Chapter 1

Introduction

This thesis investigates how the formal modeling and verification techniques of computer science can be used for the analysis of hybrid systems [7,14,22,37] — systems involving both discrete and continuous behavior. The motivation behind such research lies in the inherent similarity of the hierarchical and decentralized control strategies of hybrid systems and the communication and operation protocols used for distributed systems in computer science. As a case study, the thesis focuses on the development of techniques that use hybrid I/O automata [29, 30] to model and analyze automated vehicle transportation systems and, in particular, their various protection subsystems — control systems that are used to ensure that the physical plant at hand does not violate its various safety requirements.

The thesis is split into two major parts. In the first part, we develop an abstract model of a physical plant and its various protection subsystems — also referred to as *protectors*. The specialization of this abstract model results in the specification of a particular automated transportation system. Moreover, the proof of correctness of the abstract model leads to simple correctness proofs of the protector implementations for particular specializations of the abstract model. In this framework, the composition of independent protectors is straightforward — their composition guarantees the conjunction of the safety properties guaranteed by the individual protectors. In fact, it is shown that under certain conditions composition holds for dependent protectors also.

In the second part, we specialize the aforementioned abstract model to simplified versions of the personal rapid transit system (PRT 2000™) under development at Raytheon Corporation. We examine overspeed and collision protection for a set of vehicles traveling on straight tracks, on binary merges, and on a directed graph of tracks involving binary merges and diverges. In each case, the protectors sample the state of the physical plant and take protective actions to guarantee that the physical plant does not reach hazardous states. The proofs of correctness of such protectors involve specializing the abstract protector to the physical plant at hand and proving that the suggested protector implementations are correct. This is done by defining simulations among the states of the protector implementations and their abstract counterparts.

1.1 Hybrid Systems

The trend of system integration and automation has resulted in large and complex systems involving hierarchical and/or decentralized control structures. The higher levels of control are based on discrete algorithms and are often modeled using finite automata techniques from computer science. The lower levels of control address continuous behavior and are based on well established control theoretic techniques. The inherent complexity of the mix of continuous and discrete control and the need of a precise and efficient model for hybrid systems has encouraged research in this field.

The similarity of the hierarchical and/or decentralized control structure of hybrid systems with the distributed system setting in computer science has nurtured a distributed systems approach of analyzing hybrid systems. This approach is based on various formal modeling techniques developed for the verification and the proof of correctness of distributed systems in computer science. Such techniques use the principles of abstraction and modular decomposition to provide simple and concise models of complex systems. Once a particular system is decomposed into succinct parts, various composition theorems are used to prove that the system is functioning according to its specifications, *i.e.*, the system is *correct*.

1.1.1 Formal Framework

The formal modeling techniques that are used in this thesis are based on the *hybrid I/O automaton* model [29,30]. This model is an extension of the *timed I/O automaton* model [11, 34] and allows the explicit treatment of continuous behavior. The hybrid I/O automaton model is inspired by the phase transition models [2,4,35,36].

The hybrid I/O automaton model is a (possibly) infinite state model of a system involving both discrete and continuous behavior. The states of a hybrid I/O automaton (HIOA) are the valuations of a set of variables. The discrete behavior of a HIOA is modeled by discrete jumps in state which are described by labeled transitions. The labels of such transitions are the *actions* that carry out the transition from the initial to the final state of the jump. The continuous behavior of a HIOA is modeled by continuous changes in state which are described by sets of *trajectories*. The external interface of a HIOA is dictated by the partition of its variables and its actions into three categories: input, internal, and output. The behavior of the system being modeled over time is described by *hybrid executions* — finite or infinite alternating sequences of trajectories and actions. The externally visible part of a hybrid execution is denoted as the *hybrid trace* of the hybrid execution and involves the evolution of the input and output variables of the HIOA.

