

Framework for Knowledge-Based Fault Detection and Diagnostics in Multi-Domain Systems: Application to Heating Ventilation and Air Conditioning Systems

Parastoo Delgoshaei

Department of Civil and Environmental Engineering,
University of Maryland, College Park, MD 20742, USA
E-mail: p.delgoshaei@gmail.com

Mark A. Austin

Department of Civil and Environmental Engineering,
and Institute for Systems Research,
University of Maryland, College Park, MD 20742, USA
E-mail: austin@isr.umd.edu

Abstract—State-of-the-art fault detection methods are equipment and domain specific and non-comprehensive. As a result, the applicability of these methods in different domains is very limited and they can achieve significant levels of performance by having knowledge of the domain and the ability to mimic human thinking in identifying the source of a fault with a comprehensive knowledge of the system and its surroundings. This paper presents a comprehensive semantic framework for fault detection and diagnostics (FDD) in systems simulation and control. Our proposed methodology entails of implementation of the knowledge bases for FDD purposes through the utilization of ontologies and offers improved functionalities of such system through inference-based reasoning to derive knowledge about the irregularities in the operation. We exercise the proposed approach by working step by step through the setup and solution of a fault detection and diagnostics problem for a small-scale heating, ventilating and air-conditioning (HVAC) system.

Keywords-Fault Detection and Diagnostics; Heating Ventilating and Air-Conditioning (HVAC); Inference-Based, Knowledge Base, Ontologies; Reasoning.

I. INTRODUCTION

This paper is concerned with the development of ontology and rule-based modeling abstractions, procedures, and prototype software for automated fault detection and diagnostic (FDD) analysis of condition-based maintenance in multi-domain systems (e.g., buildings, health monitoring, power plants and aviation systems). The article builds upon our previous work [1]–[3] on behavior modeling and analysis of engineering systems with semantic web technologies.

A. Problem Statement

Automated fault detection and diagnostic (FDD) techniques provide a means of detecting unwanted conditions (i.e., “faults”) in systems by recognizing deviations in real-time or recorded data values from expected values, and then diagnosing the causes leading to the faults. Automated fault detection and diagnostic (FDD) techniques provide mechanisms for condition-based maintenance of engineered systems (e.g., buildings, health monitoring, power plants and aviation systems). Proper implementation of FDD can enable pro-active identification and remediation of faults before they become significantly deleterious to the safety, security, or efficiency of the operating system.

Within the building sector, degraded or poorly-maintained equipment currently accounts for 15 to 30 % of energy consumption in commercial buildings [4]. Approximately 50 to 67 % of air conditioners (residential and commercial) are either improperly charged or have airflow issues [5] and [6]. Faulty heating, ventilating, air conditioning, and refrigeration (HVAC&R) systems contribute to 1.5 to 2.5 % of total commercial building consumption [7]. Much of this energy usage could be prevented by utilizing automated condition-based maintenance. During the last decade, considerable research has focused on the development of FDD methods for HVAC&R systems. This work has been driven, in part, by the historically less-than-optimal operation of many state-of-the-art HVAC systems. Yet, in spite of recent advances in building simulation, automation and control (see the arrangement of ontologies, rules, reasoning and simulation software in Figure 1), automatic methods for FDD of building systems remain at a relatively immature stage of development. As a result, we require more advanced FDD techniques that leverage the untapped capabilities of building automation integrated with methods in artificial intelligence and semantic modeling. These interdisciplinary FDD systems can benefit from utilizing knowledge repositories for storing automation/simulation data and the inference-based reasoning techniques to obtain additional higher information, such as sensors location, equipment service area.

B. Objectives and Scope

This paper describes a framework for knowledge-based fault detection and diagnostics in multi-domain systems, with a focus on applications to HVAC Systems. In a departure from state-of-the-art developments in ontology engineering, which place a priority on the development and testing of ontologies alone, our objective is to create a modeling framework that supports: (1) concurrent data-driven development of domain models, ontologies and rules, and (2) inference-based reasoning for detection of faults and their causes. The proposed method employs the Web Ontology Language (OWL) [8] and Jena API [9] for the development of semantic models (ontologies and rules) spanning the building, mechanical equipment, sensor, fault detection and diagnostics (FDD), occupant and weather domains. Support for spatial reasoning among entities is provided at the meta-domain level.

The remainder of this paper proceeds as follows: Section

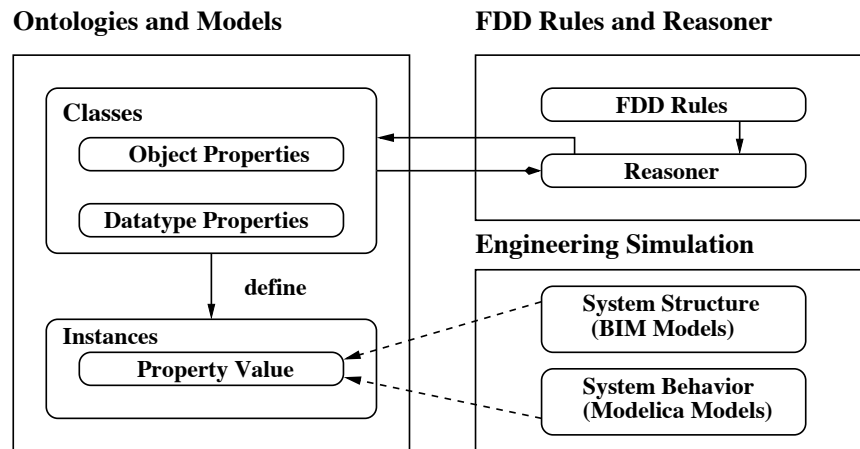


Figure 1. Architecture of engineering simulations connected to semantic models (ontologies and rules) reasoners for fault detection and diagnostic analysis (Adapted from Delgoshai, Austin and Pertzborn [2]).

II describes related work in FDD. Section III contains a brief introduction to the uses of the Semantic Web and its enabling technologies. The proposed methodology is described in Section IV. Sections V and VI cover: (1) the meta-domain and domain-specific ontologies and rules, respectively, and (2) a step-by-step procedure for detection and analysis of system faults. Section VII presents a case study problem that involves detection of faults in a simple building – procedures for reasoning across multiple domains are presented. Finally, the conclusions of this study and a discussion of next steps is presented in Section VIII.

II. RELATED WORK

Recent advances in building automation technologies provide a means for sensing and collecting the data needed for software applications to automatically detect and diagnose faults in buildings. During the past few decades a variety of FDD techniques have been developed in different domains, including model-based, rule-based, knowledge-based, and simulation-based approaches. Katipamula and Brambley summarizes FDD research for HVAC systems [4]. Their work also describes different fundamental FDD methods under the two main categories of model-based and empirical (history-based) approaches. The major difference is in the nature of the knowledge used to formulate the diagnostics. Model-based diagnostics evaluate residuals between actual system measurements and *a priori* models (e.g., first principle models). Data-driven empirical strategies, on the other hand, do not require *a priori* models. The models used in model-based methods can be quantitative or qualitative. Quantitative models represent the requisite *a priori* knowledge of the system in terms of mathematical equations, typically as explicit descriptions of the physics underlying system components. Qualitative models, conversely, combine concepts such as descriptive “states” and “rules” into statements that are axiological instead of mathematical, expressing operational correctness or desirability through an axiology, a value system, appropriate to each physical application. As a result, the building system operation can be continuously classified as being either faulty or not faulty.

Rule-based strategies are one example of qualitative model-based FDD methods. Rules can be based on first principles or they can be inferred from historical experiments, but in either case they represent expert qualitative knowledge that no purely quantitative representation could model. The first diagnostic expert systems for technical fault diagnosis were developed at MIT by Scherer and White [10]. Since then, diagnostic systems have evolved from rule-based to model-based and expert systems approaches. Semantic models offer a means for the representation of distributed and explicit knowledge and provide ways through inference-based rules to derive implicit knowledge. Berners-Lee and co-workers [11] points out to the benefits of ontology usage for knowledge representation, and utilizing high-level reasoning capabilities in the area of agent-based control solutions. Exploitation of semantics and ontologies in the area of agent-based engineering systems has become one of the hot topics recently. The main reason behind this trend is the success and promotion of Semantic Web technologies to enable languages that are both machine and human processable. Semantic Web-based applications have been developed in the areas of health care [12], biology [13], [14], and transportation [15]. In the area of fault detection and diagnostics, Batic [16] has developed an ontology-based fault detection and diagnosis systems and tested it on airport ontologies to detect the high level irregularities in the operation of airport heating/cooling plants. Also, Schumann [17] highlights the potential impacts of artificial intelligence techniques such as ontologies on tackling the challenges in obtaining a unified diagnosis framework. The benefit of this approach is that ontologies are an essential technology guaranteeing data and information interoperability in heterogeneous and content-rich environments [18] which is at heart of comprehensive fault detection and diagnostic methods.

III. THE SEMANTIC WEB

A. Semantic Models

Semantic models consist of ontologies, graphs of individuals (specific instances), and inference-based rules in the form of *if <conditions> then <consequences>*. Together, these entities and mechanisms allow for the construction and execution of domain-specific knowledge bases.

Framework for Concurrent Data-Driven Development of Domain Models, Ontologies and Rules

