Technical Report

Department of Computer Science
and Engineering
University of Minnesota
4-192 EECS Building
200 Union Street SE
Minneapolis, MN 55455-0159 USA

TR 00-003

Requirements Capture and Evaluation in NIMBUS: The Light-Control Case Study

Jeffrey M. Thompson, Michael W. Whalen, and Mats P. Heimdahl

February 01, 2000

Requirements Capture and Evaluation in NIMBUS: The Light-Control Case Study [1]

Jeffrey M. Thompson, Michael W. Whalen, Mats P.E. Heimdahl
University of Minnesota
Department of Computer Science and Engineering
Minneapolis, MN 55455
{thompson,whalen,heimdahl}@cs.umn.edu

Abstract: Evaluations of methods and tools applied to a reference problem are useful when comparing various techniques. In this paper, we present a solution to the challenge of capturing the requirements for the Light Control System case study, which was proposed before the Dagstuhl Seminar on *Requirements Capture, Documentation, and Validation* in June of 1999.

The paper focuses primarily on *how* the requirements were specified: what techniques were used, and what the results were. The language used to capture the requirements is RSML^{-e}; a state-based specification language with a fully specified formal denotational semantics. In addition, the NIMBUS environment – a toolset supporting RSML^{-e}– is used to visualize and execute the high-level requirements.

Keywords: light control system, specification-based prototyping, formal requirements modeling, state-based specification languages, requirements execution and simulation.

Category: D2.1 Requirements Specifications – Languages

1 Introduction

In this paper we present a solution to the challenge of capturing the requirements for the Light Control System case study set forth before the Dagstuhl Seminar on *Requirements Capture, Documentation, and Validation* in June of 1999[2].

The solution in this paper is captured in a fully formal, executable, state-based modeling language called RSML^{-e} (Requirements State Machine Language without events). We will demonstrate how the language is used to capture the *system requirements* for the Light Control System and how our *requirements engineering environment*, NIMBUS (based around RSML^{-e}), is used to dynamically validate and evaluate the system requirements. Furthermore, we will show how RSML^{-e} is used to refine the system requirements to *software requirements* and how NIMBUS can assist in this process—an approach we call *specification-based prototyping* [Thompson *et al.*, 1999].

The paper's focus is on *how* the requirments specification effort was completed and what was learned about the original informal requirements specification in the process. Our goal with this report is to give the reader a basic understanding of the capabilities of RSML^{-e} and the NIMBUS environment. The

[1] This work has been partially supported by NSF grants CCR-9624324 and CCR-9615088, and University of Minnesota Grant in Aid of Research 1003-521-5965.

[2] The case study is available at http://www.rn.informatik.uni-kl.de/~ecs.

completed formal requirements model and the models of the environment (sensors, actuators, and process) are too lengthy to include in this report. The full specification is nevertheless available for review on-line[3].

In the next section we will present our general view of control systems of the type presented in the Light Control System case study. We will discuss what system and software requirements are as well as a proposed structuring of the requirements models. In Section 3 we present the high-level structure of our view of the Light Control System, the system scope, and its boundaries. Sections 4 and 5 present the requirements of the Light Control System in RSML^{-e} and demonstrate how the requirements can be dynamically evaluated in the NIMBUS environment. The process of refining the system requirements to software requirements, and how to dynamically evaluate the refinement steps in NIMBUS, is discussed in Section 6. Section 7 contains an evaluation of the project and our approach as applied to the Light Control System. In this section we also provide some recommendations and discuss directions for future research.

2 The General Modeling Approach

The primary application domain for RSML^{-e} and the NIMBUS environment is safety critical applications; that is, applications where malfunction of the software may lead to death, injury, or environmental damage. Most, if not all, such systems are some form of a process control system where the software is participating in the control of a physical system. In this section we will provide a general overview of our modeling approach and point out what information needs to be captured in the system requirements specification as well as the software requirements specification.

2.1 Control Systems

A general view of a software controlled system can be seen in Figure 1. This model consists of a process, sensors, actuators, and a software controller. The process is the physical process we are attempting to control. The sensors measure physical quantities in the process. These measurements are provided as input to the software controller. The controller makes decisions on what actions are needed and commands the actuators to manipulate the process. The goal of the software control is to maintain some properties in the physical process. Thus, understanding how the sensors, actuators, and process behave is essential for the development and evaluation of correct software. The importance of this systems view has been repeatedly pointed out in the literature [Parnas and Madey, 1991, Leveson *et al.*, 1994, Heitmeyer *et al.*, 1995].

To reason about this type of software controlled systems, David Parnas and Jan Madey defined what they call the four-variable model (outside square of Figure 2) [Parnas and Madey, 1991]. In this model, the monitored variables (MON) are physical quantities we measure in the system and controlled variables (CON) are physical quantities we will control. The requirements on the control system are expressed as a mapping (REQ) from monitored to controlled variables. For instance, a requirement may be that *"when a room is occupied, there must be safe illumination."* Naturally, to implement the control software we must have

[3] Note to reviewers: If accepted we will insert a hypertext reference to a web site containing the entire system and software requirements specifications.

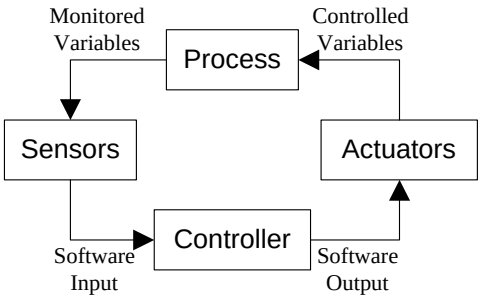

Figure 1: Traditional feedback process control model

sensors providing the software with measured values of the monitored variables (INPUT), for example, an indication if there is a person in the room. The sensors transform MON to INPUT through the IN relation; thus, the IN relation defines the sensor functions. To adjust the controlled variables, the software generates output that activates various actuators that can manipulate the physical process, for instance, a means to vary the illumination level in the room. The actuator function OUT maps OUTPUT to CON. The behavior of the software controller is defined by the SOFT relation that maps INPUT to OUTPUT.

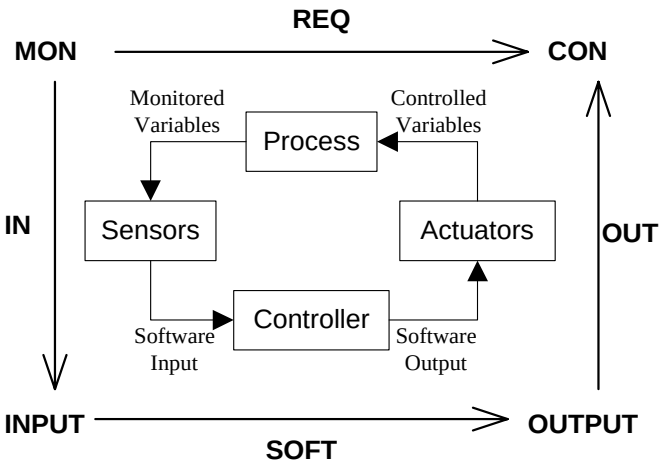

Figure 2: The four variable model for process control systems

The requirements on the control system are expressed with the REQ relation; the system requirements shall always be expressed in terms of quantities in the physical world. To develop the control software, however, we are interested in the SOFT relation. Thus, we must somehow refine the *system requirements* (the REQ relation) into the *software specification* (the SOFT relation).

2.2 Structuring SOFT

The IN and OUT relations are determined by the sensors and actuators used in the system. For example, to measure the light level in a room we may use a photo resistor coupled with an A/D converter that provides us an estimate of the light level as an integer. Similarly, to control the light level we may use dimmers and the light fixtures in the room. Armed with the REQ relation (mapping MON to CON), the IN relation (mapping IN to INPUT), and the OUT relation (mapping OUTPUT to CON) we can derive the SOFT relation. The question is, how shall we do this and how shall we structure the description of the SOFT relation in a language such as RSML^{-e}?

