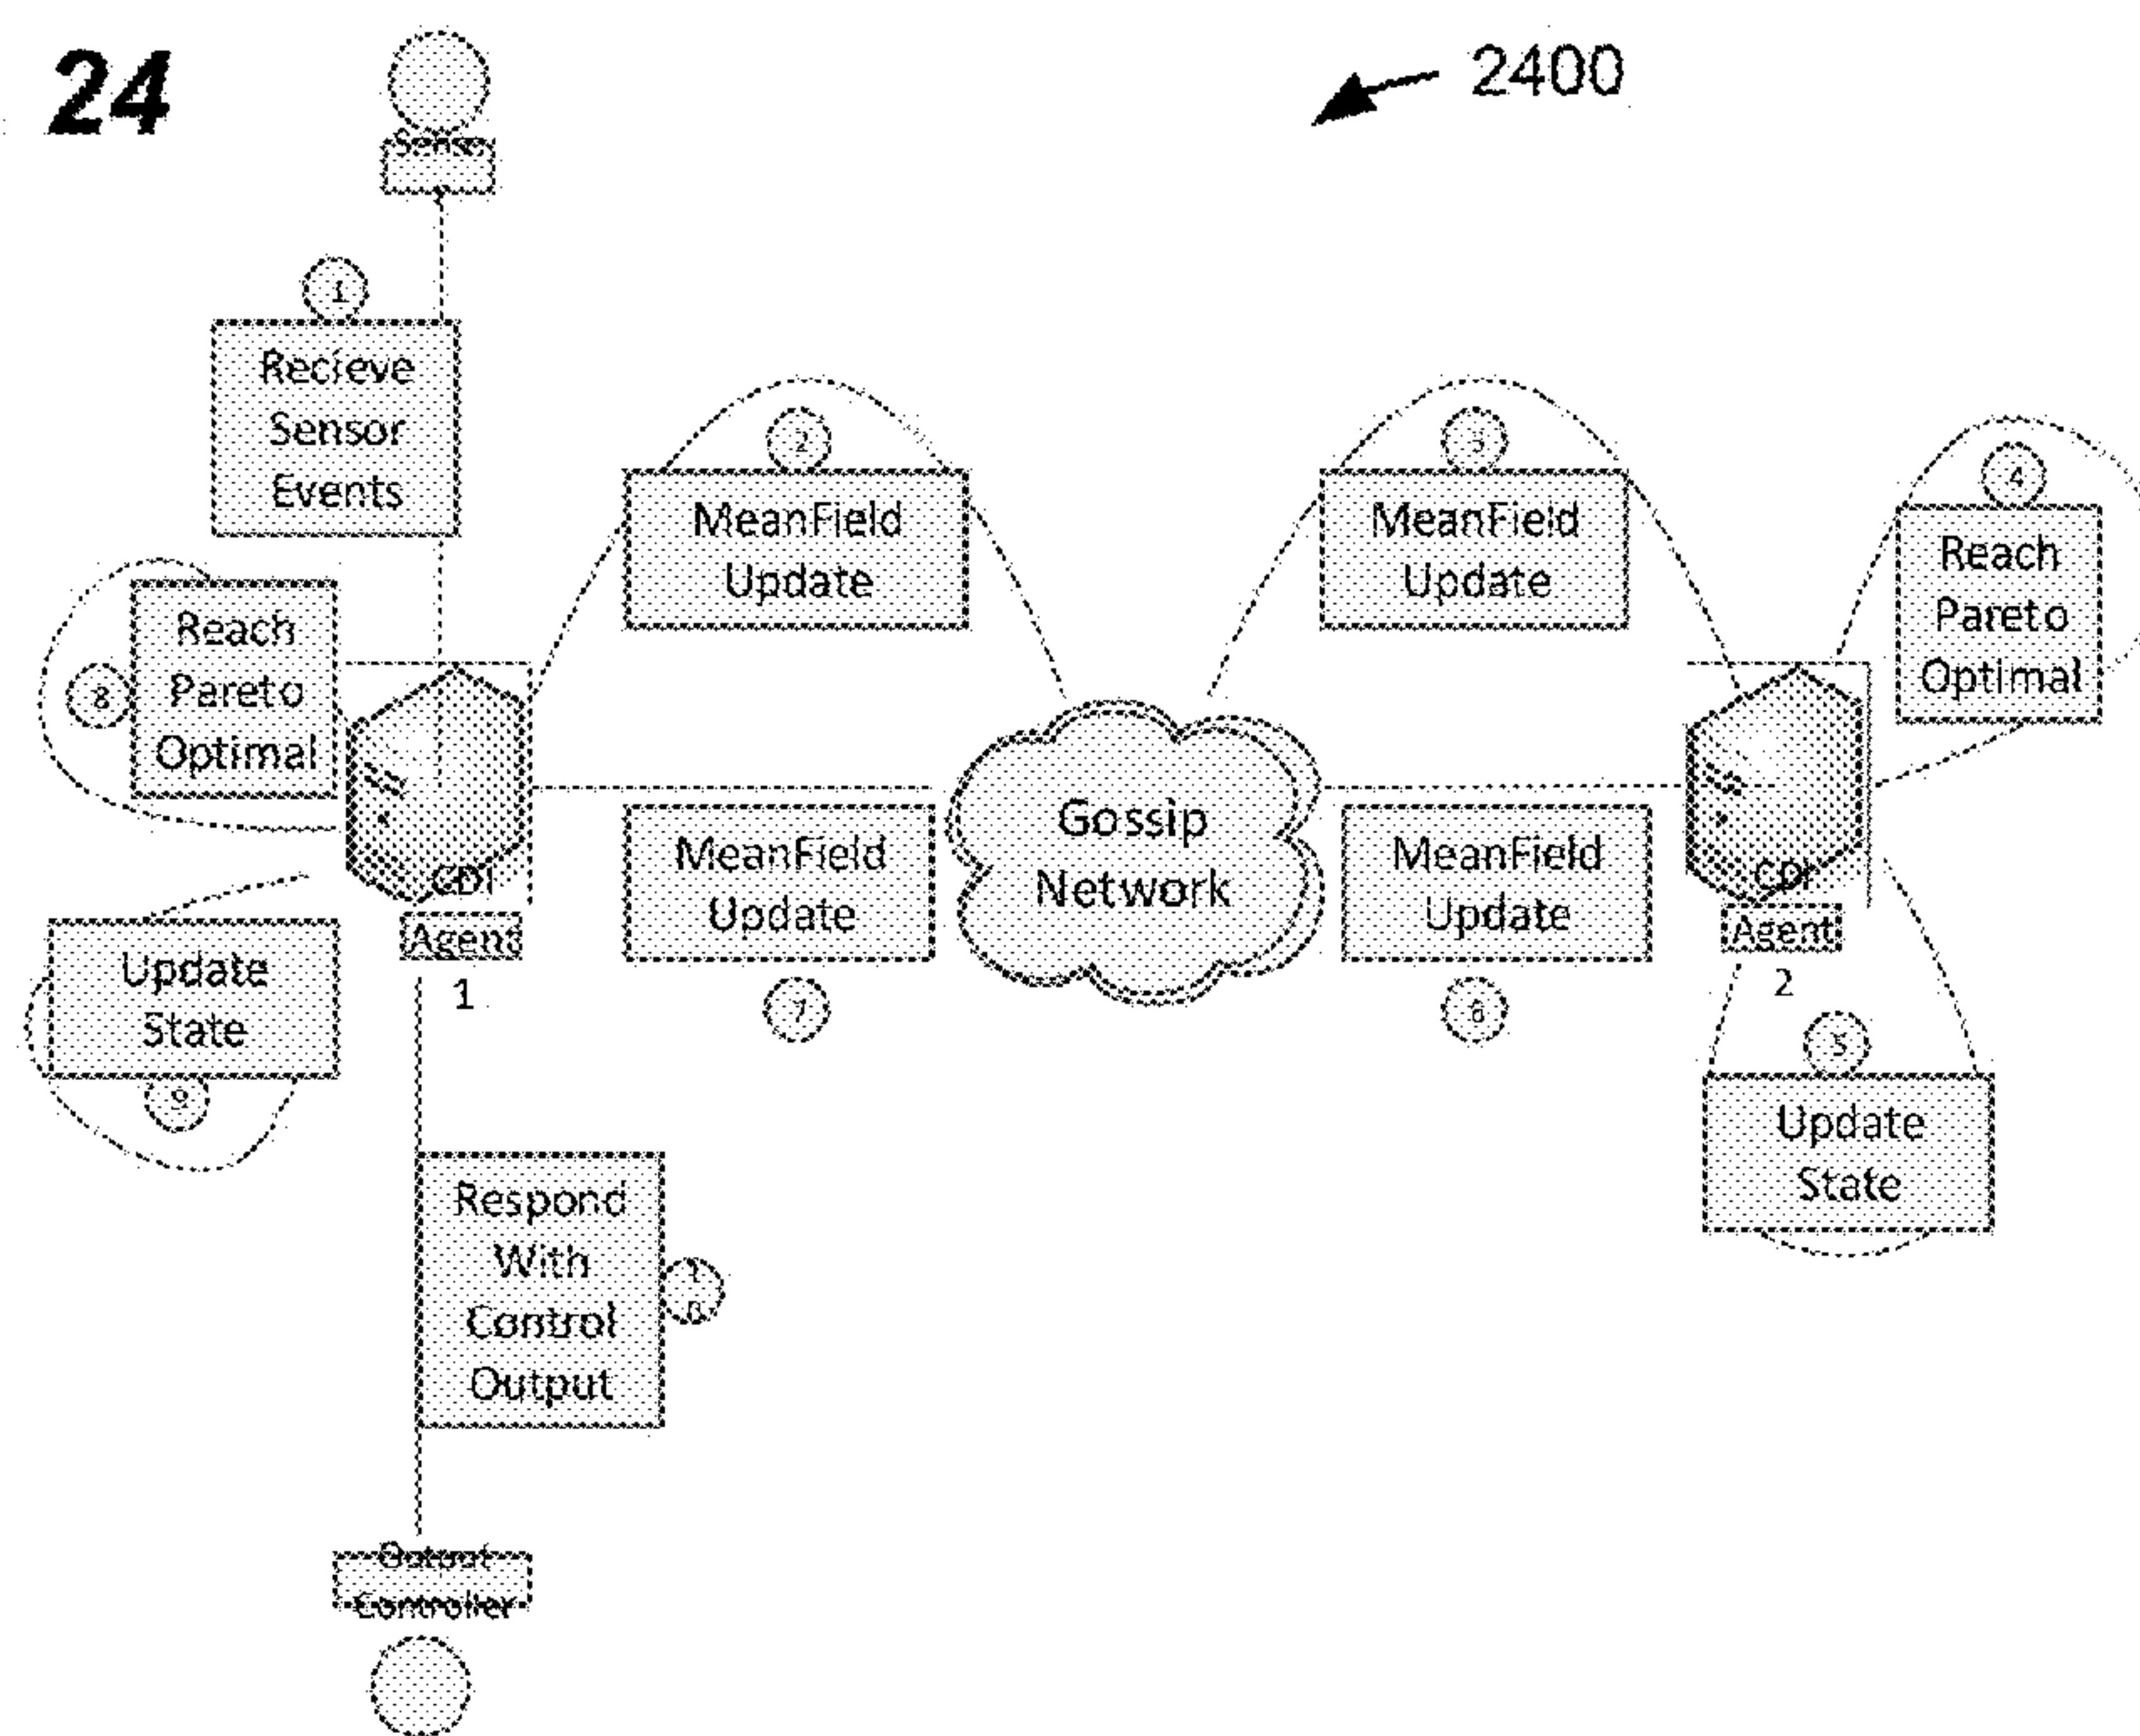
CDI Offline Workflows

2300



**FIG. 24**

2400




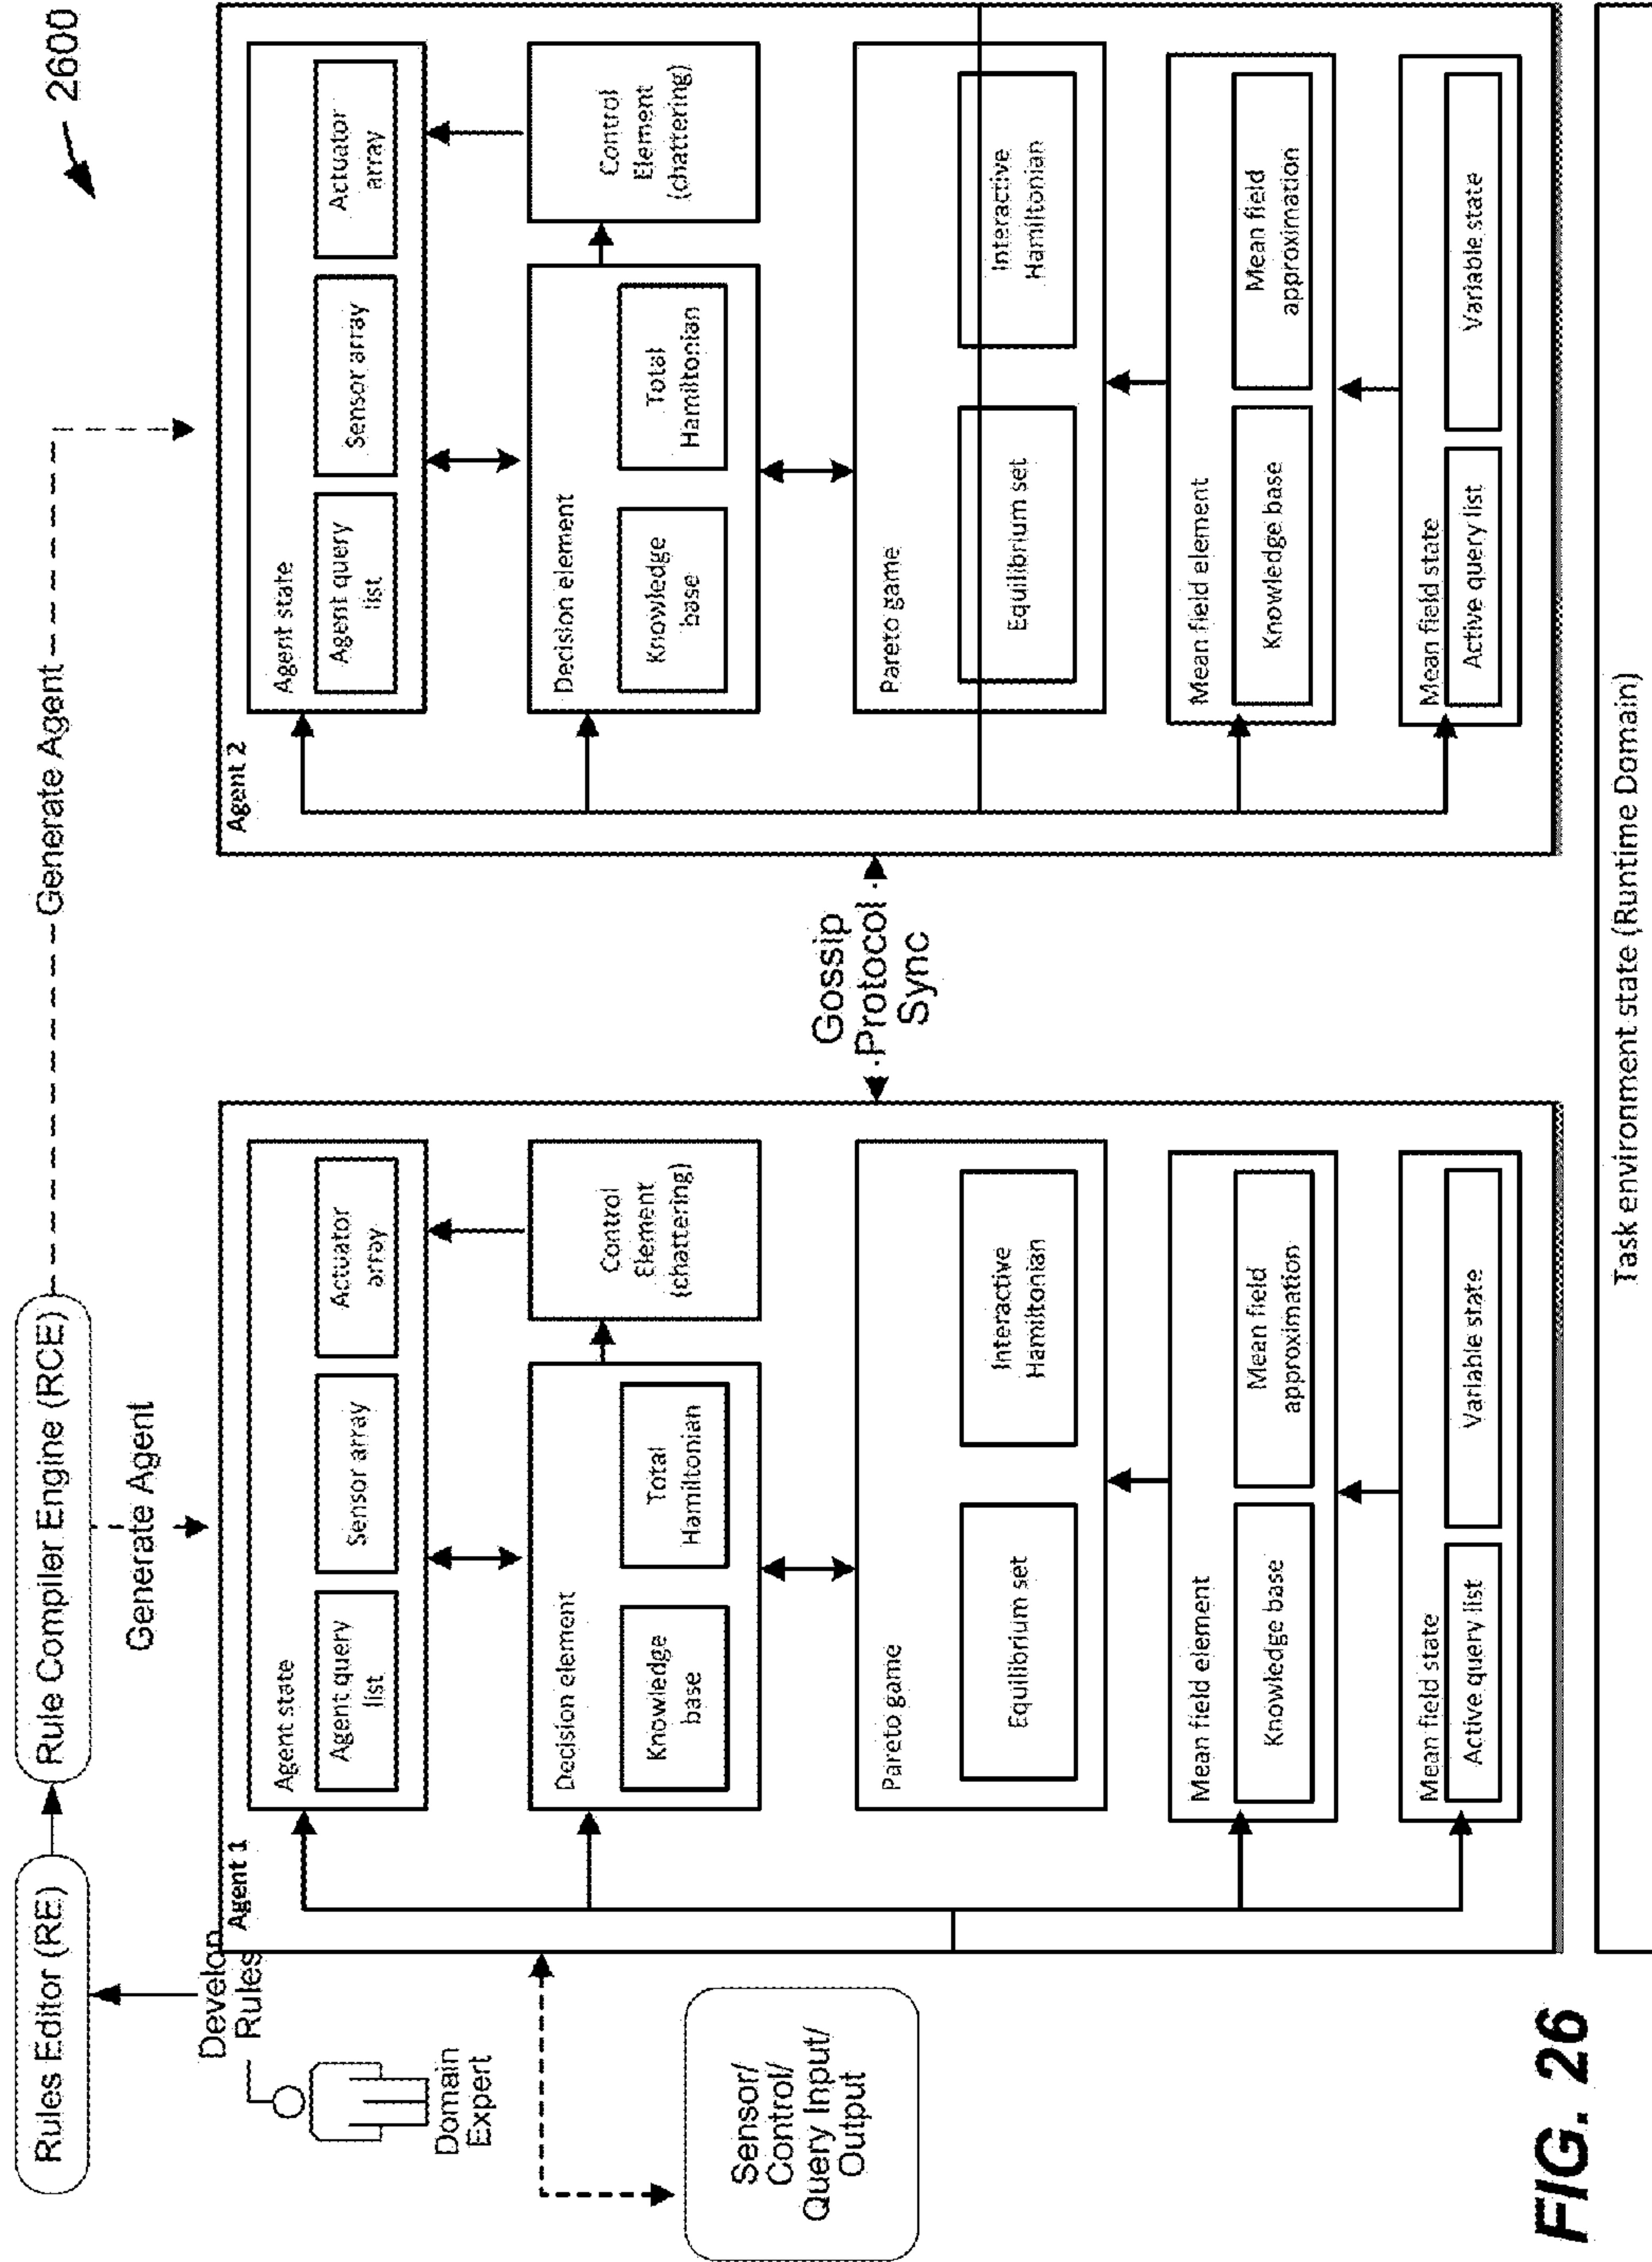


FIG. 25

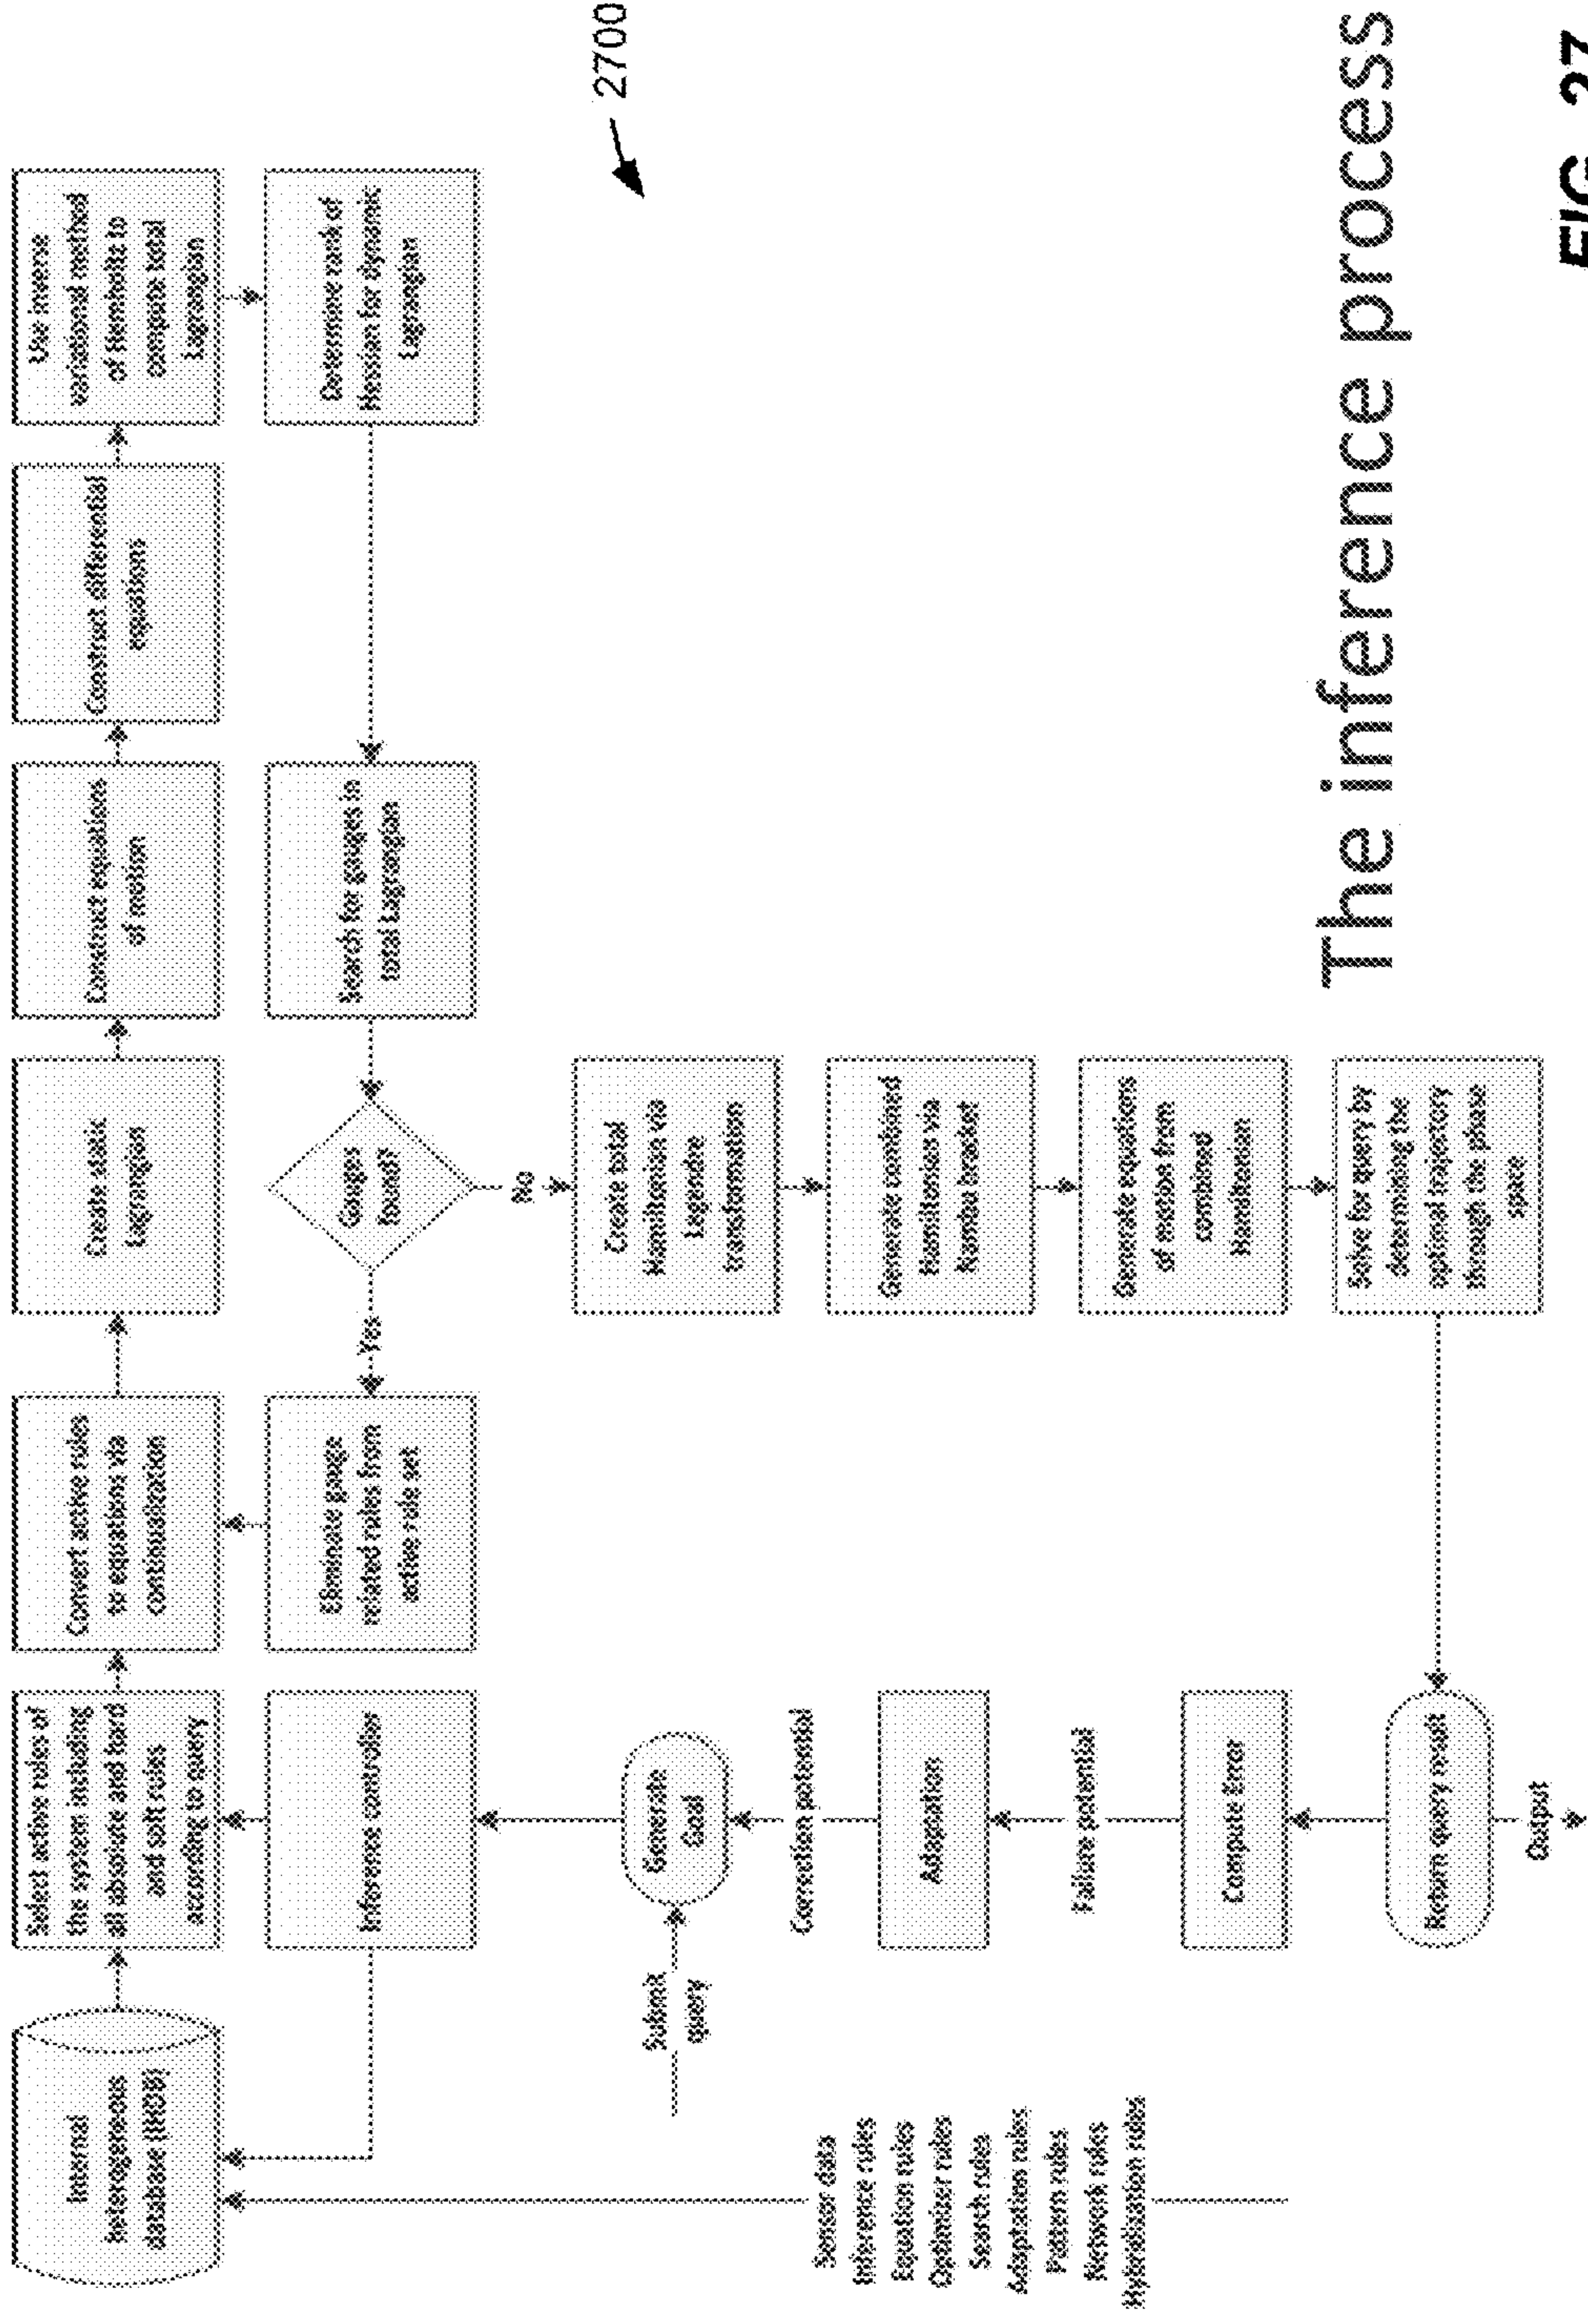


# CDI Design Diagram



**FIG. 26**





The inference process

FIG. 27





2900A

HDP Step-1 An Injection map exists




FIG. 29A



2900B

HDP Step-2a- If a Local Goal Exists




FIG. 29B

HDP Step-2b- If a Local Goal Does Not Exist ↖ 2900C





FIG. 29C

HDP STEP 3-a: -There is an optimal transition ↖ 2900D

- Constraints are satisfied
- No New rules active/deactive




**FIG. 29D**



HDP STEP 3-b: -There is an optimal transition ↖ 2900E

- Constraints are satisfied
- New rules active/deactive



**FIG. 29E**

Agent Behavior is a Trajectory in Domain Manifold. ← 2900F





FIG. 29F

# Relation Between Knowledge Base 2900G

## And Carrier Manifold



As Binary Relations on the state of the process. The rules in the knowledge base define a basis for the topology of the carrier manifold

**FIG. 29G**



# Conceptual Illustration of an Agent's $\epsilon$ -Optimality Procedure





FIG. 29H



**FIG. 291**

# Agent Behavior is a Trajectory in Domain Manifold




FIG. 29J



2900K

First Theorem of Chattering: Near optimal flow is integrand  
of Action resulting from Chattering of primitive actions




