Language-Driven System Design

S. Mauw, W.T. Wiersma, T.A.C. Willemse

Department of Mathematics and Computer Science,

Eindhoven University of Technology,

P.O. Box 513, NL–5600 MB Eindhoven, The Netherlands.

sjouke@win.tue.nl, w.t.wiersma@stud.tue.nl, timw@win.tue.nl

Abstract

The concepts and ideas of domain-specific languages (DSLs) have been well studied over the past decades. Most studies show significant benefits of the use of DSLs over traditional programming techniques, e.g. increased reliability and maintainability of the produced software and better support of reuse of code.

The use of DSLs and their benefits have been well documented. However, there is only little literature available on the relation between software engineering methods and models on the one hand and DSLs on the other hand. Second to that, the development of a DSL is often an ad-hoc process, for which there seems to be no default methodology. Our ambition is to eventually come to a methodology for the construction of DSLs and their use in a software engineering process.

We elaborate on the role a domain-specific language plays, or should play, in the software engineering process. We give an outline of an approach to the software engineering process that is in line with customary software engineering methods, models and best practices. We refer to this approach as the *language-driven approach*. The language-driven approach serves as a first step in the development of a real methodology.

The language-driven approach combines several areas of expertise. First, it relies heavily on techniques used for domain analysis. Second, it combines techniques that are used in software engineering. A domain analysis must be conducted to obtain the basic concepts and notions that are needed to describe a set of problems. The design of a language (and the possible co-design of a software product) rely heavily on techniques that are defined and used for software engineering, such as formal methods and programming paradigms. This is motivated by the fact that a language consists not only of a syntax, but also of a semantics and its pragmatics.

We focus on the language-driven approach and we illustrate the approach with a small and instructive case study.

1 Introduction

The complexity of software has steadily increased over the past decades. This necessitated the development of techniques to master the difficulties and problems due to this increase. Over the years, several process models have been introduced, for structuring the design process of software.

Many different variants of these process models exist, e.g. Boehm's spiral model [2, 3], and the incremental model [22].

Although the differences between these models can be large, all models prescribe a partitioning of the software engineering process into a number of stages. These stages are distinguished on the basis of the activities that have to be conducted. The order in which these stages must be addressed, along with the deliverables that can be expected in each stage are prescribed by the process model.

In theory, most process models are fairly general with respect to the means that must be used to obtain the deliverables in each stage, although some of the process models define best practices for a number of stages. In practice, however, the tools that are used are often determined by external influences, such as a company's policies. This traditional approach to software engineering focuses mainly on a software product that must be developed. Alternatively, the focus could be on a language (or a class of languages) that is tailored to the software product. These languages are often referred to as *Domain-Specific Languages*.

Domain-Specific Languages (DSLs) have emerged as a tool for tackling the complexity of software development projects. Many studies (e.g. [8]) have shown significant benefits of using DSLs in software development. Noteworthy are the increase in the reliability and the maintainability of the produced software, but also improved reusability of a software product's code and design (see e.g. [8, 11]).

Although the use of DSLs and their benefits have been well documented, there is only little literature available on the relation between process models and software engineering methods on the one hand and the use of DSLs on the other hand. Moreover, developing a DSL still seems to be an ad-hoc process, rather than a clearly defined process with a clearly defined methodology. In order to support the acceptance of the ideas and concepts of DSLs, a clear methodology needs to be defined. This methodology must focus on the aspects for developing a DSL, with a clear emphasis towards the intended application of the DSL for a specific problem domain.

In this paper, we discuss the *language-driven approach* to software engineering. This approach can be considered as a first step to a full methodology for designing a DSL. The emphasis of the approach is on the interplay between standard software engineering methods and its best practices, various process models and the concepts and ideas behind DSLs. The key issue in this approach is the focus on the development of a suitable DSL for writing (part of) a software product, rather than the development of the software product itself.

The language-driven approach combines and extends well-known and accepted methods from software engineering. It inherits techniques and concepts from the area of *formal methods* (in its broadest sense), but also the basic ideas and notions behind various *programming paradigms* are incorporated. Moreover, the language-driven approach relies heavily on the expertise and the techniques that are needed to conduct a *domain analysis*. Also for these techniques, there is ample literature available (e.g. [17]).

A major reason for incorporating techniques from specialist areas such as formal methods is our firm believe that a language consists not only of a syntax, but also of an (unambiguous) semantics. Moreover, every language has its own pragmatics that needs to be clear to its users. The techniques that have been studied and developed in the area of formal methods are essential in analysing and defining an understandable language, its syntax, its semantics and its pragmatics.

The tool ASF+SDF [12, 21], for instance, can be used to define and analyse the semantics of programming languages. Apart from this, the use of a formal semantics for validating programs is a key issue in formal methods' research.

The traditional scalability problems often encountered when applying formal methods in the design of software are not likely to be an issue in the language-driven approach. The problem of scalability is often caused by the large gap between the methods that are used to describe a software product and the key concepts of the software product. For the language-driven approach, this gap is relatively small, as the software product is defined in terms of its natural concepts.

In this paper, we introduce and discuss the phases of the language-driven approach. In each phase, we mention the deliverables needed and the techniques for producing them. Additional information is given for selecting alternative techniques or paradigms. This is needed if special requirements are posed on the deliverables, like the need for formal verification.

The abstract ideas in this paper are illustrated with a case study. The case study discusses the design of a language for controlling traffic lights at a junction.

This paper is organised as follows. Section 2 describes the language-driven approach in detail. In Section 3 we discuss related work and research that has already been conducted in this area. The case study is discussed in Section 4. In Section 2, we sometimes refer to the case study to exemplify some of the more abstract notions we encounter.

2 The language-driven approach

In this section we discuss and elaborate on the ingredients that play a role in the language-driven approach. These ingredients can all be found in literature. Wherever possible, we provide pointers to the literature. In our discussion, we emphasise on the formal aspects of the design approach.

2.1 Overview and rationale

In many ways, the language-driven approach resembles a standard software development process. However, there are some differences. These differences are due to the fact that in the language-driven approach, the design is centred around the development of a (formal) language. The language itself, which will be a Domain-Specific Language (DSL), constitutes the major result of the design process. Moreover, the centre of activities in the software design process shifts to earlier stages, such as the user requirements and specification phases. This has several well-known advantages, e.g. reduction of the time-to-market, early detection of errors, etc.

The development of the DSL is described by a collection of deliverables. These deliverables include the definition of its syntax and semantics, and define appropriate tool support. The language-driven design approach can be integrated in today's software process models, such as the waterfall model [19] or the spiral model [2, 3]. Using proven software process models assists the design of the language in a structured way.

Figure 1 shows an overview of the artifacts produced during the development process. We refrain from using a specific process model for the development of the deliverables. In practice, the final product will be the result of several iterations of the development process.

Due to the inherent causalities and dependencies between the deliverables, a natural ordering

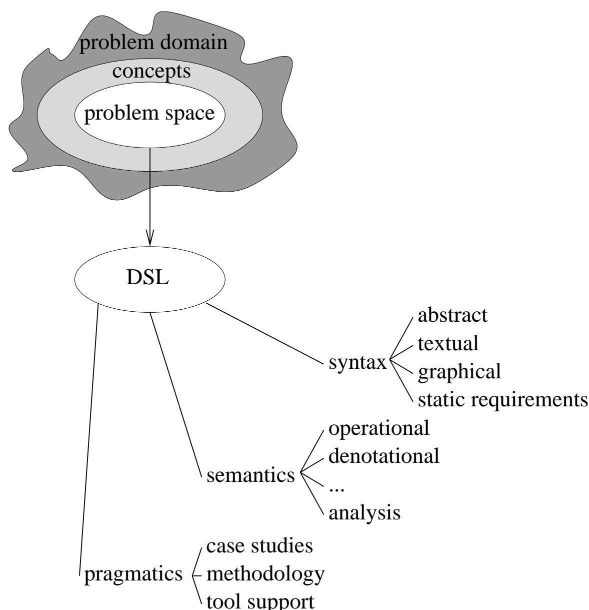

Figure 1: *The language-driven approach.*

is imposed on their development. This ordering is made explicit in the distinction between the following stages in the design process:

1. identification of the *Problem Domain*,

2. identification of the *Problem Space*,

3. formulation of the *Language Definition*.