As mentioned above, the system requirements should always be expressed in terms of the physical process. These requirements will most likely change over the lifetime of the controller (or family of similar controllers). The sensors and actuators are likely to change independently of the requirements as new hardware becomes available or the software is used in subtly different operating environments; thus, all three relations, REQ, IN, and OUT, are likely to change over time. If either one of the REQ, IN, or OUT relations change, the SOFT relation must be modified. To provide a smooth transition from system requirements (REQ) to software requirements (SOFT) and to isolate the impact of requirements, sensor, and actuator changes to a minimum, Steven Miller at Rockwell Collins has proposed to structure the software specification SOFT based heavily on the structure of the REQ relation [Miller, 1999].

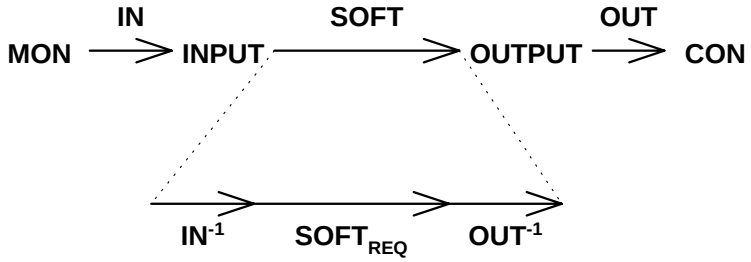

Figure 3: The SOFT relation can be split into three composed relations. The SOFT$_{REQ}$ relation is based on the original requirements (REQ) relation.

Miller proposed to achieve this by splitting the SOFT relation into three pieces, IN^{-1}, OUT^{-1}, and SOFT$_{REQ}$(Figure 3). IN^{-1} takes the measured input and reconstructs an estimate of the physical quantities in MON. The OUT^{-1} relation maps the internal representation of the controlled variables to the output needed for the actuators to manipulate the actual controlled variables. Given the IN^{-1} and OUT^{-1} relations, the SOFT$_{REQ}$ relation will now be essentially isomorphic to the REQ relation and, thus, be robust in the face of likely changes to the IN and OUT relations (sensor and actuator changes). Such changes would only effect the IN^{-1} and OUT^{-1} portions of the software requirements specification. Thus, the structuring approach outlined in this section will reduce the impact of likely changes on the software requirements specification SOFT.

In the rest of this paper we will illustrate how this framework for requirements specification and requirements refinement is used. We will demonstrate how the REQ relation is captured in a language called RSML^{-e} and how it can be validated through execution and simulation in the requirements engineering environment NIMBUS. We will also demonstrate how the REQ relation is refined to the SOFT relation and how the NIMBUS environment supports dynamic evaluation of the various models created throughout the refinement process. The result of this process will be a formal specification of SOFT that, in NIMBUS, also serves as a prototype of the control software.

3 The Light Control System

The informal requirements for the Light Control system were provided by the problem description included in the call for papers. This description provided us with a rough idea of the functionality of the system, but left many areas of the system underspecified, or, more seriously, specified in a way that contradicted other requirements in the system. What follows is our interpretation of the requirements of the system, with a few changes and additions to provide a complete and consistent description of the system. First, we will clarify the physical structure of the components of the system. Then, we can set the system boundaries and identify the monitored (MON) and controlled (CON) variables.

3.1 System Structure

The physical structure of the system and the control software boundaries are outlined in Figure 4. The control software, in the center of the diagram, is specified in RSML^{-e}. However, since RSML^{-e} was defined primarily for reactive control systems, the language does not support complex datatypes and data processing; thus, we have omitted some of the reporting features specified problem description.

The Light Control System can be thought of as a set of small control systems, each of which control the light level for a given room or hallway. Because the behaviors of the rooms and hallways are independent of one another, we can model each system separately and then "wire together" the control systems into a complete control system for a floor or a building. In RSML^{-e} we view a *system* as a collection of *components* connected by communication *channels*. A graphical representation (RSML^{-e} notation) of a collection of system components and communication channels can be seen in Figure 4. The components are connected to the channels through *interfaces* and can send *messages* over the channels. A message is a collection of *fields* holding the atomic pieces of information communicated between the components. The only information flow between the components is through the unidirectional channels.

Dividing up the specifications in this way allows us to focus on small parts of the system and makes the task of analyzing these pieces simpler. Nevertheless, the hallway and room specifications are quite similar. In general, the hallway can be viewed as providing a subset of the functionality provided by the rooms. For this reason, as well as space concerns, the remainder of this paper will focus on the specification of the system and software requirements for the rooms only (the full specifications for rooms and hallways are, of course, available on-line).

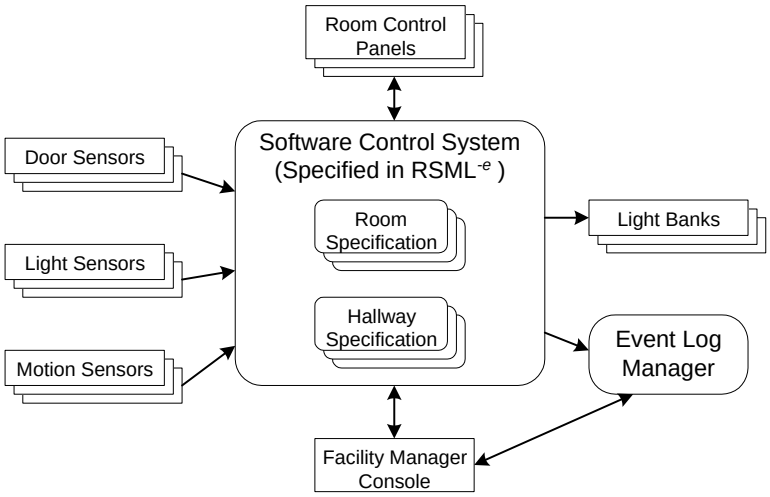

Figure 4: The software control system boundaries for the Light Control System.

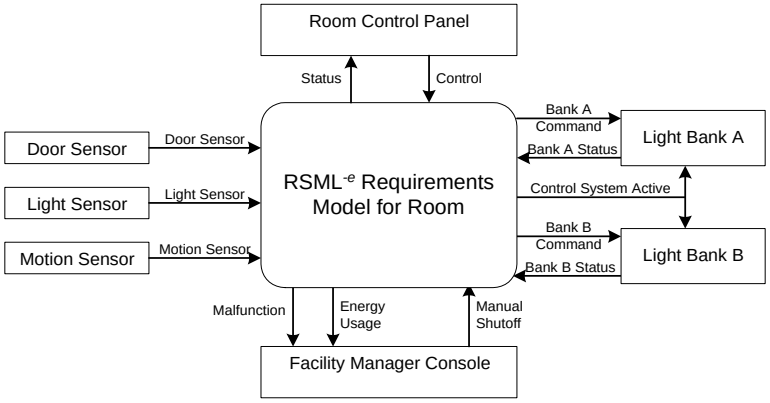

Figure 5: The RSML^{-e} room model and its interconnections.

3.2 Monitored and Controlled Variables

The first step in a requirements modeling project is to define the system boundaries and identify the monitored and controlled variables in the environment. In this paper we will not go into the details of how to scope the system requirements and identify the monitored and controlled variables—guidelines to help identify monitored and controlled variables are covered in, for example, [Miller, 1999, Faulk *et al.*, 1992, Jackson, 1995]. Here it suffices to say that the monitored and controlled variables exist in the physical system and act as the interface between the proposed controller (software and hardware) and the system to be controlled.

In the case of the Light Control System, we identified, for example, the presence of a person in a room as a monitored variable and the light level in the room as a controlled variable. Both are clearly concepts in the physical world, and thus suitable candidates as monitored and controlled variables for the requirements model. A complete list of the monitored and controlled variables we identified in the LCS for a room are defined in Table 1.

4 Modeling the REQ Relation

After we have determined the scope of our system we are positioned to capture the required control behavior in some formal notation; we are ready to capture REQ—the topic of this section.

Since our work is based around a modeling language called RSML^{-e} (Requirements State Machine Language without events), a state-based language suitable for modeling of reactive control systems, we start the section with a short introduction to the notation before we continue with a discussion of the Light Control System requirements.