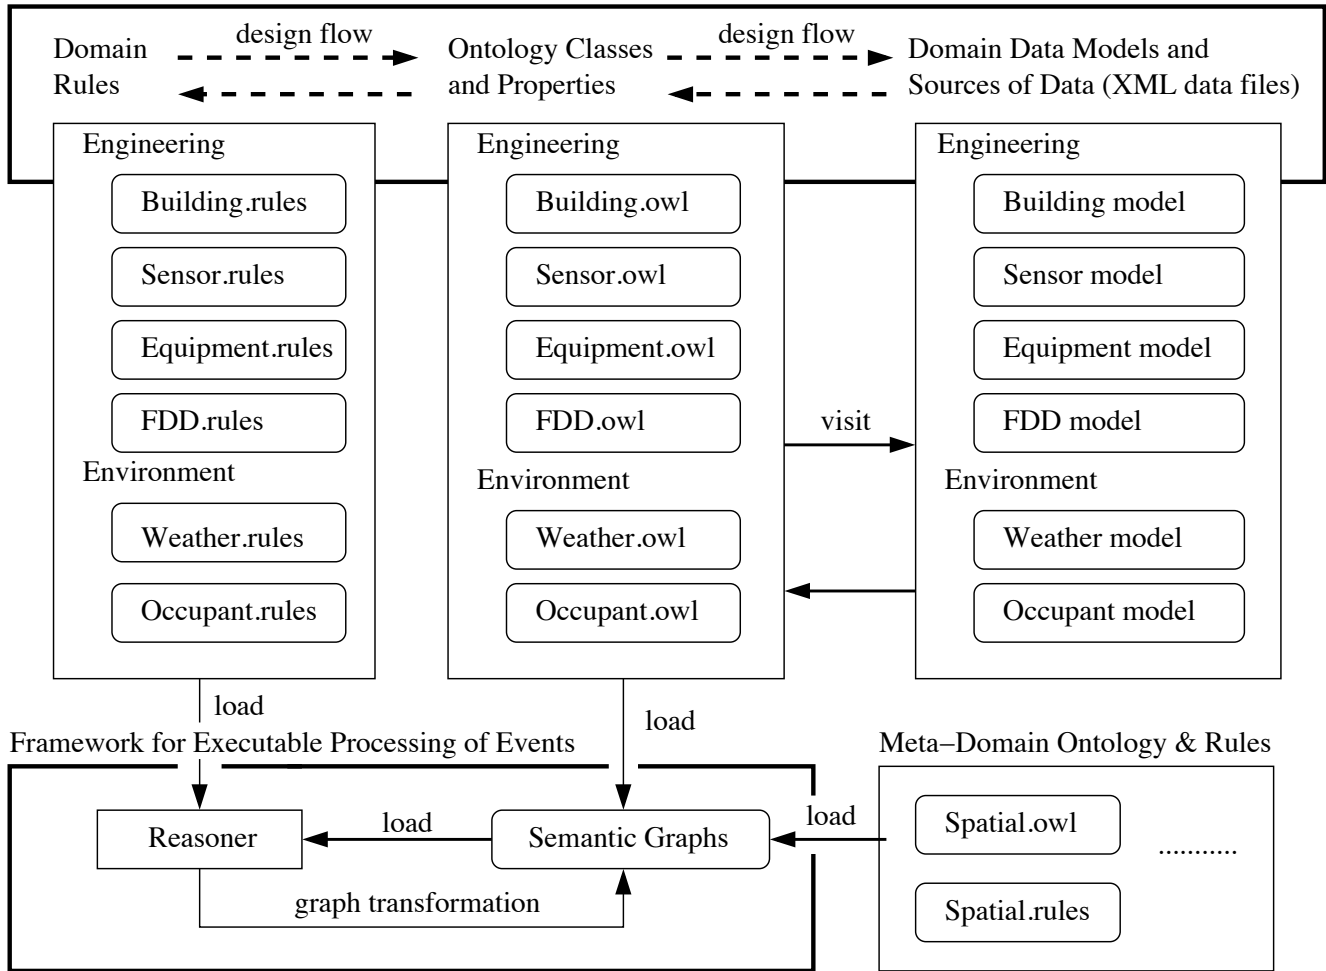


Figure 2. Proposed architecture for: (1) concurrent data-driven development of domain models, ontologies and rules, and (2) executable processing of events.

Ontologies: An ontology is a formal and explicit representation of the concepts, referred to as “classes” (e.g., cooling coil, valve), and their interrelationships in a domain. The classes may have attributes that are stored as a simple data type properties” (e.g., coil setpoint). Support for semantic relationships between classes is provided by object properties (e.g., a coil has as a valve). For the representation of domains where there are many variations to be represented, but common data properties among those variants, ontology languages provide support for the organization of similar concepts into hierarchies, and support for propagation of data and object properties through hierarchies via inheritance mechanisms. We may wish to state, for example, that a cooling coil is a type of coil. In this hierarchy, the class cooling coil is a subclass of the class coil. And the class coil is superclass of the class cooling coil. The details of the classes, data properties and object properties can be summarized as follows:

- **Classes:** Valve, Cooling Coil
- **Datatype properties:** coilTemperature (double), isClosed (Boolean), coilSetpoint(double)
- **Object Property:** hasValve

Individuals: Individuals are instances of ontology concepts, and their purpose is to represent the data in a domain, e.g.,

- **Individuals:** ValveI, ValveII, Ccoil, Hcoil
- **Storing individuals:** ;Hcoil hasValve ValveII;

One common syntax for representing facts about a domain is the triple structure <subject, predicate, object>.

Inference Rules: Inference rules and their associated reasoning mechanisms provide a way derive new information based on the existing data stored in the ontology in the form of: if <conditions> then <consequent>. For example, the script:

```

Logical Rule:
(?coil rdf:type coil) (?coil setPoint ?sp)
(?coil coilTemperature ?cp) equal(?cp,?sp)
(?coil hasValve ?valve) -> (?valve isClosed true)

Stored individuals : <Hcoil hasValve ValveII>
                    <Ccoil coilTemperature 35>
                    <Ccoil coilSetpoint 35>
Inferred Knowledge: <ValveII isClosed true>
    
```

takes existing facts and rule that covers the setpoint and temperature of a coil to infer that a valve is closed.

A key benefit of semantic modeling frameworks is that the ontologies and rules are human readable, yet they can also be compiled into code that is executable on machines.

IV. METHODOLOGY

A. Architecture Framework

In state-of-the-art development of semantic models, a common strategy is to provide classes and data properties for all possible configurations within a domain, as well as linkage to related domains. For example, in the integrated model-centric engineering ontologies (IMCE) developed at JPL (Jet Propulsion Laboratory) during the 2000-2010 era [19], [20], the electrical engineering ontology (i.e., *electrical.owl*) imports the mechanical engineering ontology (i.e., *mechanical.owl*). Both the electrical and mechanical engineering ontologies import a multitude of foundation ontologies (e.g., *analysis.owl*, *mission.owl*, *base.owl*, *project.owl*, *time.owl*) and make extensive use of multiple inheritance mechanisms in the development of new classes. The result is ontologies containing more than two hundred classes, with some classes containing three or four dozen data and object properties. Notions of “simplicity in system design” through modularity of semantic models (e.g., bundling of ontologies and rules) do not seem to exist.

In a first step toward mitigating these complexities, we propose a semantic modeling framework (see Figure 2) that supports: (1) concurrent data-driven development of domain models, ontologies and rules, and (2) executable processing of incoming faults. Instead of creating ontologies and then developing a few rules for validation of model properties, our goal is to put the development of data, ontologies and rules on an equal footing. A key advantage of this approach is that it forces designers to provide semantic representations for data that are needed in decision making, and increases the likelihood that data not needed for decision making will be left out. Rules will be developed for verification of domain properties and processing of faults through reasoning with data sources, possibly from multiple domains. Implementation of the latter goal leads to semantic graphs that will dynamically adapt to the consequences of incoming data and events (e.g., changing occupant locations and weather events) acting on the system.

Our second strategy is to minimize the use of multiple inheritance in the specification of OWL ontologies and, instead, explore opportunities for replacing inheritance relationships by object property relations. In order for the architectural framework to be both scalable and adaptable to changing external conditions, the ontologies will need to be modular, and the rules will need to act both within a domain and across domains.

B. Working with Jena and Jena Rules

Our prototype software implementation makes extensive use of Apache Jena and Jena Rules. Apache Jena [9] is an open source Java framework for building Semantic Web and linked data applications. Jena provides APIs (application programming interfaces) for developing code that handles RDF

(resource description framework), RDFS, OWL (web ontology language) and SPARQL (support for query of RDF graphs). The Jena rule-based inference subsystem is designed to allow a range of inference engines or reasoners to be plugged into Jena. Jena Rules is one such engine.

real world building environment

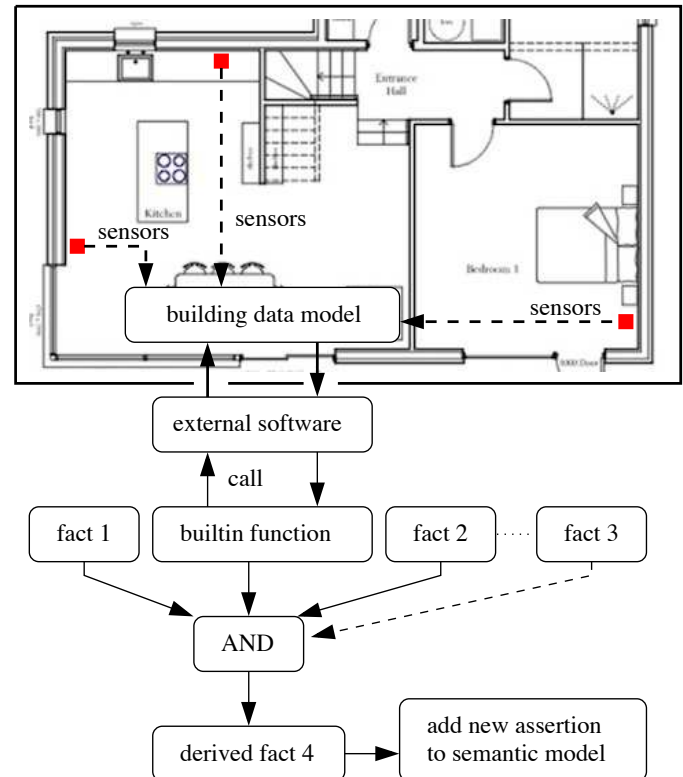


Figure 3. Framework for forward chaining of facts and results of builtin functions to new assertions (derived facts).

Jena Rules employs facts and assertions described in OWL to infer additional facts from instance data and class descriptions. As illustrated in Figure 3, it also provides support for the development of builtin functions that can link to external software programs and streams of data sensed in the real world. For the implementation of the vision implied by Figure 2, particularly support for spatial and temporal reasoning, the latter turns out to be crucially important because, by default, OWL only provides builtin datatype support for numbers (i.e., float and double), booleans (i.e., to represent true and false) and character strings (i.e., string). To combat the lack of support for complex data types, such as those needed to represent data for spatial and temporal reasoning, we adopt a strategy of embedding the relevant data in character strings, and then designing builtin functions and external software that can parse the data into spatial/temporal models, and then make the reasoning computations that are required.

C. Data-Driven Approach to Generation of Individuals in Semantic Graphs

In the proposed framework semantic models are the composition of ontologies, rules and data.

Figure 4 illustrates a data-driven approach to the generation of individuals in semantic graphs.

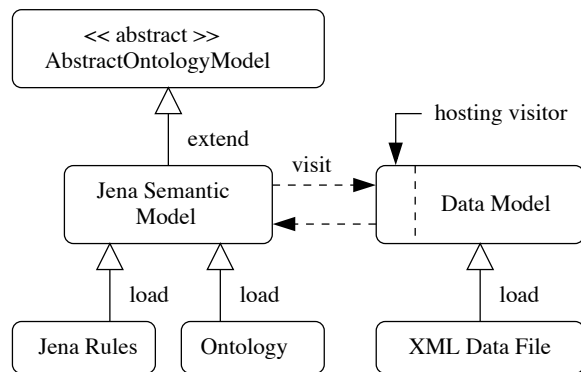


Figure 4. Data-driven approach to generation of individuals in semantic graphs.

First, data is imported into Java Object data models using JAXB, the XML binding for Java. After the ontologies and rules have been loaded into the Jena Semantic Model, the semantic model creates instances of the relevant OWL ontologies by visiting the data model and gathering information on the individuals within a particular domain (e.g., building, sensor, occupant). Once the data has been transferred to the Jena Semantic Model and used to create an ontology instance, the rules are applied.

V. META-DOMAIN ONTOLOGIES AND RULES

Meta-domain ontologies and rules have universal application across domains, and include concepts such as time, space, physical units and currency. This study employs spatial reasoning to determine the relationship of sensor and occupants to geometric entities such as rooms and building zones.

A. Spatial Ontology and Rules

Spatial logic is concerned with regions and their connectivity, allowing one to address issues of the form: what is true, and where? Formal theories for reasoning with space – points, lines, and regions – are covered by region connected calculus [21]. A robust implementation of two-dimensional spatial entities and associated reasoning procedures is provided by the Java Topology Suite (JTS) [22].

Spatial Ontology and Rules for Spatial Reasoning. Figure 5 shows an abbreviated representation of our experimental spatial (geometry) ontology and associated data and object properties. High-level classes – abstract concepts – are provided for entities that represent singular geometry (e.g., AbstractGeometry) and groups of entities (e.g., AbstractGeometryCollection). Specific types of geometry (e.g., Polygon, MultiPoint) are organized into a hierarchy similar to the Java implementation in JTS. The high-level class AbstractGeometry contains a Datatype property, hasGeometry, which stores a string representation of the JTS geometry. For example, the abbreviated string “POLYGON ((0 0, 0 5, ... 0 0))” shows the format for pairs of (x,y) coordinates defining a two-dimensional polygon. This feature allows a semantic model to

visit a domain data model, and gather a complete description of the two-dimensional geometry.