FIG. 29K

3000




Fig. 30




Fig. 31






Fig. 32




Fig. 33

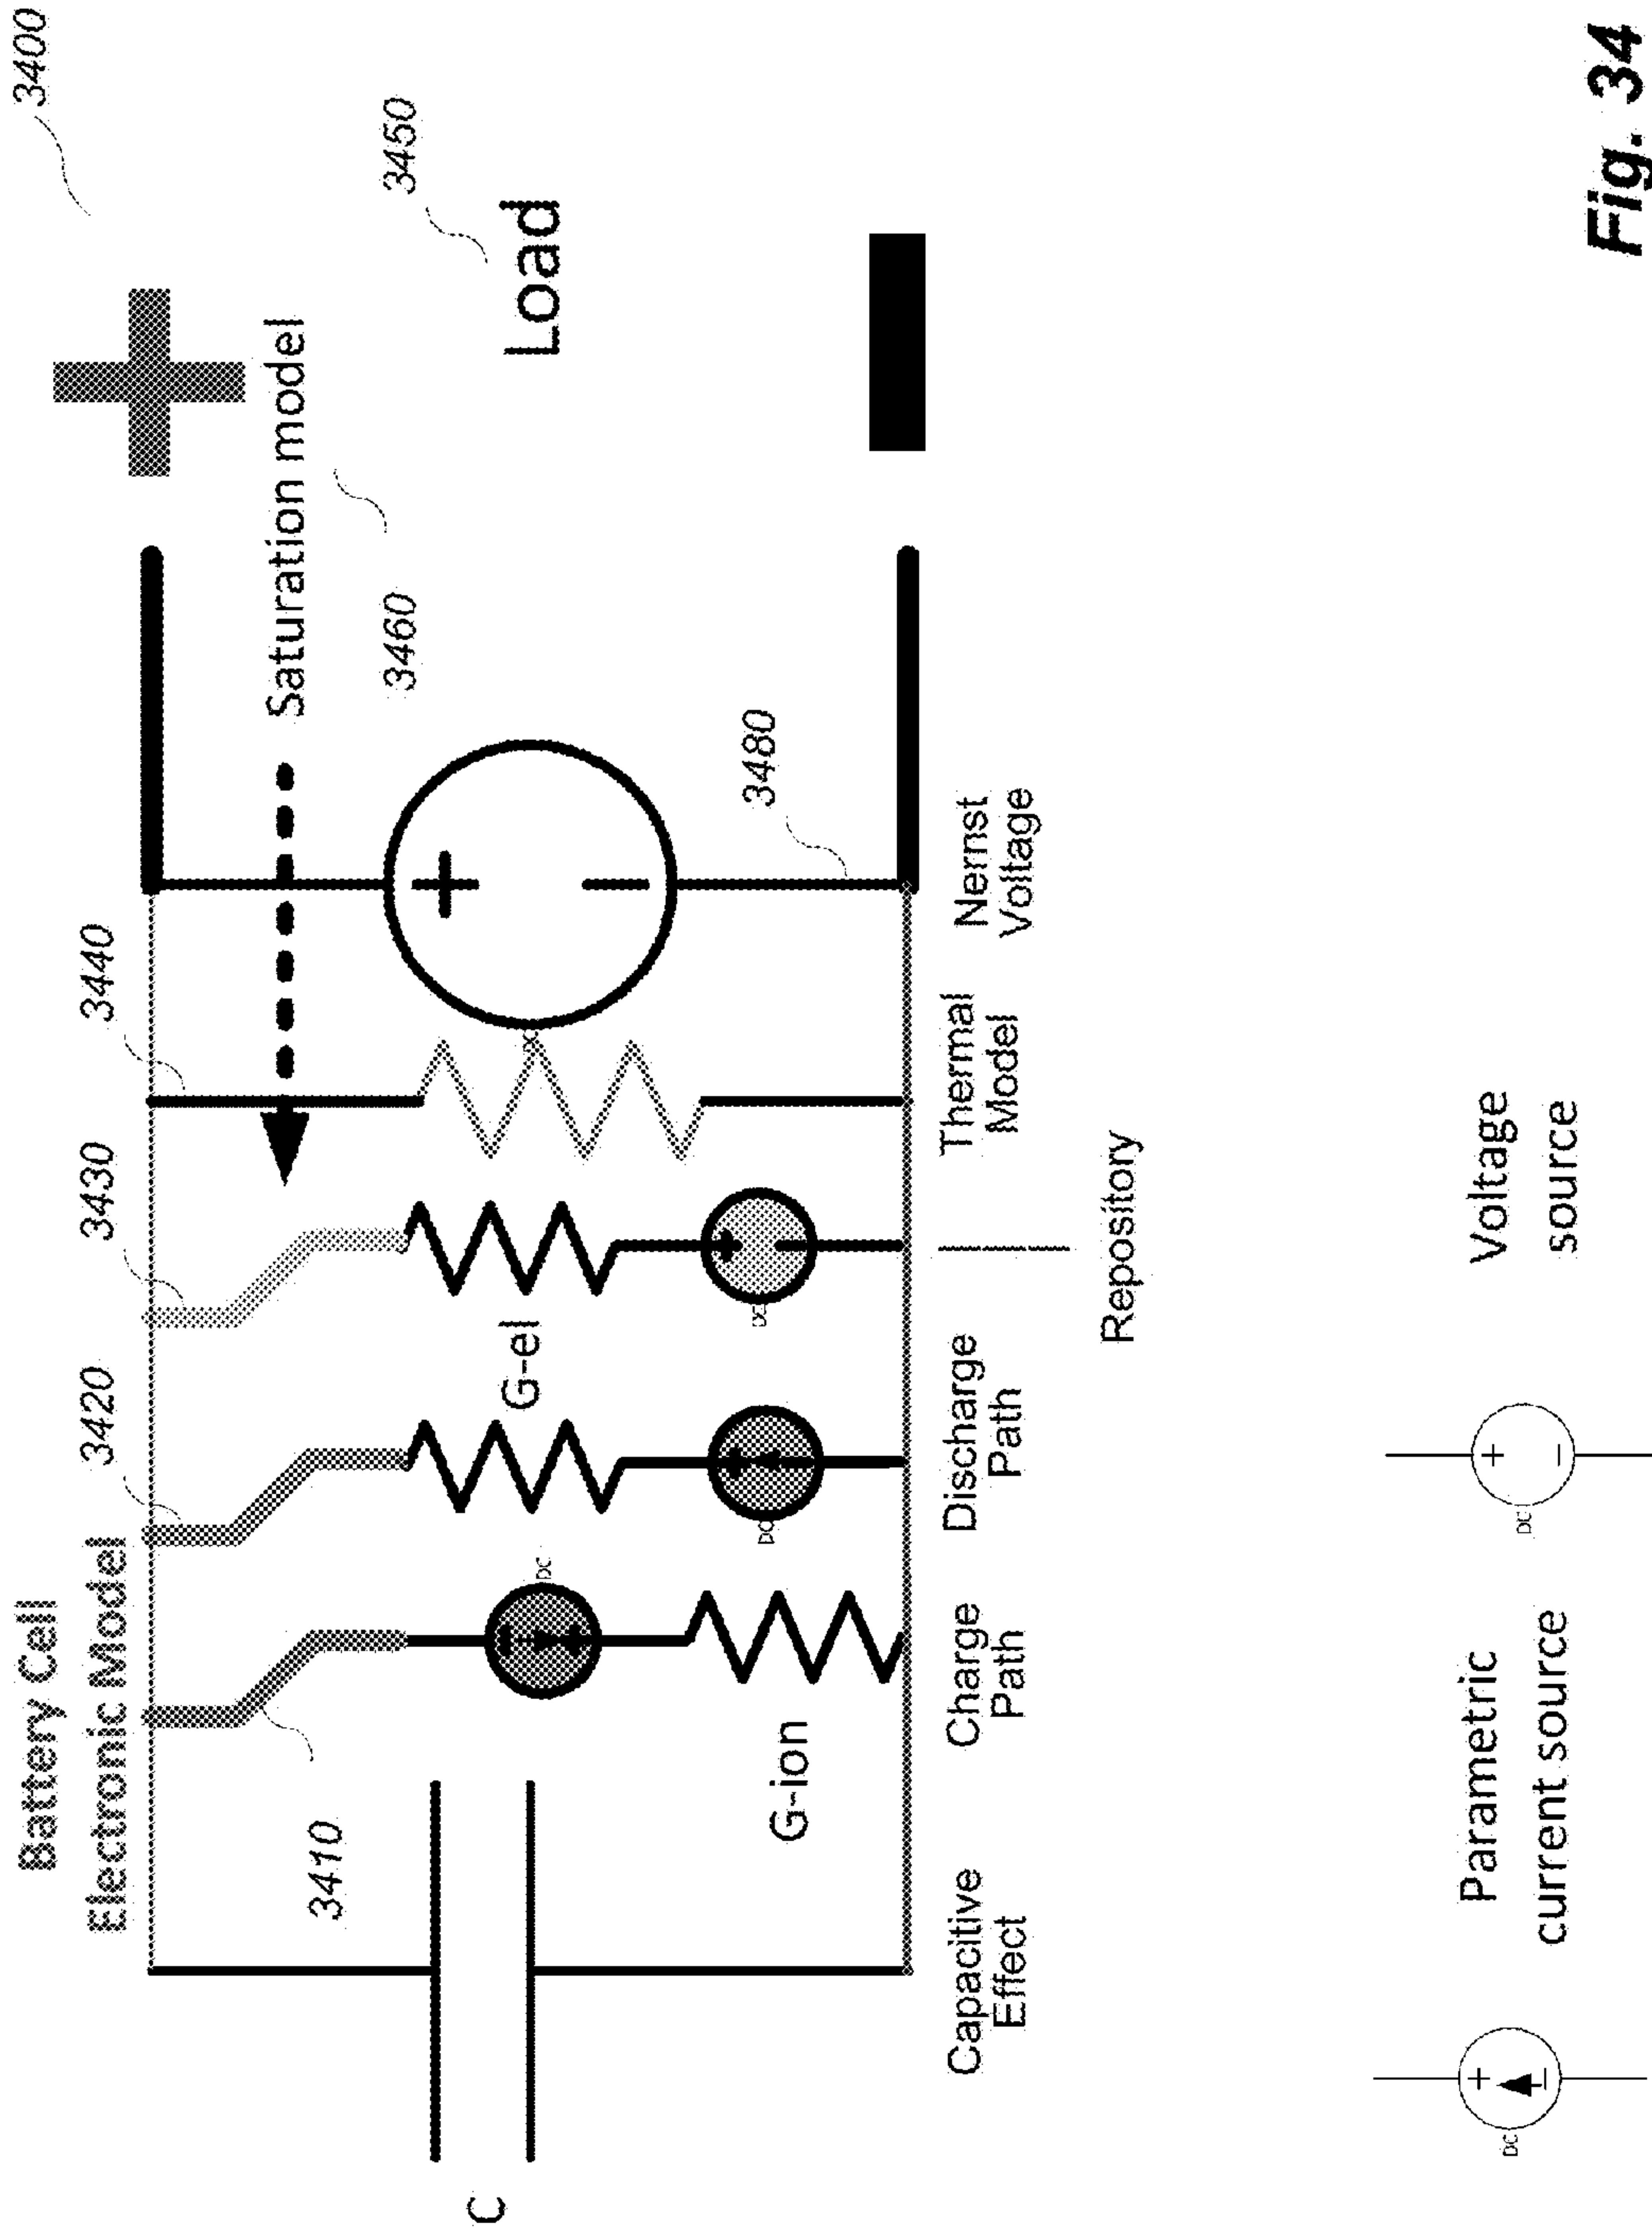


Fig. 34



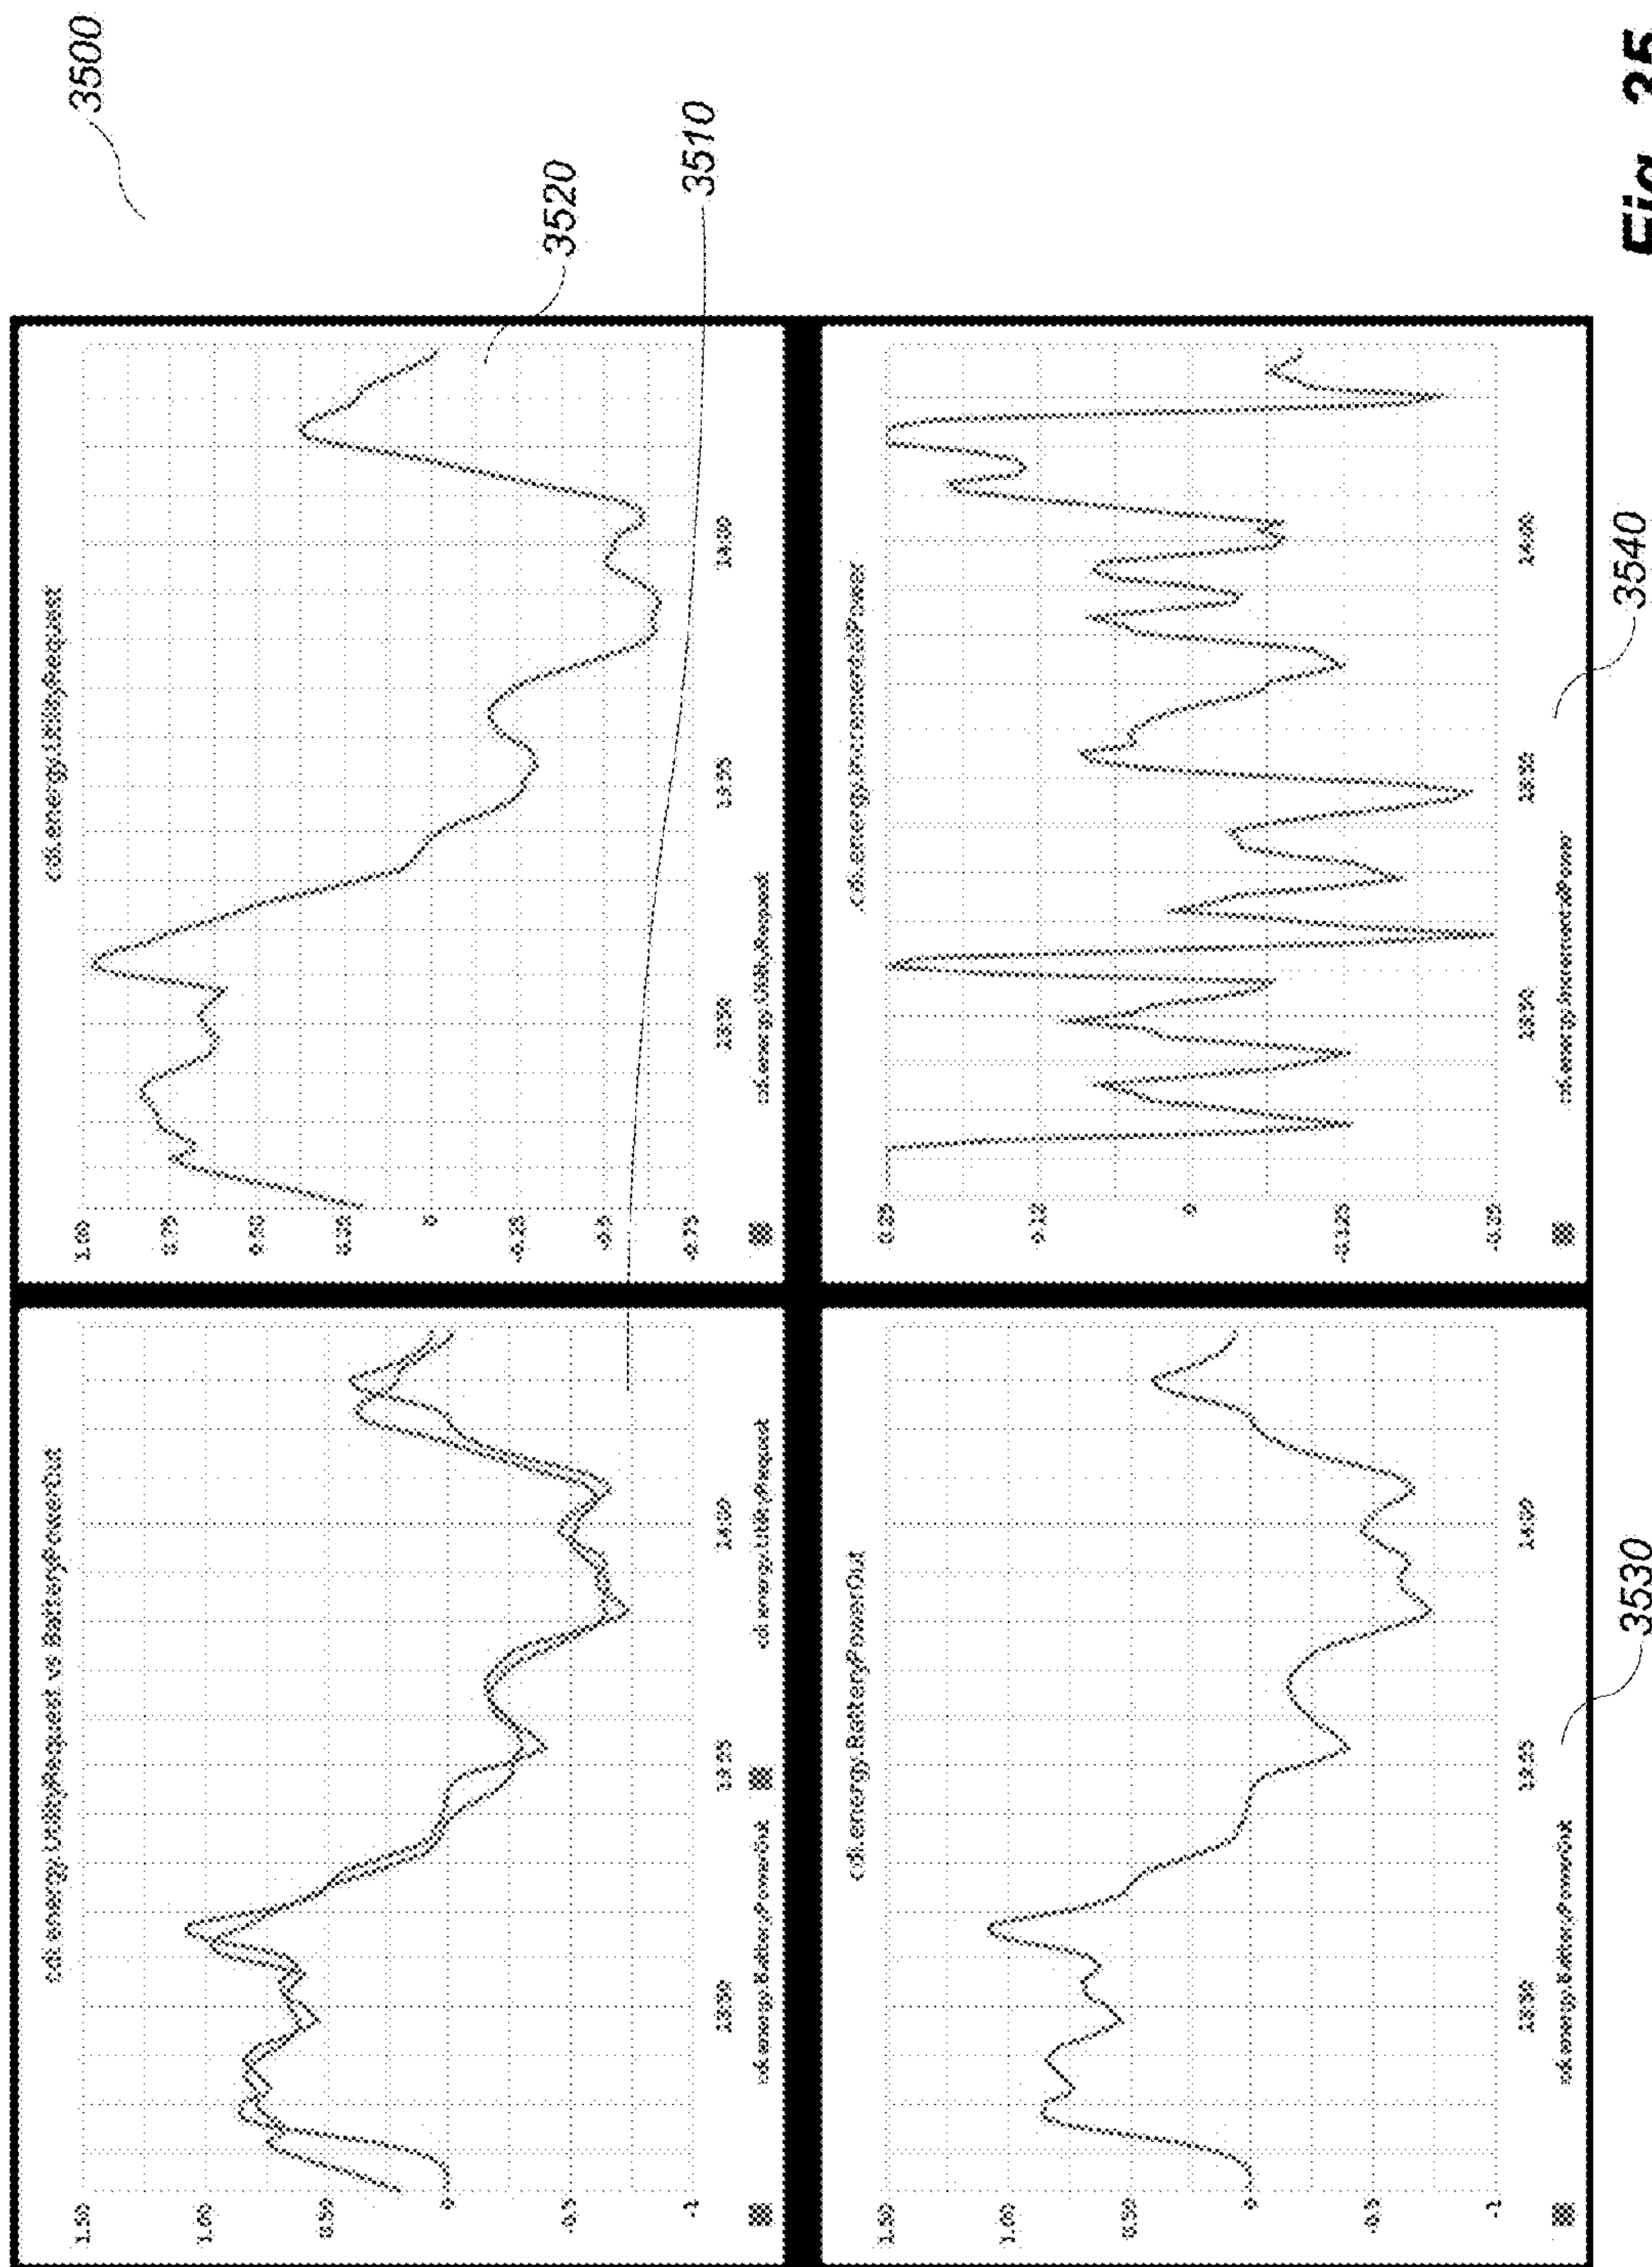


Fig. 35

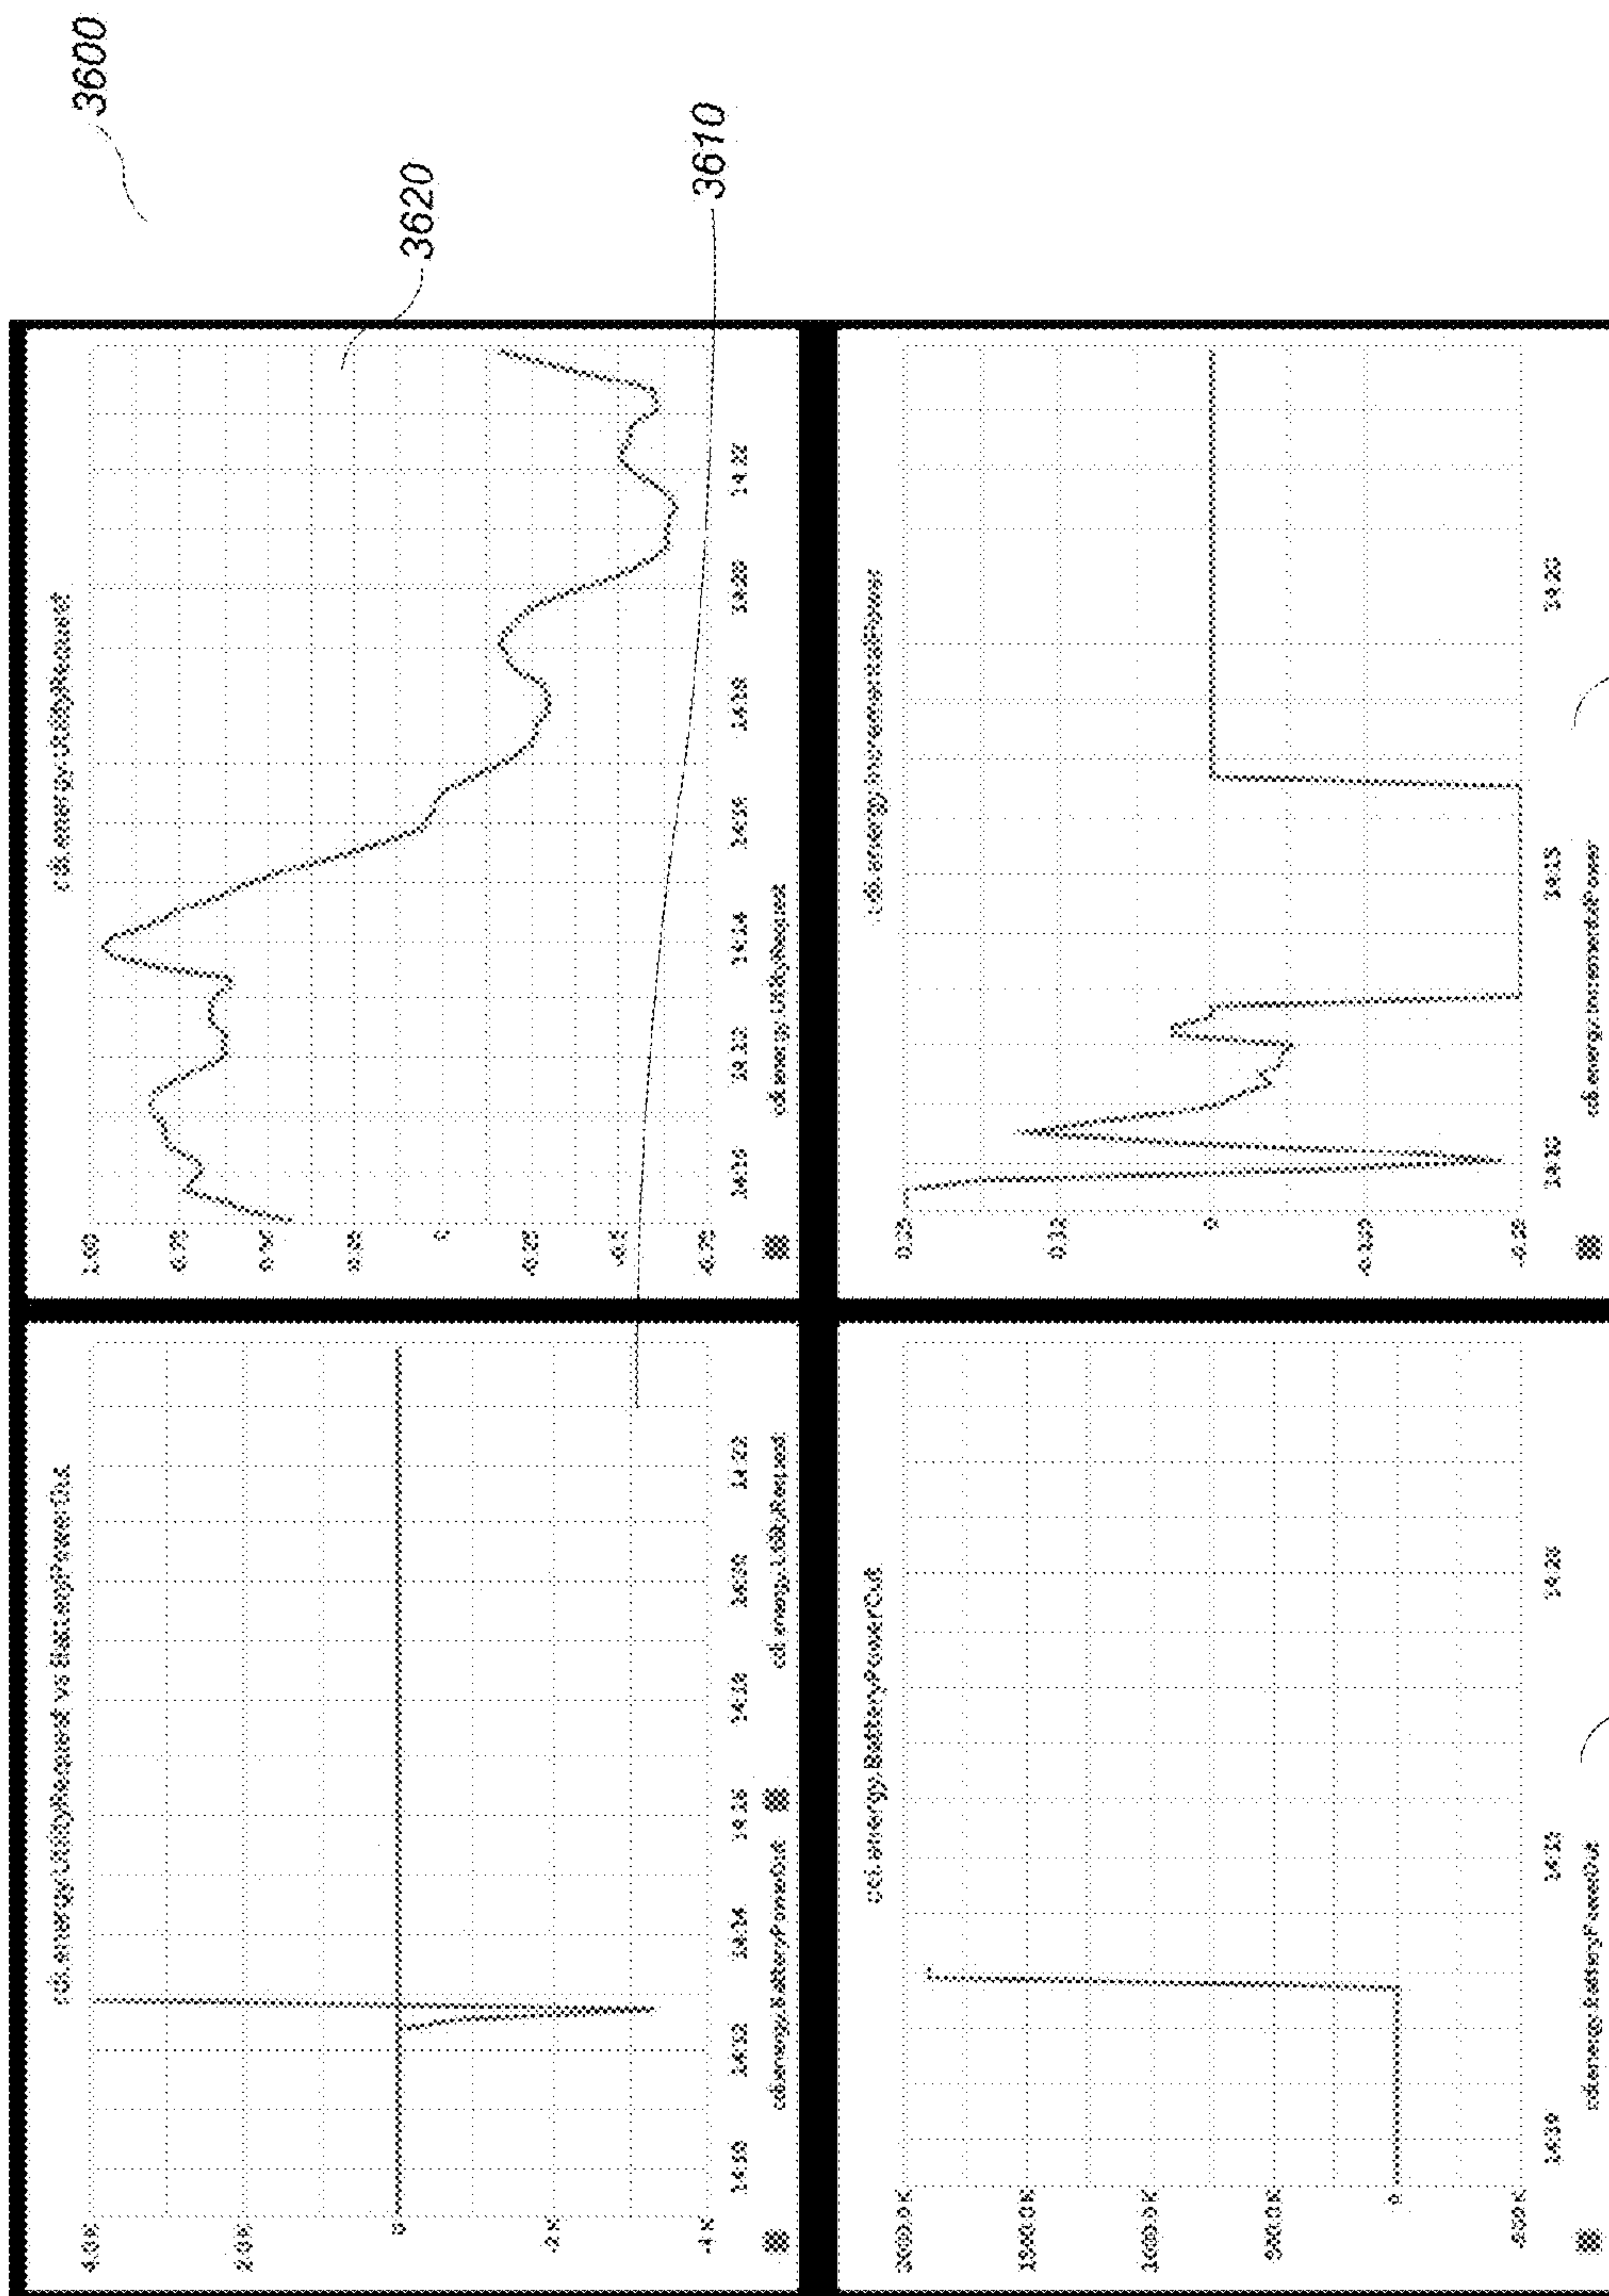


Fig. 36

3640

3630



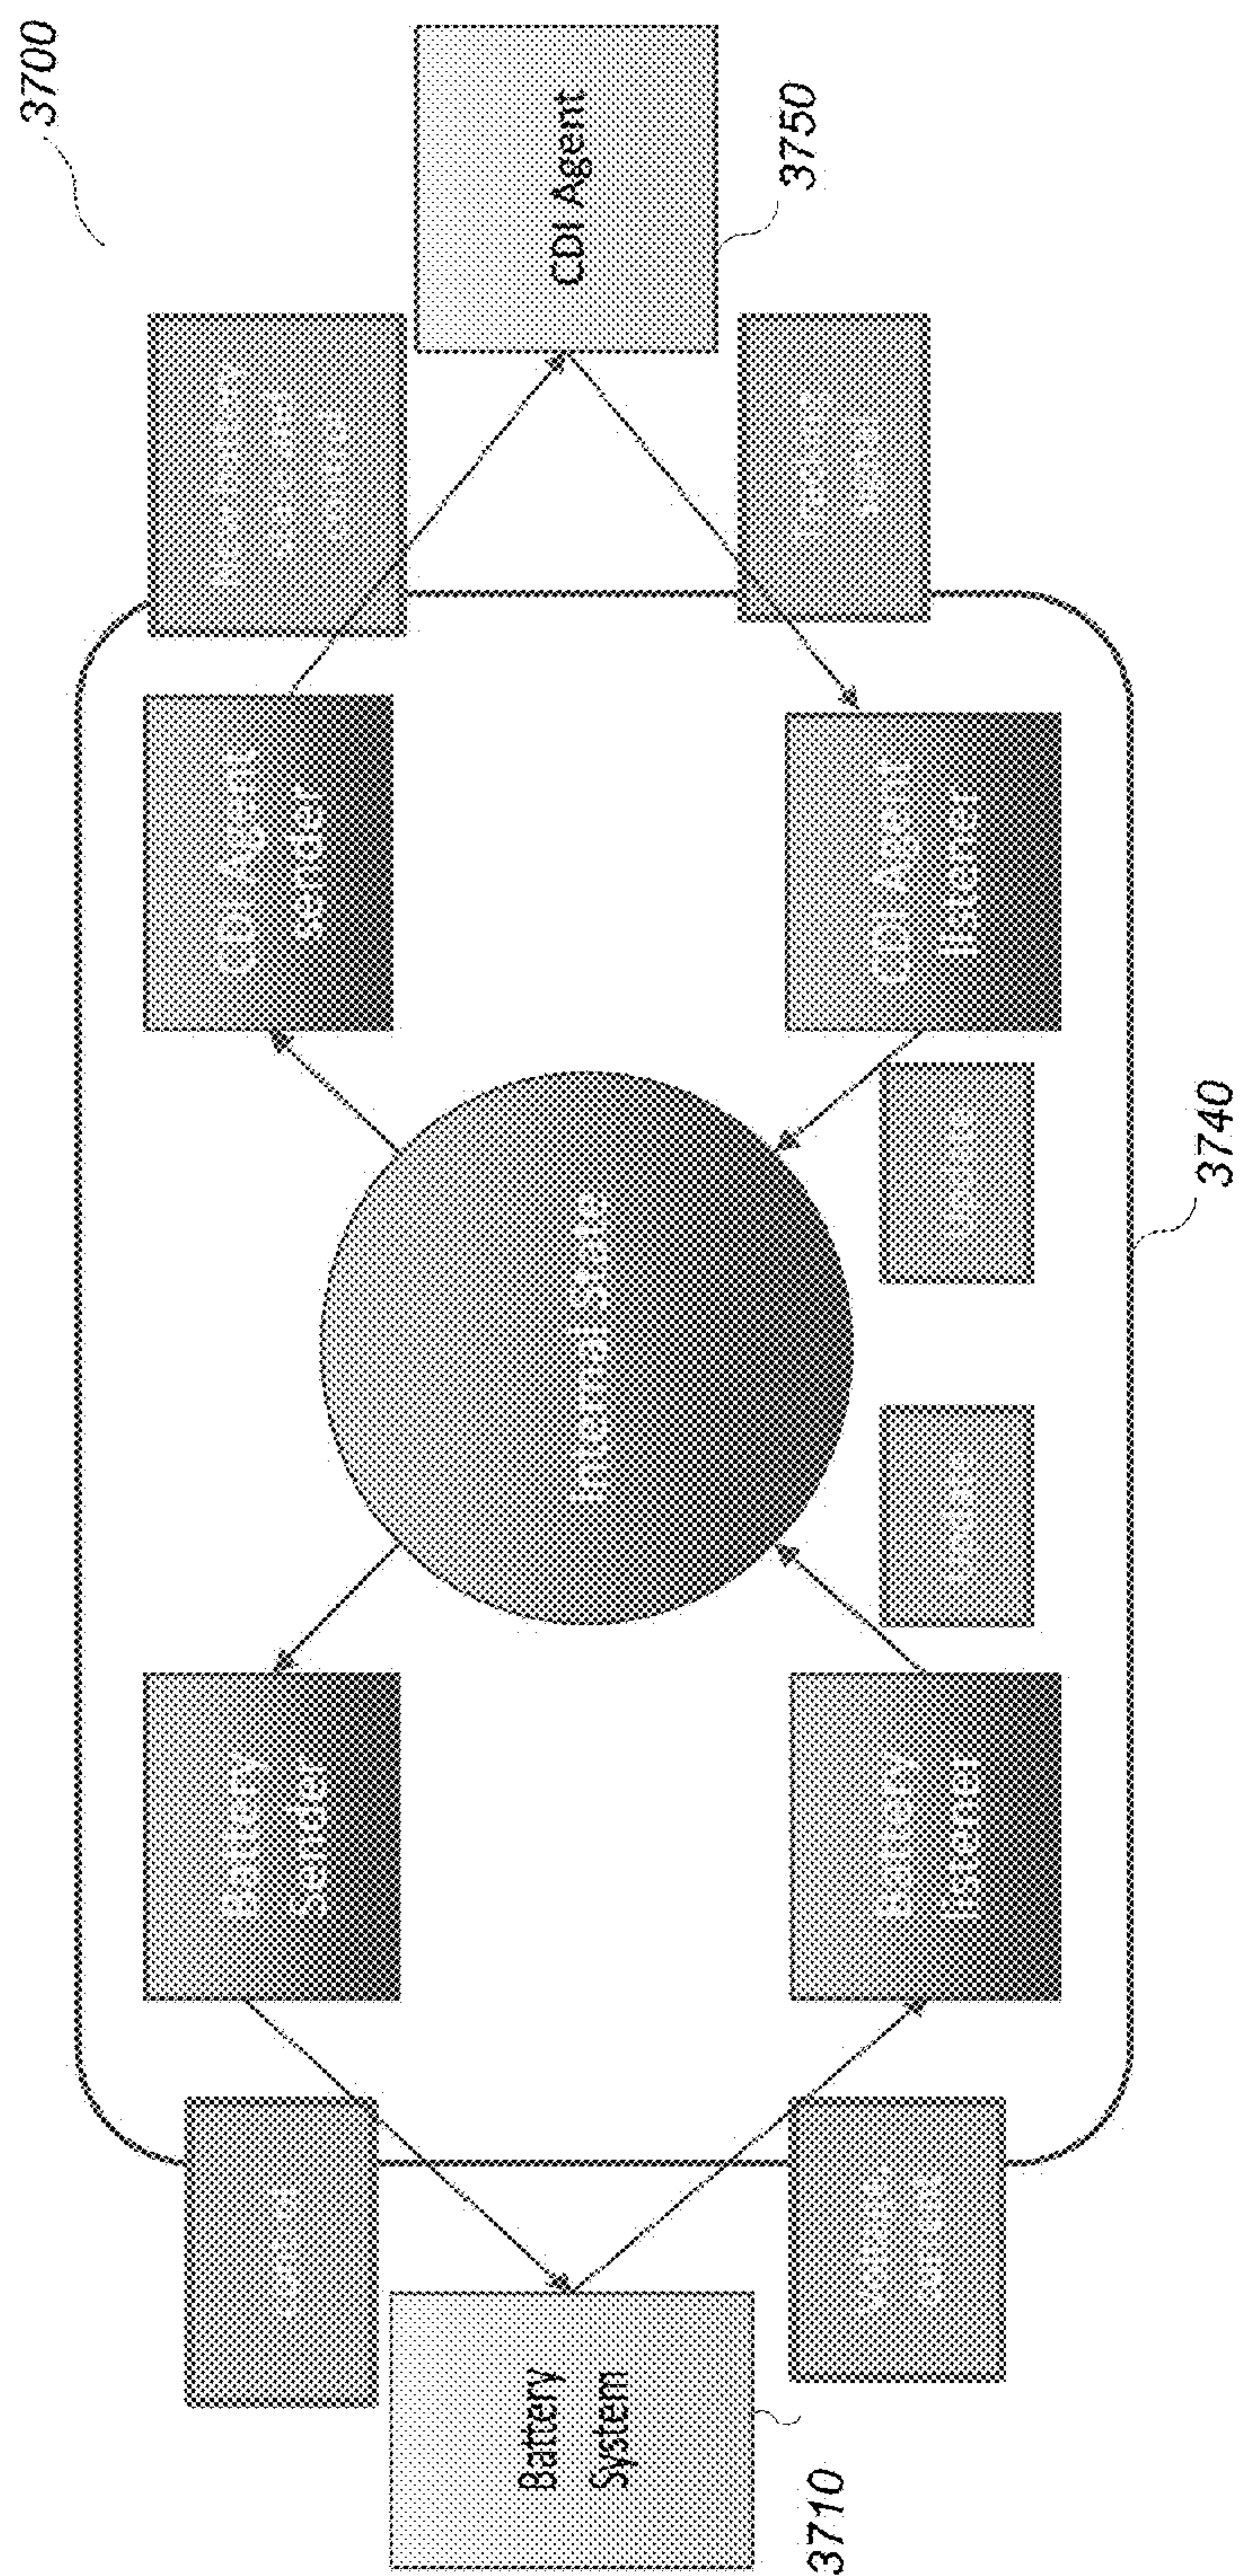


Fig. 37






Fig. 38

3900

- Battery parameters are updated dynamically to capture changes in the battery




Fig. 39



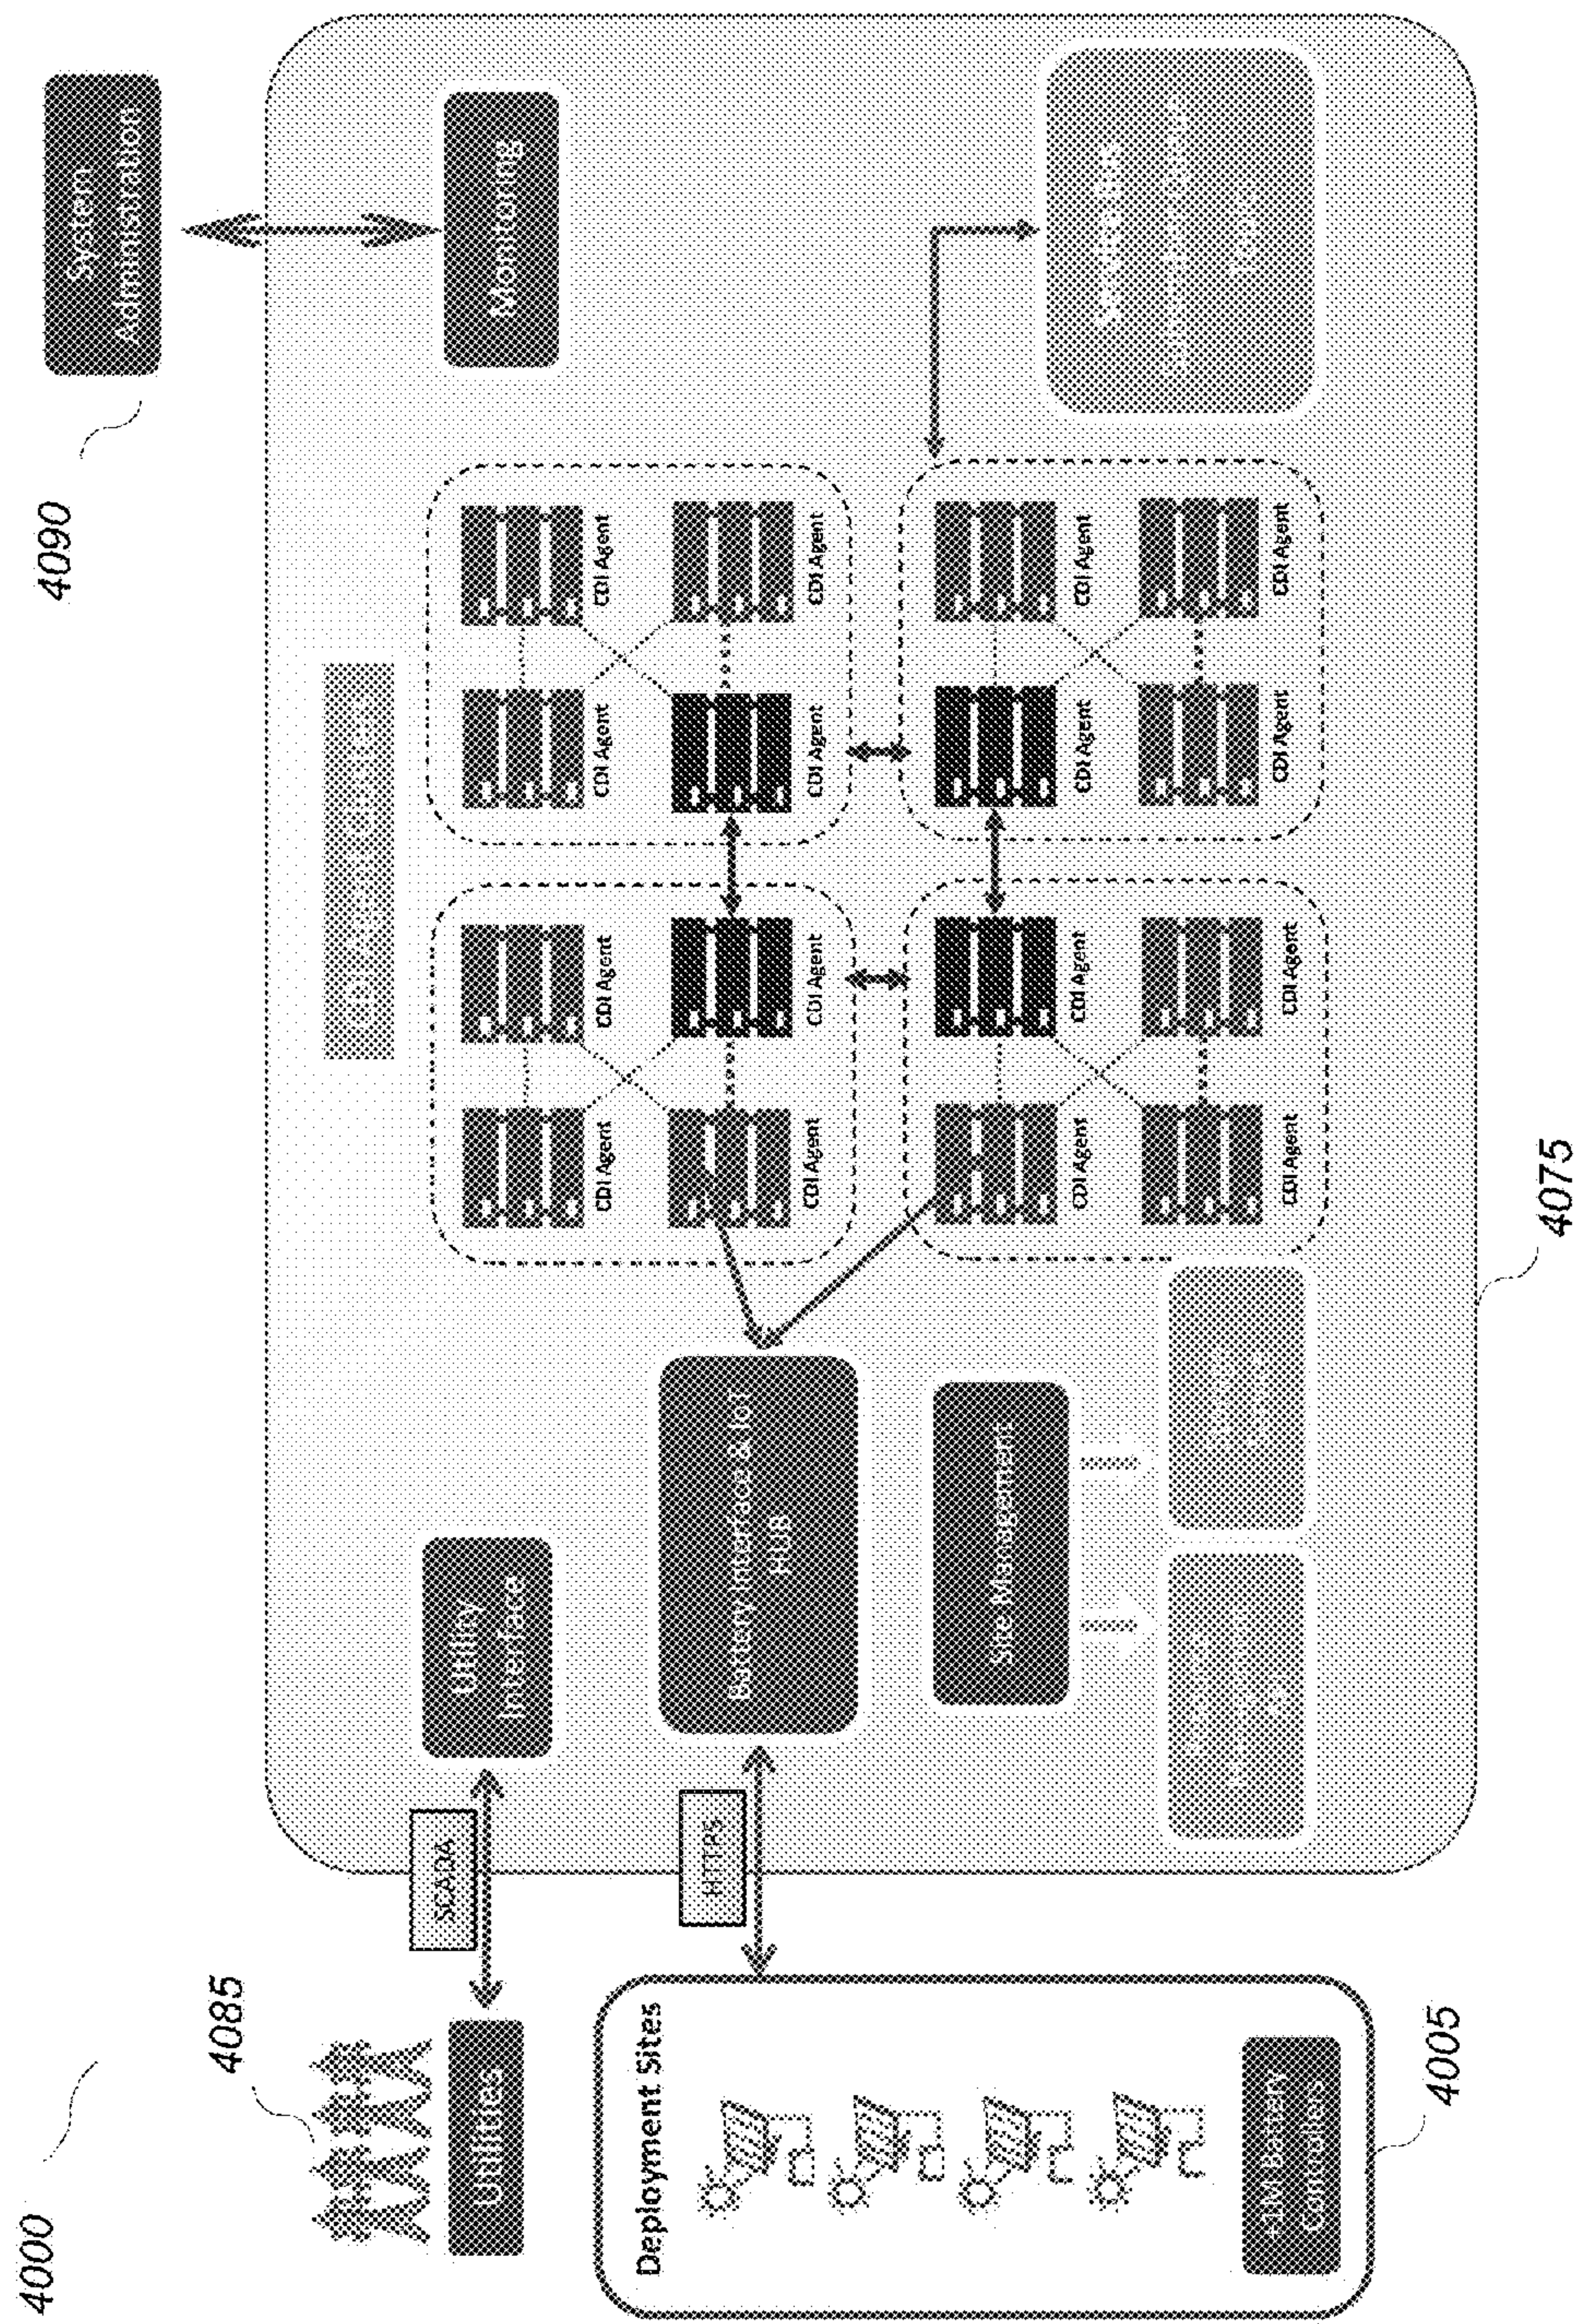


Fig. 40



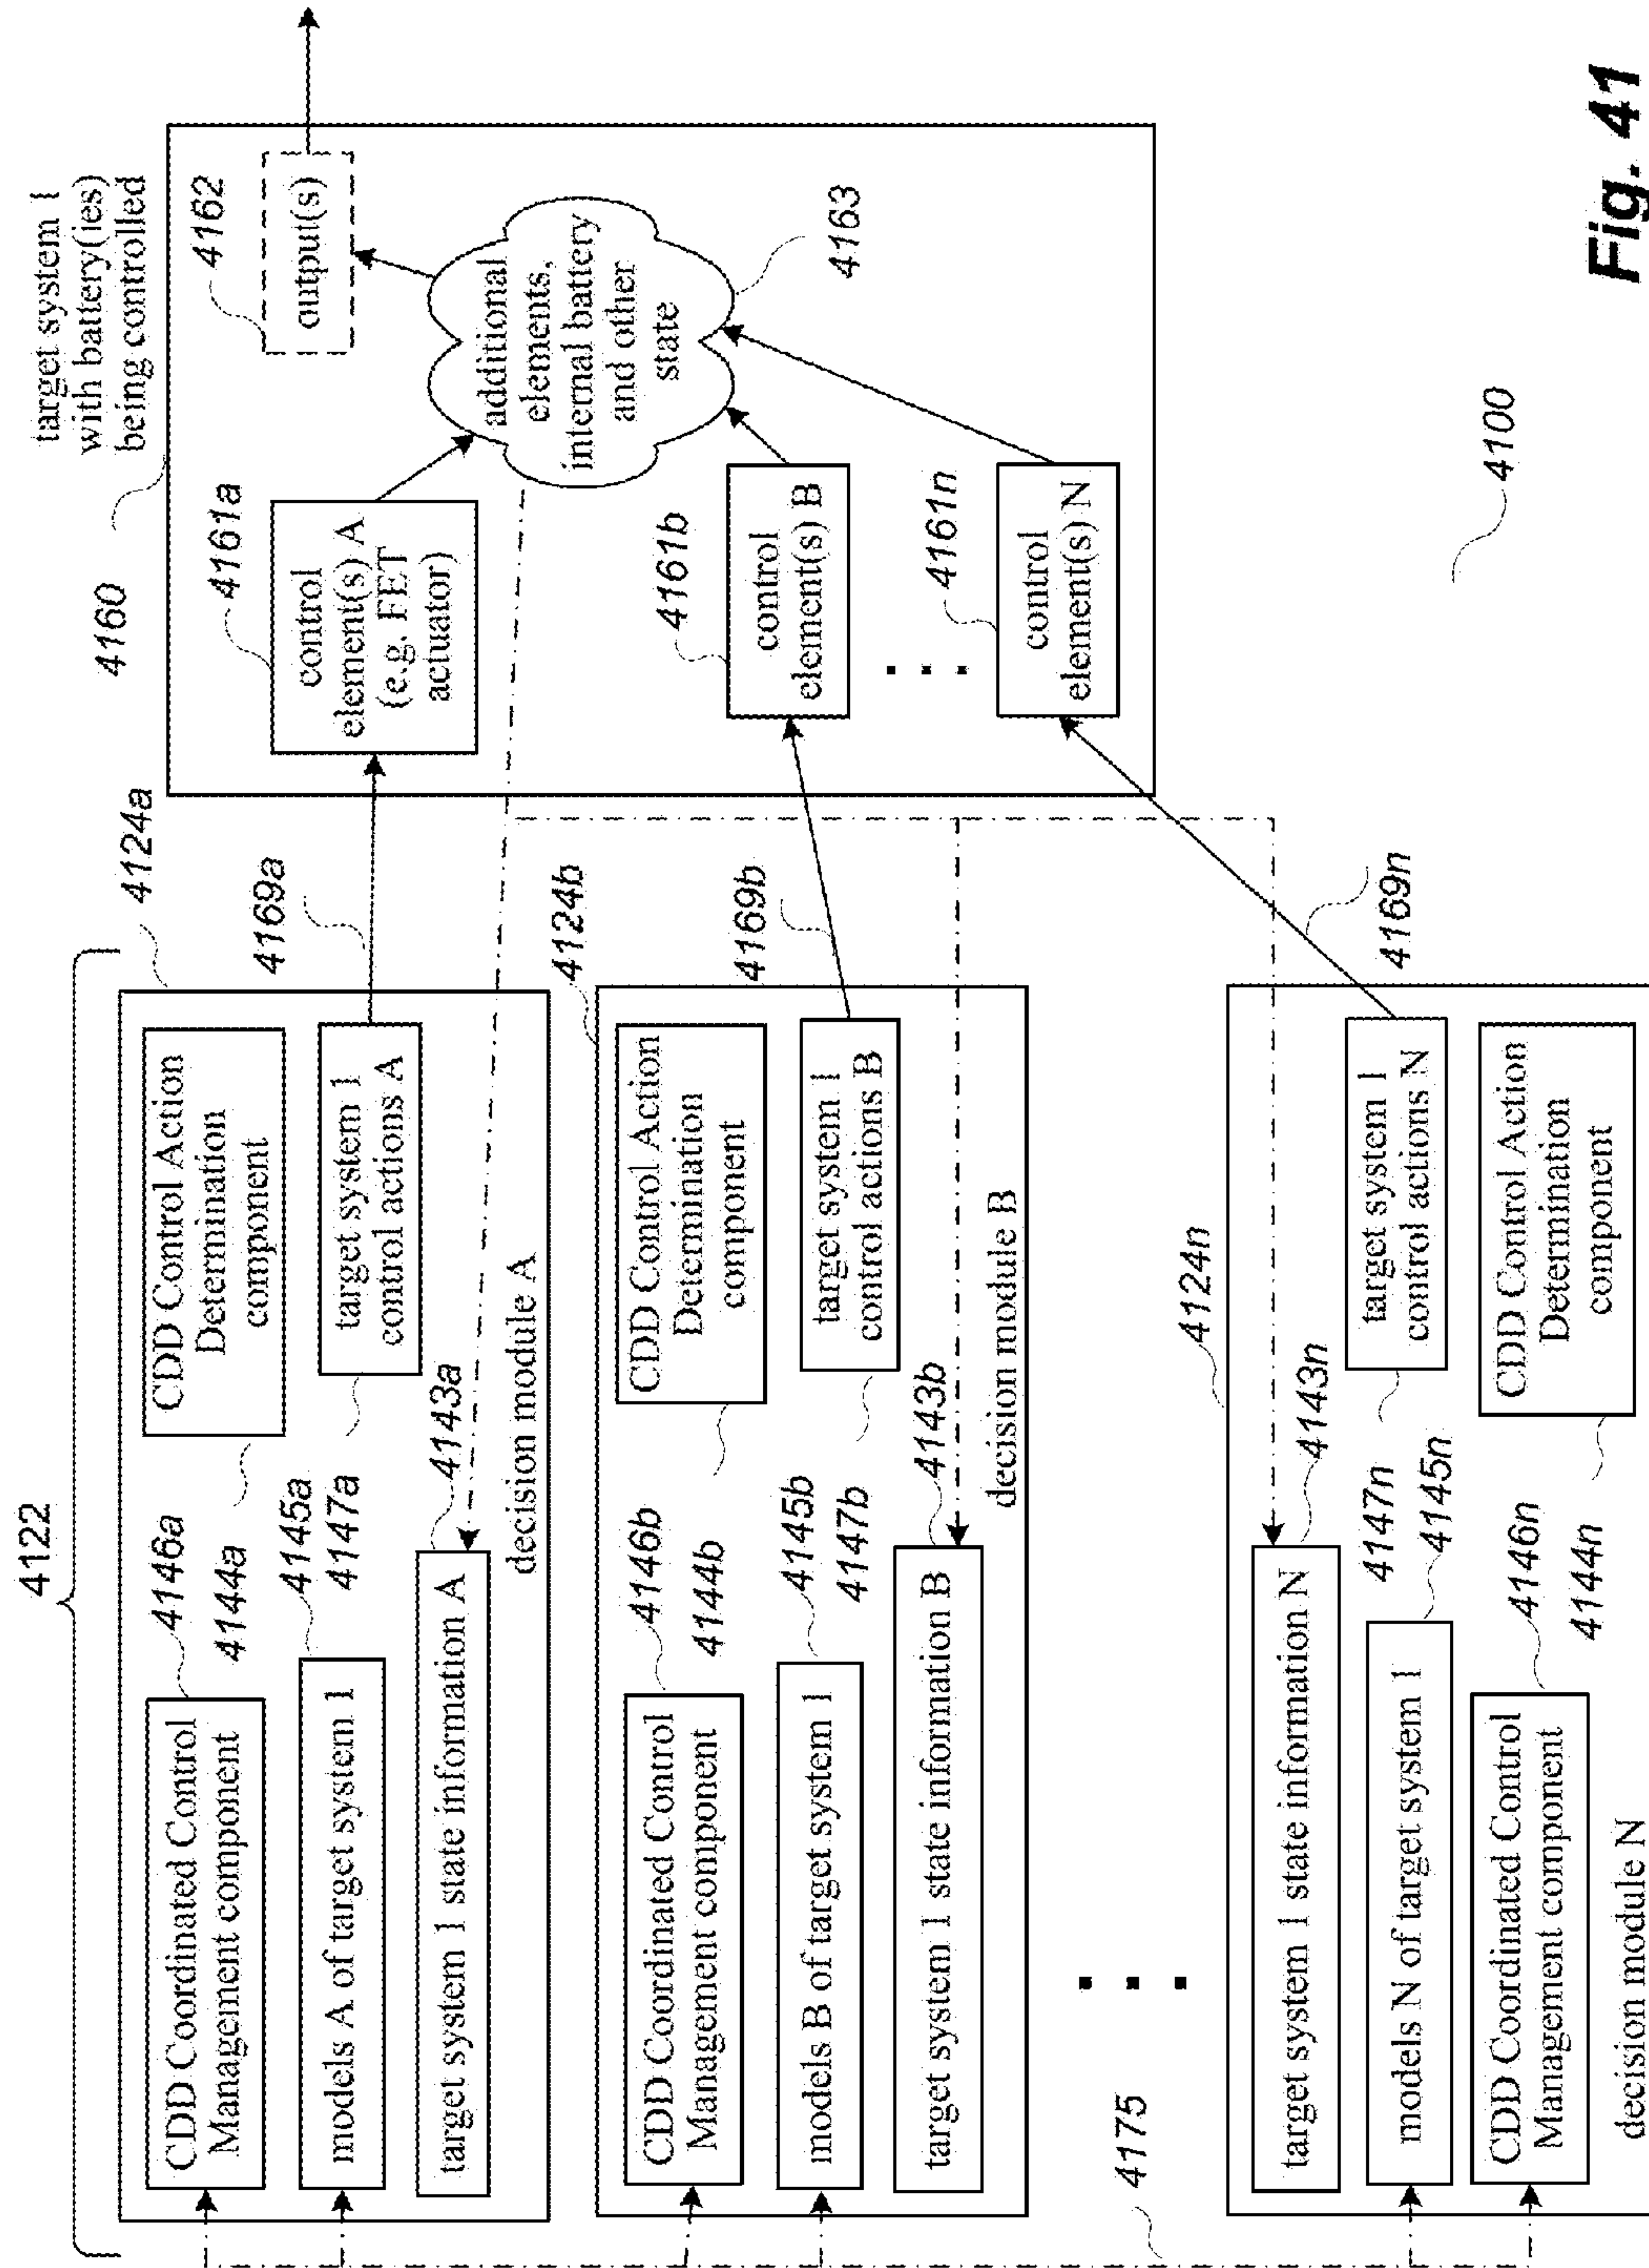
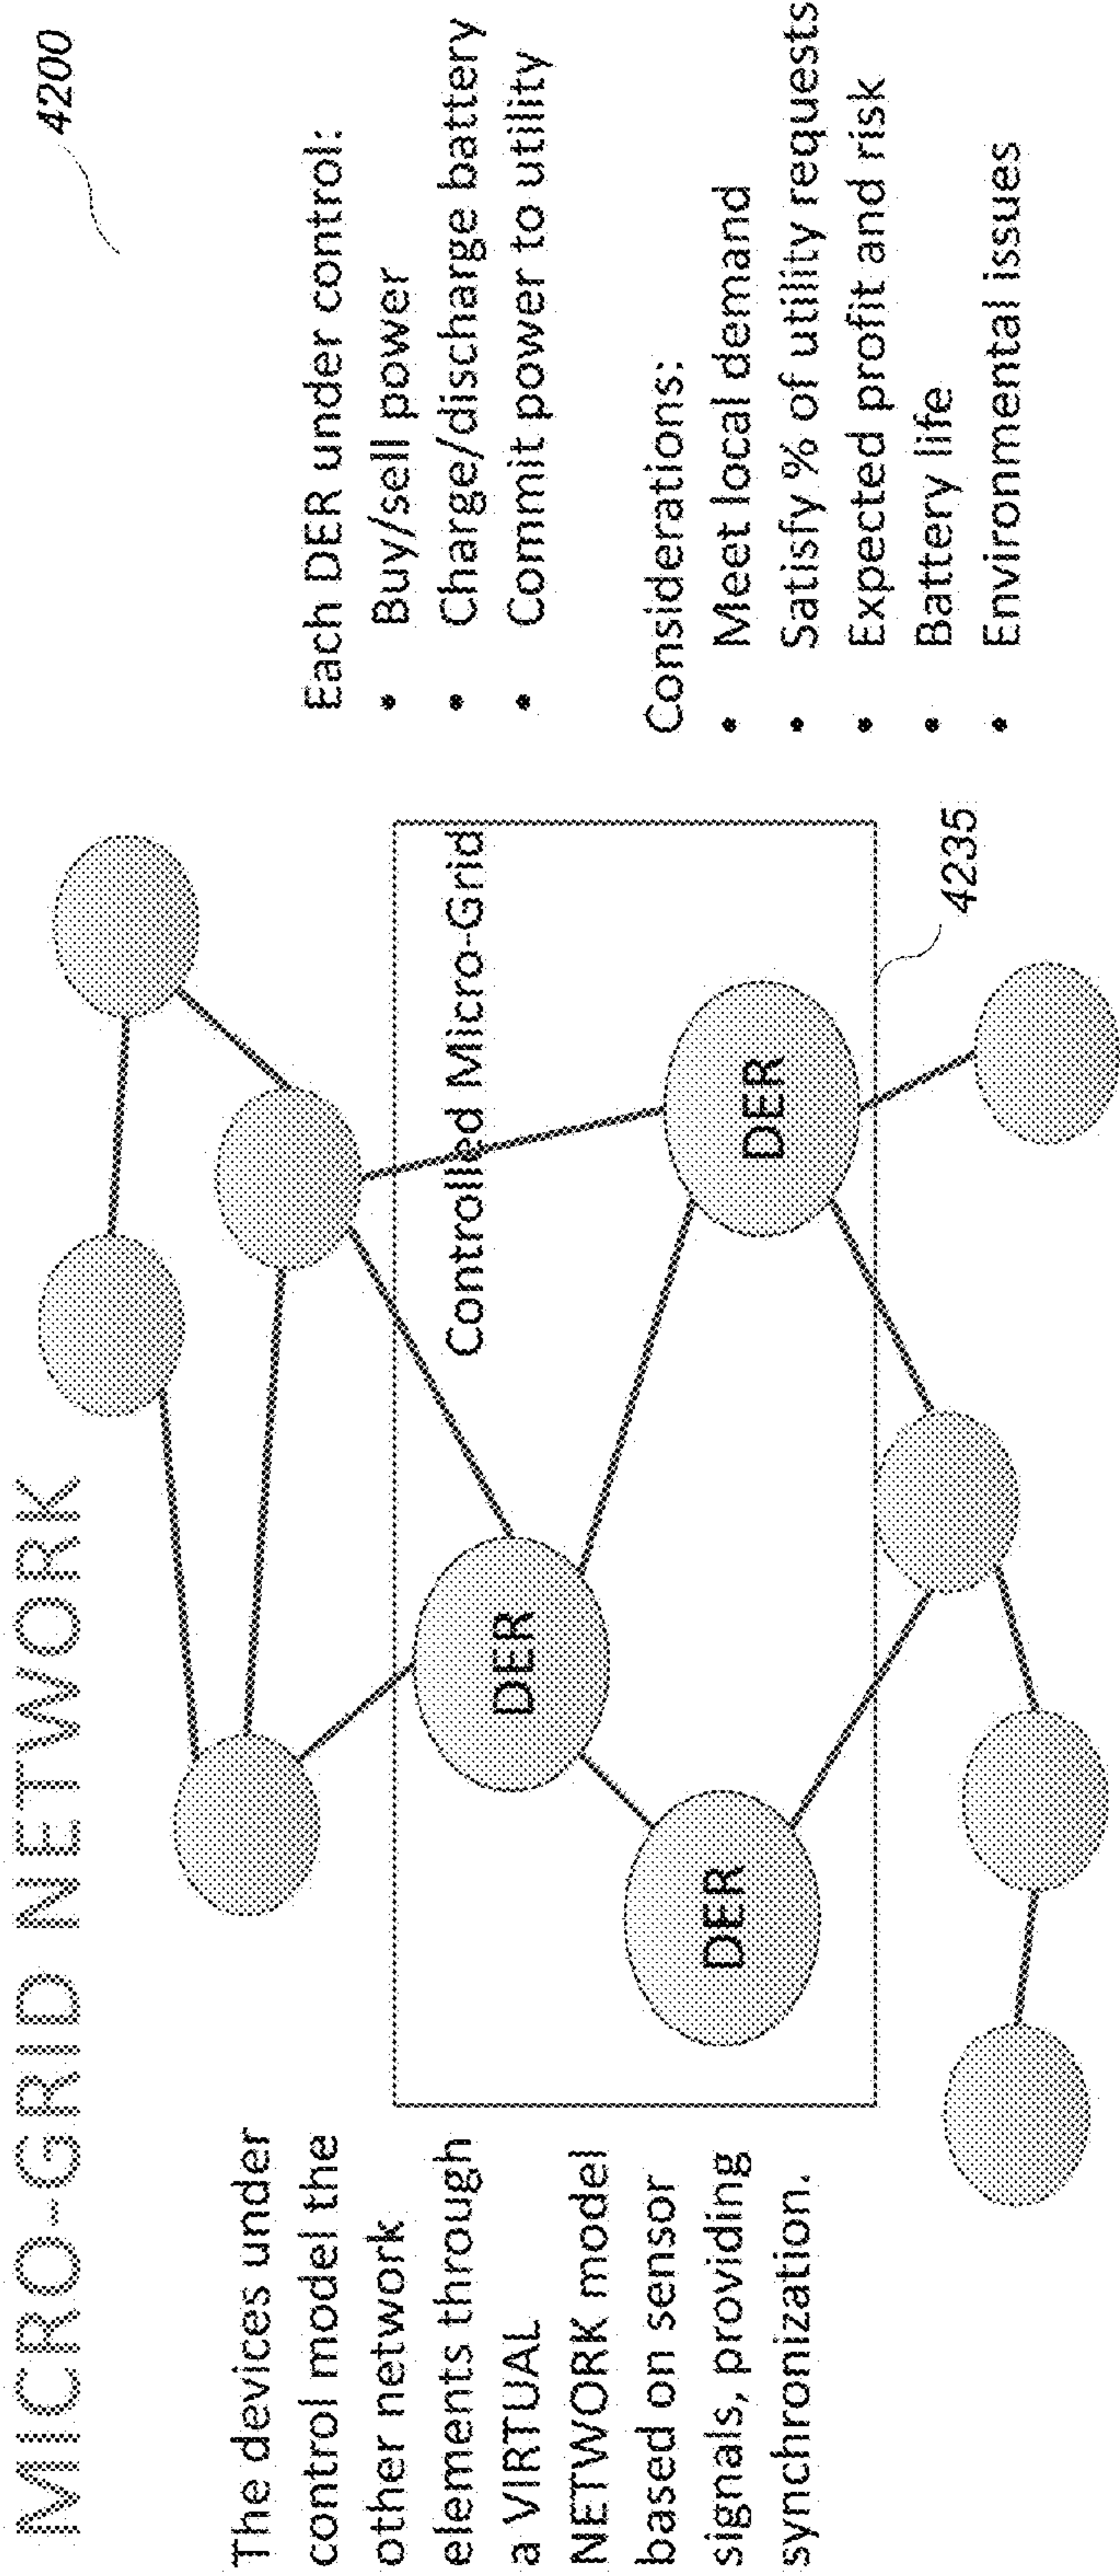


Fig. 41



- Each DER under control:
- Buy/sell power
  - Charge/discharge battery
  - Commit power to utility
- Considerations:
- Meet local demand
  - Satisfy % of utility requests
  - Expected profit and risk
  - Battery life
  - Environmental issues