4.1 Introduction to RSML^{-e}

RSML^{-e} is based on the the language Requirements State Machine Language (RSML) developed by the Irvine Safety Research group under the leadership of Nancy Leveson [Leveson *et al.*, 1994]. RSML^{-e} was developed as a requirements specification language specifically for embedded control systems. One of the main design goals of RSML^{-e} was readability and understandability by non-computer professionals such as users, engineers in the application domain, managers, and representatives from regulatory agencies. RSML^{-e} is based on hierarchical finite state machines and dataflow languages. Visually, it is somewhat similar to David Harel's Statecharts [Harel and Pnueli, 1985, Harel, 1987, Harel *et al.*, 1990]. For example, RSML^{-e} supports parallelism, hierarchies, and guarded transitions. The main differences between RSML^{-e} and RSML are the addition in RSML^{-e} of rigorous specifications of the interfaces between the environment and the control software, and the removal of internal broadcast events. The removal of events was prompted by Nancy Leveson's experiences with with RSML and a new language called SpecTRM-RL that she has evolved from RSML. These experiences have been chronicled in [Leveson *et al.*, 1999].

An RSML^{-e} specification consists of a collection of *state variables*, *I/O variables*, *interfaces*, *functions*, *macros*, and *constants*, which will be briefly discussed below.

In RSML^{-e}, the state of the model is the values of a set of *state variables*, similar to modes in SCR [Heitmeyer *et al.*, 1995]. These state variables can be

Monitored Variables:

Name	Range	Description
System Variables:		
Light_Level	0..10000 Lux	The amount of light in the room
Occupied	Boolean	TRUE if room is occupied
Light_Level_Undetectable	Boolean	Used for light sensor failure.
Occupied_Undetectable	Boolean	Used for motion or door sensor failure.
Window_Light_Bank_Intensity	0..100	Intensity of Window Light Bank
Wall_Light_Bank_Intensity	0..100	Intensity of Wall Light Bank
Operator Inputs:		
Chosen1_LS_Button_InVar	Boolean	Chooses/Replaces light scene 1
Chosen2_LS_Button_InVar	Boolean	Chooses/Replaces light scene 2
Chosen3_LS_Button_InVar	Boolean	Chooses/Replaces light scene 3
Default_LS_Button_InVar	Boolean	Chooses/Replaces default light scene
Set_LS_Button_InVar	Boolean	If TRUE and another LS button is pressed, replaces the light scene with the current light scene.
T1	1..1440 minutes	Timeout to reestablish default light scene
T3	1..1440 minutes	Timeout to shut off lights
FacM_Shutoff	Boolean	Allows fac. man. to remotely shut off lights.

Controlled Variables:

Name	Range	Description
System Variables:		
Con_Window_Light_Bank_Intensity	0..100	Intensity of Window Light Bank
Con_Wall_Light_Bank_Intensity	0..100	Intensity of Wall Light Bank
Outputs to Operator:		
Failed	Boolean	TRUE if system detects component failure.

Table 1: The monitored and controlled variables for the Light Control System in a room.

organized in parallel or hierarchically to describe the current state of the system. Parallel state variables are used to represent the inherently parallel or concurrent concepts in the system being modeled. Hierarchical relationships allow *child* state variables to present an elaboration of a particular *parent* state value. Hierarchical state variables allow a specification designer to work at multiple levels of abstraction, and make models simpler to understand.

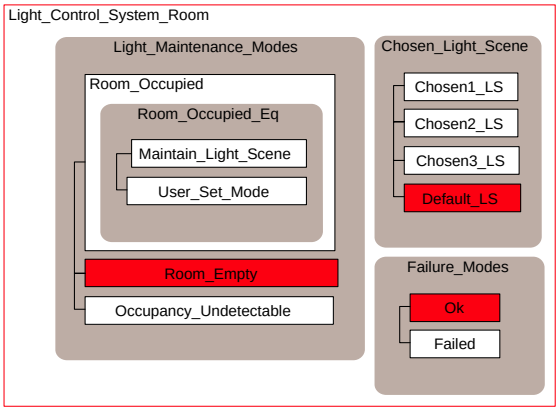

Figure 6: The state machine for the requirements model of the Light Control System in an individual room

For example, consider the Light Control System for an individual room. The state variable hierarchy used to model the requirements on this system could be represented as in Figure 6. This representation includes both parallel and hierarchical relationships of state variables. *Light_Maintenance_Modes, Chosen_Light_Scene* and *Failure_Modes* are three parallel state variables, and *Room_Occupied_Eq* is a child state variable of *Light_Maintenance_Modes*

Next state functions in RSML^{-e} determine the value of state variables. These functions can be organized as *transitions* or *condition tables*. Condition tables describe under what condition a state variables *assumes* each of its possible values. Transitions describe the condition under which a state variable is to *change* value. A transition consists of a source value, a destination value, and a guarding condition. A transition is taken (causing a state variable to change value) when (1) the state variable value is equal to the source value, and (2) the guarding condition evaluates to true. The two state function types are logically equivalent; mechanized procedures exist to ensure that both functions are complete and consistent [Heimdahl and Leveson, 1996].

The state functions are placed into a partial order based on data dependencies and the hierarchical structure of the state machine. State variables are data-dependent on any other state variables, macros, or I/O variables that are named in their transitions or condition tables. If a variable is a child variable of another state variable, then it is also dependent on its parent variable. The value of the state variable can be computed after the items on which it is data-dependent have been computed. For example, the value of the *Room_Occupied_Eq* state variable would be computed after the *Light_Maintenance_Modes* state variable,

because its value is dependent on whether or not *Light_Maintenance_Modes* is in the *Room_Occupied* state.

Conditions are simply predicate logic statement over the various states and variables in the specification. The conditions are expressed in disjunctive normal form using a notation called AND/OR tables [Leveson *et al.*, 1994] (see Figures 7, 9, 10, etc.). The far-left column of the AND/OR table lists the logical phrases. Each of the other columns is a conjunction of those phrases and contains the logical values of the expressions. If one of the columns is true, then the table evaluates to true. A column evaluates to true if all of its elements match the truth values of the associated columns. An asterisk denotes "don't care."

Transition: Occupancy_Undetectable → Room_Empty

Condition:

Occupied_InVar = TRUE	F
Occupied_Detectable_InVar = TRUE	T

Figure 7: A transition in the Light_Maintenance_Modes state machine.

I/O Variables in the specification allow the analyst to record the the monitored variables (MON) or values reported by various external sensors (INPUT) (in the case of input variables) and provide a place to capture the controlled variables (CON) or the values of the outputs (OUTPUT) of the system prior to sending them out in a message (in the case of output variables).

Interfaces, discussed briefly in Section 3.1, encapsulate the boundaries between the RSML^{-e} model and the external world.

To further increase the readability of the specification, RSML^{-e} contains many other syntactic conventions. For example, they allow expressions used in the predicates to be defined as functions (e.g., TotalIntensity()), and familiar and frequently used conditions to be defined as macros (e.g., OccupancyUndetectable()). *Functions* in RSML^{-e} are mathematical functions that are used to abstract complex calculations. A *macro* is simply a named AND/OR table that is used for frequently repeated conditions and is defined in a separate section of the document.

4.2 Overview of REQ

Our choice of the various monitored and controlled quantities places certain constraints on what can, and cannot, be specified in the REQ relation. This is natural, and desirable, because the REQ relation represents a mapping from monitored variables to controlled variables. If the monitored and controlled variables are chosen appropriately, then the specification of the REQ relation will be focused on the issues which are *central* to the requirements on the system.

In the Light Control System, some of the informal needs from the problem description will *not* be represented in the REQ relation or they will be represented in an abstract form. For example, the problem description states that

"If any outdoor light sensor or the motion detector of a room does not work correctly, the user of this room has to be informed" (U8)[4]. The REQ relation does not include any notion of a motion detector nor of the outdoor light sensor. Instead, the monitored quantities for the light level in the room and whether the room is occupied or not are used.

Often when starting to construct the REQ relation, it is helpful to examine the controlled variables of the system. It is necessary to determine what conditions partition the value of a particular controlled variables (i.e., what modes effect the controlled variable) and under what scenarios various outputs should be generated. In this light control system, the controlled variables are the intensity of the window and wall light groups in the office and the failure indication operator output. For now, we will focus on the window and wall intensity variables.

From the problem description, it becomes clear that there are two main activities which affect the way that the REQ relation must determine the values of the controlled variables. First, from needs it is clear that the REQ relation must be able to capture and use the various user set points (U6, U7, and U9 on page 9 of the problem description). Second, it is clear that the control system should somehow maintain the light level in the rooms (FM1 on page 10 and U2 on page 9). Furthermore, it would appear that these two task occur concurrently. The following two subsections consider these two tasks.