Within Jena Rules, families of builtin functions can be developed to evaluate the geometric relationship between pairs of spatial entities (e.g., to determine whether or not a point is contained within a polygon). Figure 6 shows, for example, the Jena Rule that identifies the room in which a sensor is placed. An English translation of the rule fragments is as follows: If (?r) is a room with geometry (?rg) and string representation (?rjts), and (?s) is a sensor with geometry (?sg) and string representation (?sjts), then the builtin function getPointInPolygon(?sjts,?rjts,?t) will determine if the sensor (point geometry) is inside the room (polygon geometry) and return the result as a boolean (?t). If (?t) is true, then the sensor is inside the room and a new relationship (?s bld:isInRoom ?r) is created. A similar rule would be written to establish the relationship between sensors and HVAC zones.

VI. DOMAIN ONTOLOGIES AND RULES

The domain-specific ontologies and rules are organized into two groups: (1) engineering ontologies and rules, and (2) surrounding environment ontologies and rules. In Figures 7 through 13 we use red rectangles with heavy dashed edges to highlight the classes that participate in the rule checking and/or the case study problem presented in Section VII. For a complete description of the ontologies used in this study we refer the interested reader to Delgoshaei and Austin [23].

A. Engineering Ontologies and Rules

The engineering ontologies and rules cover four domains: (1) buildings, (2) mechanical equipment, (3) sensors, and (4) procedures for fault detection and diagnosis.

Building Ontology and Rules. The prototype building ontology and rules (see Figures 7 and 8) provide computational support for the representation of two-dimensional floorplan geometry, modeling relationships between elements of floorplan geometry and sensors, zones for HVAC control, and building elements such as doors, windows and walls. The latter are modeled as subclasses of a component that has geometry described by a JTS string.

Connections to the mechanical equipment and occupancy domains are achieved through data properties for the building environment state; see, for example, hasRoomSetpoint and isOccupied. Object properties record the relationship of a room to relevant HVAC zones and sensors. Windows have the boolean data property isOpen to record whether or not a particular window is open. As we will soon in the case study problem, this parameter plays a pivotal role in diagnostic analysis of the causes leading to a fault in mechanical equipment.

The prototype software implementation has one rule for determining the spatial relationship among zones of the building. The rule systematically retrieves the JTS geometry of each zone, verifies they are not equal, and then uses the builtin function getPointInPolygon() to verify their geometric relationship. As previously noted, these backend computations are handled by the Java Topology Suite software [22].

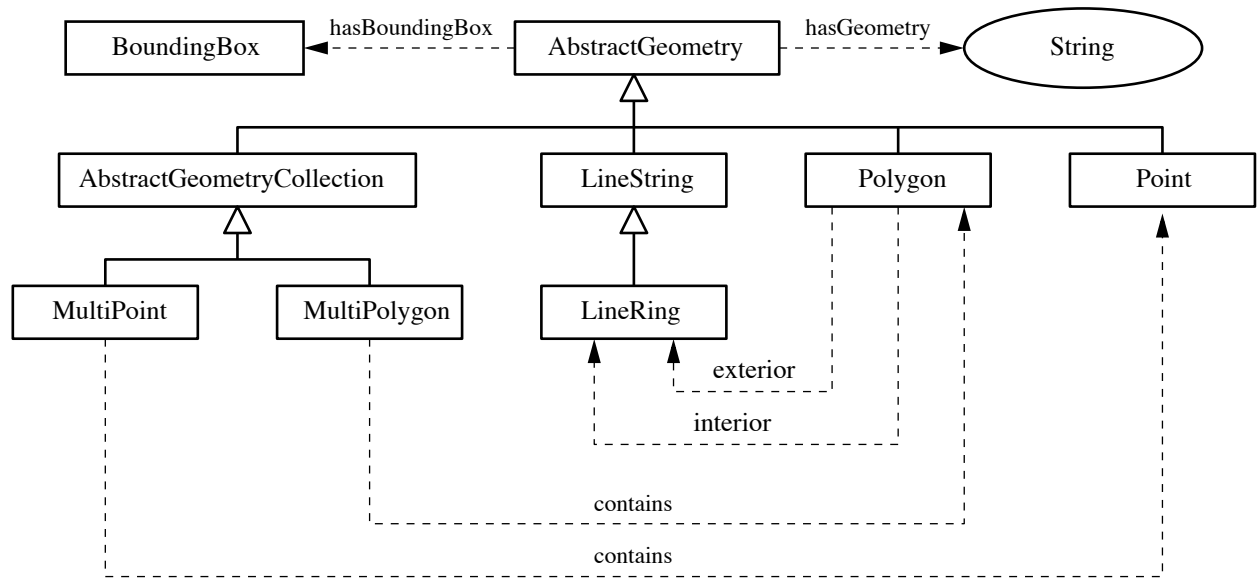


Figure 5. Abbreviated representation of spatial (geometry) ontology and associated data and object properties.

Jena Rules

```

// Rule to check if a sensor is inside a room ...
[ BuildingRule01: (?r rdf:type bld:Room) (?r bld:hasGeometry ?rg) (?rg geom:hasGeometry ?rjts)
  (?s rdf:type sen:Sensor) (?s sen:hasGeometry ?sg) (?sg geom:hasGeometry ?sjts)
  getPointInPolygon(?sjts,?rjts,?t)
  equal(?t, "true"^^xs:boolean) -> (?s bld:isInRoom ?r) ]

```

Figure 6. Rules to determine the rooms in which sensors have been placed.

Mechanical Equipment Ontology and Rules. Figures 9 and 10 illustrate the concepts (i.e., ontology classes), properties (i.e., data and object properties) and rules governing the operation and identification of faults in mechanical systems equipment. In practice, datatype property values associated with the various ontologies will be set from streams of data either performed by a simulation tool (e.g. EnergyPlus, Dymola, TRNSYS) [24]–[26], or perhaps from measurements taken in a real building, working in conjunction with BACnet protocols [27] and a co-simulation middleware.

The semantic graph shown in Figure 9 is quite broad, covering concepts from chillers and fans to zones. The scope of our investigation focuses on faults associated with valves, coils and air handling units. Basic rules (see Figure 10) are provided for: (1) controlling the flow in a valve, (2) determining if a valve is leaky, and (3) identifying situations where the normal operational status of a valve is false. Thus, we are able to determine that when a cooling coil valve is faulty, the associated air handling unit is also faulty.

Sensor Ontology and Rules. Figure 11 shows the classes and properties in our experimental sensor ontology. Our goal is to provide computational support for modeling: (1) sensor operation, including when a sensor reading might be outside an acceptable working range, and (2) determining the location of

a sensor relative to the environment in which it is embedded. These objectives are achieved with three classes: Sensor, Measurement, and the external class Geometry.

Support for modeling various types of sensor (e.g., temperature sensor, flow sensor, and CO₂ sensor) is provided through the definition of specialized sensor classes that subclass Sensor. The class Measurement has data properties to keep track of the current sensor value, the time, and the units associated with the measurement.

Two sensor rules (see Figures 6 and 12) are supported: (1) To determine if a sensor reading is beyond the acceptable range, (2) To determine the room in which the sensor is located. The first rule uses the classes Sensor and Measurement and associate properties. The second rule uses the classes Sensor and Geometry.

Fault Detection and Diagnostic Ontologies, Rules, and Procedures. The fault detection and diagnostic (FDD) ontology (see Figure 13) captures the knowledge needed for: (1) identifying that a fault exists, and (2) systematically diagnosing the fault to find the root causes. The main classes in this process are State, Fault, Hypothesis and Evidence. State is a high-level state representation that has data values – see, for example, the boolean properties hasExpectedValue and hasCurrentValu – common to many types of state representation.