**Fig. 42**

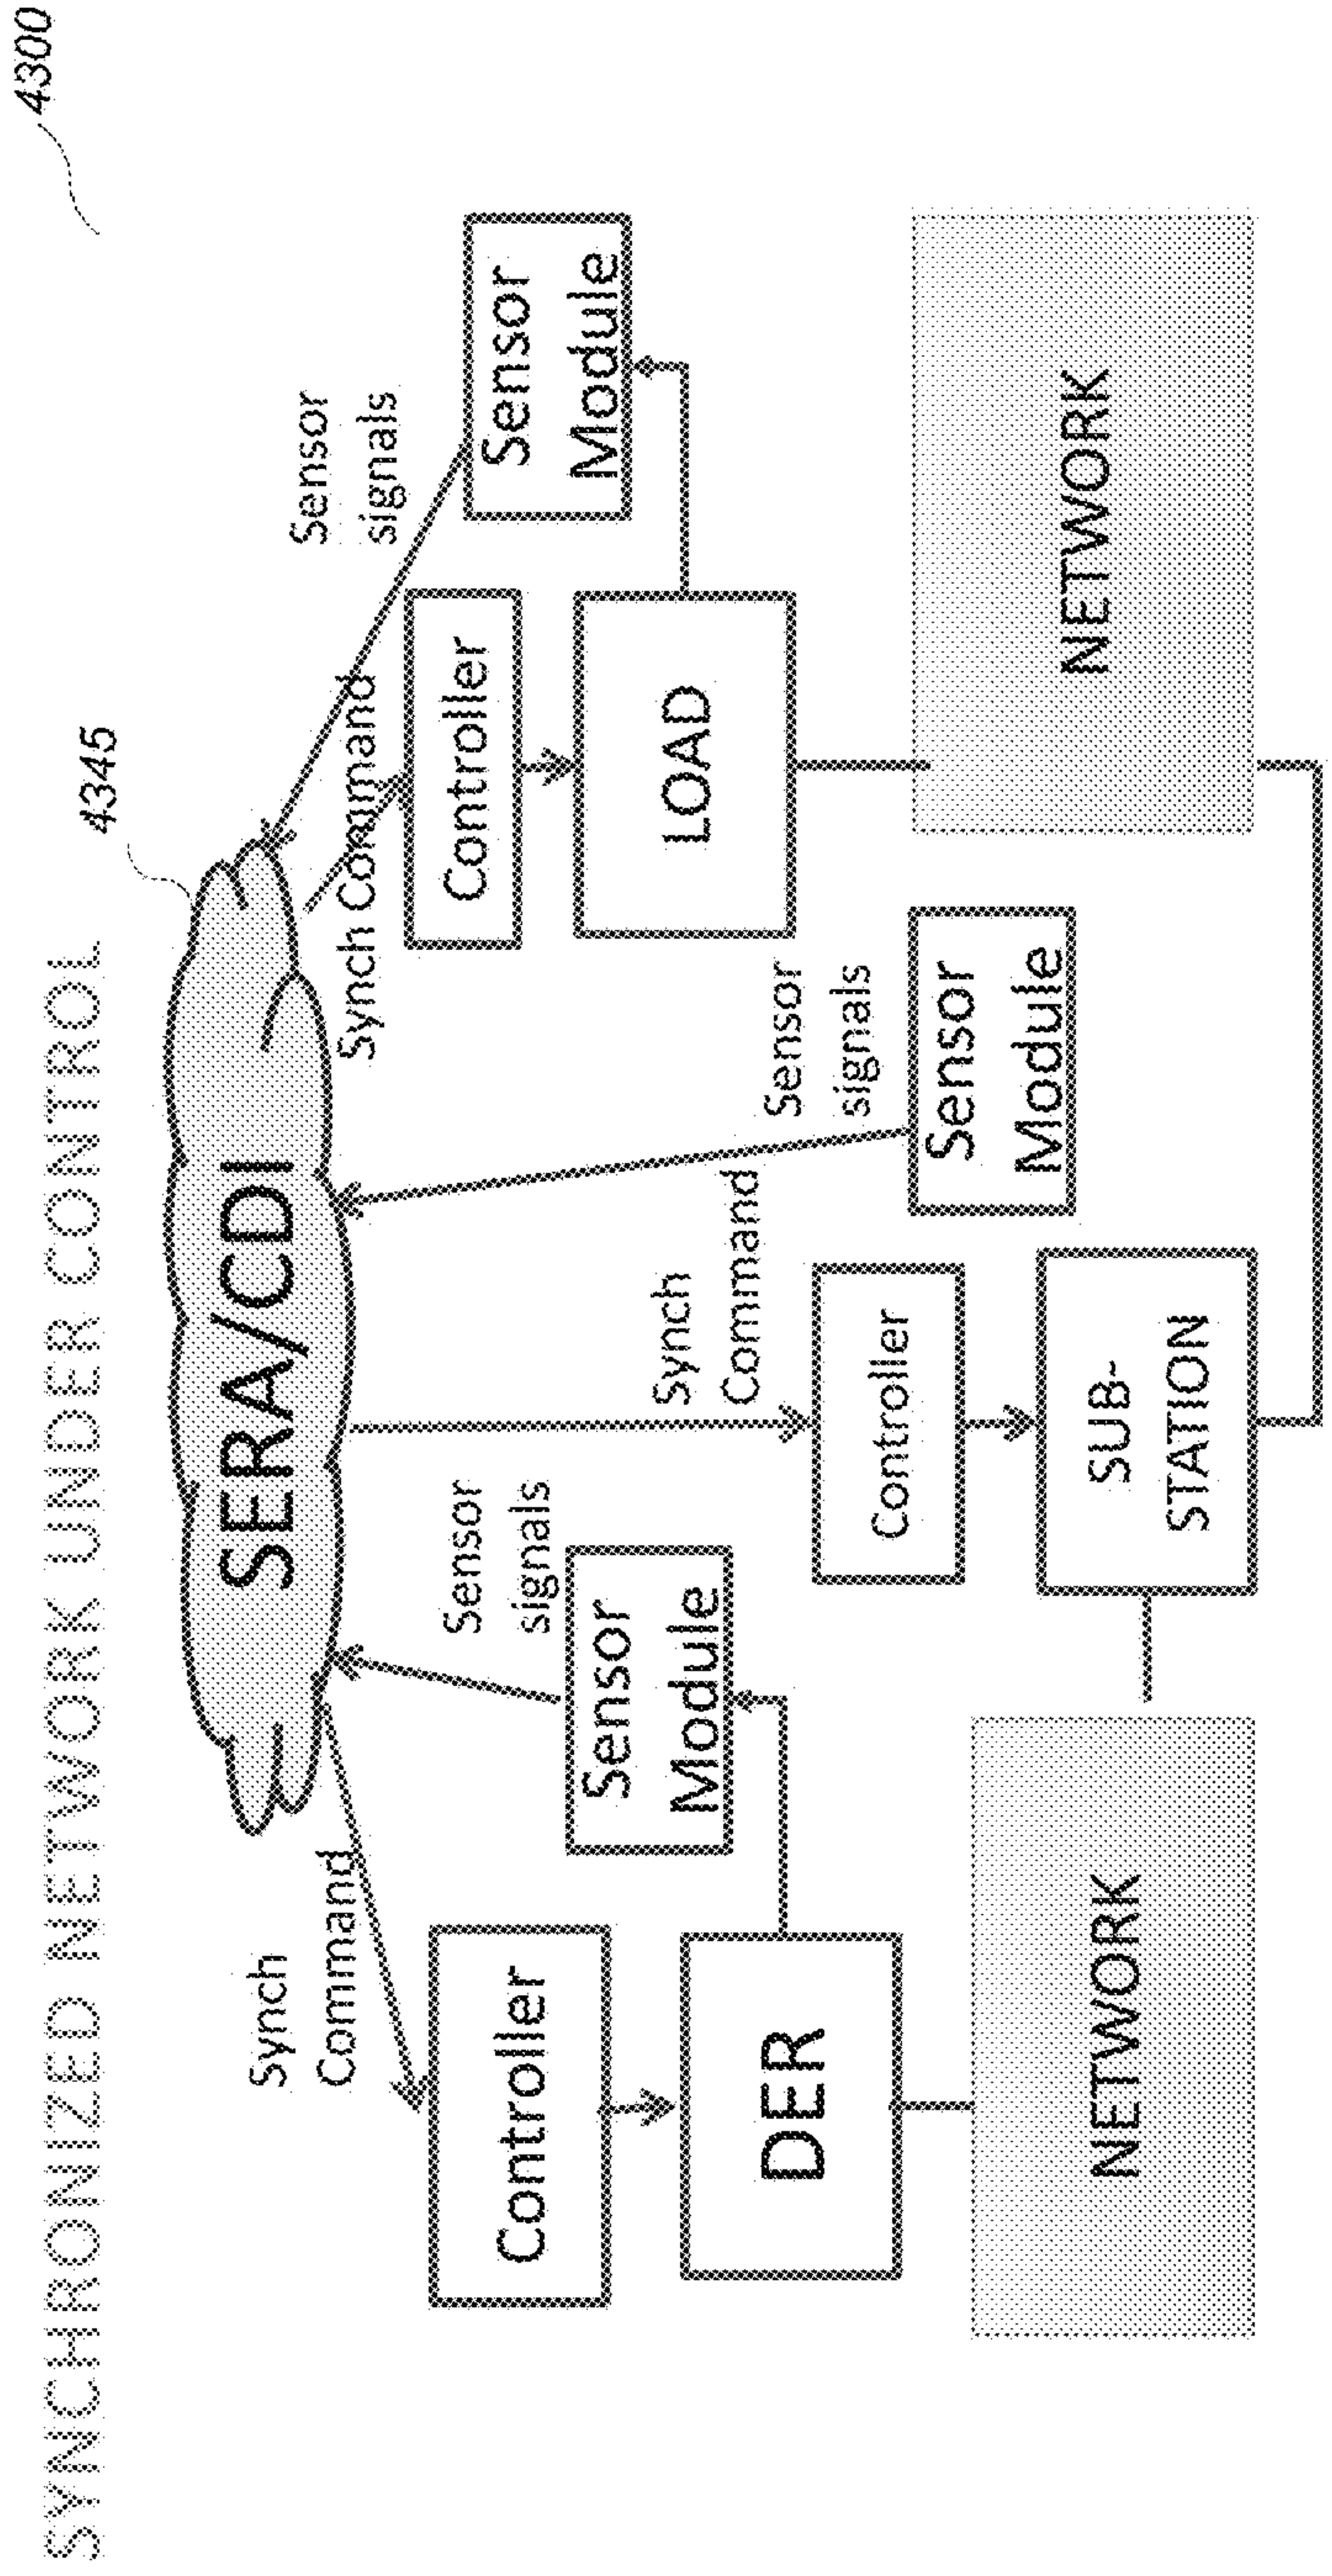


Fig. 43



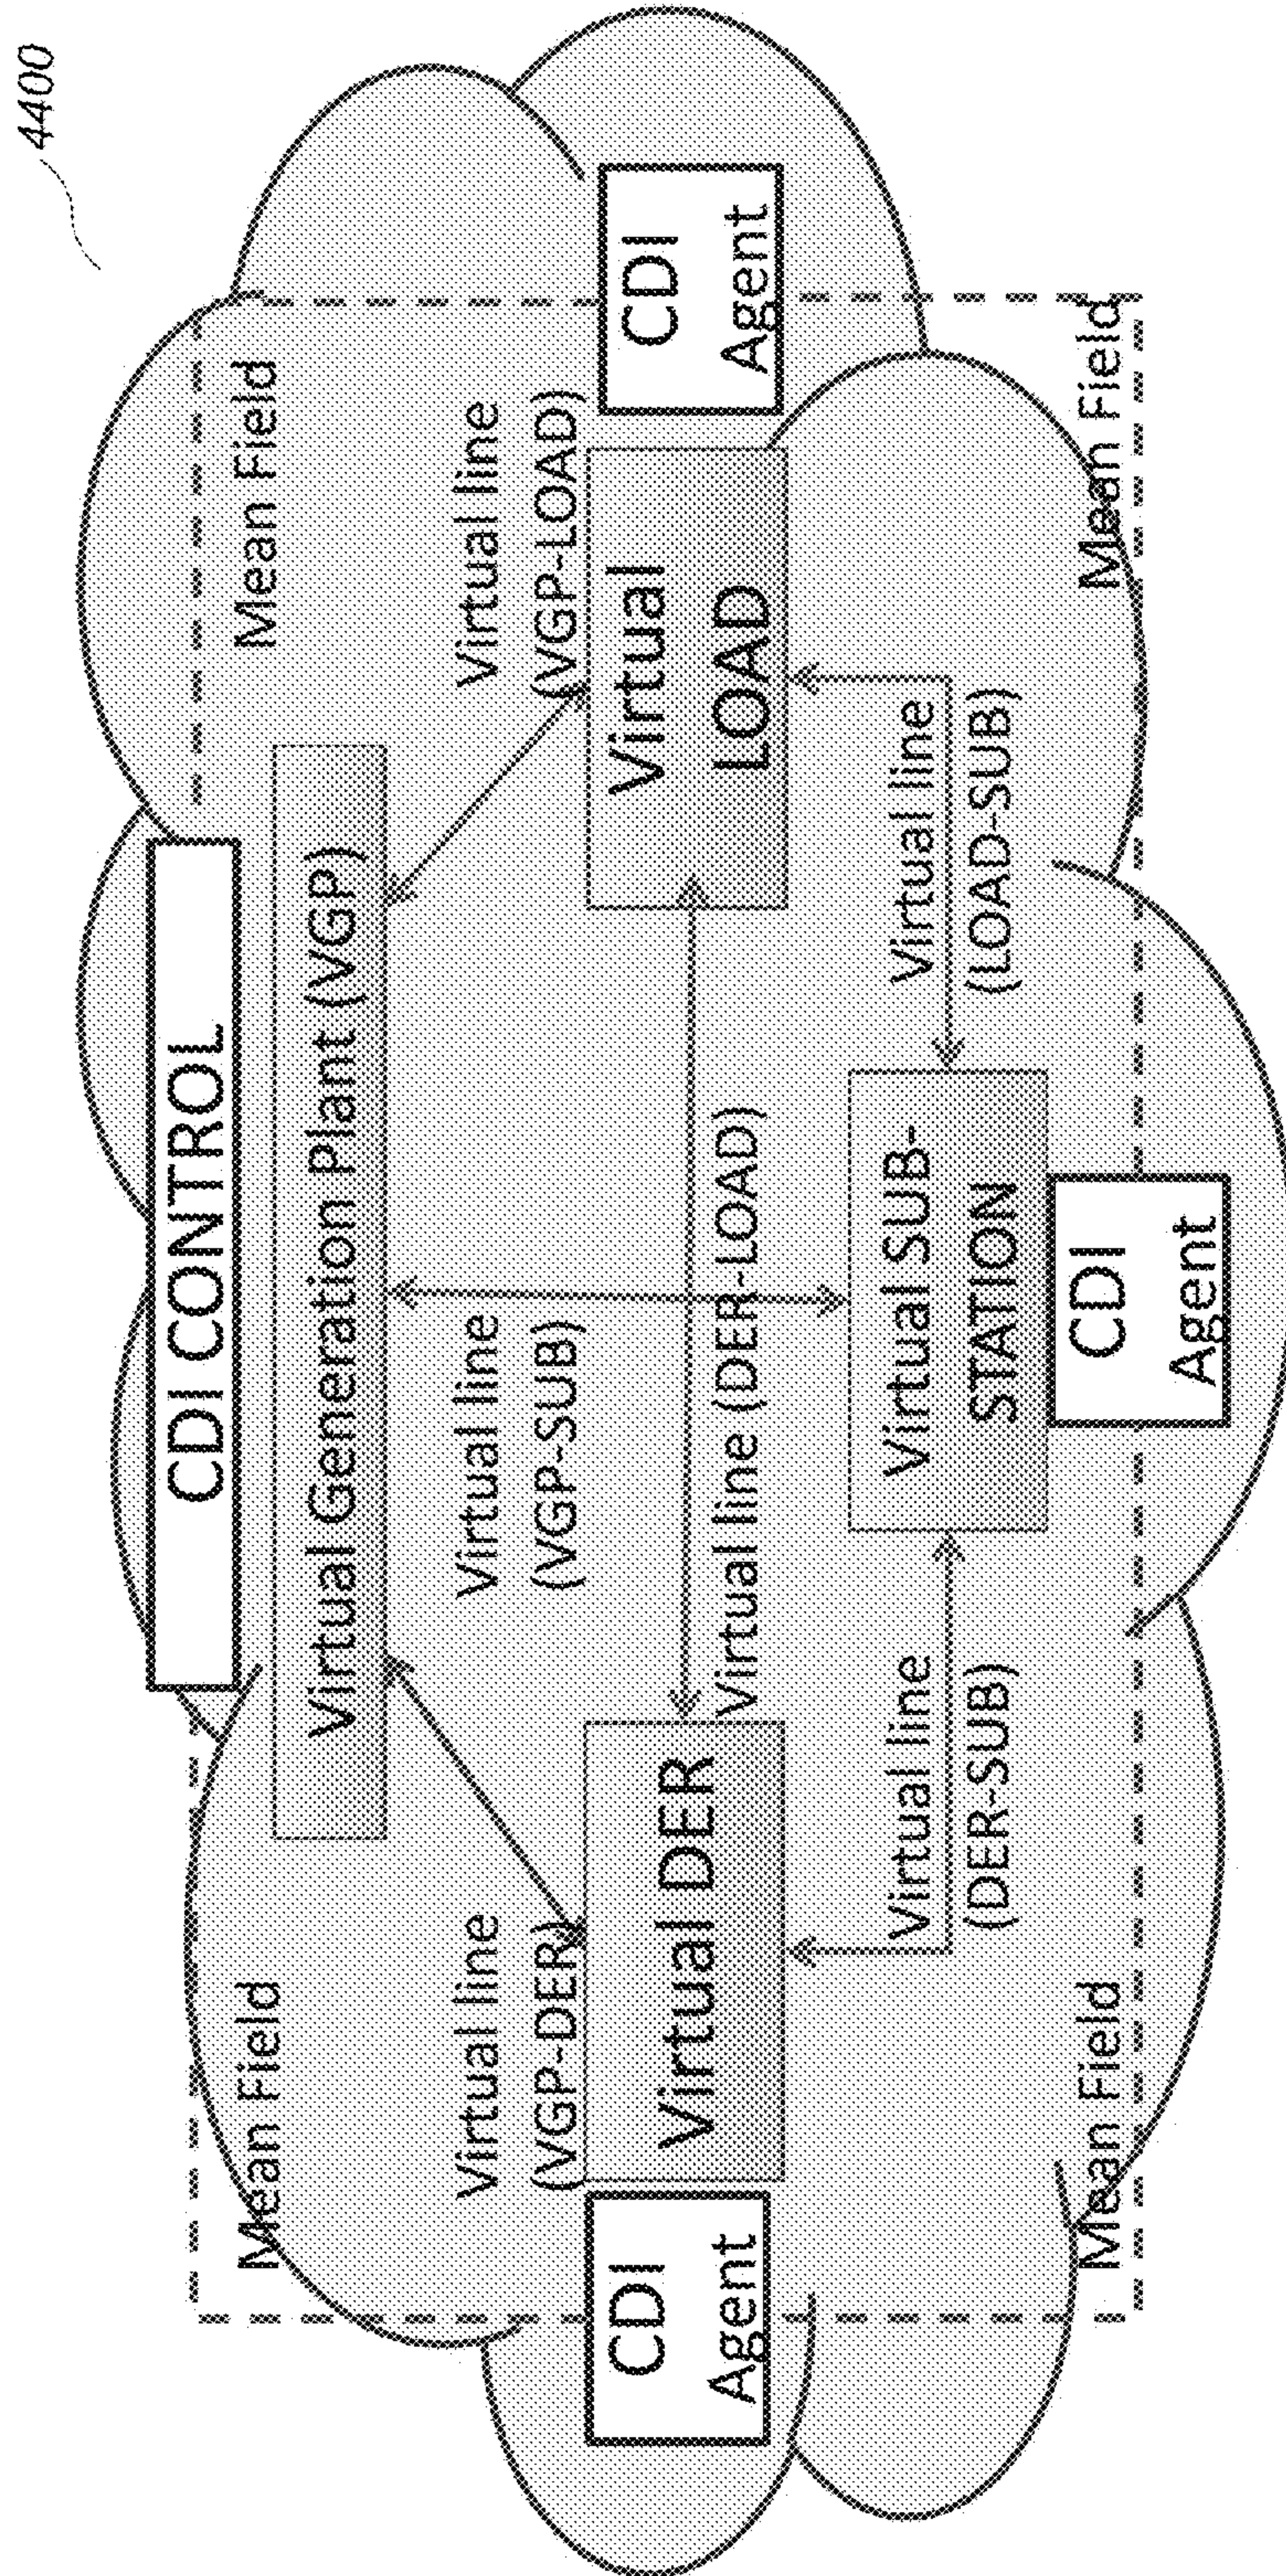


Fig. 44



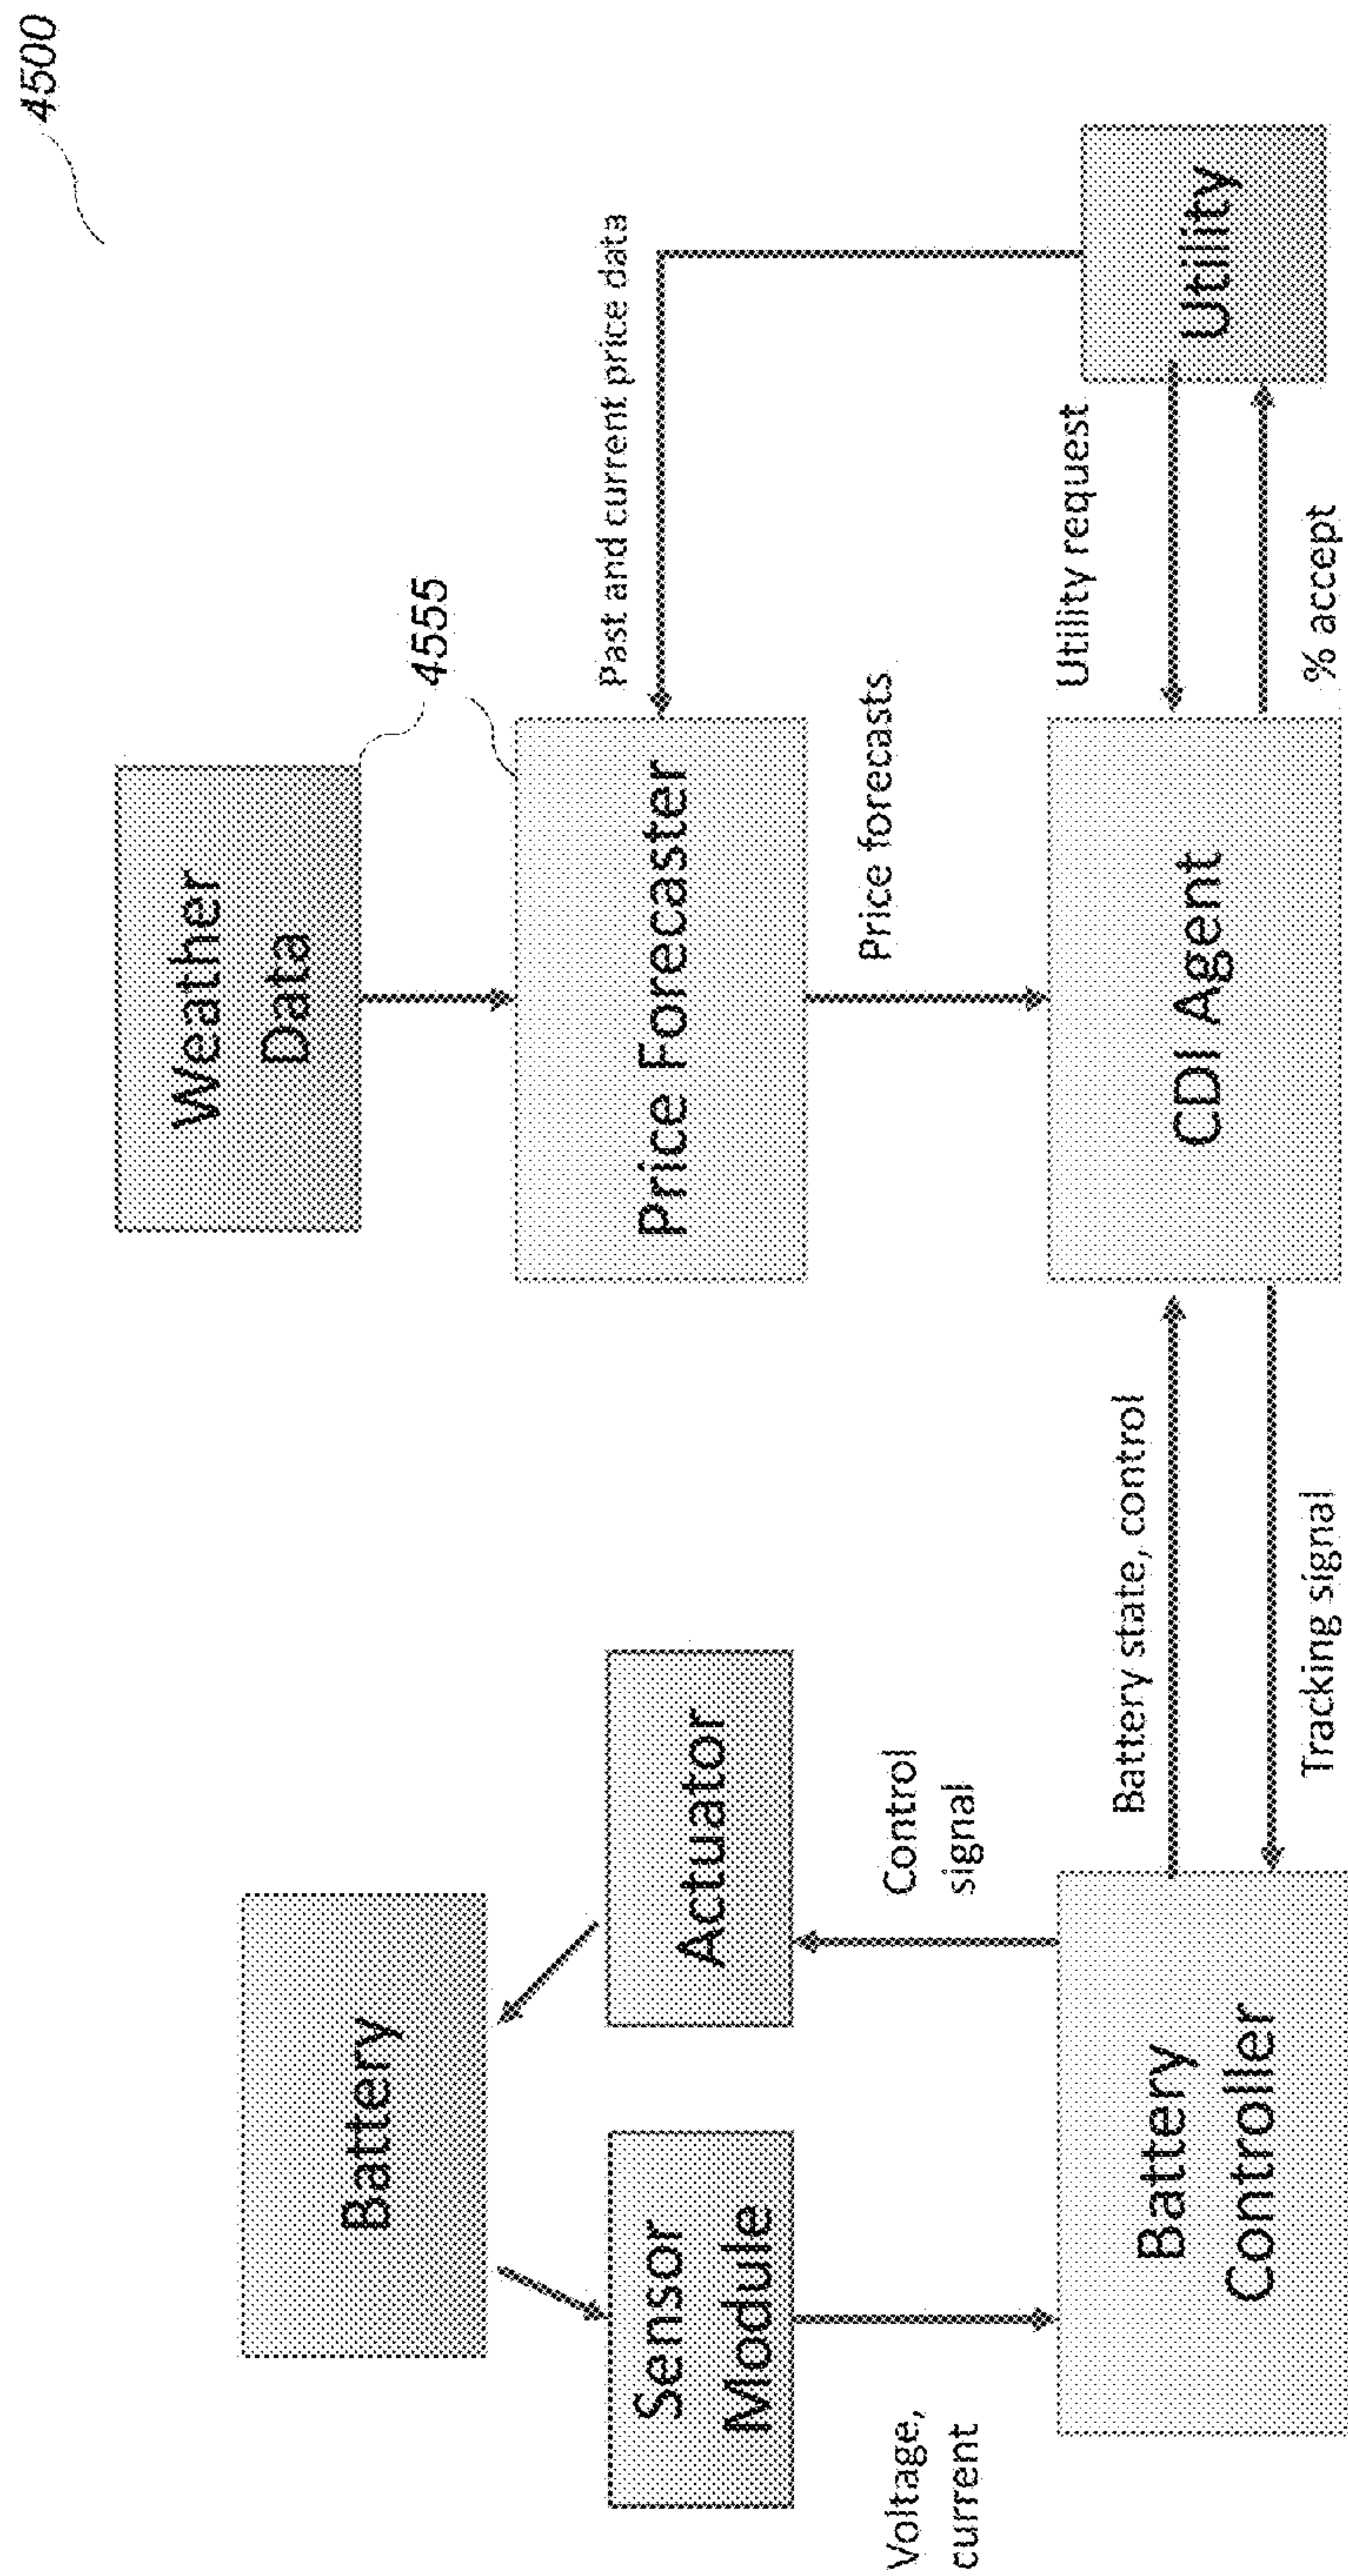


Fig. 45




Fig. 46



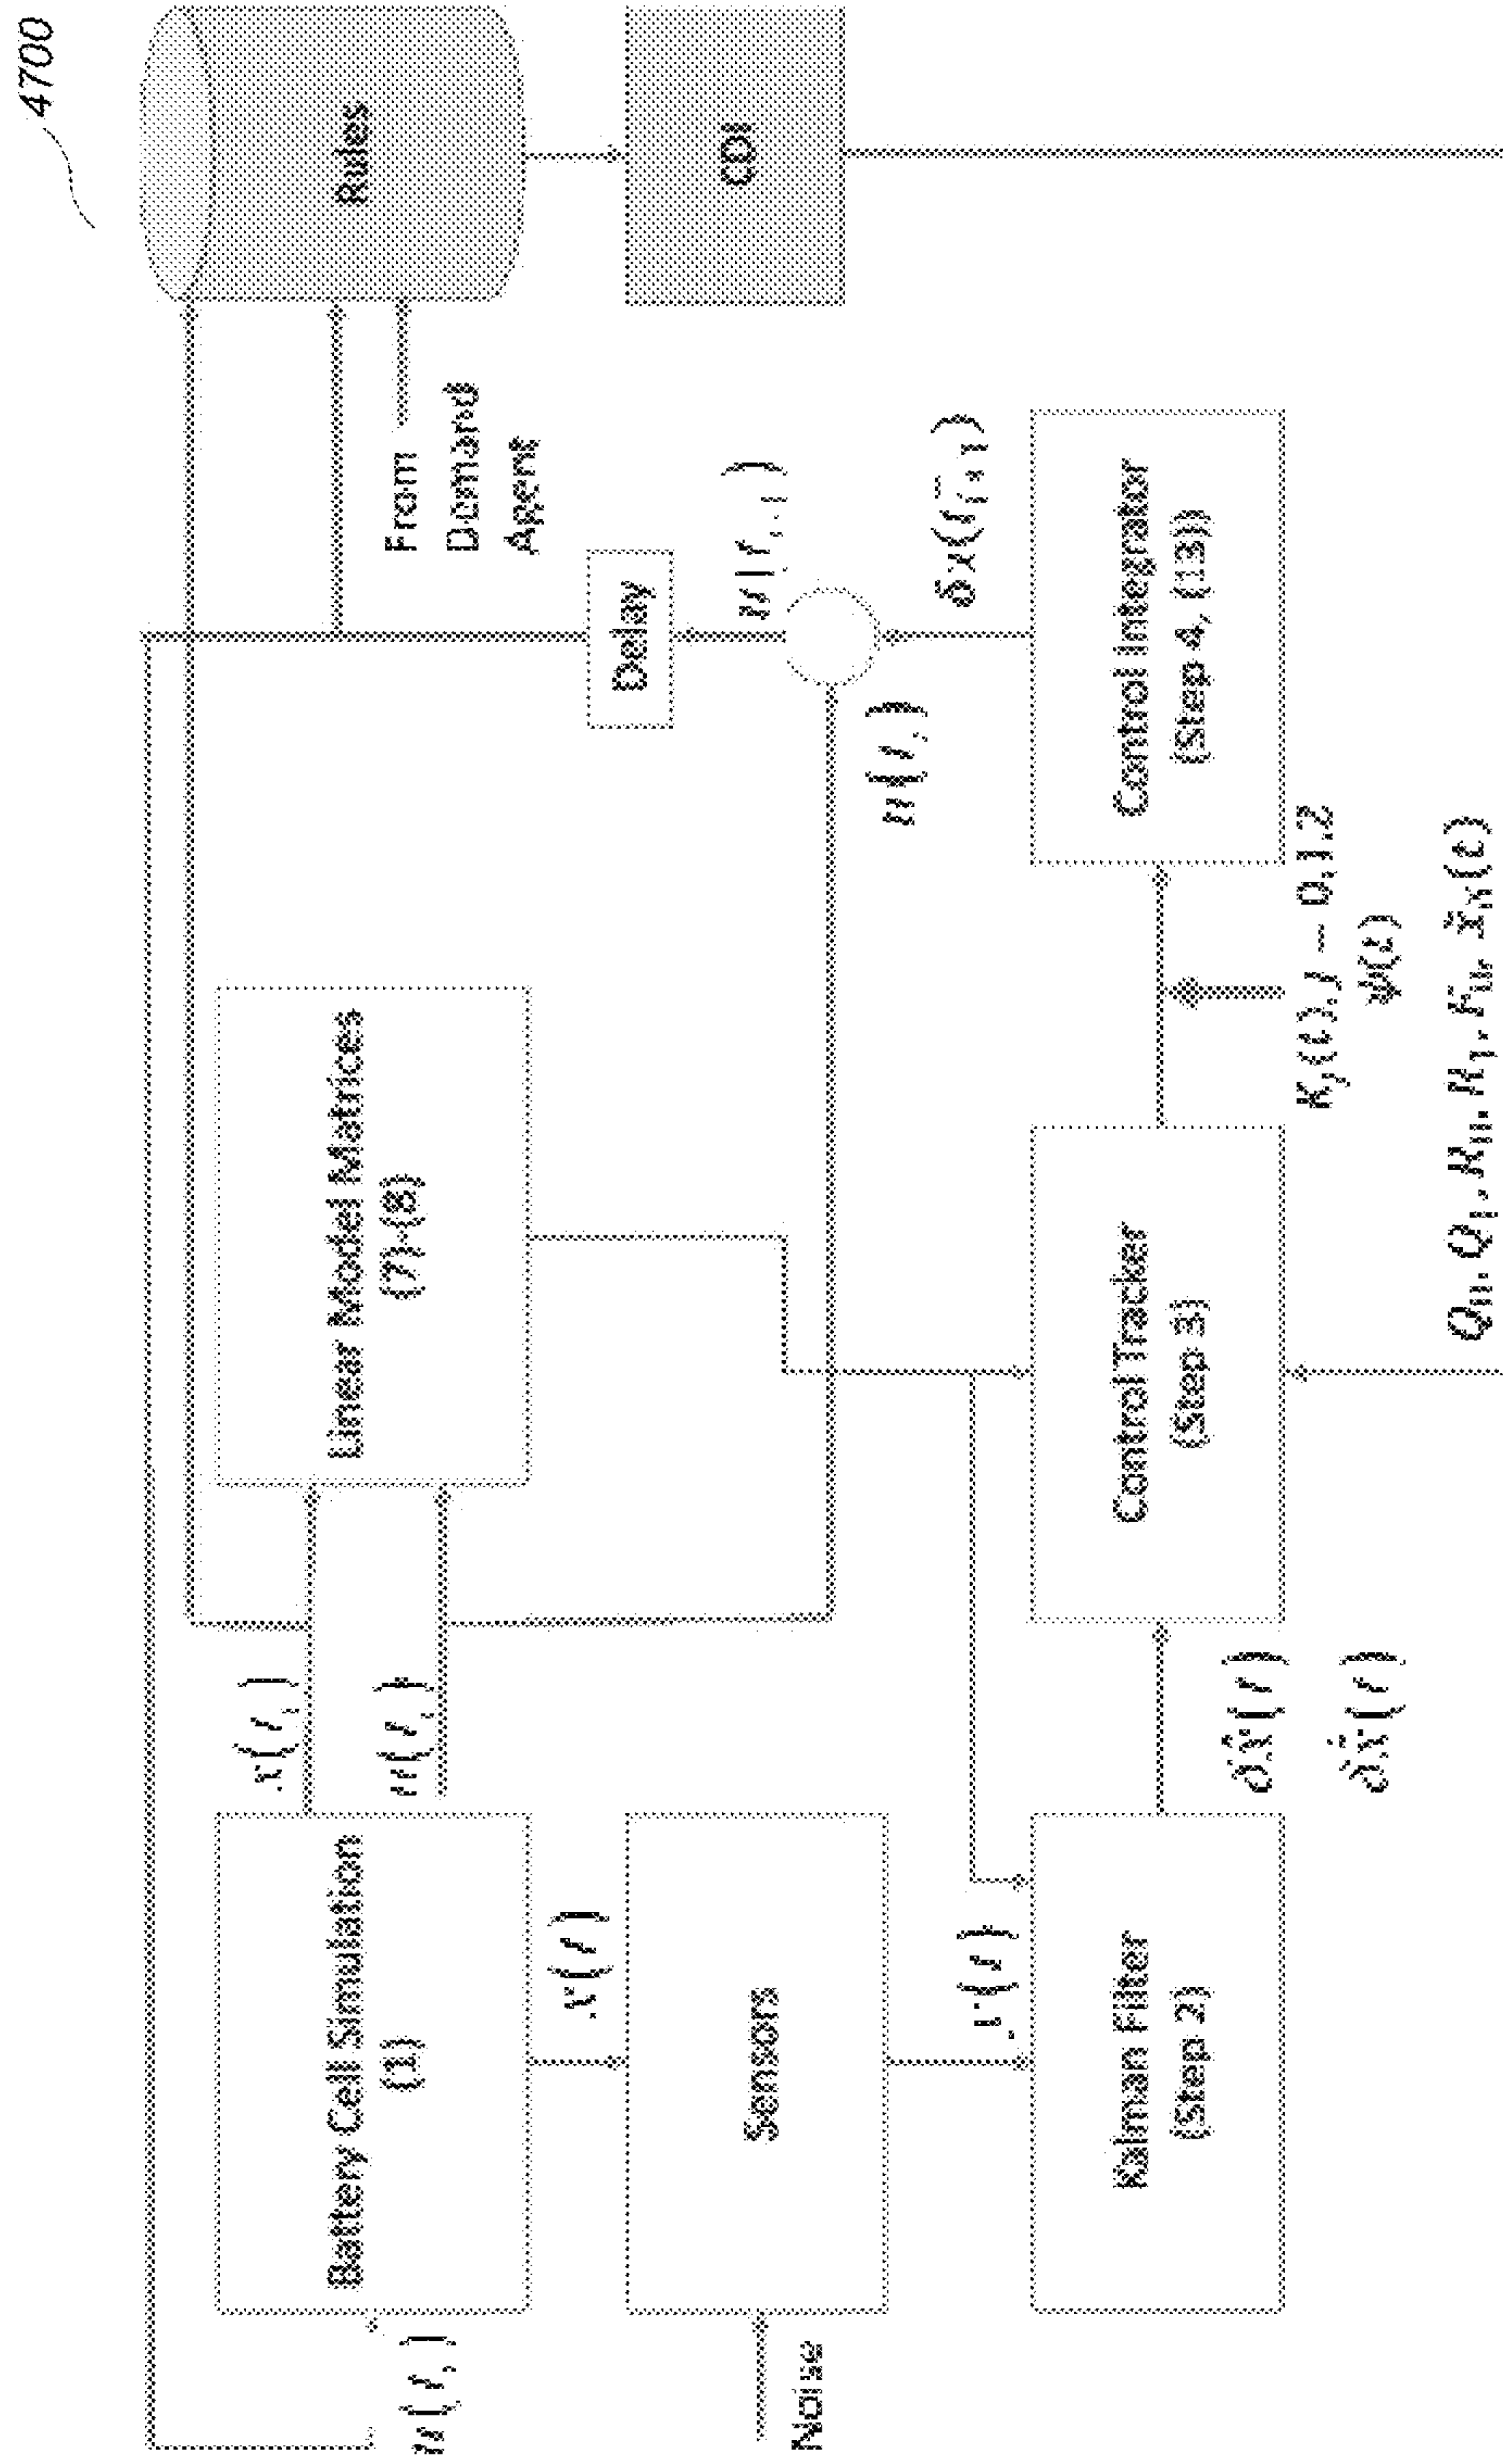


Fig. 47

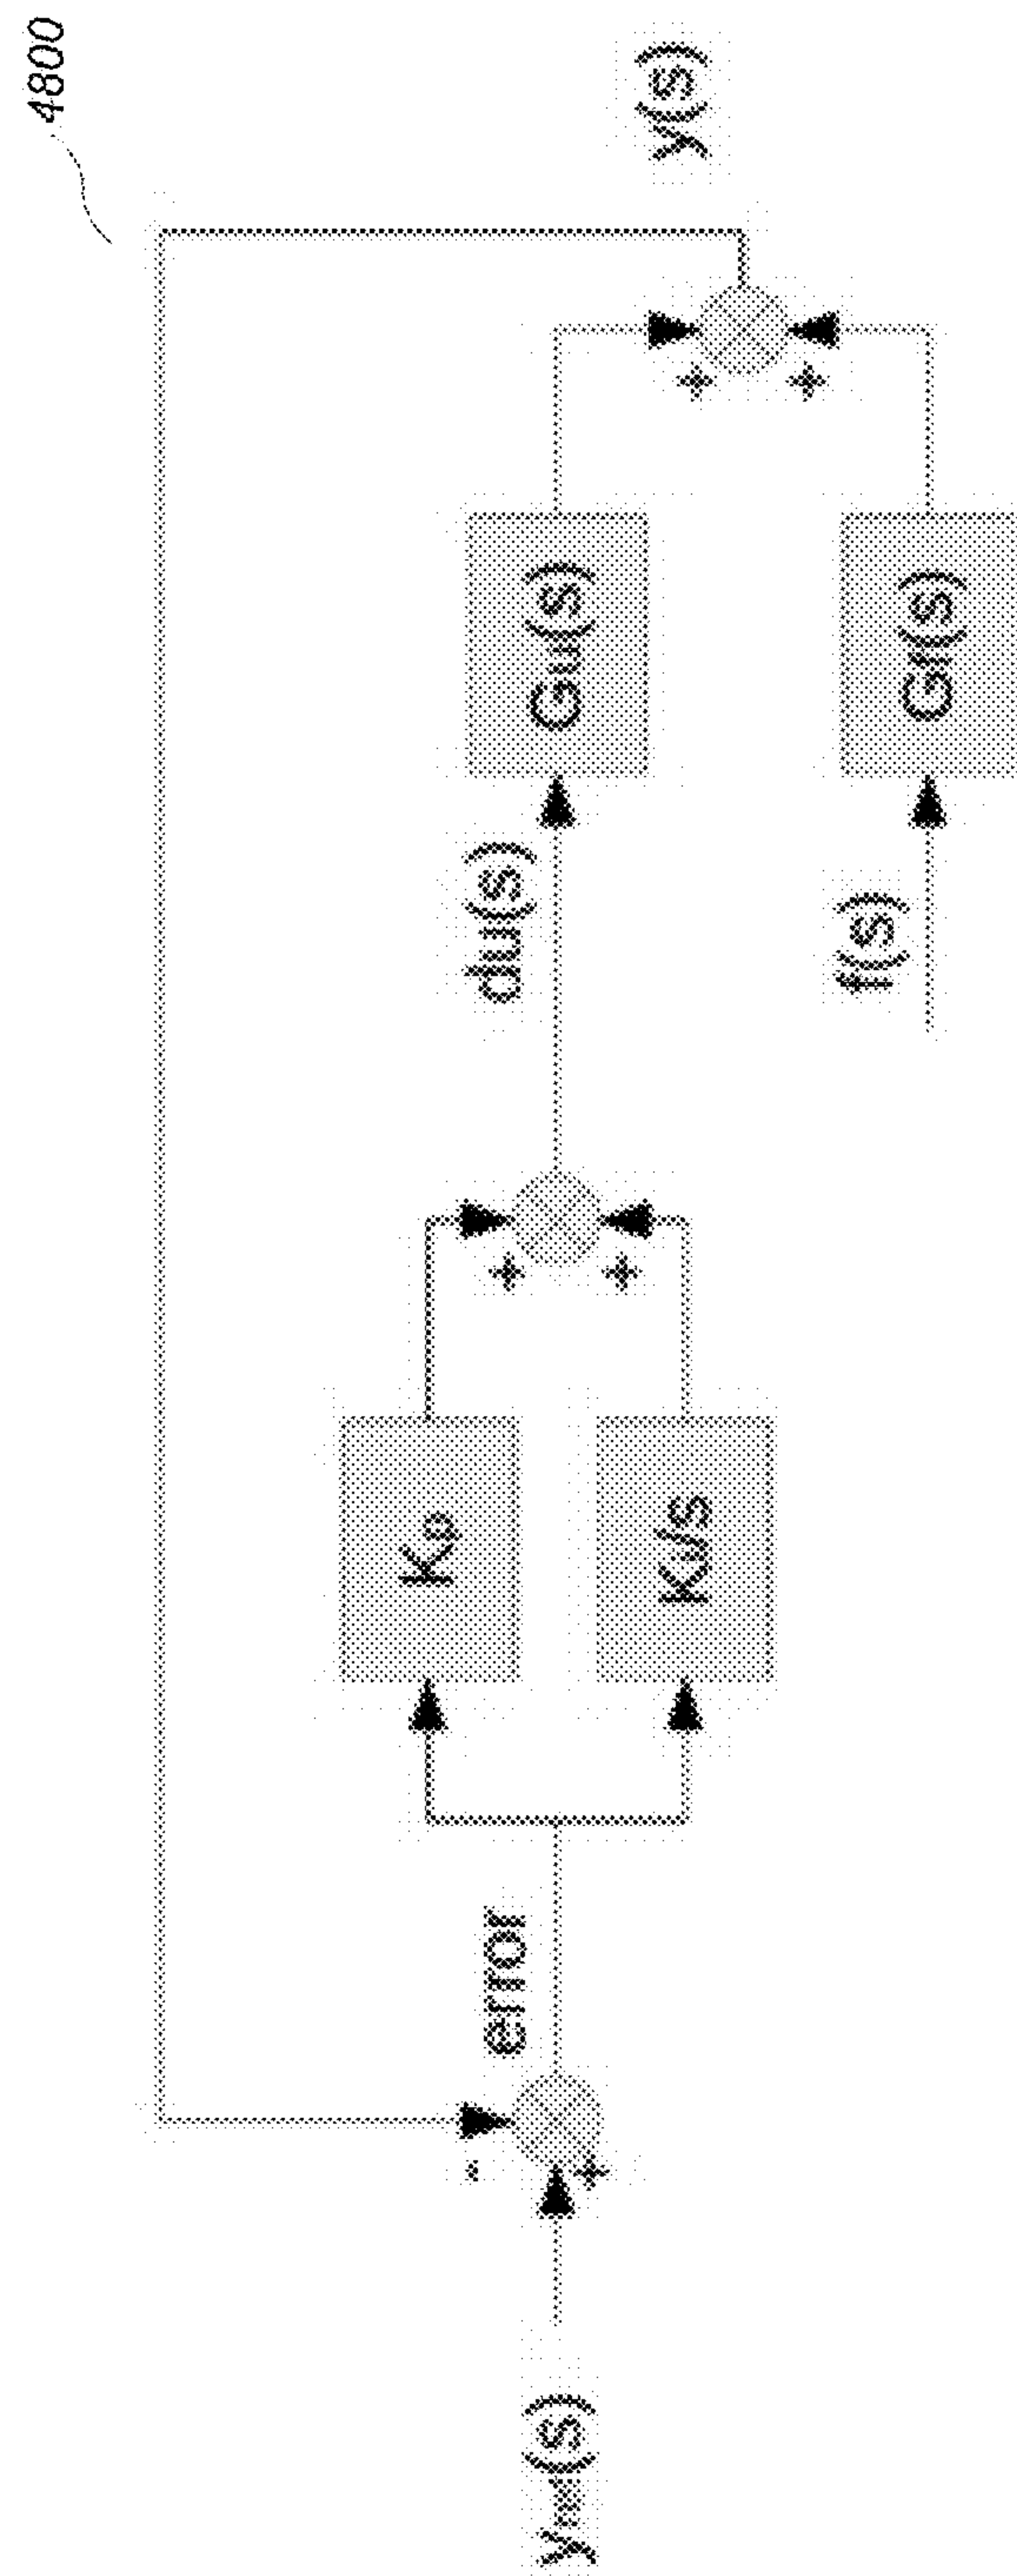


Fig. 48

4900




Fig. 49



5000

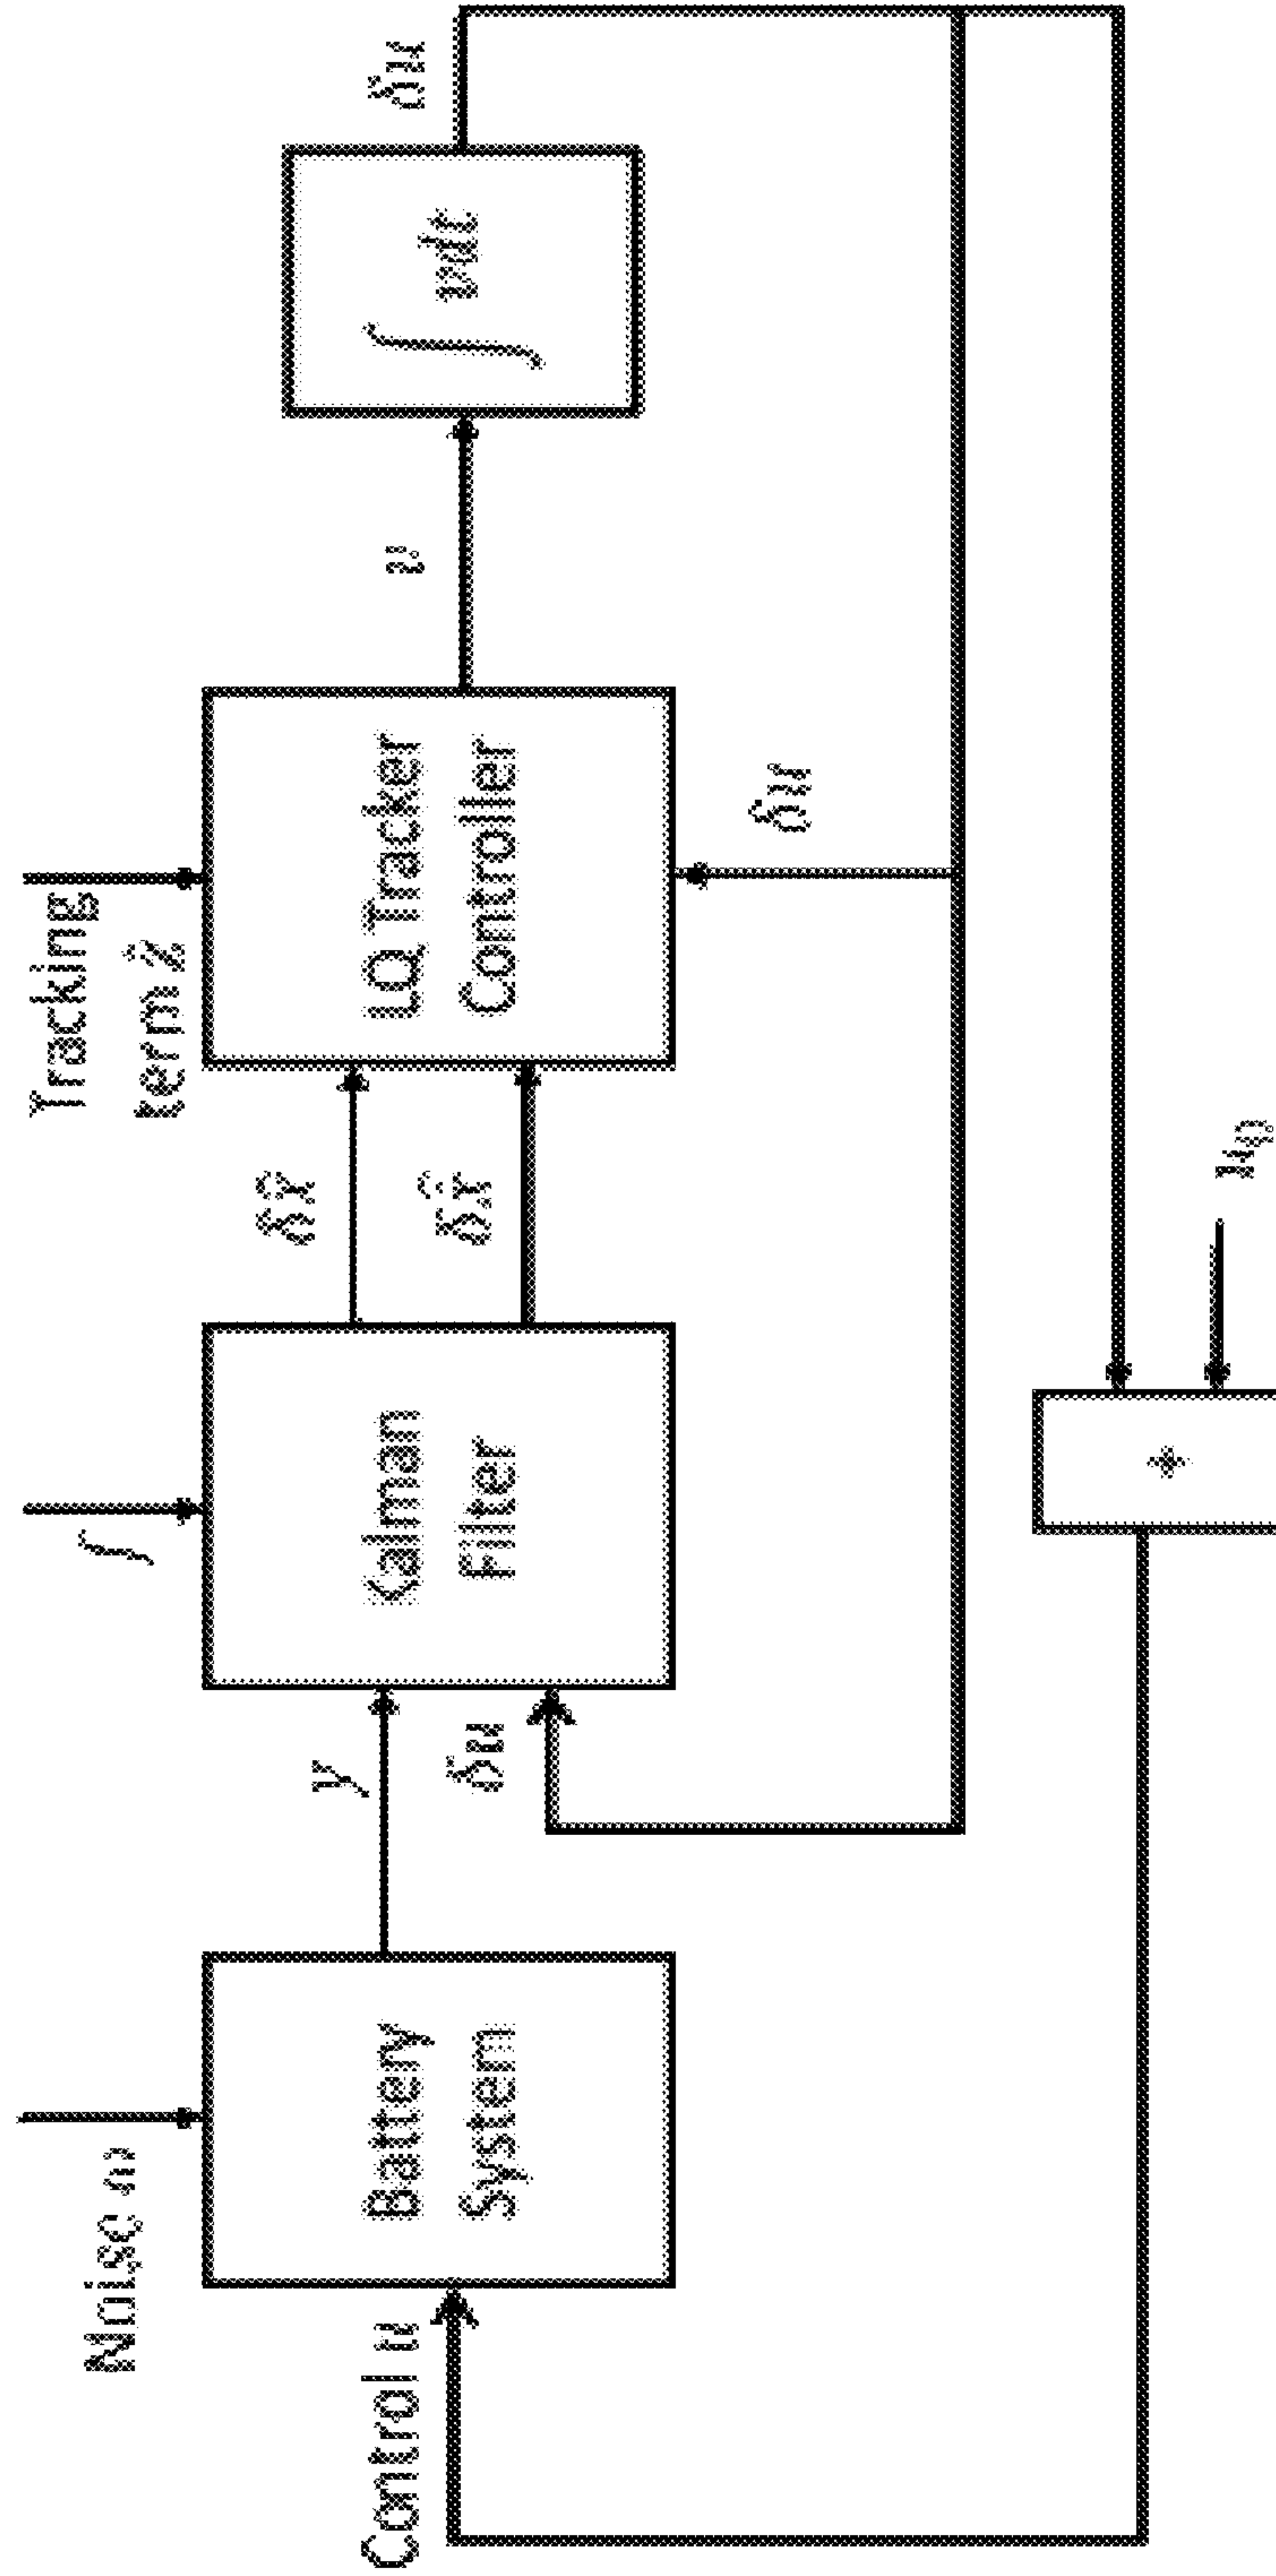


Fig. 50

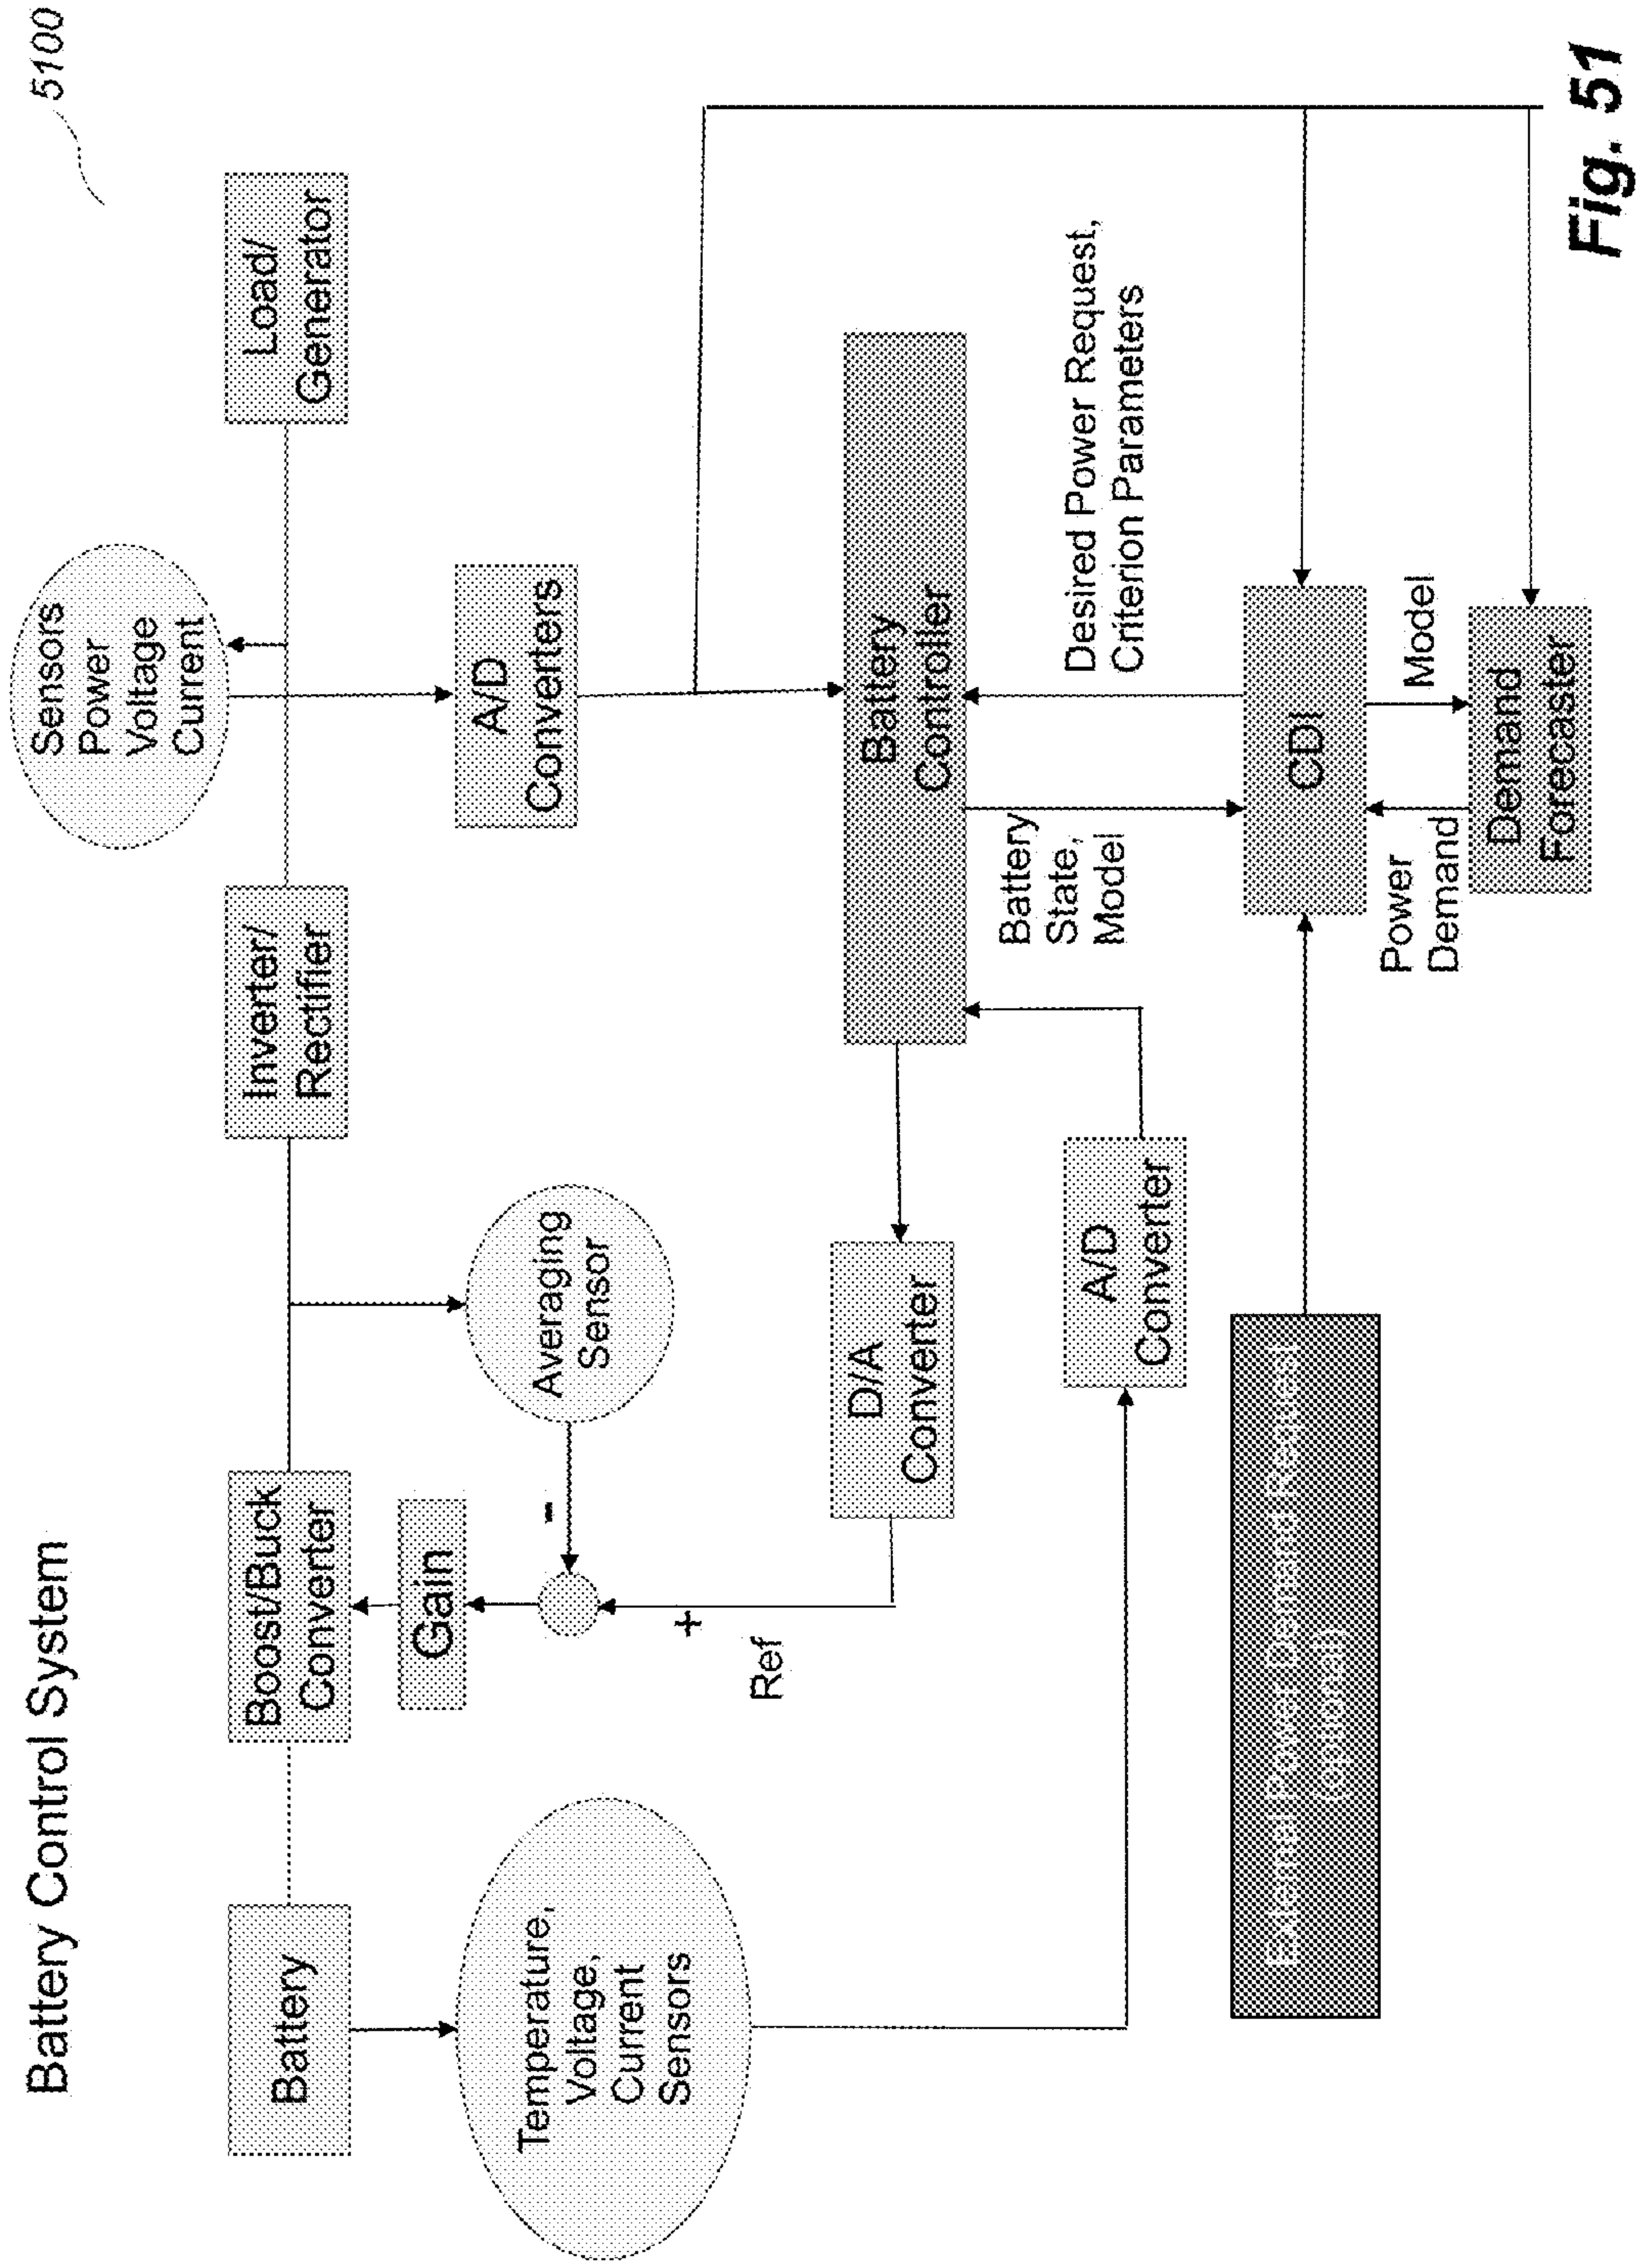


Fig. 51



## USING BATTERY DC CHARACTERISTICS TO CONTROL POWER OUTPUT

### CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of U.S. Provisional Patent Application No. 62/304,034, filed Mar. 4, 2016 and entitled “Using Battery DC Characteristics To Control Power Quality,” which is hereby incorporated by reference in its entirety.

### BACKGROUND

[0002] Various attempts have been made to implement automated control systems for various types of physical systems having inputs or other control elements that the control system can manipulate to attempt to provide desired output or other behavior of the physical systems being controlled. Such automated control systems have used various types of architectures and underlying computing technologies to attempt to implement such functionality, including to attempt to deal with issues related to uncertainty in the state of the physical system being controlled, the need to make control decisions in very short amounts of time and with only partial information, etc. One example of such an automated control system includes a system for controlling operations of a battery that is discharging electrical power to support a load and/or is charging using electrical power from a source, with uncertainty about an internal temperature and/or chemical state of the battery, and potentially with ongoing changes in load, source and/or battery internal state.

[0003] However, various difficulties exist with existing automated control systems and their underlying architectures and computing technologies, including with respect to managing large numbers of constraints (sometimes conflicting), operating in a coordinated manner with other systems, etc. Particular difficulties can arise when attempting to control one or more batteries in situations in which multiple conflicting constraints and/or goals exist.

### BRIEF DESCRIPTION OF THE DRAWINGS

[0004] The patent or application file contains at least one drawing executed in color. Copies of this patent or patent application publication with color drawing(s) will be provided by the Office upon request and payment of the necessary fee.

[0005] FIG. 1 is a network diagram illustrating an example environment in which a system for performing cooperative distributed control of target systems may be configured and initiated.

[0006] FIG. 2 is a network diagram illustrating an example environment in which a system for performing cooperative distributed control of target systems may be implemented.

[0007] FIG. 3 is a block diagram illustrating example computing systems suitable for executing an embodiment of a system for performing cooperative distributed control of target systems in configured manners.

[0008] FIG. 4 illustrates a flow diagram of an example embodiment of a Collaborative Distributed Decision (CDD) System routine.

[0009] FIGS. 5A-5B illustrate a flow diagram of an example embodiment of a CDD Decision Module Construction routine.

[0010] FIGS. 6A-6B illustrate a flow diagram of an example embodiment of a decision module routine.

[0011] FIGS. 7A-7B illustrate a flow diagram of an example embodiment of a CDD Control Action Determination routine.

[0012] FIGS. 8A-8B illustrate a flow diagram of an example embodiment of a CDD Coordinated Control Management routine.

[0013] FIG. 9 illustrates a flow diagram of an example embodiment of a routine for a target system being controlled.

[0014] FIG. 10 illustrates a network diagram of a portion of a distributed architecture of an example CDI system.

[0015] FIG. 11 illustrates a network diagram of a portion of an example Rules Builder component.

[0016] FIG. 12 illustrates a network diagram of example sub-components of an example Rules Builder component.

[0017] FIG. 13 illustrates a network diagram of interactions of portions of a Rules Builder component with an executing agent.

[0018] FIG. 14 illustrates a network diagram of interactions of portions of a Rules Builder component with chattering components.

[0019] FIG. 15 illustrates a diagram of a sliding window for use with chattering.

[0020] FIG. 16 illustrates a diagram of using Lebesgue integration to approximate a control trajectory for chattering.

[0021] FIG. 17 illustrates a diagram of various interactions of different portions of a CDI system.

[0022] FIG. 18 illustrates a diagram of example different types of rules.

[0023] FIG. 19 illustrates an example user interface related to medical/clinical auto-coding.

[0024] FIG. 20 illustrates a diagram of some components of a CDI system.

[0025] FIG. 21 illustrates a diagram of performing Pareto processing for use with mean field techniques.

[0026] FIG. 22 illustrates a network diagram of an example decision module agent.

[0027] FIG. 23 illustrates a network diagram of an example of offline workflows for knowledge capture.

[0028] FIG. 24 illustrates a network diagram of an example of workflows for mean field computation and Pareto Optimal.

[0029] FIG. 25 illustrates a network diagram of an example of an automated control system for a home solar micro-grid electrical generating system.

[0030] FIG. 26 illustrates a diagram of workflow and components of a portion of a CDI system.

[0031] FIG. 27 illustrates a diagram of a workflow for an inference process portion of a CDI system.

[0032] FIG. 28 illustrates a diagram of an overview workflow for a portion of a CDI system.

[0033] FIGS. 29A-29K illustrate examples of using a CDI system to iteratively determine near-optimal solutions over time for controlling a target system.

[0034] FIG. 30 is a block diagram illustrating example components of an embodiment of a system for using characteristics of a battery's state to perform automated control of DC power from the battery.



[0035] FIG. 31 is a block diagram illustrating example components of an embodiment of a system that performs automated control of DC power from multiple batteries in a coordinated manner.

[0036] FIG. 32 is a block diagram illustrating example components of an embodiment of a system for performing automated control of DC power from a battery that is part of a home electrical power system with solar power being generated, with the home power generation and use being monitored and synchronized by an external entity.

[0037] FIG. 33 is a block diagram illustrating example components of an embodiment of a system for performing automated control of DC power from a battery that is part of a home electrical power system with solar power being generated, in which typical monitoring and synchronization of the home power generation and use is temporarily interrupted.

[0038] FIG. 34 is a block diagram illustrating example components of an embodiment of a system for modeling a battery's state and operation to enable automated control of DC power from the battery.

[0039] FIG. 35 is a block diagram illustrating example visual displays of performance of an embodiment of a system that is using characteristics of a battery's state to perform automated control of DC power from the battery, with the home power generation and use being monitored and synchronized by an external entity.

[0040] FIG. 36 is a block diagram illustrating example visual displays of performance of an embodiment of a system that is using characteristics of a battery's state to perform automated control of DC power from the battery, in which typical monitoring and synchronization of the home power generation and use is temporarily interrupted.

[0041] FIG. 37 is a block diagram illustrating example components of an embodiment of a system for using characteristics of a battery's state to perform automated control of DC power from the battery.

[0042] FIG. 38 is a block diagram illustrating example components of an embodiment of a system for using characteristics of a battery's state to perform automated control of DC power from the battery.

[0043] FIG. 39 is a block diagram illustrating example components of an embodiment of a system for using characteristics of a battery's state to perform automated control of DC power from the battery.

[0044] FIG. 40 is a block diagram illustrating example components of an embodiment of a system that performs automated control of DC power from multiple batteries in a coordinated manner.

[0045] FIG. 41 is a block diagram illustrating an example environment in which a system for performing cooperative distributed control of target systems may be implemented.

[0046] FIG. 42 is a block diagram illustrating example components of an embodiment of a system that performs automated control of DC power from multiple batteries in a coordinated manner.

[0047] FIG. 43 is a block diagram illustrating example components of an embodiment of a system that performs automated control of DC power from multiple batteries in a coordinated manner.

[0048] FIG. 44 is a block diagram illustrating example components of an embodiment of a system that performs automated control of DC power from multiple batteries in a coordinated manner.

[0049] FIG. 45 is a block diagram illustrating example components of an embodiment of a system for using characteristics of a battery's state to perform automated control of DC power from the battery, as well as to use price forecasts and additional information to enhance financial performance of the system.

[0050] FIG. 46 is a block diagram illustrating example computing systems suitable for executing an embodiment of a system for performing cooperative distributed control of target systems in configured manners.

[0051] FIG. 47 is a block diagram illustrating an example architecture of a battery controller component.

[0052] FIG. 48 is a block diagram illustrating an example feedback controller.

[0053] FIG. 49 is a block diagram illustrating an example architecture for training a battery model for a battery.

[0054] FIG. 50 is a block diagram illustrating an example architecture for computing feedback gain for a tracking controlling system for a battery.

[0055] FIG. 51 is a block diagram illustrating an example architecture for a battery control system.

#### DETAILED DESCRIPTION

[0056] Techniques are described for implementing automated control systems to control or otherwise manipulate at least some operations of specified physical systems or other target systems. In at least some embodiments, the physical target systems include one or more batteries used to store and provide electrical power, and the automated operations to control the target systems include using characteristics of each battery's state to perform automated control of DC (direct current) power that is provided from the battery, such as by using a DC-to-DC amplifier (e.g., a FET, or field-effect transistor, amplifier) connected to the battery to control an amount of electrical current and/or voltage being output from the battery, in a real-time manner and to optimize long-term operation of the battery. The DC-to-DC amplifier may, for example, be part of a buck converter (or step-down converter) that steps down voltage while stepping up current from its input (supply) to its output (load) and/or be part of a boost converter (or step-up converter) that steps up voltage while stepping down current from its input (supply) to its output (load), referred to generally at times herein as a "boost/buck controller" or "buck/boost controller". In addition, in some embodiments and situations, multiple batteries may be controlled in such a manner by using multiple control systems that are each associated with one of the batteries, and with the overall control of the multiple batteries being coordinated in a distributed manner, such as based on interactions between the multiple associated control systems for the multiple batteries. A system that includes one or more batteries to be controlled may further include additional components in some embodiments and situations, such as one or more electrical sources and/or one or more electrical loads, with one non-exclusive example of such a type of system being one or more home or business electrical power systems that may optionally include electrical generation sources (e.g., solar panels, wind turbines, etc.) as well as electrical load from the house or business. Additional details are included below related to such techniques for implementing and using automated control systems to control target systems with one or more batteries.

[0057] The described techniques may provide a variety of benefits and advantages. Non-exclusive examples of such



benefits and advantages include controlling power output of a battery on the DC side (e.g., by varying DC current and voltage of the battery), allowing the battery to operate in its optimal or near-optimal physical state with respect to internal temperature and/or other battery parameters (e.g., by maintaining an internal chemistry of the battery within one or more chemical stoichiometric limits, such as a stoichiometric range), rather than fixing AC (alternating current) voltage and/or current being provided by an inverter connected to the battery at a single specified level, and forcing the battery to operate at a saturation level that provides maximum power but at the cost of possible non-reversible battery damage—in this manner, battery longevity and other operational performance may be optimized or otherwise enhanced by the described techniques, by allowing the battery to operate in a preferred range in which battery chemistry changes are reversible and in which short-circuits and other non-reversible damage is reduced or eliminated. In addition, in at least some embodiments, the automated control of the battery may further include active control of the battery to enhance and maintain power output amount resonance with the other components (e.g., a load and/or an external power grid), such that the amount of power provided does not exceed what is needed, while also satisfying at least a defined percentage or other amount of power output requests or load amounts (e.g., 50%, 65%, 100% or any other defined percentage or other amount). In this manner, such embodiments may be conceptualized as an automated control system to manage the internal state and operational performance (including longevity) of a battery being controlled, while satisfying power output requests if possible, rather than existing systems that fix the output voltage or current of a battery and fulfill all power requests even if it causes battery damage or other problems (e.g., excessive heating). While the benefits and operations discussed above and in some other locations herein relate to controlling power being output from a battery, it will be appreciated that the same techniques may be used to control power being stored into a battery from one or more sources, so as to cause the battery to operate in its optimal or near-optimal physical state with respect to heat and other battery parameters while storing power, and to optimize or otherwise enhance battery longevity and other operational performance by allowing the battery to operate in a preferred range in which battery chemistry changes are reversible, while satisfying at least a defined percentage or other amount of power input requests (e.g., 50%, 65%, 100% or any other defined percentage or other amount) for power being supplied by one or more sources. Additional benefits and advantages include the following, with the term ‘optimizing’ a feature or result as used herein meaning generally improving that feature or result (e.g., via partial or full optimization), and with the term ‘real-time’ as used herein meaning with respect to a time frame (e.g., fractions of a second, seconds, minutes, etc.) specific to a resulting component or system being controlled, unless otherwise indicated:

**[0058]** can improve battery lifetime by optimizing DC control variables, such as I (current), V (voltage) and R (amount of power being output)

**[0059]** can improve battery lifetime by optimizing DC control variables (I, V, R) in conjunction with a prior characterization of battery chemistry, and can optimize at DC level to improve performance and longevity

**[0060]** can optimize variables in real-time in DC domain to solve for objectives in AC phase

**[0061]** can optimize AC output in real-time to match grid frequency resulting in resonant operation, such as via control of only battery output and no other grid components

**[0062]** can improve charge/discharge cycles to improve long-term battery availability

**[0063]** can improve AC load response

**[0064]** can improve AC load response in combination with improving long-term battery availability

**[0065]** battery controller can run as embedded software on a processor in a self-sufficient manner

**[0066]** battery controller can be monitored and updated continuously from external location (e.g., the cloud or other network-accessible location)

**[0067]** battery controller can transmit battery characteristics to improve performance

**[0068]** can avoid expenses of static bar compensator hardware

Various other benefits and advantages may be further realized in at least some embodiments, as discussed in part in greater detail below.

**[0069]** FIG. 30 includes a block diagram 3000 illustrating example components of an embodiment of a system for using characteristics of a battery’s state along with other related information to perform automated control of DC power from the battery—in particular, various components of example system 3000 interact to control operations of the battery according to one or more defined goals in light of defined constraints, rules and other criteria, as discussed further below. In some embodiments, the automated activities to control the battery may be performed in a real-time manner and/or to optimize long-term operation of the battery (e.g., the life of the battery), while satisfying as many external requests for power (e.g., from a utility to which the battery can supply power) as is possible (e.g., at least a defined percentage or quantity of such requests).

**[0070]** In the illustrated example of FIG. 3000, a battery 3010 is shown that is being controlled via an actuator 3030 receiving a corresponding control signal from a battery controller component 3040 (also referred to as a “tracking controller” and/or “battery tracking controller” at times herein), such as by the battery controller specifying an amount of power to be generated as DC output of the battery. The specified power amount to be generated may include information indicating, for example, to increase or decrease the power being output by a specified amount, or to not change the power output. While not illustrated here, the output of the battery may serve to provide power to one or more loads (not shown), and in at least some embodiments may be connected to an inverter/rectifier component to convert the power output of the battery to AC power to support corresponding loads—such an inverter may, for example, control power being provided from the battery by regulating voltage and/or frequency of the AC power. Similarly, while also not illustrated here, input of the battery may serve to receive power from one or more sources (not shown), and in at least some embodiments may be connected to an inverter/rectifier component to convert AC power input from the sources to DC power for the battery—such a rectifier may, for example, control power being provided to the battery by regulating voltage and/or frequency of the AC power.



[0071] As part of determining how to control the battery, the battery controller component receives input from a sensor module 3020 regarding an internal state (not shown) of the battery, such as current values for voltage, electrical current, temperature, etc., and supplies corresponding information to a CDI agent 3050. The CDI agent, which is also referred to at times herein as a CDD (Collaborative Distributed Decision) decision module or system, receives the information from the battery controller related to the state of the battery, and also receives power supply requests from a utility component 3060, such as in a situation in which the battery supplies power at some or all times to an electrical grid (not shown) controlled by the utility. In particular, the CDI agent receives a particular request from the utility, receives and analyzes information about the state of the battery, and determines corresponding operations to take at the current time for the battery (e.g., an amount of output power to be supplied from the battery, and/or an amount of input power to be received and stored by the battery), which in at least some situations involve attempting to fully or partially satisfy the request from the utility for power in a real-time manner if the request can be satisfied in a way that also satisfies other constraints on the battery performance given the current state of the battery and the defined goal(s), such as to enable the battery to operate in a desired non-saturation range or level (e.g., with respect to an estimated internal temperature of the battery and/or estimated internal chemistry of the battery). After determining the corresponding operations to take at the current time for the battery, the CDI agent provides a corresponding tracking control signal to the battery controller, which determines how to currently modify or manipulate the actuator to effectuate the corresponding operations for the tracking control signal (e.g., an amount of positive or negative change to make in an amount of current being output from the battery), and sends a corresponding control signal to the actuator as discussed above.

[0072] While not illustrated in FIG. 30, the CDI Agent and/or battery controller may in some embodiments include a stored model of the battery that is used to estimate internal state of the battery and to select particular operations to perform based in part on that internal state. For example, in some embodiments a generic battery model may be used that is applicable to any type of battery, while in other embodiments a battery model may be used that is specific to a type of the battery (e.g., a type of chemical reaction used to store and/or generate electricity, such as lithium ion or nickel cadmium), while in yet other embodiments a battery model may be used that is designed and/or configured specifically for the particular battery in use. In addition, in at least some embodiments, a battery model that is initially employed in a particular system with a particular battery may be updated over time, such as to reflect improvements to the underlying structure of the model and/or to train the model to reflect operational characteristics specific to the particular battery and/or system in use (e.g., by monitoring how changes in observable battery state correlate to corresponding external battery electrical load and/or electrical source)—when training or otherwise adapting a model to a particular battery and/or system, the training/adaptation operations may in some embodiments be performed initially in a training phase before using the automated control system to control the battery, and/or in some embodiments may be performed continuously or periodically while the automated control

system is controlling the battery (e.g., to reflect changes over time in an impedance profile of the battery). Additional details are included elsewhere herein regarding such models, including their construction and use. In addition, while in some embodiments the battery controller and CDI agent may be implemented as separate components (e.g., with the battery controller implemented in whole or in part in hardware and/or firmware that is attached to the battery or otherwise at a location of the battery, and with the CDI agent implemented in part by software instructions executing on one or more computing systems remote from the battery location and optionally communicating with the battery controller over one or more intervening computer networks), in other embodiments the CDI agent and battery controller may be implemented as a single component (whether at the location of the battery or remote from it). Further details regarding operation of the CDI agent to determine operations to take for the battery are discussed in greater detail below.

[0073] In addition, while not illustrated with respect to FIG. 30, in some embodiments multiple batteries (e.g., tens, hundreds, thousands, millions, etc.) may each have an associated CDI agent that controls actions of that battery in a similar manner, and with the various batteries acting together in a coordinated manner to supply aggregate power to the utility or to other entities. In such embodiments, the utility or other external entity may send synchronization and monitoring signals for use by the various systems including the batteries, and the multiple CDI agents associated with the various batteries may interact to exchange information and maintain at least partial coordination between the operations of the batteries.

[0074] FIG. 31 is a block diagram illustrating one example of components of an embodiment of a system 3100 that performs automated control of DC power from multiple batteries in a coordinated manner, such as in a real-time manner and to optimize long-term operation of the batteries. In particular, multiple CDI agents 3050a-n are illustrated, which are each controlling operation of one of n associated batteries (not shown). In the example of FIG. 31, each CDI agent receives not only battery state information for its associated battery, but also may optionally receive additional information that includes requests for power from the utility, status information related to the utility, and/or predicted price information for power that is supplied to an electrical grid (not shown) of the utility from the batteries. Each CDI agent further includes rules and an optimizer component in this example, with the rules specifying the goals and constraints to satisfy for the associated battery, and the optimizer component using those rules and other input information to make corresponding automated control decisions for the associated battery, which are output by the CDI agent as target information for the associated battery. Additional details are included herein related to specification and use of such rules and operations of such an optimizer component. FIG. 31 further illustrates a combiner component 3110 that combines information from the various CDI agents to determine an aggregate response to the utility's requests, corresponding to the aggregate power being provided by the batteries being controlled.

[0075] It will also be appreciated that the described techniques may be used with a wide variety of other types of target systems, some of which are further discussed below, and that the invention is not limited to the techniques



discussed for particular target systems and corresponding automated control systems. For illustrative purposes, some embodiments are described below in which specific types of operations are performed, including with respect to using the described techniques with particular types of target systems and to perform particular types of control activities that are determined in particular manners. These examples are provided for illustrative purposes and are simplified for the sake of brevity, and the inventive techniques may be used in a wide variety of other situations, including in other environments and with other types of automated control action determination techniques, some of which are discussed below.

**[0076]** More generally, a target system to be controlled or otherwise manipulated may have numerous elements that are inter-connected in various manners, with a subset of those elements being inputs or other control elements that a corresponding automated control system may modify or otherwise manipulate in order to affect the operation of the target system. In at least some embodiments and situations, a target system may further have one or more outputs that the manipulations of the control elements affect, such as if the target system is producing or modifying physical goods or otherwise producing physical effects. For example, output of a target system involving automated control of one or more batteries may include electrical power being provided by the batteries, and inputs or other control elements may include the actuator(s) and/or battery tracking controller(s) used to manipulate the power being provided from the one or more batteries.

**[0077]** As part of implementing such an automated control system for a particular target system, an embodiment of a Collaborative Distributed Decision (CDD) system may use the described techniques to perform various automated activities involved in constructing and implementing the automated control system, including one or more CDI agents (also referred to as a CDD decision module or system, or a portion of such module or system) for use in controlling particular target systems. Some aspects of such activities of an example CDD system is provided below, with additional details included in U.S. patent application Ser. No. 14/746,738, filed Jun. 22, 2015 and entitled “Cooperative Distributed Control Of Target Systems;” in U.S. Patent Application No. 62/182,968, filed Jun. 22, 2015 and entitled “Applications Of Cooperative Distributed Control Of Target Systems;” in U.S. Patent Application No. 62/182,796, filed Jun. 22, 2015 and entitled “Gauge Systems;” and in international PCT Patent Application No. PCT/US2015/037022, filed Jun. 22, 2015 and entitled “Cooperative Distributed Control Of Target Systems;” each of which is hereby incorporated by reference in its entirety.

**[0078]** In particular, the CDD system may in some embodiments implement a Decision Module Construction component that interacts with one or more users to obtain a description of a target system, including restrictions related to the various elements of the target system, and one or more goals to be achieved during control of the target system—the Decision Module Construction component then performs various automated actions to generate, test and deploy one or more executable decision modules (also referred to at times as “decision elements” and/or “agents”) to use in performing the control of the target system. When the one or more executable decision modules are deployed and executed, the CDD system may further provide various

components within or external to the decision modules being executed to manage their control of the target system, such as a Control Action Determination component of each decision module to optimize or otherwise enhance the control actions that the decision module generates, and/or one or more Coordinated Control Management components to coordinate the control actions of multiple decision modules that are collectively performing the control of the target system.

**[0079]** As noted above, a Collaborative Distributed Decision (CDD) system may in some embodiments use at least some of the described techniques to perform various automated activities involved in constructing and implementing a automated control system for a specified target system, such as to modify or otherwise manipulate inputs or other control elements of the target system that affect its operation (e.g., affect one or more outputs of the target system). An automated control system for such a target system may in some situations have a distributed architecture that provides cooperative distributed control of the target system, such as with multiple decision modules that each control a portion of the target system and that operate in a partially decoupled manner with respect to each other. If so, the various decision modules’ operations for the automated control system may be at least partially synchronized, such as by each reaching a consensus with one or more other decision modules at one or more times, even if a fully synchronized convergence of all decision modules at all times is not guaranteed or achieved.

**[0080]** The CDD system may in some embodiments implement a Decision Module Construction component that interacts with one or more users to obtain a description of a target system, including restrictions related to the various elements of the target system, and one or more goals to be achieved during control of the target system the Decision Module Construction component then performs various automated actions to generate, test and deploy one or more executable decision modules to use in performing the control of the target system. The Decision Module Construction component may thus operate as part of a configuration or setup phase that occurs before a later run-time phase in which the generated decision modules are executed to perform control of the target system, although in some embodiments and situations the Decision Module Construction component may be further used after an initial deployment to improve or extend or otherwise modify an automated control system that has one or more decision modules (e.g., while the automated control system continues to be used to control the target system), such as to add, remove or modify decision modules for the automated control system.

**[0081]** In some embodiments, some or all automated control systems that are generated and deployed may further provide various components within them for execution during the runtime operation of the automated control system, such as by including such components within decision modules in some embodiments and situations. Such components may include, for example, a Control Action Determination component of each decision module (or of some decision modules) to optimize or otherwise determine and improve the control actions that the decision module generates. For example, such a Control Action Determination component in a decision module may in some embodiments attempt to automatically determine the decision module’s control actions for a particular time to reflect a near-optimal



solution with respect to or one more goals and in light of a model of the decision module for the target system that has multiple inter-related constraints—if so, such a near-optimal solution may be based at least in part on a partially optimized solution that is within a threshold amount of a fully optimized solution. Such determination of one or more control actions to perform may occur for a particular time and for each of one or more decision modules, as well as be repeated over multiple times for ongoing control by at least some decision modules in some situations. In some embodiments, the model for a decision module is implemented as a Hamiltonian function that reflects a set of coupled differential equations based in part on constraints representing at least part of the target system, such as to allow the model and its Hamiltonian function implementation to be updated over multiple time periods by adding additional expressions within the evolving Hamiltonian function.

**[0082]** In some embodiments, the components included within a generated and deployed automated control system for execution during the automated control system's runtime operation may further include one or more Coordinated Control Management components to coordinate the control actions of multiple decision modules that are collectively performing the control of a target system for the automated control system. For example, some or all decision modules may each include such a Control Action Determination component in some embodiments to attempt to synchronize that decision module's local solutions and proposed control actions with those of one or more other decision modules in the automated control system, such as by determining a consensus shared model with those other decision modules that simultaneously provides solutions from the decision module's local model and the models of the one or more other decision modules. Such inter-module synchronizations may occur repeatedly to determine one or more control actions for each decision module at a particular time, as well as to be repeated over multiple times for ongoing control. In addition, each decision module's model is implemented in some embodiments as a Hamiltonian function that reflects a set of coupled differential equations based in part on constraints representing at least part of the target system, such as to allow each decision module's model and its Hamiltonian function implementation to be combined with the models of one or more other decision modules by adding additional expressions for those other decision modules' models within the initial Hamiltonian function for the local model of the decision module.

**[0083]** As noted above, the described techniques may be used to provide automated control systems for various types of physical systems or other target systems. In one or more embodiments, an automated control system is generated and provided and used to control a micro-grid electricity facility, such as at a residential location that includes one or more electricity sources (e.g., one or more solar panel grids, one or more wind turbines, etc.) and one or more electricity storage and source mechanisms (e.g., one or more batteries). The automated control system may, for example, operate at the micro-grid electricity facility (e.g., as part of a home automation system), such as to receive requests from the operator of a local electrical grid to provide particular amounts of electricity at particular times, and to control operation of the micro-grid electricity facility by determining whether to accept each such request. If a request is accepted, the control actions may further include selecting

which electricity source (e.g., solar panel, battery, etc.) to use to provide the requested electricity, and otherwise the control actions may further include determine to provide electricity being generated to at least one energy storage mechanism (e.g., to charge a battery). Outputs of such a physical system include the electricity being provided to the local electrical grid, and a goal that the automated control system implements may be, for example, is to maximize profits for the micro-grid electricity facility from providing of the electricity.

**[0084]** In one or more embodiments, an automated control system is generated and provided and used to control a vehicle with a battery, a motor and in some cases an engine, such as an electrical bicycle in which power may come from a user who is pedaling and/or from a motor powered by the battery and/or the engine. The automated control system may, for example, operate on the vehicle or on the user, such as to control operation of the vehicle by determining whether at a current time to remove energy from the battery to power the motor (and if so to further determine how much energy to remove from the battery) or to instead add excess energy to the battery (e.g., as generated by the engine, and if so to further determine how much energy to generate from the engine; and/or as captured from braking or downhill coasting). Outputs of such a physical system include the effects of the motor to move the vehicle, and a goal that the automated control system implements may be, for example, to move the vehicle at one or more specified speeds with a minimum of energy produced from the battery, and/or to minimize use of fuel by the engine.

**[0085]** It will be appreciated that batteries may be used in a wide variety of other situations and may similarly be controlled by embodiments of the described techniques, such as with solar panels and other photovoltaic systems, electrical cars and other vehicles, etc.

**[0086]** Use of the described techniques may also provide various types of benefits in particular embodiments, including non-exclusive examples of beneficial attributes or operations as follows:

- [0087]** infer interests/desired content in a cold start environment where textual (or other unstructured) data is available and with minimal user history;
- [0088]** improve inference in a continuous way that can incorporate increasingly rich user histories;
- [0089]** improve inference performance with the addition of feedback, explicit/implicit, positive/negative and preferably in a real-time or near-real-time manner;
- [0090]** derive information from domain experts that provide business value, and embed them in inference framework;
- [0091]** dynamically add new unstructured data that may represent new states, and update existing model in a calibrated way;
- [0092]** renormalize inference system to accommodate conflicts;
- [0093]** immediately do inferencing in a new environment based on a natural language model;
- [0094]** add new information as a statistical model, and integrate with a natural language model to significantly improve inference/prediction;
- [0095]** integrate new data and disintegrate old data in a way that only improves performance;
- [0096]** perform inferencing in a data secure way;



[0097] integrate distinct inferencing elements in a distributed network and improve overall performance;

[0098] easily program rules and information into the system from a lay-user perspective;

[0099] inexpensively perform computer inferences in a way that is suitable for bandwidth of mobile devices; and

[0100] incorporate constraint information.

It will be appreciated that some embodiments may not include all some illustrative benefits, and that some embodiments may include some benefits that are not listed.

[0101] FIG. 41 is a block diagram illustrating an example environment in which a system 4100 for performing cooperative distributed control of target systems may be implemented, such as to control physical target systems having one or more batteries by using characteristics of each battery's state to perform automated control of DC power that is provided from the battery (e.g., in a real-time manner and to optimize long-term operation of the battery), such as in the matter discussed with respect to FIGS. 30-40 and 42-51 and elsewhere herein. In particular, the system 4100 is analogous in some respects to system 3100 of FIG. 31, such as the CDD decision modules 4124 of FIG. 41 (or their CDD Control Action Determination components 4144 and/or CDD Coordinated Control Management components 4146) correspond to the CDI agents 3050 of FIG. 31, and if the target system of FIG. 41 include batteries to be controlled (e.g., with respect to control elements 4161), but with additional details in FIG. 41 regarding operation of the CDD decision modules.

[0102] In particular, in the example environment of FIG. 41, target system 1 4160 is illustrated, with the automated control system 4122 being deployed and implemented to use in actively controlling the target system 1 4160. It will be appreciated that such target system may include a variety of mechanical, electronic, chemical, biological, and/or other types of components to implement operations of the target system in a manner specific to the target system. In the example of FIG. 41, the decision modules 4124 are represented as individual decision modules 4124a, 4124b, etc., to 4124n, and may be executing locally to the target system 1 4160 and/or in a remote manner over one or more intervening computer networks (not shown). In the illustrated example, each of the decision modules 4124 includes a local copy of a CDD Control Action Determination component 4144, such as with component 4144a supporting its local decision module 4124a, component 4144b supporting its local decision module 4124b, and component 4144n supporting its local decision module 4124n. Similarly, the actions of the various decision modules 4124 are coordinated and synchronized in a peer-to-peer manner in the illustrated embodiment, with each of the decision modules 4124 including a copy of a CDD Coordinated Control Management component 4146 to perform such synchronization, with component 4146a supporting its local decision module 4124a, component 4146b supporting its local decision module 4124b, and component 4146n supporting its local decision module 4124n.

[0103] As the decision modules 4124 and automated control system 4122 execute, various interactions 4175 between the decision modules 4124 are performed, such as to share information about current models and other state of the decision modules to enable cooperation and coordination between various decision modules, such as for a particular

decision module to operate in a partially synchronized consensus manner with respect to one or more other decision modules (and in some situations in a fully synchronized manner in which the consensus actions of all of the decision modules 4124 converge). During operation of the decision modules 4124 and automated control system 4122, various state information 4143 may be obtained by the automated control system 4122 from the target system 4160, such as initial state information and changing state information over time (e.g., with respect to batteries, not shown, corresponding to control elements 4161 of target system 1), and including outputs or other results in the target system 1 from control actions performed by the decision modules 4124.