A HIOA A_1 implements another HIOA A_2 if every external behavior of A_1 is allowed by A_2. In this setting A_1 and A_2 are referred to as the *implementation* and the *specification*, respectively. The notion of an implementation relation is given by inclusion of the sets of hybrid traces; that is, the set of hybrid traces of A_1 is a subset of the set of hybrid traces of A_2. The composition of two HIOA is defined as their synchronization on shared input/output variables and input/output actions. Under straightforward and simple conditions, the composition of two HIOA results in a HIOA. Moreover, composition respects the implementation relation, *i.e.*, supposing B is a HIOA, if the HIOA A_1 implements the HIOA A_2, then the composition of A_1 with B implements the composition of A_2 with B.

Most of the proofs in the HIOA framework use *invariant assertions* and *simulations*. In the case of invariant assertions, the proofs are by induction on the length of a hybrid execution of the HIOA at hand. Such proofs show that a particular predicate on the state of the HIOA is satisfied in every state of the execution. A simulation is a mapping between the states of the two HIOA and is used to prove that one HIOA implements another. The fact that the mapping is indeed a simulation is again done by induction on the length of a hybrid execution of the implementation. This induction matches up individual steps in the implementation with either single steps, or sequences of steps, in the specification.

1.1.2 Related Work

The recent interest in the area of hybrid systems has resulted in a number of techniques to model and analyze their behavior. In particular, models that are analogous to the timed I/O automaton model [11,34] are the models of Alur and Dill [6], Lamport [20], and Henzinger, Manna, and Pnueli [18]. As is the case with the timed I/O automaton model, these models have also been extended to the hybrid setting; for instance, the timed transition model [18] has been extended to the phase transition model [35,36]. Phase transition systems are analogous to hybrid I/O automata — the transitions and the activities of phase transition systems correspond to the discrete transitions and the trajectories of hybrid I/O automata. The hybrid system model [2,4] is similar to the phase transition model with the distinction that, as in the hybrid I/O automaton model, discrete transitions are labeled, thus allowing the appropriate synchronization of composed automata. The distinction between the hybrid system model [2,4] and the hybrid I/O automaton model lies in the latter's classification of the discrete transitions and variables into input, internal, and output.

In the realm of applications, the formal modeling techniques presented above have been used for the analysis of various problems. The railroad crossing problem [15] and the steam boiler problem [1,21] comprise two commonly used benchmark problems. The former benchmark problem considers the control of a railroad gate that prevents cars and pedestrians from crossing the railroad tracks while the train is in the vicinity of the crossing. This gate must be lowered prior to the arrival of the train and lifted once the train has passed by. The latter benchmark problem involves the control of the level of water in a steam boiler.

The success in the modeling, analysis, and controller design for the above benchmark problems has encouraged the formal modeling of more complex hybrid systems; for example, automated transportation systems [41,42], industrial and chemical processes [9,40], rail-vehicle control [39], and complex automotive suspension systems [38]. The motivation behind such research lies mostly in the safety-critical nature of the systems at hand. In the case of automated transportation systems, the safety of the passengers has greatly encouraged the use of formal techniques.

The recent interest in addressing safety concerns related to automated highway systems and, in particular, the California PATH project [41], has resulted in a surge of hybrid system problems. The goal of PATH is to increase vehicle throughput by organizing traffic into platoons of closely spaced vehicles. Godbole, Lygeros, and Sastry [12, 13, 23, 25, 27] at U.C. Berkeley have studied various problems that arise in the control of the vehicle platoons. Such problems address the control of the leader of a platoon in view of following the preceding platoon at a safe distance, tracking an optimal cruising velocity, and performing various platoon maneuvers. The platoon maneuvers that have been addressed are the *platoon join*, in which two or more adjacent platoons join to form a single platoon, the *platoon split*, in which a platoon splits in two, and the *platoon lane change*. Lygeros [22] and Lygeros *et al.* [26, 27] used a game theoretic approach to prove that all platoon maneuvers are safe. Recently, Dolginova and Lynch [8] have used hybrid I/O automata to model and verify the safety of the platoon join maneuver.