4.3 User Functionality

In this section, we will consider the maintenance of the various set points by the user and the way in which they are modeled in our version of the REQ relation. This is the simplest of the two main tasks of the REQ relation. Nevertheless, by constructing a formal model several issues with the definition of a light scene were exposed.

In the problem description, the room control panel (RCP) is described by U7, U9, and U12. U12 states that the RCP is a mobile, stand-alone device; therefore, a reasonable design for such a device seemed to be to keep the controls as simple as possible. Our initial design of the RCP is shown in Figure 8.

We envisioned the user selecting a light scene by simply pressing its associated button and setting a light scene by first adjusting the light in the room using the top two control groups and then pressing the "set" button while pressing one of the light scene buttons (i.e., similar to the functionality of a typical car radio).

This seems a simple, and intuitive solution; nevertheless, the definition of a light scene given in the problem description states that a name is associated with each light scene. In order to allow the user to enter a name, a more complex interface would seem to be in order. However, this would (1) add significantly to the cost and complexity of the RCP and (2) probably not add much, if anything, in terms of functionality and user satisfaction. Therefore, we decided to simplify the definition of the light scene so that the light scene is simply selected with a button on the RCP (*Scene 1, Scene 2, and Scene 3*).

To model the light scenes, we use a number of state variables to capture the light level in Lux. Figure 9 shows the RSML^{-e} definition for the state variable that models the light level for on of the first light scenes.

The value of the variable is updated only if the user is pressing the "set" button and the button for light scene 1. If this is the case, then the state variable

[4] All requirements references are to the original problem description made available with the call for papers.

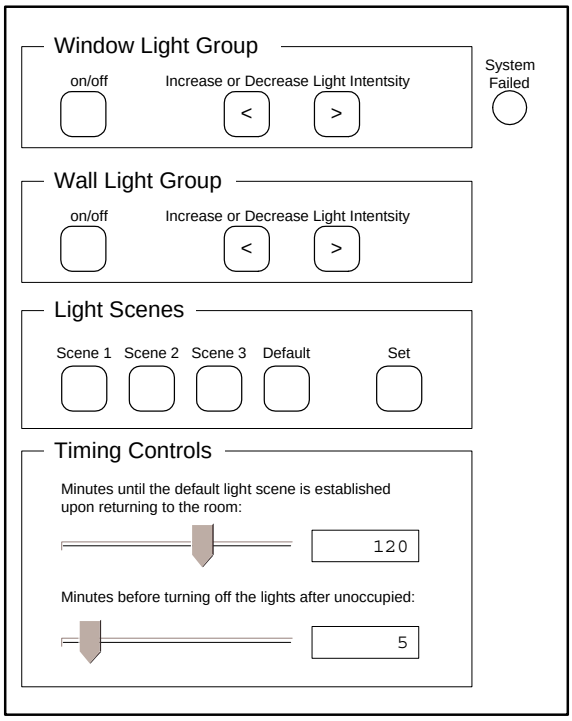

Figure 8: Initial Room Control Panel

records the value of the monitored variable for the current light level in the room. This the first case in Figure 9. Otherwise, the variable's value must stay the same. In RSML^{-e} this must be explicitly specified (the second case in Figure 9). The other light scenes (chosen2, chosen3, and default) all have similar definitions for the capture of the desired light level.

The light scene definition in the problem description states that the scene consists of the light level and an option: window, wall, or both. Furthermore, if the selection is both, then both light banks are used equally to obtain the desired light level. The control system can detect whether or not both light banks are on and thereby determine the value of this option when the set button was pressed. Nevertheless, we reasoned that the user would almost *always* have both light banks on to some degree or other.

Imagine what would happen if the user (1) adjusted the lights as desired, for example, 40% window and 60% wall; (2) set this to the first light scene; and (3) pressed the first light scene button. Clearly, this light scene uses both the wall and the window light groups. But, if the light scene stores the "both" option, then when the users presses the button for the first light scene, the window light group will be raised in illumination and the wall light group will be lowered. This behavior does not seem intuitive. Therefore, we expanded the definition of a light scene so that it consists of the light level in the room plus intensity of the window light group (0 - 100) and the intensity of the wall light group. Using this information, we can maintain the proportion (window/wall) that the user has selected and thereby, in our opinion, provide a better control behavior for the system.

State Variable

Chosen1_LS_Light_Level

Type: INTEGER

Units: lux

Expected Min: 0

Expected Max: 10000

:= Light_Level_InVar **IF**

Chosen1_LS_Button_InVar = ButtonPressType::kPressed	T
Set_Light_Scene_Button_InVar = ButtonPressType::kPressed	T

:= PREV_STEP(Chosen1_LS_Light_Level) **IF**

Chosen1_LS_Button_InVar = ButtonPressType::kPressed	T	F	F
Set_Light_Scene_Button_InVar = ButtonPressType::kPressed	F	T	F

Figure 9: Capturing the Light Level in a Light Scene

4.4 Maintaining the Light Level

After the user has chosen the light level and distribution(between window and wall) for the room, this light scene must be maintained. Also, there are a number of requirements related to the user leaving the room and the whether or not room occupancy is detectable that need to be considered when computing the values of the controlled variables. This section discusses the part of the light control specification that computes the control behavior for maintaining the light level and light distribution in the room.

The control system clearly behaves differently depending on the occupancy of the room; if there is a person in room, the control system must maintain the light level in the room. If there is no one in the room, the control system must determine whether or not to shut off the lights.

Figure 6 shows this partitioning of the Light_Maintenance_Modes. Each mode has certain conditions under which it is active. These conditions are specified with a state variable definition as shown in Figure 10. These modes depend on two monitored variables: *Occupancy_Detectable* and *Room_Occupied*, which determine whether we can detect the occupancy status of the room, and if so, whether or not the room is occupied.

The way that we have chosen to implement the control of the light level in the room is as follows: (1) the light level in the room is compared with the light level required by the current light scene, (2) if light level is not equal to the light level specified in the current light scene, the light intensity of the window/wall light banks are adjusted proportionally up or down by a small increment. Then, the system will poll the light level again within a short amount of time and

State Variable

Light_Maintenance_Modes

Location: Light_Control_System_Room

:= Room_Occupied

Condition:

Occupied_InVar = TRUE	T
Occupied_Detectable_InVar = TRUE	T

:= Occupancy_Undetectable

Condition:

Occupied_Detectable_InVar = TRUE	F

Transition: Occupancy_Undetectable → Room_Empty

Condition:

Occupied_InVar = TRUE	F
Occupied_Detectable_InVar = TRUE	T

Transition: Room_Occupied → Room_Empty

Condition:

Occupied_InVar = TRUE	F
Occupied_Detectable_InVar = TRUE	T

Figure 10: Light Maintenance Modes in the REQ Relation

eventually, the light in the room will comply with the selected light scene.

There is an issue, however, with the fact that it is *not* desirable to have the control system adjust the light intensity at the same time as the user attempts to adjust it; that is, the control system should not fight the user for control over the lights. Thus, it is necessary to partition the Room_Occupied mode into two sub-modes: one where the system is receiving user input and should produce no control actions and one where the system is responsible for maintaining the light level in the room. This partition can be seen in Figure 6.

The current light scene is the basis if the control of the light in the room. First, it is computed and then it is used to determine the values of the controlled variables. The current light scene, like any other light scene, consists of a light level (in lux) and the intensity of the window and wall light banks.

Figure 11 shows the state variable definition for the light level of the current light scene. The first case in the definition simply states that the the light level will be updated to the current light level in the room if the user is setting the controls for the room. This ensures that as the user makes changes to the lights, the changes are maintained, not reset, by the system.

The second group of cases in Figure 11 (cases 2-5) handles the user pressing one of the light scene buttons on the RCP. If the user presses one of these buttons, the light level associated with the selected light scene is used as the current light scene and will thus be maintained by the system.

The next two cases (6 and 7) determine the light level in the room if the room and unoccupied or reoccupied. The lights are shut off if the room is empty and either T3 has passed or the facility manager has issued the shutoff command. If the room is reoccupied, then the light level is determined by whether or not T1 has passed.