[0104] The target system 1 in this example includes various control elements 4161 (e.g., batteries and their power output) that the automated control system 4122 may manipulate, and in this example each decision module 4124 may have a separate group of one or more control elements 4161 that it manipulates (such that decision module A 4124a performs interactions 4169a to perform control actions A 4147a on control elements A 4161a, decision module B 4124b performs interactions 4169b to perform control actions B 4147b on control elements B 4161b, and decision module N 4124n performs interactions 4169n to perform control actions N 4147n on control elements N 4161n). Such control actions affect the internal state 4163 of other elements of the target system 1, including optionally to cause or influence one or more outputs 4162 (e.g., aggregate electrical power being produced from the multiple batteries). As discussed in greater detail elsewhere herein, control element 4161a may, for example, be a FET actuator connected to a particular first battery (not shown) of target system 1 that is being controlled, and other control elements 4161b-n may similarly be other FET actuators connected to other batteries (not shown) of target system 1. As operation of the target system 1 is ongoing, at least some of the internal state information 4163 is provided to some or all of the decision modules to influence their ongoing control actions, with each of the decision modules 4124a-124n possibly having a distinct set of state information 4143a-143n, respectively, in this example.

[0105] As discussed in greater detail elsewhere, each decision module 4124 may use such state information 4143 and a local model 4145 of the decision module for the target system to determine particular control actions 4147 to next perform, such as for each of multiple time periods, although in other embodiments and situations, a particular automated control system may perform interactions with a particular target system for only one time period or only for some time periods. For example, the local CDD Control Action Determination component 4144 for a decision module 4124 may determine a near-optimal location solution for that decision module's local model 4145, and with the local CDD Coordinated Control Management component 4146 determining a synchronized consensus solution to reflect other of the decision modules 4124, including to update the decision module's local model 4145 based on such local and/or synchronized solutions that are determined. Thus, during execution of the automated control system 4122, the automated control system performs various interactions with the target system 4160, including to request state information, and to provide instructions to modify values of or otherwise manipulate control elements 4161 of the target system 4160. For example, for each of multiple time periods, decision



module **4124a** may perform one or more interactions **4169a** with one or more control elements **4161a** of the target system, while decision module **4124b** may similarly perform one or more interactions **4169b** with one or more separate control elements **B 4161b**, and decision module **4124n** may perform one or more interactions **4169n** with one or more control elements **N 4161n** of the target system **4160**. In other embodiments and situations, at least some control elements may not perform control actions during each time period.

[0106] In other embodiments and situations (e.g., if only a single battery is being controlled), the deployed copy of the automated control system may include only a single executing decision module, such as to include a local CDD Control Action Determination component but to not include any local CDD Coordinated Control Management component (since there are not other decision modules with which to synchronize and interact).

[0107] While not illustrated in FIG. 41, the distributed nature of operations of automated control systems such as those of **4122** allow partially decoupled operations of the various decision modules, include to allow modifications to the group of decision modules **4124** to be modified over time while the automated control system **4122** is in use, such as to add new decision modules **4124** and/or to remove existing decision modules **4124** (e.g., to reflect changes to underlying batteries in use, such as in different home power systems). In a similar manner, changes may be made to particular decision modules **4124** and/or **4128**, such as to change rules or other restrictions specific to a particular decision module and/or to change goals specific to a particular decision module over time, with a new corresponding model being generated and deployed within such a decision module, including in some embodiments and situations while the corresponding automated control system continues control operations of a corresponding target system. In addition, while each automated control system is described as controlling a single target system in the examples of FIG. 41, other configurations may be used in other embodiments and situations, such as for a single automated control system to control multiple target systems (e.g., multiple inter-related target systems, multiple target systems of the same type, etc.), and/or multiple automated control systems may operate to control a single target system, such as by each operating independently to control different portions of that target control system. It will be appreciated that other configurations may similarly be used in other embodiments and situations.

[0108] FIG. 32 illustrates an embodiment of controlling a battery as part of a larger target system, which in this example is a system **3200** involving a home power system that includes a solar panel thus, the example embodiments discussed previously with respect to FIGS. 30 and/or 31 may in some situations be used as part of a larger system such as the example system **3200**. In particular, the block diagram of FIG. 32 illustrates example components of an embodiment of a system **3200** for performing automated control of DC power from a battery that is part of a home electrical power system with solar power being generated, such as in a real-time manner and/or to optimize long-term operation of the battery, and with the home power generation and use being monitored and synchronized by an external entity, such as an entity providing or managing one or more CDI agents to control the battery of the system **3200**. In the example of FIG. 32, the example home's power system is

also connected to an external electrical grid from which it receives power and provides power at various times, with the battery serving to store electrical power generated by the solar power system and to supply power to the house and/or to the electrical grid as appropriate.

[0109] In the illustrated example of FIG. 32, components similar to those of FIG. 30 continue to be illustrated, including a battery **3010**, a sensor module **3020**, an actuator **3030** for the battery, an on-site battery tracking controller **3040**, etc. In the example of FIG. 32, however, the CDI agent **3050** of FIG. 30 is not illustrated as part of the components present at the physical location of the example house, such as if the CDI agent in use with respect to FIG. 32 instead executes in a remote location (e.g., in the cloud or other computer network location) and provides tracking and/or synchronization signals to the battery tracking controller **3040** of FIG. 32 in a manner similar to that illustrated with respect to FIG. 30. Such tracking and/or synchronization signals may, for example, include desired power output of the battery and/or desired battery parameters (e.g., internal temperature, voltage, current, etc.) for a current time or immediately subsequent time. In addition, as discussed in greater detail elsewhere herein, the CDI agent(s) may generate such tracking and/or synchronization signals based on monitored information about the battery **3010** (and any other batteries being controlled), power requests from the utility managing the external electrical grid, defined constraints or other rules to be used, etc.

[0110] In addition, a number of additional components are illustrated in FIG. 32, including an inverter/rectifier module **3210** that receives output power from the battery and/or supplies electrical power to the battery for storage, a solar panel **3220** that generates electrical power and that has its own associated sensor and inverter, a distribution box **3230** that receives and/or supplies power to an external electrical grid and that controls power distribution to a load **3240** for the house, etc. In addition, two local control agents **3260** and **3270** are illustrated to assist in controlling operation of the battery tracking controller **3040** of FIG. 32, with Agent1 **3260** interacting directly with the battery tracking controller, and Agent2 **3270** performing activities to synchronize the AC phase of the power for the battery with that of the house power system and/or grid, such as to provide resonance for the power being received and/or provided. The battery tracking controller **3040** and agents **3260** and **3270** (other than the utility sensor processor) are together referred to as a 'control processor' in this example, with the battery tracking controller providing system status updates, and with communications between the agents being managed to support such a multi-agent architecture. The tomography of Agent2 tracks dynamic changes in the battery state using a non-destructive x-ray. In addition, an external entity **3280** (e.g., the utility providing or managing the external electrical grid) is providing monitoring and synchronization signals in this example to the battery tracking controller **3040**, such as coordinate the power being used and/or provided via numerous such home power systems and/or other customers.

[0111] While the example of FIG. 32 involves use of the battery **3010** in a solar panel system, it will be appreciated that batteries may be charged and/or discharged in a variety of types of environments and systems, and similar activities of a corresponding CDI agent may be used to control such activities in the manner described herein.



[0112] FIG. 35 is a block diagram 3500 illustrating example visual displays of performance of an embodiment of a system that is using characteristics of a battery's state to perform automated control of DC power from the battery, with the home power generation and use being monitored and synchronized by an external entity, such as in the example of FIG. 32. In particular, in the example of FIG. 35, chart 3520 illustrates power requests that are received over time from a utility for one or more home power systems under the automated control of CDI agent(s), chart 3530 illustrates power that is supplied by the home power system(s) in response to the requests from the utility based on the automated control of the CDI agent(s), and chart 3510 combines the charts together. As is illustrated, the automated control of the CDI agent(s) provides highly accurate responses to the utility requests over time, while also optimizing performance of the one or more batteries in the home power system(s). In addition, chart 3540 further illustrates the incremental power supplied by the one or more batteries in the home power system(s) based on the automated control of the CDI agent(s).

[0113] FIG. 33 illustrates a further example of a system 3300 similar to that of system 3200 of FIG. 32, but in which the system 3300 is operating in a restricted emergency mode due to loss of monitoring/synchronization signals from the utility or other external entity that provides power requests. In particular, the block diagram of system 3300 illustrates a number of elements similar to that of system 3200 of FIG. 32, including a battery 3010 that is being controlled, a solar panel 3220, a distribution box 3230 and a house load 3240. However, in the example of FIG. 33, the signals from entity 3280 have been disrupted, and in response the control of the battery 3010 has been switched from the normal battery tracking controller 3040 of FIG. 32 to an alternative emergency battery tracking controller 3340 in FIG. 33. In addition, the components illustrated in FIG. 32 with respect to the Agent2 3270 are not illustrated in FIG. 33, as the system does not provide the phase synchronization and resonance functionality of FIG. 32 without the monitoring/synchronization signals that are missing in the situation of FIG. 33. In this example, the emergency battery tracking controller 3340 operates in a temporary fashion to maintain operation of the battery 3010 in the absence of the monitoring/synchronization signals, such as to receive information about the battery state and to issue corresponding instructions for the battery operation, such as based on battery state but without attempting to respond to power requests from the utility. Additional details are described below with respect to operation of the emergency controller in at least some embodiments.

[0114] FIG. 36 is a block diagram 3600 similar to the block diagram 3500 of FIG. 35, with respect to illustrating example visual displays of performance of an embodiment of a system that is using characteristics of a battery's state to perform automated control of DC power from the battery, but in which typical monitoring and synchronization of the home power generation and use is temporarily interrupted, such as in the example of FIG. 33. In particular, in the example of FIG. 36, chart 3620 illustrates power requests that are made over time by a utility for one or more home power systems under the automated control of CDI agent(s) but are not received by the home power systems due to the temporary interruption, and chart 3630 illustrates power that is supplied by the home power system(s) over time. Chart

3610 combines information together about the power supplied and the home power system's information about the utility's power requests, but with those power requests set to 0 during the time of temporary interruption since they are not received. Chart 3640 further illustrates the incremental power supplied by the one or more batteries in the home power system(s) based on the automated control of the CDI agent(s)—as is shown, despite the lack of utility power requests, the CDI agent(s) still control the power output of the one or more batteries based on other factors, such as the internal state of the one or more batteries, in order to continue to optimize performance of the one or more batteries in the home power system(s) to the extent possible based on the partial information that is available.

[0115] FIG. 34 is a block diagram illustrating example components of a model 3400 of a battery's state and operation, to enable automated control of DC power from the battery, such as in a real-time manner and to optimize long-term operation of the battery. In particular, the model 3400 of FIG. 34 includes representations of a charge path 3410 for the battery, a discharge path 3420 for the battery, a thermal model 3440 for the battery, a repository 3430, as well as a load 3450 that may be placed on the battery and supported by a nernst voltage 3480 produced by the battery. As discussed in greater detail elsewhere herein, in at least some embodiments and situations, the control of the battery may include preventing or minimizing the operation of the battery at a saturation level corresponding to the saturation model 3460, such as with the battery operating outside an optimal temperature range or other set of operational characteristics that decrease the life of the battery and/or impose other detrimental effects on the battery. By instead controlling the battery to operate in a linear range at all times or as much as is possible, various benefits are obtained, including increased battery life and/or other improved operational characteristics, as discussed in greater detail elsewhere herein.

[0116] FIG. 37 is a block diagram illustrating an embodiment of an example system 3700 for using characteristics of a battery's state to perform automated control of DC power from the battery, such as in a real-time manner and to optimize long-term operation of the battery. In particular, in a manner similar to that of system 3000 of FIG. 30, a CDI agent 3750 operates to control operation of a battery system 3710, with details of the battery system (e.g., a battery, sensor(s), actuators, battery tracking controller, etc.) not illustrated. Additional information 3740 is also illustrated to show information used and exchanged as part of controlling the battery system, including to maintain information about an internal state of the battery, and monitoring components to send and listen to various information being exchanged (e.g., as part of an information publish/subscribe or other push system) and used to update the internal state information.

[0117] FIG. 38 is a block diagram illustrating example components of an embodiment of a system 3800 for using characteristics of a battery's state to perform automated control of DC power from the battery, such as in a real-time manner and to optimize long-term operation of the battery. In particular, the illustrated system shows data flow into and between various components, including a state estimator component 3855 that receives information for a battery including various battery parameters (e.g., nernst potential, incremental value to saturation, capacity, etc.), other mea-



surements (e.g., voltage, current, temperature, etc.) and incremental control applied to the battery (e.g., a desired power to provide, either as an absolute value or an incremental change), and estimates a corresponding state of the battery. The state estimator component also includes a parameter adaptation engine (PAE) **3865** in this example that adapts the incoming battery parameter information in one or more defined manners. The state estimator component provides output about the estimated state to the LQ tracker component **3845**, which operates as a battery tracking controller to receive the estimated battery state information and other tracking parameters, to compute one or more optimal or approximately optimal control actions for the battery, and to output those control actions to be applied to the battery.

[0118] FIG. 39 is a block diagram illustrating example components of an embodiment of a system **3900** for using characteristics of a battery's state to perform automated control of DC power from the battery, such as in a real-time manner and to optimize long-term operation of the battery. In particular, FIG. 39 continues the example of FIG. 38, and provides further details regarding an example embodiment of the state estimator component **3855** of FIG. 38, which in the example of FIG. 39 is a discrete Kalman filter **3955**. In this example, the filter **3955** takes as input information about the previously estimated battery parameters, initial values for the parameters, and information about incremental control actions taken for the battery, and uses the information to generate a new estimate of the battery parameters, so as to dynamically update the battery parameter estimates to capture ongoing changes in the battery.

[0119] FIG. 40 is a block diagram illustrating example components of an embodiment of a system **4000** that performs automated control of DC power from multiple batteries in a coordinated manner, such as in a real-time manner and to optimize long-term operation of the batteries. In particular, the system **4000** of FIG. 40 has some similarities to that of FIG. 32, but illustrates an example architecture of a system to support coordinated control of large numbers of batteries and associated systems (e.g., over one million such batteries and associated systems in this example, such as to correspond to one or more regions, states, countries, etc.). In particular, in the illustrated example, various batteries and associated systems **4005** (e.g., home power systems with solar panels) having on-site battery tracking controllers are illustrated, along with one or more utilities **4085** that provide power requests for the batteries and associated systems **4005**, and one or more entities **4090** serving as system administration to manage a distributed control system **4075** for the batteries and associated systems **4005**.

[0120] In this example, the distributed control system **4075** is implemented in a centralized manner in a network-accessible location, such as via an online computing environment (e.g., Microsoft Azure), although it may be implemented in other manners in other embodiments. The distributed control system **4075** includes one or more components to interface with and interact with the utilities **4085**, one or more components to interface with and interact with the batteries and associated systems **4005**, and one or more monitoring and/or configuration components with which the system administration entities **4090** may interact to monitor and/or control the distributed control system **4075**. In addition, the various CDI agents that support the batteries and associated systems **4005** (e.g., with one CDI agent per battery pack and associated system) are executed in the

network-accessible location and are clustered together, with various inter-cluster communication mechanisms used (e.g., a publish/subscribe system with various topics, a communication service bus between at least some CDI agents and/or clusters, etc.). The clusters may be formed in various manners in various embodiments, such as to group CDI agents based on having associated batteries and systems that share one or more characteristics, such as geographical location (e.g., being part of the system electrical grid substation area) and/or operating characteristics. In addition, the clusters may be used to coordinate the CDI agents in stages and/or tiers, such as to first coordinate the CDI agents within a cluster, then coordinate between two or more clusters, etc., and optionally with multiple tiers of clusters (e.g., structured in a hierarchical manner). Various additional components may be provided and used as part of the distributed control system **4075**, such as a site management system to manage changes in CDI agents and/or batteries and associated systems (e.g., to add new CDI agents for new battery systems and/or to remove existing CDI agents for existing battery systems being removed from system **4075** management; to add new CDI agents to particular clusters and/or to remove existing CDI agents from clusters; to create, remove and modify clusters; etc.), storage services available from the network-accessible location to store state information and other information being used, resource management services available from the network-accessible location to manage computing resources provided by the network-accessible location, etc.

[0121] FIG. 42 is a block diagram illustrating example components of an embodiment of a system **4200** that performs automated control of DC power from multiple batteries in a coordinated manner, such as in a real-time manner and to optimize long-term operation of the batteries. In particular, FIG. 42 illustrates that a micro-grid network or other network of batteries and associated systems may be controlled and managed in at least some embodiments via a subset **4235** of the batteries and associated systems, such as if other of the batteries and associated systems are not part of a distributed control system. In such situations, a virtual network model may be created to model and estimate information about the micro-grid network or other network as a whole, including to estimate information about other battery and associated systems (referred to as a DER, or distributed energy resource, in this example) that are not being controlled.

[0122] FIG. 43 continues the example of FIG. 42, and includes additional information about the DER components being controlled in a coordinated manner. In particular, FIG. 43 is a block diagram illustrating example components of an embodiment of a system **4300** that performs automated control of DC power from multiple batteries in a coordinated manner, such as in a real-time manner and to optimize long-term operation of the batteries. The system **4300** includes a representation of a networked system **4345** analogous to distributed control system **4075** of FIG. 40 with multiple coordinated CDI agents, with the networked system **4345** also referred to as SERA (Smart Energy Reference Architecture) in this example. In this example, the networked system **4345** receives monitoring information from various DER components, and outputs synchronization signals to battery tracking controller components of the DER systems to control their operation.



[0123] FIG. 44 continues the examples of FIGS. 42 and 43, and includes additional information about the networked system 4345 of FIG. 43, such as to include information about a virtual network as noted with respect to FIG. 42, in order to simulate or otherwise estimate information about operation of the overall system of batteries and associated CDI agents. In particular, FIG. 44 is a block diagram illustrating example components of an embodiment of a system 4400 that performs automated control of DC power from multiple batteries in a coordinated manner, such as in a real-time manner and to optimize long-term operation of the batteries. The system 4400 includes models or other visual representations of various elements that include a virtual generation plant, virtual load, virtual substation(s), and virtual DER(s), with one or more CDI agents associated with each. The various elements exchange information as shown, including to estimate or otherwise model operation of the overall system as a whole, even in situations in which only a subset of the DER components are being controlled. The various CDI agents further use a mean field representation to coordinate their actions, as discussed in greater detail elsewhere herein.

[0124] FIG. 45 illustrates a system 4500 similar to system 3000 of FIG. 30, in which various components interact to control operations of the battery according to defined criteria, but with additional elements 4555 illustrated in FIG. 45. In particular, FIG. 45 is a block diagram illustrating example components of an embodiment of a system 4500 for using characteristics of a battery's state to perform automated control of DC power from the battery, such as in a real-time manner and to optimize long-term operation of the battery, as well as to use price forecasts and additional information to enhance financial performance of the system. The additional elements 4555 of FIG. 45 allow the system 4500 to obtain and use information about forecasted prices for power supplied to the utility, such as based on weather data, past and current price data, etc., such as to further manage the control of the battery to optimize or otherwise enhance one or more financial constraints, such as in combination with other constraints related to battery life and/or other performance characteristics.

[0125] Additional details regarding an example embodiment of a typical battery tracking controller that uses monitoring/synchronization signals from an external entity (also referred to as a 'hybrid tracker') are as follows.

[0126] One example embodiment of a tracking control system for a generic lithium ion high power battery cell may be modeled by the electric circuit representation in FIG. 34. The model 3400 includes current and voltage sources to represent the chemical reactions that characterize the cell in charging and discharging operations. The control actions are mediated by an actuator driven by a tracking controller. The Controller tracks a desired power signal that is generated by an inference module determining the desired response of the cell as a function of current electric storage level, and power demand. The inverter or rectifier circuits are not modeled in this example, and an idealized actuator is used for this model.

[0127] The dynamic behavior of the circuit in FIG. 34 is given by a differential equation as follows:

$$\dot{x}(t)=G(x(t),u(t),\text{parameters}) \quad (1)$$

where the state

$$x(t) = \begin{bmatrix} \text{power}(t) \\ \text{voltage}(t) \\ \text{current}(t) \\ \text{temperature}(t) \end{bmatrix} \in R^4,$$

and the control  $u(t) \in R$ , and the time line is represented by  $t \in R$ .

The function  $G(x_1(t), x_2(t), x_3(t), x_4(t), u(t))$  is given by (2)

$$\begin{bmatrix} G_1(x_1(t), x_2(t), x_3(t), x_4(t), u(t)) \\ G_2(x_1(t), x_2(t), x_3(t), x_4(t), u(t)) \\ G_3(x_1(t), x_2(t), x_3(t), x_4(t), u(t)) \\ G_4(x_1(t), x_2(t), x_3(t), x_4(t), u(t)) \end{bmatrix} = \begin{bmatrix} \alpha_1 x_1(t)x_2(t) + \alpha_2 x_1(t)u(t) + \beta_1 x_1(t) + \lambda_1 x_4(t - \tau) \\ \beta_2 x_2(t) - \beta_3 x_1(t) + \phi_1 u(t) + \lambda_2 x_4(t - \tau) \\ \beta_4 x_3(t) + \beta_5 x_2(t) + \beta_6 x_1(t) + \lambda_3 x_4(t - \tau) \\ \beta_7 x_1(t) + \beta_8 x_2(t) + \beta_9 x_3(t) + \phi_2 u(t) + \lambda_4 x_4(t - \tau) \end{bmatrix} \quad (2)$$

[0128] For simplicity the dependence on the parameters in the argument have been suppressed. The parameters:  $\alpha_1, \alpha_2, \beta_1, \dots, \beta_9, \phi_1, \phi_2, \lambda_1, \dots, \lambda_4$  represent the physical components in the cell model (i.e., resistors, the capacitor, the voltage, current sources and the saturation limits). The parameter  $\tau$  is a time delay, which can be estimated using the historical data. The parameters:  $\alpha_1, \alpha_2, \beta_1, \dots, \beta_6, \phi$  represent the physical components in the cell model (i.e., resistors, the capacitor, the voltage, current sources and the saturation limits). The control design generates an approximate solution of (1) by a piecewise linear stochastic differential equation over small time intervals.

Thus, let  $t_0, t_1, \dots, t_i, t_{i+1}, \dots$ , be a partition of the time line. On each interval  $[t_i, t_{i+1})$ , we seek solutions of the form

$$x(t) = x(t_i) + \delta x \quad (3)$$

$$\dot{x}(t) = \delta \dot{x} \quad (4)$$

over  $[t_i, t_{i+1})$ , where  $\delta \hat{x}(t)$  is the conditional mean of  $\delta x(t)$ , and  $\delta \dot{x}(t)$  is the conditional rate, obtained from a Kalman filter based on the following piecewise linear stochastic model. The stochastic increment  $\delta x(t)$  satisfies the following stochastic differential equation:

$$d\delta x(t) = \frac{\partial G(x(t_i), u(t_i))}{\partial x} \delta x(t) dt + \frac{\partial G(x(t_i), u(t_i))}{\partial u} \delta u(t) dt + G(x(t_i), u(t_i)) dt + d\omega(t) \quad (5)$$

where the noise has zero mean and the covariance is proportional to the second order term that is taken out in the approximation, e.g.,

$$\omega(t) \sim N\left(0, \alpha \frac{\partial^2 \sum_{j=1}^3 G_j(x(t_i), u(t_i))}{\partial x^2} + \epsilon I\right),$$

where  $\epsilon > 0$  and  $I$  is the identity matrix.



Sensors provide power, voltage, and current measurements. The observations from the battery sensors are modeled by

$$y(t)=x(t)+\theta(t)$$

where the observation noise  $\theta(t)$  is characteristic of the sensors, with zero mean and covariance matrix determined from the signal to noise ratio specifications. From (3) and (5) the incremental observation is given by

$$\delta y(t)=\delta x(t)+\theta(t)$$

with  $\delta y(t)=y(t)-x(t_i)$ .

An effective of the tracking problem has a criterion of the form

$$\min_{\delta u(t)} E \int_{t_i}^{t_{i+1}} \frac{1}{2} [(\delta x(t) - \tilde{x}_\delta(t))^T Q_0 (\delta x(t) - \tilde{x}_\delta(t)) + \delta \dot{x}(t)^T Q_1 \delta \dot{x}(t) + \delta u(t)^T R_0 \delta u(t) + \delta \dot{u}(t)^T R_1 \delta \dot{u}(t)] dt + \frac{1}{2} \delta x(t_{i+1})^T F_0 \delta x(t_{i+1}). \quad (6)$$

**[0129]** The tracking value  $\tilde{x}_\delta(t)$  is generated from rules defining the desired power behavior of the battery, dynamically. Note that rules from a CDI agent may not be given in terms of the original state vector (power, current, voltage, temperature), and if not are translated to state desired behavior in terms of the desired incremental state  $\tilde{x}_\delta(t)$ . A first example of such rules is given below.

Rule 1: over a week, at least 78% of the power demand should be satisfied.

Rule 2: battery longevity  $\geq$  five years.

Rule 3: for the battery, satisfy thresholds on amount of charge and discharge.

A second example of such rules is given below.

**[0130]** 1. Maximum charge limit: Do not charge the battery if the current charge has exceeded a first defined threshold level.

**[0131]** 2. Minimum charge limit: Do not discharge the battery if the current charge is below a second defined threshold level.

**[0132]** 3. Rate limitation: Do not change the desired power in/out to battery faster than a third defined threshold limit for rate of change.

**[0133]** 4. Switching between charge and discharge: Reduce rate of charge when power is near zero to prevent switching between charging and discharging (or vice versa) too fast or frequently, such as based on one or more fourth defined threshold levels.

**[0134]** 5. Maintain battery temperature: If the temperature begins to rise, adjust desired power output based on current battery state (i.e., if battery charge is low and the temperature starts increasing, then charge it).

As part of these rules of the second example, fulfill power requests as much as possible (referred to at times herein as maximizing the “Q factor”), such as, for example, at a rate of 70%. If power requests are being satisfied at more and/or at less than a desired level (e.g., 30%, 40%, 50%, 60%, 70%, 80%, 90%, etc.), weighting with one or more of the rules above may be adjusted up or down in some embodiments to lower or raise the level of satisfied power requests, respectively, and to correspondingly increase or reduce battery life, respectively, from such changes.

The rules here pertain to the increments, and convert the average over a long time frame (weeks, years) to a running average over  $[t_i, t_{i+1}]$ . For the battery model, the temperature

typically changes much faster than power, voltage and current. Motivated by Aoki’s partitioning method, the problem may be decoupled through the state variables so that the full state space problem can be transformed into two sub-problems which are solved with different time intervals. The temperature controller is referred to as “a high speed controller” and the power/voltage/current controller “a low speed controller”, with the following details directed to the low speed controller. In particular, an optimal control tracking problem to be solved for the low speed controller may be summarized as:

$$\min_{\delta u(t)} E \int_{t_i}^{t_{i+1}} \frac{1}{2} [(\delta x(t) - \tilde{x}_\delta(t))^T Q_0 (\delta x(t) - \tilde{x}_\delta(t)) + \delta \dot{x}(t)^T Q_1 \delta \dot{x}(t) + \delta u(t)^T R_0 \delta u(t) + \delta \dot{u}(t)^T R_1 \delta \dot{u}(t)] dt + \frac{1}{2} \delta x(t_{i+1})^T F_0 \delta x(t_{i+1}) \quad (6)$$

subject to

$$d\delta x(t) = \frac{\partial G(x(t_i), u(t_i))}{\partial x} \delta x(t) dt + \quad (7)$$

$$\frac{\partial G(x(t_i), u(t_i))}{\partial u} \delta u(t) dt + G(x(t_i), u(t_i)) dt + \lambda T(t - \tau) dt + d\omega_0(t)$$

$$d^2 \delta x(t) = \frac{\partial G(x(t_i), u(t_i))}{\partial x} d\delta x(t) dt + \frac{\partial G(x(t_i), u(t_i))}{\partial u} \delta \dot{u}(t) dt + d\omega_1(t) \quad (8)$$

$$\delta \dot{u}(t) = v(t). \quad (9)$$

**[0135]** In (7),

$$\lambda = \begin{bmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \end{bmatrix},$$

and  $T(t-\tau)$  is the average temperature for the previous time interval  $t_{i-1} \leq t \leq t_i$ . So  $\lambda T(t-\tau)$  can be considered a constant for  $t_i \leq t \leq t_{i+1}$ . The optimization problem formulated in (6)-(9) satisfies the assumptions of the separation principle. This leads to the following approach.

Step 1: Determine the conditional mean  $\delta \hat{x}(t)$  and the conditional rate  $\delta \dot{\hat{x}}(t)$  for  $t_i \leq t \leq t_{i+1}$  generated by a Kalman filter, which is described in a separate note.

Step 2: Solve (6)-(9), obtaining a feedback solution of the form

$$\delta \dot{u}(t) = K_0(t) \delta x(t) + K_1(t) \delta \dot{x}(t) + K_2(t) \delta u(t) + \psi(t) \quad (10)$$

where  $K_j(t)$ ,  $j=0, 1, 2$  and  $\psi(t)$  are the gains and the affine terms resulting from the optimization.

Step 3. Replace  $\delta x(t)$  and  $\delta \dot{x}(t)$  in (10) with  $\delta \hat{x}(t)$  and  $\delta \dot{\hat{x}}(t)$  from the outputs of the Kalman filter.

Step 4. Integrate the following equation,

$$\delta \dot{u}(t) = K_0(t) \delta \hat{x}(t) + K_1(t) \delta \dot{\hat{x}}(t) + K_2(t) \delta u(t) + \psi(t). \quad (11)$$

The intervals  $[t_i, t_{i+1}]$  for all  $i$ , are chosen small enough so that the gains  $K_j(t)$ ,  $j=0, 1, 2$  and  $\psi(t)$  can be considered constant over each interval, and are evaluated at  $t_i$ .

Using the variation of constant formula, the integral of (11) at  $t_{i+1}^-$  is given by

$$\delta u(t_{i+1}) = \int_{t_i}^{t_{i+1}} e^{K_2(t)(t_{i+1}-\tau)} (K_0(\tau) \delta \hat{x}(\tau) + K_1(\tau) \delta \dot{\hat{x}}(\tau) + \omega(\tau)) d\tau \quad (12)$$



Integrate (12) using an impulsive approximation assuming that the integrands are impulses at  $t_i$ . Thus, the incremental control for the low speed controller is

$$\delta u(t_{i+1}^-) \approx e^{K_2(t_i)(t_{i+1}-t_i)} (K_0 a(t_i) \delta x(t_i) + K_1(t_i) \dot{\delta x}(t_i) + \psi(t_i)). \quad (13)$$

Solving for the high speed controller, we get the incremental control for the high speed controller,  $\delta u_H(t_{i+1}^-)$ . Taking the linear combination of the two incremental controls, we get

$$\delta u(t_{i+1}^-) = \epsilon_1 \delta u_L(t_{i+1}^-) + \epsilon_2 \delta u_H(t_{i+1}^-). \quad (14)$$

The control to the battery at  $t_{i+1}$  is given by

$$u(t_{i+1}) = u(t_i) + \delta u(t_{i+1}^-). \quad (15)$$

The state at  $t_{i+1}$  is

$$x(t_{i+1}) = x(t_i) + \delta x(t_{i+1}). \quad (16)$$

the rate of the state at  $t_{i+1}$  is

$$\dot{x}(t_{i+1}) = \dot{x}(t_i) + \dot{\delta x}(t_{i+1}). \quad (17)$$

Both (16) and (17) are sent to the rules module for future requirements to update  $\tilde{x}_\delta(t)$ . This approach is repeated on the next interval. This is illustrated in FIG. 47, showing the architecture 4700 of the battery controller.

**[0136]** The model of the cell may be trained, for example, with the following parameters, which are typical of a 1 kW rated Lithium ion battery.

Parameter Table	
$\alpha_1$	-0.001
$\alpha_2$	-0.002
$\beta_1$	-0.01
$\beta_2$	-0.01
$\beta_3$	0.001
$\beta_4$	-0.01
$\beta_5$	0.01
$\beta_6$	0.001
$\phi$	0.01

With respect to the battery operation, the controller for this example compensate very well for the uncertainty in the parameters and approximation errors, and exhibits good quality of robustness and time response.

**[0137]** Additional details regarding an example embodiment of continuously estimating state for a battery in use, such as for a generic lithium ion high power battery cell and for a low speed controller, are as follows:

The stochastic differential equation for the state update is:

$$d\delta x(t) = A(t)\delta x(t)dt + B(t)\delta u(t)dt + f(t)dt + d\omega(t) \quad (1)$$

where the state  $\delta x(t) \in \mathbb{R}^3$ ,  $u(t) \in \mathbb{R}^1$ ,  $A(t)$  is a  $3 \times 3$  matrix,  $B(t)$  is a  $3 \times 1$  vector,  $f(t)$  is a  $3 \times 1$  vector and  $\omega(t)$  is a  $3 \times 1$  vector. The noise has zero mean and the covariance matrix  $W$ , e.g.,  $\omega(t) \sim N(0, W)$ .

The observation equation is given by

$$\delta y(t) = \delta x(t) + \theta(t) \quad (2)$$

where the measurement  $\delta y(t) \in \mathbb{R}^3$ , and  $\theta(t)$  is a  $3 \times 1$  vector. The noise has zero mean and the covariance matrix  $V$ , e.g.,  $\theta(t) \sim N(0, V)$ .

The state update equation is:

$$\dot{\delta \hat{x}}(t) = A(t)\delta \hat{x}(t) + B(t)\delta u(t) + K(t)(\delta y(t) - \delta \hat{x}(t)) \quad (3)$$

where the Kalman gain  $K$  is a  $3 \times 3$  matrix, given by

$$K(t) = P(t)V^{-1} \quad (4)$$

and the covariance update equation is:

$$\dot{P}(t) = A(t)P(t) + P(t)A(t)^T + W - P(t)V^{-1}P(t) \quad (5)$$

In order to determine the conditional mean  $\delta \hat{x}(t)$  and the conditional rate  $\dot{\delta \hat{x}}(t)$  for  $t_i \leq t < t_{i+1}$ , the first step is to initialize  $\delta \hat{x}(t_i) = 0$  and  $P(t_i) = 0$ .

The second step is to solve (5) numerically (e.g., Runge-Kutta method) for  $P(t)$ , for  $t_i \leq t < t_{i+1}$ . The third step is to get the Kalman gain, using (4).

The fourth step is to solve (3) numerically, where the parameters  $A(t)$ ,  $B(t)$  and the observations  $\delta y(t)$  are known. This provides the conditional mean  $\delta \hat{x}(t)$  and the conditional rate  $\dot{\delta \hat{x}}(t)$  for  $t_i \leq t < t_{i+1}$ .

**[0138]** Additional details regarding an example embodiment of computing feedback gain for a battery in use, such as for a generic lithium ion high power battery cell and for a low speed controller as described above, are as follows: An optimal control tracking problem to be solved may be summarized as:

$$\min_{\delta u(t)} E \int_{t_i}^{t_{i+1}} \frac{1}{2} [(\delta x(t) - \tilde{x}_\delta(t_i))^T Q_0 (\delta x(t) - \tilde{x}_\delta(t_i)) + \delta \dot{x}(t)^T Q_1 \delta \dot{x}(t) + \delta u(t)^T R_0 \delta u(t) + \delta \dot{u}(t)^T R_1 \delta \dot{u}(t)] dt + \frac{1}{2} \delta x(t_{i+1})^T F_0 \delta x(t_{i+1}) \quad (1)$$

subject to

$$d\delta x(t) = \frac{\partial G(x(t_i), u(t_i))}{\partial x} \delta x(t) dt + \frac{\partial G(x(t_i), u(t_i))}{\partial u} \delta u(t) dt + G(x(t_i), u(t_i)) dt + \lambda T(t - \tau) dt + d\omega_0(t) \quad (2)$$

$$d^2 \delta x(t) = \frac{\partial G(x(t_i), u(t_i))}{\partial x} d\delta x(t) dt + \frac{\partial G(x(t_i), u(t_i))}{\partial u} \delta \dot{u}(t) dt + d\omega_1(t) \quad (3)$$

$$\delta \dot{u}(t) = v(t). \quad (4)$$

Define the new state variable

$$z(t) = \begin{bmatrix} \delta x(t) \\ \delta \dot{x}(t) \\ \delta u(t) \end{bmatrix}$$

and the new control variable  $v(t) = \delta \dot{u}(t)$ , the state equation in (2-4) can be rewritten as

$$\dot{z}(t) = \tilde{A}(t)z(t) + \tilde{B}(t)v(t) + \tilde{f}(t)$$

with the initial conditions  $z(t_i) = 0$ , and

$$\tilde{A}(t_i) = \begin{bmatrix} \frac{\partial G(x(t_i), u(t_i))}{\partial x} & 0 & \frac{\partial G(x(t_i), u(t_i))}{\partial u} \\ 0 & \frac{\partial G(x(t_i), u(t_i))}{\partial x} & 0 \\ 0 & 0 & \epsilon I \end{bmatrix},$$

$$\tilde{B}(t_i) = \begin{bmatrix} 0 \\ \frac{\partial G(x(t_i), u(t_i))}{\partial u} \\ 1 \end{bmatrix}, \tilde{f}(t_i) = \begin{bmatrix} G(x(t_i), u(t_i)) + \lambda T(t - \tau) dt \\ 0 \\ 0 \end{bmatrix}.$$

The optimal control tracking problem to be solved is summarized as:

$$\min_{v(t)} E \int_{t_i}^{t_{i+1}} \frac{1}{2} [(z(t) - \tilde{z}(t_i))^T Q (z(t) - \tilde{z}(t_i)) + v(t)^T R v(t)] dt + \frac{1}{2} z(t_{i+1})^T F z(t_{i+1}) \quad (5)$$

subject to

$$dz(t) = \tilde{A}(t_i)z(t)dt + \tilde{B}(t_i)v(t)dt + \tilde{f}(t_i)dt + d\omega(t)$$

with

$$Q = \begin{bmatrix} Q_0 & 0 & 0 \\ 0 & Q_1 & 0 \\ 0 & 0 & R_0 \end{bmatrix}, R = R_1, F = \begin{bmatrix} F_0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix},$$

$$\tilde{z}(t_i) = \begin{bmatrix} \tilde{x}_o(t_i) \\ 0 \\ 0 \end{bmatrix}.$$

The feedback law  $v(t) = -R^{-1}\tilde{B}^T p(t)$  with

$$p(t) = \Sigma(t)z(t) + \phi(t).$$

$\Sigma(t)$  is a matrix, and  $\phi(t)$  is a vector. They are provided by solving the following ordinary differential equations.

$$\dot{\Sigma}(t) = -\Sigma(t)\tilde{A}(t_i) - \tilde{A}(t_i)^T \Sigma(t) + \Sigma(t)\tilde{B}(t_i)R^{-1}\tilde{B}(t_i)^T \Sigma(t) - Q \quad (6)$$

with  $\Sigma(t_{i+1}) = F$

$$\dot{\phi}(t) = (-\tilde{A}(t_i)^T + \Sigma(t)\tilde{B}(t_i)R^{-1}\tilde{B}(t_i)^T)\phi(t) - \Sigma(t)\tilde{f}(t_i) + Q\tilde{z}(t_i) \quad (7)$$

with  $\phi(t_{i+1}) = 0$ .

From (5),

**[0139]**

$$v(t) = -R^{-1}\tilde{B}^T(\Sigma(t)z(t) + \phi(t)) = -R^{-1}\tilde{B}^T \Sigma(t)z(t) - R^{-1}\tilde{B}^T \phi(t) = -K_{LQ}(t)z(t) + \psi(t) \quad (8)$$

where  $K_{LQ}(t) = -R^{-1}\tilde{B}^T \Sigma(t)$ ,  $\psi(t) = -R^{-1}\tilde{B}^T \phi(t)$ .

By the separation principle, we can calculate the  $K_{LQ}(t)$  and  $\psi(t)$  by solving the deterministic problem first, then replace the state  $z(t)$  in (8) with

$$\hat{z}(t) = \begin{bmatrix} \delta\hat{x}(t) \\ \delta\hat{x}(t) \\ \delta u(t) \end{bmatrix}$$

to get  $v(t)$ , as shown in (9).

$$v(t) = \delta u(t) = K_0(t)\delta\hat{x}(t) + K_1(t)\delta\hat{x}(t) + K_2(t)\delta u(t) + \psi(t) \quad (9)$$

Using an impulsive approximation assuming that the integrands are impulses at  $t_i$ , yields

$$\delta u(t_{i+1}^-) \approx e^{K_2(t_i)(t_{i+1}-t_i)}(K_0(t_i)\delta\hat{x}(t_i) + K_1(t_i)\delta\hat{x}(t_i) + \psi(t_i)) \quad (10)$$

To get the state value for the deterministic problem, solve the following differential equation.

$$\dot{z}(t) = \tilde{A}(t_i)z(t) + \tilde{B}(t_i)(K_{LQ}(t_i)z(t) + \psi(t_i)) + \tilde{f}(t_i) = (\tilde{A}(t_i) + \tilde{B}(t_i)K_{LQ}(t_i))z(t) + \tilde{B}(t_i)\psi(t_i) + \tilde{f}(t_i)$$

An architecture for such operations is illustrated in system **5000** of FIG. **50**.

**[0140]** Additional details regarding an example embodiment of an emergency battery tracking controller are as follows:

The state space model of the battery system:

$$\delta\dot{x}(t) = A\delta x(t) + B\delta u(t) + f \quad (1)$$

$$y(t) = C\delta x(t) \quad (2)$$

The state observer is of the form:

$$\delta\dot{\hat{x}}(t) = A\delta\hat{x}(t) + B\delta u(t) + f + K(y(t) - C\delta\hat{x}(t)) \quad (3)$$

where  $\hat{x}(t)$  is the estimate of  $x(t)$  and  $K$  is the latest Kalman filter gain computed by the hybrid tracker before we switch to the closed-loop tracker.

Equation (3) can be rewritten as follows:

$$\delta\dot{\hat{x}}(t) = (A - KC)\delta\hat{x}(t) + B\delta u(t) + f + Ky(t) \quad (4)$$

We take the Laplace transform of both sides of the equation:

$$s\delta\hat{x}(s) - \delta\hat{x}(0) = (A - KC)\delta\hat{x}(s) + B\delta u(s) + f(s) + Ky(s) \quad (5)$$

where  $\delta\hat{x}(0) = 0$  and

$$f(s) = \frac{f}{s}.$$

Solving for  $\delta\hat{x}(s)$ , we find:

$$\delta\hat{x}(s) = (sI - A + KC)^{-1} [B\delta u(s) + f(s) + Ky(s)] \quad (6)$$

Multiplying by  $C$  on both sides of (6) and replacing  $C\delta\hat{x}(s)$  with  $y(s)$ , we get:

$$y(s) = C(sI - A + KC)^{-1} (B\delta u(s) + f(s) + Ky(s)) \quad (7)$$

Solving for  $y(s)$ , we find:

$$(I - C(sI - A + KC)^{-1}K)y(s) = C(sI - A + KC)^{-1} (B\delta u(s) + f(s)) \quad (8)$$

$$y(s) = [I - C(sI - A + KC)^{-1}K]^{-1} C(sI - A + KC)^{-1} (B\delta u(s) + f(s)) \quad (9)$$

Let

$$G_u(s) = [I - C(sI - A + KC)^{-1}K]^{-1} C(sI - A + KC)^{-1} B \quad (10)$$

$$G_f(s) = [I - C(sI - A + KC)^{-1}K]^{-1} C(sI - A + KC)^{-1} f \quad (11)$$

then,

$$y(s) = G_u(s)\delta u(s) + G_f(s)f(s) \quad (12)$$

Let  $y_{ref}(t)$  be a reference signal, which can be obtained by setting  $\delta\hat{k}(t) = 0$  and  $\delta u(t) = 0$  in (18), and then solving for  $\delta x(t)$ , yielding

$$y_{ref}(t) = -(A^T A)^{-1} A^T f \quad (13)$$

Design a PI controller, such that

$$\delta u(s) = \left( K_p + \frac{K_i}{s} \right) (y_{ref}(s) - y(s)) \quad (14)$$



where  $K_K$  is the proportional gain and  $K_i$  is the integral gain to be determined, and

$$y_{ref}(s) = -(A^T A)^{-1} A^T f/s.$$

The feedback controller **4800** is shown in FIG. **48**. Combine (12) and (24), yielding

$$y(s) = G_u(s) \left( K_p + \frac{K_i}{s} \right) (y_{ref}(s) - y(s)) + G_f(s) f(s) \left[ I + G_u(s) \left( K_p + \frac{K_i}{s} \right) \right] y(s) = G_u(s) \left( K_p + \frac{K_i}{s} \right) y_{ref}(s) + G_f(s) f(s)$$

Solving for  $y(s)$ , we get:

$$y(s) = \left[ I + G_u(s) \left( K_p + \frac{K_i}{s} \right) \right]^{-1} G_u(s) \left( K_p + \frac{K_i}{s} \right) y_{ref}(s) + \left[ I + G_u(s) \left( K_p + \frac{K_i}{s} \right) \right]^{-1} G_f(s) f(s) \quad (15)$$

Now stabilize the feedback control system with  $K_p$  and  $K_i$  to be determined. The closed-loop transfer function is:

$$T(s) = \left[ I + G_u(s) \left( K_p + \frac{K_i}{s} \right) \right]^{-1} G_u(s) \left( K_p + \frac{K_i}{s} \right) = \left[ I + \left( [I - C(sI - A + KC)^{-1} K]^{-1} C(sI - A + KC)^{-1} B \left( K_p + \frac{K_i}{s} \right) \right) \right]^{-1} \left( [I - C(sI - A + KC)^{-1} K]^{-1} C(sI - A + KC)^{-1} B \left( K_p + \frac{K_i}{s} \right) \right) \quad (16)$$

Chose  $K_K$  and  $K_i$  such that all the poles of  $T(s)$  located in the open left half of the  $s$ -plane.

### Numerical Implementation

**[0141]** The numerical algorithm to compute  $K_p$  and  $K_i$  is shown below. As previously mentioned, the state space model of the battery system is as follows:

$$\dot{\delta x}(t) = A\delta x(t) + B\delta u(t) + f \quad (18)$$

$$\delta y(t) = C\delta x(t) \quad (19)$$

The state observer is of the form:

$$\dot{\delta \hat{x}}(t) = A\delta \hat{x}(t) + B\delta u(t) + f + K(\delta y(t) - C\delta \hat{x}(t)) \quad (20)$$

where  $\hat{x}(t)$  is the estimate of  $x(t)$  and  $K$  is the latest Kalman filter gain computed by the hybrid tracker before we switch to the closed-loop tracker.  $\delta y(t)$  is obtained from the history measurement data. Equation (20) can be rewritten as follows:

$$\dot{\delta \hat{x}}(t) = (A - KC)\delta \hat{x}(t) + B\delta u(t) + f + K\delta y(t) \quad (21)$$

Equation (21) can be approximated by

$$\delta \hat{x}(t+\Delta) = e^{(A-KC)\Delta} \delta \hat{x}(t) + \int_t^{t+\Delta} e^{(A-KC)(t+\Delta-\tau)} (B\delta u(\tau) + f + K\delta y(\tau)) d\tau \quad (22)$$

Assuming the integrands are impulses at  $t$ , we have

$$\begin{aligned} \delta \hat{x}(t+\Delta) &= e^{(A-KC)\Delta} \delta \hat{x}(t) + e^{(A-KC)\Delta} (B\delta u(t) + f + K\delta y(t)) \\ &= e^{(A-KC)\Delta} (\delta \hat{x}(t) + B\delta u(t) + f + K\delta y(t)) \end{aligned} \quad (23)$$

Design a PI controller, such that

$$\delta u(t) = K_p(\delta y_{ref} - \delta y(t)) + K_i \int_0^t (\delta y_{ref} - \delta y(\tau)) d\tau \quad (24)$$

where  $K_p$  is the proportional gain and  $K_i$  is the integral gain to be determined,  $\delta y_{ref}$  is a reference signal, which can be obtained by setting  $\delta \dot{x}(t) = 0$  and  $\delta u(t) = 0$  in (18), and then solving for  $\delta x(t)$ , yielding

$$\delta y_{ref} = (A^T A)^{-1} A^T f \quad (25)$$

Solve a least squares problem for  $K_p$  and  $K_i$  such that the integral of squared residuals  $\int_{t=0}^T (\delta y_{ref} - C\delta \hat{x}(t))^T (\delta y_{ref} - C\delta \hat{x}(t)) dt$  is minimized, subject to equation (23), (24) and (25).

$$\min_{K_p, K_i} \sum_{n=0}^N (\delta y_{ref} - C\delta \hat{x}(n\Delta))^T (\delta y_{ref} - C\delta \hat{x}(n\Delta)) \quad (26)$$

$$\text{s.t. } \delta u(n\Delta) = K_p(\delta y_{ref} - \delta y(n\Delta)) + K_i \sum_{i=0}^n (\delta y_{ref} - \delta y(i\Delta)) \quad (27)$$

$$\delta \hat{x}((n+1)\Delta) = e^{(A-KC)\Delta} (\delta \hat{x}(n\Delta) + B\delta u(n\Delta) + f + K\delta y(n\Delta)) \quad (28)$$

$$\delta \hat{x}(0) = 0, \forall n = 0 \dots N. \quad (29)$$

where, by Nyquist,  $N$  is determined by the minimum eigenvalues of the state transition matrix  $A$ ,  $\lambda_{min}$ , i.e.,

$$N = (5 \sim 10) \frac{1}{\lambda_{min}} \quad (30)$$

Let  $\tilde{A} = e^{(A-KC)\Delta}$ ,  $\delta \hat{x}_n = \delta \hat{x}(n\Delta)$ ,  $\delta y_n = \delta y(n\Delta)$  and replace  $\delta u(n\Delta)$  in (28) with (27), yields:

$$\min_{K_p, K_i} \sum_{n=0}^N (\delta y_{ref} - C\delta \hat{x}_n)^T (\delta y_{ref} - C\delta \hat{x}_n) \quad (31)$$

$$\text{s.t. } \delta \hat{x}_{n+1} = \tilde{A} \left( \delta \hat{x}_n + BK_p(\delta y_{ref} - \delta y_n) + BK_i \sum_{i=0}^n (\delta y_{ref} - \delta y_i) + f + K\delta y_n \right) \quad (32)$$

$$\delta \hat{x}_0 = 0, \forall n = 0 \dots N. \quad (33)$$

Solve the problem using a dynamic programming approach. The cost-to-go function is written as:

$$V(\delta \hat{x}_n, n) = \min_{K_p, K_i} \{ (\delta y_{ref} - C\delta \hat{x}_n)^T (\delta y_{ref} - C\delta \hat{x}_n) + V(\hat{x}_{n+1}, n+1) \} \quad (34)$$

$$= \min_{K_p, K_i} \{ (\delta \hat{x}_n)^T C^T C \delta \hat{x}_n - 2(\delta \hat{x}_n)^T C^T \delta y_{ref} + (\delta y_{ref})^T \delta y_{ref} + V(\delta \hat{x}_{n+1}, n+1) \} \quad (35)$$

with  $V(\hat{x}_N, N) = (\delta \hat{x}_N)^T C^T C \delta \hat{x}_N - 2(\delta \hat{x}_N)^T C^T \delta y_{ref} + (\delta y_{ref})^T \delta y_{ref}$  and equate it to the Riccati form

$$V(\delta \hat{x}_n, n) = (\delta \hat{x}_n)^T \Phi_n \delta \hat{x}_n + \Omega_n \quad (36)$$

where  $\Phi_n$  represents a symmetric positive definite matrix,  $\Phi_n$  is a positive vector, and  $\Omega_n$  is a positive scalar. Combining the equations (35), (36) and the dynamics in (32), yields

$$\begin{aligned} V(\delta \hat{x}_n, n) &= \min_{K_p, K_i} \{ (\delta \hat{x}_n)^T C^T C \delta \hat{x}_n - 2(\delta \hat{x}_n)^T C^T \delta y_{ref} + \\ &\quad (\delta y_{ref})^T \delta y_{ref} + (\delta \hat{x}_{n+1})^T \Phi_n \delta \hat{x}_{n+1} + (\delta \hat{x}_{n+1})^T \Psi_{n+1} + \Omega_{n+1} \} \\ &= \min_{K_p, K_i} \{ (\delta \hat{x}_n)^T C^T C \delta \hat{x}_n - 2(\delta \hat{x}_n)^T C^T \delta y_{ref} + (\delta y_{ref})^T \delta y_{ref} \\ &\quad + \left( \delta \hat{x}_n + BK_p(\delta y_{ref} - \delta y_n) + BK_i \sum_{i=0}^n (\delta y_{ref} - \delta y_i) + f + K \delta y_n \right)^T \tilde{A}^T \Phi_n \\ &\quad + \left( \delta \hat{x}_n + BK_p(\delta y_{ref} - \delta y_n) + BK_i \sum_{i=0}^n (\delta y_{ref} - \delta y_i) + f + K \delta y_n \right)^T \\ &\quad \left. \left( \delta \hat{x}_n + BK_p(\delta y_{ref} - \delta y_n) + BK_i \sum_{i=0}^n (\delta y_{ref} - \delta y_i) + f + K \delta y_n \right)^T \right\} \\ &\quad \Psi_{n+1} + \Omega_{n+1} \end{aligned} \quad (37)$$

In order to minimize this expression, isolate the terms with  $K_p$  and  $K_i$  in them,

$$\begin{aligned} & (BK_p(\delta y_{ref} - \delta y_n))^T \tilde{A}^T \Phi_{n+1} \tilde{A} (BK_p(\delta y_{ref} - \delta y_n)) + \\ & \quad 2(BK_p(\delta y_{ref} - \delta y_n))^T \tilde{A}^T \Phi_{n+1} \\ & \quad \tilde{A} \left( BK_i \sum_{i=0}^n (\delta y_{ref} - \delta y_i) + \delta \hat{x}_n + f + K \delta y_n \right) + \\ & \quad (BK_p(\delta y_{ref} - \delta y_n))^T \Psi_{n+1} + \left( BK_i \sum_{i=0}^n (\delta y_{ref} - \delta y_i) \right)^T \tilde{A}^T \\ & \quad \Phi_{n+1} \tilde{A} \left( BK_i \sum_{i=0}^n (\delta y_{ref} - \delta y_i) \right) + 2 \left( BK_i \sum_{i=0}^n (\delta y_{ref} - \delta y_i) \right)^T \\ & \quad \quad \tilde{A}^T \Phi_{n+1} \tilde{A} (BK_p(\delta y_{ref} - \delta y_n) + \delta \hat{x}_n + f + K \delta y_n) \\ & \quad + \left( BK_i \sum_{i=0}^n (\delta y_{ref} - \delta y_i) \right)^T \Psi_{n+1} \end{aligned} \quad (39)$$

and take the derivative with respect to each element of  $K_p$  and  $K_i$ , and set the value to 0. This yields the solution for the optimal gain  $K_p^*$ , and  $K_i^*$  with respect to  $\Phi_n$ ,  $\Phi_n$  and  $\Omega_n$ . Equating the Riccati form (36) with the value function in (38) evaluated at  $K_p^*$  and  $K_i^*$ , solve for  $\Phi_n$ ,  $\Phi_n$  and  $\Omega_n$ , thereafter, numerical values for  $K_p^*$  and  $K_i^*$ .

**[0142]** Additional details regarding an example embodiment of controlling a battery as part of a network of multiple home solar power systems are as follows, which in this example use functionality to support Internet Of Things (IoT) capabilities in an online network environment such as Microsoft's Azure:

#### Architecture

**[0143]** In this example, the architecture is designed to control millions of batteries at multiple sites. The architecture consists of the following components:

**[0144]** Battery control units. These reside on-site and allow for control of batteries even in the case where network connection is lost

**[0145]** Battery interface (Azure IoT hub). This allows sending and receiving of data to each battery control unit

**[0146]** Utility interface (Azure IoT hub). Allows sending and receiving of data to the utility company.

**[0147]** CDI agents. There is one agent for each battery to control; these communicate to the batteries over the battery interface, and with each other to determine an optimal control.

**[0148]** Site management. This component adapts the network of CDI agents to allow for batteries to be added or removed from control automatically

**[0149]** Monitoring. This component tracks vital statistics of the CDI agents to allow a user to verify that everything is working, and to provide diagnostics to take action if a problem occurs.

#### Battery Control Units

**[0150]** The battery control units are located in hardware, on site, and are directly connected to the battery. The battery control units connect to the cloud to receive their desired control via the Azure Internet of Things (IoT) Hub. When the connection is lost, the control units provide backup control to the batteries to ensure they stay in a stable state.

#### Battery Interface

**[0151]** As mentioned above, each site will connect to this example system via the Azure IoT (Internet of Things) Hub. This will be done using the HTTPS or AMQPS protocols. The IoT hub allows us to scale the number of batteries to the millions, and will handle authentication and message routing. Each battery control unit will have a unique topic that it sends sensor data to, and a unique topic that other components can send control messages to back to the battery.

#### Utility Interface

**[0152]** Another process is responsible for reading/writing utility requests to the utility. This interface is likely a variant of SCADA, but may be adapted to suit the utility company.

#### CDI Agents

**[0153]** Each battery will have a process, known as a "CDI agent" to compute the optimal control for the battery. The agents will be implemented as service bus listeners in the Azure cloud. The agents communicate with other agents via the Azure Service Bus using AMQP. Each agent subscribes to the IoT topic from the battery it controls, and so can receive the current state of its battery.

To compute an approximate globally optimal control, the CDI agents communicate their estimation state and optimal control with each other (known as a "mean field"). To keep computation and message passing scalable, the CDI agents are clustered into a 2- or 3-level hierarchy, based on location and possibly battery type.

Clusters may, for example, be created for each substation level, and range from 100-1000 agents. Each cluster has a service bus topic to which all agents in the cluster publish and subscribe to. To share state between clusters, a particular node in each cluster is marked as a master. This node additionally subscribes and sends state to another topic shared by other master nodes.



### Site Management

**[0154]** There is another set of nodes that are used for site management, that is, to be able to adapt the network of CDI agents as batteries are added or removed, or repair the network if failures occur.

**[0155]** The network structure is encoded by the list of agents and battery control units, as well as the list of which topics each agent should publish and subscribe to. This data is saved within the Azure Storage Service. Creating, destroying, and updating agents is done using the Azure Resource Management API.

**[0156]** When a new battery is installed, the system will receive new messages from the battery control unit for a request to add to the system; this request will include data about the battery (type, location, substation, etc.) that will be needed to find the appropriate CDI agent cluster to add to.

**[0157]** When a battery goes offline for a short time (e.g., due to loss of connection) then the IoT hub manages the last estimate of the state of the battery, and the battery goes into the backup local control until it reconnects. However, if a battery goes offline for a long time (e.g., because it is permanently disconnected) then the Site Management component then stops the associated CDI agent.

**[0158]** When clusters become too unbalanced, that is, they have too many or too few nodes, the site management component splits the cluster in two or combines nearby clusters into one, such as based on proximity or battery type. The site management component sends the affected agents the new lists of topics they subscribe and publish to. The picking of CDI master nodes is also done by the Site Management component.

### Monitoring

**[0159]** The monitoring component listens for messages from each battery control unit and each CDI agent to make sure that

**[0160]** messages are being sent as expected

**[0161]** overall power is being produced within tolerances of the utility

**[0162]** battery health is maintained for each battery at each site

If a CDI agent is failing to control a battery, then the monitoring may alert a user over a dashboard, and also signal the site management component to restart that agent. If battery control units are failing to respond for a sufficient time, the monitoring may alert to have someone confirm that their battery is going offline, or have someone go out to the site if necessary. If power is not being produced to the requirements of the utility, or battery health is not maintained for some batteries, then the monitoring agent may send messages to the CDI agents to update their parameters to better meet demand or save battery life.

### Updates/Maintenance

**[0163]** For updates to Battery Control code, the battery control unit is sent a first message to download the new code, and then a second message to stop the current processes and switch to the newly downloaded code. For updates to CDI Agent code, the site management component can stop running agents one at a time, and restart them with the new version of code. To deploy new code to the monitoring, site

management, and utility listener, an updated version is started in parallel, and verified to work, then the old version is decommissioned.

**[0164]** As noted above, in some embodiments and situations, a general battery model may be trained to reflect attributes and characteristics specific to a particular battery in use, such as before controlling of the particular battery in use begins and/or during such controlling and use. Additional details regarding an example embodiment of training a model are as follows:

A parameter learning engine is described for adapting the parameters in the incremental model for the power battery cell to reflect a particular power battery cell.

The stochastic differential equation is:

$$d\delta x(t) = A(t)\delta x(t)dt + B(t)\delta u(t)dt + f(t)dt + d\omega(t) \quad (1)$$

where the state  $\delta x(t) \in \mathbb{R}^3$ ,  $u(t) \in \mathbb{R}^1$ ,  $A(t)$  is a  $3 \times 3$  matrix,  $B(t)$  is a  $3 \times 1$  vector,  $f(t)$  is a  $3 \times 1$  vector and  $\omega(t)$  is a  $3 \times 1$  vector. The noise has zero mean and the covariance matrix  $W$ , e.g.,  $\omega(t) \sim N(0, W)$ .

The observation equation is given by

$$\delta y(t) = \delta x(t) + \theta(t) \quad (2)$$

where the measurement  $\delta y(t) \in \mathbb{R}^3$ , and  $\theta(t)$  is a  $3 \times 1$  vector. The noise has zero mean and the covariance matrix  $V$ , e.g.,  $\theta(t) \sim N(0, V)$ .

The parameter learning engine estimates the  $A$  matrix, creating a vector of the nine values  $\text{vect}(\alpha_{11}, \alpha_{12}, \alpha_{13}, \alpha_{21}, \alpha_{22}, \alpha_{23}, \alpha_{31}, \alpha_{32}, \alpha_{33})$ . These parameters are expected to change slower than the incremental state dynamic updates, so a discrete Kalman filter can be used to estimate the  $A$  parameters. For this parameter learning engine, the values of  $B(t)$ ,  $\delta u(t)$ ,  $f(t)$ ,  $\omega(t)$ ,  $\delta y(t)$ , and  $\theta(t)$  are known at times  $t_i$ ,  $t_{i-1}$  and  $t_{i-2}$ .

Let  $\psi_t = \text{vect}(\alpha_{11}, \alpha_{12}, \alpha_{13}, \alpha_{21}, \alpha_{22}, \alpha_{23}, \alpha_{31}, \alpha_{32}, \alpha_{33})_t$  and the dynamics of the parameters is,  $\psi_{t+1} = \psi_t + \Lambda$  (3)

where  $\Lambda_t$  is  $N(0, \lambda)$ .

Solving (1) using the variation of constants formula, we get

$$\delta x(t_{i+1}) = e^{A(t_{i+1}-t_i)}\delta x(t_i) + \int_{t_i}^{t_{i+1}} e^{A(t_{i+1}-\tau)}B(\tau)\delta u(\tau) + f(\tau) + \omega(\tau) d\tau \quad (4)$$

and using an impulsive approximation, yields,

$$\delta x(t_{i+1}) = e^{A(t_{i+1}-t_i)}\delta x(t_i) + e^{A(t_{i+1}-t_i)}(B(t_i)\delta u(t_i) + f(t_i) + \omega(t_i)) \quad (5)$$

and applying a first order Taylor series expansion on the exponential term yields,

$$\delta x(t_{i+1}) = (I + A(t_{i+1}-t_i))(\delta x(t_i) + B(t_i)\delta u(t_i) + f(t_i) + \omega(t_i)). \quad (6)$$

Substituting  $\delta x(t_i) = \delta y(t_i) - \theta(t_i)$  from (2) yields,

$$\delta Y(t_{i+1}) = (I + A(t_{i+1}-t_i))(\delta y(t_i) + B(t_i)\delta u(t_i) + f(t_i) + \omega(t_i) - \theta(t_i)) + \theta(t_{i+1}). \quad (7)$$

To simplify the notation, let

$$\gamma(t_i) = \delta y(t_i) + B(t_i)\delta u(t_i) + f(t_i) \quad (8)$$

then,

$$\delta y(t_{i+1}) = (I + A(t_{i+1}-t_i))(\gamma(t_i) + \omega(t_i) - \theta(t_i)) + \theta(t_{i+1})\delta y(t_{i+1}) \\ = A(t_{i+1}-t_i)\gamma(t_i) + (\gamma(t_i) + \omega(t_i) - \theta(t_i)) + \theta(t_{i+1}) + \Omega(t_{i+1})$$

where  $\Omega(t_{i+1})$  is a noise term. Then we have,

$$\delta y(t_{i+1}) - \gamma(t_i) = A(t_{i+1}-t_i)\gamma(t_i) + \omega(t_i) - \theta(t_i) + \theta(t_{i+1}) + \Omega(t_{i+1}) = \\ A(t_{i+1}-t_i)(\gamma(t_i)) + \Xi(t_{i+1}) \quad (9)$$

Let the sum of the independent noise vectors be  $\Xi(t_i) = \omega(t_i) - \theta(t_i) + \theta(t_{i+1}) + \Omega(t_{i+1})$  which has zero mean and because the noises are independent, the covariance matrix is the sum of the individual variances.

Now, to estimate the nine parameter values in A, which are constant over long intervals, we write (9) at times  $t_i$ ,  $t_{i-1}$  and  $t_{i-2}$ ,

$$\delta y = (t_i) - \gamma(t_{i-1}) = A(t_i - t_{i-1})(\gamma(T_{i-1})) + \Xi(t_i)$$

$$\delta y = (t_{i-1}) - \gamma(t_{i-2}) = A(t_i - t_{i-2})(\gamma(T_{i-2})) + \Xi(t_{i-1})$$

$$\delta y = (t_{i-2}) - \gamma(t_{i-3}) = A(t_i - t_{i-3})(\gamma(T_{i-2})) + \Xi(t_{i-2})$$

and rewrite in matrix form for the parameter observation equation for the parameters, as,

$$\begin{bmatrix} \delta y_1(t_i) - \gamma_1(t_{i-1}) \\ \delta y_2(t_i) - \gamma_2(t_{i-1}) \\ \delta y_3(t_i) - \gamma_3(t_{i-1}) \\ \delta y_1(t_{i-1}) - \gamma_1(t_{i-2}) \\ \delta y_2(t_{i-1}) - \gamma_2(t_{i-2}) \\ \delta y_3(t_{i-1}) - \gamma_3(t_{i-2}) \\ \delta y_1(t_{i-2}) - \gamma_1(t_{i-3}) \\ \delta y_2(t_{i-2}) - \gamma_2(t_{i-3}) \\ \delta y_3(t_{i-2}) - \gamma_3(t_{i-3}) \end{bmatrix} = \mathcal{G}(t_i) \begin{bmatrix} a_{11} \\ a_{12} \\ a_{13} \\ a_{21} \\ a_{22} \\ a_{23} \\ a_{31} \\ a_{32} \\ a_{33} \end{bmatrix} + \tilde{\Xi}(t_i) \quad (10)$$

where  $\tilde{\Xi}(t_i)$  is a  $9 \times 1$  Gaussian vector with zero mean and a diagonal covariance matrix, with diagonal entries equal to the diagonal elements of  $W + 2V + \text{Cov}(\Omega)$ .

Now, we have,

$$\mathcal{G}(t_i) = \begin{bmatrix} \mathcal{G}_{11}(t_i) & \mathcal{G}_{12}(t_i) & \mathcal{G}_{13}(t_i) \\ \mathcal{G}_{21}(t_i) & \mathcal{G}_{22}(t_i) & \mathcal{G}_{23}(t_i) \\ \mathcal{G}_{31}(t_i) & \mathcal{G}_{32}(t_i) & \mathcal{G}_{33}(t_i) \end{bmatrix}$$

where

$$\mathcal{G}_{11}(t_i) = \begin{bmatrix} \gamma_1(t_{i-1})(t_i - t_{i-1}) & \gamma_2(t_{i-1})(t_i - t_{i-1}) & \gamma_3(t_{i-1})(t_i - t_{i-1}) \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$\mathcal{G}_{12}(t_i) = \begin{bmatrix} 0 & 0 & 0 \\ \gamma_1(t_{i-1})(t_i - t_{i-1}) & \gamma_2(t_{i-1})(t_i - t_{i-1}) & \gamma_3(t_{i-1})(t_i - t_{i-1}) \\ 0 & 0 & 0 \end{bmatrix}$$

$$\mathcal{G}_{13}(t_i) = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ \gamma_1(t_{i-1})(t_i - t_{i-1}) & \gamma_2(t_{i-1})(t_i - t_{i-1}) & \gamma_3(t_{i-1})(t_i - t_{i-1}) \end{bmatrix}$$

and

$$\mathcal{G}_{21}(t_i) = \begin{bmatrix} \gamma_1(t_{i-2})(t_{i-1} - t_{i-2}) & \gamma_2(t_{i-2})(t_{i-1} - t_{i-2}) & \gamma_3(t_{i-2})(t_{i-1} - t_{i-2}) \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$\mathcal{G}_{22}(t_i) = \begin{bmatrix} 0 & 0 & 0 \\ \gamma_1(t_{i-2})(t_{i-1} - t_{i-2}) & \gamma_2(t_{i-2})(t_{i-1} - t_{i-2}) & \gamma_3(t_{i-2})(t_{i-1} - t_{i-2}) \\ 0 & 0 & 0 \end{bmatrix}$$

-continued

$$\mathcal{G}_{23}(t_i) = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ \gamma_1(t_{i-2})(t_{i-1} - t_{i-2}) & \gamma_2(t_{i-2})(t_{i-1} - t_{i-2}) & \gamma_3(t_{i-2})(t_{i-1} - t_{i-2}) \end{bmatrix}$$

and

$$\mathcal{G}_{31}(t_i) = \begin{bmatrix} \gamma_1(t_{i-3})(t_{i-2} - t_{i-3}) & \gamma_2(t_{i-3})(t_{i-2} - t_{i-3}) & \gamma_3(t_{i-3})(t_{i-2} - t_{i-3}) \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

$$\mathcal{G}_{32}(t_i) = \begin{bmatrix} 0 & 0 & 0 \\ \gamma_1(t_{i-3})(t_{i-2} - t_{i-3}) & \gamma_2(t_{i-3})(t_{i-2} - t_{i-3}) & \gamma_3(t_{i-3})(t_{i-2} - t_{i-3}) \\ 0 & 0 & 0 \end{bmatrix}$$

$$\mathcal{G}_{33}(t_i) = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ \gamma_1(t_{i-3})(t_{i-2} - t_{i-3}) & \gamma_2(t_{i-3})(t_{i-2} - t_{i-3}) & \gamma_3(t_{i-3})(t_{i-2} - t_{i-3}) \end{bmatrix}$$

Now we use a discrete Kalman filter to estimate the parameters

$$\psi_i = \begin{bmatrix} a_{11} \\ a_{12} \\ a_{13} \\ a_{21} \\ a_{22} \\ a_{23} \\ a_{31} \\ a_{32} \\ a_{33} \end{bmatrix}_{t_i}$$

As shown in (3) and (10), the parameter dynamics equation and the parameter observation equation are

$$\psi_{i+1} = \psi_i + \Lambda_i$$

$$Z_i = G_i \psi_i + \tilde{\Xi}_i$$

where  $\Lambda$  is  $N(0, \lambda)$ ,  $\tilde{\Xi}_i$  is  $N(0, W + 2V + \text{Cov}(\Omega))$ ,  $G_i = g(t_i)$  as shown in (11), and the observation

$$Z_i = \begin{bmatrix} \delta y_1(t_i) - \gamma_1(t_{i-1}) \\ \delta y_2(t_i) - \gamma_2(t_{i-1}) \\ \delta y_3(t_i) - \gamma_3(t_{i-1}) \\ \delta y_1(t_{i-1}) - \gamma_1(t_{i-2}) \\ \delta y_2(t_{i-1}) - \gamma_2(t_{i-2}) \\ \delta y_3(t_{i-1}) - \gamma_3(t_{i-2}) \\ \delta y_1(t_{i-2}) - \gamma_1(t_{i-3}) \\ \delta y_2(t_{i-2}) - \gamma_2(t_{i-3}) \\ \delta y_3(t_{i-2}) - \gamma_3(t_{i-3}) \end{bmatrix}$$

as shown in (10).