Note that due to the iterative nature of many process models, the actual order in which these stages are addressed is not fixed, even though there is a clear and intended dependence between the stages. Also, these stages should not be considered to be *atomic*, i.e. it is possible to start a parallel trajectory on the next stage if sufficient information from another stage is available.

In the subsequent sections, the three stages are explained in greater detail. We occasionally refer to the traffic light case study described in Section 4 to explain some of the more abstract notions discussed in these sections.

2.2 Identification of the Problem Domain

The identification of the problem domain is the first stage in the language-driven approach. Rather than focusing on the single problem that needs to be investigated, the language-driven approach focuses on a class of problems stemming from a common problem domain. A thorough domain analysis is necessary to give a complete and precise definition of all essential concepts in the problem domain. Fortunately, there are many existing techniques that support a domain analysis and domain specification. A proper demarcation of the problem domain is vital, as all subsequent artifacts depend on the concepts captured and described by the problem domain. We refer to [17] for an overview of best practices and techniques for conducting the domain analysis.

The problem domain can be (and usually is) much larger –both in generality of the concepts and in the number of concepts– than is strictly needed to solve the actual problem. This has several advantages, e.g. reuse of domain knowledge and self-containment of the problem domain. A restricted problem domain often implies that some design decisions have alread been made. An example of the concepts that are revealed in the domain analysis of the traffic light case study is given below.

Example 2.1. In the traffic light case study, we describe the domain model by means of basic mathematical constructs, such as sets and relations. Typical concepts in our case study are plain entities like *road user* and *lane* and relations like *conflict*, which describes which lanes may have potentially conflicting traffic.

2.3 Identification of the Problem Space

The second stage of the language-driven approach is the identification of the problem space. As already remarked, the previously determined problem domain is a mostly exhaustive collection of all concepts used and related to the actual problem. As mentioned, this has definite advantages; however, for solving the actual problem the problem domain is often too general and too large. Therefore, a restriction of the problem domain is necessary.

A first observation is that in order to provide a direction for solving the actual problem, *design decisions* must be made. These design decisions possibly lead to concepts that have not been identified in the problem domain. The concepts thus introduced play a pivotal role in solving the actual problem. The fact that these concepts are not part of the problem domain follows straightforwardly from the fact that the identification of the problem domain is not a design-driven activity.

A second observation is that with respect to the actual problem, the problem domain contains some inherent redundancy. Whenever concepts do not appear to play any part in solving the actual problem, we can consider them irrelevant for solving this problem. Hence, we only need to consider the concepts that are relevant to our problem. The last observation is that with respect to the actual problem, many of the identified concepts are too general. Therefore, a natural second classification of the concepts is to make a distinction between the concepts that can be constrained in some sense, and the concepts that are inherently variable.

These observations lead to a classification of all concepts into the following three categories:

- concepts that are *irrelevant* to the actual problem,
- concepts that are *variable*, and
- concepts that have been *fixed* for the actual problem.

As already remarked, a concept is *irrelevant* whenever it does not play any part in the solution to the actual problem. Moreover, concepts can often be classified as irrelevant due to abstraction and aggregation.

A concept is called *variable* whenever it varies depending on the actual problem instance, or it varies within the actual problem instance. The variable concepts that vary depending on the actual problem instance can often be considered as problem parameters; every (allowed) instantiation of

the problem parameters calls for its own solution. These problem parameters are the part of the actual problem which will be specified by means of an expression in the DSL. As a result, this collection of variable concepts determines the syntax of the language. The variable concepts that vary within the actual problem determine the behaviour of the system. In an operational semantics, these concepts reappear as a part of the state space. They cannot be specified by means of expressions in the syntax.

The category of the *fixed* concepts consists of the concepts which are identical for all problems considered. This can be caused by the fact that the notion is inherently constant (e.g. a law of nature), but more often it concerns a variable notion which is restricted to simplify the problem setting.

The class containing the variable concepts and the fixed concepts is referred to as the *problem space*. Obviously, this part of the problem space is more concrete compared to the problem domain. This, however, reduces the complexity of the basic notions that are relevant to the actual problem, while still retaining enough information to describe the actual problem accurately. Since the notions discussed in this section are rather abstract, the four types of concepts are exemplified:

Example 2.2. In our case study, we have introduced the concept of *priorities*. This concept is derived from our desire to model the traffic lights as a competitive system. However, the notion of priorities is not a notion that could have been derived during the identification of the problem domain. Note that different design decisions might have led to the introduction of other concepts.

The concept of a *road user* turns out to be irrelevant in our case study. Although a road user plays a vital role in the domain analysis, it does not re-occur in any of the subsequent phases. This is because road users are not important to the goals set for solving our actual problem; their presence can only be detected indirectly by sensors.

An example of a variable concept, is the notion of a *conflict matrix*. The conflict matrix describes the actual situation at a particular traffic intersection. In order to deal with more than one fixed intersection, we require this concept to be variable. This means that this concept also reappears in the syntax of the language, as it is needed there to describe the intersection in terms of its conflicts. An example of a concept which is variable within a problem instance is the *current colour* of a traffic light. This will change during operation and is, therefore, included in the state of the system.

When we fix the order in which a traffic light displays its colours this concept can be considered as an example of a fixed concept. Not determining the order of the colours in advance would support any ordered list of colours and therefore would be much more general. However, this would also increase the complexity of the syntax, because it would need constructs for specifying the order of the colours. Since the order for traffic light colours is more or less standardized, fixing the order is not a severe restriction.

2.4 Formulation of the Language Definition

The design of the language concretises the notions and concepts that can be found in the problem domain. In this section, we discuss the constituent parts of a DSL. We advocate a formal treatment for the specification of both the syntax and semantics of a DSL.

The language definition stage is subdivided into three sub-phases in which the *syntax*, the *semantics* and the *pragmatics* of the language are defined.

2.4.1 Syntax

The appearance of a language is defined by means of its *syntax*. In the language-driven approach, the constructs of the language are related to the concepts that have been identified in the domain space. The syntax of the language consists of expressions of the variable concepts that have been identified in the problem space. Unlike the variable concepts, fixed concepts are not defined in the syntax.

The syntax serves several purposes. It supports the user in expressing the properties of the problems the user wants to solve using the language. Second, the semantics is based on the syntactical expressions. Moreover, the language constructs serve as a basis for applying analysis techniques on both the language and the problems described using the language.

The format of the language is constrained in several ways. Most importantly, it must be susceptible to interpretation and/or transformation by means of a computer. Moreover, the syntax is often constrained by a number of generally accepted requirements such as readability and writeability of the language constructs. Finally, we stress the importance of choosing syntax expressions for which the mathematical semantics correspond to the intuitive semantics.

In general, a language can have one or more syntactical descriptions. These descriptions depend on the required use of the language. Three of the more popular formats are the following:

- the *abstract* syntax,
- the *textual* or *linear* syntax,
- the *graphical* syntax.

The *abstract* syntax is often used to express all semantically relevant information in a minimal way (e.g. without keywords or superfluous transitions in the defining grammar). Moreover, the data structure that is used by computers to store the information that is obtained while processing programs is often strongly related to the abstract syntax. The expressions in the abstract syntax are generally not meant for human processing or usage, but they are useful during design of the language.

The *textual* or *linear* syntax is the description that is most often encountered in language descriptions. The information in the textual syntax is essentially the same as in the abstract syntax. However, the textual syntax is easier to read and use. This is best exemplified by constructs such as the *if-then-else* construct. This construct will have all the appropriate keywords in the textual syntax, but in the abstract syntax it will simply be a triplet.

The third format, the *graphical* syntax, is gradually gaining popularity. Graphical, or *visual* languages have several definite benefits over linear languages, such as the ability to express spatial properties or complex relations in a more intuitive fashion. The general availability of graphical workstations makes it possible for regular users to work with visual languages. Although the graphical syntax may seem capable of expressing more than the textual or abstract syntax, the semantically relevant information should be identical.

The abstract syntax and the textual syntax can be partly defined by means of BNF grammars, or BNF-like grammars (i.e. BNF grammars enhanced with simple (mathematical) structuring mechanisms such as sets or records). For the graphical syntax, no generally accepted format for defining the language exists. The most popular way of defining the graphical language is by means of *graph grammars* [18].