Finally, the light level will remain the same if the user is not making changes to the light level and room has not been recently reoccupied.

Once the light level is computed, it is compared with the monitored light level in the room. If the light needs to be increased/decreased, it is done so while maintaining the proportion of window light to wall light specified in the current light scene. The details of this part of the specification will not be included in this report due to space considerations.

5 Execution of the REQ Relation in NIMBUS

Now that we have an initial version of the REQ specification, it is possible to analyze it and simulate it in a realistic environment using the NIMBUS tools for requirements modeling. This paper will focus on the execution and simulation capabilities of the environment. This next section introduces the NIMBUS environment and presents our results for the light control case study.

5.1 The NIMBUS Environment

Assurance that the requirements model (system or software) possesses desired properties can be achieved through (1) manual inspections, (2) formal verification of the desired properties, or (3) simulation and testing of the specification. To achieve the high level of confidence in the correctness required in a safety-critical system, all three approaches must be used in concert. One environment, called NIMBUS, under development at the University of Minnesota provides support for all these activities [Thompson *et al.*, 1999, Thompson and Heimdahl, 1999].

Output Variable

Current_LS_Light_Level

Type: INTEGER

Units: lux

Expected Min: 0

Expected Max: 10000

:= Light_Level_InVar **IF**

..Room_Occupied_Eq **IN_STATE** User_Set_Mode	T

:= Chosen1_LS_Light_Level **IF**

Chosen1_LS_Button_InVar = ButtonPressType::kPressed	T
Set_Light_Scene_Button_InVar = ButtonPressType::kNotPressed	T

:= Chosen2_LS_Light_Level **IF**

Chosen2_LS_Button_InVar = ButtonPressType::kPressed	T
Set_Light_Scene_Button_InVar = ButtonPressType::kNotPressed	T

:= Chosen3_LS_Light_Level **IF**

Chosen3_LS_Button_InVar = ButtonPressType::kPressed	T
Set_Light_Scene_Button_InVar = ButtonPressType::kNotPressed	T

:= Default_LS_Light_Level **IF**

Default_LS_Button_InVar = ButtonPressType::kPressed	T
Set_Light_Scene_Button_InVar = ButtonPressType::kNotPressed	T

:= 0 **IF**

..Light_Maintenance_Modes **IN_STATE** Room_Empty	T	T
TIME >= ..Light_Maintenance_Modes **TIME_ENTERED** Room_Empty + T3_InVar	T	*
MESSAGE_AT(FacM_Shutoff)	*	T

:= Reoccupied_Light_Level() **IF**

..Room_Occupied_Eq **IN_STATE** User_Set_Mode	F
PREV_STEP (..Light_Maintenance_Modes **IN_STATE** Room_Occupied)	F

:= PREV_STEP(Current_LS_Light_Level) **IF**

..Room_Occupied_Eq **IN_STATE** User_Set_Mode	F
PREV_STEP (..Light_Maintenance_Modes **IN_STATE** Room_Occupied)	T

Figure 11: Current Light Scene Light Level in the REQ relation

This is the environment which we have used to capture and evaluate the required behavior of the Light Control System.

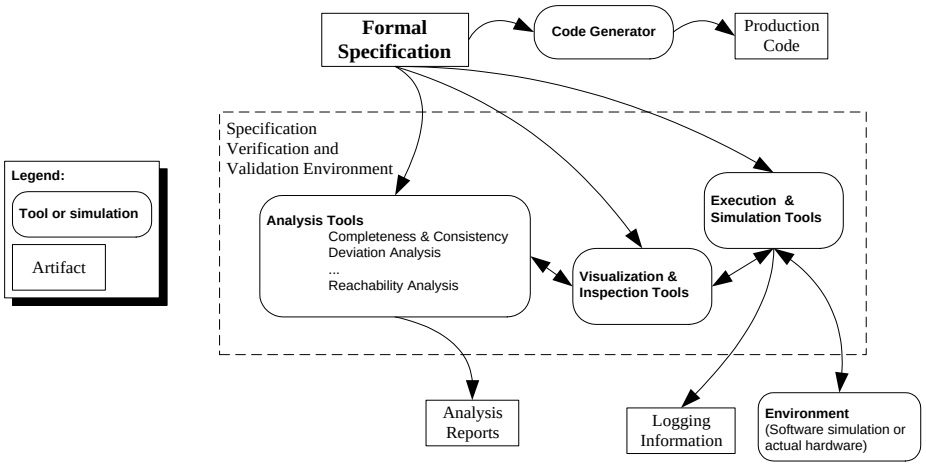

Figure 12: An overview of the support needed to effectively take advantage of code generation from formal high-level requirements models.

The three V&V techniques fill complementary roles within the validation and verification process. Manual inspections and visualization provide the specification team, customers, systems engineers, and regulatory representatives the means to informally verify that the behavior described formally in the specification matches the desired "real world" behavior of the system. Formal analysis is helpful to determine if the specification possesses desirable properties, for instance, if the specification is complete and consistent, and whether unsafe states are reachable. Simulation and testing are necessary to provide additional assurance that the specification captures the desirable behavior and is free of faults. In this report we will focus on the uniquely flexible capabilities for execution and simulation available in NIMBUS.

The execution and simulation capabilities NIMBUS environment are based on the ideas that (1) the engineers would like to have an executable specification of the system early in the project and that (2) as the specification is refined it is desirable to integrate it with more detailed models of the environment to enable accurate validation of the requirements model. Therefore, in the initial stages of the project, we want the executions to take their input from simple models of the embedding environment, for example, text files or user input. As the specification is refined, the analyst can add more detailed models of how the controlled system behaves, for example, additional RSML^{-e} specifications or software simulations of the system. As the requirements specification (REQ) is refined to a software specification (SOFT), models of the sensors and actuators can be incorporated into the executions. In order to have a closed loop simulation, a model of the process can be added between the sensor and actuator models. Finally, when the specification has been refined to the point of defining the software inputs

and outputs (INPUT and OUTPUT), the analyst can execute it directly with the hardware. This hardware-in-the-loop simulation closes the gap between the prototype and the actual hardware. These ideas are illustrated in Figure 13.

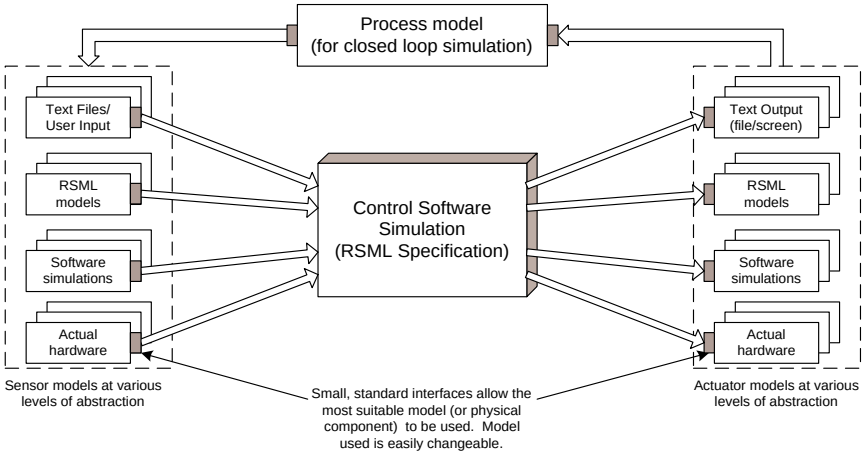

Figure 13: The NIMBUS Environment

The flexibility to quickly and easily connect different models of the components in a system provides new opportunities when creating and validating formal specifications. An environment such as ours can be used to aid in requirements based prototyping, in refinement of the specification as well as the environmental models, the evaluation of the operator interface, and in testing both the specification and the implementation derived from it. The detailed discussion of the various capabilities of our environment is beyond the scope of this paper and the interested reader is referred to [Thompson *et al.*, 1999].

5.2 Executing the Requirements

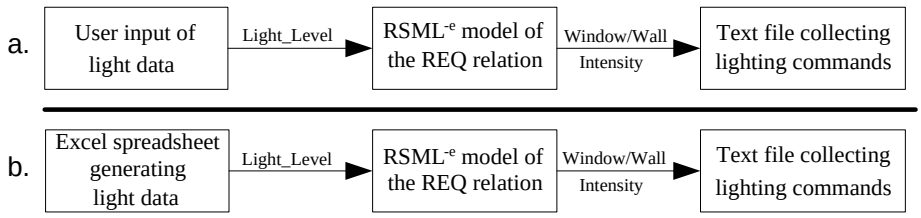

Figure 14: The REQ relation can be evaluated using text files or user input (a) or interacting with a simulation of the environment (b).