On a similar note, Weinberg [43] has analyzed a deceleration maneuver in which a discrete controller slows a train down to a target velocity range within a given distance. In further research, Weinberg *et al.* [42] have modeled the personal rapid transit system (PRT 2000™) under development at Raytheon Corporation and verified the correct operation of the emergency control components used to guarantee that the vehicles neither exceed a prespecified speed limit, nor collide among themselves.

1.2 Automated Transportation Systems

Among the hybrid systems that are being analyzed using formal methods, systems in transportation are particularly common. This is due to the fact that such systems are safety-critical and, therefore, their correct analysis and verification is of uttermost importance.

An important feature of the design of the various autonomous transportation systems is

Figure 1.1 Separation of system functionality into operation and protection.

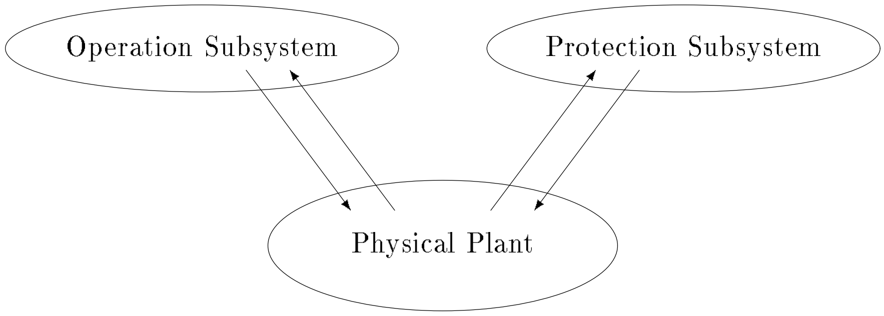

their absolute safety requirements. These requirements translate to stringent design criteria and have led to the complete separation of the system functionality into the parallel components of *operation* and *protection* as shown in Figure 1.1. The operation component is responsible for the "normal" control of the system and can be composed of complex software and hardware. The protection component is responsible for the "emergency" control of the system and is designed to be simple and reliable. In ordinary operation, the protection component is not supposed to take any action — it merely monitors the system. In a potentially hazardous situation, however, the protection component must react strongly enough to guarantee that, regardless of the behavior of the operation component, the safety requirements are met. In the interest of making the protection component reliable, designers keep it simple; instead of having complex control abilities, the protection component depends only on the correct execution of a few decisive emergency commands.

The separation of operation and protection functions is a generally recognized engineering paradigm for the design of safety-critical systems. In the realm of transportation systems, this structure was initially used in the design of railroad systems. Automatic safety systems were added to human-controlled railroad systems to protect against human error and mechanical malfunctions. As railroad and mass transit systems have evolved to become more automated, this division of labor has been retained in the form of Automatic Train Operation (ATO) and Automatic Train Protection (ATP) systems. This paradigm occurs in most existing automated train systems, including the Washington Metro, the Miami People Mover, the O'Hare People Mover, the Detroit People Mover, and systems in Toronto, Vancouver, and Jacksonville. The use of this split migrated to automated vehicle transportation systems with the pioneering Morgantown PRT system in the late sixties; this system has been in continuous active use for over 20 years with no serious accidents.

1.2.1 The PRT 2000™

Raytheon engineers are currently working on the design and development of a new personal rapid transit (PRT) system called PRT 2000™. This system uses 4-passenger vehicles that travel on an elevated guideway with Y-shaped merges and diverges. Passengers on this system board at stations and travel directly to their desired destination stations without intermediate stops. Compared to conventional transportation systems, the PRT 2000™ can provide shorter average trip times and shorter average waiting times with equivalent passenger throughput. These performance improvements are achieved because the vehicles are separated on the guideway by only a few seconds, instead of the minutes typical of a conventional transit system. The vehicles are controlled by a distributed network of computers, which receive data from sensors on the vehicles and in the tracks.