In most cases BNF-like grammars are not expressive enough to exactly describe which expressions in the language are *well-formed*. Context-sensitive properties, such as the *declare-before-use* property of variables, must be expressed in a different way. These additional requirements on well-formedness of expressions are often referred to as the *static semantics*. This title is somewhat misleading, since it deals with syntactical properties of the language. Therefore, we prefer to use the term *static requirements* for this purpose. Most often, attribute grammars [13] are used for specifying the static requirements, but logical predicates can also be applied. The notion of static semantics also has a different interpretation, namely the semantics of the static (i.e. non-behavioural) part of the language.

To conclude this section on syntax, we mention that there are several other syntactical aspects which can be specified. A requirement that is often posed on expressions in a graphical languages is to have a layout that is transparent to tools. It is not likely that the detailed layout will have any semantical meaning, so one may not expect that the textual syntax is capable of expressing such properties. This is resolved either by extending the textual syntax with information that is semantically irrelevant, or by defining an additional syntax which is tailored to expressing such details. The latter approach is often called a *tool interchange format* (see e.g. the Common Interchange Format CIF for the SDL language [10]).

Not all three formats may be necessary: one might skip e.g. the textual syntax. Two issues must be kept in mind, namely, that there has to be a syntactical representation which covers all semantically relevant issues and that a formal definition of the syntactical ingredients should be given.

2.4.2 Semantics

A semantics for a language is a mathematical model that reflects the intended computational behaviour of expressions in the language. In essence, one can classify the characteristics of a language in two classes:

- Language components dealing with dynamic behaviour,
- Language components describing purely static information.

This distinction is also reflected in the semantics of the language.

As for general-purpose languages, various approaches exist to defining a semantics for a language. The choice of a suitable semantical approach depends largely on the characteristics of the language itself, i.e. the class to which the language belongs. However, the practical use of the semantics is important as well. Dependent on the type of semantics, techniques such as behavioural analysis, invariant analysis or simulation of expressions in the language can be used. Most designers of a language are biased towards certain approaches.

Basic to most semantical approaches is the existence of a semantical domain. Such a domain often consists of a set (or collection of sets) with additional structure defined by relations. Expressions in the language relate to entities in this domain and obtain their meaning via the properties of the related entities.

Although the different approaches are all variations on a similar theme, each of these approaches emphasises on a different aspect and has its own benefits. We subsequently give a short overview of the main advantages of these approaches in the next paragraphs.

Operational semantics Operational semantics is used to give meaning to the dynamic part of a language. It is centred around the notions of a state and the transitions between the states (see e.g. [7]). The transitions between the state can be described by means of a transition function. Various ways exist for defining the operational semantics, e.g. by means of SOS-rules [16].

The operational semantics of a language is quite close to the intuition behind the language. It is often used by implementors. An operational semantics provides the means for performing simulations of expressions in the language by considering runs of the transition function. This is useful in areas of testing, or even automated testing. Moreover, there is also the possibility of analysing the transition graph that is induced by a language expression. This is often used in verification efforts. Finally, tools, such as ASF+SDF [12, 21] or Maude [4] may be used to develop prototypes of the language.

Denotational semantics A denotational semantics is centred around the idea of a mathematical function that describes the meaning of an expression by means of a translation to a well-understood mathematical model, as described in the beginning of this section (see e.g. [20]). Its virtue is the use of this mathematical model for the analysis and comparison of expressions in the language.

The theory of the denotational semantics is mathematically very rigorous. It is often used by language designers, as it precisely expresses the requirements on the language. Techniques to prove two expressions in the language equivalent are easily formulated using the underlying mathematical model. Such techniques can also be automated, using theorem provers.

Axiomatic semantics The axiomatic semantics is given by means of a number of axioms relating expressions in the language. The axioms can be based on some underlying logic. The axiomatic semantics is often used in combination with a denotational or operational semantics to provide for an underlying mathematical model and a suitable notion of equivalence.

The axioms defining the semantics of a language provide the possibility to interpret the axioms as a set of rewrite rules. This allows for rapid prototyping of the language. Moreover, based on the axiomatic system, there is an option for theorem proving. Examples of an axiomatic semantics are the pre-and post conditions used for programming languages [9] or the use of axioms in the context of process algebras [1].

The semantics of the language are mostly defined on the abstract syntax representation. In case the language has a graphical syntax, its semantics can be defined directly on the graphical

syntax, but it is often more convenient to define a mapping from the graphical syntax onto the abstract syntax and formalise the semantics of the latter.

As may be expected, the definition of a (formal) semantics is crucial to unambiguously understand the programs and to define analysis techniques, together with proper support tools. More than one semantics may be defined, as long as these are consistent.

Analysis techniques are used for the semantical analysis of expressions in a language. We consider these techniques as part of the semantical development of the language, as these techniques are largely dependent on the choices made in defining the semantics of the language. The analysis techniques provide an increased insight into the meaning of possible expressions. Moreover, correctness of the language is better understood by determining the properties of expressions in the language.

Often, the analysis techniques follow some standard mathematical approach. However, it is conceivable that new theory needs to be developed for performing some desired analysis.

2.4.3 Pragmatics

The pragmatics of a language deals with all aspects of the use of the language. Obviously, a language design is not finished without guidelines on how to properly use the language. A collection of examples may show the application of typical features, case studies will prove usefulness for real examples. Moreover, documentation, including tutorials and educational material, together with rules of thumb, etc. are needed to advocate the proper use of the language. These guidelines are called the *methodology* of the language.

Apart from the methodology of the language, tools need to be defined for interpreting or compiling the language, and to support the analysis of programs written in the language. Ideally, these tools should follow from the semantical definitions. For instance, an interpreter of a language needs to show exactly the behaviour described by the operational semantics. Several *meta-tools* support the generation of a parser based on the formally defined syntax and as already mentioned, the generation of an interpreter based on the operational semantics is also viable (see e.g. [12, 21, 4]).

3 Related Work

There are many publications describing the development of some specific DSL or which describe a set of (meta-) tools to support such development. However, there is only little literature on methodological aspects of the design of domain-specific languages. We discuss some relevant work below.

In [5], Consel and Marlet describe a methodology for developing DSLs. It relates two orthogonal perspectives (a programming language perspective and a software architecture perspective) and describes a staged development of DSLs. The methodology is based on the formal framework of denotational semantics, and uses techniques to obtain dedicated abstract machines from the denotational semantics of a language.

Our approach resembles the methodology outlined in [5]. However, our approach does not use

the denotational semantics as a formal framework. Instead, dependent on the desired use of the language, we allow for different methods for specifying the semantics, e.g. operational semantics. Moreover, the methodology outlined in [5] does not recognise the importance of the *problem space* definition, and assumes a proper domain analysis has already been conducted to obtain the basic concepts that are needed to come to the notion of a *problem family*. Finally, the implementation of the language receives much attention in [5], whereas this issue is of lesser importance in our approach, since our focus is not restricted to executable languages.

Montages and its graphical tool environment Gem-Mex (see [14]) form a suite for describing several aspects of programming languages, such as syntax, static analysis and semantics, and dynamic semantics. Syntax is described by BNF rules, and Abstract State Machines (formerly known as evolving algebras) are used to define the semantics of a language. The system is able to generate a visual programming environment for the specified language. The Montages methodology has no support for domain analysis.

In [6], Gupta and Pontelli start reasoning from the observation that any software system can be understood in terms of how it interacts with the outside world. Thus, every system is in essence defined by its input language, which in turn can be considered a domain-specific language. They use Horn logic to give a denotational definition of such DSL, which automatically yields a parser, an interpreter and tools to support verification. The focus of their research is on applying Horn logic for these purposes, without developing a more generally applicable methodology.

Pfahler and Kastens [15] discuss issues related to the maintenance of a DSL. Rather than updating a language by going through a new language development cycle again, they propose to develop DSLs in such a way that small maintenance can be performed easily. Thereto, they consider a language design based on a collection of components, which can be glued together in different ways, thus making for a more flexible language definition, or rather a language family. This DSL life-cycle is called the *Jacob* approach. Corresponding tool support makes it possible to automatically generate substantial parts of an implementation. The authors do not describe a methodology for designing a language family (i.e. an appropriate set of components). We expect that the methodology outlined here will also be applicable to language families.