The NIMBUS environment allows us to execute and simulate the model we discussed in Section 4 using input data representing the monitored variables and collect output representing the controlled variables. Input data could come from several sources. The simplest option for input is, of course, to have the user specify the values (either interactively, or by putting the values into a text file ahead of time). This scenario is illustrated in Figure 14(a).

Unfortunately, it is often difficult to create appropriate input scenarios since the physical characteristics of the environment enforce constraints and inter-relationships over the monitored and controlled variables. For example, if we increase the light output from one of the light groups, how much will the illumination in the room increase? Thus, to create a valid (i.e., physically realistic) input sequence, the analyst must have a model of the environment. Initially, this model may be an informal mental model of how the environment operates. As the evaluation process progresses, however, a more detailed model is most likely needed. Therefore, in this stage of the modeling we may develop a simulation of the physical environment. The NIMBUS architecture lets us easily replace the inputs read from text files with a software simulation emulating the environment. This refinement can be done without any modifications to the REQ model.

For the Light Control System, we created a spreadsheet in Microsoft Excel to emulate the behavior of light groups and the illumination level in the room (Figure 14(b)) This simple environmental model allows us to interactively modify the traffic through the room, the external light available, etc., and to easily explore many possible scenarios.

Developing a model of the room control panel is another way to enhance the simulation of the REQ relation. There are a number of reasons why a mockup of the RCP is better than providing simply input from text files.

First, the RCP provides a significant amount of data to the REQ specification: The user settings of the window and wall intensity, the values of T1 and T3, and the selection and setting of four different light scenes. There are many different combinations of input sequences that can be generated from the room control panel (RCP) and to generate test files for even a small number of them would be frustrating at best.

Second, constructing an interface mockup provides invaluable opportunities to evaluate the control system with the intended users. For example, in our case we constructed a mockup of the RCP as it was pictured in Figure 8. When we tried to use this mockup, however, it became clear that the user would desire a separate control to set the hours and minutes so as to more easily enter the times T1 and T3 into the system. Our mockup was done using Visual Basic and is pictured in Figure 15. In the figure, both the window and the wall light groups have been turned on and the user has selected a value for T1 of 1 hour and 30 minutes and a value for T3 of 5 minutes.

Figure 16 shows data flow between the applications used to do the system simulation of the REQ relation. The values for the room occupancy and the facility manager shutoff signal are still represented by user inputs because a more complicated model is, in out opinion, not necessary at this stage . The values for light level and the RCP are represented as described above. Note that the user settings of the window and wall intensity must be passed both the the REQ specification and to the light model of the room.

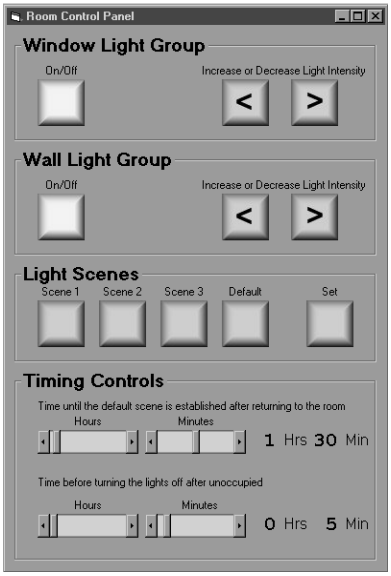

Figure 15: The Room Control Panel with both light banks on.

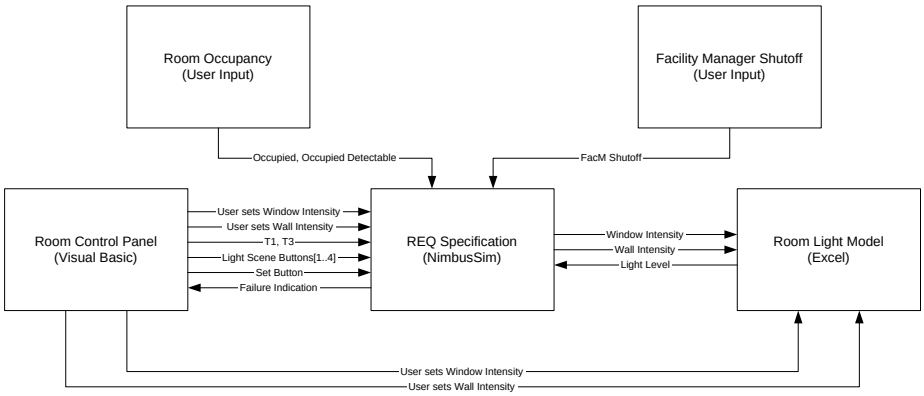

Figure 16: Overview of the REQ simulation

5.3 Results

Simulating the requirements provides the opportunity to discover conceptual errors in the requirements model as well as gain a greater understanding into the requirements themselves. In the case of the light control system, we discovered a number of mistakes in our formal requirements model. For example, we discovered numerous erroneous state transition definitions. Many, but not all, of these errors could be found be using the completeness and consistency analysis procedures that are also a part of the NIMBUS tools. In the interestes of space, we will not go into detail on these types of errors.

More interesting are errors which can be traced back to inconsistencies or underspecification of the original requirements document. For example, consider the case where the user has set the values of T1 and T3 so that T3 is greater than T1. In other words, suppose the user has adjusted the lights to a chosen light scene and then left the room. The user is out of the room long enough for T1 to pass, i.e., when the user re-enters the room (according to 9:U4) the default light scene has to be established. Because the user did not turn out the lights and T3 has not passed yet the lights in the room will still be on as the user left them. Nevertheless, when the control system detects that the room is occupied is *will* change the lighting in the room to comply with the default scene. This is the behavior that is specified in the problem description, but it might not be the behavior that users' expect.

We have found that simulation of the high-level requirements in a *realistic* environment is incredibly valuable for finding these and other types of conceptual errors. Experimentation in this fashion provides a specification of REQ that is a solid foundation from which to refine a specification of the SOFT relation: a topic which is addressed in the following section.

6 Refining System Requirements to Software Requirements

Once the model of REQ has been thoroughly simulated and analyzed, it is ready to be refined to a model of SOFT_{REQ}. This section discusses how to refine REQ to SOFT_{REQ} and how this was done for several monitored variables in the light control system.

6.1 Refine REQ to SOFT_{REQ}

From the start of the modeling effort, we know that we will not be able to directly access the monitored and controlled variables–we must use sensors and actuators. Thus, when refining REQ to SOFT, we will not be able to use variables such as *Occupied*. At this early stage, we may not know exactly what hardware will be used for sensors and actuators; but, we do know that we must use something and we may as well prepare for it early. By simply encapsulating the monitored and controlled variables we can get a model that is essentially isomorphic to the requirements model; the only difference is that this model is more suited for the refinement steps that will follow as the surrounding system is completed.

In our case, using a macro, *IsOccupied*(), instead of the monitored variable *Occupied* will shield the specification from possible changes in how the final

software will determine that a room is occupied (See Figure 17). By performing this encapsulation for all monitored and controlled variables we refine REQ to SOFT_{REQ}, a mapping from estimates of the monitored variables to an internal representation of the controlled variables.

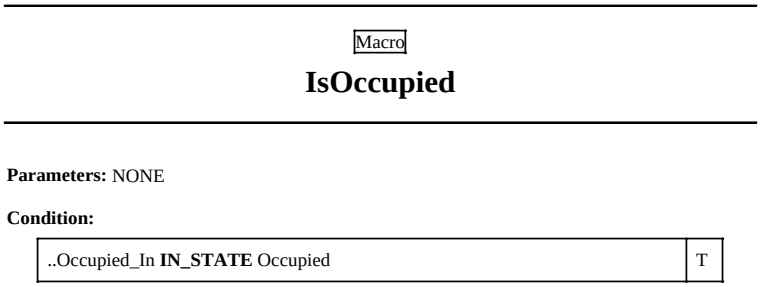

Figure 17: The IsOccupied() macro from the refined light control system

6.2 IN, OUT, IN^{-1}, and OUT^{-1}

As the hardware components of the system are defined (either developed in house or procured), the IN and OUT relations can be rigorously specified. The IN and OUT models represent our assumptions about how the sensors and actuators operate.