Once again, the control of the PRT 2000™ is split into the Automated Vehicle Operation System (AVOS) and the Automated Vehicle Protection System (AVPS). The AVOS is in charge of the normal operation of the system and the AVPS is used to protect the system against hazards.

1.2.2 Formal Modeling of the PRT 2000™

The safety-critical nature of the PRT 2000™ has lead to an interest in modeling its protection system using formal modeling techniques from computer science. The advantage of using such modeling methods is twofold. First, they formalize the safety concerns addressed by the protection system and, second, they are used to prove the correctness of the protection system at hand. The safety properties that are addressed are those of overspeed and collision avoidance, *i.e.*, either the property that the vehicles comprising the system do not exceed the speed limit, or the property that they do not collide among themselves. These are by no means the only safety requirements enforced by the AVPS of the PRT 2000™, but they are among the most important and complex.

The approach to modeling this automated transportation system is based on abstraction and modular decomposition. Abstraction is used to mask all inessential implementation details from the model of the system. Modular decomposition is used either to model each of the safety properties in isolation, or to model a particular safety property as the conjunction of several less complex safety properties. As shown in Figure 1.2, the protection system is defined as the composition of a set of simpler modules referred to as *protectors*. The composition of all these protectors results in a protection system that enforces the conjunction of the safety properties enforced by the individual protectors being composed. For instance, in the case of a protection system that prevents the vehicles from exceeding the speed limit, each of the protectors would correspond to protection subsystems that prevent individual vehicles from exceeding the speed limit. However, their composition would constitute an overspeed protection system for all the vehicles.

This thesis extends the work by Weinberg, Lynch, and Delisle [42] on modeling the AVPS of the personal rapid transit system (PRT 2000™) under development at Raytheon Corpo-

Figure 1.2 Modular decomposition of the AVPS of the PRT 2000™.

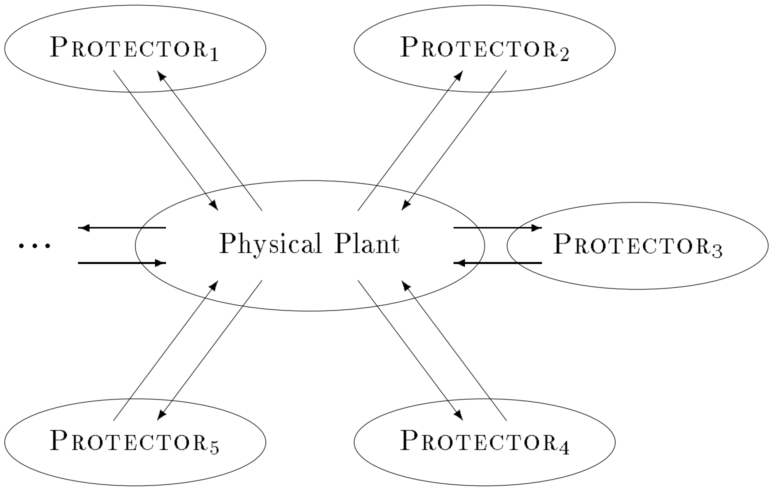

ration. Weinberg *et al.* [42] model the PRT 2000™ as a transportation system where:

- vehicles are traveling on a single track,
- vehicle velocities are non-negative,
- vehicles can stop instantaneously, as if they could hit a brick wall,
- collisions among vehicles are pairwise,
- brakes are binary, *i.e.*, the braking of a particular vehicle results in a vehicle deceleration equal to a prespecified value,
- the acceleration is constrained to a particular range of values, and
- the vehicle brakes comprise monotonic system constraints, *i.e.*, the instruction of a vehicle to brake can never be revoked.

In addition to the above assumptions, the communication among the various subsystems of the PRT 2000™ is assumed to be reliable, periodic, and timely.