4 The Case Study

We illustrate the language-driven approach by developing a domain-specific language for the regulation of traffic lights. The case study is discussed in great detail in the subsequent sections.

The problem deals with traffic junctions and the traffic passing the junction. We can distinguish between traffic junctions that do not need any control and traffic junctions that do need control. The former are often traffic junctions that have only little traffic passing it, whereas the latter are often junctions that have many conflicting traffic streams. Regulation of traffic streams is done by means of traffic lights and division of roads into lanes.

A standard approach to controlling these traffic lights is to fix an order in which these traffic lights allow traffic to cross the junction. This, however, leads to sub-optimal throughput, traffic congestion, etc. To overcome such problems, sensors are used that register the presence of traffic per lane. The sensors' information is the basis for the order in which these traffic lights allow traffic to cross the junction. Notice that the addition of sensors renders systems that can respond to

events from their environments, i.e. the traffic light controllers we consider are dynamic, reactive systems.

In order to cope with high-priority vehicles (e.g. police vehicles), special care must be taken to make sure these vehicles are allowed to cross the junction as soon as possible. However, it is not allowed to have unsafe situations at the traffic junction at any moment in time. Hence, conflicting traffic streams are not allowed to cross the traffic junction at the same time. Moreover, we cannot *a priori* assume that a traffic stream has cleared the junction immediately after a traffic light has changed to red. Therefore, to each traffic light a *clearance duration* is associated. The clearance duration of a traffic light is the time that is needed to clear the junction from traffic. In order to prevent a traffic light from switching colours too fast, we associate a minimal duration to each colour of the traffic light.

The goal is to obtain autonomous traffic junction regulators that are more efficient than the controllers defined by the standard approach and still guarantee safety. To achieve this goal, we develop a DSL that is tailored to the control of traffic lights as we envision it.

Our presentation of the case study will be in a linear way. However, it must be noticed that the results described in this section are the product of several iterations. In our presentation of the deliverables for our case study, we closely follow the order prescribed by the language-driven approach. In Section 4.1 we focus on the problem domain. The concretisation of the problem domain into the problem space is discussed in Section 4.2. Finally, in Sections 4.3 to 4.6, the language definition is presented.

4.1 The Problem Domain

The first step of the language-driven approach to software engineering is a proper identification of all concepts that are essential in the problem domain. As most people are familiar with traffic junctions, obtaining an initial set of concepts is rather straightforward (e.g. by means of a brainstorm session and interviews). We assume that the introductory text in the previous section is sufficient for a basic understanding of the problem domain.

The concepts that are a natural consequence of the characterisation of a traffic junction, made in the previous section, are discussed in the subsequent sections. Notice that we have already marked the concepts (using ^{f} for *fixed* and ^{v} for *variable*) that are part of the problem space, as to avoid duplication of information. In Section 4.2 we provide a motivation for our choice for these concepts.

Time. A traffic light controller is a dynamic, time-dependent system. The concept of time is therefore indispensable. There are two different types of time, i.e. relative time and absolute time. For traffic light controllers, both types of time are possible. Relative time is usually employed if one wants to refer to periods of time (e.g. a traffic light must be green for at least ten seconds). Absolute time is used to model that events must occur at specific moments (e.g. a traffic light is out of order until May 1, 2001). Therefore, we introduce the following concepts:

1^f.	*Duration*	Lapse of time (relative)
2^v.	*Time*	‘Calendar’ time (absolute)

Participants. As mentioned before, the traffic light controller is a reactive system. It responds to the events it receives from its environment. These environmental events are triggered by traffic participants, i.e. road users. These road users can be of a specific type, e.g. car, pedestrian, train, etc.

1.	*Roadusers*	Set of all possible traffic participants
2.	*T_Roadusers*	Set of all possible types of traffic participants
3.	*UserType*	$: Roadusers \rightarrow T_Roadusers$

Different views on junctions. One of the most natural concepts of our problem domain is the concept of a junction. For junctions, we can recognise three levels of concreteness: the physical topology of the intersection, the traffic rules that apply to the intersection and the logical characteristics of the intersection. These three levels are explained in greater detail below.

The physical topology of the intersection consists of several crossing roads. Roads can be divided into a number of lanes. We define lanes as stretches of road that have identical behaviour (lanes can also be sidewalks or rail tracks), i.e. a lane consists of a number of (parallel) strips. The users of a lane are supposed to follow the same route (or set of routes) on an intersection. For traffic junctions, we consider two types of lanes, viz. lanes entering and lanes leaving an intersection. Since we are interested in traffic crossing an intersection, we must consider the possibilities for doing so. From the perspective of the physical topology, we arrive at the notion of possible continuations for every lane entering an intersection.

1^v.	*InLanes*	Set of all traffic lanes entering the intersection
2^v.	*OutLanes*	Set of all traffic lanes leaving the intersection
3.	*Lanes*	$= InLanes \cup OutLanes$
		requirement: $InLanes \cap OutLanes = \emptyset$
4.	*PossibleContinuations*	$: InLanes \rightarrow (\mathcal{P}(OutLanes) - \{\emptyset\})$
5.	*PossibleLaneUsers*	$: Lanes \rightarrow \mathcal{P}(T_Roadusers)$

Observing the traffic laws that hold for an intersection, we see that these laws restrict traffic in an essential way. Rather than considering all possible continuations of a lane entering an intersection, we should in fact consider a subset thereof. This is motivated by the fact that, although the physical possibilities are there, the law forbids these continuations. We thus arrive at the notion of continuations. Note that a lane entering an intersection must always have at least one continuation.

6.	*LaneUsers*	$\subseteq PossibleLaneUsers$
7^v.	*Continuations*	$\subseteq PossibleContinuations$
		requirement: $\forall_{a \in InLanes, b \in OutLanes}$
		$b \in Continuations(a) \Rightarrow LaneUsers(a) \subseteq LaneUsers(b)$

From a logical point of view, the intersection can still exhibit unsafe behaviour. This unsafe behaviour has two causes. On the one hand, traffic entering the intersection via one lane can be in conflict with traffic entering the intersection via another lane. This conflict is dependent on the physical location of the lanes and the continuations of lanes entering the intersection. In order to reason about such lanes, we describe which lanes are conflicting, i.e. which lanes cannot simultaneously have a green light. Such a conflict relation is often called a *conflict matrix*.

On the other hand, we can observe that it takes some time for a traffic stream to clear the intersection after it has received a red light. This period needs to be taken into account in order

to guarantee safety. We refer to this period as the clearance duration. Clearance duration is a binary function on the lanes entering and the lanes leaving an intersection. We can consider a more abstract notion of clearance duration, i.e. one that determines for an incoming lane the maximum clearance duration over all its continuations.

8.	*Conflict*	$\subseteq (Inlanes \times Outlanes)^2$
		requirement: *Conflict* is symmetric and irreflexive
9^v.	*Conflict*	$\subseteq InLanes^2$
		where *Conflict* is derived as:
		$\{(i_1, i_2) \mid \exists_{o_1 \in Continuations(i_1), o_2 \in Continuations(i_2)}\ Conflict((i_1, o_1), (i_2, o_2))\}$
10.	*ClearanceDuration*	$: InLanes \times OutLanes \to Duration$
		requirement: *ClearanceDuration* is a partial function
		defined on all (i, o), such that $i \in InLanes, o \in Continuations(i)$
11^v.	*ClearanceDuration*	$: InLanes \to Duration$
		where $ClearanceDuration(l)$ is derived as:
		$\max\{ClearanceDuration(l, o) \mid o \in Continuations(l)\}$

Traffic lights. Various important characteristics of traffic lights can be identified. A main characteristic is the set of colours the traffic light has. A traffic light usually changes colour in a fixed order, i.e. a notion of state cycle can be identified. The state of a traffic light is tightly coupled to the traffic light itself, i.e. a *current* state can be identified. We are also interested in how long the light is already in this state. Traffic lights are often required to be in a state for a minimum time (e.g. a traffic light is required to show a green light for at least three seconds).