With the information about the sensor (IN) and actuator (OUT) relations, we can start adding to the SOFT_{REQ} relation to move towards SOFT. To achieve this, we refine the IN^{-1} relation in our model. In the Light Control System there are several sensors and actuators that must be considered. In the following sections, we will focus our attention on the sensors needed to determine if a room is occupied and to detect the light level in the room.

6.2.1 Refining Occupied

The control system must compute whether or not the room is occupied based on the input from the motion detector and the door sensors. For simplicity, we assume that the door sensors are wired as one input to the system so that if any of the doors are open, then the door sensor indicates "open" and if all the doors are closed, then the door sensor reads "closed". When computing whether or not the room is occupied, it is necessary to have some state information (i.e., we must know whether or not the room was occupied in the previous instance in time); therefore, a state variable needs to be added so that IN^{-1} can be properly computed.

The refined state machine can be seen in Figure 18. Instead of the idealistic true occupancy of the room used when evaluating REQ, the specification now

takes the input from one motion detector and the door sensors. Thanks to the structuring of the SOFT relation, this refinement could be done with minimal changes to the SOFT_{REQ} relation.

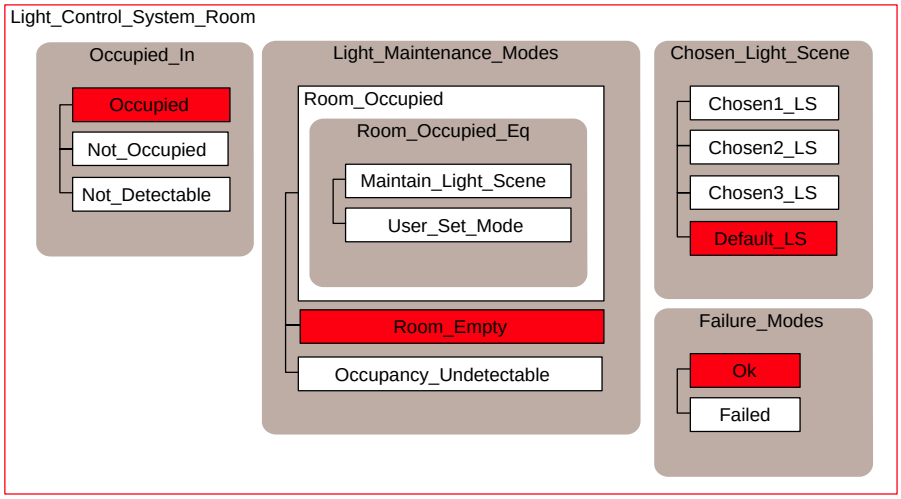

Figure 18: The state machine from Figure 6 refined to include the IN^{-1} portions for Occupied

The computation of the occupied quantity is shown in Figure 19. The first two cases determine whether or not the room is occupied based on the value of the motion detector. The last case, the condition to be in Not_Detectable, defines the conditions under which there may be a sensor failure or malfunction. If the room was not occupied and the doors have remained closed and then motion is detected; there must be a problem with the sensors.

When attempting to complete this refinement, we discovered conflicts in the problem statement. For instance, consider the following elements from the problem description: (1) "If any motion detector of a room or hallway section does not work correctly, the control system should behave as if the room or hallway section were occupied" (NF4); (2) the motion detector can detect even small motions within the room and it covers the whole room (sensor description); and (3) the occupancy of a room cannot change when the doors are closed (customer feedback 25).

On the surface, these all seem reasonable statements. The first two statements imply that the control software should attempt to detect sensor failures. The last statement says that if the doors in the room are closed, it doesn't matter what the reading from the motion detector is, the occupancy of the room will be unchanged. However, if this were added to, for example, the condition to be in Not_Occupied then it would overlap with the condition to be in Not_Detectable and an inconsistency is introduced in the model.

State Variable

Occupied_In

Location: Light_Control_System_Room

:= Occupied

Condition:

Motion_Detected_InVar = TRUE	T

:= Not_Occupied

Condition:

Motion_Detected_InVar = FALSE	T

:= Not_Detectable

Condition:

PREV_STEP(DoorSensor_InVar = DoorSensorType::kClosed)	T
PREV_STEP(..Occupied **IN_STATE** Not_Occupied)	T
Motion_Detected_InVar = TRUE	T
DoorSensor_InVar = DoorSensorType::kClosed	T

Figure 19: The definition for the Occupied_In state variable

6.2.2 Refining Light Level

In the REQ specification, we assume we know the correct level of light in the room (Light_Level). In reality, this is not the case; we must add sensing capabilities to determine an approximation of the light level in the room.

The specification states that we should attempt to compute the light level given an outdoor light sensor and the amount of illumination from the two light banks (in a feed-forward type fashion). However, this is extraordinarily difficult because of several factors:

- In most office buildings, offices are equipped with blinds or curtains. If the user closes the blinds, then little sunlight will enter the room. Even if there are no blinds, if the light coming into the room bothers the occupant he or she will most likely find some way to cover the windows (e.g., by putting up paper). Thus, transmission of light through the glass cannot be assumed to behave according to some set function.
- Users may have desk lamps or other sources of illumination. The light-level algorithm will not be able to account for these alternate light sources.

- The amount of illumination measured by the outdoor light source depends on the angle of the sun relative to the sensor, as well as the intensity of the light. The amount of light actually entering a room depends on the position of the window relative to the sun. All these factors vary with the time of day, weather conditions, and time of year. Providing an accurate calculation of the light level in the room based on the light level at an external sensor would be prohibitively expensive (if not impossible).
- The algorithm must assume that all filaments in the light banks are functioning correctly (i.e., no burned out filaments). This assumption may not be valid.

Because these concerns make the light-level computation error-prone at best, we have chosen to introduce a light sensor into each room or hallway that we monitor. Using this approach, we can directly measure the light level in the room. The monitored variable *Light_Level* can be viewed as a state variable whose value is a scaling function of the light sensor value. For now, we assume that the light sensor is similar to type as described for the outdoor light sensor, so the scaling factor is 1.

In the specification, we introduce a new input variable *Light_Sensor_Level* that records the raw sensor value. Then we (trivially) convert it to the *Light_Level* monitored variable. Although in this case *Light_Sensor_Level* is always the same as *Light_Level*, both are useful, because they decouple the REQ relation from the particular sensors. If we change the sensors, we just have to change the definition of the *Light_Level* variable, without impacting the rest of REQ.

Given that we have a light sensor in the room, one problem is that we only know the light level at the location of the light sensor. Therefore, where the light sensor is placed in the room is important. If the light sensor is obscured, or if it is placed very close to one of the light banks, then the light level of the room may be inaccurately measured. Depending on how accurate we required the light level to be, we could create an environment model of the room in which we could move the the light sensor around, then connect it to a RSML^{-e} simulation to investigate the behavior of the system.

6.3 Refining the Simulations

When evaluating RSML specifications in NIMBUS, the analyst has great freedom in how he or she models the environment. When we evaluated the REQ model in Section 4, we used a user or a software simulations to provide the RSML^{-e} model with monitored variables and to evaluate the controlled variables. As the IN^{-1} and OUT^{-1} relations are added to the RSML^{-e} model, the data provided (and consumed) by the model of the embedding environment must be refined to reflect the software inputs and outputs (INPUT and OUTPUT) instead of the monitored and controlled variables (MON and CON).

This can be achieved in two ways; (1) refine the model of the physical process to produce INPUT and consume OUTPUT (incorporate sensors and actuators into the model of the environment), or (2) add explicit and separate models of the sensors and actuators to the simulation. In reality, the refinement of the environmental model and the SOFT relation progress in parallel and is an iterative process. The sensor and actuator models may be added one at a time and the interaction with different components may merit different refinement strategies. NIMBUS naturally allows any combination of the approaches mentioned above to be used.