The Discrete Kalman filter equations follow.

The predictor equation of the state estimate is

$$\psi_{i+1|f} = \psi_{i|f} \quad (12)$$



and the corrector equation with measurement is

$$\Psi_{i+1|i} = \Psi_{i+1|i} + K_i^{gain} (Z_{i+1} - \Psi_{i+1|i}) \quad (13)$$

where  $K_i^{gain}$  is the gain matrix.

The covariance matrix for the Kalman filter, denoted as  $\Sigma$ , is computed by the predictor equation,

$$\Sigma_{i+1|i} = \Sigma_{i|i} + \lambda \quad (14)$$

and the corrector equation,

$$\Sigma_{i+1|i+1} = [I - K_i^{gain}] \Sigma_{i+1|i} \quad (15)$$

The gain is

$$K_i^{gain} = \Sigma_{i+1|i} (\Sigma_{i+1|i} + W + 2V + Cov(\Omega))^{-1} \quad (16)$$

The initial condition is

$$\Sigma_{0|0} = 0_0 \quad (17)$$

FIG. 49 illustrates an architecture 4900 of the Parameter Learning Engine, with the Parameter Adaptive Engine (PAE) being a central component. The parameters of the model change slower than the state of the system.

[0165] As noted above, in some embodiments and situations, the internal state of an operating battery may be estimated, such as to estimate an internal temperature of the battery. Rather than managing battery temperature in the manner discussed herein, it is noted that prior systems, if they use temperature at all, typically merely shut down a battery's operation if the external temperature is too high, rather than control the battery's operations to manage the temperature. In addition, while not described in the example temperature model below, in some embodiments additional operations may be performed to control the internal temperature of the battery based at least in part on managing an external temperature surrounding the battery, such as by controlling the external temperature based at least in part on an estimated internal state of the battery.

[0166] Additional details regarding an example embodiment of using a battery temperature model to estimate such internal battery temperature are as follows:

[0167] The incremental dynamic behavior of the battery temperature (running at a higher speed) is given by a differential equation,

$$\delta \text{(?)}(t) = \text{(?)} \delta T(t - \tau) + v(t) \quad (1)$$

(?) indicates text missing or illegible when filed

[0168] where  $\tau$  is a time delay representing diffusion effect, and  $v(t)$  is a linear function of the incremental state vector  $\delta r(t)$  (power, voltage, current) and the incremental control  $\delta u(t)$  for the lower speed control tracker

$$v(t) = b \cdot \delta z(t) + c \cdot \delta u(t) + v(t) \quad (2)$$

[0169] The parameters a,b,c and  $\tau$  can be estimated by a non-linear least squares estimator using the historical data, and  $v(t)$  is a Gaussian noise.

[0170] The Laplace transform of (1) is:

$$\text{(?)} \delta T(s) = \alpha \text{(?)} \delta T(s) + v(s) \quad (3)$$

$$(s - \alpha \text{(?)}) \delta T(s) = v(s) \quad (4)$$

(?) indicates text missing or illegible when filed

[0171] We approximate the term **([text missing or illegible when filed])** in (4) with a Pade approximate [1] of the form

$$\frac{P(\text{?})}{Q(\text{?})}$$

(?) indicates text missing or illegible when filed

where P and Q are polynomials in s, and the order  $o(P) > o(Q)$ . The Laplace transform of (1) is approximated by

$$\frac{P(a)}{Q(a)} \delta T(a) = v(t) \quad (5)$$

[0172] For a low frequency approximation, we select a Pade approximate of the form

$$\frac{P(\text{?})}{Q(\text{?})}$$

(?) indicates text missing or illegible when filed

with

$$P(s) = \text{(?)} + \alpha_1 s + \text{(?)} \quad (6)$$

$$Q(s) = \beta_1 s + \text{(?)} \quad (7)$$

(?) indicates text missing or illegible when filed

[0173] We define

$$w(s) = P(s) \delta T(s) \quad (8)$$

[0174] Replacing P(s) in (8) with (6), yields

$$\text{(?)} \delta T(s) + \alpha_1 \text{(?)} \delta T(s) + \text{(?)} \delta T(s) = w(s) \quad (9)$$

(?) indicates text missing or illegible when filed

[0175] The inverse Laplace transform of (9) is

$$\delta \text{(?)}(t) + \alpha_1 \delta \text{(?)}(t) + \text{(?)} \delta T(t) = w(t) \quad (10)$$

(?) indicates text missing or illegible when filed

Define **[text missing or illegible when filed]**, we have

$$\textcircled{?} = \textcircled{?} = \textcircled{?} - \textcircled{?} + \textcircled{?} = \textcircled{?} - \textcircled{?} + \textcircled{?} \quad (11)$$

$$\begin{bmatrix} \textcircled{?} \\ \textcircled{?} \end{bmatrix} = \begin{bmatrix} \textcircled{?} \\ \textcircled{?} \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \end{bmatrix} v(t)$$

⓪ indicates text missing or illegible when filed

From (5) and (8), we also have

$$w(s) = Q(s)\textcircled{?}(s) \quad (12)$$

⓪ indicates text missing or illegible when filed

Replacing Q(s) in (12) with (7), yields

$$w(s) = \beta_1 \textcircled{?}(s) + \textcircled{?}(s) \quad (13)$$

⓪ indicates text missing or illegible when filed

The inverse Laplace transform of (13) is:

$$w(t) = \beta_1 \textcircled{?}(t) + \textcircled{?}(t) \quad (14)$$

⓪ indicates text missing or illegible when filed

Replacing v(t) in (14) with (2), yields

$$w(t) = \beta_1 b \cdot \textcircled{?}(t) + \beta_1 c \cdot \textcircled{?}(t) + \textcircled{?} \cdot \textcircled{?}(t) + \textcircled{?} c \cdot \textcircled{?}(t) + v_1(t) \quad (15)$$

⓪ indicates text missing or illegible when filed

where **[text missing or illegible when filed]** is a Gaussian noise.

In (15), we set the values of **[text missing or illegible when filed]** to be constant during a unit time interval of the lower speed controller, in:

$$\textcircled{?}(t) = \textcircled{?}(\textcircled{?}), \textcircled{?}(t) = \textcircled{?}(\textcircled{?}), \textcircled{?}(t) = \textcircled{?}(\textcircled{?}) \quad (16)$$

⓪ indicates text missing or illegible when filed

and let **[text missing or illegible when filed]** to be the control variable for the light speed controller, yielding

$$\begin{bmatrix} \textcircled{?} \\ \textcircled{?} \end{bmatrix} = \begin{bmatrix} \textcircled{?} \\ \textcircled{?} \end{bmatrix} + \begin{bmatrix} \textcircled{?} \\ \textcircled{?} \end{bmatrix} + \begin{bmatrix} \textcircled{?} \\ \textcircled{?} \end{bmatrix} \quad (17)$$

⓪ indicates text missing or illegible when filed

The state equation in (17) one be rewritten as:

$$\textcircled{?} = \textcircled{?} + \textcircled{?} + f + \Omega(c) \quad (18)$$

-continued

with

$$A = \begin{bmatrix} \textcircled{?} \\ \textcircled{?} \end{bmatrix},$$

$$B = \begin{bmatrix} \textcircled{?} \\ \textcircled{?} \end{bmatrix}$$

and

$$f = \begin{bmatrix} \textcircled{?} \\ \textcircled{?} \end{bmatrix}.$$

⓪ indicates text missing or illegible when filed

and  $\Omega(t)$  is a Gaussian noise.

The optimal control tracking problems for the high speed controller is summarized as:

$$\textcircled{?} E\textcircled{?} + \textcircled{?} dt \quad (19)$$

s.t.

$$\textcircled{?} = \textcircled{?} + \textcircled{?} + f + \textcircled{?} \quad (20)$$

⓪ indicates text missing or illegible when filed

By the separation principle, we can get the incremental control for the high speed controller.

$$\textcircled{?}(t) = K\textcircled{?}(t) + \textcircled{?} \quad (21)$$

⓪ indicates text missing or illegible when filed

where **[text missing or illegible when filed]** is the state estimate by running a discrete Kalman filter, K and  $\Phi$  are obtained by solving the deterministic trading problem.

**[0176]** In the end, we take the linear combination of the two incremental controls **[text missing or illegible when filed]** to get the incremental control  $\delta u(t)$ ,

$$\delta u(t) = \textcircled{?} \textcircled{?}(t) + \textcircled{?} \textcircled{?}(t) \quad (22)$$

⓪ indicates text missing or illegible when filed

**[0177]** where **[text missing or illegible when filed]** = 1 and  $0 \leq \textcircled{?} \leq 1$ .

**[0178]** Notes: the (2,1) pade approximation for

$$\textcircled{?} = \textcircled{?}$$

⓪ indicates text missing or illegible when filed

then the (2,1) pade approximation for **[text missing or illegible when filed]** is

$$\frac{P(a)}{Q(a)} = \quad (23)$$

$$\textcircled{?} \frac{\textcircled{?} + 4\textcircled{?}}{-2(\textcircled{?}) + 6} = \frac{(\textcircled{?})\textcircled{?} + (\textcircled{?})\textcircled{?} - 6a}{\textcircled{?} + \textcircled{?}} = \frac{\textcircled{?} + \textcircled{?} - \textcircled{?}}{\textcircled{?} + \textcircled{?}}$$

⓪ indicates text missing or illegible when filed



[0179] From (6) and (7), we have

$$\begin{aligned} \textcircled{?} &= \textcircled{?}, & (24) \\ \textcircled{?} &= \textcircled{?}, \\ \textcircled{?} &= \textcircled{?}, \\ \textcircled{?} &= \textcircled{?} \end{aligned}$$

⓪ indicates text missing or illegible when filed

[0180] In general, if P and Q are on an order of m and n respectively, (11) and (14) can be generalized as follows

$$\begin{bmatrix} \textcircled{?} \\ \textcircled{?} \\ \vdots \\ \textcircled{?} \\ \textcircled{?} \end{bmatrix} = \begin{bmatrix} 0 & 1 & 0 & \dots & 0 \\ 0 & 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \dots & 1 \\ \textcircled{?} & \textcircled{?} & \textcircled{?} & \dots & \textcircled{?} \end{bmatrix} \begin{bmatrix} \textcircled{?} \\ \textcircled{?} \\ \vdots \\ \textcircled{?} \\ \textcircled{?} \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 1 \end{bmatrix} \textcircled{?} \quad (25)$$

$$\textcircled{?} = \textcircled{?}(t) + \textcircled{?} + \dots + \textcircled{?} \quad (26)$$

⓪ indicates text missing or illegible when filed

[0181] FIG. 51 is a block diagram illustrating an example architecture 5100 for a battery control system. The example architecture includes, for example, the following components:

[0182] a Boost/Buck Converter, such as a DC-to-DC converter used to match the average pulsed output to the commanded power from the battery controller;

[0183] an Inverter/Rectifier, such as with the power inverter converting DC to AC when battery is discharging, and with the rectifier converting AC to DC when battery is charging;

[0184] a A/D converter, such as an analog-to-digital converter that converts a continuous physical quantity to a digital number that represents the quantity's amplitude;

[0185] a D/A converter, such as a digital-to-analog converter that converts digital data (usually binary) into an analog signal (current, voltage, or electric charge);

[0186] a Battery Controller, such as an incremental feedback hybrid tracking control system (e.g., that includes a state estimator, a parameter learning engine, and a state and state rate tracker controller, as discussed in greater detail elsewhere herein) that drives/controls the Boost/Buck converter so that the battery operates in near resonance with a desired power signal generated by a CDI agent;

[0187] a CDI Agent, such as to infer battery parameters and to use longevity and performance rules to generate the tracking signal for the desired power signal; and

[0188] a Demand Forecaster, such as to generate the desired output power/charge power forecast of the battery.

[0189] FIG. 46 is a block diagram illustrating example computing systems suitable for performing techniques for implementing automated control systems to control or otherwise manipulate at least some operations of specified physical systems or other target systems in configured

manners, such as to control physical target systems having one or more batteries by using characteristics of each battery's state to perform automated control of DC power that is provided from the battery (e.g., in a real-time manner and to optimize long-term operation of the battery), such as in the matter discussed above with respect to FIGS. 30-45 and 47-51 and elsewhere herein. In particular, FIG. 46 illustrates a server computing system 4600 suitable for providing at least some functionality of a CDD system, although in other embodiments multiple computing systems may be used for the execution (e.g., to have distinct computing systems executing the CDD Decision Module Construction component for initial configuration and setup before run-time control occurs, and one or more copies of the CDD Control Action Determination component 4644 and/or the CDD Coordinated Control Managements component 4646 for the actual run-time control). FIG. 46 also illustrates various client computer systems 4650 that may be used by customers or other users of the CDD system 4640, as well as one or more target systems with batteries to be controlled (in this example, target system 1 4660 and target system 2 4670, which are accessible to the CDD system 4640 over one or more computer networks 4690).

[0190] The server computing system 4600 has components in the illustrated embodiment that include one or more hardware CPU ("central processing unit") computer processors 4605, various I/O ("input/output") hardware components 4610, storage 4620, and memory 4630. The illustrated I/O components include a display 4611, a network connection 4612, a computer-readable media drive 4613, and other I/O devices 4615 (e.g., a keyboard, a mouse, speakers, etc.). In addition, the illustrated client computer systems 4650 may each have components similar to those of server computing system 4600, including one or more CPUs 4651, I/O components 4652, storage 4654, and memory 4657, although some details are not illustrated for the computing systems 4650 for the sake of brevity. The target systems 4660 and 4670 may also each include one or more computing systems (not shown) having components that are similar to some or all of the components illustrated with respect to server computing system 4600, but such computing systems and components are not illustrated in this example for the sake of brevity.

[0191] The CDD system 4640 is executing in memory 4630 and includes components 4642-4646, and in some embodiments the system and/or components each includes various software instructions that when executed program one or more of the CPU processors 4605 to provide an embodiment of a CDD system as described elsewhere herein. The CDD system 4640 may interact with computing systems 4650 over the network 4690 (e.g., via the Internet and/or the World Wide Web, via a private cellular network, etc.), as well as the target systems 4660 and 4670 in this example. In this example embodiment, the CDD system includes functionality related to generating and deploying decision modules in configured manners for customers or other users, as discussed in greater detail elsewhere herein. The other computing systems 4650 may also be executing various software as part of interactions with the CDD system 4640 and/or its components. For example, client computing systems 4650 may be executing software in memory 4657 to interact with CDD system 4640 (e.g., as part of a Web browser, a specialized client-side application program, etc.), such as to interact with one or more interfaces (not shown)



of the CDD system **4640** to configure and deploy automated control systems (e.g., stored automated control systems **4625** that were previously created by the CDD system **4640** for use in controlling one or more physical target systems with batteries) or other decision modules **4629**, as well as to perform various other types of actions, as discussed in greater detail elsewhere. Various information related to the functionality of the CDD system **4640** may be stored in storage **4620**, such as information **4621** related to users of the CDD system (e.g., account information), and information **4623** related to one or more target systems that have batteries to be controlled.

[0192] It will be appreciated that computing systems **4600** and **4650** and target systems **4660** and **4670** are merely illustrative and are not intended to limit the scope of the present invention. The computing systems may instead each include multiple interacting computing systems or devices, and the computing systems/nodes may be connected to other devices that are not illustrated, including through one or more networks such as the Internet, via the Web, or via private networks (e.g., mobile communication networks, etc.). More generally, a computing node or other computing system or device may comprise any combination of hardware that may interact and perform the described types of functionality, including without limitation desktop or other computers, database servers, network storage devices and other network devices, PDAs, cell phones, wireless phones, pagers, electronic organizers, Internet appliances, television-based systems (e.g., using set-top boxes and/or personal/digital video recorders), and various other consumer products that include appropriate communication capabilities. In addition, the functionality provided by the illustrated CDD system **4640** and its components may in some embodiments be distributed in additional components. Similarly, in some embodiments some of the functionality of the CDD system **4640** and/or CDD components **4642-4646** may not be provided and/or other additional functionality may be available.

[0193] It will also be appreciated that, while various items are illustrated as being stored in memory or on storage while being used, these items or portions of them may be transferred between memory and other storage devices for purposes of memory management and data integrity. Alternatively, in other embodiments some or all of the software modules and/or systems may execute in memory on another device and communicate with the illustrated computing systems via inter-computer communication. Thus, in some embodiments, some or all of the described techniques may be performed by hardware means that include one or more processors and/or memory and/or storage when configured by one or more software programs (e.g., by the CDD system **4640** and/or the CDD components **4642-4646**) and/or data structures, such as by execution of software instructions of the one or more software programs and/or by storage of such software instructions and/or data structures. Furthermore, in some embodiments, some or all of the systems and/or components may be implemented or provided in other manners, such as by using means that are implemented at least partially or completely in firmware and/or hardware, including, but not limited to, one or more application-specific integrated circuits (ASICs), standard integrated circuits, controllers (e.g., by executing appropriate instructions, and including microcontrollers and/or embedded controllers), field-programmable gate arrays (FPGAs), complex programmable logic devices (CPLDs), etc. Some or all of

the components, systems and data structures may also be stored (e.g., as software instructions or structured data) on a non-transitory computer-readable storage medium, such as a hard disk or flash drive or other non-volatile storage device, volatile or non-volatile memory (e.g., RAM), a network storage device, or a portable media article to be read by an appropriate drive (e.g., a DVD disk, a CD disk, an optical disk, etc.) or via an appropriate connection. The systems, components and data structures may also in some embodiments be transmitted as generated data signals (e.g., as part of a carrier wave or other analog or digital propagated signal) on a variety of computer-readable transmission mediums, including wireless-based and wired/cable-based mediums, and may take a variety of forms (e.g., as part of a single or multiplexed analog signal, or as multiple discrete digital packets or frames). Such computer program products may also take other forms in other embodiments. Accordingly, the present invention may be practiced with other computer system configurations.

[0194] For illustrative purposes, some additional details are included below regarding some embodiments in which specific types of operations are performed in specific manners, including with respect to particular types of target systems and for particular types of control activities determined in particular manners. These examples are provided for illustrative purposes and are simplified for the sake of brevity, and the inventive techniques may be used in a wide variety of other situations, including in other environments and with other types of automated control action determination techniques, some of which are discussed below.

[0195] FIG. 1 is a network diagram illustrating an example environment in which a system for performing cooperative distributed control of one or more target systems may be configured and initiated. In particular, an embodiment of a CDD system **140** is executing on one or more computing systems **190**, including in the illustrated embodiment to operate in an online manner and provide a graphical user interface (GUI) (not shown) and/or other interfaces **119** to enable one or more remote users of client computing systems **110** to interact over one or more intervening computer networks **100** with the CDD system **140** to configure and create one or more decision modules to include as part of an automated control system to use with each of one or more target systems to be controlled.

[0196] In particular, target system **1160** and target system **2170** are example target systems illustrated in this example, although it will be appreciated that only one target system or numerous target systems may be available in particular embodiments and situations, and that each such target system may include a variety of mechanical, electronic, chemical, biological, and/or other types of components to implement operations of the target system in a manner specific to the target system. In this example, the one or more users (not shown) may interact with the CDD system **140** to generate an example automated control system **122** for target system **1**, with the automated control system including multiple decision modules **124** in this example that will cooperatively interact to control portions of the target system **1160** when later deployed and implemented. The process of the users interacting with the CDD system **140** to create the automated control system **122** may involve a variety of interactions over time, including in some cases independent actions of different groups of users, as discussed in greater detail elsewhere. In addition, as part of the process of creating



and/or training or testing automated control system **122**, it may perform one or more interactions with the target system **1** as illustrated, such as to obtain partial initial state information, although some or all training activities may in at least some embodiments include simulating effects of control actions in the target system **1** without actually implementing those control actions at that time.

[0197] After the automated control system **122** is created, the automated control system may be deployed and implemented to begin performing operations involving controlling the target system **1** **160**, such as by optionally executing the automated control system **122** on the one or more computing systems **190** of the CDD system **140**, so as to interact over the computer networks **100** with the target system **1**. In other embodiments and situations, the automated control system **122** may instead be deployed by executing local copies of some or all of the automated control system **122** (e.g., one or more of the multiple decision modules **124**) in a manner local to the target system **1**, as illustrated with respect to a deployed copy **121** of some or all of automated control system **1**, such as on one or more computing systems (not shown) that are part of the target system **1**.

[0198] In a similar manner to that discussed with respect to automated control system **122**, one or more users (whether the same users, overlapping users, or completely unrelated users to those that were involved in creating the automated control system **122**) may similarly interact over the computer network **100** with the CDD system **140** to create a separate automated control system **126** for use in controlling some or all of the target system **2** **170**. In this example, the automated control system **126** for target system **2** includes only a single decision module **128** that will perform all of the control actions for the automated control system **126**. The automated control system **126** may similarly be deployed and implemented for target system **2** in a manner similar to that discussed with respect to automated control system **122**, such as to execute locally on the one or more computing systems **190** and/or on one or more computing systems (not shown) that are part of the target system **2**, although a deployed copy of automated control system **2** is not illustrated in this example. It will be further appreciated that the automated control systems **122** and/or **126** may further include other components and/or functionality that are separate from the particular decision modules **124** and **128**, respectively, although such other components and/or functionality are not illustrated in FIG. 1.

[0199] The network **100** may, for example, be a publicly accessible network of linked networks, possibly operated by various distinct parties, such as the Internet, with the CDD system **140** available to any users or only certain users over the network **100**. In other embodiments, the network **100** may be a private network, such as, for example, a corporate or university network that is wholly or partially inaccessible to non-privileged users. In still other embodiments, the network **100** may include one or more private networks with access to and/or from the Internet. Thus, while the CDD system **140** in the illustrated embodiment is implemented in an online manner to support various users over the one or more computer networks **100**, in other embodiments a copy of the CDD system **140** may instead be implemented in other manners, such as to support a single user or a group of related users (e.g., a company or other organization), such as if the one or more computer networks **100** are instead an

internal computer network of the company or other organization, and with such a copy of the CDD system optionally not being available to other users external to the company or other organizations. The online version of the CDD system **140** and/or local copy version of the CDD system **140** may in some embodiments and situations operate in a fee-based manner, such that the one or more users provide various fees to use various operations of the CDD system, such as to perform interactions to generate decision modules and corresponding automated control systems, and/or to deploy or implement such decision modules and corresponding automated control systems in various manners. In addition, the CDD system **140**, each of its components (including component **142** and optional other components **117**, such as one or more CDD Control Action Determination components and/or one or more CDD Coordinated Control Management components), each of the decision modules, and/or each of the automated control systems may include software instructions that execute on one or more computing systems (not shown) by one or more processors (not shown), such as to configure those processors and computing systems to operate as specialized machines with respect to performing their programmed functionality.

[0200] FIG. 2 is a network diagram illustrating an example environment in which a system for performing cooperative distributed control of target systems may be implemented, and in particular continues the examples discussed with respect to FIG. 1. In the example environment of FIG. 2, target system **1** **160** is again illustrated, with the automated control system **122** now being deployed and implemented to use in actively controlling the target system **1** **160**. In the example of FIG. 2, the decision modules **124** are represented as individual decision modules **124a**, **124b**, etc., to **124n**, and may be executing locally to the target system **1** **160** and/or in a remote manner over one or more intervening computer networks (not shown). In the illustrated example, each of the decision modules **124** includes a local copy of a CDD Control Action Determination component **144**, such as with component **144a** supporting its local decision module **124a**, component **144b** supporting its local decision module **124b**, and component **144n** supporting its local decision module **124n**. Similarly, the actions of the various decision modules **124** are coordinated and synchronized in a peer-to-peer manner in the illustrated embodiment, with each of the decision modules **124** including a copy of a CDD Coordinated Control Management component **146** to perform such synchronization, with component **146a** supporting its local decision module **124a**, component **146b** supporting its local decision module **124b**, and component **146n** supporting its local decision module **124n**.

[0201] As the decision modules **124** and automated control system **122** execute, various interactions **175** between the decision modules **124** are performed, such as to share information about current models and other state of the decision modules to enable cooperation and coordination between various decision modules, such as for a particular decision module to operate in a partially synchronized consensus manner with respect to one or more other decision modules (and in some situations in a fully synchronized manner in which the consensus actions of all of the decision modules **124** converge). During operation of the decision modules **124** and automated control system **122**, various state information **143** may be obtained by the automated control system **122** from the target system **160**, such as



initial state information and changing state information over time, and including outputs or other results in the target system **1** from control actions performed by the decision modules **124**.

[0202] The target system **1** in this example includes various control elements **161** that the automated control system **122** may manipulate, and in this example each decision module **124** may have a separate group of one or more control elements **161** that it manipulates (such that decision module A **124a** performs interactions **169a** to perform control actions A **147a** on control elements A **161a**, decision module B **124b** performs interactions **169b** to perform control actions B **147b** on control elements B **161b**, and decision module N **124n** performs interactions **169n** to perform control actions N **147n** on control elements N **161n**). Such control actions affect the internal state **163** of other elements of the target system **1**, including optionally to cause or influence one or more outputs **162**. As operation of the target system **1** is ongoing, at least some of the internal state information **163** is provided to some or all of the decision modules to influence their ongoing control actions, with each of the decision modules **124a-124n** possibly having a distinct set of state information **143a-143n**, respectively, in this example.

[0203] As discussed in greater detail elsewhere, each decision module **124** may use such state information **143** and a local model **145** of the decision module for the target system to determine particular control actions **147** to next perform, such as for each of multiple time periods, although in other embodiments and situations, a particular automated control system may perform interactions with a particular target system for only one time period or only for some time periods. For example, the local CDD Control Action Determination component **144** for a decision module **124** may determine a near-optimal location solution for that decision module's local model **145**, and with the local CDD Coordinated Control Management component **146** determining a synchronized consensus solution to reflect other of the decision modules **124**, including to update the decision module's local model **145** based on such local and/or synchronized solutions that are determined. Thus, during execution of the automated control system **122**, the automated control system performs various interactions with the target system **160**, including to request state information, and to provide instructions to modify values of or otherwise manipulate control elements **161** of the target system **160**. For example, for each of multiple time periods, decision module **124a** may perform one or more interactions **169a** with one or more control elements **161a** of the target system, while decision module **124b** may similarly perform one or more interactions **169b** with one or more separate control elements B **161b**, and decision module **124n** may perform one or more interactions **169n** with one or more control elements N **161n** of the target system **160**. In other embodiments and situations, at least some control elements may not perform control actions during each time period.

[0204] While example target system **2** **170** is not illustrated in FIG. 2, further details are illustrated for decision module **128** of automated control system **126** for reference purposes, although such a decision module **128** would not typically be implemented together with the decision modules **124** controlling target system **1**. In particular, the deployed copy of automated control system **126** includes only the single executing decision module **128** in this

example, although in other embodiments the automated control system **126** may include other components and functionality. In addition, since only a single decision module **128** is implemented for the automated control system **126**, the decision module **128** includes a local CDD Control Action Determination component **244**, but does not in the illustrated embodiment include any local CDD Coordinated Control Management component, since there are not other decision modules with which to synchronize and interact.

[0205] While not illustrated in FIGS. 1 and 2, the distributed nature of operations of automated control systems such as those of **122** allow partially decoupled operations of the various decision modules, include to allow modifications to the group of decision modules **124** to be modified over time while the automated control system **122** is in use, such as to add new decision modules **124** and/or to remove existing decision modules **124**. In a similar manner, changes may be made to particular decision modules **124** and/or **128**, such as to change rules or other restrictions specific to a particular decision module and/or to change goals specific to a particular decision module over time, with a new corresponding model being generated and deployed within such a decision module, including in some embodiments and situations while the corresponding automated control system continues control operations of a corresponding target system. In addition, while each automated control system is described as controlling a single target system in the examples of FIGS. 1 and 2, in other embodiments and situations, other configurations may be used, such as for a single automated control system to control multiple target systems (e.g., multiple inter-related target systems, multiple target systems of the same type, etc.), and/or multiple automated control systems may operate to control a single target system, such as by each operating independently to control different portions of that target control system. It will be appreciated that other configurations may similarly be used in other embodiments and situations.

[0206] FIG. 3 is a block diagram illustrating example computing systems suitable for performing techniques for implementing automated control systems to control or otherwise manipulate at least some operations of specified physical systems or other target systems in configured manners. In particular, FIG. 3 illustrates a server computing system **300** suitable for providing at least some functionality of a CDD system, although in other embodiments multiple computing systems may be used for the execution (e.g., to have distinct computing systems executing the CDD Decision Module Construction component for initial configuration and setup before run-time control occurs, and one or more copies of the CDD Control Action Determination component **344** and/or the CDD Coordinated Control Managements component **346** for the actual run-time control). FIG. 3 also illustrates various client computer systems **350** that may be used by customers or other users of the CDD system **340**, as well as one or more target systems (in this example, target system **1** **360** and target system **2** **370**, which are accessible to the CDD system **340** over one or more computer networks **390**).

[0207] The server computing system **300** has components in the illustrated embodiment that include one or more hardware CPU ("central processing unit") computer processors **305**, various I/O ("input/output") hardware components **310**, storage **320**, and memory **330**. The illustrated I/O components include a display **311**, a network connection



**312**, a computer-readable media drive **313**, and other I/O devices **315** (e.g., a keyboard, a mouse, speakers, etc.). In addition, the illustrated client computer systems **350** may each have components similar to those of server computing system **300**, including one or more CPUs **351**, I/O components **352**, storage **354**, and memory **357**, although some details are not illustrated for the computing systems **350** for the sake of brevity. The target systems **360** and **370** may also each include one or more computing systems (not shown) having components that are similar to some or all of the components illustrated with respect to server computing system **300**, but such computing systems and components are not illustrated in this example for the sake of brevity.

[0208] The CDD system **340** is executing in memory **330** and includes components **342-346**, and in some embodiments the system and/or components each includes various software instructions that when executed program one or more of the CPU processors **305** to provide an embodiment of a CDD system as described elsewhere herein. The CDD system **340** may interact with computing systems **350** over the network **390** (e.g., via the Internet and/or the World Wide Web, via a private cellular network, etc.), as well as the target systems **360** and **370** in this example. In this example embodiment, the CDD system includes functionality related to generating and deploying decision modules in configured manners for customers or other users, as discussed in greater detail elsewhere herein. The other computing systems **350** may also be executing various software as part of interactions with the CDD system **340** and/or its components. For example, client computing systems **350** may be executing software in memory **357** to interact with CDD system **340** (e.g., as part of a Web browser, a specialized client-side application program, etc.), such as to interact with one or more interfaces (not shown) of the CDD system **340** to configure and deploy automated control systems (e.g., stored automated control systems **325** that were previously created by the CDD system **340**) or other decision modules **329**, as well as to perform various other types of actions, as discussed in greater detail elsewhere. Various information related to the functionality of the CDD system **340** may be stored in storage **320**, such as information **321** related to users of the CDD system (e.g., account information), and information **323** related to one or more target systems.

[0209] It will be appreciated that computing systems **300** and **350** and target systems **360** and **370** are merely illustrative and are not intended to limit the scope of the present invention. The computing systems may instead each include multiple interacting computing systems or devices, and the computing systems/nodes may be connected to other devices that are not illustrated, including through one or more networks such as the Internet, via the Web, or via private networks (e.g., mobile communication networks, etc.). More generally, a computing node or other computing system or device may comprise any combination of hardware that may interact and perform the described types of functionality, including without limitation desktop or other computers, database servers, network storage devices and other network devices, PDAs, cell phones, wireless phones, pagers, electronic organizers, Internet appliances, television-based systems (e.g., using set-top boxes and/or personal/digital video recorders), and various other consumer products that include appropriate communication capabilities. In addition, the functionality provided by the illustrated CDD system **340** and its components may in some embodiments be distrib-

uted in additional components. Similarly, in some embodiments some of the functionality of the CDD system **340** and/or CDD components **342-346** may not be provided and/or other additional functionality may be available.

[0210] It will also be appreciated that, while various items are illustrated as being stored in memory or on storage while being used, these items or portions of them may be transferred between memory and other storage devices for purposes of memory management and data integrity. Alternatively, in other embodiments some or all of the software modules and/or systems may execute in memory on another device and communicate with the illustrated computing systems via inter-computer communication. Thus, in some embodiments, some or all of the techniques may be performed by hardware means that include one or more processors and/or memory and/or storage when configured by one or more software programs (e.g., by the CDD system **340** and/or the CDD components **342-346**) and/or data structures, such as by execution of software instructions of the one or more software programs and/or by storage of such software instructions and/or data structures. Furthermore, in some embodiments, some or all of the systems and/or components may be implemented or provided in other manners, such as by using means that are implemented at least partially or completely in firmware and/or hardware, including, but not limited to, one or more application-specific integrated circuits (ASICs), standard integrated circuits, controllers (e.g., by executing appropriate instructions, and including microcontrollers and/or embedded controllers), field-programmable gate arrays (FPGAs), complex programmable logic devices (CPLDs), etc. Some or all of the components, systems and data structures may also be stored (e.g., as software instructions or structured data) on a non-transitory computer-readable storage medium, such as a hard disk or flash drive or other non-volatile storage device, volatile or non-volatile memory (e.g., RAM), a network storage device, or a portable media article to be read by an appropriate drive (e.g., a DVD disk, a CD disk, an optical disk, etc.) or via an appropriate connection. The systems, components and data structures may also in some embodiments be transmitted as generated data signals (e.g., as part of a carrier wave or other analog or digital propagated signal) on a variety of computer-readable transmission mediums, including wireless-based and wired/cable-based mediums, and may take a variety of forms (e.g., as part of a single or multiplexed analog signal, or as multiple discrete digital packets or frames). Such computer program products may also take other forms in other embodiments. Accordingly, the present invention may be practiced with other computer system configurations.

[0211] FIG. 4 is a flow diagram of an example embodiment of a Collaborative Distributed Decision (CDD) system routine **400**. The routine may, for example, be provided by execution of the CDD system **340** of FIG. 3 and/or the CDD system **140** of FIG. 1, such as to provide functionality to construct and implement automated control systems for specified target systems.

[0212] The illustrated embodiment of the routine begins at block **410**, where information or instructions are received. If it is determined in block **420** that the information or instructions of block **410** include an indication to create or revise one or more decision modules for use as part of an automated control system for a particular target system, the routine continues to block **425** to initiate execution of a



Decision Module Construction component, and in block **430** obtains and stores one or more resulting decision modules for the target system that are created in block **425**. One example of a routine for such a Decision Module Construction component is discussed in greater detail with respect to FIGS. **5A-5B**.

[**0213**] After block **430**, or if it is instead determined in block **420** that the information or instructions received in block **410** are not to create or revise one or more decision modules, the routine continues to block **440** to determine whether the information or instructions received in block **410** indicate to deploy one or more created decision modules to control a specified target system, such as for one or more decision modules that are part of an automated control system for that target system. The one or more decision modules to deploy may have been created immediately prior with respect to block **425**, such that the deployment occurs in a manner that is substantially simultaneous with the creation, or in other situations may include one or more decision modules that were created at a previous time and stored for later use. If it is determined to deploy one or more such decision modules for such a target system, the routine continues to block **450** to initiate the execution of those one or more decision modules for that target system, such as on one or more computing systems local to an environment of the target system, or instead on one or more remote computing systems that communicate with the target system over one or more intermediary computer networks (e.g., one or more computing systems under control of a provider of the CDD system).

[**0214**] After block **450**, the routine continues to block **460** to determine whether to perform distributed management of multiple decision modules being deployed in a manner external to those decision modules, such as via one or more centralized Coordinated Control Management components. If so, the routine continues to block **465** to initiate execution of one or more such centralized CDD Coordinated Control Management components for use with those decision modules. After block **465**, or if it is instead determined in block **460** to not perform such distributed management in an external manner (e.g., if only one decision module is executed, if multiple decision modules are executed but coordinate their operations in a distributed peer-to-peer manner, etc.), the routine continues to block **470** to optionally obtain and store information about the operations of the one or more decision modules and/or resulting activities that occur in the target system, such as for later analysis and/or reporting.

[**0215**] If it is instead determined in block **440** that the information or instructions received in block **410** are not to deploy one or more decision modules, the routine continues instead to block **485** to perform one or more other indicated operations if appropriate. For example, such other authorized operations may include obtaining results information about the operation of a target system in other manners (e.g., by monitoring outputs or other state information for the target system), analyzing results of operations of decision modules and/or activities of corresponding target systems, generating reports or otherwise providing information to users regarding such operations and/or activities, etc. In addition, in some embodiments the analysis of activities of a particular target system over time may allow patterns to be identified in operation of the target system, such as to allow a model of that target system to be modified accordingly

(whether manually or in an automated learning manner) to reflect those patterns and to respond based on them. In addition, as discussed in greater detail elsewhere, distributed operation of multiple decision modules for an automated control system in a partially decoupled manner allows various changes to be made while the automated control system is in operation, such as to add one or more new decision modules, to remove one or more existing decision modules, to modify the operation of a particular decision module (e.g., by changing rules or other information describing the target system that is part of a model for the decision module), etc. In addition, the partially decoupled nature of multiple such decision modules in an automated control system allows one or more such decision modules to operate individually at times, such as if network communication issues or other problems prevent communication between multiple decision modules that would otherwise allow their individualized control actions to be coordinated in such situations, some or all such decision modules may continue to operate in an individualized manner, such as to provide useful ongoing control operations for a target system even if optimal or near-optimal solutions cannot be identified from coordination and synchronization between a group of multiple decision modules that collectively provide the automated control system for the target system.

[**0216**] After blocks **470** or **485**, the routine continues to block **495** to determine whether to continue, such as until an explicit indication to terminate is received. If it is determined to continue, the routine returns to block **410**, and otherwise continues to block **499** and ends.

[**0217**] FIGS. **5A-5B** illustrate a flow diagram of an example embodiment of a CDD Decision Module Construction routine **500**. The routine may, for example, be provided by execution of the component **342** of FIG. **3** and/or the component **142** of FIG. **1**, such as to provide functionality to allow users to provide information describing a target system of interest, and to perform corresponding automated operations to construct one or more decision modules to use to control the target system in specified manners. While the illustrated embodiment of the routine interacts with users in particular manners, such as via a displayed GUI (graphical user interface), it will be appreciated that other embodiments of the routine may interact with users in other manners, such as via a defined API (application programming interface) that an executing program invokes on behalf of a user. In some embodiments, the routine may further be implemented as part of an integrated development environment or other software tool that is available for one or more users to use, such as by implementing an online interface that is available to a variety of remote users over a public network such as the Internet, while in other embodiments a copy of the CDD system and/or particular CDD components may be used to support a single organization or other group of one or more users, such as by being executed on computing systems under the control of the organization or group. In addition, the CDD Decision Module Construction component may in some embodiments and situations be separated into multiple sub-components, such as a rules editor component that users interact with to specify rules and other description information for a target system, and a rules compiler engine that processes the user-specified rules and other information to create one or more corresponding decision modules.

[**0218**] The illustrated embodiment of the routine **500** begins at block **510**, where the routine provides or updates



a displayed user interface to one or more users, such as via a request received at an online version of component that is implementing the routine, or instead based on the routine being executed by one or more such users on computing systems that they control. While various operations are shown in the illustrated embodiment of the routine as occurring in a serial manner for the purpose of illustration, it will be appreciated that user interactions with such a user interface may occur in an iterative manner and/or over multiple periods of time and/or user sessions, including to update a user interface previously displayed to a user in various manners (e.g., to reflect a user action, to reflect user feedback generated by operation of the routine or from another component, etc.), as discussed further below.

[0219] After block 510, the routine continues to block 520 to receive information from one or more such users describing a target system to be controlled, including information about a plurality of elements of the target system that include one or more manipulatable control elements and optionally one or more outputs that the control elements affect, information about rules that specify restrictions involving the elements, information about state information that will be available during controlling of the system (e.g., values of particular elements or other state variables), and one or more goals to achieve during the controlling of the target system. It will be appreciated that such information may be obtained over a period of time from one or more users, including in some embodiments for a first group of one or more users to supply some information related to a target system and for one or more other second groups of users to independently provide other information about the target system, such as to reflect different areas of expertise of the different users and/or different parts of the target system.

[0220] After block 520, the routine continues to block 525 to identify any errors that have been received in the user input, and to prompt the user(s) to correct those errors, such as by updating the display in a corresponding manner as discussed with respect to block 510. While the identification of such errors is illustrated as occurring after the receiving of the information in block 520, it will be appreciated that some or all such errors may instead be identified as the users are inputting information into the user interface, such as to identify syntax errors in rules or other information that the users specify. After block 525, the illustrated embodiment of the routine continues to block 530 to optionally decompose the information about the target system into multiple subsets that each correspond to a portion of the target system, such as with each subset having one or more different control elements that are manipulatable by the automated control system being created by the routine, and optionally have overlapping or completely distinct goals and/or sets of rules and other information describing the respective portions of the target system. As discussed in greater detail elsewhere, such decomposition, if performed, may in some situations be performed manually by the users indicating different subgroups of information that they enter, and/or in an automated manner by the routine based on an analysis of the information that has been specified (e.g., based on the size of rules and other descriptive information supplied for a target system, based on inter-relationships between different rules or goals or other information, etc.). In other embodiments, no such decomposition may be performed.

[0221] After block 530, the routine continues to block 535 to, for each subset of target system description information

(or for all the received information if no such subsets are identified), convert that subset (or all the information) into a set of constraints that encapsulate the restrictions, goals, and other specified information for that subset (or for all the information). In block 540, the routine then identifies any errors that occur from the converting process, and if any are identified, may prompt the user to correct those errors, such as in a manner similar to that described with respect to blocks 525 and 510. While not illustrated in this example, the routine may in some situations in blocks 525 and/or 540 return to block 510 when such errors are identified, to display corresponding feedback to the user(s) and to allow the user(s) to make corrections and re-perform following operations such as those of blocks 520-540. The errors identified in the converting process in block 540 may include, for example, errors related to inconsistent restrictions, such as if the restrictions as a group are impossible to satisfy.

[0222] After block 540, the routine continues to block 545 to, for each set of constraints (or a single constraint set if no subsets were identified in block 530), apply one or more validation rules to the set of constraints to test overall effectiveness of the corresponding information that the constraints represent, and to prompt the one or more users to correct any errors that are identified in a manner similar to that with respect to blocks 525, 540 and 510. Such validation rules may test one or more of controllability, observability, stability, and goal completeness, as well as any user-added validation rules, as discussed in greater detail elsewhere. In block 550, the routine then converts each validated set of constraints to a set of coupled differential equations that model at least a portion of the target system to which the underlying information corresponds.

[0223] After block 550, the routine continues to block 553 to perform activities related to training a model for each set of coupled differential equations, including to determine one or more of a size of a training time window to use, size of multiple training time slices within the time window, and/or a type of training time slice within the time window. In some embodiments and situations, the determination of one or more such sizes or types of information is performed by using default or pre-specified information, while in other embodiments and situations the users may specify such information, or an automated determination of such information may be performed in one or more manners (e.g., by testing different sizes and evaluating results to find sizes with the best performance). Different types of time slices may include, for example, successions of time slices that overlap or do not overlap, such that the training for a second time slice may be dependent only on results of a first time slice (if they do not overlap) or instead may be based at least in part on updating information already determined for at least some of the first time slice (if they do overlap in part or in whole). After block 553, the routine continues to block 555 to, for each set of coupled differential equations representing a model, train the model for that set of coupled differential equations using partial initial state information for the target system, including to estimate values of variable that are not known and/or directly observable for the target system by simulating effects of performing control actions over the time window, such as for successive time slices throughout the time window, and to test the simulated performance of the trained model. Additional details related to training and testing are included elsewhere herein.



[0224] After block 555, the routine continues to block 560 to determine whether the training and testing was successful, and if not returns to block 510 to display corresponding feedback information to the users to allow them to correct errors that caused the lack of success. If it is instead determined in block 560 that the testing and training were successful, however, the routine continues instead to block 570 to generate an executable decision module for each trained and tested model that includes that model, as well as a local CCD Control Action Determination component that the decision module will use when executed to determine optimal or near-optimal control actions to perform for the target system based on the information included in the model, and in light of the one or more goals for that decision module. The generated executable decision module may in some embodiments and situations further include a local CCD Coordinated Control Management component to coordinate control actions of multiple decision modules that collectively will provide an automated control system for the target system, such as by synchronizing respective models of the various decision modules over time. After block 570, the routine continues to block 580 to provide the generated executable decision modules for use, including to optionally store them for later execution and/or deployment.

[0225] After block 580, the routine continues to block 595 to determine whether to continue, such as until an explicit indication to terminate is received. If it is determined to continue, the routine returns to block 510, and otherwise continues to block 599 and ends.

[0226] FIGS. 6A-6B illustrate a flow diagram of an example embodiment of a routine 600 corresponding to a generic representation of a decision module that is being executed. The routine may, for example, be provided by execution of a decision module 329 or as part of an automated control system 325 of FIG. 3 and/or a decision module 124 or 128 of FIG. 1 or 2, such as to provide functionality for controlling at least a portion of a target system in a manner specific to information and a model encoded for the decision module, including to reflect one or more goals to be achieved by the decision module during its controlling activities. As discussed in greater detail elsewhere, in some embodiments and situations, multiple decision modules may collectively and cooperatively act to control a particular target system, such as with each decision module controlling one or more distinct control elements for the target system or otherwise representing or interacting with a portion of the target system, while in other embodiments and situations a single decision module may act alone to control a target system. The routine 600 further reflects actions performed by a particular example decision module when it is deployed in controlling a portion of a target system, although execution of at least portions of a decision module may occur at other times, such as initially to train a model for the decision module before the decision module is deployed, as discussed in greater detail with respect to the CDD Decision Module Construction routine 500 of FIGS. 5A-5B.

[0227] The illustrated embodiment of the routine 600 begins at block 610, where an initial model for the decision module is determined that describes at least a portion of a target system to be controlled, one or more goals for the decision module to attempt to achieve related to control of the target system, and optionally initial state information for the target system. The routine continues to block 615 to

perform one or more actions to train the initial model if needed, as discussed in greater detail with respect to blocks 553 and 555 of FIGS. 5A-5B—in some embodiments and situations, such training for block 615 is performed only if initial training is not done by the routine 500 of FIGS. 5A-5B, while in other embodiments and situations the training of block 615 is performed to capture information about a current state of the target system at a time that the decision module begins to execute (e.g., if not immediately deployed after initial creation and training) and/or to re-train the model at times as discussed in greater detail with respect to routine 700 of FIGS. 7A-7B as initiated by block 630.

[0228] After block 615, the routine continues to block 617 to determine a time period to use for performing each control action decision for the decision module, such as to reflect a rate at which control element modifications in the target system are needed and/or to reflect a rate at which new incoming state information is received that may alter future manipulations of the control elements. The routine then continues to block 620 to start the next time period, beginning with a first time period moving forward from the startup of the execution of the decision module. Blocks 620-680 are then performed in a loop for each such time period going forward until execution of the decision module is suspended or terminated, although in other embodiments a particular decision module may execute for only a single time period each time that it is executed.

[0229] In block 625, the routine optionally obtains state information for the time period, such as current state information that has been received for the target system or one or more related external sources since the last time period began, and/or by actively retrieving current values of one or more elements of the target system or corresponding variables as needed. In block 630, the routine then initiates execution of a local CCD Control Action Determination component of the decision module, with one example of such a routine discussed in greater detail with respect to routine 700 of FIGS. 7A-7B. In block 635, the results of the execution of the component in block 630 are received, including to either obtain an updated model for the decision module with a local solution for the current time period and decision module that includes one or more proposed control action determinations that the decision module may perform for the current time period, or to receive an indication that no local solution was found for the decision module in the allowed time for the execution of the component in block 630. It is then determined in block 640 whether a solution was found, and if so continues to block 642 to store the updated model for the decision module, and otherwise continues to block 643 to use the prior model for the decision module to determine one or more control action determinations that are proposed for the current time period based on a previous model (e.g., that does not reflect recent changes in state information and/or recent changes in activities of other decision modules, if any), as discussed in greater detail with respect to routine 700 of FIGS. 7A-7B.

[0230] After blocks 642 or 643, the routine continues to block 644 to determine if other decision modules are collectively controlling portions of the current target system, such as part of the same automated control system as the local decision module, and if so continues to block 645. Otherwise, the routine selects the local proposed control actions of the decision module as a final determined control



action to perform, and continues to block **675** to implement those control actions for the current time period.

[0231] If there are other operating decision modules, the routine in block **645** determines if the local decision module includes a local copy of a CDD Coordinated Control Management (CCM) component for use in synchronizing the proposed control action determinations for the decision module's local solutions with activities of other decision modules that are collectively controlling the same target system. If so, the routine continues to block **647** to provide the one or more proposed control action determinations of the decision module and the corresponding current local model for the decision module to the local CDD CCM component, and otherwise continues to block **649** to provide the one or more proposed control action determinations for the decision module and the corresponding local model of the decision module to one or more centralized CDD CCM components.

[0232] After blocks **647** or **649**, the routine continues to block **655** to obtain results of the actions of the CDD CCM component(s) in blocks **647** or **649**, including to either obtain a further updated model resulting from synchronization of the local model for the current decision module with information from one or more other decision modules, such that the further updated model indicates one or more final control action determinations to perform for the time period for the current decision module, or an indication that no such synchronization was completed in the allowed time. The routine continues to block **660** to determine whether the synchronization was completed, and if so continues to block **665** to store the further updated model from the synchronization, and otherwise continues to block **670** to use the prior proposed control action determinations locally to the decision module as the final control action determinations for the time period.

[0233] After blocks **665** or **670**, the routine continues to block **675** to implement the one or more final determined control actions for the decision module in the target system, such as by interacting with one or more effectuators in the target system that modify values or otherwise manipulate one or more control elements of the target system, or by otherwise providing input to the target system to cause such modifications or other manipulations to occur. In block **680**, the routine optionally obtains information about the results in the target system of the control actions performed, and stores and/or provides information to the CDD system about such obtained results and/or about the activities of the decision module for the current time period.

[0234] After block **680**, the routine continues to block **695** to determine whether to continue, such as until an indication to terminate or suspend is received (e.g., to reflect an end to current operation of the target system or an end of use of the decision module to control at least a portion of the target system). If it is determined to continue, the routine returns to block **620** to start the next time period, and otherwise continues to block **699** and ends.

[0235] FIGS. 7A-7B are a flow diagram of an example embodiment of a CDD Control Action Determination routine **700**. The routine may, for example, be provided by execution of the component **344** of FIG. 3 and/or components **144a-n** or **244** of FIG. 2, such as to determine control actions for a decision module to propose and/or implement for a target system during a particular time period, including in some embodiments to perform an optimization to deter-

mine near-optimal actions (e.g., within a threshold of an optimal solution) to perform with respect to one or more goals if possible. While the illustrated embodiment of the routine is performed in a manner local to a particular decision module, such that some or all decision modules may each implement a local version of such a routine, in other embodiments the routine may be implemented in a centralized manner by one or more components with which one or more decision modules interact over one or more networks, such as with a particular decision module indicated to be used at a particular time rather than acting on behalf of the local decision module.

[0236] The illustrated embodiment of the routine **700** begins at block **703**, where information or a request is received. The routine continues to block **705** to determine a type of the information or request, and to proceed accordingly. In particular, if a request is received in block **703** to attempt to determine a solution for a current time period given a current model of the local decision module, the routine continues to block **710** to begin to perform such activities, as discussed in greater detail with respect to block **710-790**. If it is instead determined in block **705** that a request to relax one or more rules or other restrictions for the current model of the local decision module is received, such as discussed in greater detail with respect to blocks **760** and **765**, the routine continues to block **765**. If it is determined in block **705** that a request is received to repair one or more rules or other restrictions for the current model of the local decision module, such as discussed in greater detail with respect to blocks **775** and **780**, the routine continues to block **780** to obtain user input to use during the rule repair process (e.g., to interact with a CDD Decision Module Construction component, or to instead interact with one or more users in another manner), such as to allow the current model for the local decision module to later be updated and replaced based on further resulting user actions, or if operation of the target system can be suspended, to optionally wait to further perform the routine **700** until such an updated model is received. If it is instead determined in block **705** that the information or request is of another type, the routine continues instead to block **708** to perform one or more other indicated operations as appropriate, and to then proceed to block **799**. Such other indicated operations may include, for example, receiving information about current models and/or control actions proposed or performed by one or more other decision modules that are collectively controlling a target system with the local decision module (such as for use in synchronizing the model of the local decision module with such other decision modules by generating a consensus or converged shared model, as discussed in greater detail with respect to routine **800** of FIGS. 8A-8B), to receive updates to a model or underlying information for the model for use in ongoing operation of the routine **700** (e.g., from a CDD Decision Module Construction component, such as results from interactions performed in block **780**), to receive current state information for the target system, such as for use as discussed in routine **600** of FIGS. 6A-6B, etc.

[0237] If it is determined in block **705** that a request for a solution was received in block **703** for a current time period and based on a current model of the local decision module, the routine continues to block **710** to receive a current set of coupled differential equations that represent the current model for the local decision module of at least a portion of the target system, optionally along with additional state



information for the target system for the current time. The routine then continues to block **715** to determine whether to train or re-train the model, such as if the routine is called for a first time upon initial execution of a corresponding decision module or if error measurements from ongoing operations indicate a need for re-training, as discussed in greater detail with respect to blocks **755**, **770** and **730**. If it is determined to train or re-train the model, the routine continues to block **720** to determine one or more of the size of a training time window, size of training time slices within the time window, and/or type of training time slices within the training time window, such as in a manner similar to that previously discussed with respect to block **553** of routine **500** of FIGS. **5A-5B**. After block **720**, the routine continues to block **725** to use partial initial state information for the target system to train the model, including to estimate values of state variables for the target system that are not known and/or directly observable, by simulating effects of performing control actions over the time window for each of the time slices, as discussed in greater detail with respect to block **555** of routine **500** of FIGS. **5A-5B**.

[**0238**] After block **725**, or if it is instead determined in block **715** not to train or re-train the model, the routine continues to block **730** to perform a piecewise linear analysis to attempt to determine a solution for the current model and any additional state information that was obtained in block **710**, with the solution (if determined) including one or more proposed control action determinations for the local decision module to take for a current time period, as well as in some embodiments to use one or more model error gauges to make one or more error measurements with respect to the current model, as discussed in greater detail elsewhere. The routine then continues to block **735** to determine if the operations in block **730** determined a solution within a amount of time allowed for the operation of block **730** (e.g., a defined subset or fraction of the current time period), and if so continues to block **740** to update the current set of coupled differential equations and the resulting current model for the local decision module to reflect the solution, with the resulting updated information provided as an output of the routine **700**.

[**0239**] If it is instead determined in block **735** that the operations in block **730** did not determine a solution, the routine continues to block **745** to determine if additional time is available within the current time period for further attempts to determine a solution, and if not continues to block **790** to provide output of the routine **700** indicating that no solution was determined for the current time period.

[**0240**] If additional time is available within the current time period, however, the routine continues to perform blocks **755-780** to perform one or more further attempts to identify the solution—it will be appreciated that one or more of the operations of blocks **755-780** may be repeatedly performed multiple times for a given time period if sufficient time is available to continue further solution determination attempts. In particular, the routine continues to block **755** if additional time is determined to be available in block **745**, where it determines whether the measurements from one or more gauges indicate model error measurements that are over one or more thresholds indicating modifications to the model are needed, such as based on the model error measurements from the gauges discussed with respect to block **730**. If not, the routine continues to block **760** to determine whether there are one or more rules or other restrictions in

the current model that are available to be relaxed for the current time period (that have not previously attempted to be relaxed during the time period, if this is not the first pass through this portion of the routing for the current time period), and if so continues to block **765** to relax one or more such rules or other restrictions and to return to block **730** to re-attempt the piecewise linear analysis with the revised model based on those relaxed rules or other restrictions.

[**0241**] If it is instead determined in block **755** that the model error measurements from one or more of the gauges are sufficient to satisfy one or more corresponding thresholds, the routine continues instead to block **770** to determine whether to re-train the model based on one or more of the gauges indicating sufficient errors to do so, such as based on accumulated errors over one or more time periods of updates to the model. If so, the routine returns to block **720** to perform such re-training in blocks **720** and **725**, and then continues to block **730** to re-attempt the piecewise linear analysis with the resulting re-trained model.

[**0242**] If it is instead determined in block **770** not to re-train the model (or if the model was re-trained already for the current time period and the resulting re-attempt in block **730** again failed to find a solution), the routine continues to block **775** to determine whether the model error measurements from one or more of the gauges indicate a subset of one or more rules or other restrictions in the model that potentially have errors that need to be repaired. If so, the routine continues to block **780** to provide information to one or more users via the CDD Decision Module Construction component, to allow the users to revise the rules or other restrictions as appropriate, although in other embodiments some or all such rule repair activities may instead be attempted or performed in an automated manner. After block **780**, or if it is instead determined in block **775** not to repair any rules, the routine continues to block **790** to provide an indication that no solution was determined for the current time period. After blocks **740**, **708**, or **790**, the routine continues to block **799** and ends. It will be appreciated that if the routine **700** was instead implemented as a centralized routine that supports one or more decision modules remote from the executing component for the routine, the routine **700** may instead return to block **703** to await further information or requests.

[**0243**] FIGS. **8A-8B** are a flow diagram of an example embodiment of a CDD Coordinated Control Management routine **800**. The routine may, for example, be provided by execution of the component **346** of FIG. **3** and/or the components **146a-n** of FIG. **2**, such as to attempt to synchronize current models and their proposed control actions between multiple decision modules that are collectively controlling a target system. In the illustrated embodiment of the routine, the synchronization is performed in a pairwise manner between a particular local decision module's local current model and an intermediate shared model for that decision module that is based on information about the current state of one or more other decision modules, by using a Pareto game technique to determine a Pareto equilibrium if possible that is represented in a consensus shared model, although in other embodiments other types of synchronization methods may be used. In addition, in the illustrated embodiment, the routine **800** is performed in a local manner for a particular local decision module, such as by being included within that local decision module, although in other embodiments the routine **800** may be



implemented in a centralized manner to support one or more decision modules that are remote from a computing system implementing the component for the routine and that communicate with those decision modules over one or more computer networks, such as with a particular decision module indicated to be used at a particular time rather than acting on behalf of the local decision module.

[0244] The illustrated embodiment of the routine **800** begins at block **805**, where it waits to receive information or another indication. The routine continues to block **810** to determine if a consensus model or other updated information for another decision module has been received, such as from a copy of the routine **800** executing for that other decision module, and if so continues to block **815** to use the received information to update local intermediate shared model information for use with the local decision module on whose behalf the current copy of the routine **800** is executing, as discussed in greater detail with respect to block **830**. If it is instead determined in block **810** that the information or request received in block **805** is not information related to one or more other decision modules, or after block **815**, the routine continues to block **820** to determine whether to currently perform a synchronization for the current local model of the local decision module by using information about an intermediate shared model of the local decision module that includes information for one or more other decision modules, such as to do such synchronization each time that an update to the local decision module's model is received (e.g., based on operation of the routine **700** for a copy of the CDD Control Action Determination component local to that decision module) in block **805** and/or each time that information to update the local decision module's intermediate shared model is received in block **805** and used in block **815**, or instead as explicitly indicated in block **805**—if the synchronization is to currently be performed, the routine continues to block **825** and begins to perform blocks **820-880** related to such synchronization activities. Otherwise, the routine continues to block **885** to perform one or more other indicated operations as appropriate, such as to receive requests from the CDD system or other requestor for current information about operation of the routine **800** and/or to provide corresponding information to one or more entities (e.g., to reflect prior requests), etc.

[0245] If it is determined in block **820** that synchronization is to be currently performed, such as based on updated model-related information that is received in block **805**, the routine continues to block **825** to obtain a current local model for the local decision module to use in the synchronizing, with the model including one or more proposed control actions to perform for a current time period based on a local solution for the local decision module. The routine then continues to block **830** to retrieve information for an intermediate shared model of the local decision module that represents information for one or more other decision modules (e.g., all other decision modules) that are collectively participating in controlling the target system, with that intermediate shared model similarly representing one or more other proposed control actions resulting from local solutions of those one or more other decision modules, optionally after partial or complete synchronization has been performed for those one or more other decision modules between themselves.

[0246] The routine then continues to block **835** to attempt to determine a consensus shared model that synchronizes the

current model of the local decision module and the intermediate shared model by simultaneously providing solutions to both the local decision module's current model and the intermediate shared model. In some embodiments, the operations of block **835** are performed in a manner similar to that discussed with respect to blocks **710-730** of routine **700** of FIG. 7A-7B, such as if the local model and the intermediate shared model are combined to create a combination model for whom one or more solutions are to be identified. As discussed in greater detail elsewhere, in some embodiments, the local current model and intermediate shared model may each be represented by a Hamiltonian function to enable a straightforward creation of such a combined model in an additive manner for the respective Hamiltonian functions, with the operations of routines **600** and/or **700** of FIGS. 6A-6B and 7A-7B, respectively, similarly representing the models that they update and otherwise manipulate using such Hamiltonian functions.

[0247] After block **835**, the routine continues to block **840** to determine whether the operations of block **835** succeeded in an allowed amount of time, such as a fraction or other portion of the current time period for which the synchronization is attempted to be performed, and if so the routine continues to block **845** to update both the local model and the intermediate shared model of the local decision module to reflect the consensus shared model. As earlier noted, if sufficient time is allowed for each decision module to repeatedly determine a consensus shared model with changing intermediate shared models representing one or more other decision modules of a collective group, the decision modules of the collective group may eventually converge on a single converged shared model, although in other embodiments and situations there may not be sufficient time for such convergence to occur, or other issues may prevent such convergence. After block **845**, the routine continues to block **850** to optionally notify other decision modules of the consensus shared model determined for the local decision module (and/or of a converged shared model, if the operations of **835** were a last step in creating such a converged shared model), such as if each of the notified decision modules is implementing its own local version of the routine **800** and the provided information will be used as part of an intermediate shared model of those other decision modules that includes information from the current local decision module's newly constructed consensus shared model.

[0248] If it is instead determined in block **840** that a synchronization did not occur in the allowed time, the routine continues to perform blocks **860-875** to re-attempt the synchronization with one or more modifications, sometimes repeatedly if sufficient time is available, and in a manner similar to that discussed with respect to blocks **745-780** of routine **700** of FIGS. 7A-7B. In the illustrated example, the routine determines in block **860** if additional time is available for one or more such re-attempts at synchronization, and if not the routine continues to block **880** to provide an indication that no synchronization was performed within the allowed time. Otherwise, the routine continues to block **870** to take one or more actions to perform one or more of relaxing rules or other restrictions, repairing rules or other restrictions, and/or re-training a model, with respect to one or both of the current model of the local decision module and/or one or more other decision modules whose information is represented in the intermediate shared model of the local decision module. If it is determined in block **870** to



proceed in this manner, the routine continues to block **875** to perform corresponding actions, sometimes one at a time in a manner similar to that discussed with respect to routine **700**, including to cause resulting updates to the current model of the local decision module and/or to the local intermediate shared model of the local decision module, after which the routine returns to block **835** to re-attempt to synchronize the local model and the intermediate shared model of the local decision module.

[0249] If it is instead determined in block **870** that no further actions are to be performed with respect to relaxation, repair and/or re-training, the routine continues instead to block **880**. After blocks **850**, **880** or **885**, the routine continues to block **895** to determine whether to continue, such as until an explicit indication to terminate or suspend operation of the routine **800** is received, such as to reflect an end to operation of the target system and/or an end to use of the local decision module and/or a collective group of multiple decision modules to control the target system. If it is determined to continue, the routine returns to block **805**, and otherwise continues to block **899** and ends.

[0250] FIG. 9 illustrates a flow diagram of an example embodiment of a routine **900** performed for a representative generic target system, with respect to interactions between the target system and one or more decision modules that are controlling at least a portion of the target system. The routine may, for example, be provided by execution of a target system **360** and/or **370** of FIG. 3, and/or a target system **160** and/or **170** of FIGS. 1 and 2, such as to implement operations specific to the target system. It will be appreciated that the illustrated embodiment of the routine focuses on interactions of the target system with the one or more decision modules, and that many or all such target systems will perform many other operations in a manner specific to those target systems that are not illustrated here for the purpose of brevity.

[0251] The routine begins at block **910**, where it optionally provides initial state information for the target system to a CDD system for use in an automated control system of the CDD system for the target system, such as in response to a request from the CDD system or its automated control system for the target system, or instead based on configuration specific to the target system (e.g., to be performed upon startup of the target system). After block **910**, the routine continues to block **920** to receive one or more inputs from a collective group of one or more decision modules that implement the automated control system for the target system, including one or more modified values for or other manipulations of one or more control elements of a plurality of elements of the target system that are performed by one or more such decision modules of the automated control system. As discussed in greater detail elsewhere, the blocks **920**, **930**, **940** may be repeatedly performed for each of multiple time periods, which may vary greatly in time depending on the target system (e.g., a microsecond, a millisecond, a hundredth of a second, a tenth of a second, a second, 2 seconds, 5 seconds, 10 seconds, 15 seconds, 30 seconds, a minute, 5 minutes, 10 minutes, 15 minutes, 30 minutes, an hour, etc.).

[0252] After block **920**, the routine continues to block **930** to perform one or more actions in the target system based on the inputs received, including to optionally produce one or more resulting outputs or other results within the target system based on the manipulations of the control elements.

In block **940**, the routine then optionally provides information about the outputs or other results within the target system and/or provides other current state information for the target system to the automated control system of the CDD system and/or to particular decision modules of the automated control system. The routine then continues to block **995** to determine whether to continue, such as until an explicit indication to terminate or suspend operation of the target system is received. If it is determined to continue, the routine returns to block **920** to begin a next set of control actions for a next time period, and otherwise continues to block **999** and ends. As discussed in greater detail elsewhere, state information that is provided to a particular decision module may include requests from external systems to the target system, which the automated control system and its decision modules may determine how to respond to in one or more manners.

[0253] The following sections describe a variety of specific, non-exclusive embodiments in which some or all of the techniques may be implemented. It will be appreciated that particular details of particular embodiments may not be included in or true for all embodiments, and that the described embodiments may be implemented individually or in combination with any and all other combinations of other embodiments.

[0254] The following discusses several non-exclusive example embodiments, in which one or more specified embodiments of the Collaborative Distributed Decision system are each referred to as a Collaborative Distributed Inferencing (CDI) system or a Cooperative Distributed Inferencing Platform (CDIP), in which one or more specified embodiments of the Decision Module Construction component are each referred to as the “Rules Builder” or including or being performed by a “Rule Conversion Engine” (RCE) or a “CDL Compiler” or a “Rule(s) Entry Interface” (REI) or a “Rules and Goals UI”, in which one or more specified embodiments of the Control Action Determination component are each referred to as having “chattering” subcomponents or performing “chattering” functionality or including or being performed by a “Query Response Engine” (QRE) or an “Optimization Element” or via “Hamiltonian-Chattering Control”, in which one or more specified embodiments of the Coordinated Control Management component are each referred to as having or performing “mean field” information or functionality or including or being performed by a “Tellegen” agent or a “Pareto Multi-Criteria Optimization Engine” (PMOE) or a “Pareto element”, in which decision modules are referred to as “agents” or “peer agents” or “agent nodes” or “decision elements”, in which an automated control system may include a “cluster” of agents and in some cases is referred to as a “distributed architecture”, in which a target system is referred to as an “application domain”, in which a decision module’s stored information about a target system includes an “Internal Heterogeneous Database” (IHDB) and/or a “Domain Rules Knowledge Base”, in which a decision module’s consensus shared model is referred to as an “internal optimum” generated using mean field information, in which changes propagated from one decision module to others is referred to as a “delta”, etc.

[0255] CDI is built using a Peer-to-Peer (P2P) based distributed architecture that allows for partitioning the overall optimization problem into smaller sub-tasks or workloads between peer agents. CDI Peer Agents are equally



privileged participants in the application domain. They are configured to form a peer-to-peer network of nodes. This allows each agent in the network to solve the problem independently, using its own internal knowledge base. Each agent also internally engages in Pareto game playing to reach internal optimum before sharing changes with other agents. The agents then communicate the mean field changes with other agents in the network using a gossip protocol to synchronize their internal mean field approximation in an eventually consistent fashion. FIG. 10 illustrates a network diagram 1000 of a portion of a distributed architecture of an example CDI system.

#### CDI Cluster Setup

[0256] When system is initially configured certain agents are tagged as seed-nodes in the network. The seed agent nodes can be started in any order and it is not necessary to have all the seed agent nodes running. When initializing the CDI cluster at least one seed agent node is preferably running, otherwise the other CDI seed nodes may not become initialized and other CDI agent nodes might not join the cluster.

#### CDI Agent Cluster Registry

[0257] Each agent has a built-in registry of active CDI agent nodes in the cluster. This agent registry is started on all CDI agent nodes and is updated with every change in cluster membership including new agents joining the cluster, agents leaving the cluster, agent timeouts etc. Agents inform each other of active membership through a heartbeat mechanism. The registry is eventually consistent, i.e. changes may not immediately be visible at other nodes, but typically they will be fully replicated to all other nodes after a few seconds. The changed are propagated using the change deltas and are disseminated in a scalable way to other nodes with a gossip protocol.

#### CDI Cluster Agent Failure Detection

[0258] The CDI agent nodes in the cluster monitor each other by sending heartbeats to detect if a CDI agent node is unreachable from the rest of the CDI cluster. The heartbeat arrival times is interpreted by an implementation of the Phi Accrual Failure Detector (Hayashibara, Dēfago, Yared, & Katayama, 2004).

[0259] The suspicion level of failure is given by a value called phi. The basic idea of the phi failure detector is to express the value of phi on a scale that is dynamically adjusted to reflect current network conditions.

The value of phi is calculated as:

$$\text{Phi} = -\log_{10}(1 - F(\text{timeSinceLastHeartbeat}))$$

F is the cumulative distribution function of a normal distribution with mean and standard deviation estimated from historical heartbeat inter-arrival times.

#### CDI Cluster Agent Network Partition

[0260] Once a CDI agent node becomes network partitioned it will stop receiving mean field updates from the other agents in the Cluster. This however, should not prevent it from continuing to perform local optimizations based on its internal knowledge base, local mean field and sensor inputs that it will continue to receive.

[0261] The Rules Builder contains the following components:

- [0262] Rules Entry Interface
- [0263] Validator
- [0264] CDL Compiler
- [0265] Step Processor
- [0266] Workflow Remote Control
- [0267] Compiled Rules Engine Artifact

And interacts directly with the following components:

- [0268] Master Agent.

[0269] The Master Agent is responsible for interaction with all Chattering components, including the following:

- [0270] System Trainer & Bootstrapper
- [0271] Runner

The System also utilizes a Persistent Store to access state and necessary models/data. FIG. 11 illustrates a network diagram 1100 of a portion of an example Rules Builder component.

#### Rules Entry Interface:

[0272] The Rules Entry Interface is responsible for receiving rules in a syntax familiar to a domain expert. The entry interface is a text entry tool where a domain expert would insert the rules, goal and system information (rules script(s)) to provide continuous control suggestions for a given problem. The entered rules script(s) would be composed in the CDI Domain-Specific Language (DSL). The CDI DSL facilitates functional translation of the entered rules script(s) into a problem definition which houses the control definitions (measurable variable with ranges allowed and any associated constraints on the control variable).

For example:

[0273] In the entry interface using the CDI DSL one would specify an upper and lower bound using the following syntax first defining a rule, with the parameters and the rule logic that evaluates against the system state.

[0274] rule id: "max production", params: ["production"], {it<=15.0}

[0275] rule id: "min production", params: ["production"], {it>=-50.0}

[0276] control var: "production", min: "min production", max: "max production"

[0277] In the entry interface using the CDI DSL one would specify the dynamics (variables that change with relation to other measurable variables and the definition of this change)

[0278] delta var: "inventory", params: ["production", "demand"], {p, d->p-d}

[0279] In the entry interface using the CDI DSL one would specify the goal (an objective i.e., to minimize or maximize a relationship between the control and controllable dynamics).