1^f.	*TL_State*	Non-empty set of all possible traffic light states
2^f.	*StateCycle*	$\in TL_State^+$
3.	*TLights*	Set of all traffic lights
4^v.	*CurrentTLightState*	$: TLights \to TL_State$
5.	*TL_Loc*	$: TLights \to InLanes$
6^v.	*MinStateTime*	$: TL_State \to Duration$
7^v.	*CurrentDuration*	$: TL_State \to Duration$

Sensors. To obtain information about their environment, traffic lanes must be equipped with sensors. When a sensor is triggered, it produces an input event that changes the state of that sensor. Thus, sensors have a notion of state. Moreover, at each moment in time, we can inspect the state of a sensor, hence, we can identify the *current* state for sensors.

Sensors can be placed at lanes for detecting specified types of road users. This is convenient for detecting speeding ambulances or police vehicles.

1^v.	*Sensors*	Set of all sensors
2^f.	*SensorState*	Set of all possible sensor states
3^v.	*CurrentSensorState*	$: Sensors \to SensorState$
4^v.	*SensorLoc*	$: Sensors \to InLanes$
5.	*SensorRecog*	$: Sensors \to \mathcal{P}(T_Roadusers)$

4.2 The Problem Space

An important step in the language-driven approach is the identification of the problem space. As mentioned before, the problem space is both a restriction of concepts of the problem domain and an extension of the problem domain with concepts due to design decisions. The restriction of the problem domain is discussed in Section 4.2.2. First, the design decisions (i.e. the extensions of the problem domain) are discussed in Section 4.2.1

4.2.1 Design Decisions

We will model the traffic lights as a *competitive system*. This means that every traffic light competes with other lights for the right to change colour. So we have to keep information local to the traffic stream to reach a global decision which lights can change colour. For this reason we introduce the concept of assigning *priorities* to traffic streams. These priorities can dynamically change, based on the progress of time or the detection of traffic. We assume a totally ordered set *Prio* of priority values. Then we have for every sensor the priority value to which the corresponding lane will be initialised if traffic is detected by that sensor. For sensor s, this will be denoted by *InitPrio(s)*. Finally, we have a priority update function, which determines the new priority value of a lane after the elapse of one time unit. This function will be denoted by *UpdatePrio*.

1^f.	*Prio*	Totally ordered set of priority values
2^v.	*InitPrio*	: $\mathit{Sensors} \rightarrow \mathit{Prio}$
3^v.	*UpdatePrio*	: $\mathit{InLanes} \times \mathit{Prio} \rightarrow \mathit{Prio}$
4^v.	*CurrentLanePrio*	: $\mathit{Inlane} \rightarrow \mathit{Prio}$

4.2.2 Reduction of the Problem Domain

The concepts that are irrelevant to traffic light control are the concepts of Section 4.1 that are not marked with ^{f} or ^{v}. The concepts that turn out to be relevant, but can be fixed have been marked with ^{f}, whereas the concepts that need to be variable are marked with ^{v}. In this section, we restrict our discussion to only a few examples of irrelevant, fixed and variable concepts.

Irrelevant concepts. In our problem domain we have identified the concept of *Roaduser*. This concept is irrelevant to our actual problem, since we have decided to build a system in which individual road users do not play a role. They can only be detected indirectly by a sensor. They might play a role, however, in case one of the goals was to build a simulator showing behaviour of individual road users.

A second example of an irrelevant notion is *TLights*. Although this notion is at the right level of abstraction, we observe that it would not be a severe restriction if there is exactly one element of *TLights* for every element of *InLanes*. Therefore, we can simply identify these two notions and discard *TLights*. Henceforth, we will speak, for instance, of the colour of a lane, instead of the colour of a traffic light.

Variable concepts. The variable concepts that depend on the actual problem instance can be defined using the syntax. An example of such a variable concept is the conflict matrix (i.e. the concept *Conflict*). In our goal to describe traffic light control for more than a single fixed traffic junction, we need to take the conflict matrix into account. This is due to the fact that, dependent on the junction, the conflict matrix can differ. Hence, fixing the conflict matrix would be unwise, as it would restrict our language to describing only junctions with identical conflicting traffic streams.

Another example of a variable concept is the concept of clearance duration. As we have seen, the notion of clearance duration is important to guarantee safety of the traffic junction. Hence, the concept cannot be considered irrelevant. If we consider this clearance duration as a fixed concept, then we restrict our language to describing intersections that have a single (fixed)

clearance duration for all traffic streams. This, of course, is too restrictive. Hence, the concept of clearance duration must be defined as a variable concept and can thus be defined using the syntax of the language.

For the other variable concepts (*InLanes*, *Continuations*, *Outlanes*, *MinStateTime*, *Sensors*, *SensorLoc*, *InitPrio*, and *UpdatePrio*), a similar reasoning holds.

An example of a variable concept that is variable within the problem instance is the concept *CurrentSensorState*. This concept defines the relation between the concepts of *Sensors* and the concepts of *State*. This relation, however, is not static, as the state of a sensor can change over time (e.g. by means of traffic passing the sensor). In fact, this relation clearly illustrates the dynamic nature of traffic light control.

Similarly, the concept of *CurrentTLightState* is a variable concept.

Fixed concepts. We have fixed several notions to a concrete value in order to make the problem less abstract. First of all, we will restrict the colours that a traffic light can have by defining $\mathit{TL_State} = \{\mathit{green}, \mathit{yellow}, \mathit{red}\}$, which also determines the standard order $\mathit{StateCycle} = \mathit{green} \circ \mathit{yellow} \circ \mathit{red}$. Such a decision may come from the fact that the system is only to be applied in countries where this is the standard order of operation. It might be considered a severe restriction that this also implies that special operation of traffic lights (e.g. a flashing yellow light) is not supported.

For ease of reasoning, we will take $\mathit{Prio} = \mathbb{N}$. We will assume a discrete time domain, and set $\mathit{Time} = \mathit{Duration} = \mathbb{N}$

4.3 Syntax

This section will describe the syntax of the traffic regulation language. We provide a definition of the abstract syntax and we give examples of expressions in the concrete and graphical syntax. Due to space limitations we will not give the syntax definitions in full detail.

The abstract syntax serves to express in a minimal format the semantically relevant information which a designer of an intersection should provide in order to obtain an operational system. The abstract syntax has a clear correspondence with the variable concepts identified in the problem space.

Words between angular brackets, $\langle\ \rangle$, are the non-terminals of the language. We assume that the non-terminals $\langle \mathit{inlaneid} \rangle$, $\langle \mathit{outlaneid} \rangle$, and $\langle \mathit{sensorid} \rangle$ produce disjoint sets of identifier symbols. Furthermore, $\langle \mathit{updateprio} \rangle$ produces a natural expression (possibly containing occurrences of a variable, say x) which represents the priority update function. The initial priority of a sensor is captured by $\langle \mathit{initprio} \rangle$. Non-terminals $\langle \mathit{initprio} \rangle$, $\langle \mathit{clearance} \rangle$, $\langle \mathit{greentime} \rangle$, $\langle \mathit{yellowtime} \rangle$, and $\langle \mathit{redtime} \rangle$ produce a natural numeric constant.

$\langle \mathit{junction} \rangle ::= \langle \mathit{lane} \rangle^* \ \langle \mathit{conflict} \rangle^* \ \langle \mathit{mintime} \rangle$
$\langle \mathit{lane} \rangle ::= \langle \mathit{inlaneid} \rangle \ \langle \mathit{continuation} \rangle^* \ \langle \mathit{sensor} \rangle^* \ \langle \mathit{updateprio} \rangle$
$\langle \mathit{continuation} \rangle ::= \langle \mathit{outlaneid} \rangle \ \langle \mathit{clearance} \rangle$
$\langle \mathit{sensor} \rangle ::= \langle \mathit{sensorid} \rangle \ \langle \mathit{initprio} \rangle$
$\langle \mathit{conflict} \rangle ::= \langle \mathit{inlaneid} \rangle \ \langle \mathit{inlaneid} \rangle$
$\langle \mathit{mintime} \rangle ::= \langle \mathit{greentime} \rangle \ \langle \mathit{yellowtime} \rangle \ \langle \mathit{redtime} \rangle$

As an example of a static requirement defined on this abstract syntax, we will specify the predicate *irreflexive-conflict*. This static requirement follows from the requirement on the conflict matrix as specified in Section 4.1.

$$\textit{irreflexive-conflict}(\textit{lane-list conflict-list mintime}) = \textit{irreflexive-conflict}(\textit{conflict-list})$$
$$\textit{irreflexive-conflict}(\varepsilon) = \textit{true}$$
$$\textit{irreflexive-conflict}(\textit{inlaneid1 inlaneid2 conflict-list}) = (\textit{inlaneid1} \neq \textit{inlaneid2}) \wedge \textit{irreflexive-conflict}(\textit{conflict-list})$$

We have overloaded the predicate name such that it accepts expressions of type $\langle \textit{junction} \rangle$ and $\langle \textit{conflict} \rangle^*$. With ε we denote the empty list. Typing of the other variables follows from their naming scheme.