Weinberg *et al.* [42] verify the correctness of the overspeed and the collision protection subsystems. First, it is shown that the overspeed protector guarantees that none of the vehicles exceed the speed limit and that the collision protector prohibits vehicle collisions provided that none of the vehicles exceed the speed limit. Then using a one-way dependence protector composition theorem it is shown that the composition of the overspeed and collision protectors guarantees that the vehicles neither exceed the speed limit, nor collide among themselves. It should be noted that the model of the physical plant is simplified to the point that abrupt changes of the vehicle velocities, due to collisions for example, are not modeled. The advantage of this simplification is that the overspeed protector does not depend on the collision protector and, therefore, the one-way dependence protector composition theorem suffices. The disadvantage is that the simplified model might not be a truthful representation of the real physical plant.

In this thesis, we extend the protector composition results of Weinberg *et al.* [42] and relax their modeling assumptions about the PRT 2000™. Regarding the composition of protection systems, we present theorems that dictate the conditions under which the composition of independent, one-way dependent, and even two-way dependent protectors guarantees the conjunction of the safety properties guaranteed by the individual protectors being composed. Regarding the transportation system model, two of the aforementioned assumptions are relaxed. First, the constraint on the track topology is gradually relaxed from that of a single track to that of a general track topology involving a directed graph of tracks comprised of Y-shaped merges and diverges. Second, the monotonicity constraint on the instruction of the vehicles to brake is relaxed such that the instruction of a vehicle to brake may be revoked, provided the vehicle in question is out of risk. Moreover, in an effort to truthfully model the transportation system, we extend the model of the physical plant to allow vehicle collisions that can adversely affect the velocity and the acceleration of the vehicles involved in a collision. Thus, since collisions may cause instantaneous jumps in vehicle velocities, the overspeed protector must require that no collisions ever occur in the physical plant; that is, the overspeed and the collision protectors are two-way dependent. Subsequently, it is shown that the two-way dependence composition conditions are met by the proposed overspeed and collision protectors and that their composition results in a protection system that guarantees that the vehicles neither exceed the speed limit, nor collide among themselves.

1.3 Thesis Overview

In order for this thesis to be self contained, Chapter 2 gives a short and terse treatment of the hybrid I/O automaton model [30] and describes the conventions used in the specification of HIOA in this thesis. In order to facilitate the modeling of complex system properties, we introduce notation to allow the explicit restriction of the states of a hybrid I/O automaton to sets of states that are comprised of all states satisfying complex state properties of the HIOA. In Chapter 3, we present an abstract model of a physical plant interacting with various protection systems. Both the physical plant and the protection systems are modeled as hybrid I/O automata. Provided that protectors are independent, they can be composed and their composition guarantees the conjunction of the safety properties guaranteed by the individual protectors being composed. Under certain conditions, the same applies for the composition of protectors that rely on the correct operation of each other. The abstract protector is defined as the composition of a sensor automaton and a discrete controller automaton. The sensor samples the state of the physical plant at regular intervals of time and the discrete controller issues protective actions so as to guarantee that the physical plant exhibits a particular safety property.

In subsequent chapters, we present a simple model of the PRT 2000™ and introduce overspeed and collision protectors. This is done for increasingly complicated track topologies. First we consider a single track, then a Y-shaped merge, and, finally, a general track topology comprised of Y-shaped merges and diverges. Chapter 4 defines a system of n vehicles traveling on a single track and Chapters 5 and 6 define its overspeed and collision protectors. Chapter 7 extends the model of the physical plant to involve a Y-shaped merge and defines a collision protector for the new model. Chapter 8 augments the model of the physical plant to involve a general track topology comprised of Y-shaped merges and diverges and defines a collision protector for the new model. In Chapter 9, we prove that the overspeed and collision protectors of the various track topologies can be composed so as to guarantee that the vehicles neither exceed the speed limit, nor collide among themselves. Finally, in Chapter 10 we give a summary of the thesis, an evaluation of the research presented, and directions in which such research could be extended or continued.