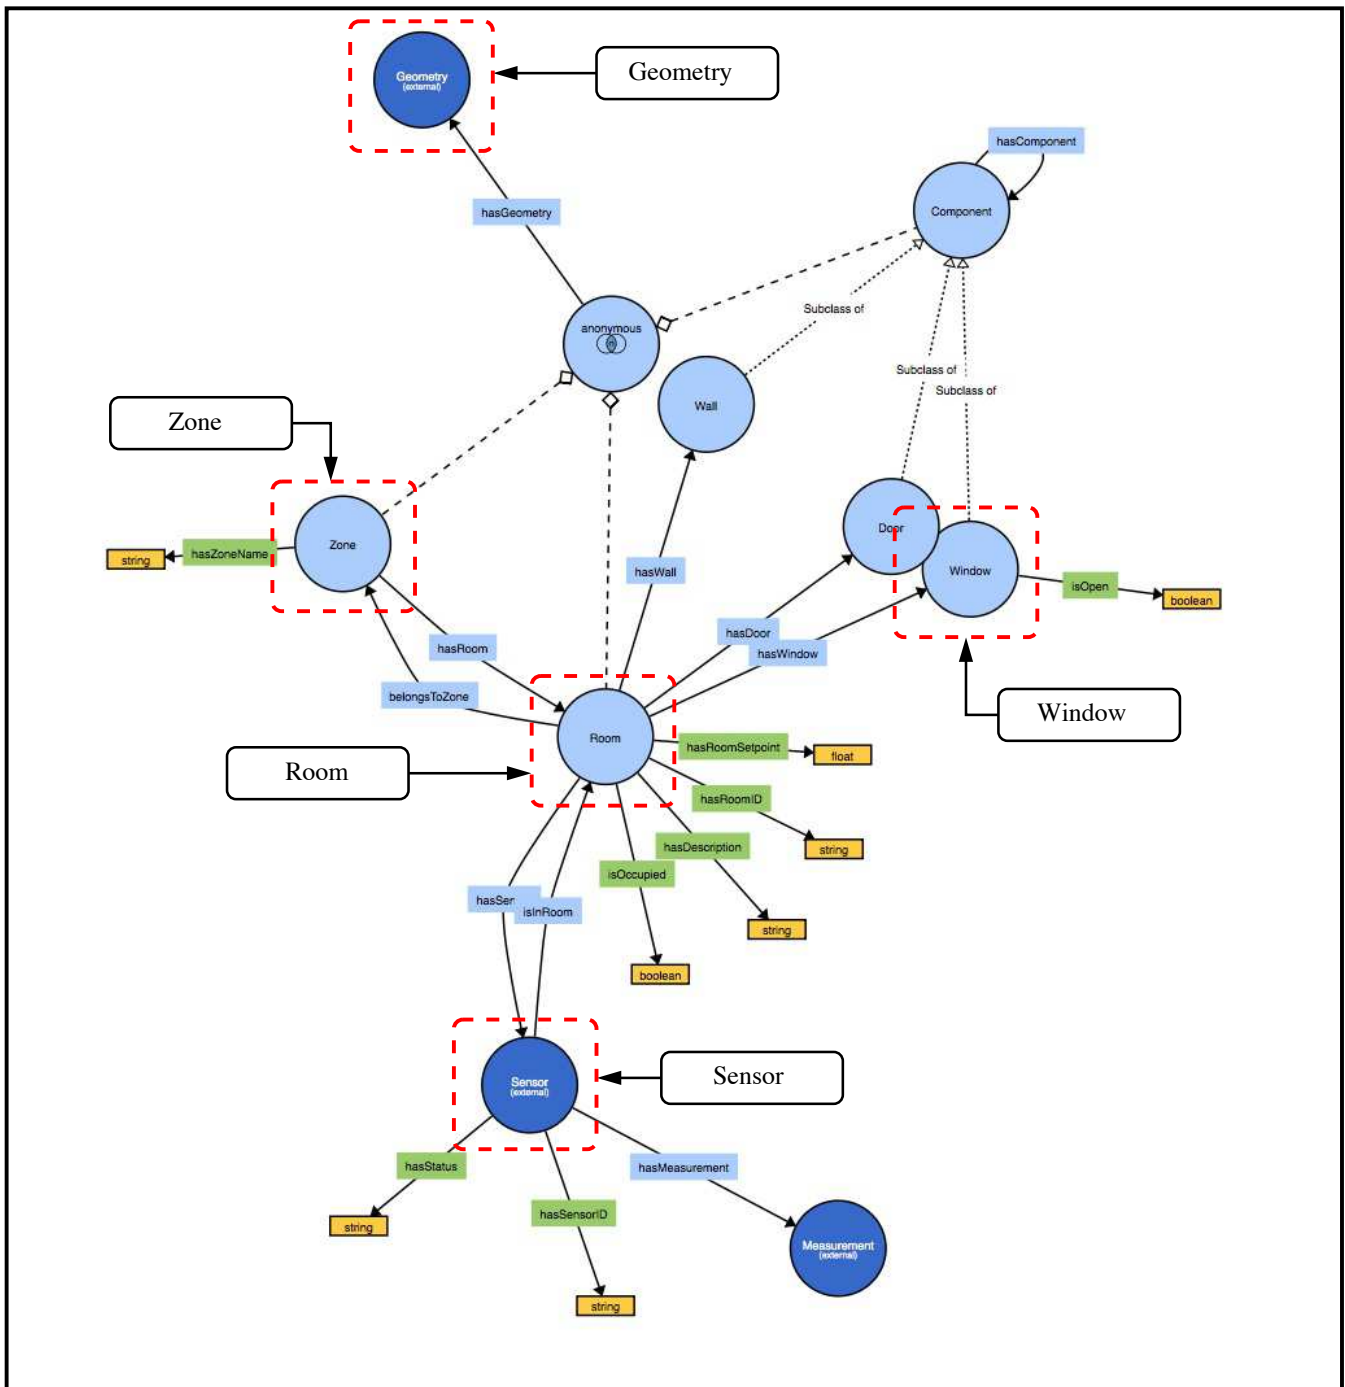


Figure 7. Schematic of building ontology classes and properties.

Jena Rules

```
// Rule to check if two zones intersect ...
[ BuildingRule02: (?r1 rdf:type bld:Zone) (?r1 bld:hasGeometry ?r1g) (?r1g geom:hasGeometry ?r1jts)
  (?r2 rdf:type bld:Zone) (?r2 bld:hasGeometry ?r2g) (?r2g geom:hasGeometry ?r2jts)
  notEqual( ?r1jts, ?r2jts ) getPointInPolygon( ?r1jts, ?r2jts, ?t)
  equal(?t, "true"^^xs:boolean) -> (?r1 bld:intersects ?r2)]
```

Figure 8. Rule for Zone Intersect.

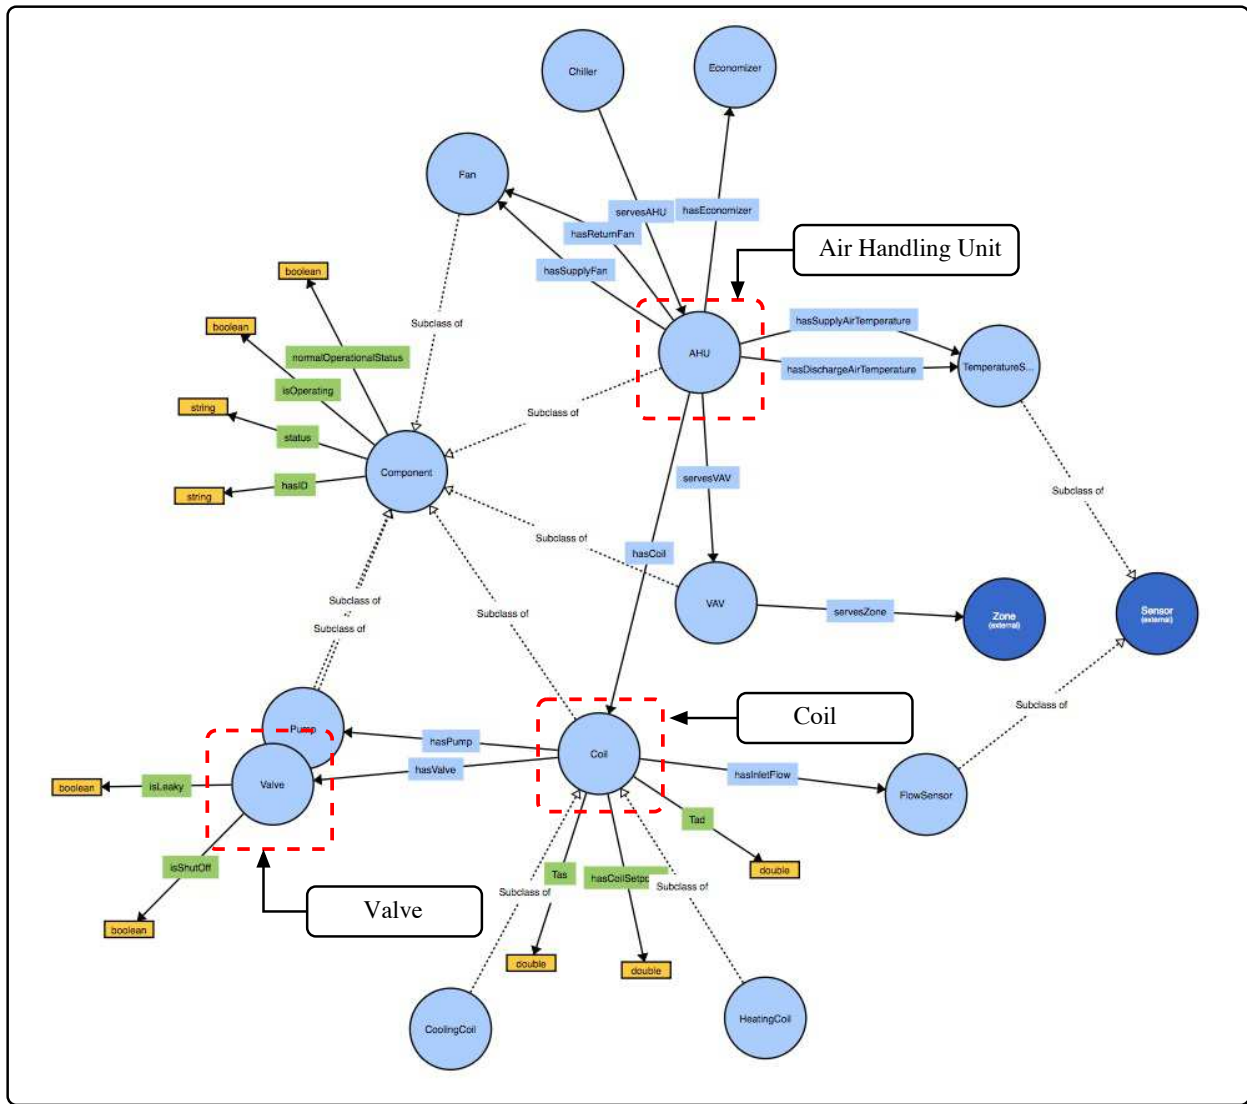


Figure 9. Schematic of equipment ontology classes and properties.

Jena Rules

```
// Close the valve when the coil temperature is the same as coil setpoint.
[ EquipmentRule01: (?coil rdf:type eq:Coil) (?coil eq:hasCoilSetpoint ?sp)
  (?coil eq:hasCoilTemperature ?cp) equal(?sp,?cp) (?coil eq:hasValve ?valve) ->
  (?valve eq:isShutOff "true"^^xs:boolean) print('valve is shut')]

// If the valve is shut, the temperature of the air that passes through the coil
// has to be the same. Otherwise, the valve is leaky
[ EquipmentRule02: (?hvw rdf:type eq:Valve) (?hvw eq:isShutOff "true"^^xs:boolean)
  (?c rdf:type eq:Coil)(?c eq:hasValve ?hvw) (?c eq:Tad ?t1)
  (?c eq:Tas ?t2) notEqual(?t2 ?t1) -> (?hvw eq:isLeaky "true"^^xs:boolean)
  (?hvw eq:hasNormalOperationalStatus "false"^^xs:boolean) print('valve is Leaky') ]

// If the a valve fails, the AHU fails too ...
[ EquipmentRule03: (?hvw rdf:type eq:Valve) (?AHU eq:hasCoil ?c) (?c eq:hasValve ?v)
  (?v eq:hasNormalOperationalStatus "false"^^xs:boolean) ->
  print('AHU malfunction') (?AHU eq:hasNormalOperationalStatus "false"^^xs:boolean)]
```

Figure 10. Rules for establishing the operational status and simple operations of mechanical equipment.

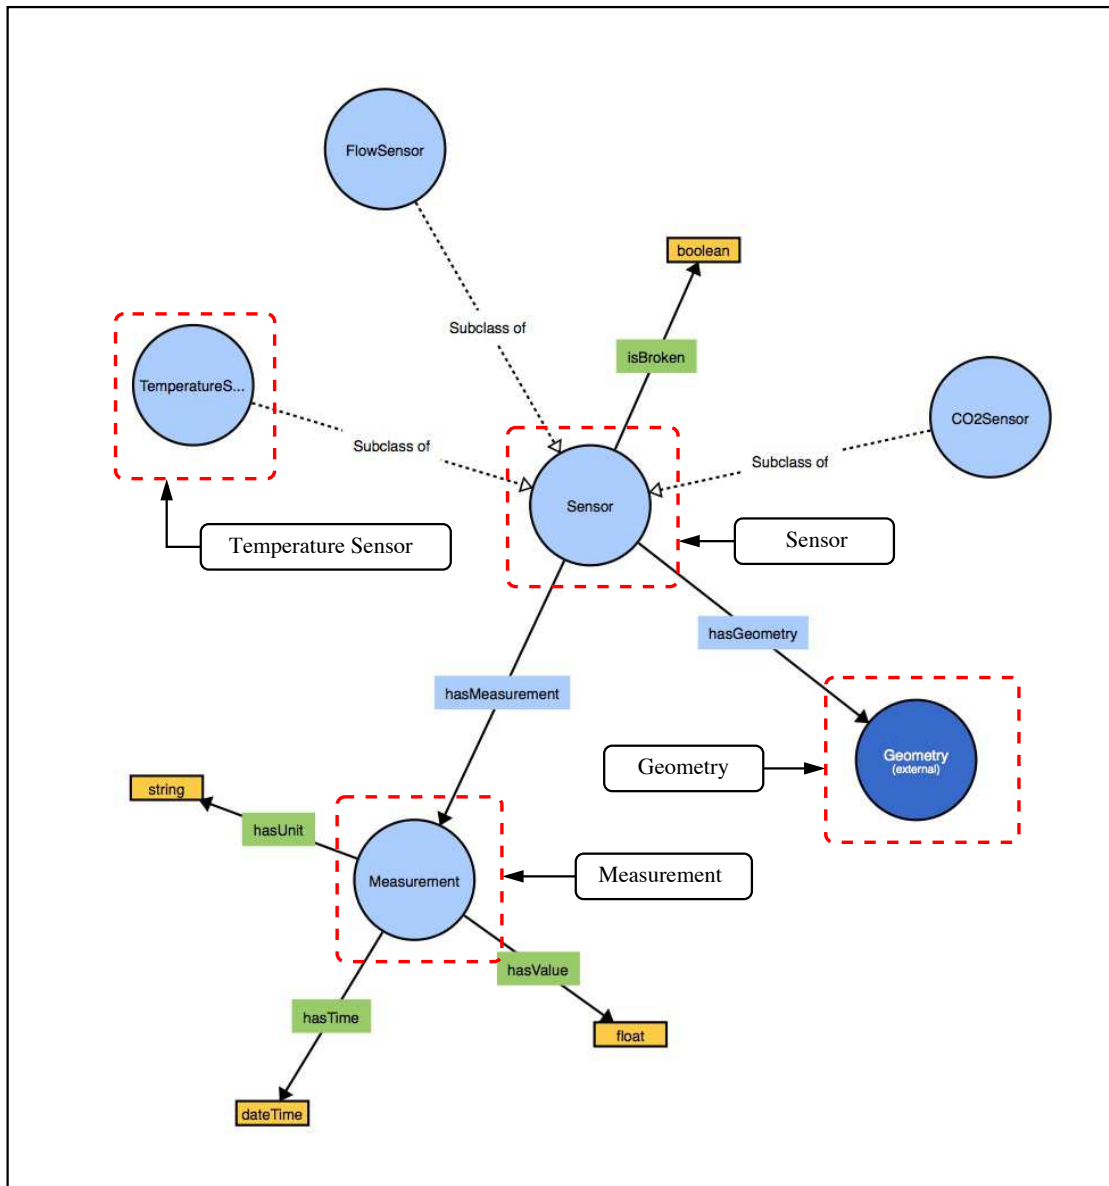


Figure 11. Sensor ontology classes and properties.