In the same way we can define auxiliary functions to extract information from the abstract syntax, such as the function *sensors* which determines for each lane l the set of available sensors.

$$\textit{sensors}(l, \textit{lane-list conflict-list mintime}) = \textit{sensors}(l, \textit{lane-list})$$
$$\textit{sensors}(l, \textit{inlaneid continuation-list sensor-list update lane-list}) = \begin{cases} \textit{sensors}(\textit{sensor-list}) & \textit{if } l = \textit{inlaneid} \\ \textit{sensors}(l, \textit{lane-list}) & \textit{if } l \neq \textit{inlaneid} \end{cases}$$
$$\textit{sensors}(\varepsilon) = \emptyset$$
$$\textit{sensors}(\textit{sensorid initprio sensor-list}\ ) = \{\textit{sensorid}\} \cup \textit{sensors}(\textit{sensor-list})$$

This function will be used in the definition of the semantics.

There are many ways in which the abstract syntax can be represented in a more readable format. The textual representation of the example in Figure 2 is a bit more verbose. This example describes a junction with two incoming lanes (a and b) and two outgoing lanes (c and d). Lane a continues at lanes c and d. The clearance duration of the path from lane a to lane c is 3. Lane a has two sensors, called *normal* and *bus*. The initial priority of the normal sensor is 10, while detection of a bus sets the priority to 100. The sensor at lane b cannot make a distinction between the type of traffic detected. The priorities of the lanes a and b are updated every time unit with the update functions[1] $\lambda x.x + 1$ and $\lambda x.x + 2$, respectively. The two lanes a and b have a conflict. Finally, the minimal state time of the traffic light colours is set to 1, 1, 3 (for red, yellow, and green).

We leave it to the reader to interpret the graphical symbols in Figure 2.

4.4 Semantics

In this section we provide an operational semantics for the traffic light control system. After motivating the design of the semantics, we define a *state space* (see Section 4.4.1) and *transition functions* (see Section 4.4.2).

4.4.1 State

In order to give an operational semantics based on state transitions, we will first define the *state* of the system. The state is based on the concepts which are variable within a problem instance, i.e. the priorities of the lanes, the states of the traffic lights, the time that the traffic lights have their current colour, the states of the sensors, and the absolute time. We define state domain $\Sigma = \Pi \times \Gamma \times \Delta \times \Xi \times \mathit{Time}$, where

[1] We will use the notation $\lambda x.f(x)$ to describe a function f with parameter x.

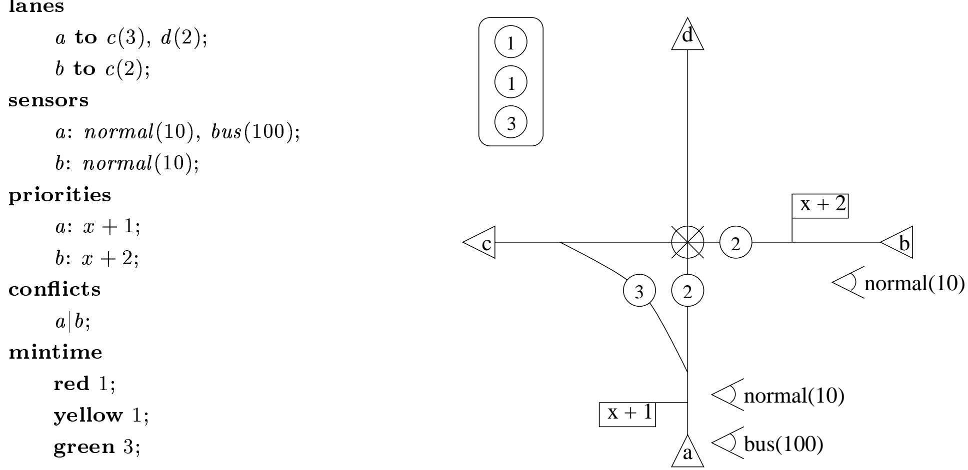

Figure 2: *Example junction in textual and graphical syntax.*

1. $\Pi \quad = InLanes \rightarrow Prio$
2. $\Gamma \quad = InLanes \rightarrow TL_State$
3. $\Delta \quad = InLanes \rightarrow Duration$
4. $\Xi \quad = Sensors \rightarrow SensorState$
5. *Time* is the absolute time domain

Then the state of the system is represented by a five-tuple $\sigma = (\pi, \gamma, \delta, \xi, \tau) \in \Sigma$. Henceforth, we will use the notations σ and $(\pi, \gamma, \delta, \xi, \tau)$ interchangebly. Without making it explicit, e.g. the function γ_2 will denote the second component of state σ_2. The initial state of the system is $\sigma_0 = (\pi_0, \gamma_0, \delta_0, \xi_0, \tau_0)$, where

$\pi_0 = \lambda l.0$
$\gamma_0 = \lambda l.red$
$\delta_0 = \lambda l.0$
$\xi_0 = \lambda s.false$
$\tau_0 = 0$

4.4.2 Transition rules

Following the domain analysis, there are three ways in which the state of the system can be changed. First, there can be an input event (i.e. a sensor changes its state). Secondly, an output event can be generated (i.e. a traffic light should change colour). And thirdly, time can progress. These three events give rise to three types of transitions.

To simplify matters, we assume a slotted operation. By this, we mean that during each time interval first all sensor inputs are collected (if any), then all traffic light outputs are generated (if any), and finally, time progresses to the next time slot. Henceforth, we will consider only transition graphs which satisfy this restriction on the order of transitions.

A transition generated by an input event is denoted by $\stackrel{i}{\rightarrow}$, where $i \subseteq Sensors$ denotes the set of sensors which have detected traffic during the time slot. A transition based on an output event is denoted by $\stackrel{\langle tr,ty,tg \rangle}{\longrightarrow}$, where tr, ty, $tg \subseteq InLanes$ denote the sets of lanes whose traffic lights should

advance to *red*, *yellow*, and *green*, respectively. Finally, $\xrightarrow{t}$ denotes progress of time with one time unit.

Now, we will discuss the three transition rules which define the transition system.

Check for input. We consider an input from a sensor as the indication that traffic has been detected during the just finished time slot. This does not mean that the traffic is (still) waiting. Sensor information is kept in the state variable ξ, which is set to *true* for a given sensor every time that the sensor yields an input. Since the system cannot control its inputs, every possible collection of sensor inputs should be accepted. This is modelled by having a transition for every subset of $Sensors$.

For every $i \subseteq Sensors$, we have the following transition:

$$(\pi, \gamma, \delta, \xi, \tau) \xrightarrow{i} (\pi', \gamma, \delta, \xi', \tau),$$

where

$$\pi' = \lambda l. \begin{cases} \max(\pi(l), \max\{InitPrio(s) \mid s \in sensors(l) \wedge (\xi(s) \vee s \in i)\}) & \text{if } \gamma(l) = red \\ 0 & \text{otherwise} \end{cases}$$

$$\xi' = \lambda s.(\xi(s) \vee s \in i)$$

The priority of lane l is set to the maximum value of the initial priorities from all triggered sensors that belong to the lane. However, if the priority is already larger than the initial priority, the old value remains. The sensor status is changed as explained above. All other components in the state remain unchanged.

Generate output. The transition rule in which the output to the traffic lights is generated looks as follows.

$$(\pi, \gamma, \delta, \xi, \tau) \xrightarrow{\langle TR_\sigma, TY_\sigma, TG_\sigma \rangle} (\pi', \gamma', \delta', \xi', \tau),$$

The sets TR_σ, TY_σ, and TG_σ are defined below in such a way that the lane with highest priority can go first, while allowing non-conflicting traffic of lower priority to pass too.