[0280] goal objective: "minimize", params: ["inventory", "production"], {inventory, production->((inventory-10)\*(inventory-10)+(production\*production)}

[0281] Finally, the entry interface is also responsible for facilitating a domain expert to provide certain settings that are used by the system as well as providing user workflow steps. The settings provided represent key: value terms used by the system and are available throughout by the Settings Provider encapsulated within the Rules Engine. An domain expert might write initial values of the following form:

[0282] settings [initialState: "20", initialCoState: "0", terminalCoState: "0", numChatteringLevels: "9", horizon: "3", delta: "0.1", iterations: "15"]



## CDL Compiler:

**[0283]** The CDL Compiler is responsible for the conversion of said rules into evaluable constraints, system information and an optimization goal. The CDI Compiler translates the entered script(s) into a problem definition. A problem definition is composed of labeled rules. A rule is a tuple comprised of a unique name and a Term. A Term is an evaluable function where the logic was authored as previously described above, the input to which contains a representation of the system state (StateMap).

## Validator:

**[0284]** The Validator is responsible for the validation of converted constraints with regards to restrictions the optimization problem needs to satisfy as a prerequisite to its solution and/or reaching a solution quality threshold.

**[0285]** The converted constraints defined in the problem definition, along with the settings provider, is validated for controllability, observability, stability and goal completeness.

**[0286]** controllability—this check ensures that every control variable has a defined rule relating it to one or more dynamic state variables. The above example satisfies this check since our single control variable ‘production’ is used in the dynamic definition of ‘inventory’;

**[0287]** observability—checks the statemap against the defined rules in the script(s). An observability check does not impede the solution of a problem, however if failing, represents that our problem may not be well defined;

**[0288]** stability—this check ensures that our system will converge over time on a solution, which can be tested in the testing stage; and

**[0289]** goal completeness—ensures that every control variable appears in the goal.

## Compiled Rules Engine:

**[0290]** The conversion of rule scripts into mathematical expressions that can be yielded through an explicit contract to the chattering agent for use in solving the optimization goal, manifests in the compiled artifact that we label the Compiled Rules Engine

**[0291]** The Rules Engine defines and implements several interfaces for acquiring the values of these mathematical expressions, as later described with respect to “chattering” components. The Rules Engine is compiled as an artifact available to the Chattering component by the Compiler and made available for use in continuous processing. FIG. 12 illustrates a network diagram 1200 of example sub-components of an example Rules Builder component, in which compilation is dependent upon successful Validation of the observability, completeness, controllability, and workflow step validation. Any errors/warnings are returned to the domain expert for correction. Once corrections/improvements are complete, a new Compiled Rules Engine artifact is produced along with new commands delivered to the Step Processor.

## Steps Processor and Remote Control:

**[0292]** Furthermore, the rules builder is also responsible for facilitating the workflow steps an individual domain expert would need to undertake before deploying such a system. These steps include:

**[0293]** training an optimization system using the converted rules, with bootstrapped sensor data;

**[0294]** testing the resulting control actuator inputs (the end results of the chattering optimization); and

**[0295]** persisting the trained model for use in the running of the overall system.

**[0296]** The workflow component allows for training and testing a system. This is integrated into the DSL in the form of a workflow step builder.

**[0297]** For example,

---

```

workflow {
  train.onComplete {
    persist: model, '/tmp/model'
    run.withSensorData('/tmp/simulatedData').after(10 minutes)
  }
  {
    persist: results, '/tmp/results'
    shutdown
  }
}

```

---

**[0298]** Here a domain expert instructs the default bootstrapper to be used in loading predefined sensor data, train the system on this data, persist the trained model output to the file ‘/tmp/model’, then once trained to begin running the system (as a single agent). The inputs are tested by analyzing the results from the run. The artifacts and model information are persisted to a persistent store. The Steps Processor facilitates the dispatch of interpreted workflow steps to the Remote Control for delegation to the Master Agent of a running system for processing.

**[0299]** FIG. 13 illustrates a network diagram 1300 of interactions of portions of a Rules Builder component with an executing agent, including the Steps processor, the Remote Control and interactions with a running agent (the Step Processor handles interpretation of workflow steps, and the steps are delivered for dispatch by the Remote Control, which passes commands to the Master Agent; the Master Agent handles communication between a running system’s controlled elements and the Remote Control, and examples of the Controlled Elements are the Trainer and Runner).

**[0300]** FIG. 14 illustrates a network diagram 1400 of interactions of portions of a Rules Builder component with chattering components, including a Running Agent that delivers messages and interacts with the chattering components to facilitate training and running the system (the Trainer persists its model once training has completed; The Trainer invokes a Bootstrapper for acquisition of training data; The Trainer utilizes the Chattering Library; The Trainer utilizes the Chattering Library; Chattering utilizes the compiled rules; Chattering persists its control suggestions; Bootstrapper persists training data).

**[0301]** With respect to the discussion below, FIG. 15 illustrates a diagram 1500 of a sliding window for use with chattering, and is referred to in the text for the example embodiment as “FIG. 1”, while FIG. 16 illustrates a diagram 1600 of using Lebesgue integration to approximate a control trajectory for chattering, and is referred to in the text for the example embodiment as “FIG. 2”.

## 1. Parametric Chattering Control Over a Rolling Horizon

**[0302]** The real-time problem we would like to solve is really over an infinite time horizon in real-time, but our



approach is to consider a window of fixed length, and then apply invariant imbedding as we increase the window length to achieve an optimal relaxed control. The chattering control is applied to the window of fixed length, and results in a relatively easily solved problem. Then the window slides to a point where, perhaps, a new measurement is available, or a new control action is applicable. And the algorithm iterates in this manner.

**[0303]** To develop our chattering control method, we first start with a classical, deterministic control problem of the form,

$$\begin{aligned} \min_{u(t) \in U} \int_{t_0}^T \mathcal{L}(x(t), u(t)) dt + \psi(x(T)) \quad (1) \\ \text{s.t.} \\ x(t) = x(t_0) + \int_{t_0}^t f(x(\tau), u(\tau)) d\tau \\ \text{or alternatively} \\ \text{s.t.} \\ \dot{x}(t) = f(x(t), u(t)) \\ \text{for} \\ t_0 \leq t \leq T \end{aligned}$$

where  $x(t) \in \mathbb{R}^n$ ,  $u(t) \in U \subset \mathbb{R}^m$ ,  $\mathcal{L}$  is twice continuously differentiable with respect to  $x$  and continuous with respect to  $u$ ,  $\psi(x(T)): \mathbb{R}^n \rightarrow \mathbb{R}$  and is continuous and twice continuously differentiable with respect to  $x$ . Also,  $f(x(t), u(t))$  is twice continuously differentiable with respect to  $x$  and  $u$ . The time horizon  $T$  is assumed to be finite and known. We also assume initial conditions  $x(t_0)$  are known.

**[0304]** The objective is to get the behavior at  $T$  based on information from the recent past. We develop discrete time windows with dynamics of recent past, update the state with sensory information, and construct an open loop feedback strategy.

**[0305]** We reformulate Problem (1) to a standard form. First we introduce a new variable  $x_{n+1}(t)$ , with  $x_{n+1}(t_0) = 0$ , and defined by

$$\dot{x}_{n+1}(t) = \mathcal{L}(x(t), u(t)). \quad (2)$$

Note that, for  $t_0 \leq t \leq T$ ,

$$x_{n+1}(t) = \int_{t_0}^t \mathcal{L}(x(\tau), u(\tau)) d\tau.$$

Further, convert the terminal term in the criterion by defining a state variable  $y(t) \in \mathbb{R}^1$ , with  $y(t_0) = 0$ , and

$$\dot{y}(t) = \frac{\partial \psi(x(t))}{\partial x(t)} f(x(t), u(t)) \quad (3)$$

for  $t_0 \leq t \leq T$ . Note that  $y(T) = \psi(x(T))$  for any trajectory since

$$\dot{y}(t) = \frac{\partial \psi(x(t))}{\partial x(t)} \dot{x}(t).$$

**[0306]** The original control problem (1) is now of the form

$$\begin{aligned} \min_{u(t) \in U} x_{n+1}(T) + y(T) \\ \text{s.t. } x_{n+1}(t) = \int_{t_0}^t \mathcal{L}(x(\tau), u(\tau)) d\tau \\ y(t) = \int_{t_0}^t \frac{\partial \psi}{\partial x} f(x(\tau), u(\tau)) d\tau \\ x(t) = x(t_0) + \int_{t_0}^t f(x(\tau), u(\tau)) d\tau \end{aligned}$$

**[0307]** For ease of notation, we create an  $n+2$  dimension state vector,

$$z(t) = \begin{bmatrix} x_{n+1}(t) \\ x(t) \\ y(t) \end{bmatrix}$$

and denote the dynamics in integral form as

$$z(t) = z(t_0) + \int_{t_0}^t F(z(\tau), u(\tau)) d\tau$$

or, if the dynamics are differentiable, as

$$\dot{z}(t) = F(z(t), u(t)) = \begin{bmatrix} \dot{x}_{n+1}(t) \\ \dot{x}(t) \\ \dot{y}(t) \end{bmatrix} = \begin{bmatrix} \mathcal{L}(x(t), u(t)) \\ f(x(t), u(t)) \\ \frac{\partial \psi}{\partial x} f(x(t), u(t)) \end{bmatrix}.$$

**[0308]** Note that the state vector  $z(t)$  captures the original dynamics, criterion and terminal cost.

**[0309]** Without loss of generality, we now consider the following optimal control problem

$$\begin{aligned} \min_{u(t) \in U} Qz(T) \quad (4) \\ \text{s.t. } z(t) = z(t_0) + \int_{t_0}^t F(z(\tau), u(\tau)) d\tau \end{aligned}$$

where  $Q$  is a row vector of dimension  $n+2$  with  $Q = [1 \ 0_{1 \times n} \ 1]$ , so that  $Qz(T) = x_{n+1}(T) + y(T)$ . We assume  $z(t_0)$  is known (since  $x(t_0)$  is assumed known,  $x_{n+1}(t_0) = 0$ , and  $y(t_0) = 0$ ).

**[0310]** For the rolling horizon approach, we consider the initial problem on a window  $[t_0, \gamma]$ , partitioned into  $N$  sufficiently small intervals of length  $\Delta$ . We let  $\Delta$  be the same for all intervals for ease of notation. Note that  $t_0 + N\Delta = \gamma$ . We get a sequence of windows as the window moves, and the  $k$ th window is  $[t_0 + k\Delta, \gamma + k\Delta]$ . See FIG. 1 for an illustration of the sliding window.

**[0311]** We reformulate the control problem specified in (4) for the  $k$ th window, as follows,

$$\begin{aligned} \min_{u(t) \in U} Qz(\gamma + k\Delta) \quad (5) \\ \text{s.t. } z(t) = z(t_0 + k\Delta) + \int_{t_0 + k\Delta}^t F(z(\tau), u(\tau)) d\tau \\ \text{for } t_0 + k\Delta \leq t \leq \gamma + k\Delta. \end{aligned}$$

**[0312]** We let  $k=0; 1; \dots$ , and consider the end of each window as corresponding to real-time.

**[0313]** Our approach is to chatter on the control, in a similar manner as in Kohn et al. (2010). In this paper, we use a probability distribution to derive a chattering control in a probability space, and find an approximation to the relaxed control problem.

**[0314]** The relaxed form of the control problem for the  $k$ th window specified in (5) is

$$\begin{aligned} & \min_{0 \leq \alpha_i(t) \leq 1} Qz(\gamma + k\Delta) & (6) \\ \text{s.t. } & z(t) = z(t_0 + k\Delta) + \int_{t_0+k\Delta}^t \int_{\Omega} F(z(\tau), u(\tau)c) d\alpha d\tau \\ & \int_{\Omega} d\alpha = 1 \end{aligned}$$

where  $c$  is the control distribution defined by constructing the control probability distribution  $\alpha$ . The resulting  $\alpha^*$  provides the optimal distribution that solves the problem in (5).

**[0315]** Problem (6) can be viewed as a relaxed optimal control problem with respect to the Young measure.

**[0316]** Note that  $\Delta$  and  $\gamma$  should be sufficiently small, so that  $F$  remains continuous and differentiable in  $z$ , and measurable in the control.

**[0317]** The idea underlying the chattering control is to approximate the control trajectory using Lebesgue integration instead of the more traditional Riemann integration. See FIG. 2.

**[0318]** The chattering approximation to the relaxed problem (6) for the  $k$ th window is

$$\begin{aligned} & \min_{0 \leq \alpha_i(t) \leq 1, \forall i \text{ and } t} Qz(\gamma + k\Delta) & (7) \\ \text{s.t. } & z(t + \Delta) = z(t) + \Delta \sum_{i=1}^I F(z(\tau), c_i) \alpha_i(t) \\ & \sum_{i=1}^I \alpha_i(t) = 1 \\ & \text{for } t = t_0 + k\Delta, t_0 + k\Delta + \Delta, \dots, \gamma + k\Delta - \Delta \end{aligned}$$

where  $c_i$  for  $i=1; \dots; I$  is the  $i^{\text{th}}$  quantization level of  $u(t)$  at time  $t$ , providing specific control levels for the Lebesgue integration, and  $I$  is the number of levels in the interval. For ease of notation, we let the number of levels  $I$  be the same over the whole time interval. The sum of  $\alpha_i$  over  $i$  should equal one, for all  $t$ , since  $\alpha_i$  is representing a probability. Also, upper and lower bounds on  $\alpha_i$  should be satisfied, i.e.,  $0 \leq \alpha_i \leq 1$  for all  $i$  and all  $t$ .

**[0319]** Notice that  $c_i(t)$  has the same dimension as  $u(t)$ . And the number of levels  $I$  should be large enough to cover the range of  $u(t)$ . For example, if  $u(t)$  is in two dimensions, and the number of levels on dimension 1 is three, and the number of levels on dimension 2 is two, then the number of levels in is six.

**[0320]** If we know  $z(t_0+k\Delta)$ , we could solve the chattering problem (7) for  $\alpha_i$  at discrete times  $t_0+k\Delta, t_0+k\Delta+\Delta, \dots, \gamma+k\Delta-\Delta$  using an optimization engine, such as FMINCON. This involves evaluating  $F$  at all discrete times and a levels,

so we only want to solve this chattering problem occasionally. Instead of solving it frequently, we develop continualized incremental dynamics and an incremental optimization that has less computation. Then we only solve (7) when we have a large change in values or the accumulated error grows to be too large.

**[0321]** Also, if we know  $z(t_0+k\Delta)$ , we can easily solve the one-time period chattering problem for  $\alpha_i z(t_0+k\Delta)$ . For a single time period, the chattering problem (7) reduces to

$$\begin{aligned} & \min_{\alpha_i(t_0+k\Delta) \forall i} Q \left( z(t_0 + k\Delta) + \Delta \sum_{i=1}^I F(z(t_0 + k\Delta), c_i) \alpha_i(t_0 + k\Delta) \right) & (8) \\ & \sum_{i=1}^I \alpha_i(t_0 + k\Delta) = 1 \\ & 0 \leq \alpha_i(t_0 + k\Delta) \leq 1. \end{aligned}$$

**[0322]** Notice that one-time period chattering problem (8) is very easy to solve, since it is a type of relaxed knapsack problem, so the optimal  $\alpha$  is determined by rank ordering the coefficients in the objective (since the coefficients in the constraint are all equal to one). Also note that  $Q$  only has two non-zero values, so only the first and last elements of  $F$  need to be evaluated (at  $z(t_0+k\Delta)$  and all  $c_i$ ).

**[0323]** Hence we know analytically that the optimal  $c_i$  is given by,

$$\alpha_i(t_0 + k\Delta)^* = \begin{cases} 1 & \text{if } i \text{ provides the smallest value} \\ & \text{of all } QF(z(t_0 + k\Delta), c_j), j = 1, \dots, I, \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

## 2. Incremental Optimization and Continualized Incremental Dynamics

**[0324]** We start by adding an argument to the variables in (7) to indicate they are associated with the  $k$ th window problem. For instance,  $z(t, k)$  indicates the state at time  $t$  in the  $k$ th window. We also linearize with respect to  $F$ , yielding

$$z(t, k) = z(t, k-1) + \delta z(t, k) + O(\Delta^2) \quad (10)$$

for  $t_0+k\Delta \leq t \leq \gamma+(k-1)\Delta$ . We use a continualization approach ?? to yield

$$\delta z(t, k) = \sum_{i=1}^I \frac{\partial F(z(t_0 + k\Delta, k), c_i)}{\partial z} \alpha_i(t, k) \delta z(t, k)$$

where  $\delta z(t, k)$  is the change in  $z(t, k)$ .

**[0325]** For most of the systems we are considering, the value of the Jacobian matrix,

$$\frac{\partial F(z(t_0 + k\Delta, k), c_i)}{\partial z}$$

changes slowly with the state, so we can approximate it by its value at the state at the beginning of the time window



$z(t_0+k\Delta, k)$ . Thus, the matrix is only computed at the beginning of the time window. It also could be computed using finite differencing.

**[0326]** With this approximation, the incremental optimization problem with continualized incremental dynamics is given by,

$$\min_{\{\alpha_i(t,k)\}} Q\delta z(\gamma+k\Delta, k) \quad (11)$$

$$\text{s.t. } \delta \dot{z}(t, k) = \sum_{i=1}^l \frac{\partial F(z(t_0+k\Delta, k), c_i)}{\partial z} \alpha_i(t, k) \delta z(t, k)$$

$$\sum_{i=1}^l \alpha_i(t, k) = 1$$

$$0 \leq \alpha_i(t, k) \leq 1.$$

**[0327]** From Pontryagin's minimum principle for the incremental problem (11), the three necessary conditions for optimality are,

$$1. \delta \dot{z}^*(t, k) = \left( \sum_{i=1}^l \frac{\partial F(z(t_0+k\Delta, k), c_i)}{\partial z} \alpha_i(t, k) \right) \delta z^*(t, k) \quad (12)$$

with initial condition  $\delta z(t_0+k\Delta)$

$$2. p(t, k) = - \left( \sum_{i=1}^l \frac{\partial F(z(t_0+k\Delta, k), c_i)}{\partial z} \alpha_i(t, k) \right) p(t, k) \quad (13)$$

with terminal condition  $p(\gamma+k\Delta)=Q^T$ . Notice that the p equation is not incremental (p, not  $\delta p$ ), but it includes an approximation because the Jacobian matrix is evaluated at the beginning of the window.

3.

$$p^T(t, k) \sum_{i=1}^l \frac{\partial F(z(t_0+k\Delta, k), c_i)}{\partial z} \alpha_i^*(t, k) \delta z^*(t, k) \leq$$

$$p^T(t, k) \sum_{i=1}^l \frac{\partial F(z(t_0+k\Delta, k), c_i)}{\partial z} \alpha_i(t, k) \delta z(t, k)$$

for  $\alpha$  satisfying  $\sum_{i=1}^l \alpha_i(t, k) = 1$  and  $0 \leq \alpha_i(t, k) \leq 1$ .

**[0328]** Solving (12) and (13) will allow us to propagate over the length of the window, and then update the Jacobian matrix when the error reaches a threshold value.

**[0329]** The following discusses an example implementation that may be used for such chattering.

Setup

**[0330]** The optimization problem we're solving is the following:

$$\min_u \int_{t_0}^T \mathcal{L}(x(t), u(t)) dt + \psi(x(T))$$

$$\text{s.t. } \dot{x}(t) = f(x(t), u(t))$$

**[0331]** The function  $\mathcal{L}$  is the objective function; this comes from adding the original objective function specified in the DSL with penalty terms for the rules.

**[0332]** The function  $f$  specifies the dynamics of the system; these are specified in the DSL.

**[0333]** The variable  $x$  represents the state of the system at time  $t$ ; the variable  $u$  is the control (to be found) at time  $t$ .

**[0334]** Finally,  $\psi$  represents the cost at being in the terminal state  $x(T)$ . This term is typically set to 0, but does not have to be

**[0335]** The rules-builder engine creates code that allows us to compute functions used in this system, specifically

$$\mathbf{[0336]} \quad \mathcal{L}(x(t), u(t))$$

$$\mathbf{[0337]} \quad f(x(t), u(t))$$

$$\mathbf{[0338]} \quad \omega(x(t))$$

**[0339]** For these algorithms, we introduce two new (scalar) variables

$$x_{n+1}(t) = \int_{t_0}^t \mathcal{L}(x(\tau), u(\tau)) d\tau$$

$$y(t) = \int_{t_0}^t \frac{\partial \psi}{\partial x} f(x(\tau), u(\tau)) d\tau$$

**[0340]** Let  $z$  be the vector

$$z(t) = \begin{bmatrix} x_{n+1}(t) \\ x(t) \\ y(t) \end{bmatrix}$$

**[0341]** Define the function  $F(x, u, t)$  as

$$F(z(t), u(t)) = \begin{bmatrix} \mathcal{L}(x(t), u(t)) \\ f(x(t), u(t)) \\ \frac{\partial \psi(x(t))}{\partial x(t)} f(x(t), u(t)) \end{bmatrix}$$

**[0342]** Also define the matrix  $Q=[10_{1 \times n} \ 1]$ .

**[0343]** Observe that all entries of  $F$  may be computed at any time based on the functions provided by the rule builder. (Use finite differencing to compute

$$\frac{\partial \psi(x(t))}{\partial x(t)}.$$

## Chattering—Offline Training

[0344] The notation here is the same as in the paper Parametric

[0345] Chattering Control of Dynamic Systems.

[0346] Inputs:

[0347] the functions  $\mathcal{L}$ ,  $f$ , and  $\psi$ .

[0348] a time horizon  $T$

[0349] the time step  $\Delta$ .

[0350] a time window size  $\gamma$

[0351] the initial state  $x(0)$

[0352] a guess  $\delta z_0$  for the initial rate of change in the state  $z$

[0353] the terminal costate  $p(T)$ , typically assumed to be 0.

[0354]  $u_{min}$  and  $u_{max}$ , the min and max possible values for each control  $u$

[0355]  $c$ , the number of control levels for each control  $u$

[0356] Outputs:

[0357] Let  $N=T/\Delta$  be the number of intervals.

[0358] the optimal controls  $u(t)$  for each  $t=0, \Delta, 2\Delta, \dots, N\Delta$ .

[0359] the costate (assuming these controls)  $p(t)$  for each  $t=0, \Delta, 2\Delta, \dots, N\Delta$ .

[0360] the state (corresponding to picking optimal controls)  $x(t)$  for each  $t=0, \Delta, 2\Delta, \dots, N\Delta$ .

[0361] status information such as the number of iterations, whether there was a failure, etc.

[0362] Steps:

---

 Algorithm 1 Finding optimal control within window  $k$ 


---

1: for each component  $j = 1; \dots; k$ , construct  $r_j$  control levels by linearly interpolating from  $u_{min_j}$  to  $u_{max_j}$   
 2: create a list all possible control combinations of each level for each control  $j = 1; \dots; k$ . There are  $I = \prod_{j=1}^k r_j$  total such combinations. Let  $c_i$  be the  $i$ th such control combination, for  $i = 1; \dots; I$ .  
 3: Solve the following problem

$$\min_{\alpha} Qz(\gamma + k\Delta)$$

$$\text{s.t. } z(t + \Delta) = z(t) + \Delta \sum_{i=1}^I F(z(t), c_i) \alpha_i(t)$$

$$\sum_{i=1}^I \alpha_i(t) = 1 \text{ for } t = t_0 + k\Delta, \dots, \gamma + k\Delta - \Delta$$

$$0 \leq \alpha_i(t) \leq 1 \text{ for all } i, t$$

Note this is a nonlinear optimization problem and is solved using a nonlinear solver such as FMINCON.

4: the optimal control at each time is given by  $\sum_{i=1}^I \alpha_i c_i$

5: find  $p(t)$  at every time from  $t = t_0$  to  $\gamma$  via numerically integrating the following ordinary differential equation backwards in time (using any standard solver):

$$p(\gamma, 0) = Q$$

$$\dot{p}(t, k) = - \left( \sum_{i=1}^I \frac{\partial F(z(t_0 + k\Delta, k), c_i)}{\partial z} \alpha_i(t, k) p(t, k) \right)$$

-continued

---

 Algorithm 1 Finding optimal control within window  $k$ 


---

6: find  $\delta z(t)$  at every time from  $t = t_0$  to  $\gamma$  by numerically integrating the following ordinary differential equation forwards in time:

$$\delta z(0, 0) = \delta z_0$$

$$\delta \dot{z}(t, k) = \sum_{i=1}^I \frac{\partial F(z(t_0 + k\Delta, k), c_i)}{\partial z} \alpha_i(t, k) \delta z(t, k)$$

Note that the matrix  $\frac{\partial F(z(t_0 + k\Delta, k), c_i)}{\partial z}$

is the same between this step and the previous, so it doesn't have to be recomputed.

7: update the state  $z$  at every time using

$$z(t, k) = z(t, k-1) + \delta z(t, k)$$

8: return the state  $z$ , the change in state  $\delta z$ , the chattering levels  $\alpha$ , the controls  $u$  and the costate  $p$  at every time in the window

---



---

 Algorithm 2 Chattering offline training
 

---

1: Solve the chattering problem on the first window with  $k = 0$  using Algorithm 1.

2: for  $k = 2; 3; \dots$  (where  $k$  is the window number) do

3: Solve the incremental optimization problem to get the optimal controls at the end of the window

$$\min_{\alpha} Q\delta z(\gamma + k\Delta, k)$$

$$\text{s.t. } \delta \dot{z} = \sum_{i=1}^I \frac{\partial F(z(t_0 + k\Delta, k), c_i)}{\partial z} \alpha_i(t, k) \delta z(t, k)$$

$$\sum_{i=1}^I \alpha_i(t) = 1$$

$$0 \leq \alpha_i(t) \leq 1 \text{ for all } i$$

(This is a much easier optimization problem than finding the controls in Algorithm 1 as it is only for a single time step.)

4: update our estimates of  $p(t)$  at every time in the window by numerically integrating the following ODE backwards:

$$p(\gamma, 0) = Q$$

$$\dot{p}(t, k) = - \left( \sum_{i=1}^I \frac{\partial F(z(t_0 + k\Delta, k), c_i)}{\partial z} \alpha_i(t, k) p(t, k) \right)$$

5: update  $\delta z(t)$  at every time in the window by numerically integrating the following ODE forwards:

$$\delta z(0, 0) = \delta z_0$$

$$\delta \dot{z}(t, k) = \sum_{i=1}^I \frac{\partial F(z(t_0 + k\Delta, k), c_i)}{\partial z} \alpha_i(t, k) \delta z(t, k)$$

6: update the state  $z$  at every time using

$$z(t, k) = z(t, k-1) + \delta z(t, k)$$

7: end for

8: return the state  $z$ , the change in state  $\delta z$ , the chattering levels  $\alpha$ , the controls  $u$  and the costate  $p$  at every time from 0 to  $T$

---



**[0363]** Chattering—Online Prediction

**[0364]** Inputs:

**[0365]** the hamiltonian function  $H$

**[0366]** the time step  $\Delta$

**[0367]** the current state  $x(t)$

**[0368]** the current costate  $p(t)$

**[0369]**  $u_{min}$  and  $u_{max}$ , the min and max possible values for each control  $u$

**[0370]**  $c$ , the number of control levels for each control  $u$

**[0371]** Outputs:

**[0372]** the optimal control  $u(t)$

**[0373]** the new state  $x(t+\Delta)$

**[0374]** the new costate  $p(t+\Delta)$

**[0375]** Steps:

---

Algorithm 3 Chattering online updating

---

- 1: given  $x(t)$  and  $p(t)$  at the current time  $t$ .
- 2: compute optimal control  $u(t)$  by solving:

$$\min_{\alpha} Q \left( z(t_0 + k\Delta) + \Delta \sum_{i=1}^l F(z(t_0 + k\Delta)\alpha_i(t_0 + k\Delta)) \right)$$

$$\text{s.t. } \sum_{i=1}^l \alpha_i(t_0 + k\Delta) = 1$$

$$0 \leq \alpha_i(t_0 + k\Delta) \leq 1$$

This problem has a closed-form solution, as mentioned in the paper.

- 3:  $u(t) = \sum_{i=1}^l \alpha_i c_i$
- 4: update  $\delta z(t)$  over the last  $\Delta$  using

$$\delta z(0, 0) = \delta z_0$$

$$\delta \dot{z}(t, k) = \sum_{i=1}^l \frac{\partial F(z(t_0 + k\Delta, k)c_i)}{\partial z} \alpha_i(t, k) \delta z(t, k)$$

- 5: update the state  $z$  at every time using  
 $z(t, k) = z(t, k-1) + \delta z(t, k)$
  - 6: return  $z(t_{i+1}), u(t)$
- 

**[0376]** Chattering—Computing optimal values for  $T$  and  $\Delta$

**[0377]** Inputs

**[0378]** The function  $F(x, u, t)$  (from the rule builder).

**[0379]** A required error tolerance (provided in the DSL by the customer).

**[0380]** an initial guess for  $T$  and  $\Delta$

**[0381]** Outputs

**[0382]**  $\Delta$ , the time step to use

**[0383]**  $T$ , the time horizon

**[0384]** Steps

---

Algorithm 4 Finding  $\Delta$  and  $T$

---

- 1: Solve the offline chattering problem from  $t = 0$  to  $t = T$  to get controls and state for every time from  $t = 0$  to  $T$ .

- 2: Compute the Jacobian matrices  $A(t) = \frac{\partial F(z(t))}{\partial z(t)}$

at each time  $t = 0$  to  $t = T$  using finite differencing.

-continued

---

Algorithm 4 Finding  $\Delta$  and  $T$

---

- 3: for  $i = 1; 2; \dots; N$  do
- 4: let  $t = i\Delta$
- 5: Compute the error matrix  $E(t) = \int_0^t (A(t) - A(0))e^{A(0)(t-\tau)} d\tau$  via numerical integration.
- 6: Compute the total error as the sum of the absolute values of the entries of  $E(t)$
- 7: if  $E >$  error tolerance then
- 8: let  $T = (i - 1)\Delta$
- 9: break
- 10: end if
- 11: end for
  
- 12: Compute the Jacobian matrix  $J = \frac{\partial F(z)}{\partial z}$

at the current time using finite differencing.

- 13: Find the eigenvalues  $\lambda_1, \dots, \lambda_n$  of  $J$
  - 14: Take the eigenvalue  $\lambda_i$  with largest magnitude  $|\lambda_i|$ .
  - 15: Let  $\Delta = 1/|\lambda_i|$
- 

**[0385]** The following discusses further details regarding possible use with such chattering.

**[0386]** In the chattering algorithm, there are two timing parameters, the length of the window  $T$ , and the time slice  $\Delta$  (where  $N\Delta=T$ ). For accuracy purposes, these two parameters should be sufficiently small, but for computational efficiency, these two parameters should be large. Therefore, we want to keep the parameters as large as possible, as long as the error is tolerable.

**[0387]** The incremental state equation (for  $\delta z$ ) relies on the Jacobian matrix evaluated at the beginning of the time window. To characterize the error, we compare two linear systems, one that has a constant matrix, denoted  $A(t_0)$ , and the other has a time varying matrix, denoted  $A(t)$ .

Consider both linear systems:

$$\dot{x} = A(t)x(t) \quad (1)$$

$$\dot{y} = A(t_0)y(t) \quad (2)$$

with the same initial conditions,  $x(t_0)=y(t_0)$ .

The solution to (1) is of the form

$$x(t) = \phi(t, t_0)x(t_0)$$

where  $\phi(t, t_0)$  satisfies several properties,

$$\dot{\phi}(t, t_0) = A(t)\phi(t, t_0)$$

$$\phi(t_0, t_0) = I$$

$$\phi^{-1}(t, t_0) = \phi(t_0, t).$$

The solution to (2) is of the form

$$y(t) = e^{A(t-t_0)}x(t_0).$$

The error due to approximating  $A(t)$  with  $A(t_0)$  is:

$$E(t) = \phi(t, t_0) - e^{A(t-t_0)}.$$

**[0388]** Taking the derivative yields

$$\dot{E}(t) = \dot{\phi}(t, t_0) - A(t_0)e^{A(t-t_0)}$$

and replacing  $\dot{\phi}(t, t_0)$  with  $A(t)\phi(t, t_0)$  yields

$$\dot{E}(t) = A(t)\phi(t, t_0) - A(t_0)e^{A(t-t_0)}$$

and substituting for  $\phi(t, t_0)$  yields

$$\begin{aligned}\dot{E}(t) &= A(t)(E(t) + e^{A(t-t_0)}) - A(t_0)e^{A(t-t_0)} \\ &= A(t)E(t) + (A(t) - A(t_0))e^{A(t-t_0)}.\end{aligned}$$

This provides the error as a function of  $(A(t)-A(t_0))$ , and if the eigenvalues of  $A(t)$  for all  $t$  have real values less than zero, the error is linearly proportional to  $(A(t)-A(t_0))$ .

**[0389]** The procedure to choose  $T$  is to ask the customer what size of error is tolerable. The customer may say an error of 20% is tolerable, in which case the window length  $T$  can increase until reaching the associated threshold. The measure of the error is related to the difference,

$$\sum_{i=1}^l \left. \frac{\partial F}{\partial z} \alpha_i \right|_t - \sum_{i=1}^l \left. \frac{\partial F}{\partial z} \alpha_i \right|_{t_0}$$

The size of the time slice  $\Delta$  is related to the magnitude of the largest eigenvalue of

$$\frac{\partial F}{\partial z}.$$

Let  $|\lambda| = \sqrt{\lambda_{real}^2 + \lambda_{image}^2}$  and suppose  $|\lambda|_i$  has the largest value. Then

$$\Delta = \frac{1}{|\lambda|_i}.$$

**[0390]** The following discusses further details regarding an example of synchronizing a decision module's model and current information with that of information from one or more other decision modules.

**[0391]** In the multi-agent self-organizing architecture of CDI, all agents synchronize in some way. This is analogous in the worst case to the many-body problem, which Isaac Newton first formulated, and is unsolvable for three or more bodies. However, good approximations for systems of more than two bodies are computationally tractable. The approach in CDI is to solve a series of two-body problems that emulate a mean field aggregation, and update sequentially through a Pareto game.

**[0392]** Consider  $N$  agents, and each agent  $i$ ,  $i=1, \dots, N$ , has its own optimization problem with criterion, state and control:

$$\begin{aligned}\text{Min}_{u_i} J_i(x_i, u_i) \\ \text{s.t. } \dot{x}_i = f(x_i, u_i)\end{aligned}$$

The  $N$  problem in a single optimization problem is not algorithmically solvable, but it is possible to solve algorithmically a Pareto-optimization problem with two players (the Pareto game), and approximate the solution to the  $N$  player problem. The two-agent (agent 1 and agent 2) Pareto-optimization problem is

$$\text{Min}_{u_1, u_2, \alpha_1, \alpha_2} \alpha_1 J_1(x_1, u_1) + \alpha_2 J_2(x_2, u_2)$$

subject to

$$\dot{x}_1 = f(x_1, u_1)$$

$$\dot{x}_2 = f(x_2, u_2)$$

$$\alpha_1(t) + \alpha_2(t) = 1$$

**[0393]** The CDI approach is that each agent  $i$  plays a two-agent game with the CDI Mean Field agent composed of the criteria from all other agents except agent  $i$ . Instead of solving the two-agent optimization problem directly, we convert the formulation to a Hamiltonian problem (using Pontryagin's minimum principle).

**[0394]** The Hamiltonians are additive for Pareto optimization. Therefore, we can add local Hamiltonians for individual agents to create an aggregate mean field Hamiltonian.

**[0395]** Let  $H_i$  be the local Hamiltonian for agent  $i$  and let  $H_i^{MF}$  be the mean field Hamiltonian composed of the Hamiltonians for all other agents, excluding  $i$ . That is, the mean field Hamiltonian for agent  $i$  is a functional form of the Hamiltonians of the other agents.

Then the two-agent optimization problem is

$$\text{Min}_{u_i, u_i^{MF}, \alpha_1, \alpha_2} \alpha_1 H_i(x_i, u_i) + \alpha_2 H_i^{MF}(x_i^{MF}, u_i^{MF})$$

subject to the state and costate equations of the combined Hamiltonian, and

$$\alpha_1(t) + \alpha_2(t) = 1.$$

The state consists of  $[x_i, x_i^{MF}]^T$ , the costate is  $[p_i, p_i^{MF}]^T$ , and the control is  $[u_i, u_i^{MF}]^T$ . The initial conditions for  $x_i^{MF}$  and  $p_i^{MF}$  come from the previous solution, and  $u_i$  is the local solution from the previous pass.

**[0396]** The solution to the two-agent Pareto game provides  $\alpha_1(t)$  and  $\alpha_2(t)$ , and the total Hamiltonian for agent  $i$  is updated:

$$H_i^T = \alpha_1 H_i + \alpha_2 H_i^{MF}$$

The modified Hamiltonian for agent  $i$  is constructed by projecting the total Hamiltonian into the local state space. The modified mean field Hamiltonian is constructed by projecting the total Hamiltonian into the mean field state space.

**[0397]** Now, the another agent plays the game, solving the two-agent Pareto game, and updating locally, and each agent updates its mean field agent when it is ready to play its game.

**[0398]** The following discusses further details regarding example embodiments.

#### Overview

**[0399]** The Cooperative Distributed Inference (CDI) platform is a unique advanced technology offering to enable near-optimal and near-real-time decision making on vast amount of heterogeneous and distributed information, in complex knowledge-based decision support systems, combining absolute, hard and soft rules to handle various requirements from natural or governing laws, policies, and best practices. FIG. 17 illustrates a diagram 1700 of various interactions of different portions of a CDI system.

**[0400]** The CDI platform features a Distributed Architecture (DA) for resolving queries by accessing information from both an Internal Heterogeneous Database (IHDB) and external data sources referred to as Sensors. CDI utilizes a network of computing devices as its nodes—called Decision Elements (DE)—that cooperate to resolve queries given to them.



**[0401]** The Decision Elements (DE) in a given DA can work together to reach best outcome for the whole group, i.e., reaching Pareto Efficiency (also referred to as Pareto equilibrium)—a stable state where no change by any individual can be made to make the sum of whole group better.

**[0402]** Each DE can solve the problem independently if provided with complete knowledge and data. DE solves the query with a very unique approach, using Optimal Control Theory, starting with a technique called analytic continuation (transforming query and rules into differential equations whose dependent variables represent internal variables and parameters of the rules), then using its own internal knowledge to solve the query, providing an outcome.

**[0403]** In distributed environments, group of Decision Elements synchronize in an iterative process utilizing updates from other DEs and provide a final result, via a Pareto multi criteria optimization strategy.

**[0404]** The CDI platform needs rules, not necessarily exact criteria, to reach the objective state, producing the near-optimal solution. This is a very attractive feature when exact quantitative criteria cannot be provided in advance due to uncertainty or other reasons.

#### Platform Features

**[0405]** CDI is perfect for big data analytics, supporting many data types:

**[0406]** Structured: databases, ontologies

**[0407]** Semi-structured: spreadsheets, CSV/TSV, forms, emails

**[0408]** Unstructured: web pages, blogs, text documents

**[0409]** Symbolic: business rules, math formulas

CDI's distributed computing architecture also make it very scalable to handle large amount (peta-size) of heterogeneous data ingestion while performing real-time analytic results even in microseconds (depending in part on resources available and the complexity of queries).

#### Rules Support

**[0410]** CDI can integrate different types of business rules: absolute, hard and soft rules, from natural or government laws, operational or policy requirements, and practice guidelines. Absolute rules and hard rules always take logic value 0 (false) or 1 (true) when instantiated. Soft rules, however, may take any value in the interval [0,1], or more generally more than two values. Absolute rules reflects a should-satisfy situation, such as FDA/USDA requirements; hard rules are operational requirements such as “no serious persistent side effects”, but which may be temporarily relaxed in specific situations; and soft rules can come from guidelines or experiences, such as “better not give drug XYZ to diabetes patients with heart problems”. FIG. 18 illustrates a diagram 1800 of examples of different types of rules. Please note that all chaining of the rules may happen automatically in a dynamic manner during the optimization process. CDI handles this complexity with said techniques above by solving a distributed continuous-space optimization problem.

#### Self-Adapting and Learning

**[0411]** CDI platform features a self-learning design. As CDI converts the original query solving into an optimal control problem, it can use feedbacks from the environment (external sensor or internal updates from peer decision

elements) to refine its internal model: a Hamilton-Jacobi-Bellman equation will be updated to reflect new information and automatically form soft-rule like constraints internally. The process to update internal mathematical model is similar to reconstructing 3D images from CT or MRI imaging by tomographic reconstruction, but used for dynamic systems.

#### Scalability and Performance

**[0412]** CDI platform enables functionality to:

**[0413]** Integrate data that may include 1,000,000+ variables and 100,000+ of constraints

**[0414]** Provide data integration over evolving distributed network

**[0415]** Specify queries over a broad range of languages

**[0416]** Specify queries of a broad range of complexity

**[0417]** Provide best known response to queries at the local level

**[0418]** Operate in a variety of environments including cloud-based or local deployments

**[0419]** Real-time processing of queries

#### Comparisons

**[0420]** CDI stands out in being able to do the following things:

**[0421]** Approximate the optimal solution of NP-hard problems (such as planning and scheduling) by mapping criteria and constraints onto a continuous space and solve it as an Optimal Control problem with polynomial-time algorithms. The solutions offered by CDI may be near optimal rather than exactly optimal but the running time will be greatly shortened and some large-scale intractable problems can become solvable by CDI.

**[0422]** Blend adaptive learning together with rules composition. Traditional rule engines might support hard and soft rules (constraints), but the scoring mechanism and weights need be manually adjusted to meet the goals, while in CDI, via feedback mechanism, these weights can be optimized as internal configurations.

The unique abilities make CDI an ideal choice for intelligent decision support systems.

#### Sample Applications

**[0423]** CDI platform works great in areas where there are a lot of heterogeneous data, government compliances and business requirements, such as healthcare and energy.

#### Clinical Auto-Coding Application

**[0424]** A Clinical Auto-Coding (C.A.C.) application can be used to detect medical under-coding, over-coding, and miscoding, highlighting potential opportunities where higher billing is justified. It automatically generates clinical codes, including ICD-10 directly from clinical encounter notes such as physician notes, lab results, and discharge records, while supporting workflows accommodating the roles played by administrators, coders and doctors in coding. FIG. 19 illustrates an example user interface 1900 related to medical/clinical auto-coding. C.A.C. incorporates user feedback to learn coding best practices as it processes records. With its sophisticated adaptive technology, C.A.C. improves over time, optimizing coding for each organization to improve the efficiency, accuracy and revenue capture of the medical coding activity



### Clinical Intelligence Web Services

**[0425]** CDI powers the following intelligent healthcare services that can be easily integrated:

**[0426]** Clinical Record Intelligence—For analysis of EMRs (electronic medical records) and encounter notes, identification of actionable clinical terms and concepts in those documents. The services can extract concepts, understand issues, and analyze documents.

**[0427]** Clinical Coding Intelligence—For inferring clinical codes from a set of clinical concepts, grouping codes, correlating codes, analyze documents, audit and justification, and compliance.

**[0428]** Patient-Centric Intelligence—Secure delivery of end-user applications to predict, recommend, and explain personalized actions to improve patient outcomes. It provides a patient centric view, can provide medical risk prediction and proactive monitoring for patients, automated data abstraction, actionable recommendations, and more.

All these web services run in secure cloud with full compliance measurement in place.

### Fraud, Waste and Abuse Detection

**[0429]** In healthcare, CDI helps assess and monitor risk of medical fraud, waste and abuse (FWA) by uncovering providers, and to a lesser extent pharmacies, who are suspected of having committed fraud via a variety of schemes. Similar application domains include financial frauds.

**[0430]** By ingesting a wide variety of data that is difficult to link (including live public data), a FWA service analyzes data to find evidence and patterns of fraud, and provides a dashboard application for agents/detectives to prioritize and act on the discovered suspected cases of medical fraud. The service features strict and fuzzy rules-based detection as well as automatic pattern discovery, and runs in real-time and continuous mode to support proactive monitoring and action.

### Energy Intelligence

**[0431]** Smart Grid can involve uncertain bi-directional exchange in distributed grids, so intelligent control can be used to provide active synchronization between the network of element controllers and the outside grid management system allows high quality of service to be maintained in a cost-effective manner, hence the dynamics can be learned from sensory observations. CDI enables distributed micro-grid control, supports inductive modeling with “soft” rules that are learned and continuously updated, to optimize the grid.

### Distributed Architecture (DA)

**[0432]** The DA is a network of interacting components called decision elements (DEs). The DEs collaborate in the resolution of a query posed by one of them. The DA’s block diagram is shown in FIG. 18, where Sensors refer to external input data in general. There is an additional translation layer to process external data to be consumed by DEs.

### Decision Element (DE)

**[0433]** A decision element is a higher-level functional component solving queries locally. It can have subcomponents such as a programmable search engine, internal het-

erogeneous database, Inference engine, Inference rule base, API/user interface, and network interface. A decision element is capable of providing a quick and near-optimal solution to a complex query with complete input of data.

### Internal Heterogeneous Database (IHDB) and External Knowledge Base (EKB)

**[0434]** IHDB is data preprocessed and stored by a specific decision element (DE). FIG. 20 illustrates a diagram 2000 of some components of a CDI system, including Knowledge Bases containing data sources ranging from domain knowledge to general facts. IHDB encodes knowledge into knowledge components (KC’s). Each KC is used and updated by a DE in the DA, and multiple KCs may share rules.

**[0435]** External Knowledge Base (EKB) refers to data, including rules, as input into specific decision elements, such as patient’s body temperature, blood pressure, or instantiated rule to determine if patient’s cholesterol level is high. EKB can also contain communicated information from other DEs. Domains for variables include: real, complex, integer, binary numbers and symbolic token on finite domains.

### Interfacing Components

**[0436]** The following components act between users and the data via API and/or GUI.

**[0437]** Rule Entry Interface—It provides the entry of rules into the IHDB, validates the specification of rules before insertion, and route the rules to the appropriate DE for insertion to their respective knowledge components.

**[0438]** Sensor Ingestion Interface—A sensor is a machine or service where external data exist. Sensor Ingestion Interface (SII) enables the system to add or remove a sensor, poll a sensor and submit data to the network.

**[0439]** Query Language Interface—Query Language Interface accepts the query, submits it to the system (Distributed Architecture) and provides response. It is an API and flexible UI can be built on top of it.

### Minimization Function Generator (MFG)

**[0440]** The minimization function generator converts a query to a minimization function (i.e. analytic continuation). This is useful because the problem is converted from search problem in a large discrete space (like graph search problems, which are usually NP-complete) into an optimization problem in continuous space (polynomial algorithms exist). An analogy is NP-hard integer programming, while continuous linear programming has a very efficient solution.

### Query Response Engine (QRE)

**[0441]** The Query Response Engine is the core of the whole system responsible solving the query (locally). A mathematical model is constructed based on these equations obtained from previous steps, containing the current state and object state (goal state). The continuous-space optimization automatically handles forward and backward rule chaining by moving along the trajectory toward target state. It also manages uncertainty by keeping a large set of possible states and reducing the solution space only when



more information becomes available. Standard optimization techniques (e.g. Newton-Raphson Method) can be employed to solve the problem.

**[0442]** Feedbacks and updates (data from EKB or other DE in the architecture) will be used to refine the mathematical model over time; therefore, the core engine is self-adapting.

#### Pareto Multi-Criteria Optimization Engine

**[0443]** The Pareto Multi-Criteria Optimization Engine (PMOE) is the aggregation step where all DE in the network settle to obtain a good stable solution—a state where no improvements can be made to any individual DE without reducing the whole team’s performance—a state belongs Pareto Optimal Set in Game Theory. It is like each DE is playing a game and they communicate and interact and work together to make the best outcome for the criteria (query). To efficiently synchronize all decision elements, Mean Field Theory is applied for dimension reduction using knowledge obtained. FIG. 21 illustrates a diagram 2100 of performing Pareto processing for use with mean field techniques, including to reach Pareto efficiency.

#### Data Exchange Specifications

**[0444]** The system can take many different data types via ingestion API and has adapters from different public data sources. The supported data types include common ones from XML, CSV, TSV, SQL, Spreadsheet, JSON, and more. OData support is also available. Output types can be API (XML or JSON), as well as exporting to CSV, SQL or directly to services such as SOLR and Cassandra.

#### Running Environment

**[0445]** CDI can run in the cloud as Software as a Service platform. We can provide whole end-to-end support by setting up the infrastructure in a Virtual Private Cloud or deploy and configure it with the cluster clients provide. The system features SaaS architecture and provides APIs to be used by third parties. For advanced integration requests, Java/Scala and Python libraries are available.

**[0446]** To sum up, CDI platform utilizes many advanced techniques from Mean Field Theory, Lagrangian and Hamiltonian functions, Pareto Optimal Set, Gauge Theory, and so on derived from modern mathematics, quantum physics, optimal control and game theory to achieve high performance. In the detailed computation process, lots of transformation, approximation, optimization, caching and other advanced computing techniques are used to improve accuracy, speed and scalability. Due to this unique approach, CDI is able to solve complex decision making problem in a smoothly, efficient manner and achieve near optimal results.

**[0447]** The Cooperative Distributed Inference (CDI) platform is a unique advanced technology offering to enable near-optimal and near real-time decision making on vast amount of heterogeneous and distributed information, in complex knowledge-based decision support systems, combining absolute, hard and soft rules to handle various requirements from natural or governing laws, policies, and best practices.

**[0448]** Each agent in the Distributed Architecture is a Decision Element that can solve the problem independently, provided with knowledge base and data. In a very unique manner, the agent uses Optimal Control Theory to obtain a

near optimal solution, starting with a technique called analytic continuation (transforming query and rules into differential equations whose dependent variables represent internal variables and parameters of the rules), then using its own internal knowledge to solve the problem as an optimization problem in continuous space, to allow for efficient solving.

**[0449]** The outcome will be fed back to the Rules Editor and optimization process (the Chattering algorithm) to enable automatic adjustment of weights of soft rules (constraints) and achieve optimal score of the objective function.

**[0450]** A CDI agent is an independent decision element that can take in sensor data (input from the environment, streaming or in batches), as well as the knowledge-bases (rules composed by domain experts, including absolute, hard and soft rules) to generate a set of partial differential functions (Lagrangian constraints), through continualization. Similarly, the objective is also converted into a minimization function, as part of the Lagrangians which will be solved via the Hamilton-Jacobi-Bellman equation. Each agent will talk to other agents via a “mean field” abstraction layer to greatly reduce the communication and computation overhead, and incorporate additional information to reach the global optimal state (a stable, near optimal set of states across all agents). Finally, the agent exercises control over the system. FIG. 22 illustrates a network diagram 2200 of an example decision module agent.

**[0451]** Various domain knowledge is captured from experts in the form of Domain Specific Language. Some are constraints; logic forms need be converted to Boolean equations; and variables in soft rules take values  $[0 \dots 1]$ . This creates a set of equations to be solved by the optimization algorithm. The optimization process takes data from a time range and finds the best state configuration to reach optimality. It starts with a “learning” process to find a good initial configuration by taking in a short history of data, before processing real-time streams. FIG. 23 illustrates a network diagram 2300 of an example of offline workflows for knowledge capture.

**[0452]** Each CDI agent computes a “mean field” view of the system via its neighbors, and responds to queries with the latest updates. The approximate mean field view of a group greatly reduces computational dimensions. The agents synchronize with others and understand the global state via the mean-field approximation. They engage in games to reach Pareto Optimal (also referred to as Pareto equilibrium)—the best output. FIG. 24 illustrates a network diagram 2400 of an example of workflows for mean field computation and Pareto Optimal.

**[0453]** An example home solar micro-grid system illustrates one example embodiment of a CDI application or automated control system, which takes the sensor data from the solar panel (in the house), substation, and power network, and decides whether or not to fulfill the utility requests in real time using a set of complex rules. FIG. 25 illustrates a network diagram 2500 of an example of an automated control system for a home solar micro-grid electrical generating system.

**[0454]** FIGS. 26-28 provide further details regarding operations of portions of example CDI systems and their sub-components, and FIGS. 29A-29K illustrate examples of using a CDI system to iteratively determine near-optimal solutions over time for controlling a target system in diagrams 2900A-2900K. In particular, FIG. 26 illustrates a further diagram 2600 of workflow and components of a



portion of a CDI system, FIG. 28 illustrates a diagram 2800 of an overview workflow for a portion of a CDI system, and FIG. 27 illustrates a diagram 2700 of workflow for an inference process portion of a CDI system.

[0455] Further details related to an example CDI system are shown below and included in provisional U.S. Patent Application No. 62/015,018, filed Jun. 20, 2014 and entitled “Methods And Systems For Cooperative Distributed Inferencing,” which is hereby incorporated by reference in its entirety. In addition, further details related to example details of using gauges to perform model error measurements are included in provisional U.S. Patent Application No. 62/182,796, filed Jun. 22, 2015 and entitled “Gauge Systems,” which is also hereby incorporated by reference in its entirety. Furthermore, further details related to examples of using the CDD system in particular manners with particular types of target systems are included in provisional U.S. Patent Application No. 62/182,968, filed Jun. 22, 2015 and entitled “Applications Of Cooperative Distributed Control Of Target Systems,” which is also hereby incorporated by reference in its entirety—the example types of target systems discussed include, for example, a vehicle being controlled, controlling inventory management, controlling a micro-grid electrical generation facility, controlling activities in a day trader financial setting, etc.

Current notation	Comments
t	Notation for algorithmic time.
q(t)	Notation of canonical coordinate vector for entire system.
q	Notation of canonical coordinate vector dropping time dimension.
q <sup>f</sup>	Notation of canonical coordinate vector for specific function f.
$\dot{q}$	Notation of first time derivative of canonical coordinate vector.
$\ddot{q}$	Notation of second time derivative of canonical coordinate vector.
h	Notation of HEAD of Horn clause.
$\phi(q)$	Notation of generic proposition.
$\sigma(q)$	Notation of generic proposition alternate to $\phi$ .
$T_i$	Notation of the TV of a soft rule.
$\tilde{i}(q; \phi, \sigma)$	Generic equational form relating two propositions.
$\check{\phi}(q)$	Notation of the equational form of $\phi(q)$ .
$\Phi_Q(q)$	Notation for proposition defined by the query.
$\check{\Phi}_Q(q)$	Notation for equation defined by the query.
J(q)	Notation for minimization function for the query.
$\mathcal{L}$	Notation for static Lagrangian
$\mathcal{L}_k^{(\alpha, T)}$	Notation for total static Lagrangian for $DE_k$ .
q	
$\{p_\alpha\}$	
$u^{(k)}$	
$H_k^{(\alpha)}$	Primary Hamiltonian for the absolute rules for $DE_k$ .
$H_k^{(A)}$	Hamiltonian for the Tellegen agent of the total Hamiltonian's rules.
$H_k^{(T)}$	Total Hamiltonian for $DE_k$ .

## Overview

[0456] This document introduces and specifies the architecture for the Cooperative Distributed Inferencing Platform (CDIP). The primary instance of this is the Distributed Architecture (DA) for resolving queries by accessing both an Internal Heterogeneous Database (IHDB) populated by a special class of Horn Clause rules and external data sources referred to as sensors.

[0457] The architecture implements a network of active devices at its nodes. These devices are called Decision Elements (DE's). The DE's cooperate in the resolution of a query posed to one or several of them. The DE's in a given DA are referred to as the team.

[0458] Every DE in a team is programmed to transform rules in its domain, determined by a posed query, into an ordinary differential equation (ODE), whose dependent variables represent internal variables and parameters. The dependent variables include unknowns of the query posed to the DE. The DE's in the architecture are synchronized via a Pareto multi criteria optimization strategy.

[0459] The components of the CDIP include:

[0460] Application requirements that the system is designed to accommodate.

[0461] Functional requirements that satisfy the application requirements and pertain directly to the construction and operation of the system components.

[0462] Subcomponents which are necessary to implement the functional requirements

[0463] Limitations which highlight noteworthy constraints which are inherent the specified implementation of the architecture.

[0464] Architectural flow describes key aspects of the architecture that indicate how the system is to be constructed given the specified essence and key behavior of the subcomponents.

[0465] Software realization of the architecture describes the key pieces of software necessary for system implementation.

[0466] Data describes the kinds of data the system is expected to accept as input and produce as output.

[0467] Data exchange protocols reference key data types and structures that need to be exchanged across the system and the protocols for exchange.

[0468] Environment describes the particulars of the environments that the system will be able to operate in and therefore should be tested in.

[0469] Testing describes how the system should be tested given the data and operating environments.

## Subcomponents

[0470] Subcomponents are fundamental parts of the architecture that perform particular roles. This section contains descriptions for each of the subcomponents of the architecture. The subcomponents are:

[0471] The Distributed Architecture (DA).

[0472] The Internal Heterogeneous Database (IHDB).

[0473] The Rule Entry Interface (REI).

[0474] The Rule Editor (RE).

[0475] The External Knowledge Base (EKB).

[0476] The Sensor Ingestion Interface (SII).

[0477] The Rule Conversion Engine (RCE).

[0478] The Decision Element (DE).

[0479] The Query Language Interface (QLI).



- [0480] The Minimization Function Generator (MFG).
- [0481] The Query Response Engine (QRE).
- [0482] The Pareto Multi-Criteria Optimization Engine (PMOE).

#### Distributed Architecture (DA)

[0483] The DA is a platform of interacting components called DE's. The DE's collaborate in the resolution of a query posed by one of them. The DE's implement a distributed, dynamic optimization process, herein referred as the optimization process (OP). OP implements an optimization process that computes an answer to the active queries as a function of data stored in two categories of repositories: IHDB and EKB's. These repositories of the data needed to implement OP given a query.

[0484] The EKB's are a collection of public or private repositories of knowledge relevant to the DE posing a query. A DE has a list of external repositories (LER). Each entry in an LER includes 1) a protocol, 2) a heading sub-list, and 3) a translation grammar. Each protocol entry prescribes the access procedure to the corresponding knowledge repository. Each heading sub-list entry contains a summary of the knowledge contents of the corresponding repository. Finally, each translation grammar entry provides a procedure for converting knowledge elements of the corresponding repository in to the rule representation in the IHDB of the DE.

[0485] Functional characteristics of this architecture and in particular, the DE's, IHDB, and the sensors are described, including the following concepts:

- [0486] The DA
- [0487] A process for resolving queries by accessing the IHDB and External Knowledge Bases (EKB's) through sensors
- [0488] The constitution of DE's
- [0489] A query and corresponding rules transformation into an ODE
- [0490] The orchestration of a team of DE's through a Pareto multi criteria optimization strategy
- [0491] The Internal Heterogeneous Database (IHDB)
- [0492] Composition of IHDB as a Set of Knowledge Components
- [0493] The IHDB encodes knowledge about the implemented application. The IHDB is divided into knowledge components (KC's). Each KC is consulted and updated by a DE in the DA. Any pair of KC's may have an overlapping set of rules by which they operate, but there is no a priori constraint on intersections or inclusion. The collection of KC's constitutes the existing knowledge of the system.
- [0494] Algorithmic Formulation of a Rule
- [0495] A KC is a collection of rules, written in a restrictive Horn clause format. The rules are logic entities. When a rule is instantiated, it has a logic value. The logic values a rule can have are taken from the interval [0,1]. The entire system of rules is evaluated using variables and parameters which are collectively referred to as the generalized coordinates of the system and are indexed as follows:

$$q(t)=\{q^{(1)}(t), \dots, q^{(N)}(t)\}. \quad (3.2-1)$$

[0496] The time argument  $t$  refers to the algorithmic time of the system which means that it has no other meaning than as a continuous index with respect to the evolution of the system. There is therefore no requirement that it correspond

to a physical aspect of the system although this may naturally occur. Physical time may be represented specifically by a canonical coordinate of choice  $q^{(i)}(t)$ . Alternatively, we may refer to the  $q$ 's without expressly stating the independent time argument and write

$$q(t)=\{q^{(1)}, \dots, q^{(N)}\}. \quad (3.2-2)$$

Then we should also note that the time derivatives are referred to notationally as

$$\dot{q} = \frac{dq(t)}{dt}, \quad \ddot{q} = \frac{d^2q(t)}{dt^2} \quad (3.2-3)$$

[0497] These coordinates are referred to variously as  $q$  depending on the context and the expected arguments of the function to which they are applied. When it is necessary to distinguish between more than one  $q$  in equational form we generally write  $q_f$  where  $f$  denotes the reference function or appropriate domain. Typically, we assume without loss of generality the entire set of canonical coordinates  $q$  is an argument to any function, term or proposition. In practice, we may further assume it is possible to apply the particular required coordinates as need to mathematical construct in question.

[0498] The rules in each knowledge component are of three types: absolute rules, hard rules, and soft rules. Absolute rules and hard rules take logic value 0 (false) or 1 (true) when instantiated. Soft rules take any value in the interval [0,1].

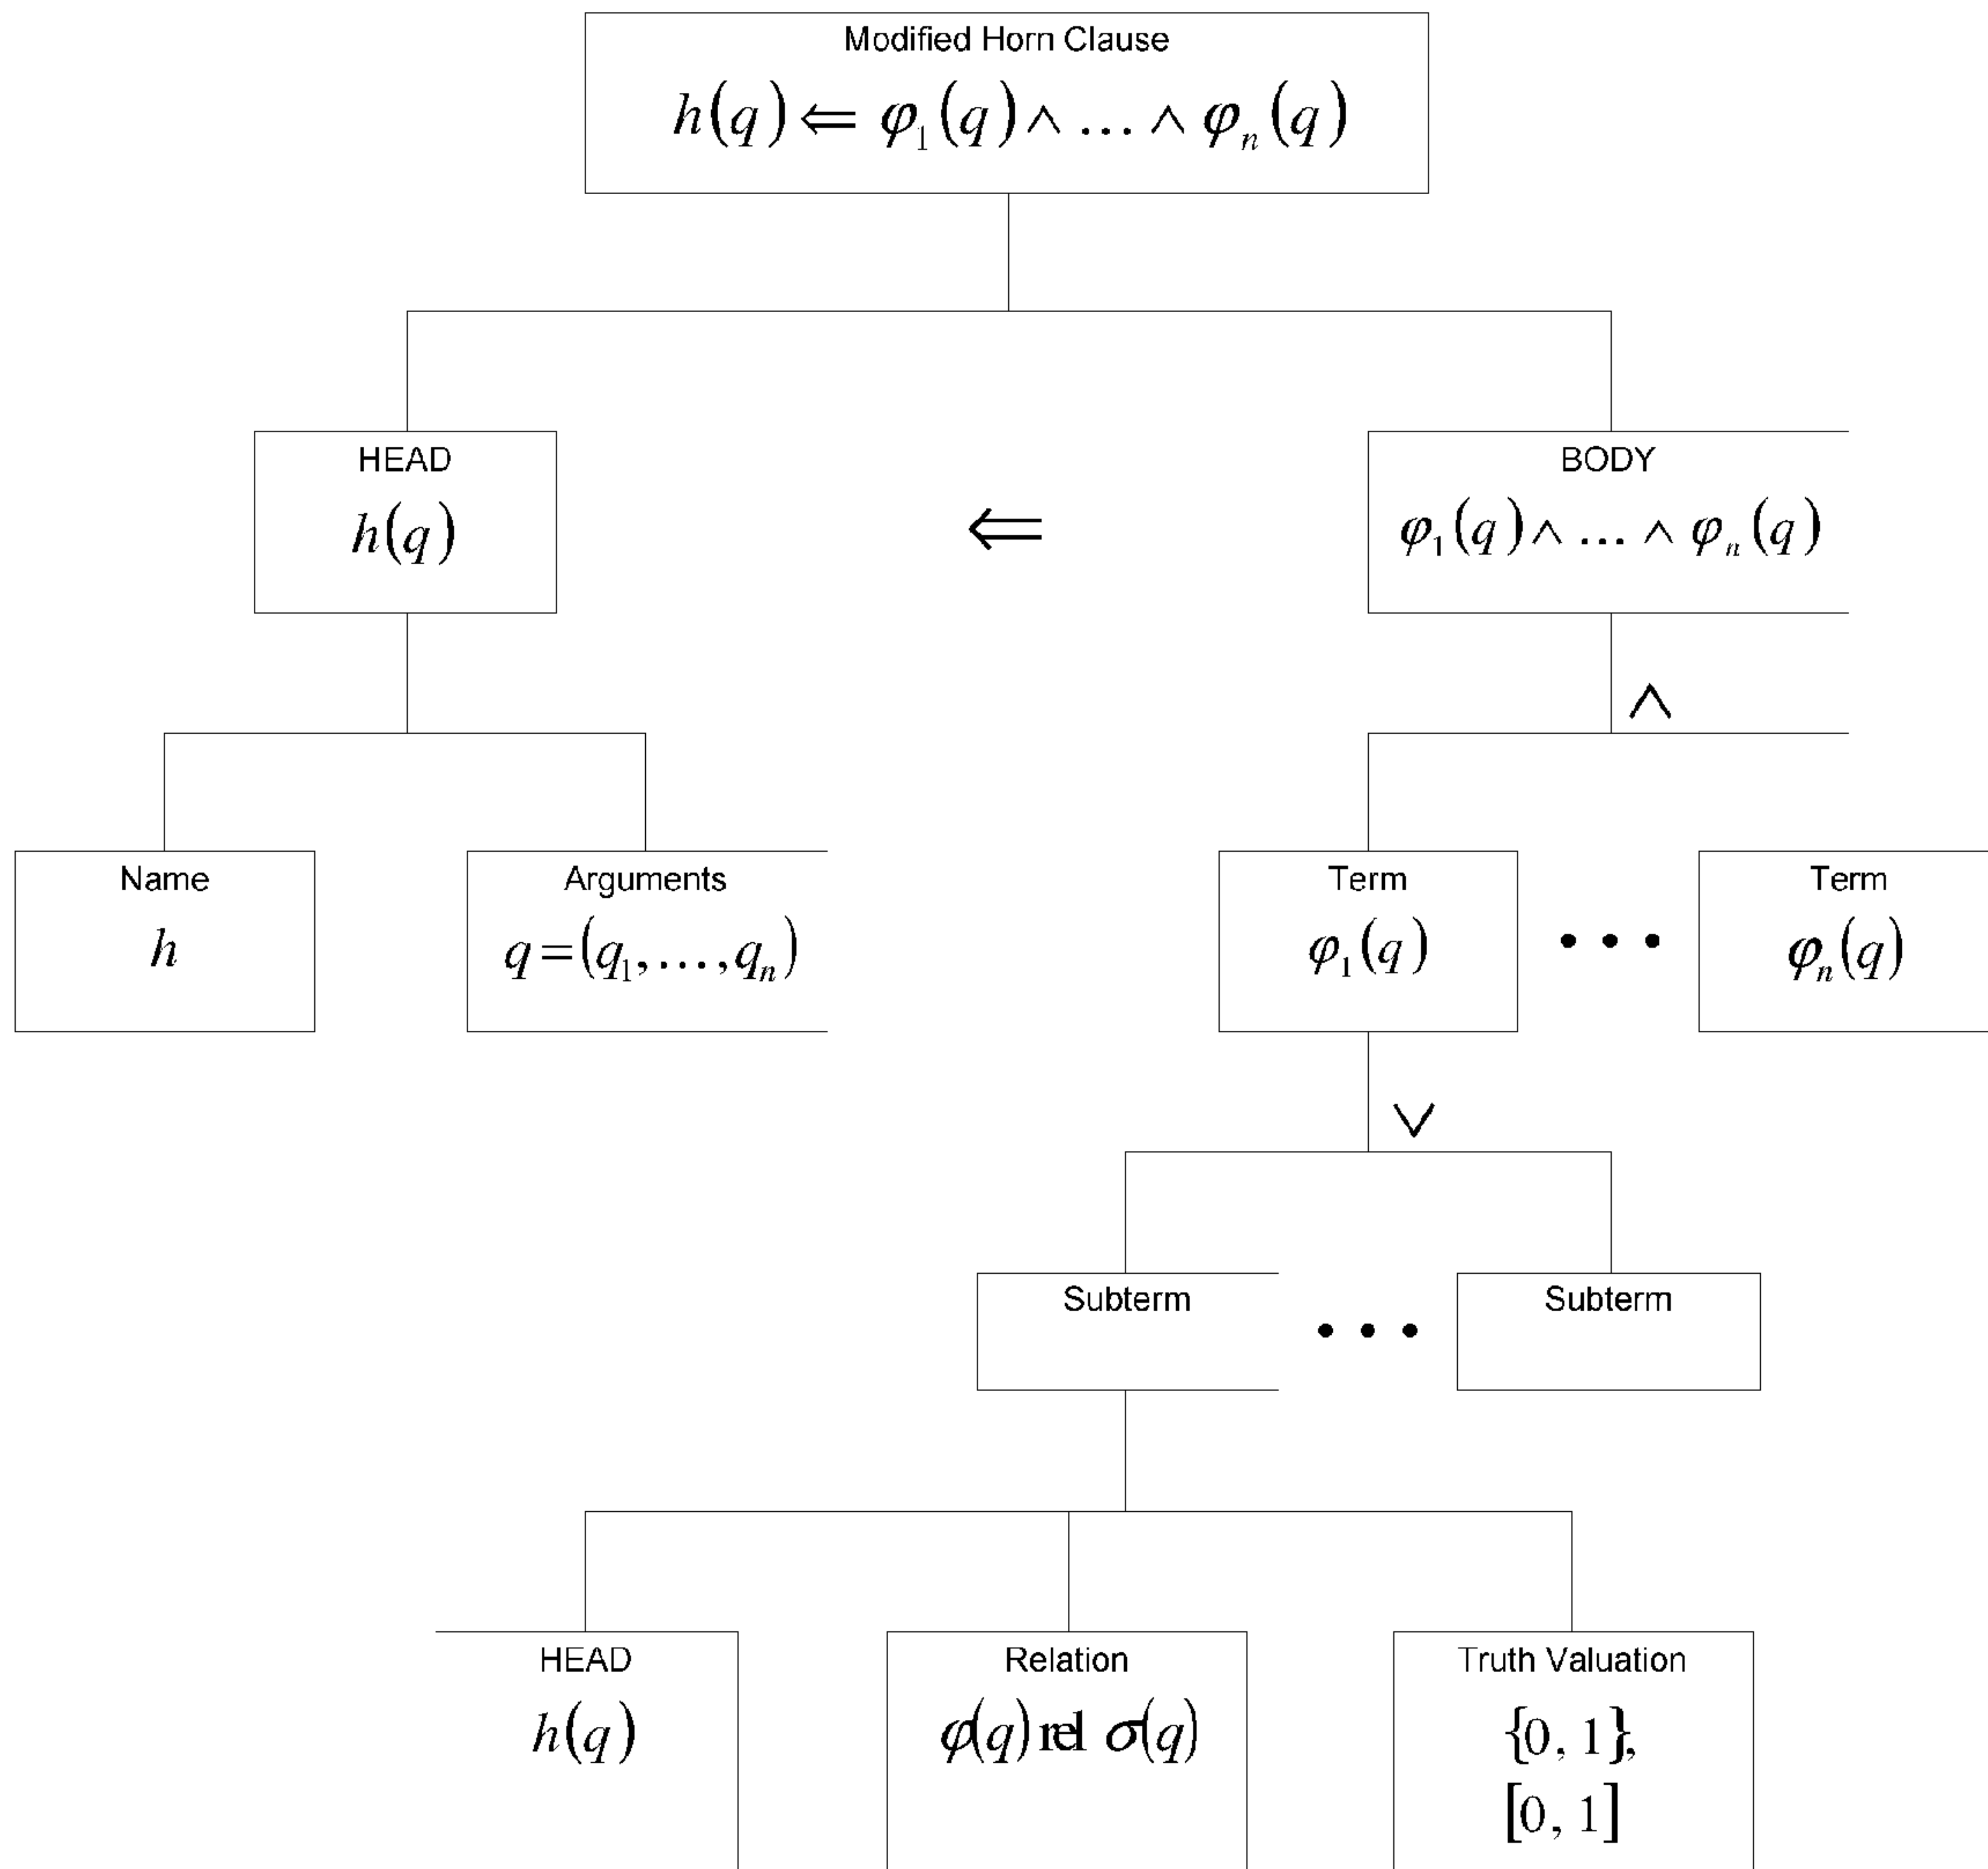
[0499] The format of the restrictive Horn Clauses in the IHDB is illustrated in Equation 3.2-2. A Horn Clause is an object composed of two objects a HEAD and a BODY connected by backward Implication ( $\leftarrow$ ) The logic implication transfers the logic value of the BODY to the HEAD. If the rule is an absolute rule or a hard rule, the logic value is 1 (if the BODY is logically true) or 0 (if the BODY is logically false). If the rule is a soft rule, the logic value transferred by the body is any number in [0,1].