Jena Rules

```
// Simple rule to check if a sensor is broken ...
[ SensorRule01: (?s rdf:type sen:Sensor) (?s sen:hasMeasurement ?m) (?m sen:hasValue ?r)
  isOutOfRange(?m ?t) -> (?s sen:isBroken ?t) ]
```

Figure 12. Rule for Zone Intersect.

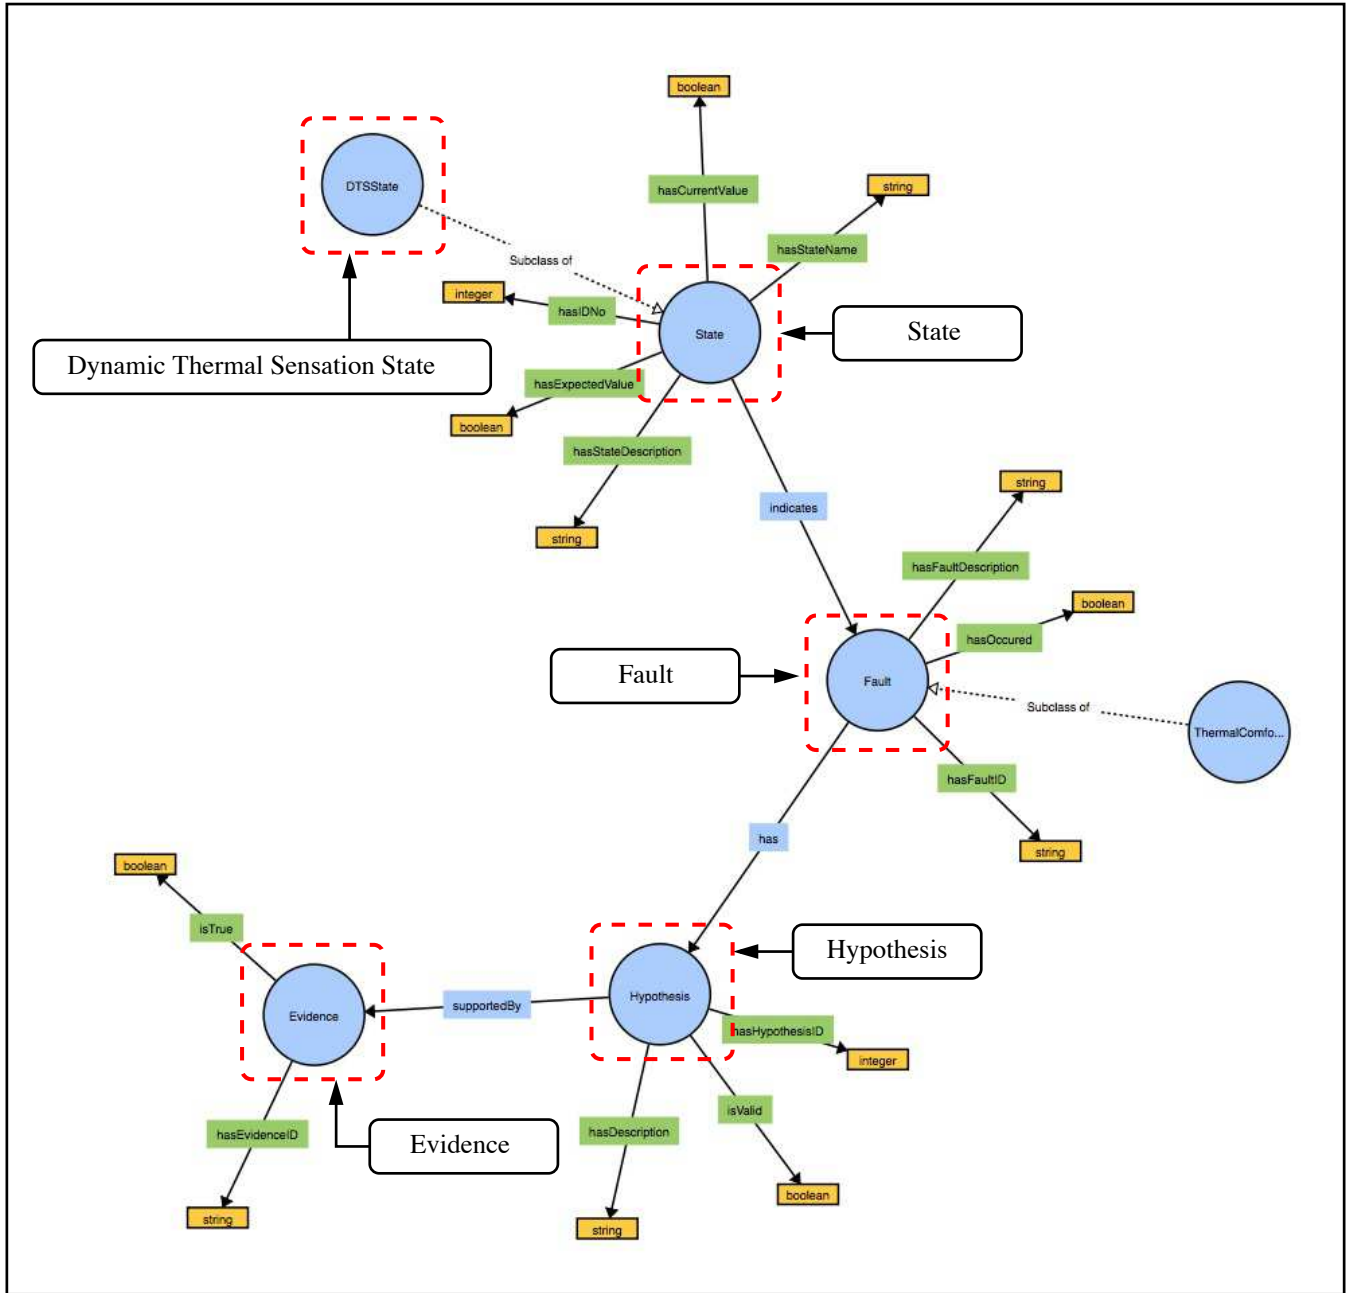


Figure 13. Fault detection and diagnostic ontology classes and properties.

Jena Rules

```
// General purpose rule for recording when a fault has occurred.
[FDDRule01: (?st rdf:type fdd:State) (?st fdd:hasCurrentValue ?csv)
(?st fdd:belongsToFault ?F) (?st fdd:hasExpectedValue ?esv)
notEqual(?csc,?esv) -> (?F fdd:hasOccured 'true') print('faultoccured')]
```

Figure 14. Rule for detecting a faulty state.

Jena Rules

```
// Determine room in which an occupant is located.

[ OccupantRule01: (?r rdf:type bld:Room) (?o rdf:type occ:Occupant)
  (?o occ:hasOccupantGeometry ?og) (?og geom:hasGeometry ?ojts)
  (?r bld:hasGeometry ?rg) (?rg geom:hasGeometry ?rjts)
  getPointInPolygon(?ojts,?rjts,?t) equal(?t, "true"^^xs:boolean) ->
  (?r bld:hasOccupant ?o) print(?o,'OccupantisInRoom',?r,?t)]

// When positive values of DTSSIndex are greater than 0.3, an occupant is not comfortable.

[ OccupantRule02: (?oc rdf:type occ:Occupant) (?oc occ:hasDTSSIndex ?v) greaterThan(?v,0.3)
  (?oc occ:hasDTSSState ?dts) -> print(?oc,'isComfortable' "false"^^xs:boolean)
  (?oc occ:isComfortable "false"^^xs:boolean)
  (?dts fdd:hasCurrentValue "false"^^xs:boolean)]
```

Figure 16. Rule for occupants location and thermal comfort.

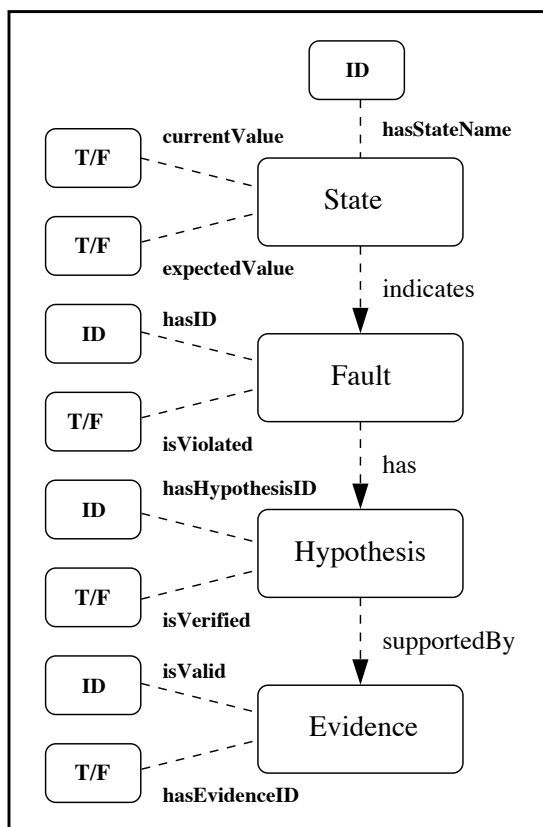


Figure 15. Flow chart for identification of faults and identification and verification of hypotheses and supporting evidence.

Our experimental FDD ontology also supports DTSSState, a subclass of State, designed to represent states associated with dynamic thermal sensation (DTS).

Figure 15 is a flowchart for fault detection and the identification and verification of relevant hypotheses and supporting evidence. The step-by-step procedure for detecting a fault and diagnosing its causes corresponds to a traversal through the classes State, Fault, Hypothesis and Evidence. A fault is indicated when the current and expected values of a state

are in conflict. Each fault has a hypothesis that needs to be supported by evidence. The evaluation procedure works backwards. Verification of the evidence is a prerequisite to validating a hypothesis. In an implementation of the procedure, data properties indicate whether or not a fault has been verified, whether or not an hypothesis has been verified, and whether or not supporting evidence is valid. This procedure is mirrored by set of rules shown in Figure 14.

B. Surrounding Environment Ontologies and Rules

The surrounding environment ontologies and rules include model support for the building occupants and weather phenomena.

Occupant Ontology and Rules. While several studies [28], [29] have recently identified the importance of including inhabitants as an integral part of simulation and control of energy systems and indoor environments, present-day procedures rely on predetermined occupancy schedules and/or empirical estimates based on sensors. For fault detection and diagnostic analysis of mechanical equipment in buildings, solutions are complicated by the strong coupling of human presence, comfort and behavior, to details of the building state (e.g., whether or not a window is open) and surrounding environment (e.g., what side of the building is in the sun).