In order to calculate these sets, we define $M_\sigma \subseteq InLanes$ as the set of all lanes that, based on their priorities, should receive green light. Possibly, not all lights from M_σ will be set to green in the current state, because conflicting streams may be crossing the junction. M_σ is a *conflict-free* subset of *InLanes* with *maximal priority*. This means that:

1. $\forall_{k,l \in M_\sigma} \neg Conflict(k, l)$
2. M_σ is maximal w.r.t. the total order $\geq$ on $\mathcal{P}(InLanes)$ as defined for any $A, B \subseteq InLanes$ by:
   - $A \geq \emptyset$;
   - if $A, B \neq \emptyset$ then $A \geq B \Leftrightarrow maxprio(A) > maxprio(B) \vee (maxprio(A) = maxprio(B) \wedge remmax(A) \geq remmax(B))$,
     where $maxprio(X) = \max\{\pi(x) \mid x \in X\}$ and $remmax(X) = X - \{x\}$ for some $x \in X$ s.t. $\pi(x) = maxprio(X)$.[2]

[2]Note that the definition of $\geq$ is independent of the choice of element x with highest priority.

Please notice that M_σ is not uniquely defined by this definition. There may be different sets of *InLanes* with equal priorities. Therefore, we will treat M_σ as a non-deterministic function and allow a transition $\stackrel{\langle TR_\sigma, TY_\sigma, TG_\sigma \rangle}{\longrightarrow}$ for every maximal set M_σ.

The set TR_σ contains all yellow lights for which the minimal yellow time has elapsed. The set TY_σ contains all green lights of which the minimal green time has elapsed and which are in conflict with one of the *InLanes* in M_σ. The set TG_σ contains all red lights from M_σ of which the minimal red time has elapsed and which are not in conflict with any of the currently crossing traffic streams.

$TR_\sigma = \{l \in \mathit{InLanes} \mid \gamma(l) = \mathit{yellow} \wedge \delta(l) \geq \mathit{MinStateTime}(\mathit{yellow})\}$
$TY_\sigma = \{l \in \mathit{InLanes} \mid \gamma(l) = \mathit{green} \wedge \delta(l) \geq \mathit{MinStateTime}(\mathit{green}) \wedge \exists_{m \in M_\sigma} \mathit{Conflict}(l, m)\}$
$TG_\sigma = \{l \in M_\sigma \mid \gamma(l) = \mathrm{red} \wedge \delta(l) \geq \mathit{MinStateTime}(\mathit{red}) \wedge \neg\exists_{m \in \mathit{InLanes}}(\mathit{Conflict}(l, m) \wedge \mathit{crossing}_\sigma(m))\}$

We used the predicate $\mathit{crossing}_\sigma$, which is defined as follows:

$$\mathit{crossing}_\sigma(m) = (\gamma(m) = \mathit{red} \Rightarrow \delta(m) < \mathit{ClearanceDuration}(m))$$

Next, we define the state resulting after an output transition.

The priorities of the lights that switch to green are reset to zero:

$$\pi' = \lambda l. \begin{cases} 0 & \text{if } l \in TG_\sigma \\ \pi(l) & \text{otherwise} \end{cases}$$

The switching lights receive their new colours:

$$\gamma' = \lambda l. \begin{cases} \mathit{red} & \text{if } l \in TR_\sigma \\ \mathit{yellow} & \text{if } l \in TY_\sigma \\ \mathit{green} & \text{if } l \in TG_\sigma \\ \gamma(l) & \text{otherwise} \end{cases}$$

If a light changes its colour, the colour duration of this light should be reset to zero:

$$\delta' = \lambda l. \begin{cases} 0 & \text{if } l \in TR_\sigma \cup TY_\sigma \cup TG_\sigma \\ \delta(l) & \text{otherwise} \end{cases}$$

The sensor states of the lights that switch to yellow are reset:

$$\xi' = \lambda s. \begin{cases} \mathit{false} & \text{if for some } l \in TY_\sigma, s \in \mathit{sensors}(l) \\ \xi(s) & \text{otherwise} \end{cases}$$

Delay. When time advances, we have to update the priorities and the duration of the current colours. This is expressed in the following transition rule.

$$(\pi, \gamma, \delta, \xi, \tau) \xrightarrow{t} (\pi', \gamma, \delta', \xi, \tau + 1),$$

where

$\pi' = \lambda l. \begin{cases} \mathit{UpdatePrio}(l, \pi(l)) & \text{if } \gamma(l) = \mathit{red} \\ 0 & \text{otherwise} \end{cases}$

$\delta' = \lambda l. \delta(l) + 1$

4.5 Analysis

This phase deals with the development of (mathematical) techniques which aid in validating expressions in our language. We have already mentioned some important properties of a good traffic regulation system. Two of these properties will be discussed in some detail: *safety* and *fairness*.

Safety. The basic safety property of a regulated junction is that never two conflicting traffic streams are allowed to pass the intersection at the same time. Using the predicate *crossing* from Section 4.4, we can formally define this property as follows:

$$\forall_{\sigma \in \Sigma_R} \forall_{l,m \in \mathit{InLanes}} \mathit{Conflict}(l,m) \Rightarrow (\neg \mathit{crossing}_\sigma(l) \vee \neg \mathit{crossing}_\sigma(m))$$

Here, Σ_R denotes the set of reachable states. These are all states that can be reached from the initial state of the system by a series of transitions. Using $\rightarrow^*$ for a series of transitions, it is defined as follows:

$$\Sigma_R = \{\sigma \in \Sigma \mid \sigma_0 \rightarrow^* \sigma\}$$

The safety property holds for every expression in our language. We give an outline of the inductive proof. It clearly holds for the initial state. If it holds for state σ_1 then a successor state σ_2 is reached by one of the three transitions defined in Section 4.4. An input transition does not change the functions γ and δ, so the property also holds for state σ_2. If σ_2 was reached via an output transition, we observe that the only way in which a lane can become crossing is because it is in the set TG_{σ_1}. Then this lane must be in M_σ, which is conflict-free by definition, and it can not be in conflict with an lane that is at the moment crossing. Finally, if σ_2 was reached via a delay transition, remark that δ is incremented (but γ stays the same). This will at best cause some incoming lanes not to be crossing any more, which will not violate the invariant.

Fairness. We will consider a weak notion of fairness: Every traffic light will always eventually become green. Stronger notions could also be defined.

$$\forall_{\sigma_1 \in \Sigma_R} \forall_{l \in \mathit{InLanes}} \exists_{\sigma_2 \in \Sigma} (\sigma_1 \rightarrow^* \sigma_2 \wedge \gamma_2(l) = \text{green})$$

In contrast to the safety property, fairness does not hold for all expressions. Inspection of the definitions reveals that with a bad choice for the function *UpdatePrio* it can be the case that a lane will never receive a green light. Given the fact that we assumed a discrete time domain, requiring that the *UpdatePrio* function is a strictly increasing unbounded function is sufficient. We will not give the proof that under this condition the weak fairness property holds.

There are many other interesting properties. We will mention *throughput*. This is a measure for the efficiency of the intersection. A stochastic analysis, including the dependency of the throughput on *InitPrio* and *UpdatePrio*, is beyond the scope of this paper.

4.6 Pragmatics

Now that we have constructed the actual language, we can have a look at the pragmatics of using the language. As stated before, the pragmatics is concerned with all aspects of using the language. We restrict our discussion to a few interesting aspects.

4.6.1 Methodology

An important part of the methodology is the documentation of the language. This is needed to make clear what well-formed programs are and what they mean. This has already been described in some detail in the previous sections on syntax and semantics. However, for proper use of the language, a designer of a regulated intersection will need more information.

The first step a designer has to take is to determine the physical layout of the intersection. One can derive this from the existing situation, or in some cases it must be developed from scratch. A designer will need guidelines in order to be able to determine the physical structure, e.g. concerning optimal throughput for a given traffic intensity, side conditions due to legislation, cost, etc. Although this is not part of the language proper, the language cannot be effectively used without such methodological issues. The layout of the intersection is (in an abstract way) represented in the program. Since also the grouping of lanes is taken into account, there must be guidelines of how to sensibly form such groups.