[0500] The HEAD is a data structure composed of two objects: A name,  $h$ , and a list of arguments described by the argument vector  $q=(q^{(1)}, \dots, q^{(m)})$ . The list of arguments includes variables and parameters. The variables take values in the domain of the rule and the parameters are constants passed to the rule and unchanged by the instantiation of the rule. The domain of the rule is a set of values that each of its variables can take. In general, variables can take values over numerical or symbolic domains. As an example, a symbolic domain can be a list of diseases. A numeric domain can be a set of pairs of numbers representing blood pressure.

[0501] For the applications of CDI, the domains for variables are: real numbers, complex numbers (floating point and floating point complex numbers), integer numbers, binary numbers and symbolic token on finite domains.

[0502] The BODY of a clause is a data structure composed of one or more terms, denoted  $\phi_i(q)$ . The composition operation is extended- and, denoted by:  $\wedge$ . The extended- and works as a regular and in absolute rules and hard rules and as a functional product<sup>2</sup> on soft rules.

[0503] A rule with a head but not a body is called a fact. A fact's truth value is determined on the basis of the instantiation of its variables.





**[0504]** Each term in the body of a rule is an extended disjunction (or denoted by  $\vee$ ) of sub-terms. The  $\vee$  operator behaves like the standard- or for absolute and hard rules and behaves in a functional form, described later, when connecting sub-terms encoding heads of soft rules.

**[0505]** A sub-term is either the HEAD of a rule, a relation or a truth valuation (TV). When it is a HEAD it may have the same name as the one in the HEAD of the rule but with different arguments. This provides a recursive mechanism for rule evaluation.

**[0506]** When a rule has a sub-term that is the head of another rule it is said that the two rules are chained together by the corresponding sub-term. Note that a rule can be chained to several rules via corresponding sub-terms.

**[0507]** Constraint Domains

**[0508]** Constraint domains augment the BODY clause of Horn clauses to facilitate dynamic programming. Constraints are specified as a relationship between terms. Define the relationship between two terms

$$\phi(q) \text{rel} \sigma(q). \quad (3.2-4)$$

as a member of the following set

$$\text{rel} \in \{=, \neq, \leq, \geq, \text{statistical propagation}\}. \quad (3.2-5)$$

**[0509]** A relation can be of two types numeric or symbolic. Numeric relations establish equational forms between two functional forms. (For the initial phase only polynomial and affine linear functional forms will be considered.)

**[0510]** In general, an equational form is a set of one or more relations. For numeric relations,  $\phi(q) \text{rel} \sigma(q)$ ,  $\text{rel} \in \{=, \neq, \leq, \geq, <, >, \text{statistical propagation}\}$ . The table below gives the relations considered and their symbols.

TABLE EE

Numeric Relation	Symbol	Code Form
Equality	=	$\phi = \sigma$
Disequation	$\neq$	$\phi \neq \sigma$
Less-inequality	<	$\phi < \sigma$
Less-Equal	$\leq$	$\phi \leq \sigma$
Great-equality	>	$\phi > \sigma$
Great-inequality	$\geq$	$\phi \geq \sigma$

**[0511]** The adopted code forms are the ones used in constraint logic programming.

**[0512]** A symbolic relation can be of two types: inclusion and constraint. Inclusion relations are of the form:

$$q \in \text{Set} \quad (3.2-6)$$

Where  $x$  is a variable or a parameter,  $\epsilon$  is the inclusion symbol and Set is a set of symbolic forms or a set of numbers or a composite set of the form shown in the table below.

TABLE FF

Composite Set	Symbol	Code Form
Intersection	$\cap$	Set1 $\wedge$ Set2
Union	$\cup$	Set1 $\vee$ Set2
Complement	$\setminus$	$\setminus \text{Set}$

Constraint forms of the symbolic relational type may be one or a set of the forms presented in the table below For numeric relations,  $\phi(q) \text{rel} \sigma(q)$ ,  $\text{rel} \in \{=, \neq, <, >, \leq, \geq\}$ .

TABLE GG

Symbolic Relation	Symbol	Code Form
Equal	=	$\phi = \sigma$
Not Equal	$\neq$	$\phi \neq \sigma$
Is Contained	$\subset$	$\phi < \sigma$
Contains	$\supset$	$\phi > \sigma$
Is Contained or Equal	$\subset \cup =$	$\phi \leq \sigma$
Contains or Equal	$\supset \cup =$	$\phi \geq \sigma$

**[0513]** A TV is either a variable or a constant with values in the interval [0,1]. The TV of a Horn Clause that is an absolute rule or a hard rule can only take two values: 1 or 0. The TV when instantiated is 0 or 1. If the TV for an absolute or hard rule is 1, the rule is said to be in inactive state; if the TV is 0, the rule is said to be in active state.

**[0514]** The TV,  $T_i$ , of a soft rule satisfies

$$0 \leq T_i \leq 1. \quad (3.2-7)$$

If  $T_i$  above satisfies,

$$T_i \geq T_{\text{threshold}} \quad (3.2-8)$$

the soft clause is said to be in inactive state. If

$$T_i < T_{\text{threshold}} \quad (3.2-9)$$

the soft clause is said to be in active state, where  $T_{\text{threshold}}$  is a constant in [0,1] defined for each soft clause. The default value is 0.5.

**[0515]** This concludes the description of the knowledge representation. The instantiation process of the goal in a DE, as function of its knowledge base, is carried out by the inference engine of the DE. This process is the central component of CDI and will be described later on the document.

## Summary of Terminology

**[0516]** The following table summarizes the terminology we have just reviewed.

TABLE HH

Reference term	Definition
proposition	Defined as a construct as in the propositional calculus where the proposition takes on the value of true or false.
term	Recursively according to its assigned sub-term.
sub-term	A sub-term may be a Horn clause, a relation between two other sub-terms or an extended truth valuation depending on the context of absolute, hard or soft rules. In the case of absolute and hard rules it may be evaluated as a proposition. In the case of soft rules it takes a value on the interval [0, 1] and is considered to be active or true in the case that it exceeds its specific threshold.
Horn clause	A disjunction of terms with at most one positive term.
definite clause	A Horn clause with exactly one positive term.
goal clause	A Horn clause with no positive terms.
fact	A definite clause with no negative terms.
head	The positive term of a definite clause.
inactive state	The case when a rule will not apply for constrained optimization.

TABLE HH-continued

Reference term	Definition
active state	The case when a rule will apply for constrained optimization.
truth value, TV	The value that is used to determine whether a rule is active or inactive.

## Horn Clause Example

**[0517]** The following example illustrates a Horn clause:

**[0518]** has fever(name, temperature, white\_count, heart-rate, blood\_pressure)

$\Leftarrow$  (temperature >37)

$\wedge$  ((heartrate  $\geq$ 70)  $\vee$  bp (name, temperature, blood\_pressure)

$\vee$  wc(name, white\_count)) (3.2-10)

The clause establishes under which conditions the patient of name name, has a fever. The variables in clause are:

**[0519]** name, temperature, white\_count, heartrate, blood\_pressure.

When instantiated they represent, respectively, the name of the patient, his current body temperature, his white blood cell count, his heart rate range, and his blood pressure.

**[0520]** The clause body includes other clauses: bp (blood pressure) and wc (white count).

**[0521]** This completes the specification of the rule-based framework. The next step is to specify a complete process for converting all rules of this form to a set of equations.

**[0522]** Rule Entry Interface (REI)

**[0523]** The Rule Entry Interface provides a mechanism for:

**[0524]** Providing an API for the entry of rules into the IHDB.

**[0525]** Validating the specification of rules to be inserted into the IHDB.

**[0526]** Routing the rules to the appropriate DE's for insertion to their respective KC's.

**[0527]** Rule Editor (RE)

**[0528]** The Rule Editor allows users to specify rules associated with the systems to be interrogated.

**[0529]** External Knowledge Base (EKB)

**[0530]** It may be distributed or not

**[0531]** It may be persisted or not

**[0532]** It may be persisted locally or remotely to an agent

**[0533]** It may or may not be architecturally co-located with the IHDB

**[0534]** A sensor may include any source of data used by the agent

**[0535]** It may use various types of buses for data communication

**[0536]** Sensors may or may not be co-located with agents/DEs

**[0537]** Rule Conversion Engine (RCE)

**[0538]** The rule conversion engine converts rules of the IHDB into equations.

**[0539]** Method for Specification of a Simple Term as an Equation

**[0540]** Consider the term  $\phi(q)$  with the following truth assignment.

$$\varphi(q) = \begin{cases} T & q \in \mathcal{D}_\varphi \\ F & q \in \mathcal{D}_\varphi \subset \mathcal{D} \end{cases} \quad (3.6-1)$$

Then we can define the set of arguments which yield positive truth assignment.

$$\{q \in \mathcal{D}_\varphi \mid \phi(q) \leftarrow T\}. \quad (3.6-2)$$

Define the corresponding equation  $\check{\phi}(q)$  of the term  $\phi(q)$  as

$$\check{\phi}(q) = \begin{cases} 1 & \varphi(q) \leftarrow T \\ 0 & \varphi(q) \leftarrow F \end{cases} \quad (3.6-3)$$

Then extend the range of  $\check{\phi}(a)$  to the closed unit interval

$$\check{\phi}(q) \rightarrow [0,1]. \quad (3.6-4)$$

Revisiting the taxonomy of absolute, hard and soft rules, we recognize that hard and soft rules (terms in this example) can take values along the interval

$$0 \leq \check{\phi}(q) \leq 1. \quad (3.6-5)$$

And for absolute rules we add the constraint  $\check{\phi}(a) \rightarrow \{0,1\}$

$$\check{\phi}(q)(1 - \check{\phi}(q)) = 0. \quad (3.6-6)$$

**[0541]** Conversion of Fundamental Clauses of Propositional Calculus to Equations

**[0542]** Define the following notation for the propositional calculus.

TABLE II

Symbol	Function
$\wedge$	And
$\vee$	Or
$\Rightarrow$	Implication
$\sim$	Not
$\exists$	Exists
$\forall$	All
Fuzzy rule	

**[0543]** Theorem 3.6.1.

**[0544]** Given the method for the specification of equations from propositions, we prove the following transformations.

TABLE JJ

Proposition	Equation
$\sim\phi(q)$	$1 - \check{\phi}(q)$
$\phi(q) \wedge \sigma(q)$	$\check{\phi}(q) \cdot \check{\sigma}(q)$
$\phi(q) \vee \sigma(q)$	$\check{\phi}(q) + \check{\sigma}(q) - \check{\phi}(q) \cdot \check{\sigma}(q)$
$\phi(q) \Rightarrow \sigma(q)$	$1 - \phi(q) + \phi(q) \cdot \sigma(q)$
$\phi_1(q) \wedge \phi_2(q) \wedge \dots \wedge \phi_{k-1}(q) \wedge \phi(q) \Rightarrow \phi(q)$ (tail recursive)	$\check{\phi}(n, q) = \frac{\check{h}(n-1, q)}{\check{\sigma}(n, q) \check{\phi}(n-1, q) - 1}$

**[0545]** Proof by enumeration for equational representation of conjunction



**[0546]** Define the function  $\check{r}(q; \phi, \sigma)$  which represents the equation corresponding to conjunction ( $\wedge$ ). Verify by enumeration the correspondence of the mathematical equation values corresponding to the mapping  $T \rightarrow 1$  and  $F \rightarrow 0$ .

TABLE KK

$\phi(q)$	$\wedge$	$\sigma(q)$	$\check{r}(q; \phi, \sigma) = \check{\phi}(q) \cdot \check{\sigma}(q)$
T	T	T	1 = 1 · 1
T	F	F	0 = 1 · 0
F	F	T	0 = 0 · 1
F	F	F	0 = 0 · 0

**[0547]** Proof by enumeration for equational representation of disjunction

**[0548]** Define the function  $\check{r}(q; \phi, \sigma)$  which represents the equation corresponding to disjunction ( $\vee$ ). Verify by enumeration the correspondence of the mathematical equation values corresponding to the mapping  $T \rightarrow 1$  and  $F \rightarrow 0$ .

TABLE LL

$\phi(q)$	$\vee$	$\sigma(q)$	$\check{r}(q; \phi, \sigma) = \check{\phi}(q) + \check{\sigma}(q) - \check{\phi}(q) \cdot \check{\sigma}(q)$
T	T	T	1 = 1 + 1 - 1 · 1
T	T	F	1 = 1 + 0 - 1 · 1
F	T	T	1 = 0 + 1 - 0 · 1
F	F	F	0 = 0 + 0 - 0 · 0

**[0549]** Proof by enumeration for equational representation of negation

**[0550]** Define the function  $\check{r}(q; \phi, \sigma)$  which represents the equation corresponding to negation ( $\neg$ ). Verify by enumeration the correspondence of the mathematical equation values corresponding to the mapping  $T \rightarrow 1$  and  $F \rightarrow 0$ .

TABLE MM

$\phi(q)$	$\neg\phi(q)$	$\check{r}(q; \phi, \sigma) = 1 - \check{\phi}(q)$
T	F	0 = 1 - 1
F	T	1 = 1 - 0

**[0551]** Proof by enumeration for equational representation of implication

**[0552]** Define the function  $\check{r}(q; \phi, \sigma)$  which represents the equation corresponding to disjunction ( $\Rightarrow$ ). First note the equivalence of

$$\phi(q) \Rightarrow \sigma(q) \text{ and } \sim\phi(q) \vee \sigma(q) \quad (3.6-7)$$

Verify by enumeration the correspondence of the mathematical equation values corresponding to the mapping  $T \rightarrow 1$  and  $F \rightarrow 0$ .

TABLE NN

$\sim\phi(q)$	$\vee$	$\sigma(q)$	$\check{r}(q; \phi, \sigma) = 1 - \check{\phi}(q) + \check{\phi}(q) \cdot \check{\sigma}(q)$
T	T	T	1 = 1 - 1 + 1 · 1
T	F	F	0 = 1 - 1 + 1 · 0
F	T	T	1 = 1 - 0 + 0 · 1
F	T	F	1 = 1 - 0 + 0 · 0

Proof for equational representation of tail recursion

**[0553]** Tail recursion is propositionally defined as

$$\phi(q) \Leftarrow \phi_1(q) \wedge \phi_2(q) \wedge \dots \wedge \phi_{k-1}(q) \wedge \phi(q) \quad (3.6-8)$$

where  $s$  represent the current state. To develop an equational representation of the recursive formulation, first define the general function  $\check{\phi}(n, q)$  where  $n$  represents the  $n^{\text{th}}$  iteration of the tail recursion and  $\check{\phi}(n, q)$  is the logical consequent. Then rewrite the above formulation using the recursive step.

$$\check{\phi}_1(n, q) \wedge \check{\phi}_2(n, q) \wedge \dots \wedge \check{\phi}_{k-1}(n, q) \wedge \check{\phi}(n-1, q) \Rightarrow \check{\phi}(n, q) \quad (3.6-9)$$

Define

**[0554]**

$$\check{\sigma}(n, q) \Leftarrow \check{\phi}_1(n, q) \wedge \check{\phi}_2(n, q) \wedge \dots \wedge \check{\phi}_{k-1}(n, q) \quad (3.6-10)$$

**[0555]** Then the tail recursion is rewritable as

$$\check{\sigma}(n, q) \wedge \check{\phi}(n-1, q) \Rightarrow \check{\phi}(n, q) \quad (3.6-11)$$

According to the equational representation of implication, let

$$\check{h}(n-1, q) = 1 - \check{\sigma}(n, q) \cdot \check{\phi}(n-1, q) + \check{\sigma}(n, q) \cdot \check{\phi}(n-1, q) \quad (3.6-12)$$

Since by definition  $\check{r}(n-1, q) = \check{\sigma}(n, q) \cdot \check{\phi}(n-1, q)$ . Then

$$\check{\phi}(n, q) = \frac{\check{h}(n-1, q) + \check{\sigma}(n, q) \cdot \check{\phi}(n-1, q) - 1}{\check{\sigma}(n, q) \cdot \check{\phi}(n-1, q)} \quad (3.6-13)$$

with boundary condition  $n=0$ .

**[0556]** Converting rules based system of inference to the problem of constrained minimization

**[0557]** Consider the following example.

**[0558]** Converting rules to constraints

The preceding discussion has established an algorithm for convert rules of the form

$$h(q) \Leftarrow \phi_1(q) \wedge \phi_2(q) \wedge \dots \wedge \phi_m(q) \quad (3.6-14)$$

To constraints of the form

$$\check{h}(q) = \check{\phi}_1(q) \cdot \check{\phi}_2(q) \cdot \dots \cdot \check{\phi}_m(q). \quad (3.6-15)$$

**[0559]** Decision Element (DE)

**[0560]** A diagram of the Decision Element Architecture is shown in illustration OO.

**[0561]** It is composed six elements:

**[0562]** Programmable search engine (PSE)

**[0563]** Internal heterogeneous database (IHDB)

**[0564]** Inference engine (IE)

**[0565]** Inference rule base (IRB)

**[0566]** API/user interface

**[0567]** Network interface (NI)

**[0568]** List of External Repositories (LER)

**[0569]** A DE has a List of External Repositories (LER). Each entry in an LER includes 1) a protocol, 2) a heading sub-list, and 3) a translation grammar. Each protocol entry prescribes the access procedure to the corresponding external knowledge repository. Each heading sub-list entry contains a summary of the knowledge contents of the corresponding repository. Finally, each translation grammar entry provides a procedure for converting knowledge elements of the corresponding repository in to the rule representation in the IHDB of the DE.



**[0570]** Programmable Search Engine

**[0571]** The programmable search engine implements a standard hashing algorithm for detecting active rules as a function of the current instantiation of the variables in a variable buffer (VB) of the IE, and the contents of the active rule buffer (ARB). The VB contains the variables that form part of the query and all additional variables incorporated to this buffer during the inference process (IP). The VB includes all relevant data from the EKB beneficial to perform the query. The IP will be described below. The ARB contains all the currently active rules in the IP.

**[0572]** The search hashing algorithm is characterized by the search rules in the Inference Rule Base.

**[0573]** Internal Heterogeneous Database

**[0574]** The IHDB is the repository of the application clauses associated with the DE. These encode the domain of knowledge characterizing the expertise of the DE. For example in a medical application, a decision element may deal with expertise on heart illnesses, and the corresponding clauses might encode diagnoses and treatments for these diseases.

**[0575]** Inference Ermine

**[0576]** The IE encodes an algorithm, the IP, for assigning values to the variables appearing in the query.

**[0577]** Inference Rule Types

**[0578]** The DE incorporates inference rules (IR) that are a collection of rules for transforming and inferring instantiations of the goal. These rules provide the Inference Engine with directives for processing database rules to give a satisfactory instantiation to a given query or to request additional information so that a satisfactory instantiation can be generated. They are organized according to their functionality as follows. (See Illustration PP)

**[0579]** Equation Rules

**[0580]** These rules include the formal rules for inference. This includes all rules for natural language modeling from first principles.

**[0581]** Optimizer Rules

**[0582]** These rules include rules for finding the interior point in optimization.

**[0583]** Search Rules

**[0584]** These rules include rules for identifying the nature of insufficient potential. The goal is to apply these rules to acquire additional information required to satisfy the optimization goal.

**[0585]** Adaptation Rules

**[0586]** Adaptation rules are used to update the soft rules to relax them further to reduce the complexity and constrains of the optimization problem. The adaptation also serves to update the search rules to improve information acquisition.

**[0587]** Language Statistics and Pattern Rules

**[0588]** These rules embody the machine learning models.

**[0589]** Network Rules

**[0590]** These rules define how information is distributed over the network and what information is available from which resources.

**[0591]** Hybridization Rules

**[0592]** The rules define how other rules may be combined.

**[0593]** User Interface

**[0594]** The UI provides the utilities for entering queries, pragma rules, displaying query answers, status and for general interaction with the IE.

**[0595]** Network Interface

**[0596]** The NI provides a generic mechanism for interacting with other DE's via a procedure termed companionship. The companionship procedure implements the active coupling for the cooperation of the DE's in query resolution. This procedure is not hierarchical and implements a Pareto Agreement set strategy as the mechanism for CDI.

**[0597]** Query Language Interface (QLI)

**[0598]** Information about the QLI is available elsewhere herein.

**[0599]** Process for Determining Active Constraints

**[0600]** The process for determining active constraints is available elsewhere herein.

**[0601]** Minimization Function Generator (MFG) and Determining Active Constraints

**[0602]** The minimization function generator converts a query to a minimization function. Again, we assume without loss of generality the entire set of canonical coordinates  $q$  is an argument to any proposition  $\phi_i$ . In practice, we may further assume it is possible to apply the particular required coordinates as need to the proposition or function in question. Then let  $\phi^{(k)}$  be the set of propositions associated with  $DE_k$  in the context of query  $Q$ . These propositions are composed of the proposition associated with the query  $\phi_Q(q)$ , and other propositions  $\phi_i(q)$ , comprising the constraints of the system. The proposition  $\phi_Q(q)$  associated with a given query  $Q$  can be converted to an equation  $\check{\phi}_Q(q)$ . Queries that are satisfiable specify a set.

$$\{q | \phi_Q(q) \leftarrow T\} \quad (3.10-1)$$

Similarly, a satisfied query represented as an equation is also a set

$$\{q | \check{\phi}_Q(q) = 1\}. \quad (3.10-2)$$

Relaxing the values that  $\check{\phi}_Q(\bullet)$  can take to include the unit interval so that soft rules are incorporated yields the following constrained optimization expression. Let  $J(q) = (\check{\phi}_Q(q) - 1)^2$ . Then specify the optimization

$$\min_q j(q) \bar{X}_o \quad (3.10-3)$$

Subject to:

**[0603]** 1.  $\check{\phi}_Q(q) \leq 1$

**[0604]** 2.  $\check{\phi}_Q(q) \geq 0$

**[0605]** 3. A knowledge base on the set  $\{\check{\phi}_i(q), \dots, \check{\phi}_n(q), \dots, \check{\phi}_{n+s}(q)\} \subseteq \check{\phi}^{(k)}$  which represents a further set of active constraints specific to the problem:

**[0606]** a.  $\check{\phi}_i(q) \geq 0$  for  $1 \leq i \leq n$ ,

**[0607]** b.  $\check{\phi}_i(q) \leq 1$  or, equivalently  $-(\check{\phi}_i(q) - 1) \geq 0$  for  $1 \leq i \leq n$ ,

**[0608]** c. and in the case of absolute rules  $\check{\phi}_i(q)(1 - \check{\phi}_i(q)) = 0$  for  $n < 1 \leq n + s$ .

Introduce the indicator functions

$$V_{\check{\phi}_i} = \begin{cases} 0 & \check{\phi}_i(q) \geq 0 \\ \infty & \check{\phi}_i(q) < 0 \end{cases} \quad (3.10-4)$$



-continued

and

$$V_{\varphi_i}^+ = \begin{cases} 0 & 1 - \check{\varphi}_i(q) \geq 0 \\ \infty & 1 - \check{\varphi}_i(q) < 0 \end{cases} \quad (3.10-5)$$

which yields the two logarithmic barrier functions

$$\check{V}_{\check{\Phi}_i^-} = -\log(\check{\Phi}_i^-(q)) \quad (3.10-6)$$

and

$$\check{V}_{\check{\Phi}_i^+} = -\log(\check{\Phi}_i^+(q)) \quad (3.10-7)$$

According to the method of Lagrange multipliers, combine this with the equality constraints to form the static Lagrangian function

$$\begin{aligned} \mathcal{L}(q; \check{\varphi}_Q, \check{\varphi}^{(k)}, \omega_1^{(+)}, \dots, \omega_n^{(+)}, \omega_{n+1}^{(-)}, \dots, \\ \omega_{2n}^{(-)}, \omega_{2n+1}^{(\lambda)}, \dots, \omega_{2n+s}^{(\lambda)}, \omega_{2n+s+1}^{(Q)}, \omega_{2n+s+2}^{(Q)}) = \\ \check{\varphi}_Q(q) + \sum_{i=1}^n [\omega_i^{(+)} \check{V}_{\check{\varphi}_i^+} + \omega_{n+i}^{(-)} \check{V}_{\check{\varphi}_i^-}] + \sum_{l=1}^s \omega_{2l+1}^{(\lambda)} \check{\varphi}_1(q) (1 - \check{\varphi}_1(q)) - \\ \omega_{2n+s+1}^{(Q)} \log(\check{\varphi}_Q(q)) - \omega_{2n+s+2}^{(Q)} \log(1 - \check{\varphi}_Q(q)), \end{aligned} \quad (3.10-8)$$

the roots of which can be found using a formulation of Newton-Raphson. Since  $\mathcal{L}$  here includes absolute, hard and soft rules we may call it the total static Lagrangian for  $DE_k$  and refer to it as  $\mathcal{L}^{(T)}$ .

**[0609]** Construct Equations of Motion**[0610]** Information for equations of motion is available elsewhere herein.**[0611]** Query Response Engine (QRE) which Includes Process for Constructing Differential Equations**[0612]** Application of Newton-Raphson**[0613]** Consider a continuous analog of the independent variables of  $\mathcal{L}(\bullet)$ 

$$q = q(t) = \begin{bmatrix} q^{(1)}(t) \\ \vdots \\ q^{(v)}(t) \end{bmatrix} \quad (3.12-1)$$

where each of the  $v$  total independent variables of  $\mathcal{L}(\bullet)$  is mapped to its corresponding position in  $q(t)$ , the column vector that is represented with a lower-case  $q$ . To reiterate, the independent variable  $t$  refers algorithmic time as opposed to physical time which may also be represented in the system. The corresponding unconstrained optimization goal can be written as

$$\min_{q_1, \dots, q_v} \mathcal{L}(q^{(1)}(t), \dots, q^{(v)}(t)) \quad (3.12-2)$$

so that  $\nabla \mathcal{L}(q)$ 

$$\nabla \mathcal{L}(q(t)) = \begin{bmatrix} \frac{\partial \mathcal{L}}{\partial q^{(1)}} \\ \vdots \\ \frac{\partial \mathcal{L}}{\partial q^{(v)}} \end{bmatrix} = \begin{bmatrix} \nabla \mathcal{L}_1 \\ \vdots \\ \nabla \mathcal{L}_v \end{bmatrix} = 0, \quad (3.12-3)$$

with positive definite Hessian matrix

$$\begin{aligned} \nabla^2 \mathcal{L}(q(t)) = \begin{bmatrix} \frac{\partial^2 \mathcal{L}}{\partial q^{(1)} \partial q^{(1)}} & \cdots & \frac{\partial^2 \mathcal{L}}{\partial q^{(1)} \partial q^{(v)}} \\ \vdots & \ddots & \vdots \\ \frac{\partial^2 \mathcal{L}}{\partial q^{(v)} \partial q^{(1)}} & \cdots & \frac{\partial^2 \mathcal{L}}{\partial q^{(v)} \partial q^{(v)}} \end{bmatrix} = \begin{bmatrix} \nabla \mathcal{L}_{11} & \cdots & \nabla \mathcal{L}_{1v} \\ \vdots & \ddots & \vdots \\ \nabla \mathcal{L}_{v1} & \cdots & \nabla \mathcal{L}_{vv} \end{bmatrix} = \\ > 0. \end{aligned} \quad (3.12-4)$$

Write the recursion for Newton's method

$$q_{(k+1)}(t) = q_{(k)}(t) - (\nabla^2 \mathcal{L}(q_{(k)}(t)))^{-1} \nabla \mathcal{L}(q_{(k)}(t)). \quad (3.12-5)$$

This is equivalently rewritten

$$\frac{q_{(k+1)}(t) - q_{(k)}(t)}{\delta} = -\frac{1}{\delta} (\nabla^2 \mathcal{L}(q_{(k)}(t)))^{-1} \nabla \mathcal{L}(q_{(k)}(t)). \quad (3.12-6)$$

Via continualization we approximate the derivative

$$\dot{q}(t) = \frac{dq(t)}{dt} = -(\nabla^2 \mathcal{L}(q(t)))^{-1} \nabla \mathcal{L}(q(t)). \quad (3.12-7)$$

**[0614]** Translation of Inverted Matrix**[0615]** Consider  $M$ , an invertible and positive definite matrix. Then we make the following provable assertions.**[0616]** 1.  $A^T A$  is symmetric.**[0617]** 2.  $-A^T A$  has negative eigenvalues.

Define

**[0618]**

$$\frac{dM(t)}{dt} = -A^T A M(t) + A^T \quad (3.12-8)$$

Then as  $t \rightarrow \infty$ ,  $M(t) \rightarrow A^{-1} = \nabla^2 \mathcal{L}(q_{(k)}(t))^{-1}$ . Using (3.12-3) and (3.12-4) approximate  $\dot{q}(t)$  by rewriting the derivative in the context of  $M(t)$ . This yields the following two equations.

(3.12-9)

$$\dot{q}(t) = -M(t)\nabla L(q(t)) = \begin{bmatrix} m_{11} & \dots & m_{1v} \\ \vdots & \ddots & \vdots \\ m_{v1} & \dots & m_{vv} \end{bmatrix} \begin{bmatrix} \nabla \mathcal{L}_1 \\ \vdots \\ \nabla \mathcal{L}_v \end{bmatrix} = \begin{bmatrix} m_{11}\nabla \mathcal{L}_1 + \dots + m_{1v}\nabla \mathcal{L}_v \\ \vdots \\ m_{v1}\nabla \mathcal{L}_1 + \dots + m_{vv}\nabla \mathcal{L}_v \end{bmatrix}$$

(3.12-10)

$$\frac{dM(t)}{dt} =$$

$$\begin{aligned} & -(\nabla^2 \mathcal{L}(q(t)))^T (\nabla^2 \mathcal{L}(q(t))) M(t) + (\nabla^2 \mathcal{L}(q(t)))^T = - \begin{bmatrix} \nabla \mathcal{L}_{11} & \dots & \nabla \mathcal{L}_{1v} \\ \vdots & \ddots & \vdots \\ \nabla \mathcal{L}_{v1} & \dots & \nabla \mathcal{L}_{vv} \end{bmatrix} \begin{bmatrix} \nabla \mathcal{L}_{11} & \dots & \nabla \mathcal{L}_{1v} \\ \vdots & \ddots & \vdots \\ \nabla \mathcal{L}_{v1} & \dots & \nabla \mathcal{L}_{vv} \end{bmatrix} \begin{bmatrix} m_{11} & \dots & m_{1v} \\ \vdots & \ddots & \vdots \\ m_{v1} & \dots & m_{vv} \end{bmatrix} + \begin{bmatrix} \nabla \mathcal{L}_{11} & \dots & \nabla \mathcal{L}_{1v} \\ \vdots & \ddots & \vdots \\ \nabla \mathcal{L}_{v1} & \dots & \nabla \mathcal{L}_{vv} \end{bmatrix} = \\ & - \begin{bmatrix} \nabla \mathcal{L}_{11}^2 + \dots + \nabla \mathcal{L}_{v1}^2 & \dots & \nabla \mathcal{L}_{11}\nabla \mathcal{L}_{1v} + \dots + \nabla \mathcal{L}_{v1}\nabla \mathcal{L}_{vv} \\ \vdots & \ddots & \vdots \\ \nabla \mathcal{L}_{11}\nabla \mathcal{L}_{1v} + \dots + \nabla \mathcal{L}_{v1}\nabla \mathcal{L}_{vv} & \dots & \nabla \mathcal{L}_{1v}^2 + \dots + \nabla \mathcal{L}_{vv}^2 \end{bmatrix} \begin{bmatrix} m_{11} & \dots & m_{1v} \\ \vdots & \ddots & \vdots \\ m_{v1} & \dots & m_{vv} \end{bmatrix} + \begin{bmatrix} \nabla \mathcal{L}_{11} & \dots & \nabla \mathcal{L}_{1v} \\ \vdots & \ddots & \vdots \\ \nabla \mathcal{L}_{v1} & \dots & \nabla \mathcal{L}_{vv} \end{bmatrix} = \\ & - \begin{bmatrix} (\nabla \mathcal{L}_{11}^2 + \dots + \nabla \mathcal{L}_{v1}^2)m_{11} + \dots + & (\nabla \mathcal{L}_{11}^2 + \dots + \nabla \mathcal{L}_{v1}^2)m_{1v} + \dots + \\ (\nabla \mathcal{L}_{11}\nabla \mathcal{L}_{1v} + \dots + \nabla \mathcal{L}_{v1}\nabla \mathcal{L}_{vv})m_{v1} + & \dots & (\nabla \mathcal{L}_{11}\nabla \mathcal{L}_{1v} + \dots + \nabla \mathcal{L}_{v1}\nabla \mathcal{L}_{vv})m_{vv} + \\ \nabla \mathcal{L}_{11} & & \nabla \mathcal{L}_{v1} \\ \vdots & \ddots & \vdots \\ (\nabla \mathcal{L}_{11}\nabla \mathcal{L}_{1v} + \dots + \nabla \mathcal{L}_{v1}\nabla \mathcal{L}_{vv})m_{11} + \dots + & (\nabla \mathcal{L}_{11}\nabla \mathcal{L}_{1v} + \dots + \nabla \mathcal{L}_{v1}\nabla \mathcal{L}_{vv})m_{1v} + \dots + \\ (\nabla \mathcal{L}_{1v}^2 + \dots + \nabla \mathcal{L}_{vv}^2)m_{v1} + \nabla \mathcal{L}_{1v} & \dots & (\nabla \mathcal{L}_{1v}^2 + \dots + \nabla \mathcal{L}_{vv}^2)m_{vv} + \nabla \mathcal{L}_{vv} \end{bmatrix} = \begin{bmatrix} \nabla m_{11} & \dots & \nabla m_{1v} \\ \vdots & \ddots & \vdots \\ \nabla m_{v1} & \dots & \nabla m_{vv} \end{bmatrix} \end{aligned}$$

The approximation proceeds as follows, applying the Magnus expansion to compute the integral:

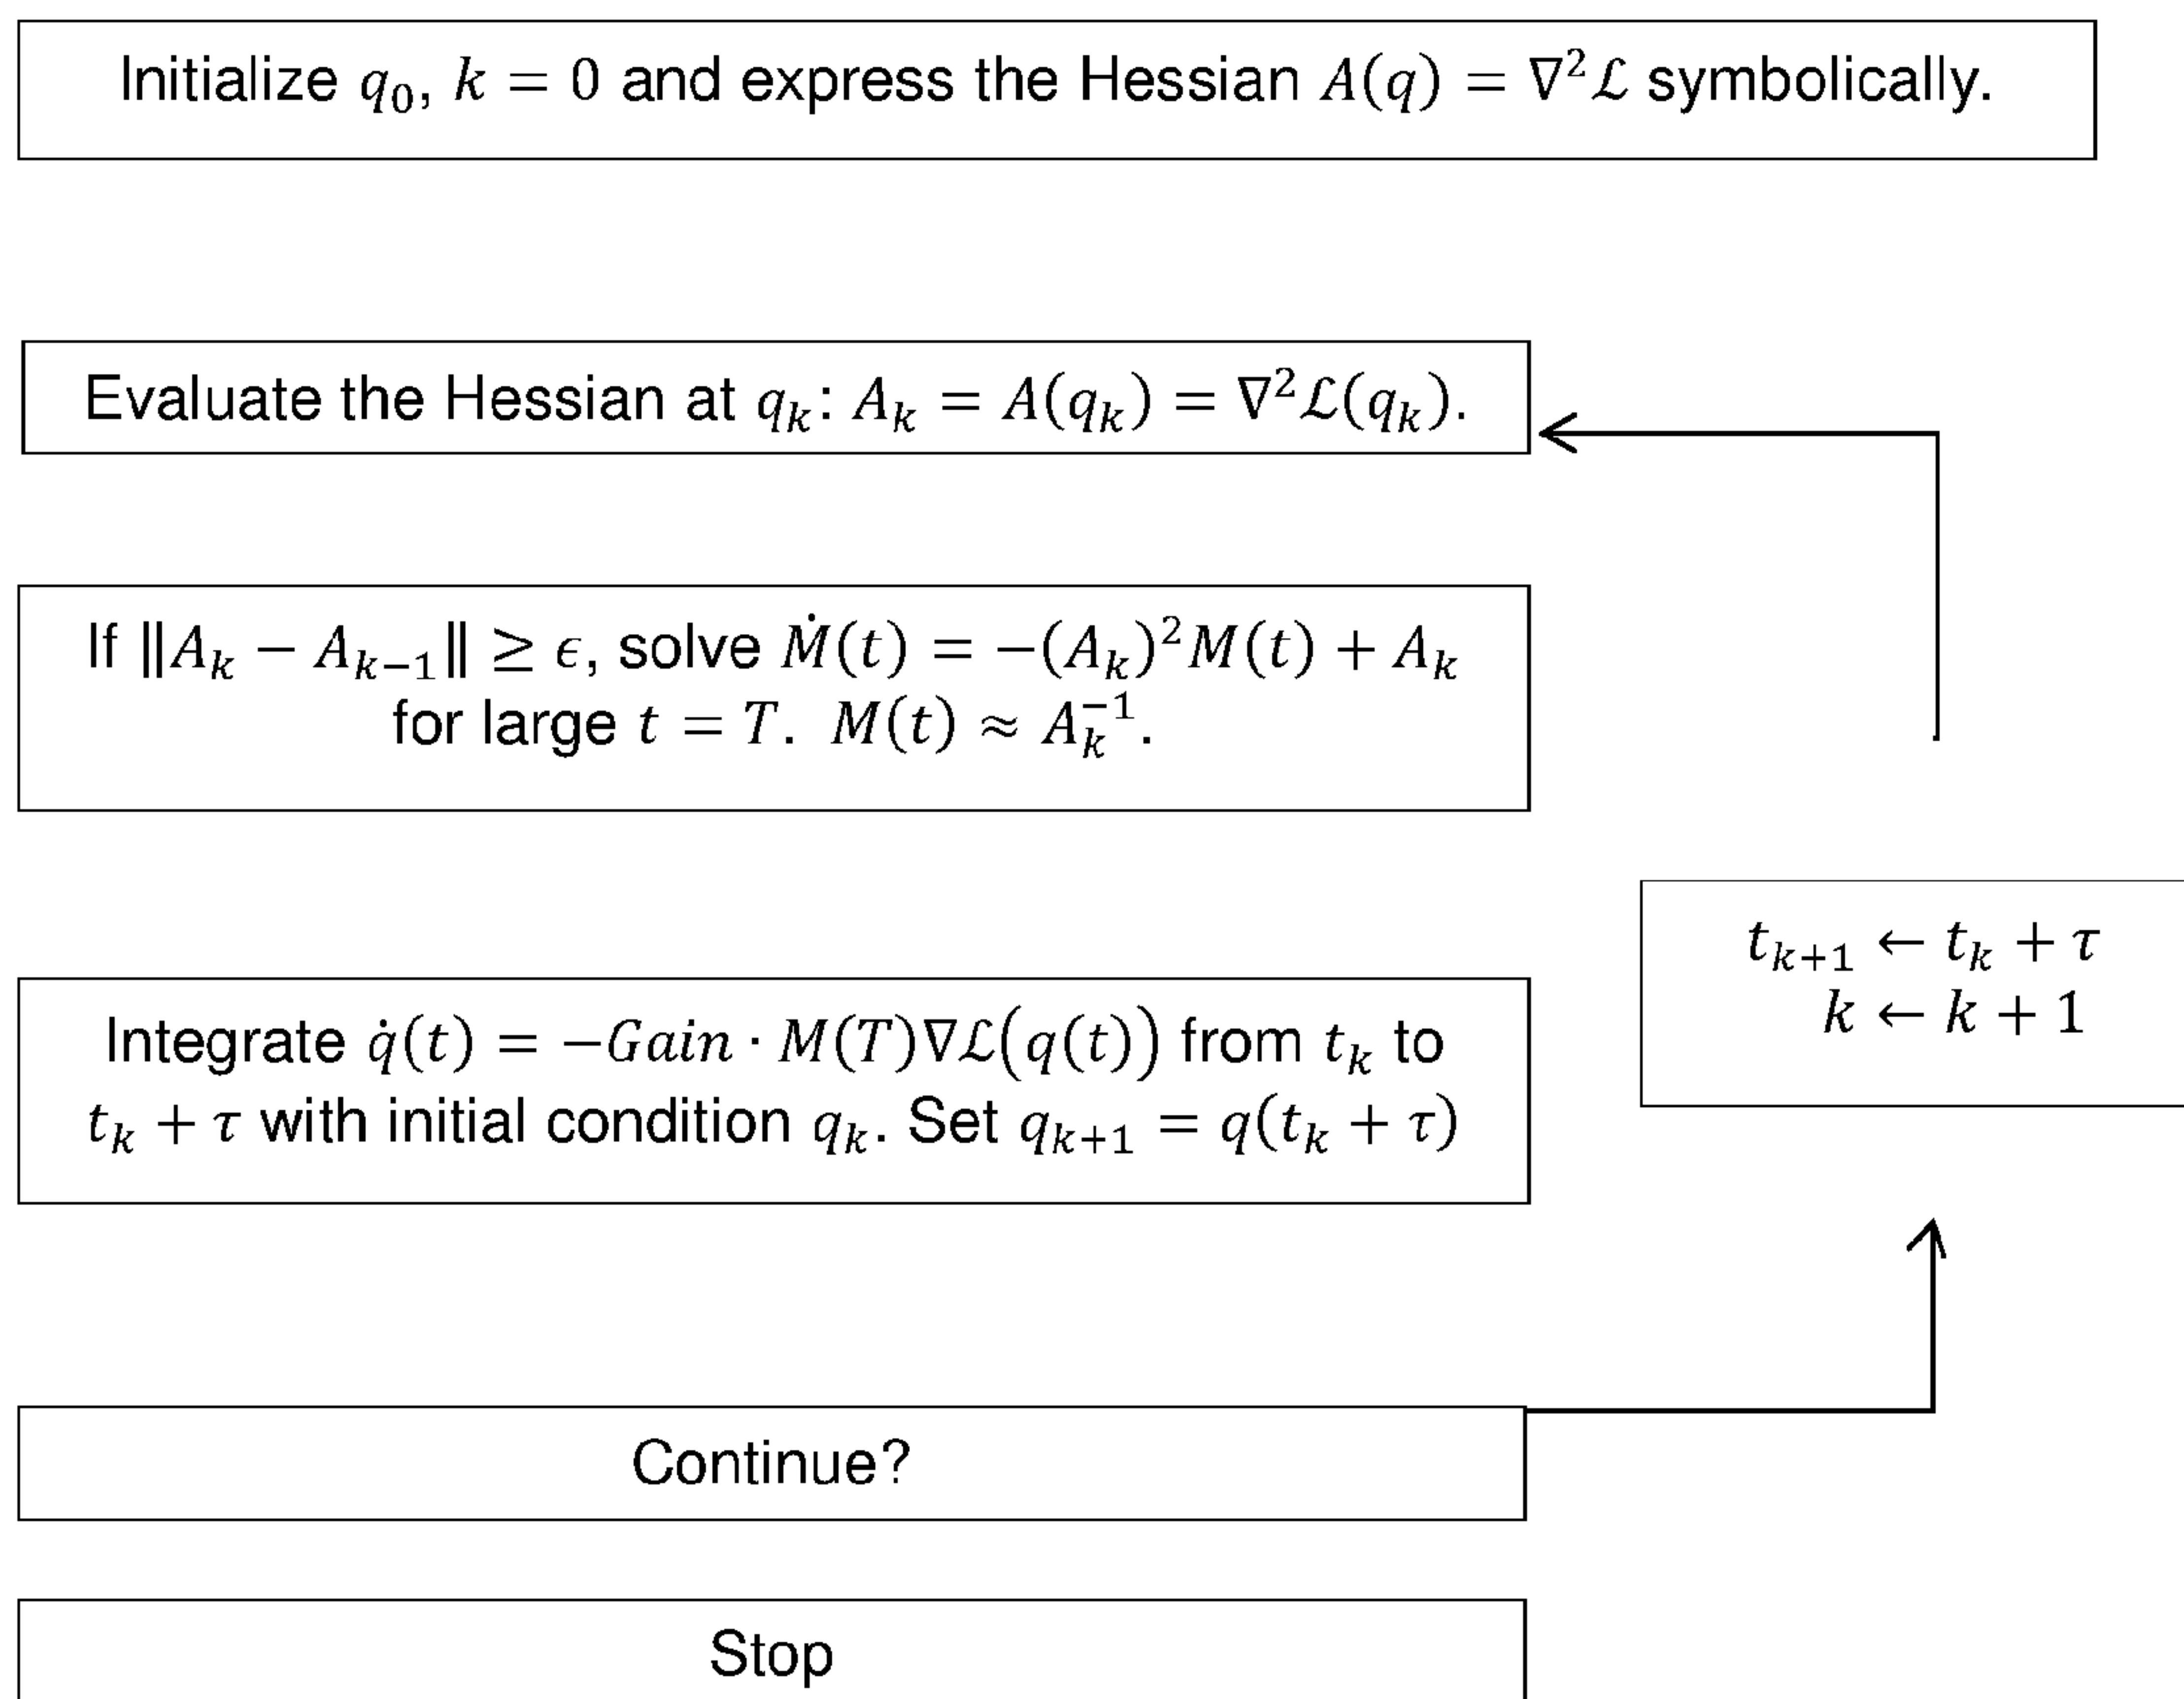
**[0619]** 1. Fix  $M(0) = \nabla^2 \mathcal{L}(q(0))$  and  $\nabla^2 \mathcal{L}(q(t))$ .

**[0620]** 2. Use the variation of constants formula to solve

$$M(T) = e^{-[\nabla^2 \mathcal{L}(q(T))]^2 T} M(0) + \left[ \int_0^T e^{-[\nabla^2 \mathcal{L}(q(\tau))]^2 (T-\tau)} d\tau \right] \nabla^2 \mathcal{L}(q(T))$$

**[0621]** The following is computation flow, flowing down unless otherwise indicated.





**[0622]** Process for Determining Dynamic Lagrangian Via Hemholtz Equations

Given

**[0623]**

$$G_i(\ddot{q}, \dot{q}, q) = \sum_{j=1}^n W_{i,j}(\dot{q}, q)(\dot{q}, q)\ddot{q}^{(j)} + K_i(\dot{q}, q) = 0 \quad (3.12-11)$$

$$j = 1, \dots, n$$

If the three conditions

$$\frac{\partial G_i}{\partial \ddot{q}^{(i)}} = \frac{\partial G_j}{\partial \ddot{q}^{(i)}}, \quad (3.12-12)$$

$$\frac{\partial G_i}{\partial \ddot{q}^{(j)}} + \frac{\partial G_j}{\partial \ddot{q}^{(i)}} = \frac{d}{dt} \left( \frac{\partial G_i}{\partial \dot{q}^{(j)}} + \frac{\partial G_j}{\partial \dot{q}^{(i)}} \right),$$

$$\frac{\partial G_i}{\partial \dot{q}^{(j)}} - \frac{\partial G_j}{\partial \dot{q}^{(i)}} = \frac{1}{2} \frac{d}{dt} \left( \frac{\partial G_i}{\partial \dot{q}^{(j)}} + \frac{\partial G_j}{\partial \dot{q}^{(i)}} \right),$$

with  $i, j=1, \dots, n$  hold, then

$$\sum_{j=1}^n \frac{\partial^2 L}{\partial \dot{q}^{(i)} \partial \dot{q}^{(j)}} \dot{q}^{(j)} + \frac{\partial^2 L}{\partial q^{(i)} \partial \dot{q}^{(i)}} - \frac{\partial L}{\partial q^{(i)}} = G_i, \quad (3.12-13)$$

$$i = 1, \dots, n$$

This is a second order, linear hyperbolic differential equation on the Lagrangian L. It can be solved efficiently by the method of characteristics.

Let

**[0624]**

$$G(\ddot{q}, \dot{q}, q) = \begin{bmatrix} q(t) \\ \dot{q}(t) \\ \dot{M}(t) \end{bmatrix} = \begin{bmatrix} q^{(1)} \\ \vdots \\ q^{(v)} \\ m_{11} \nabla \mathcal{L}_1 + \dots + m_{1v} \nabla \mathcal{L}_v \\ \vdots \\ m_{v1} \nabla \mathcal{L}_1 + \dots + m_{vv} \nabla \mathcal{L}_v \\ \nabla m_{11} \\ \vdots \\ \nabla m_{v1} \\ \vdots \\ \nabla m_{1v} \\ \vdots \\ \nabla m_{vv} \end{bmatrix} \quad (3.12-14)$$

**[0625]** Process for Determining Hessian Rank of Dynamic Lagrangian

**[0626]** Information for determining Hessian rank of dynamic Lagrangian is available elsewhere herein.

**[0627]** Converting the Lagrangian to the Hamiltonian Via the Legendre Transformation.

**[0628]** In our formulation the Lagrangian,  $L_k^{(T)}(q, \dot{q}; \omega)$ , may be converted to the Hamiltonian using the Legendre transformation, so that

$$H_k^{(T)}(q, p; \omega) = \frac{\partial L_k^{(T)}}{\partial \dot{q}} \dot{q} - L_k^{(T)}(q, \dot{q}; \omega) = p^T \dot{q} - L_k^{(T)}(q, \dot{q}; \omega) \quad (3.12-15)$$

**[0629]** Pareto Multi-Criteria Optimization Engine (PMOE)

**[0630]** Consider the problem of determining the relaxed Pareto optimal solution to a given system query at a given time step. There are N decision elements,  $k=1, \dots, N$ . A given decision element,  $DE_k$ , has the following associated parameters which are constituent to the ARB:

**[0631]** A generalized set of coordinates relevant to  $DE_k$ ,  $q$ .

**[0632]** A generalized set of linearly independent momenta  $\{p_a\}$  where the index a refers the linearly independent momenta selected from the canonical set p.

**[0633]** A set of control parameters  $\omega$  for hard a soft rules of the system, where  $0 \leq \omega_i \leq 1$ .

**[0634]** The ARB has the following components which determine the constraints of  $DE_k$ :

**[0635]** The Hamiltonian which identifies the fundamental dynamics of the system of the system for the k'th decision element denoted

$$H_k^{(\circ)}(q, \{p_a\}). \quad (3.13-1)$$

**[0636]** The summation of the first class constraints of the system, which is

$$\sum_i \omega_i f_i(q^{(i)}, \omega_i) \quad (3.13-2)$$

**[0637]** The summation of the second class constraints of the system which is

$$\sum_i g_i(q^{(i)}, \omega_i) \quad (3.13-3)$$

**[0638]** The Tellegen agent which is a function of the Hamiltonians of the absolute rules of the other N-1 decision elements in the system

$$H_k^{(A)} = F_k^{(A)}(H_1^{(T)}, \dots, H_{k-1}^{(T)}, H_{k+1}^{(T)}, \dots, H_N^{(T)}) \quad (3.13-4)$$

**[0639]** The total Hamiltonian of the system is denoted  $F^{(T)}$ .

**[0640]** Approximations to the various Hamiltonian's are denoted  $\hat{H}_k^{(A)}$ ,  $\hat{H}_k^{(T)}$  and  $\hat{H}_k^{(\circ)}$  for the Tellegen, total, and DE-level Hamiltonians respectively.

**[0641]** System Initialization

**[0642]** Determining the relaxed Pareto optimal point of the system is a process which includes:

**[0643]** Initialization of N decision elements.

**[0644]** Synchronization through companionship of each of the N decision elements with its respective Tellegen agent.



**[0645]** System Operation

**[0646]** Illustration SS shows how decision elements interact with the network, receive queries, and return results. In this example, the distributed system effectively implements an abstract classifier that has no real implementation. The DE's receive sensor data from the network which includes new available information which may benefit classification. The user submits a query that is received by a DE which then returns a result.

**[0647]** Gauge Systems in a Hamiltonian Domain

**[0648]** The time integral of the Lagrangian  $L(q, \dot{q})$  is the action  $S_L$  defined as

$$S_L = \int_{t_1}^{t_2} L(q, \dot{q}) dt$$

where

$$\dot{q} = \frac{dq(t)}{dt}.$$

The Lagrangian conditions for stationarity are first that

$$\frac{d}{dt} L_{\dot{q}^{(n)}} - L_{q^{(n)}} = 0 \quad (3.14-1)$$

where  $n=1, \dots, N$ ,

$$L_{q^{(n)}} = \frac{\partial L}{\partial \dot{q}^{(n)}}, \text{ and}$$

$$L_{\dot{q}^{(n)}} = \frac{\partial L}{\partial q^{(n)}}.$$

And, secondarily

$$\left[ \sum_{n'=1}^N \ddot{q}^{(n')} \right] L_{\dot{q}^{(n)} \dot{q}^{(n)}} = L_{q^{(n)}} - \dot{q}^{(n)} L_{q^{(n)} \dot{q}^{(n)}} \quad (3.14-2)$$

where

$$\ddot{q}^{(n')} = \frac{d^2 q^{(n')}}{dt^2} \text{ and}$$

$$L_{\dot{q}^{(n)} \dot{q}^{(n)}} = \frac{\partial^2 L}{\partial (\dot{q}^{(n)})^2}.$$

The generalized accelerations  $\ddot{q}^{(n)}$  are immediately determined if  $L_{\dot{q}^{(n)} \dot{q}^{(n)}}$  is invertible, or equivalently

$$\det(L_{\dot{q}^{(n)} \dot{q}^{(n)}}) \neq 0 \quad (3.14-3)$$

for  $i=1, \dots, N$ . If for some  $n$ ,  $\det(L_{\dot{q}^{(n)} \dot{q}^{(n)}}) = 0$ , the acceleration vector  $\ddot{q}^{(n)}$  will not be uniquely determined.

**[0649]** The departing point for the Hamiltonian approach is the definition of conjugate momentum

$$p_n = L_{\dot{q}^{(n)}} \quad (3.14-4)$$

where  $n=1, \dots, N$ . We will see that (3.14-3) is the condition of non-invertibility of

$$L_{\dot{q}\dot{q}} = \begin{bmatrix} L_{\dot{q}^{(1)} \dot{q}^{(1)}} & \dots & L_{\dot{q}^{(1)} \dot{q}^{(N)}} \\ \vdots & \ddots & \vdots \\ L_{\dot{q}^{(N)} \dot{q}^{(1)}} & \dots & L_{\dot{q}^{(N)} \dot{q}^{(N)}} \end{bmatrix}$$

of the velocities of the functions of the coordinates  $q$  and momenta  $p$ . In other words, in this case, the momenta defined in (3.14-4) are not all independent. Define the relations that follow from (3.14-4) as

$$\Phi_m(q, p) \quad (3.14-5)$$

where  $m=1, \dots, M$ . Write (3.14-4) in vector notation as

$$p = L_{\dot{q}}(q, \dot{q}).$$

Then compatibility demands

$$\Phi_m(q, L_{\dot{q}}(q, \dot{q})) = 0$$

is an identity with  $m=1, \dots, M$ .

**[0650]** Relations specified in (3.14-5) are called primary constraints. For simplicity let's assume that  $\text{rank}(L_{\dot{q}\dot{q}})$  is constant throughout the phase space,  $(q, \dot{q})$ , so that (3.14-5) defines a submanifold smoothly embedded in the phase space. This manifold is known as the primary constraint surface. Let

$$\text{rank}(L_{\dot{q}\dot{q}}) = N - M' \quad (3.14-6)$$

Then there are  $M'$  independent constraints among (3.14-5) and the primary constraint surface is a phase space submanifold of dimension  $2N - M'$ .

**[0651]** We do not assume that all the constraints are linearly independent so that

$$M' \leq M. \quad (3.14-7)$$

It follows from (3.14-5) that the inverse transformation from the  $p$ 's to the  $q$ 's is multivalued. That is, given  $q, p$  that satisfies (3.14-5), the inverse image  $(q, \dot{q})$  that satisfies

$$p = \left( \frac{\partial L}{\partial \dot{q}} \right)^T \quad (3.14-8)$$

is not unique, since (3.14-8) defines a map from a  $2N$ -dimensional manifold  $(q, \dot{q})$  to the smaller  $(2N - M')$ -dimensional manifold. Thus the inverse image of the points of (3.14-5) form a manifold of dimension  $M'$ .

**[0652]** Conditions on the Constraint Function

**[0653]** There exist many equivalent ways to represent a given surface by means of equations of the form of (3.14-5). For example the surface  $p_1 = 0$  can be represented equivalently by  $p_1^2 = 0$ ,  $\sqrt{|p_1|} = 0$ , or redundantly by  $p_1 = 0$  and  $p_1^2 = 0$ . To use the Hamiltonian formalism, it is necessary to impose some restrictions which the regularity conditions for the constraints.

**[0654]** Regularity Conditions

**[0655]** The  $(2N - M')$ -dimensional constraint surface  $\phi_m(q, p)$  should be covered of open region: in each region the constraints can be split into independent constraints

$$\{\Phi_m, m=1, \dots, M'\}.$$

Their Jacobian matrix

$$\left\{ \frac{\partial \phi_{m'}}{\partial p_n, q^{(n)}} \right\} = \begin{bmatrix} \frac{\partial \phi_1}{\partial p_1, q^{(1)}} & \cdots & \frac{\partial \phi_1}{\partial p_n, q^{(n)}} \\ \vdots & \ddots & \vdots \\ \frac{\partial \phi_{M'}}{\partial p_1, q^{(1)}} & \cdots & \frac{\partial \phi_{M'}}{\partial p_n, q^{(n)}} \end{bmatrix}$$

with  $m'=1, \dots, M'$  and  $n=1, \dots, N$ , is of rank  $M'$ .

**[0656]** The dependent constraints  $\phi_m$ ,  $m=M'+1, \dots, M$  of the other  $\phi_{m'}, =0 \Rightarrow \phi_m = 0$ . Alternatively the condition on the Jacobian.

**[0657]** 1. The function  $\phi_{m'}$  can be taken locally as the first  $M'$  coordinates of a new regular system in the vicinity of the constraint surface or the differentials  $d\phi_1, \dots, d\phi_{M'}$  are locally linearly independent:

$$d\phi_1 \wedge \dots \wedge d\phi_{M'} \neq 0 \quad (3.14-9)$$

**[0658]** 2. The variations  $\delta\phi_{m'}$  are of order  $\epsilon$  for arbitrary variations  $\delta q^{(i)}, \delta p_i$  of order  $\epsilon$  (Dirac's approach).

Theorem 3.14.1. If a smooth, phase space function  $G$  vanishes on  $\{\phi_m=0\}$  then

$$G = \sum_{m=1}^M G^{(m)} \phi_m \quad (3.14-10)$$

Proof:

**[0659]** (local proof). Set  $\phi_{m'}$ ,  $m'=1, \dots, M'$  as coordinates  $(y_{m'}, x_\alpha)$  with  $y_{m'} = \phi_{m'}$ . In these coordinates  $G(0, x) = 0$  and

$$\begin{aligned} G(y, x) &= \int_0^1 \frac{d}{dt} G(ty, x) dt \\ &= \sum_{m'=1}^{M'} y_{m'} \int_0^1 \frac{\partial}{\partial y_{m'}} G(ty, x) dt \\ &= \sum_{m'=1}^{M'} g^{(m')}(y, x) \phi_{m'}(y, x) \end{aligned} \quad (3.14-11)$$

with

$$g^{(m')}(y, x) = \int_0^1 \frac{\partial}{\partial y_{m'}} G(ty, x) dt.$$

Theorem 3.14.2.

**[0660]** If the sum  $\sum (\lambda^{(n)} \delta q^{(n)} + \mu_n \delta p_n) = 0$  for arbitrary variations  $\delta q^{(i)}, \delta p_i$  tangent to the constraint surface  $\{\phi_m(q, p) = 0 | m=1, \dots, M\}$ , then

$$\lambda^{(n)} = \sum_{m=1}^M u^{(m)} \frac{\partial \phi_m}{\partial q^{(n)}} \quad (3.14-12)$$

-continued

$$\mu_n = \sum_{m=1}^M u^{(m)} \frac{\partial \phi_m}{\partial p_n} \quad (3.14-13)$$

Proof.

**[0661]** The dimension of  $\{\phi_m\}$  is  $2N - M'$ . Thus the variations at a point  $(p, q)$  forms a  $2N - M'$  dimensional space

$$\sum_{n=1}^N (\lambda^{(n)} \delta q^{(n)} + \mu_n \delta p_n) = 0 \quad (3.14-14)$$

By the singularity assumption, there exists exactly  $M'$  solutions to (3.14-14). Clearly, the gradients

$$\left\{ \frac{\partial \Phi_{m'}}{\partial q^{(n)}} \right\} \text{ and } \left\{ \frac{\partial \Phi_{m'}}{\partial p_n} \right\}$$

are linearly independent. They are the basis for solutions to (3.14-14).

**[0662]** Note that in the presence of redundant constraints, the functions  $u^{(m)}$  exist but are not unique.

**[0663]** Canonical Hamiltonian

**[0664]** The Hamiltonian in canonical coordinates is

$$H(q, p) = \sum_{n=1}^N \dot{q}^{(n)} p_n - L(q, \dot{q}) \quad (3.14-15)$$

The rate  $\dot{q}$  enters through the combination through conjugate momenta defined for each coordinate

$$p_n(q, \dot{q}) = L_{\dot{q}^{(n)}}(q, \dot{q}) \quad (3.14-16)$$

This remarkable property is essential for the Hamiltonian approach. It is verified by evaluating the change  $\delta H$  involved by arbitrary independent variations of position and velocities.

$$\begin{aligned} \delta H &= \sum_{n=1}^N (\dot{q}^{(n)} \delta p_n + \delta \dot{q}^{(n)} p_n) - \delta L \\ &= \sum_{n=1}^N (\dot{q}^{(n)} \delta p_n + \delta \dot{q}^{(n)} p_n) - \sum_{n=1}^N (L_{q^{(n)}} \delta q^{(n)} + L_{\dot{q}^{(n)}} \delta \dot{q}^{(n)}) \end{aligned} \quad (3.14-17)$$

Utilizing (3.14-16) in (3.14-17) yields

$$\delta H = \sum_{n=1}^N (\dot{q}^{(n)} \delta p_n - L_{q^{(n)}} \delta q^{(n)}) \quad (3.14-18)$$

The Hamiltonian defined by (3.14-15) is not unique as a function of  $p, q$ . This can be inferred from (3.14-18) by noticing that  $\{\delta p_n | n=1, \dots, N\}$  are not all independent. They



are restricted to preserve the primary constraints  $\phi_m \approx 0$  which are identities when the  $p$ 's are expressed as functions of  $q$ 's via (3.14-16).

[0665] Using the definition of the differential in several variables applied to  $\delta H = \delta H(\{q^{(n)}\}, \{p_n\})$ , (3.14-18) can be rewritten

$$\sum_{n=1}^N \left( \frac{\partial H}{\partial q^{(n)}} \delta q^{(n)} + \frac{\partial H}{\partial p_n} \delta p_n \right) = \sum_{n=1}^N \left( \dot{q}^{(n)} \delta p_n - \delta q^{(n)} \frac{\partial L}{\partial q^{(n)}} \right) \text{ or} \quad (3.14-19)$$

$$\sum_{n=1}^N \left( \frac{\partial H}{\partial q^{(n)}} + \frac{\partial H}{\partial q^{(n)}} \right) \delta q^{(n)} + \sum_{n=1}^N \left( \frac{\partial H}{\partial p_n} - \dot{q}^{(n)} \right) \delta p_n = 0$$

From theorem 2 we then conclude for each  $n$  that.

$$\frac{\partial H}{\partial q^{(n)}} + \frac{\partial L}{\partial q^{(n)}} = \sum_{m=1}^M u^{(m)} \frac{\partial \Phi_m}{\partial q^{(n)}} \text{ and} \quad (3.14-20)$$

$$\frac{\partial H}{\partial p_n} - \dot{q}^{(n)} = \sum_{m=1}^M u^{(m)} \frac{\partial \Phi_m}{\partial p_n}.$$

So for each  $n$ :

$$\dot{q}^{(n)} = \frac{\partial H}{\partial p_n} + \sum_{m=1}^M u^{(m)} \frac{\partial \Phi_m}{\partial p_n}, \quad (3.14-21)$$

$$n = 1, \dots, N$$

and

$$-\frac{\partial L}{\partial q^{(n)}} = \frac{\partial H}{\partial q^{(n)}} + \sum_{m=1}^M u^{(m)} \frac{\partial \Phi_m}{\partial q^{(n)}}, \quad (3.14-22)$$

$$n = 1, \dots, N.$$

Note that if the constraints are independent, the vectors

$$\sum_{n=1}^N \frac{\partial \Phi_m}{\partial p_n},$$

$m=1, \dots, M$  are also independent because of the regularity conditions (this is proved later). Hence no two sets of  $\{u^{(m)}\}_{m=1, \dots, M}$  can yield the same velocities via (3.14-21).

[0666] Thus, using

$$\dot{q}^{(n)} = \frac{\partial H}{\partial p_n} + \sum_{m=1}^M u^{(m)}(q, \dot{q}) \frac{\partial \Phi_m}{\partial p_n}(q, p(q, \dot{q}))$$

we can find  $u^{(m)}(p, \dot{q})$ . If we define the transformation from  $(q, \dot{q})$  to the manifold  $\{\phi_m(q, p)=0 \mid m=1, \dots, M\}$ , from  $q, \dot{q}, u \rightarrow q, p, u$  by

$$q = q, n=1, \dots, N$$

$$p_n = L_{q^{(n)}}(q, \dot{q}), n=1, \dots, N-M'$$

$$u^{(m)} = u^{(m)}(q, \dot{q}), m=1, \dots, M'$$

We see that this transformation is invertible since one has from  $q, p, u \rightarrow q, \dot{q}, u$

$$q = q$$

$$\dot{q}^{(n)} = \frac{\partial H}{\partial p_n} + \sum_{m=1}^M u^{(m)} \frac{\partial \Phi_m}{\partial p_n}$$

$$\phi_m(q, p) = 0$$

Thus invertibility of the Legendre transformation when

$$\det(L_{\dot{q}\dot{q}}) \neq 0$$

can be regained at the prices of adding extra variables.

[0667] Action Principle of the Hamiltonian Form

[0668] With (3.14-21) and (3.14-22) we can rewrite (3.14-1) in the equivalent Hamiltonian form

$$\dot{q}^{(n)} = \frac{\partial H}{\partial p_n} + \sum_{m=1}^M u^{(m)} \frac{\partial \Phi_m}{\partial p_n} \quad (3.14-23)$$

$$\dot{p}^{(n)} = -\frac{\partial H}{\partial p_n} - \sum_{m=1}^M u^{(m)} \frac{\partial \Phi_m}{\partial p_n}$$

$$\Phi_m(q, p) = 0,$$

$$m = 1, \dots, M'$$

The Hamiltonian Equations (3.14-23) can be derived from the following variational principle:

$$\delta \int_{t_1}^{t_2} \left[ \sum_{n=1}^N \dot{q}^{(n)} p_n - H - \sum_{m=1}^M u^{(m)} \phi_m \right] = 0$$

for arbitrary variations of  $\delta q^{(n)}$ ,  $\delta p_n$ , and  $\delta u^{(m)}$  subject to

$$\delta q(t_1) = \delta q(t_2) = 0$$

where the  $u^{(m)}$  appear now as Lagrange multipliers enforcing the primary constraints

$$\Phi_m(q, p) = 0, m=1, \dots, M.$$

Let  $F(p, q)$  be an arbitrary function of the canonical variables, then

$$\frac{dF}{dt} = \sum_{n=1}^N \frac{\partial F}{\partial q^{(n)}} \dot{q}_n + \sum_{n=1}^N \frac{\partial F}{\partial p_n} \dot{p}_n \quad (3.14-25)$$

$$= \sum_{n=1}^N \frac{\partial F}{\partial q^{(n)}} \left[ \frac{\partial H}{\partial p_n} + \sum_{m=1}^M u^{(m)} \frac{\partial \Phi_m}{\partial p_n} \right] +$$

$$\sum_{n=1}^N \frac{\partial F}{\partial p_n} \left[ -\frac{\partial H}{\partial p_n} - \sum_{m=1}^M u^{(m)} \frac{\partial \Phi_m}{\partial p_n} \right]$$

$$= [F, H] + \sum_{m=1}^M u^{(m)} [F, \phi_m]$$

The equation (3.14-25) introduces the new binary operator  $[\bullet, \bullet]$  which is the Poisson bracket and has the form

$$[F, G] = \sum_{n=1}^N \left[ \frac{\partial F}{\partial q^{(n)}} \frac{\partial G}{\partial p_n} + \frac{\partial F}{\partial p_n} \frac{\partial G}{\partial q^{(n)}} \right] \quad (3.14-26)$$

$$= \sum_{n=1}^N [F_{q^{(n)}} G_{p_n} + F_{p_n} G_{q^{(n)}}]$$

**[0669]** Secondary Constraints

**[0670]** The basic consistency condition is that the primary constraints be preserved in time. So for

$$F(p, q) = \Phi_m(q, p)$$

we should have that  $\dot{\Phi}_m = 0$ .  $\{\Phi_m(q, p) = 0\}$ . So this means

$$[\Phi_m, H] + \sum_{m'=1}^{M'} u^{(m')} [\Phi_m, \Phi_{m'}] = 0 \quad (3.14-27)$$

This equation can either reduce to a relation independent of the  $u^{(m')}$ , or, it may impose a restriction on the  $u$ 's.

$$u = -\{[\Phi_m, \Phi_m]\} [\Phi_m, H](q, p) \quad (3.14-28)$$

In the case (3.14-27) is independent of the  $u$ 's (3.14-27) is called a secondary constraint. The fundamental difference of secondary constraints with respect to primary constraints is that primary constraints are the consequence of the definition (3.14-8) while secondary constraints depend on the dynamics.

**[0671]** If  $X(q, p) = 0$  is an external constraint, we most impose a compatibility condition

$$[X, H] + \sum_{m=1}^{M'} u^{(m)} [X, \Phi_m] = 0 \quad (3.14-29)$$

Next we need to test whether this constraint:

$$\Phi(p, q) = [X, H] + \sum_{m=1}^{M'} u^{(m)} [X, \Phi_m] = 0 \quad (3.14-30)$$

$$(3.14-31)$$

Implies new secondary constraints or whether it only restricts the  $u$ 's. After the process is finished we are left with a number of secondary constraints which will be denoted by

$$\Phi_k = 0, k = M+1, \dots, M+K$$

where  $K$  is the total number of secondary constraints. In general, it will be useful to denote all the constraints (primary and secondary) in a uniform way as

$$\Phi_j(q, p) = 0, j = 1, \dots, M+K = J \quad (3.14-32)$$

We make the same regularity assumptions on the full set of constraints.

**[0672]** Weak and Strong Equations

**[0673]** Equation (3.14-32) can be written as

$$\Phi_j(\bullet) \approx 0 \quad (3.14-33)$$

**[0674]** To emphasize, the quantity  $\Phi_j$  is numerically restricted to be zero but does not vanish throughout the space. What this means is that  $\Phi_j$  has non-zero Poisson brackets with the canonical variables.

**[0675]** Let  $F, G$  be functions that coincide on the manifold  $\{\Phi_j = 0 | j = 1, \dots, K\}$  are said to be weakly equal and denoted by  $F \approx G$ . On the other hand, an equation that holds throughout the entire phase space and not just on the submanifold  $\{\Phi_j = 0\}$  is called strong. Hence, by theorem 1

$$F \approx G \Leftrightarrow F - G = \sum_{j=1}^J c^{(j)}(p, q) \Phi_j \quad (3.14-34)$$

**[0676]** Restrictions on the Lagrange Multipliers

**[0677]** Assume that we have found a complete set of constraints

$$\{\Phi_j = 0 | j = 1, \dots, J\} \quad (3.14-35)$$

$$[\Phi_j, H] + \sum_{m=1}^M u^{(m)} [\Phi_j, \Phi_m] \approx 0 \quad (3.14-36)$$

We consider (3.14-36) as a set of non-homogeneous linear equations with  $M \leq J$  unknowns with coefficients that are functions of the  $q$ 's and  $p$ 's.