Figure 16 takes a first step toward the development of rules for modeling and evaluation of occupant location and thermal comfort. The occupant ontology (see reference [23] for details) expands upon the work of Mahdavi and Taheri [30], and considers four subcategory problems: (1) location, (2) actions (e.g., open/close window), (3) attitudes (e.g., thermal sensation) and (4) preferences in terms of temperature and moisture of the air. Occupant location is modeled as point geometry in the building.

Weather Ontology and Rules. Based upon the work of Staroch [31], the weather ontology and rules cover concepts such as Weather Phenomenon, Weather Report, and Weather State. The weather state is composed of different Weather phenomenon class holds the physical attributes regarding the weather such as the temperature, pressure, solar radiation, wind

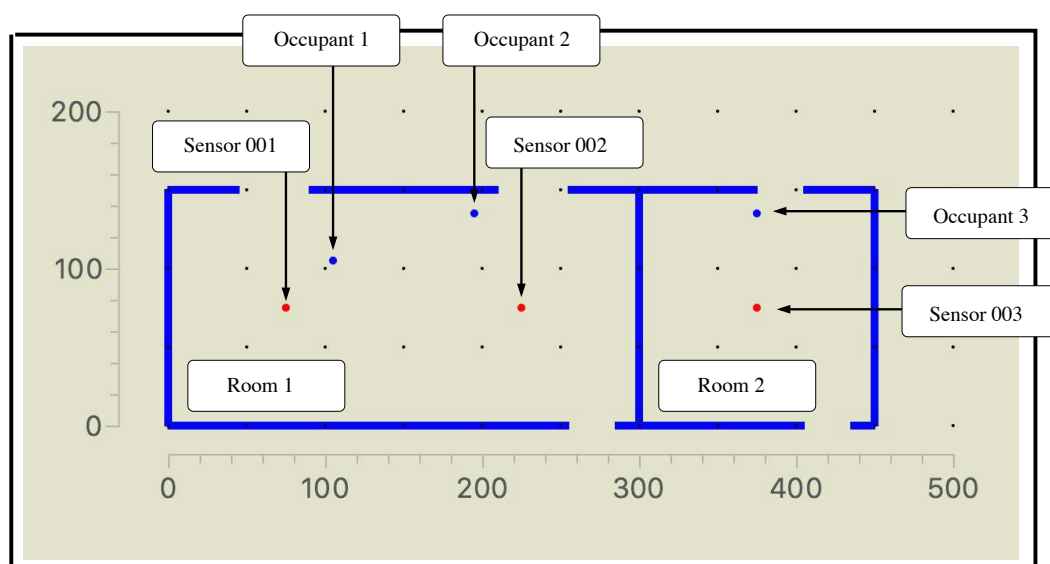


Figure 17. Plan view of two-room building architecture, sensors, and building occupants.

and cloud. Weather data is obtained from current Weather [32], a free and open source API (application programming interface) that provides access to historical as well as current and future forecast weather data from an online server. A Weather report can include data about the current weather or a forecast, specified in terms of start time and duration. For example, a medium range weather report has duration of more than 3 hours, with a start time of less than 12 hours into the future. Weather rules use current temperature values to identify a frosty and heat temperature conditions. For example, a Frost temperature condition occurs when observed temperature is below 0 C. A Heat temperature condition occurs when observed temperature is above 30 C. Similar intervals of temperature range can be defined for cold, below room temperature (at least 10 C and less than 20 C), and so forth.

VII. CASE STUDY PROBLEM

To examine capabilities of the framework for knowledge-based fault detection and diagnostic analysis, this section presents a case study test problem where faults in HVAC equipment are triggered by occupant discomfort in a conditioned space. The case study shows how heterogeneous data and knowledge from a variety of sources and domains can be integrated into a single semantic graph, how ontologies and rules can work together to detect the existence of a fault, and then diagnose the causes by systematically considering hypotheses and the supporting evidence.

A. Problem Description

Figure 17 is a plan view of the case study problem setup, consisting a small two-room building architecture, three sensors and three building occupants. Not shown is the mechanical equipment responsible for conditioning the room temperature and achieving acceptable levels of occupant comfort. The mechanical equipment consists of an air handling unit (AHU). The AHU has a coil (i.e., for heating and cooling). The water temperature that flows to the coil is managed by a valve.

Three rules are responsible for the operation and classification of faults in the mechanical equipment:

- Close the valve when the coil temperature is the same as coil setpoint.
- If the valve is shut, the temperature of the air that passes through the coil has to be the same. Otherwise, the valve is leaky
- If the a valve fails, the AHU fails too.

One measure to evaluate thermal comfort for the occupants is through computing the thermal sensation as a function of environmental factors such as outdoor and indoor temperature and some personal factors such as clothing levels. A dynamic model to compute thermal sensation (DTS) index to was introduced by Chen and co-workers [33]. According to thermal sensation scale suggested by ASHRAE [34], an acceptable range for occupancy comfort is the interval $[-0.3, 0.3]$. By comparing the current and expected values in a DTS state, the rules in Figure 14 will infer the existence of a faulty state, and then systematically examine the evidence associated with each hypothesis to find a root cause.

B. Snapshot of Semantic Graph Model Assembly

Figure 18 shows a snapshot of the building, equipment, sensor, weather, and FDD ontologies integrated together, and populated with system data. The semantic graph model contains instances of ontologies (individuals), relationships among individuals (often spanning domains), and data values associated with various individuals.

From a fault detection and diagnostics standpoint, the main points to note are as follows:

- Occupant 1 is located in Room 1.
- Room 1 has window, a temperature sensor (Sensor 001), and a carbon dioxide sensor (Sensor 002). HVAC services are provided to Room 1 by air handling unit

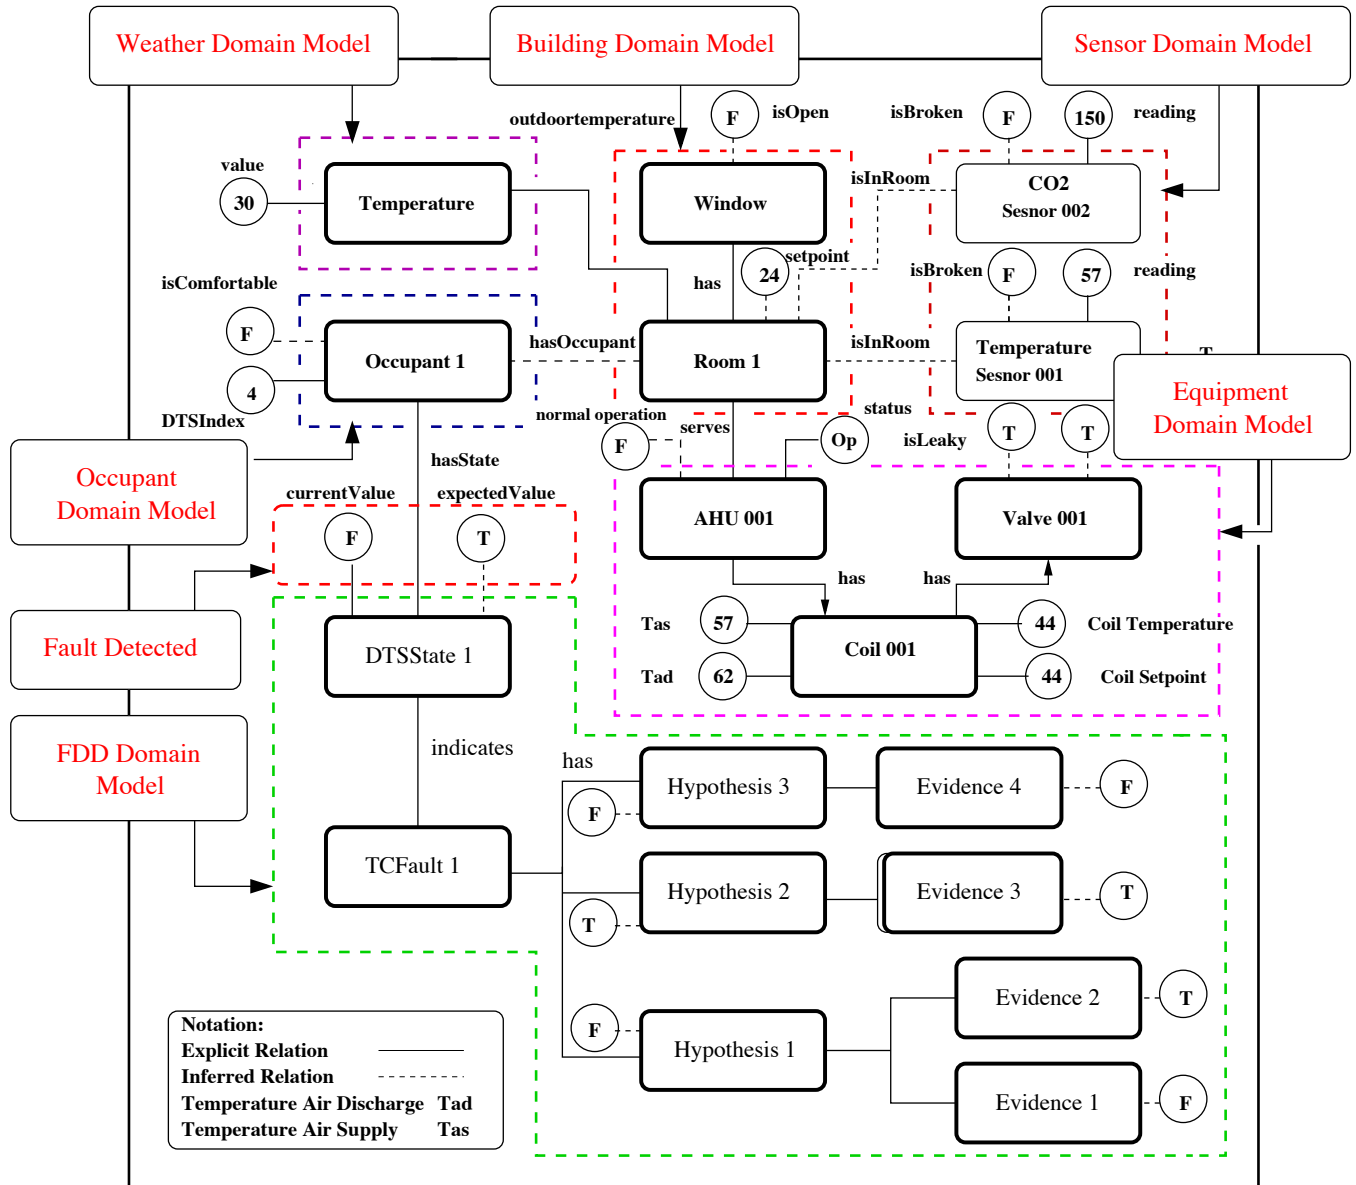


Figure 18. Snapshot of fully assembled semantic graph model. The data values will be computed and filled by the rules.

TABLE I. Instances of states, hypotheses, and evidence for identifying the cause for abnormal occupant thermal comfort value.

Class	Individual	Description
State	DTSSState 1	The DTS index in between $[-0.3, 0.3]$.
Fault	TCFault 1	The DTS index lies outside the interval $[-0.3, 0.3]$ when the air-handling unit is operating.
Evidence	Evidence 1	The CO2 sensor reading is above the normal range the and that shows the window is open.
	Evidence 2	The outdoor temperature is greater than room setpoint.
	Evidence 3	A sensor's reading is outside the range that indicates the sensor is broken.
	Evidence 4	A component is AHU is malfunctioning that results in an abnormal operation of AHU.
Hypothesis	Hypothesis 1	Warm outside air is leaking into the room through an open window -> Supported by Evidence 1 and Evidence 2.
	Hypothesis 2	The serving air-handling unit has abnormal operation. -> Supported by Evidence 3.
	Hypothesis 3	The room sensor that provides feed-back to AHU reaching its target setpoint is broken -> Supported by Evidence 4.