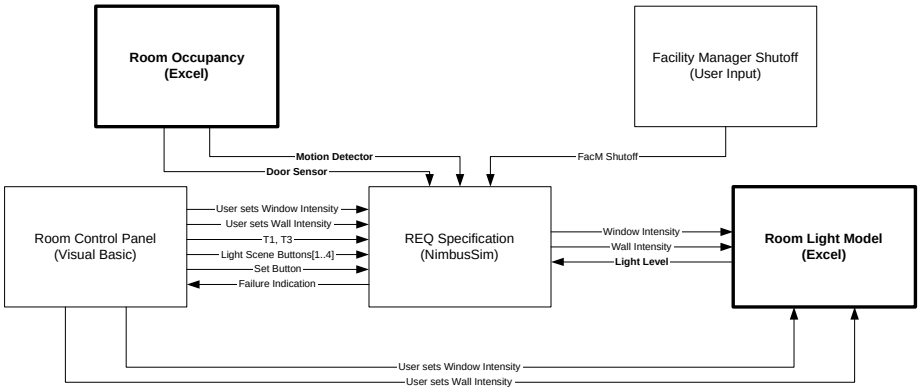

Figure 20: The simulation with the refined notion of occupied

Figure 20 show the refined simulation overview for the light control system. We have replaced the user input for the room occupancy monitored variable with an Excel spreadsheet to model the motion detectors and the doors in the room. This spreadsheet now supplies the required motion detector and door status inputs to the specification. Also, the light model spreadsheet has been refined.

As the refinement of the SOFT relation and the models of the environment progresses, we may at some point desire to perform hardware-in-the-loop simulation. Such simulations are easily accommodated in the NIMBUS framework. If we want to evaluate the Light Control System software requirements interacting with the actual hardware components, we can use any standard data acquisition card to access the hardware[5]. NIMBUS provides a collection of sample interfaces to the data acquisition card that can be easily modified to communicate with the desired hardware components. In the case of the Light Control System, we may want to take actual input from the light sensor in the room and physically control the two light groups while we use software simulations for the motion sensors and door sensors[6].

The use of hardware-in-the-loop simulation does not only provide a powerful evaluation of the proposed software system, we can also use NIMBUS to evaluate the physical system itself. For instance, by forcing the RSML^{-e} model of the software requirements into unexpected and/or hazardous states, we can inject simulated software failures into the hardware system.

To summarize, NIMBUS provides a flexible framework in which a software requirements model expressed in RSML^{-e} can be executed while it interacts with various models of the other components in a proposed system. NIMBUS supports the refinement of the REQ relation to a SOFT relation and allows easy interchange of components in the environment as the refinement takes place. It is important to recognize the difference between models which are good for representing the physical process versus models, like RSML^{-e} models, which are good for modeling the software control of the process. Modeling the process itself accurately may require complex numerical functions, for example, to generate normally distributed random errors. These types of functions are not and should not be within the scope of RSML^{-e}. However, an accurate model

[5] Currently, we are using the National Instruments DAQ 1200 series modules.

[6] Note that we have not yet had the opportunity to use any hardware from the Light Control System in our simulations.

of the process is key to the success of specification-based prototyping. This is the reason that NIMBUS provides the flexibility to integrate with many different models expressed in various ways, and one of the primary contributions of the NIMBUS environment.

7 Evaluation and Discussion

In this report we have summarized our experiences with using RSML^{-e} and the NIMBUS environment to model the required control behavior of of the Light Control System. Our experiences were generally positive and the modeling effort went by without any major complications.

The mode was developed by two graduate students over approximately three weeks time (part time). Both students were very familiar with the language, its formal semantics, and the NIMBUS environment. Unfortunately, we did not compile any accurate data on the effort required to complete the model.

The complete formal model is approximately 50 pages in length. This may seem like a large specification for such a simple problem, but the specification is formatted for readability (a lot of white space and page breaks to make it visually appealing) as well as informal English descriptions of the various parts of the model.

The main problem we encountered during the specification effort was the incompleteness of the informal requirements provided in the problem description. The formal model we developed forced resolution of many issues that might otherwise have been passed over until the detailed design or implementation stages. The simulation and execution allowed us to evaluate different panel designs and helped clarify the requirements. In addition, the detailed nature of an RSML^{-e} model forces early adoption of a control strategy. In this effort we found the notion of an external light sensor wholly unacceptable and modified the requirements accordingly. A less detailed modeling approach may defer this issue until later and end up requiring a behavior that cannot be realized in a a physical system.

Naturally, the modeling effort exposed some areas where our modeling approach needs improvements.

7.1 Issues for Future Work

The main issue raised during the modeling effort was the lack of an array construct (similar to what is available in Statecharts) in RSML^{-e}. The light scenes are concepts that are naturally modeled as an array of identical models (the three user programmable light scenes and the default light scene are identical). In our model, however, we had to explicitly include four sets of variables to model the light scenes.

There is no technical reason why arrays have not been included in RSML^{-e}. We originally developed the tool support for RSML^{-e} to prototype and evaluate various static analysis procedures, for instance, completeness and consistency checking, reachability analysis, etc. Since arrays do not add any modeling power—they are simply a syntactic nicety—but add considerable effort when implementing a tool supporting the language, we deemed them superfluous for the more theoretical work we were involved with at the time; to make the tool development easier we omitted arrays from the language. Our subsequent experiences, however, have convinced us that arrays are an absolute necessity in practical modeling and we are currently extending our tool to support arrays.

References

[Faulk *et al.*, 1992] S. Faulk, J. Brackett, P. Ward, and J Kirby, Jr. The CoRE method for real-time requirements. *IEEE Software*, 9(5), September 1992.

[Harel and Pnueli, 1985] D. Harel and A. Pnueli. On the development of reactive systems. In K.R. Apt, editor, *Logics and Models of Concurrent Systems*, pages 477–498. Springer-Verlag, 1985.

[Harel *et al.*, 1990] D. Harel, H. Lachover, A. Naamad, A. Pnueli, M. Politi, R. Sherman, A. Shtull-Trauring, and M. Trakhtenbrot. Statemate: A working environment for the development of complex reactive systems. *IEEE Transactions on Software Engineering*, 16(4):403–414, April 1990.

[Harel, 1987] D. Harel. Statecharts: A visual formalism for complex systems. *Science of Computer Programming*, pages 231–274, 1987.

[Heimdahl and Leveson, 1996] Mats P. E. Heimdahl and Nancy G. Leveson. Completeness and consistency in hierarchical state-base requirements. *IEEE Transactions on Software Engineering*, pages 363–377, June 1996.

[Heitmeyer *et al.*, 1995] C. L. Heitmeyer, B. L. Labaw, and D. Kiskis. Consistency checking of SCR-style requirements specifications. In *Proceedings of the Second IEEE International Symposium on Requirements Engineering*, March 1995.

[Jackson, 1995] Michael Jackson. The world and the machine. In *Proceedings of the 1995 Internation Conference on Software Engineering*, pages 283–292, 1995.

[Leveson *et al.*, 1994] N.G. Leveson, M.P.E. Heimdahl, H. Hildreth, and J.D. Reese. Requirements specification for process-control systems. *IEEE Transactions on Software Engineering*, pages 684–706, September 1994.

[Leveson *et al.*, 1999] Nancy G. Leveson, Mats P.E. Heimdahl, and Jon Damon Reese. Designing specification languages for process control systems: Lessons learned and steps to the future. In *Seventh ACM SIGSOFT Symposium on the Foundations on Software Engineering*, September 1999.

[Miller, 1999] Steven P. Miller. Modeling software requirements for embedded systems. Technical report, Advanced Technology Center, Rockwell Collins, Inc., 1999. In Progress.

[Parnas and Madey, 1991] David L. Parnas and Jan Madey. Functional documentation for computer systems engineering (volume 2). Technical Report CRL 237, McMaster University, Hamilton, Ontario, September 1991.

[Thompson and Heimdahl, 1999] Jeffrey M. Thompson and Mats P.E. Heimdahl. An integrated development environment prototyping safety critical systems. In *Tenth IEEE International Workshop on Rapid System Prototyping (RSP) 99*, 1999.

[Thompson *et al.*, 1999] Jeffrey M. Thompson, Mats P.E. Heimdahl, and Steven P. Miller. Specification based prototyping for embedded systems. In *Seventh ACM SIGSOFT Symposium on the Foundations on Software Engineering*, September 1999.