**[0678]** The general solution of (3.14-36) for each  $j$  is of the form

$$u^{(m)} = U^{(m)} + V^{(m)}, m = 1, \dots, M \quad (3.14-37)$$

with  $V^{(m)}$  the solution of the homogeneous equation

$$\sum_{m=1}^M V^{(m)} [\Phi_j, \Phi_m] \approx 0 \quad (3.14-38)$$

The most general solution of (3.14-38) is a linear combination of linearly independent solutions of  $V_\alpha^{(m)}$  where  $\alpha = 1, \dots, A$  with  $A \leq M$ . Under the assumption that the matrix

$$\begin{bmatrix} [\Phi_1, \Phi_1] & \dots & [\Phi_1, \Phi_M] \\ \vdots & \ddots & \vdots \\ [\Phi_J, \Phi_1] & \dots & [\Phi_J, \Phi_M] \end{bmatrix} \quad (3.14-39)$$

is of constant rank, the number of independent solutions  $A$  is the same for all  $p, q$ . Thus the general solution to (3.14-36) can be written as

$$u^{(m)} \approx U^{(m)} + \sum_{\alpha=1}^A V_\alpha^{(m)}, m = 1, \dots, M \quad (3.14-40)$$

**[0679]** Irreducible and Reducible Cases

**[0680]** If the equations  $\{\Phi_j = 0 | j = 1, \dots, J\}$  are not independent, one says that the constraints are reducible. The system is irreducible when the constraints are independent. However the separation of constraints into dependent and



independent ones might be difficult to perform. It also may disturb invariance properties under some important symmetry. In some cases it may be impossible to separate irreducible from irreducible contexts. Reducible cases arise for example when the dynamical coordinates include p-form gauge fields.

**[0681]** Any irreducible set of constraints can always be replaced by a reducible set by introducing constraints of the ones already at hand. The formalism should be invariant under such replacements.

**[0682]** Total Hamiltonian

**[0683]** We now discuss details of the dynamic equation (3.14-25)

$$\dot{F} \approx \left[ F, H' + \sum_{\alpha=1}^A V^{(\alpha)} \phi_{\alpha} \right] \quad (3.14-41)$$

where from (3.14-40)

$$H' = H + \sum_{m=1}^M U^{(m)} \phi_m \quad (3.14-42)$$

and

$$\phi_{\alpha} = \sum_{m=1}^M V_{\alpha}^{(m)} \phi_m, \alpha = 1, \dots, A$$

This is the result of theorem 3 (see below).

**[0684]** Theorem 3.

$$\left[ F, \sum_{m=1}^M U^{(m)} \phi_m \right] \approx \sum_{m=1}^M U^{(m)} [F, \phi_m] \quad (3.14-43)$$

$$\left[ F, \sum_{\alpha=1}^A V_{\alpha}^{(m)} \phi_m \right] \approx \sum_{\alpha=1}^A V_{\alpha}^{(m)} [F, \phi_m] \quad (3.14-44)$$

Proof.

**[0685]**

$$\left[ F, \sum_{m=1}^M U^{(m)} \phi_m \right] = \quad (3.14-45)$$

$$\sum_{i=1}^N \left\{ \frac{\partial F}{\partial q^{(i)}} \frac{\partial}{\partial p_i} \sum_{m=1}^M U^{(m)} \phi_m - \frac{\partial F}{\partial p_i} \frac{\partial}{\partial q^{(i)}} \sum_{m=1}^M U^{(m)} \phi_m \right\} =$$

$$\sum_{i=1}^N \left\{ \frac{\partial F}{\partial q^{(i)}} \left[ \sum_{m=1}^M \frac{\partial U^{(m)}}{\partial p_i} \phi_m + \sum_{m=1}^M U^{(m)} \frac{\partial \phi_m}{\partial p_i} \right] \right\} -$$

$$\sum_{i=1}^N \left\{ \frac{\partial F}{\partial p_i} \left[ \sum_{m=1}^M \frac{\partial U^{(m)}}{\partial q^{(i)}} \phi_m + \sum_{m=1}^M U^{(m)} \frac{\partial \phi_m}{\partial q^{(i)}} \right] \right\} =$$

-continued

$$\sum_{m=1}^M \{ [F, U^{(m)}] \phi_m + U^{(m)} [F, \phi_m] \}$$

So

$$\left[ F, \sum_{m=1}^M U^{(m)} \phi_m \right] - \sum_{m=1}^M U^{(m)} [F, \phi_m] = \sum_{m=1}^M \{ [F, U^{(m)}] \phi_m \}$$

and from (3.14-34) in (3.14-45), (3.14-43) follows. By a similar process we show (3.14-44). We now prove the validity of (3.14-41).

**[0686]** Theorem 4.

**[0687]** Let  $F(q, p)$  be a regular function, then  $F(p, q)$  propagates in time according to the approximate equation (3.14-41).

Proof.

From (3.14-25),

**[0688]**

$$\frac{dF}{dt} = [F, H] + \sum_{m=1}^M U^{(m)} [F, \phi_m]. \quad (3.14-46)$$

From (3.14-40) into (3.14-46) we obtain,

$$\frac{dF}{dt} \approx [F, H] + \sum_{m=1}^M \left\{ U^{(m)} + \sum_{\alpha=1}^A V^{(\alpha)} V_{\alpha}^{(m)} \right\} [F, \phi_m] \quad (3.14-47)$$

or

$$\frac{dF}{dt} \approx [F, H] + \sum_{m=1}^M U^{(m)} [F, \phi_m] + \sum_{m=1}^M \sum_{\alpha=1}^A V^{(\alpha)} V_{\alpha}^{(m)} [F, \phi_m]$$

Thus from (3.14-43) and (3.14-44) of theorem 3, we get

$$\frac{dF}{dt} \approx [F, H] + \sum_{m=1}^M [F, U^{(m)} \phi_m] + \quad (3.14-48)$$

$$\sum_{\alpha=1}^A V^{(\alpha)} \left[ F, \sum_{m=1}^M V_{\alpha}^{(m)} \phi_m \right] \approx$$

$$\left[ F, H + \sum_{m=1}^M U^{(m)} \phi_m + \sum_{\alpha=1}^A V^{(\alpha)} \sum_{m=1}^M V_{\alpha}^{(m)} \phi_m \right] \approx$$

$$\left[ F, H' + \sum_{m=1}^M U^{(m)} \phi_m + \sum_{\alpha=1}^A V^{(\alpha)} \phi_{\alpha} \right]$$

$$H' = H + \sum_{m=1}^M U^{(m)} \phi_m \quad (3.14-49)$$

$$\phi_{\alpha} = \sum_{m=1}^M V_{\alpha}^{(m)} \phi_m \quad (3.14-50)$$

Now define

$$H_T = H' + \sum_{\alpha=1}^A v^{(\alpha)} \phi_{\alpha}. \quad (3.14-51)$$

So we obtain

$$\frac{dF}{dt} \approx [F, H_T] \quad (3.14-52)$$

**[0689]** First and Second Class Functions

**[0690]** The distinction between primary and secondary constraints is of little importance. We now consider a fundamental classification. It depends on the concept of first class and second class functions.

**[0691]** Definition 1.

**[0692]** A function  $F(q, p)$  is said to be first class if its Poisson bracket with every constraint vanishes weakly,  $[F, \phi_j] \approx 0$ ,  $j=1, \dots, J$ . A function of the canonical variables that is not first class is called second class. Thus  $F$  is second class if  $[F, \phi_k] \neq 0$  for at least one  $k$ ,  $k=1, \dots, M$ .

**[0693]** Theorem 5.

**[0694]** If  $F$  and  $G$  are first class functions, then their Poisson bracket is also a first class function.

**[0695]** Proof:

**[0696]** By Hypothesis,

$$[F, \phi_j] = \sum_{k=1}^M f_j^{(k)} \phi_k \quad (3.14-53)$$

$$[G, \phi_j] = \sum_{l=1}^M g_j^{(l)} \phi_l \quad (3.14-54)$$

Applying the Jacobi identity, we get

$$\begin{aligned} [[F, G], \phi_j] &= [F, [G, \phi_j]] - [G, [F, \phi_j]] \\ &= \left[ F, \sum_{l=1}^M g_j^{(l)} \phi_l \right] - \left[ G, \sum_{k=1}^M f_j^{(k)} \phi_k \right] \\ &= \sum_i \left\{ \frac{\partial F}{\partial q^{(i)}} \frac{\partial}{\partial p_i} \sum_{l=1}^M g_j^{(l)} \phi_l - \frac{\partial F}{\partial p_i} \frac{\partial}{\partial q^{(i)}} \sum_{l=1}^M g_j^{(l)} \phi_l \right\} - \\ &\quad \sum_n \left\{ \frac{\partial G}{\partial q^{(n)}} \frac{\partial}{\partial p_n} \sum_{k=1}^M f_j^{(k)} \phi_k - \frac{\partial G}{\partial p_n} \frac{\partial}{\partial q^{(n)}} \sum_{k=1}^M f_j^{(k)} \phi_k \right\} \\ &\quad \sum_i \left\{ \frac{\partial F}{\partial q^{(i)}} \sum_{l=1}^M \left\{ \frac{\partial g_j^{(l)}}{\partial p_i} \phi_l + g_j^{(l)} \frac{\partial \phi_l}{\partial p_i} \right\} - \right. \\ &\quad \left. \frac{\partial F}{\partial p_i} \sum_{l=1}^M \left\{ \frac{\partial g_j^{(l)}}{\partial q^{(i)}} \phi_l + g_j^{(l)} \frac{\partial \phi_l}{\partial q^{(i)}} \right\} \right\} - \\ &= \sum_n \left\{ \frac{\partial G}{\partial q^{(n)}} \sum_{k=1}^M \left\{ \frac{\partial f_j^{(k)}}{\partial p_n} \phi_k + f_j^{(k)} \frac{\partial \phi_k}{\partial p_n} \right\} - \right. \end{aligned}$$

-continued

$$\begin{aligned} &\left. \frac{\partial G}{\partial p_n} \sum_{k=1}^M \left\{ \frac{\partial f_j^{(k)}}{\partial q^{(n)}} \phi_k + f_j^{(k)} \frac{\partial \phi_k}{\partial q^{(n)}} \right\} \right\} \\ &\quad \sum_{l=1}^M \left\{ \phi_l \sum_i \left\{ \frac{\partial F}{\partial q^{(i)}} \frac{\partial g_j^{(l)}}{\partial p_i} - \frac{\partial F}{\partial p_i} \frac{\partial g_j^{(l)}}{\partial q^{(i)}} \right\} + \right. \\ &\quad \left. g_j^{(l)} \sum_i \left\{ \frac{\partial F}{\partial q^{(i)}} \frac{\partial \phi_l}{\partial q^{(i)}} - \frac{\partial F}{\partial p_i} \frac{\partial \phi_l}{\partial p_i} \right\} \right\} - \\ &= \sum_{k=1}^M \left\{ \phi_k \sum_n \left\{ \frac{\partial G}{\partial q^{(n)}} \frac{\partial f_j^{(k)}}{\partial p_n} - \frac{\partial G}{\partial p_n} \frac{\partial f_j^{(k)}}{\partial q^{(n)}} \right\} + \right. \\ &\quad \left. + f_j^{(k)} \sum_n \left\{ \frac{\partial G}{\partial q^{(n)}} \frac{\partial \phi_k}{\partial p_n} - \frac{\partial G}{\partial p_n} \frac{\partial \phi_k}{\partial q^{(n)}} \right\} \right\} \\ &= \sum_{l=1}^M \{ \phi_l [F, g_j^{(l)}] + g_j^{(l)} [F, \phi_l] \} - \\ &= \sum_{k=1}^M \{ \phi_k [G, f_j^{(k)}] + f_j^{(k)} [G, \phi_k] \} \\ &\quad \sum_{l=1}^M [F, g_j^{(l)}] \phi_l - \sum_{k=1}^M [G, f_j^{(k)}] \phi_k + \\ &= \sum_{l'=1}^M \left\{ \sum_{l=1}^M g_j^{(l)} f_l^{(l')} \right\} \phi_{l'} - \sum_{k'=1}^M \left\{ \sum_{k=1}^M f_j^{(k)} g_k^{(k')} \right\} \phi_{k'} \approx 0 \end{aligned}$$

We now use theorem 5 to show the following.

**[0697]** Theorem 6.

**[0698]**  $H'$  defined by (3.14-49) and  $\phi_{\alpha}$  defined by (3.14-50) are first class functions.

**[0699]** Proof:

**[0700]** This follows directly from (3.14-36) and (3.14-38).

**[0701]** We learn from theorem 6 that the total Hamiltonian defined by (3.14-51) is the sum of the first class Hamiltonian  $H'$  and the first class primary constraints  $\phi_{\alpha}$  multiplied by arbitrary coefficients.

**[0702]** First Class Constraints as Generators of Gauge Transformations

**[0703]** Gauge transformations are transformations that do not change the physical state.

**[0704]** The presence of arbitrary functions of time  $v^{(\alpha)}$ ,  $\alpha=1, \dots, A$  in the total Hamiltonian,  $H_T$  (see (3.14-51)) imply that not all the  $q$ 's and  $p$ 's are observable given a set of  $q$ 's and  $p$ 's where the state of the physical system is uniquely determined. However the converse is not true: there is more than one set of values of the canonical variables that defines a state. To illustrate this, we see that if we give an initial set of values of physical state at time  $t$ , we expect the equations of motion to fully determine the state at other times. Thus any ambiguity in the value of the canonical variables at  $t_2 \neq t_1$  should be irrelevant from the physical point of view.

#### A Derivation Example

**[0705]** We propose here an alternate formulation of Dirac's formalism.



**[0706]** Primary Constraints

**[0707]** Recall that the momenta, canonically conjugate to the generalized “coordinates”  $q^{(j)}$ ,  $j=1, \dots, N$  is given by

$$p_j = \frac{\partial L(q, \dot{q})}{\partial \dot{q}^{(j)}}, \quad j = 1, \dots, N. \quad (\text{E-1})$$

For non-singular systems the equations allows us to express  $\dot{q}^{(j)}$ ,  $j=1, \dots, N$  in terms of the canonical variables,

$$\dot{q}^{(i)} = f_i(q, p), \quad i=1, \dots, N \quad (\text{E-2})$$

By performing a Legendre transformation

$$H_c(p, q) = \sum_{i=1}^N p_i f_i(q, p) + L(q, f(p, q))$$

We obtain the Hamiltonian of the system  $H_c$ . And from this function we obtain the standard equations of motion of the system.

$$\dot{q} = \frac{\partial H_c}{\partial p} \quad (\text{E-3})$$

$$\dot{p} = -\frac{\partial H_c}{\partial q}$$

For (E-2) to be well-defined we need to have the Hessian  $W$  of satisfy

$$\det W \neq 0 \quad (\text{E-4})$$

In this case the accelerations  $\ddot{q}^{(i)}$  are uniquely determined by the  $q^{(j)}$  and  $\dot{q}^{(i)}$ .

**[0708]** When  $\det W \neq 0$ , the Hamiltonian equations of motion do not take the standard form, and we speak of a singular Lagrangian. For illustration purposes, consider a Lagrangian of the form

$$L(q, \dot{q}) = \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N W_{ij}(q) \dot{q}^{(i)} \dot{q}^{(j)} + \sum_{i=1}^N \eta_i(q) \dot{q}^{(i)} - V(q) \quad (\text{E-5})$$

with  $W$  a symmetric matrix. From (E-1), the canonical momentum for (E-5) is given by

$$p_i = \frac{1}{2} \sum_{j=1}^N W_{ij}(q) \dot{q}^{(j)} + \eta_i(q), \quad i = 1, \dots, n. \quad (\text{E-6})$$

If  $W$  is singular of rank  $R_w$ , then it possesses  $N-R_w$  eigenvectors with corresponding zero eigenvalues. Then for eigenvectors  $v_j^{(\alpha)}$

$$\sum_{j=1}^N W_{ij}(q) v_j^{(\alpha)}(q) = 0, \quad \alpha = 1, \dots, N - R_w$$

So pre-multiplying (E-6) by  $v_j^{(\alpha)}$  and summing over  $i$  we get

$$\begin{aligned} \sum_{i=1}^N v_i^{(\alpha)}(q) p_i &= \sum_{i=1}^N \left[ \sum_{j=1}^N (v_i^{(\alpha)}(q) W_{ij}(q) \dot{q}^{(j)}) + v_i^{(\alpha)}(q) \eta_i(q) \right] \\ &= \sum_{i=1}^N v_i^{(\alpha)}(q) \eta_i(q), \quad \alpha = 1, \dots, N - R_w \end{aligned} \quad (\text{E-7})$$

So

$$\sum_{i=1}^N v_i^{(\alpha)}(q) (p_i - \eta_i(q)) = 0, \quad \alpha = 1, \dots, N - R_w.$$

Let  $\{p_\alpha\}$ ,  $\alpha=1, \dots, N-R_w$ , denote the linearly dependent elements of  $p$ . Let  $\{p_a\}$ ,  $a=1, \dots, R_a$  be the momenta satisfying (E-1). Then the constraint equations are of the form

$$\sum_{\beta=1}^{N-R_w} M_{\alpha\beta}(q) p_\beta - F_\alpha(q, \{p_a\}) = 0, \quad \alpha = 1, \dots, N - R_w \quad (\text{E-8})$$

and

$$M_{\alpha\beta}(q) = v_\beta^{(\alpha)}$$

$$F_\alpha(q, \{p_\beta\}) = \sum_{i=1}^N v_i^{(\alpha)}(q) \eta_i(q) + \sum_{b=1}^{R_w} v_b^{(\alpha)}(q) p_b \quad (\text{E-9})$$

The matrix  $\{M_{\alpha\beta}\}$  is necessarily invertible because otherwise  $M$  would possess eigenvectors with zero eigenvalues, implying existence of additional constraints. Note that (E-8) can be written as

$$p_\alpha - g_\alpha(q, \{p_a\}) = 0, \quad \alpha=1, \dots, N-R_w$$

with

$$g_\alpha(q, \{p_a\}) = \sum_{\beta=1}^{N-R_w} M_{\alpha\beta}^{-1} F_\beta(q, \{p_a\}) \quad (\text{E-10})$$

with  $\dim\{p_a\} = R_w$ . So we can define,

$$\Phi_\alpha(q, p) = p_\alpha - g_\alpha(q, \{p_a\}) = 0, \quad \alpha=1, \dots, N-R_w \quad (\text{E-11})$$

In Dirac's terminology, constraints of the form of (E-11) are referred to as primary constraints. Although the derivation above is based on a Lagrangian, quadratic in the velocity terms, it is generally valid for Lagrangians which depend on  $q$  and  $\dot{q}$  but not on higher derivatives.

**[0709]** Note: Primary constraints follow exclusively from the definition of canonical momenta.

**[0710]** The derivation above is valid for general Lagrangians and their Hessian. Let's assume  $\{W_{ij}(q, \dot{q})\}$  is the Hessian of a given Lagrangian  $L$ . Let  $\{W_{ab} | a, b=1, \dots, R_w\}$  be the largest sub-matrix of  $\{W_{ij}\}$  with suitable rearrangement if necessary. We then solve (E-1) for  $R_w$  velocities  $\dot{q}^{(\alpha)}$  in terms of  $\{q^{(i)} | i=1, \dots, n\}$ ,  $\{p_a | a=1, \dots, R_w\}$  and  $\{q^{(\alpha)} | \alpha=1, \dots, N-R_w\}$ . That is

$$\dot{q}^{(\alpha)} = f_\alpha(q, \{p_b\}, \{\dot{q}^{(\beta)}\}) \quad (\text{E-12})$$

with  $a, b=1, \dots, R_w$  and  $\beta=R_w+1, \dots, N$ .

[0711] Inserting these relations into (E-1), we get relations of the form

$$p_j = h_j(q, \{p_a\}, \{\dot{q}^{(\alpha)}\}) \quad (\text{E-13})$$

with  $a, j=1, \dots, R_w$  and  $\alpha=R_w+1, \dots, N$ . This relation reduces to an identity by construction. The remaining equations are of the form

$$p_\alpha = h_\alpha(q, \{p_a\}, \{\dot{q}^{(\beta)}\}) \quad (\text{E-14})$$

with  $\alpha=1, \dots, N-R_w$ . However, the right hand side cannot depend on  $\{q(O)\}$  since otherwise we could express more velocities in terms of the momenta of the coordinates of the momenta and the remaining velocities.

[0712] Hamiltonian Equations of Motion for Constrained Systems

[0713] Theorem 3.16.1.

[0714] In the space  $\Gamma_p$  define by  $\Gamma_p = \{\phi_\alpha(p, q) | \alpha=1, \dots, N-R_w\}$  where  $\phi_\alpha$  is defined as (E-11). The Hamiltonian is only a function of  $\{q^{(i)} | i=1, \dots, N\}$  and momenta  $\{p_a | a=1, \dots, R_w\}$  and does not depend on  $\{\dot{q}^{(\alpha)} | \alpha=1, \dots, N-R_w\}$

[0715] Proof.

[0716] On  $\Gamma_p$  the Hamiltonian is given by

$$H_o = H_c | \Gamma_p = \sum_{a=1}^{R_w} p_a f_a - \sum_{\alpha=1}^{N-R_w} f_\alpha \dot{q}^{(\alpha)} - L(q, \{f_b\}, \{\dot{q}^{(\beta)}\}) \quad (\text{E-15})$$

where  $f_a, a=1, \dots, N-R_w$  is given by (E-12) and  $g_\alpha, \alpha=1, \dots, R_w$  is given by (E-10). We want to show that  $H_o$  does not depend on  $\dot{q}^{(\beta)}, \beta=1, \dots, N-R_w$ . We compute

(E-16)

$$\begin{aligned} \frac{\partial H_o}{\partial \dot{q}^{(\beta)}} &= \sum_{a=1}^{R_w} p_a \frac{\partial f_a}{\partial \dot{q}^{(\beta)}} - g_\beta - \sum_{\alpha=1}^{R_w} \frac{\partial L}{\partial \dot{q}^{(\alpha)}} \bigg|_{\dot{q}^{(\alpha)}=f_\alpha} \frac{\partial f_a}{\partial \dot{q}^{(\beta)}} - \frac{\partial L}{\partial \dot{q}^{(\beta)}} \bigg|_{\dot{q}^{(\alpha)}=f_\alpha} \\ &= \sum_{a=1}^{R_w} \left( p_a - \frac{\partial L}{\partial \dot{q}^{(\alpha)}} \bigg|_{\dot{q}^{(\alpha)}=f_\alpha} \right) \frac{\partial f_a}{\partial \dot{q}^{(\beta)}} - g_\beta - \frac{\partial L}{\partial \dot{q}^{(\beta)}} \bigg|_{\dot{q}^{(\alpha)}=f_\alpha} \end{aligned}$$

Since by definition

$$p_a = \frac{\partial L}{\partial \dot{q}^{(a)}}, \quad a = 1, \dots, R_w$$

And from (E-11)

$$g_\beta = p_\beta = \frac{\partial L}{\partial \dot{q}^{(\beta)}} \bigg|_{\dot{q}^{(a)}=f_a} \quad (\text{E-17})$$

So

$$\frac{\partial H_o}{\partial \dot{q}^{(\beta)}} = 0, \quad \beta = 1, \dots, N-R_w.$$

and therefore

$$H_o(q, \{p_a\}, \{\dot{q}^{(\alpha)}\}) = H_o(q, \{p_a\}).$$

[0717] Theorem 3.16.2.

[0718] In the presence of primary constraints (E-11), the Hamilton equations of motion are given by

$$\dot{q}^{(i)} = \frac{\partial H_o}{\partial p_i} + \sum_{\beta=1}^N \dot{q}^{(\beta)} \frac{\partial \phi_\beta}{\partial p_i}, \quad i = 1, \dots, N \quad (\text{E-18})$$

$$\dot{p}_i = -\frac{\partial H_o}{\partial q^{(i)}} + \sum_{\beta=1}^n \dot{q}^{(\beta)} \frac{\partial \phi_\beta}{\partial q^{(i)}}, \quad i = 1, \dots, N$$

$$\phi_\alpha(p, q) = 0, \quad \alpha = 1, \dots, N-R_w$$

where  $\dot{q}^{(\beta)}$  are a priori underdetermined velocities.

[0719] Proof: From (E-15) we obtain and the application of Theorem 3.16.1

$$\begin{aligned} \frac{\partial H_o}{\partial p_a} &= f_a + \sum_{b=1}^{R_w} p_b \frac{\partial f_b}{\partial p_a} + \sum_{\beta=1}^{N-R_w} \frac{\partial g_\beta}{\partial p_a} \dot{q}^{(\beta)} - \sum_{b=1}^{R_w} \frac{\partial L}{\partial \dot{q}^{(b)}} \frac{\partial f_b}{\partial p_a} \quad (\text{E-19}) \\ &= \dot{q}^{(a)} + \sum_{b=1}^{R_w} \left( p_b - \frac{\partial L}{\partial \dot{q}^{(b)}} \right) \frac{\partial f_b}{\partial p_a} + \sum_{\beta=1}^{N-R_w} \frac{\partial g_\beta}{\partial p_a} \dot{q}^{(\beta)} \\ &= \dot{q}^{(a)} + \sum_{\beta=1}^{N-R_w} \frac{\partial g_\beta}{\partial p_a} \dot{q}^{(\beta)} \end{aligned}$$

with  $a=1, \dots, n-R_w$ . Further

$$\begin{aligned} \frac{\partial H_o}{\partial q^{(i)}} &= \sum_{b=1}^{R_w} p_b \frac{\partial f_b}{\partial q^{(i)}} + \sum_{\beta=1}^{N-R_w} \dot{q}^{(\beta)} \frac{\partial g_\beta}{\partial q^{(i)}} - \quad (\text{E-20}) \\ &= \frac{\partial L}{\partial q^{(i)}} \bigg|_{\dot{q}^{(a)}=f_a} - \sum_{b=1}^{R_w} \frac{\partial L}{\partial \dot{q}^{(b)}} \bigg|_{\dot{q}^{(b)}=f_b} \frac{\partial f_b}{\partial q^{(i)}} \\ &= \sum_{b=1}^{R_w} \left( p_b - \frac{\partial L}{\partial \dot{q}^{(b)}} \bigg|_{\dot{q}^{(b)}=f_b} \right) \frac{\partial f_b}{\partial q^{(i)}} + \\ &= \sum_{\beta=1}^{N-R_w} \dot{q}^{(\beta)} \frac{\partial g_\beta}{\partial q^{(i)}} - \frac{\partial L}{\partial q^{(i)}} \bigg|_{\dot{q}^{(a)}=f_a} \\ &= \sum_{\beta=1}^{N-R_w} \dot{q}^{(\beta)} \frac{\partial g_\beta}{\partial q^{(i)}} - \frac{\partial L}{\partial q^{(i)}} \bigg|_{\dot{q}^{(a)}=f_a} \\ &= \sum_{\beta=1}^{N-R_w} \dot{q}^{(\beta)} \frac{\partial g_\beta}{\partial q^{(i)}} - \frac{d}{dt} \left( \frac{\partial L}{\partial \dot{q}^{(i)}} \right) \bigg|_{\dot{q}^{(a)}=f_a} \end{aligned}$$

from (add reference).

$$\frac{\partial H_o}{\partial q^{(i)}} = -\dot{p}_i + \sum_{\beta=1}^{N-R_w} \dot{q}^{(\beta)} \frac{\partial g_\beta}{\partial q^{(i)}} \quad (\text{E-21})$$



From (E-19) and (E-20) we get:

[0720]

$$\begin{aligned} \dot{q}^{(a)} &= \frac{\partial H_o}{\partial p_a} - \sum_{\beta=1}^{N-R_w} \frac{\partial g_\beta}{\partial p_a} \dot{q}^{(\beta)}, \quad a = 1, \dots, R_w \\ \dot{p}_i &= -\frac{\partial H_o}{\partial q^{(i)}} + \sum_{\beta=1}^{N-R_w} \dot{q}^{(\beta)} \frac{\partial g_\beta}{\partial q^{(i)}}, \quad i = 1, \dots, N \end{aligned} \quad (\text{E-22})$$

Since

[0721]

$$\frac{\partial H_o}{\partial p_\alpha} = 0 \quad \text{and} \quad \frac{\partial \phi_\beta}{\partial p_\alpha} = \delta\beta_\alpha$$

we can supplement these equations with

$$\dot{q}^{(\alpha)} = \frac{\partial H_o}{\partial p_\alpha} - \sum_{\beta=1}^{N-R_w} \frac{\partial g_\beta}{\partial p_\alpha} \dot{q}^{(\beta)}, \quad \alpha = 1, \dots, N - R_w \quad (\text{E-23})$$

So we can write

$$\begin{aligned} \dot{q}^{(i)} &= \frac{\partial H_o}{\partial p_i} + \sum_{\beta=1}^{N-R_w} \frac{\partial g_\beta}{\partial p_i} \dot{q}^{(\beta)}, \quad i = 1, \dots, N \\ \dot{p}_i &= -\frac{\partial H_o}{\partial q^{(i)}} - \sum_{\beta=1}^{N-R_w} \dot{q}^{(\beta)} \frac{\partial g_\beta}{\partial q^{(i)}}, \quad i = 1, \dots, N \end{aligned} \quad (\text{E-24})$$

For consistency with (E-11) we should write

$$\dot{q}^{(\alpha)} = \frac{d}{dt} - g_\alpha(q, \{p_\alpha\}), \quad \alpha = 1, \dots, N - R_w \quad (\text{E-25})$$

where  $\dot{p}_\alpha$  is given by the right hand side of (E-22).

[0722] Streamlining the Hamiltonian Equation of Motion (EOM)

[0723] Definition 3.16-1.

[0724] A function  $f$  is weakly equal to  $g$  denoted by  $f \approx g$ , if  $f$  and  $g$  are equal on the subspace defined by the primary constraints,

$$\Phi_\beta = 0 \quad \text{when} \quad f|_{\Gamma_p} = g|_{\Gamma_p}$$

and

$$f(q,p) \approx g(q,p) \Leftrightarrow f(q,p) = g(q,p) \quad \text{when} \quad \{\Phi_\alpha(q,p) = 0\}$$

[0725] Theorem 3.16.3.

[0726] Assume  $f, g$  are defined over the entire space spanned by  $\{q^{(i)}\}, \{p_i\}$ . Then if

$$f(q,p)|_{\Gamma_p} = g(q,p)|_{\Gamma_p} \quad (\text{E-26})$$

Then

[0727]

$$\frac{\partial}{\partial q^{(i)}} \left( f - \sum_{\beta} \phi_\beta \frac{\partial f}{\partial p_\beta} \right) \approx \frac{\partial}{\partial q^{(i)}} \left( h - \sum_{\beta} \phi_\beta \frac{\partial h}{\partial p_\beta} \right) \quad (\text{E-27})$$

and

$$\frac{\partial}{\partial p_i} \left( f - \sum_{\beta} \phi_\beta \frac{\partial f}{\partial p_\beta} \right) \approx \frac{\partial}{\partial p_i} \left( h - \sum_{\beta} \phi_\beta \frac{\partial h}{\partial p_\beta} \right)$$

for  $i=1, \dots, N$ .

[0728] Proof: Consider the two functions  $f(q, \{p_\alpha\}, \{p_\beta\})$  and  $h(q, \{p_\alpha\}, \{p_\beta\})$ . Using (E-11) and from the hypothesis of the theorem,

$$f(q, \{p_\alpha\}, \{g_\alpha\}) = h(q, \{p_\alpha\}, \{g_\alpha\}) \quad (\text{E-28})$$

Thus it follows

$$\left( \frac{\partial f}{\partial q^{(i)}} + \sum_{\alpha} \frac{\partial f}{\partial p_\alpha} \frac{\partial p_\alpha}{\partial q^{(i)}} + \sum_{\beta} \frac{\partial f}{\partial p_\beta} \frac{\partial g_\beta}{\partial q^{(i)}} \right)_{\Gamma_p} = \quad (\text{E-29})$$

$$\left( \frac{\partial h}{\partial q^{(i)}} + \sum_{\alpha} \frac{\partial h}{\partial p_\alpha} \frac{\partial p_\alpha}{\partial q^{(i)}} + \sum_{\beta} \frac{\partial h}{\partial p_\beta} \frac{\partial g_\beta}{\partial q^{(i)}} \right)_{\Gamma_p}$$

and

$$\left( \frac{\partial f}{\partial p_i} + \sum_{\alpha \neq i} \frac{\partial f}{\partial p_\alpha} \frac{\partial p_\alpha}{\partial p_i} + \sum_{\beta} \frac{\partial f}{\partial p_\beta} \frac{\partial g_\beta}{\partial p_i} \right)_{\Gamma_p} = \quad (\text{E-30})$$

$$\left( \frac{\partial h}{\partial p_i} + \sum_{\alpha \neq i} \frac{\partial h}{\partial p_\alpha} \frac{\partial p_\alpha}{\partial p_i} + \sum_{\beta} \frac{\partial h}{\partial p_\beta} \frac{\partial g_\beta}{\partial p_i} \right)_{\Gamma_p}$$

Note since  $\Phi_\alpha(q,p) = p_\alpha - g_\alpha(q, \{p_\alpha\})$ , we have

$$\frac{\partial g_\beta}{\partial q^{(i)}} = -\frac{\partial \Phi_\beta(q,p)}{\partial q^{(i)}}$$

and

$$\frac{\partial g_\beta}{\partial p_i} = -\frac{\partial \Phi_\beta(q,p)}{\partial p_i}$$

and

$$\partial \Phi_\alpha(q,p) = 0$$

for  $\alpha=1, \dots, N-R_w$ . We have

$$\left( \frac{\partial f}{\partial q^{(i)}} - \sum_{\beta} \frac{\partial f}{\partial p_\beta} \frac{\partial \Phi_\beta}{\partial q^{(i)}} \right)_{\Gamma_p} = \left( \frac{\partial h}{\partial q^{(i)}} - \sum_{\beta} \frac{\partial h}{\partial p_\beta} \frac{\partial \Phi_\beta}{\partial q^{(i)}} \right)_{\Gamma_p}$$

which can be written as

$$\frac{\partial}{\partial q^{(i)}} \left( f - \sum_{\beta} \phi_{\beta} \frac{\partial f}{\partial p_{\beta}} \right) \approx \frac{\partial}{\partial q^{(i)}} \left( h - \sum_{\beta} \phi_{\beta} \frac{\partial h}{\partial p_{\beta}} \right)$$

since  $\phi_{\beta} \partial^2 f / \partial p_{\beta}^2 = 0$  because  $\phi_{\beta} = 0$ . Similarly,

$$\frac{\partial}{\partial p_i} \left( f - \sum_{\beta} \phi_{\beta} \frac{\partial f}{\partial p_{\beta}} \right) \approx \frac{\partial}{\partial q^{(i)}} \left( h - \sum_{\beta} \phi_{\beta} \frac{\partial h}{\partial p_{\beta}} \right)$$

[0729] Corrolary 3.16-1.

$$q^{(i)} = \frac{\partial H}{\partial p_i} + \sum_{\beta} v^{(\beta)} \frac{\partial \phi_{\beta}}{\partial p_i}$$

$$\dot{p}_i = -\frac{\partial H}{\partial q^{(i)}} - \sum_{\beta} v^{(\beta)} \frac{\partial \phi_{\beta}}{\partial q^{(i)}}$$

for  $i=1, \dots, N$ .

[0730] Proof.

[0731] We consider two Hamiltonians  $H(\{q^{(i)}\}, \{p_i\})$  and  $H_o(\{q^{(i)}\}, \{p_a\})$ . Define  $H(\{q^{(i)}\}, \{p_i\})$  as follows

$$H(\{q^{(i)}\}, \{p_i\}) \approx H_o(\{q^{(i)}\}, \{p_a\}).$$

Then using the result of Theorem 3.16.1, from (E-29) with  $f=H$  and  $h=H_o$

$$\frac{\partial H_o}{\partial q^{(i)}} \approx \frac{\partial}{\partial q^{(i)}} \left( H - \sum_{\beta=1}^{N-R_W} \phi_{\beta} \frac{\partial H}{\partial p_{\beta}} \right) \quad (E-31)$$

$$\frac{\partial H_o}{\partial p_i} \approx \frac{\partial}{\partial p_i} \left( H - \sum_{\beta=1}^{N-R_W} \phi_{\beta} \frac{\partial H}{\partial p_{\beta}} \right) \quad (E-32)$$

Using (E-31) and (E-32) in (E-24), we get

[0732]

$$\dot{q}^{(i)} = \frac{\partial}{\partial p_i} \left( H - \sum_{\beta} \phi_{\beta} \frac{\partial H}{\partial p_{\beta}} \right) + \sum_{\beta} \dot{q}^{(\beta)} \frac{\partial \phi_{\beta}}{\partial p_i} \quad (E-33)$$

and

$$\dot{p}_i \approx -\frac{\partial}{\partial q^{(i)}} \left( H - \sum_{\beta} \phi_{\beta} \frac{\partial H}{\partial p_{\beta}} \right) - \sum_{\beta} \dot{q}^{(\beta)} \frac{\partial \phi_{\beta}}{\partial q^{(i)}}$$

or

$$\dot{q}^{(i)} \approx \frac{\partial}{\partial p_i} \left( H - \sum_{\beta} \phi_{\beta} \left( \frac{\partial H}{\partial p_{\beta}} - \dot{q}^{(\beta)} \right) \right)$$

-continued

and

$$\dot{p}_i \approx -\frac{\partial}{\partial q^{(i)}} \left( H - \sum_{\beta} \phi_{\beta} \left( \frac{\partial H}{\partial p_{\beta}} - \dot{q}^{(\beta)} \right) \right)$$

Define

[0733]

$$v_{\beta} \equiv \dot{q}^{(\beta)} - \frac{\partial H}{\partial p_{\beta}}$$

$$H_T \equiv H + \sum_{\beta} v^{(\beta)} \phi_{\beta}$$

So (E-33) becomes

$$\dot{q}^{(i)} \approx \frac{\partial H_T}{\partial p_i} \quad (E-34)$$

$$\dot{p}_i \approx -\frac{\partial H_T}{\partial q^{(i)}}$$

[0734] Constrained Hamiltonian Systems

[0735] Local symmetries on a Lagrangian based model.

Consider

$$q^{(i)} \rightarrow q^{(i)}(t) + \delta q^{(i)}(t)$$

$$\dot{q}^{(i)} \rightarrow \dot{q}^{(i)}(t) + \delta \dot{q}^{(i)}(t)$$

with  $i=1, \dots, N$ . The action of the system is given by

$$S(q, \dot{q}) = \int L(q, \dot{q}) dt$$

where  $q$  and  $\dot{q}$  are  $n$ -dimensional column vectors. The action differential

$$\begin{aligned} \delta S &= \int L(q + \delta q, \dot{q} + \delta \dot{q}) dt - \int L(q, \dot{q}) dt \\ &= \int L(q + \delta q, \dot{q} + \delta \dot{q}) dt - \int L(q, \dot{q}) dt \\ &= \int \left[ \sum_i \frac{\partial L}{\partial q^{(i)}} \delta q^{(i)} + \sum_i \frac{\partial L}{\partial \dot{q}^{(i)}} \delta \dot{q}^{(i)} \right] dt \\ &= - \int \sum_i \left[ \frac{d}{dt} \frac{\partial L}{\partial \dot{q}^{(i)}} - \frac{\partial L}{\partial q^{(i)}} \right] \delta q^{(i)} dt \\ &= - \sum_i dt \sum_j E_i^{(o)}(q, \dot{q}, \ddot{q}) \delta q^{(i)} \end{aligned}$$

where we define the Euler-Lagrange differential operator

$$E_i^{(o)}(q, \dot{q}, \ddot{q}) = \frac{d}{dt} \frac{\partial L}{\partial \dot{q}^{(i)}} - \frac{\partial L}{\partial q^{(i)}}. \quad (3.17-1)$$

Note that

$$\int \sum_{i=1}^N E_i^{(o)}(q, \dot{q}, \ddot{q}) \delta^{(i)} dt \equiv 0$$



on shell. Expanding  $E_i^{(o)}$

$$\begin{aligned} E_i^{(o)}(q, \dot{q}, \ddot{q}) &= \sum_j \left[ \frac{\partial^2 L(q, \dot{q})}{\partial \dot{q}^{(i)} \partial \ddot{q}^{(j)}} \ddot{q}^{(j)} + \frac{\partial^2 L(q, \dot{q})}{\partial \dot{q}^{(i)} \partial q^{(j)}} \dot{q}^{(j)} \right] - \frac{\partial L(q, \dot{q})}{\partial q^{(i)}} \\ &= \sum_j W_{ij}(q, \dot{q}) \ddot{q}^{(j)} + \sum_j \frac{\partial^2 L(q, \dot{q})}{\partial \dot{q}^{(i)} \partial q^{(j)}} \dot{q}^{(j)} - \frac{\partial L(q, \dot{q})}{\partial q^{(i)}} \\ &= \sum_j W_{ij}(q, \dot{q}) \ddot{q}^{(j)} + k_i(q, \dot{q}) \end{aligned}$$

If  $L$  is singular,  $W_{(N \times N)}$  is not invertible so (3.17-1) cannot be solved for  $\ddot{q}_i$ ,  $i=1, \dots, N$ . If  $\text{Rank}(W(q, \dot{q}))=R_W$  on shell, then there exist  $N-R_W$  in the theory. There exist  $N-R_W$  independent left (or right) zero mode eigenvectors  $w_i^{(o,k)}$ ,  $i=1, \dots, N-R_W$  such that

$$\sum_i w_i^{(o,k)}(q, \dot{q}) W_{ij}(q, \dot{q}) = 0, \quad (3.17-2)$$

$$k = 1, \dots, N - R_W$$

Thus

**[0736]**

$$\phi^{(o,k)} = \sum_{i=1}^N w_i^{(o,k)}(q, \dot{q}) E_i^{(o)}(q, \dot{q}, \ddot{q})$$

depend on  $q$  and  $\dot{q}$  only. The  $\phi^{(o,k)}$  also vanish on shell:

$$\phi^{(o,k)}(q, \dot{q}) = 0, k=1, \dots, N-R_W$$

**[0737]** The set  $\{\phi^{(o,k)} | k=1, \dots, N-R_W\}$  are the zero generation constraints. It is possible that not all the  $\{\phi^{(o,k)}\}$  are linearly independent. So we may find linear combinations of the zero mode eigenvectors)

$$v_i^{(o,n_o)} = \sum_k c_k^{(n_o)} w_i^{(o,k)}$$

such that we have

$$G^{(o,n_o)} = v^{(o,n_o)} E^{(o)} = 0, n_o = \dots, N_o \quad (3.17-3)$$

These are called gauge identities.

**[0738]** Any variation  $\delta q_i$ ,  $i=1, \dots, N$ , of the form

$$\delta q_i = \sum_{n_o} \epsilon_{n_o} v_i^{(o,n_o)}$$

Is action invariant by (3.17-1).

Given this definition of  $\delta q_i$  and (3.17-3), we conclude

$$\begin{aligned} \delta S &= \int dt \sum_{i=1}^N E_i^{(o)}(q, \dot{q}, \ddot{q}) \sum_{n_o} \epsilon_{n_o}(t) v_i^{(o,n_o)} \\ &= \int dt \sum_{i=1}^N \epsilon_{n_o} \sum_{n_o} E_i^{(o)}(q, \dot{q}, \ddot{q}) v_i^{(o,n_o)}(q, \dot{q}) \\ &= \int dt \sum_{i=1}^N \epsilon_{n_o} G^{(o,n_o)} \\ &\equiv 0 \end{aligned}$$

everywhere. The remaining zero generating modes which we denote by  $u^{(o,n_o)}$  lead to genuine constraints. They are of the form  $\phi^{(o,n_o)}(q, \dot{q})=0$  on shell, where

$$\phi^{(o,n_o)} = u^{(o,n_o)} E^{(o)}. \quad (3.17-4)$$

**[0739]** The algorithm now proceeds as follows. We separate the gauge identities (3.17-3) from the nontrivial constraints (3.17-4) and will list them separately. They will be used for determining local symmetry transformations.

**[0740]** Next we want to search for additional constraints. We do this by searching for further functions of the coordinates and velocities which vanish in the space of physical trajectories. To this effect consider the following  $N+N_o$  vector constructed from  $E^{(o)}$  and the time derivative of the constraints (3.17-4)

$$[E^{(1)}] = \begin{bmatrix} E^{(o)} \\ \frac{d}{dt}(u^{(o,1)} E^{(o)}) \\ \vdots \\ \frac{d}{dt}(u^{(o,n_o)} E^{(o)}) \end{bmatrix} = \begin{bmatrix} E^{(o)} \\ \frac{d}{dt} \phi^{(o)} \end{bmatrix} \quad (3.17-5)$$

by construction. The constraint  $\phi^{(o)}$  is valid for all time and therefore

$$\frac{d}{dt} \phi^{(o)} = 0$$

on shell, but

$$\frac{d\phi^{(o,i)}}{dt} = \nabla \dot{q}(u^{(o,i)} E^{(o)}) \dot{q} + \nabla q(u^{(o,i)} E^{(o)}) \dot{q} \quad (3.17-6)$$

So

$$[E_{i_1}^{(1)}] = \sum_{j=1}^n W_{i_1 j}^{(1)}(q, \dot{q}) \dot{q}^{(j)} + k_{i_1}^{(1)}(q, \dot{q})$$

where  $i_1=1, \dots, N+N_o$ , and

$$[W_{i_1 j}^{(1)}] = \begin{bmatrix} W^{(o)} \\ \nabla \dot{q}(u^{(o,i)} E^{(o)}) \\ \vdots \\ \nabla \dot{q}(u^{(o,N_o)} E^{(o)}) \end{bmatrix} \quad (3.17-7)$$

-continued

$$[k_{i_1}^{(1)}] = \begin{bmatrix} k^{(o)} \\ \sum_j \frac{\partial}{\partial q^{(j)}} (u^{(o,i)} E^{(o)}) \dot{q}^{(j)} \\ \vdots \\ \sum_j \frac{\partial}{\partial q^{(j)}} (u^{(o,N_o)} E^{(o)}) \dot{q}^{(j)} \end{bmatrix}$$

We next look for the zero modes of  $W^{(1)}$ . By construction, these zero modes include the 0 modes of the previous level. The gauge identities at level 1 are.

$$G^{(1,n_1)} = v^{(1,n_1)} E^1 - \sum_{n_o=1}^{N_o} M_{n_1 n_o}^{(1,0)} (u^{(o,n_o)} E^{(o)}) \equiv 0 \quad (3.17-8)$$

where  $n_1=1, \dots, N_1$  and the genuine constraints are of the form

$$\phi^{(1,n_i)} = \phi^{(1,n_i)} E^1 = 0 \quad (3.17-9)$$

with  $n_1=1, \dots, N_1$  on shell.

**[0741]** We next adjoin the new identities (3.17-8) to the ones determined earlier (3.17-3) with the remaining constraints (3.17-9) we proceed as before, adjoining their time derivatives to (3.17-5) and construct  $W_{i_i}^{(1)}$  and  $k_{i_1}^{(1)}$ .

**[0742]** The iterative process will terminate at some level M if either i) there is not further zero modes, or ii) the new constraints can be expressed as linear combinations of previous constraints.

**[0743]** The Maximal Set of Linearly Independent Gauge Identities Generated by the Algorithm

**[0744]** Note that the algorithm steps are of the form)

$$G^{(o,n_o)} = u^{(o,n_o)} E^{(o)} \equiv 0 \quad (3.17-10)$$

$$G^{(l,n_l)} = u^{(l,n_l)} E^{(l)} - \sum_{l'=0}^{l-1} \sum_{n_{l'}=0}^{N_{l'}} M_{n_l n_{l'}}^{(l,l')} \phi^{(l',n_{l'})} \quad (3.17-11)$$

with  $L=1, \dots, N_1$ . The  $M_{n_l n_{l'}}^{(l,l')}$  are only functions of  $q$  and  $\dot{q}$ . And

$$\phi^{(l,n_l)} = u^{(l,n_l)} E^{(l)}, \quad n_l=1, \dots, N_l, \quad (3.17-12)$$

$$E^{(l)} = \begin{bmatrix} E^{(o)} \\ \frac{d\phi^{(o)}}{dt} \\ \vdots \\ \frac{d\phi^{(l-1)}}{dt} \end{bmatrix} \quad (3.17-13)$$

where  $\phi^{(l)}$  is a column vector with  $N_l$  components  $\phi^{(l,n_l)}$ . Thus we conclude from (3.17-13) and (3.17-11) that the general form of the gauge identity given by (3.17-11) is of the form

$$G^{(l,n_l)} = \sum_{i=1}^{N_l} \sum_{l=1}^M \sum_{m=1}^l \zeta_{mi}^{(l,m_l)} \frac{d^m}{dt^m} E_i^{(o)} \equiv 0 \quad (3.17-14)$$

where  $\zeta_{mi}^{(l,m_l)}(q, \dot{q})$  and  $N_l < M$ . From (3.17-14) it also follows that

$$\sum_{l=1}^M \sum_{n_l=1}^{N_l} \zeta^{(l,n_l)} G^{(l,n_l)} \equiv 0 \quad (3.17-15)$$

This identity can also be written as

$$\sum \delta q^{(i)} E_i^{(o)} - \frac{d}{dt} F \quad (3.17-16)$$

where

$$\delta q^{(i)} = \sum_{l=1}^M \sum_{n_l=1}^{N_l} \sum_{m=q}^l (-1)^m \frac{d^m}{dt^m} \zeta_{mi}^{(l,m_l)} \varepsilon^{(l,m_l)}(t)$$

and F is a complicated function of  $q$  and  $\dot{q}$ . By collecting indices  $l, n_l$  together

$$\delta q_i = \sum_{l=1}^M \sum_{n_l=1}^{N_l} \sum_{m=q}^l (-1)^m \zeta_{mi}^{(a)} \varepsilon^{(a)}(t)$$

**[0745]** Example of Constrained Hamiltonian System in Lagrangian Form

**[0746]** Let

$$L(q, \dot{q}) = \frac{1}{2} \dot{q}^2 + \dot{q}^{(1)} \dot{q}^{(2)} + \frac{1}{2} (q^{(1)} - q^{(2)})^2 \quad (3.17-17)$$

$$E^{(o)} = \begin{bmatrix} \frac{d}{dt} \frac{\partial}{\partial \dot{q}^{(1)}} - \frac{\partial L}{\partial q^{(1)}} \\ \frac{d}{dt} \frac{\partial}{\partial \dot{q}^{(2)}} - \frac{\partial L}{\partial q^{(2)}} \end{bmatrix} = \begin{bmatrix} \ddot{q}^{(1)} + 2\dot{q}^{(2)} - q^{(1)} \\ q^{(1)} - q^{(2)} \end{bmatrix} \quad (3.17-18)$$

$$W = \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \quad (3.17-19)$$

$$k = \begin{bmatrix} \dot{q}^{(2)} - q^{(1)} + q^{(2)} \\ -\dot{q}^{(1)} - q^{(2)} + q^{(1)} \end{bmatrix} \quad (3.17-20)$$

The only 0 mode is

$$u^{(o)} = [0, 1]$$

Then

**[0747]**

$$E^{(o)} = W^{(o)} \ddot{q} + k^{(o)} \quad (3.17-21)$$

$$= \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix} \begin{bmatrix} \ddot{q}^{(1)} \\ \ddot{q}^{(2)} \end{bmatrix} + \begin{bmatrix} \dot{q}^{(2)} - q^{(1)} + q^{(2)} \\ -\dot{q}^{(1)} - q^{(2)} + q^{(1)} \end{bmatrix}$$



-continued

Then

$$\begin{aligned} u^{(o)}E^{(o)} &= [0 \quad 1] \left[ \begin{array}{c} \left[ \begin{array}{cc} 1 & 0 \\ 0 & 0 \end{array} \right] \begin{bmatrix} \ddot{q}^{(1)} \\ \ddot{q}^{(2)} \end{bmatrix} + \begin{bmatrix} \dot{q}^{(2)} - q^{(1)} + q^{(2)} \\ -\dot{q}^{(1)} - q^{(2)} + q^{(1)} \end{bmatrix} \end{array} \right] \\ &= -\dot{q}^{(1)} - q^{(2)} + q^{(1)} \\ &= 0 \end{aligned}$$

on shell. Then there are no gauge identities for  $E^{(o)}$ . Now construct  $E^{(1)}$ .

$$E^{(1)} = \begin{bmatrix} E^{(o)} \\ \frac{t}{dt} u^{(o)} E^{(o)} \end{bmatrix} = \begin{bmatrix} \dot{q}^{(2)} - q^{(1)} + q^{(2)} \\ -\dot{q}^{(1)} - q^{(2)} + q^{(1)} \\ -\ddot{q}^{(1)} - \dot{q}^{(2)} + \dot{q}^{(1)} \end{bmatrix}$$

which can be written

$$\begin{aligned} E^{(1)} &= W^{(1)}\ddot{q} + k^{(1)} \\ &= \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ -1 & 0 \end{bmatrix} \begin{bmatrix} \ddot{q}^{(1)} \\ \ddot{q}^{(2)} \end{bmatrix} + \begin{bmatrix} \dot{q}^{(2)} - q^{(1)} + q^{(2)} \\ -\dot{q}^{(1)} - q^{(2)} + q^{(1)} \\ -\dot{q}^{(2)} + \dot{q}^{(1)} \end{bmatrix} \end{aligned}$$

There zero modes of  $W^{(1)}$  are

$$W^{(1)} \left\{ \begin{array}{l} [0 \quad 1 \quad 0] \\ [1 \quad 0 \quad 1] \end{array} \right.$$

**[0748]** The first zero mode is the previous one augmented by one dimension and reproduces the previous constraint. The second mode reproduces the negative of the constraint (3.17-21). That is,

$$v^{(1)}E^{(1)} = -u^{(o)}E^{(o)}$$

with  $v^{(1)} = [1 \quad 0 \quad 1]$ . This leads to the gauge identity

$$G^{(1)} = v^{(1)}E^{(1)} + u^{(o)}E^{(o)} = 0$$

Companionship: Reconciling Agents in the Network.

**[0749]** The outline of the companionship process is as follows for a system of N agents.

**[0750]** Determine the state action space of the system for N-1 agents to create a Tellegen decision element.

**[0751]** Update the remaining agent with the Tellegen DE.

**[0752]** Repeat process so that all N agents are updated with respect to their Tellegen DEs.

**[0753]** User submits query.

**[0754]** System used KB to establish equations of motion for system in Lagrangian or Hamiltonian form.

**[0755]** System determines optimal trajectory via optimization algorithm of the equations of motion that conform to the principle of least action.

**[0756]** System returns solution which is a point in the phase space and also serves as an answer to the query.

**[0757]** It will also be appreciated that in some embodiments the functionality provided by the routines discussed above may be provided in alternative ways, such as being split among more routines or consolidated into fewer routines. Similarly, in some embodiments illustrated routines may provide more or less functionality than is described, such as when other illustrated routines instead lack or include such functionality respectively, or when the amount of functionality that is provided is altered. In addition, while various operations may be illustrated as being performed in a particular manner (e.g., in serial or in parallel, synchronously or asynchronously, etc.) and/or in a particular order, those skilled in the art will appreciate that in other embodiments the operations may be performed in other orders and in other manners. Those skilled in the art will also appreciate that the data structures discussed above may be structured in different manners, such as by having a single data structure split into multiple data structures or by having multiple data structures consolidated into a single data structure. Similarly, in some embodiments illustrated data structures may store more or less information than is described, such as when other illustrated data structures instead lack or include such information respectively, or when the amount or types of information that is stored is altered.

**[0758]** From the foregoing it will be appreciated that, although specific embodiments have been described herein for purposes of illustration, various modifications may be made without deviating from the spirit and scope of the invention. Accordingly, the invention is not limited except as by the appended claims and the elements recited therein. In addition, while certain aspects of the invention are presented below in certain claim forms, the inventors contemplate the various aspects of the invention in any available claim form. For example, while only some aspects of the invention may currently be recited as being embodied in a computer-readable medium, other aspects may likewise be so embodied.

1. A computer-implemented method comprising:

receiving, by an automated control system implemented at least in part by one or more computing systems, information for use in controlling electrical power being output by a battery via electrical direct current (DC) supplied by the battery, wherein the received information includes a model based on multiple rules that each has one or more conditions to evaluate and that specify restrictions involving manipulating a DC-to-DC amplifier connected to the battery in a manner to achieve one or more defined goals including maintaining an internal state of the battery that includes an internal temperature of the battery being in a defined range during the controlling of the electrical power being output, wherein the defined range includes a range of internal temperatures of the battery at which the battery operates without causing premature damage; and

controlling, by the automated control system and based on the received information, the electrical power being output by the battery, including:

obtaining, by the automated control system and at one or more times, sensor information identifying current values for one or more attributes of the battery at the one or more times, and information about one or more electrical loads to be satisfied at the one or more times;



generating, by the automated control system and based at least in part on the obtained sensor information, one or more estimates of the internal state of the battery for the one or more times, including one or more estimates of the internal temperature of the battery for the one or more times;

determining, by the automated control system, and based at least in part on the model and on the one or more estimates of the internal state of the battery and on the one or more electrical loads, one or more amounts of electrical power for the battery to supply at the one or more times to satisfy at least some of the one or more electrical loads while maintaining the internal temperature of the battery in the defined range during the one or more times; and

implementing, by the automated control system and for each of the one or more determined amounts of electrical power, one or more settings of the DC-to-DC amplifier to cause electrical power being output to satisfy the determined amount of electrical power.

2. The computer-implemented method of claim 1 wherein the one or more times include a first time and a later second time,

wherein the one or more electrical loads include a first electrical load at the first time and a second electrical load at the second time,

wherein the determining of the one or more amounts of electrical power for the battery to supply at the one or more times includes determining to supply a first amount of electrical power at the first time that is less than the first electrical load and determining to supply a second amount of electrical power at the second time that satisfies the second electrical load, and

wherein the implementing of the one or more settings of the DC-to-DC amplifier includes implementing at least one first setting of the DC-to-DC amplifier at the first time to cause the first amount of electrical power to be supplied at the first time, and implementing at least one second setting of the DC-to-DC amplifier at the second time to cause the second amount of electrical power to be supplied at the second time.

3. The computer-implemented method of claim 2 wherein the determining to supply the first amount of electrical power at the first time that is less than the first electrical load includes determining that supplying a larger amount of electrical power at the first time to satisfy the first electrical load will cause the internal state of the battery to exceed the defined range.

4. The computer-implemented method of claim 1 wherein the determining of the one or more amounts of electrical power for the battery to supply at the one or more times includes:

determining a first amount of electrical power at a first time that satisfies a first electrical load at the first time and that the battery is capable of supplying;

determining, based at least in part on one of the one or more estimates of the internal state of the battery that is for the first time, that the internal state of the battery exceeds the defined range; and

determining, based at least in part on the determining that the internal state of the battery exceeds the defined range, a second amount of electrical power that is less than the first amount and is not sufficient to satisfy the first electrical load,

and wherein the implementing of the one or more settings of the DC-to-DC amplifier includes implementing at least one first setting of the DC-to-DC amplifier at the first time to cause the second amount of electrical power to be supplied at the first time.

5. The computer-implemented method of claim 1 wherein the DC-to-DC amplifier is a field-effect transistor (FET) amplifier, and wherein the implementing of the one or more settings of the DC-to-DC amplifier includes determining an amount of current to supply and applying a voltage to a gate of the FET amplifier to produce the determined amount of current.

6. The computer-implemented method of claim 1 wherein the DC-to-DC amplifier is part of at least one of a buck converter or a boost converter, and wherein the implementing of the one or more settings of the DC-to-DC amplifier includes determining an amount of voltage to supply, and modifying the at least one of the buck converter or the boost converter to produce the determined amount of voltage.

7. The computer-implemented method of claim 1 wherein the obtained sensor information identifies current values for electrical current, voltage and one or more temperatures associated with the battery, and wherein the generating of the one or more estimates of the internal state of the battery includes using a defined battery temperature model to estimate the internal temperature of the battery based at least in part on the current values.

8. The computer-implemented method of claim 1 wherein the obtained sensor information identifies a current value for at least one of electrical current, voltage or a temperature associated with the battery, and wherein the generating of the one or more estimates of the internal state of the battery also includes estimating an internal chemical reaction state for the battery.

9. The computer-implemented method of claim 1 wherein the obtained sensor information identifies a current value for at least one of electrical current, voltage or a temperature associated with the battery.

10. The computer-implemented method of claim 1 wherein the battery is part of a system having one or more electrically powered devices, and wherein the obtained information about the one or more electrical loads includes an amount of electrical demand from the one or more electrically powered devices.

11. The computer-implemented method of claim 1 wherein the battery is connected to an electrical grid, and wherein the obtained information about the one or more electrical loads includes requests from an operator of the electrical grid for electrical power to be supplied from the battery.

12. The computer-implemented method of claim 1 wherein the model further includes one or more rules that specify restrictions involving manipulating the DC-to-DC amplifier connected to the battery to maintain the internal state of the battery in the defined range during controlling of electrical power being supplied to the battery for charging, and wherein the method further comprises controlling, by the automated control system, the electrical power being supplied to the battery for charging, including:

obtaining, by the automated control system and at one or more additional times, additional sensor information identifying current values for one or more attributes of the battery at the one or more additional times, and



information about one or more electrical supply amounts to be provided to the battery at the one or more additional times;

generating, by the automated control system and based at least in part on the obtained additional sensor information, one or more additional estimates of the internal state of the battery for the one or more additional times;

determining, by the automated control system, and based at least in part on the model and on the one or more additional estimates of the internal state of the battery and on the one or more electrical supply amounts, one or more additional amounts of electrical power for the battery to receive at the one or more additional times to accept at least some of the one or more electrical supply amounts while maintaining the internal state of the battery in the defined range during the one or more additional times; and

implementing, by the automated control system and for each of the one or more determined additional amounts of electrical power, one or more settings of the DC-to-DC amplifier to cause electrical power being supplied to the battery to satisfy the determined additional amount of electrical power.

**13.** The computer-implemented method of claim **12** wherein the battery is part of a system having a solar power generator, and wherein the obtained information about the one or more electrical supply amounts to be provided to the battery includes an amount of electrical supply available from the solar power generator.

**14.** The computer-implemented method of claim **1** wherein the model included in the received information is configured for a battery type of the battery, and wherein the method further comprises adapting the model to information specific to the battery by, during an initial training period before the controlling of the electrical power being output, monitoring changes in values of the one or more attributes of the battery and changes in electrical loads being supplied by the battery, and modifying the model to correspond to the monitored changes in values and the monitored changes in electrical loads.

**15.** The computer-implemented method of claim **1** wherein the controlling of the electrical power being output further includes adapting the model included in the received information to additional information specific to the battery by monitoring changes over time in values of the one or more attributes of the battery that are based at least in part on increasing impedance of the battery, and modifying the model to correspond to the monitored changes, and wherein at least one amount of electrical power for the battery to supply that is determined and is caused to be output by the implementing is based at least in part on the modified model.

**16.** The computer-implemented method of claim **1** wherein the automated controller system includes a battery controller component connectively coupled to the battery and further includes a control action determination component that is executing remotely from the battery on at least one of the one or more computing systems and that communicates with the battery controller component over one or more computer networks, wherein the control action determination component performs at least the determining of the one or more amounts of electrical power and further performs sending instructions over the one or more computer networks to the battery controller component about the determined one or more amounts of electrical power, and

wherein the battery controller component performs at least the implementing of the one or more settings of the DC-to-DC amplifier based at least in part on the sent instructions.

**17.** The computer-implemented method of claim **16** further comprising, controlling, by the battery controller component and at an additional time without receiving any instructions over the one or more computer networks from the control action determination component, the electrical power being output by the battery, including:

obtaining, by the battery controller component and at the additional time, additional sensor information identifying current values for the one or more attributes of the battery at the additional time;

generating, by the battery controller component and based at least in part on the obtained additional sensor information, an additional estimate of the internal state of the battery for the additional time;

determining, by the battery controller component, and based at least in part on the additional estimate of the internal state of the battery, an additional amount of electrical power for the battery to supply at the additional time to maintain the internal state of the battery in the defined range for the additional time; and

implementing, by the battery controller component, one or more settings of the DC-to-DC amplifier to cause electrical power being output to satisfy the determined additional amount of electrical power.

**18.** The computer-implemented method of claim **1** further comprising, controlling, by the automated control system and at an additional time without a current electrical load, the electrical power being output by the battery, including:

obtaining, by the automated control system and at the additional time, additional sensor information identifying current values for the one or more attributes of the battery at the additional time;

generating, by the automated control system and based at least in part on the obtained additional sensor information, an additional estimate of the internal state of the battery for the additional time;

determining, by the automated control system, and based at least in part on the model and on the additional estimate of the internal state of the battery, an additional amount of electrical power for the battery to supply at the additional time to maintain the internal state of the battery in the defined range for the additional time; and

implementing, by the automated control system, one or more settings of the DC-to-DC amplifier to cause electrical power being output to satisfy the determined additional amount of electrical power.

**19.** The computer-implemented method of claim **1** wherein the determining of the one or more amounts of electrical power for the battery to supply at the one or more times further includes retrieving information about at least one previous amount of electrical power that the battery supplied, and selecting the determined one or more amounts of electrical power to further control a rate of change in amount of electrical power that the battery supplies, as part of maintaining the internal state of the battery in the defined range during the one or more times.

**20.** The computer-implemented method of claim **1** further comprising identifying at least one previous electrical load, determining one or more amounts of change in electrical load from the at least one previous electrical load to the one



or more electrical loads, and further performing the selecting of the determined one or more amounts of electrical power based at least in part on the determined one or more amounts of change in electrical load.

**21.** The computer-implemented method of claim **1** further comprising identifying at least one previous estimate of the internal state of the battery that is based at least in part on an internal impedance of the battery, determining one or more amounts of change in estimated internal impedance of the battery from the at least one previous estimate to the one or more estimates, and further performing the selecting of the determined one or more amounts of electrical power based at least in part on the determined one or more amounts of change in the estimated internal impedance of the battery.

**22.** The computer-implemented method of claim **1** further comprising identifying at least one previous estimate of the internal state of the battery that is based at least in part on the internal temperature of the battery, determining one or more amounts of change in estimated internal temperature of the battery from the at least one previous estimate to the one or more estimates, and further performing the selecting of the determined one or more amounts of electrical power based at least in part on the determined one or more amounts of change in the estimated internal temperature of the battery.

**23.** A non-transitory computer-readable medium having stored contents that cause one or more devices implementing at least part of an automated control system to perform a method, the method comprising:

receiving, by the automated control system, information for use in controlling electrical power from a battery that supplies electrical direct current (DC), wherein the received information includes a model of the battery that specifies restrictions involving manipulating a DC-to-DC amplifier connected to the battery to achieve one or more defined goals including maintaining an internal temperature of the battery in a defined range during the controlling of the electrical power from the battery, wherein the defined range includes a range of internal temperatures of the battery at which the battery operates to reduce premature damage; and

controlling, by the automated control system and based on the received information, the electrical power from the battery, including:

obtaining, by the automated control system, sensor information identifying values for one or more attributes of the battery at multiple times, and information about electrical loads to be satisfied at the multiple times;

generating, by the automated control system and based at least in part on the obtained sensor information, estimates of the internal temperature of the battery for the multiple times;

determining, by the automated control system for one or more of the multiple times, one or more amounts of electrical power for the battery to supply at the one or more times to satisfy at least some of the electrical loads while maintaining the internal temperature of the battery in the defined range, wherein the determining is based at least in part on at least one previous amount of electrical power supplied by the battery in order to control a rate of change in amount of electrical power that the battery supplies, and wherein the determining is further based at least in

part on the model and on the estimates of the internal temperature of the battery and on the electrical loads; and

implementing, by the automated control system and for each of the one or more determined amounts of electrical power, one or more settings of the DC-to-DC amplifier to cause electrical power being output to satisfy the determined amount of electrical power.

**24.** The non-transitory computer-readable medium of claim **23** wherein the model of the battery includes one or more rules that limit changing an amount of electrical power being output from the battery to be less than a defined threshold limit for the rate of change, wherein the stored contents include software instructions that, when executed, further cause at least one of the one or more devices to identify the at least one previous amount of electrical power supplied by the battery and to determine a target rate of change that maintains the internal temperature of the battery in the defined range and that is within the defined threshold limit, and wherein the determining of the one or more amounts of electrical power for the battery to supply at the one or more times is further performed to satisfy the determined target rate of change and based at least in part on the one or more rules.

**25.** The non-transitory computer-readable medium of claim **23** wherein the determining of the one or more amounts of electrical power for the battery to supply at the one or more times further includes identifying at least one previous electrical load, determining one or more amounts of change in electrical load from the at least one previous electrical load to the electrical loads to be satisfied at the multiple times, and selecting at least one of the determined one or more amounts of electrical power based at least in part on the determined one or more amounts of change in electrical load.

**26.** The non-transitory computer-readable medium of claim **23** wherein the determining of the one or more amounts of electrical power for the battery to supply at the one or more times further includes identifying at least one previous estimate of an internal state of the battery that is based at least in part on an internal impedance of the battery, determining one or more amounts of change in estimated internal impedance of the battery from the at least one previous estimate to subsequent estimates of the internal state of the battery for the multiple times, and selecting at least one of the determined one or more amounts of electrical power based at least in part on the determined one or more amounts of change in the estimated internal impedance of the battery.

**27.** The non-transitory computer-readable medium of claim **23** wherein the determining of the one or more amounts of electrical power for the battery to supply at the one or more times further includes identifying at least one previous estimate of an internal state of the battery that is based at least in part on the internal temperature of the battery, determining one or more amounts of change in estimated internal temperature of the battery from the at least one previous estimate to the estimates of the internal temperature of the battery for the multiple times, and selecting at least one of the determined one or more amounts of electrical power based at least in part on the determined one or more amounts of change in the estimated internal temperature of the battery.



**28.** A system comprising:  
 one or more hardware processors of one or more devices;  
 and  
 one or more memories storing instructions that, when executed by at least one of the one or more hardware processors, cause the at least one hardware processor to implement an automated control system that controls electrical power supplied from or to a battery, by:  
 receiving a model of the battery having specified restrictions for manipulating an actuator that controls direct current (DC) characteristics of the electrical power supplied from or to the battery;  
 obtaining sensor information identifying values for one or more attributes of the battery at an indicated time, and additional information about at least one of an electrical load to be satisfied at the indicated time or an electrical supply amount available to be provided to the battery at the indicated time;  
 generating, based at least in part on the obtained sensor information, one or more estimates for the indicated time of an internal state of the battery that is not directly observable;  
 determining, as part of maintaining the internal state of the battery in a defined range for the indicated time, and based at least in part on the model and on the one or more estimates of the internal state of the battery, an amount of electrical power for the battery to supply at the indicated time to satisfy at least some of the electrical load, or an amount of electrical power for the battery to receive at the indicated time by accepting at least some of the electrical supply amount; and

implementing one or more settings of the actuator to cause electrical power being output for the indicated time to satisfy the determined amount of electrical power for the battery to supply, or to cause electrical power being supplied to the battery for the indicated time to satisfy the determined amount of electrical power for the battery to receive.

**29.** The system of claim **28** wherein the actuator is a DC-to-DC amplifier, and wherein the automated control system controls electrical power supplied from the battery by determining the amount of electrical power for the battery to supply at the indicated time to satisfy at least some of the electrical load and by implementing one or more settings of the DC-to-DC amplifier to cause electrical power being output for the indicated time to satisfy the determined amount of electrical power for the battery to supply.

**30.** The system of claim **28** wherein the received model of the battery includes multiple rules that specify restrictions involving manipulating the actuator in a manner to achieve one or more defined goals including maintaining the internal state of the battery in the defined range during controlling of the electrical power supplied from or to the battery, wherein the multiple rules include one or more rules to not charge the battery if the battery has a current charge that exceeds a first defined threshold, and one or more rules to not discharge the battery if the battery has a current charge that is below a second defined threshold, and one or more rules to not charge or discharge the battery if doing so would result in an estimated internal temperature of the battery being outside a defined range.

\* \* \* \* \*