Jena Rules

```

// Evidence Rule 01: A window is open base on CO2 concentration in the room.
// -----
[ EvidenceRule01: (?cs rdf:type sen:CO2Sensor) (?cs bld:isInRoom ?room)
  (?r bld:hasWindow ?w)(?cs bld:hasReading ?m) lessThan(?m,600)
  greaterThan(?m,400) (?e fdd:hasEvidenceID ?n) equal("1"^^xs:integer,?n) ->
  (?w building:isOpen "true"^^xs:boolean) (?e fdd:isTrue "true"^^xs:boolean) ]

// Evidence Rule 02: Outside temperature is warmer than the setpoint.
// -----
[ EvidenceRule02: (?r rdf:type bld:Room) (?r bld:hasSetpoint ?sp)
  (?t rdf:type we:Temperature) (?t we:hasTemperatureValue ?tv)
  greaterThan(?tv,?sp) equal("2"^^xs:integer,?n) (?e rdf:type fdd:Evidence)
  (?e fdd:hasEvidenceID ?n) -> (?e fdd:isTrue "true"^^xs:boolean) ]

// Evidence Rule 03: Temperature sensor in a room is broken.
// -----
[EvidenceRule03: (?ts rdf:type sen:TemperatureSensor) (?ts bld:isInRoom ?room)
  (?ts bld:isBroken ?t) equal(?t, "true"^^xs:boolean) equal("3"^^xs:integer,?n)
  (?e rdf:type fdd:Evidence)
  (?e fdd:hasEvidenceID ?n ->(?e fdd:isTrue "true"^^xs:boolean) ]

// Evidence Rule 04: Malfunction is in the Air Handling Unit.
// -----
[EvidenceRule04: (?AHU rdf:type eq:AHU) (?v eq:hasNormalOperationalStatus "false"^^xs:boolean)
  equal(?t, "true"^^xs:boolean) equal("4"^^xs:integer,?n)
  (?e rdf:type fdd:Evidence)-> (?e fdd:isTrue "true"^^xs:boolean) ]

// FDD Rule 02: Indicate when thermal comfort in a conditioned room has expected value.
// -----
[FDDRule02: (?AHU rdf:type eq:AHU)(?AHU eq:servesRoom ?r)(?r bld:hasOccupant ?oc)
  (?oc occ:hasDTSSState ?dts) (?AHU eq:status ?s)
  equal(?s "Operating") -> print('Expected DTS',?oc)(?dts fdd:hasExpectedValue "true"^^xs:boolean)]

```

Figure 19. Fault detection diagnostic rules for operation of a heating coil and for checking evidence 3 and evidence 4.

AHU 001. AHU 001 has a coil (Coil 001); Coil 001 has a valve (Valve 001).

- The datatype property for AHU001 “normal Operation” is set to false. This setting is based on the system data and the result of equipment rules 01 through 03 being triggered.
- The setpoint temperature for Room 1 is 24 C, but the current temperature reading for Sensor 001 is 57 C.
- OccupantRule02 sets the “isComfortable” datatype property for Occupant1 to “false” as the result of a DTSSindex value of 4.
- Occupant 1 has dynamic thermal sensation (DTS) state DTSSState 1. DTSSState 1 indicates a thermal comfort fault (TCFault1), which will be diagnosed by looking at three hypotheses and their supporting evidence.
- The relationship between Hypotheses 1 through 3 and supporting evidence is shown along the bottom of Figure 18. Users may query the semantic graph to find the correct hypotheses and valid supporting evidence.

C. Test Problem Scenario and Hypothesis Evaluation Procedure

The test problem scenario assumes that the numerical value of occupant thermal comfort in a conditioned room has fallen outside the acceptable range. This is detected by FDD Rule 01. With this scenario in place, any one of three hypotheses could potentially be true. To correctly identify the correct hypothesis, the system requires to reason among the facts and identify the evidence existing in different domains,

- The outdoor temperature is higher than the setpoint (weather) and the window in the room is open (building, sensor, weather).
- The air-handling unit is malfunctioning (mechanical equipment),
- The room sensor providing feed-back to the air-handling unit to reach its target setpoint is broken (sensor).

As a result, this task will require comprehensive reasoning over multiple domains and identifying the supporting evidence to the most probable hypothesis. To achieve this, we used the

proposed framework and implemented ontologies for weather, building, occupant, sensor and equipment domains. The ontologies are populated with data. However, in general this data will be obtained from simulations or real buildings.

D. Synthesis of Multi-domain Rules

Table I describes the instances for key concepts of FDD ontology as they apply to the test case problem, and explains details of the individuals for FDD ontology. For the case study problem, the chain of dependency relationships between hypotheses and supporting evidence is as follows:

- Hypothesis 1 is that warm outside air is leaking into the room through an open window. Evaluation of this hypothesis is supported by execution of two evidence rules, EvidenceRule01 and EvidenceRule02.
- Hypothesis 2 is that the serving air-handling unit has abnormal operation. Evaluation of this hypothesis is supported execution of EvidenceRule03.
- Hypothesis 3 states that the room sensor that provides feedback to AHU reaching its target setpoint is broken. Supporting evidence is provided by the execution of EvidenceRule04.

Figure 19 presents the fault detection diagnostic rules for: (1) Operation of a heating coil, (2) Checking evidence 3 and evidence 4, and (3) Detecting when the thermal comfort in a conditioned room matches its expected value.

E. Multi-domain Rule Evaluation

Figure 20 shows a snapshot of multi-domain evaluation and forward chaining of rules. From an evaluation standpoint, the eight rules can be clustered into two pathways, the first focusing on fault detection and the second focusing on diagnostic investigation of probable causes, represented as hypotheses and supporting evidence.

Fault Detection: The first pathway identifies the existence of a fault and is covered by rules 1 through 4:

- Rule 01: Use OccupantRule01 (see Figure 16) to determine when an occupant is located in a room.
- Rule 02: Use FDDRule02 (see Figure 19) to determine the expected comfort of an occupant.
- Rule 03: Use OccupantRule02 (see Figure 16) to determine the current comfort of an occupant.
- Rule 04: Use OccupantRule02 (see Figure 16) to compute when a fault has occurred.

determine in which room an occupant is located and whether or not the current value of occupant comfort matches the expected value of comfort. In the snapshot, activation of Rule 01 determines that: Occupant1 is located in Room1. A separate execution would also determine that Occupant2 is also located in Room1. Activation of Rule 02 is based upon the output of Rule 01, state data from the building domain, the relationship of the air handling unit to Room1. In the snapshot trace, the output of Rule 02 states that DTSSState for Occupant1 is true

and that Occupant1 has a DTSSIndex of 4. A fault occurs when there is a discrepancy between the current and expected values of comfort (see F7 and F9), as indicated by the values of current and expected values of DTSSState.

Fault Diagnostics: By systematically examining hypotheses and supporting evidence, the second pathway diagnoses the causes of a fault. For the scenario outlined in Figure 20, this procedure is covered by rules 5 through 8:

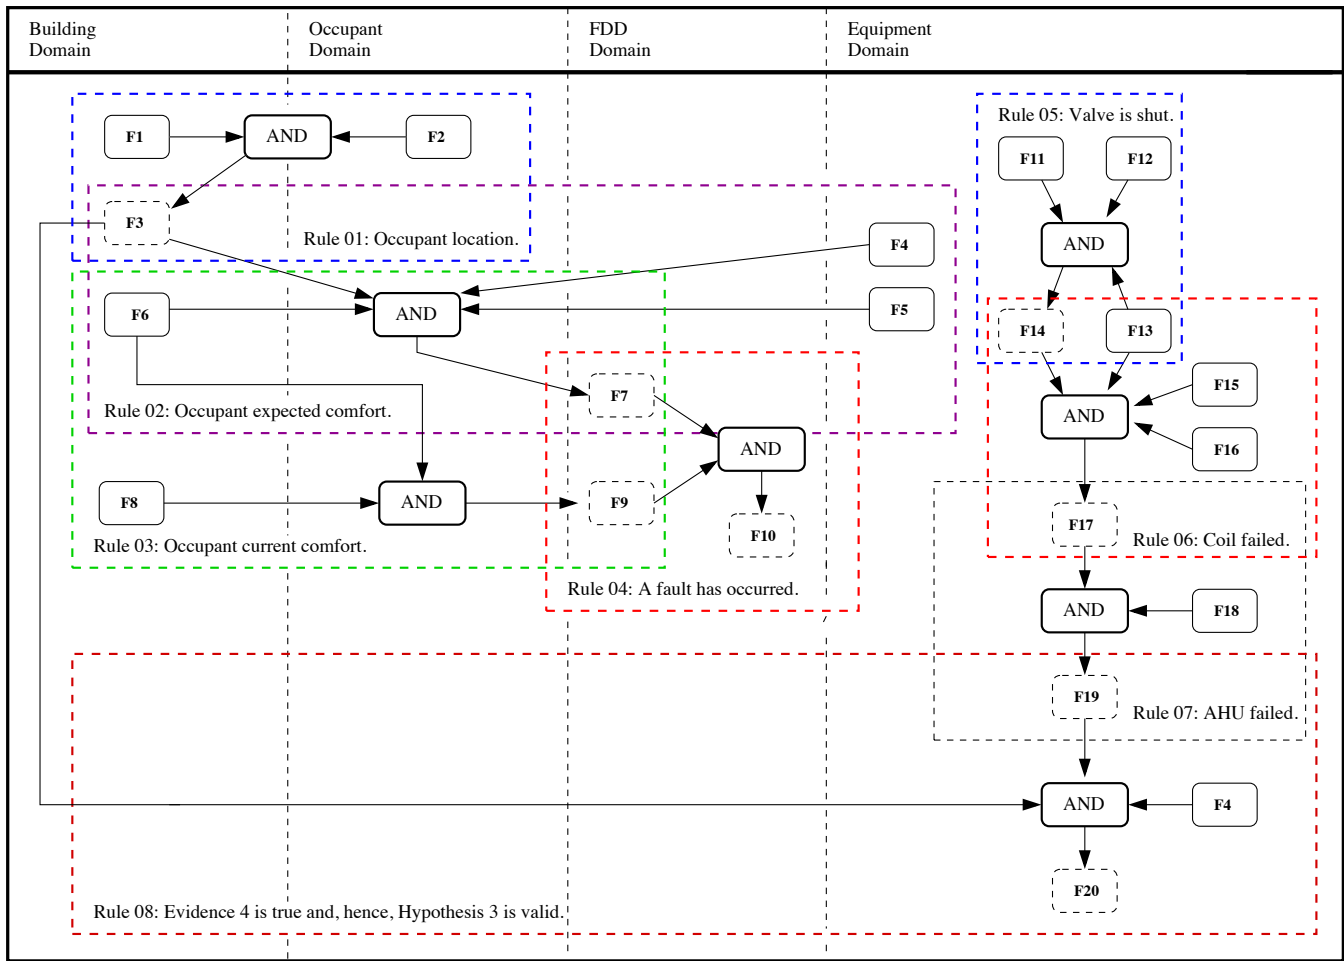
- Rule 05: Use EquipmentRule01 (see Figure 10) to determine if a valve is shut.
- Rule 06: Use EquipmentRule02 (see Figure 10) to determine if the coil has failed.
- Rule 07: Use EquipmentRule03 (see Figure 10) to determine whether or not the air handling unit has failed.
- Rule 08: If EvidenceRule04 (see Figure 19) evaluates to true then Hypothesis 3 is true.

The rule for determining whether or not the valve is shut takes input values from the Coil001 CoilSetpoint (44) and CoilTemperature (44) (see F12 and F13), and checks to verify that the coil has a valve. In our scenario, the rule output (F14) is true, indicating that Valve001 is shut, and hence in Rule 06 normal operation evaluates to false. A simple check to verify that the coil belongs to air handling unit AHU001 generates the conclusion that normal operation of the AHU is false (see F19). Finally, input from the room occupancy test and a test to verify that AHU001 is connected to Room1, leads to the conclusion Evidence 4 is supported and Hypothesis 3 is valid. Finally, we note that except for the room occupancy information feeding into Rule 08, the fault detection and diagnosis pathways operate independently.

VIII. CONCLUSIONS AND FUTURE WORK

We have proposed in this paper a knowledge-based framework for fault detection and diagnostics. The underlying process closely mimics the “thinking process” that humans follow in identifying and diagnosing the causes of a fault. Thus, the steps of gathering data for the participating domains, populating ontologies with individuals, and using rules to detect and diagnose faults and their causes is easy for humans to understand and generally applicable to other domains (e.g., building energy, automotive, health care) for FDD purposes. Capabilities of the prototype implementation has been demonstrated by working step by step through the procedure of detecting and diagnosing the source of faults in an HVAC system.

Key advantages of this approach include: (1) it is decoupled from the system simulation, (2) it is comprehensive, and (3) it is scalable. In fact, the process for expanding an application to include new domains as they come along is very straight forward. The inference-based rules are guaranteed to check at anytime a changed occurred in an ontology resulting in event-driven fault detection and diagnostic. Finally, inference-based rules provide mechanisms in capturing chain effects that exists in the nature of system failure – for example, if a valve is not



Legend:

- | | | |
|-----------------------------------|------------------------------------|---|
| F1 = Room hasGeometry | F8 = Occupant1 hasDTSIndex 4 | F15 = Coil001 Tas 62 |
| F2 = Occupant hasGeometry | F9 = DTSSstate currentValue false | F16 = Coil001 Tad 57 |
| F3 = Room1 hasOccupant1 | F10 = DTSSstate indicates DTSFault | F17 = Valve001 isShut normalOperation false |
| F4 = AHU001 serves Room1 | F11 = Coil001 CoilSetpoint 44 | F18 = AHU001 hasCoil Coil001 |
| F5 = AHU001 status operating | F12 = Coil001 CoilTemperature 44 | F19 = AHU001 normalOperation false |
| F6 = Occupant1 hasState DTSSstate | F13 = Coil001 hasValve Valve001 | F20 = Evidence1 isValid true |
| F7 = DTSSstate expectedValue true | F14 = Valve001 isShut true | |

Figure 20. Snapshot of multi-domain evaluation and forward chaining of rules.

operational, the evidence that AHU is not operating properly also holds true.

In our prototype implementation, the small two-room building model extracted data from a custom “system data model” currently under development. We expect that a more mature version of this ontology would extract semantic information from instances of building information models (BIM) such as the Industry Foundation Class (IFC). Future work will also include deployment in real building systems. We anticipate that the proposed methodology will be integrated into building automation systems (BAS) and support investigations where analytic built-in functions are implemented in the condition part of inference-based rules. These functions will perform time-history analyses to identify a faulty state for the system. We anticipate a trend where formal approaches to analysis are

used to irregularities in building performance, which are indicators of possible system faults. Moreover, we will investigate strategies for taking control actions based on recognized faults of the system.

IX. ACKNOWLEDGMENT

The first author was supported by a fellowship award from the NIST Graduate Student Measurement Science and Engineering (GMSE) Program.

REFERENCES

[1] P. Delgoshaei, M.A. Austin, and D. Veronica, “Semantic Models and Rule-based Reasoning for Fault Detection and Diagnostics: Applications in Heating, Ventilating and Air Conditioning Systems,” The Twelfth

- International Conference on Systems (ICONS 2017), April 23-27 2017, pp. 48–53.
- [2] P. Delgoshaei, M. A. Austin and A. Pertzborn, "A Semantic Framework for Modeling and Simulation of Cyber-Physical Systems," *International Journal On Advances in Systems and Measurements*, vol. 7, no. 3-4, December 2014, pp. 223–238.
 - [3] M.A. Austin, P. Delgoshaei and A. Nguyen, "Distributed System Behavior Modeling with Ontologies, Rules, and Message Passing Mechanisms," *Procedia Computer Science*, vol. 44, 2015, pp. 373 – 382, 2015 Conference on Systems Engineering Research. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1877050915002951>
 - [4] S. Katipamula and M. R. Brambley, "Review Article: Methods for Fault Detection, Diagnostics, and Prognostics for Building SystemsA Review, Part I," *HVAC&R Research*, vol. 11, no. 1, 2005, pp. 3–25.
 - [5] J. A. Siegel and C. P. Wray, "An Evaluation of Superheat-based Refrigerant Charge Diagnostics for Residential Cooling Systems/Discussions," *ASHRAE Transactions* 108(1), 2002, p. 965.
 - [6] W. Kim and J. E. Braun, "Impacts of refrigerant charge on air conditioner and heat pump performance," in *Impacts of Refrigerant Charge on Air Conditioner and Heat Pump Performance*, July 10–15 2010, pp. 2433–2441.
 - [7] M. Wiggins and J. Brodrick, "Emerging Technologies: HVAC Fault Detection," *ASHRAE Journal*, April 2012, pp. 78–80.
 - [8] OWL., "Web Ontology Language Overview, W3C Recommendation from February, 2004. For details, see <http://www.w3.org/TR/owl-features/> (Accessed, April 2017)."
 - [9] Apache Jena, "An Open Source Java Framework for building Semantic Web and Linked Data Applications, Accessible at <https://jena.apache.org> (Accessed on 12/12/16)," 2016.
 - [10] W. T. Scherer and C. C. White, *A Survey of Expert Systems for Equipment Maintenance and Diagnostics*. Boston, MA: Springer US, 1989, pp. 285–300.
 - [11] T. Berners-Lee, J. Hendler and O. Lassila, "The Semantic Web," *Scientific American*, May 2001, pp. 35–43.
 - [12] T. Q. Dung and W. Kameyama, *Ontology-based Information Extraction and Information Retrieval in Health Care Domain*, ser. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2007, vol. 4654 LNCS, pp. 323–333.
 - [13] C. Taswell, "DOORS to the Semantic Web and Grid with a PORTAL for Biomedical Computing," *IEEE Trans Inf Technol Biomed*, vol. 12, no. 2, 2008, pp. 191–204.
 - [14] P. Lord, S. Bechhofer, M. D. Wilkinson, G. Schiltz, D. Gessler, D. Hull, C. Goble, and L. Stein, *Applying Semantic Web Services to Bioinformatics: Experiences Gained, Lessons Learnt*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 350–364.
 - [15] D. Corsar, D. Milan, P. Edwards, and J. D. Nelson, *The Transport Disruption Ontology*. *Lecture Notes in Computer Science*, vol 9367, Springer, 2015, pp. 329–336.
 - [16] M. Batic, N. Tomasevic and S. Vranes, "Ontology-based Fault Detection and Diagnosis System Querying and Reasoning Examples," in *ICK-DDM 2015 : 17th International Conference on Knowledge Discovery and Data Mining*, vol. 2, no. 1. International Science Index, Industrial and Manufacturing Engineering, 2015.
 - [17] S. Schumann, J. Hayes, P. Pompey and O. Verscheure, "Adaptable Fault Identification for Smart Buildings," in *2011 AAAI Workshop (WS-11-07)*, 2011.
 - [18] M. Merdan, "Knowledge-based Multi-Agent Architecture Applied in the Assembly Domain," Ph.D. Dissertation, Vienna University of Technology, 2009.
 - [19] T. Bayer, D. Dvorak, S. Friedenthal, S. Jenkins, C. Lin and S. Mandutianu, "Foundational Concepts for Building System Models," in *SEWG MBSE Training Module 3*, see <http://nen.nasa.gov/web/se/mbse/documents>, California Institute of Technology, CA, USA, 2012.
 - [20] D.A. Wagner, M.B. Bennett R. Karban, N. Rouquette, S. Jenkins and M.o Ingham, "An Ontology for State Analysis: Formalizing the Mapping to SysML," in *Proceedings of 2012 IEEE Aerospace Conference*, Big Sky, Montana, March 2012.
 - [21] D.A. Randell, Z. Cui, and A.G. Cohn, "A Spatial Logic based on Regions and Connectivity," 1994, Division of Artificial Intelligence, School of Computer Studies, Leeds University.
 - [22] Java Topology Suite (JTS). See <http://www.vividsolutions.com/jts/> (Accessed August 4, 2017).
 - [23] P. Delgoshaei, and M.A. Austin, "Framework for Knowledge-Based Fault Detection and Diagnostics in Multi-Domain Systems: Application to HVAC Systems," *Institute for Systems Research, University of Maryland, College Park, MD 20742, USA*, Tech. Rep. 2017-4, November 2017.
 - [24] D.B. Crawley, L.K. Lawrie, F.C. Winkelmann, W.F. Buhl, Y.J. Huang, C.O. Pedersen, R.K. Strand, R.J. Liesen, D.E. Fisher, M.J. Witte and J. Glazer, "EnergyPlus: Creating a New-Generation Building Energy Simulation Program," *Energy and Buildings*, vol. 33, no. 4, 2001, pp. 319 – 331, Special Issue: {BUILDING} SIMULATION'99.
 - [25] S.A. Klein, W.A. Beckman, et al., 1994, *TRNSYS: A Transient Simulation Program*, Engineering Experiment Station Report 38-12, University of Wisconsin, Madison.
 - [26] "TRNSYS: The Transient Energy System Simulation Tool. See: <http://www.trnsys.com/> (Accessed September 8, 2017)." 2017.
 - [27] 2004, *BACnet: A Data Communication Protocol for Building Automation and Control Networks*, ANSI/ASHRAE 135.
 - [28] Y. Agarwal, B. Balaji, R. Gupta, J. Lyles, M. Wei and T. Weng, "Occupancy-Driven Energy Management for Smart Building Automation," in *Proceedings of the 2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building (BuildSys 2010)*, Zurich, Switzerland, November 3-5 2010, pp. 1–6.
 - [29] J. Lu, T.Sookoor, V. Srinivasan, G. Gao, B. Holben, J. Stankovic, E. Field and K. Whitehouse, "The Smart Thermostat: Using Occupancy Sensors to Save Energy in Homes," in *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems (SenSys 2010)*, Zurich, Switzerland, November 3-5 2010, pp. 211–224.
 - [30] A. Mahdavi and M. Taheri, "An Ontology for Building Monitoring," *Journal of Building Performance Simulation*, October 2016, pp. 1–10.
 - [31] P. Staroch, "A Weather Ontology for Predictive Control in Smart Homes," 2013, M.S. Thesis in Software Engineering and Internet Computing, Vienna University of Technology.
 - [32] Weather API. See <https://openweathermap.org/api> (Accessed September 14, 2017).
 - [33] X. Chen, and Q. Wang, and J. Srebric, "Occupant Feedback-based Model Predictive Control for Thermal Comfort and Energy Optimization: A Chamber Experimental Evaluation," *Applied Energy*, vol. 164, 2016, pp. 341 – 351.
 - [34] 2010, *ASHRAE Standard 552010 Thermal Environmental Conditions for Human Occupancy*, American Society of Heating, Refrigerating and Air-Conditioning Engineers.