Given the physical layout, it is not necessary that all traffic streams that cross each other are in conflict. Some crossings may be considered harmless, e.g. because there is only little traffic. Therefore, the developer also needs to develop the conflict matrix, as a subset of all possible conflicts. Guidelines with respect to this issue should also be covered in a methodology handbook. The notion of a priority update function is very specific to our developed language and not likely to be generally known by developers of intersections. Furthermore, the selected function will have quite some impact on the actual behaviour of the system. A priority function which grows linearly will make a traffic stream less important than one with exponential growth. Therefore, a number of guidelines on which functions to use in which situations are necessary.

From a different perspective, one should not only describe proper use of the language, but also discourage improper use. One could, for instance, use the value of the clearance duration of one traffic stream to regulate the relative priority of conflicting streams. This could be considered bad style.

Another important function of the methodology is to explain when to use which tools to obtain certain results.

4.6.2 Tool support

Clearly, a set of tools should come with the language in order to facilitate the development of a regulated intersection. Ideally, these tools form an *Integrated Development Environment* (IDE). Here, we just mention some interesting tools, without stepping into any details concerning the development of these tools.

The first step is to design tool support for producing expressions in our language. Normally, one would expect tools for (syntax directed) editing, parsing, static checking, etc. Since part of the language is concerned with describing a two-dimensional layout, a graphical editor will show useful. There could be standard components like traffic lights, lanes, and sensors that can be "dragged-and-dropped" to the desired locations. This tool should be able to transform the visual model into the correct program text, and vice versa.

After a model of a traffic junction has been constructed or an existing model is loaded, we want to be able to analyse the model with respect to some of the issues that were raised in Section 4.5.

Two concepts that influence the throughput of the junction are the update functions and clearance duration. To optimise the throughput of the junction we have to build a tool that can assign stochastic distribution functions to the traffic streams so the mean, variance, and other statistical data for the waiting times of these streams can be calculated. Apart from writing a new tool to do this, we can also export the parameters of the junction into an existing statistical tool. Once we have the statistical data we can use this information to recommend values for clearance durations and the update functions. This can be done by building an expert system to calculate the optimal throughput for the given junction.

A simulator is an important tool which can help to assess the strengths and weaknesses of a model before physically building it. This tool will use a pre-constructed model and can visually show the user how the values of clearance durations, update functions, and the offer of traffic streams do influence the behaviour of the junction. There should be a possibility to manipulate time, sudden increases/decreases of traffic, and the occurrence of sensor-input.

Finally, the given model must be implemented to control the target junction. Ideally, the behaviour represented in the model could be compiled into the command language of the device actually controlling a junction. Alternatively, an interpreter of the language could be run which controls the sensors and actuators via an appropriate interface. The latter option would allow for remote control of intersections (e.g. via Internet).

5 Closing Remarks

The purpose of this paper was to promote the use of domain-specific languages as a regular part of the software engineering process. Therefore, based on well known material and published case studies, we described a language-driven approach for software development.

We identified three phases in this approach: formulation of the problem domain, the identification of the problem space and the development of the language. The problem domain follows from a domain analysis. The problem domain is necessarily general and abstract and therefore does not focus on the actual problem exclusively. The problem space adds concepts (concepts due to design decisions) to the concepts of the problem domain. Moreover, the problem space separates the relevant concepts from the irrelevant concepts and considers instantions of the relevant concepts, as to better accomodate for the problem or class of problems that must be solved. The domain-specific language being developed must exactly span the problem space. This language is developed in three sub-phases: syntax, semantics, and pragmatics.

This approach is illustrated by means of a conceptually simple case study. Although the case study is presented in a linear way, the process of developing the case study was iterative. It was our experience that one of the main factors with respect to the quality of the language design was the consistency of the deliverables involved. For instance, in our case study, the priority function as a concept was introduced only *after* developing the semantics, i.e. it was not obtained as a concept in the initial domain analysis. An integrated set of support tools covering all phases of the approach should take care of this consistency checking.

It is a generally accepted fact that it is preferable to detect errors during the early and more abstract phases of system design. The language-driven approach focuses the attention of the developers on the basic concepts of the language and therefore on (the building blocks of) the

semantics. When developing a software system these abstract building blocks must be completely understood. In our case study, we experienced that the discussion focussed mainly on the concepts and semantics. A traditional development process would have focussed more on the design and the implementation of the system.

Because the traffic light case study focuses on a relatively small domain, we cannot assess the applicability of the language-driven approach in a large domain. It must be investigated whether techniques such as top-down design and modularisation, which have a natural place in traditional life-cycle models, also have their counterpart in our approach. A natural way of dealing with larger problems is to identify substructures in the problem domain, which can be dealt with in isolation. The sub-domains give rise to a number of problem spaces, each with their own language. The composition of these partial solutions can e.g. be defined with a co-ordination language.

Acknowledgements. The authors would like to thank Paul Derksen, Ronald Middelkoop, Felix Ogg, and Robert Spee for their discussions and help on the case study. Marc Voorhoeve is acknowledged for his help in clarifying the ideas that led to this paper. Thanks are due to Michael van Hartskamp for proof reading.

References

[1] J.C.M. Baeten and W.P. Weijland. *Process Algebra*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 1990.

[2] B.W. Boehm. *Software Engineering Economics*. Prentice-Hall, 1981.

[3] B.W. Boehm. A spiral model of software development and enhancement. *IEEE Computer*, 21(5):61–72, 1988.

[4] M. Clavel, F. Duràn, S. Eker, and J. Meseguer. Maude as a formal meta-tool. In J. Wing and J. Woodcock, editors, *The World Congress On Formal Methods*, volume 1709 of *LLNCS*, pages 1684–1703. Springer-Verlag, 1999.

[5] C. Consel and R. Marlet. Architecturing software using a methodology for language development. In C. Palamidessi, H. Glaser, and K. Meinke, editors, *Proceedings of the 10th International Symposium on Programming Languages, Implementations, Logics and Programs (PLIP/ALP '98)*, volume 1490, pages 170–194. Springer-Verlag, 1998.

[6] G. Gupta and E. Pontelli. A Horn logic denotational framework for specification, 1999.

[7] M. Hennessy. *The Semantics of Programming Languages: An Elementary Introduction using Structured Operational Semantics*. Wiley, New York, 1991.

[8] R.M. Herndon and V.A. Berzins. The realizable benefits of a language prototyping language. *IEEE Transactions on Software Engineering*, 14:803–809, 1988.

[9] C.A.R. Hoare. An axiomatic basis for computer programming. *Communications of the ACM*, 12(10):576–583, October 1969.

[10] ITU-T. *ITU-T Recommendation Z.106: Common Interchange Format for SDL*. ITU-T, Geneva, 1996.

[11] R. Kieburtz, L. McKinney, J. Bell, J. Hook, A. Kotov, J. Lewis, D. Oliva, T. Sheard, I. Smith, and L. Walton. A software engineering experiment in software component generation. In *Proceedings of the 18th IEEE International Conference on Software Engineering ICSE-18*, pages 542–553. IEEE Computer Society Press, 1996.

[12] P. Klint. A meta-environment for generating programming environments. *ACM Transactions of Software Engineering and Methodology*, 2(2):176–201, 1993.

[13] D.E. Knuth. *Semantics of Context-Free Languages*, volume 2, pages 127–145. Springer-Verlag, New York, 1968.

[14] P.W. Kutter and A. Pierantonio. Montages specifications of realistic programming languages. *Journal of Universal Computer Science*, 3(5):416–442, 1997.

[15] P. Pfahler and U. Kastens. Configuring component-based specifications for domain-specific languages. In *Proceedings of the 34th Annual Hawaii International Conference on System Sciences*. IEEE Computer Society Press, 2000.

[16] G.D. Plotkin. A structural approach to operational semantics. Technical Report DIAMI FN-19, Computer Science Department, Aarhus University, 1981.

[17] R. Prieto-Díaz. Domain analysis: An introduction. *Software Engineering Notes*, 15(2):47–54, 1990.

[18] J. Rekers and A. Schürr. A graph grammar approach to graphical parsing. In *Proceedings of the 1995 IEEE Symposium on Visual Languages*, 1995.

[19] W.W. Royce. Managing the development of large software systems. In *Proceedings of the IEEE WESCON*, 1970.

[20] D.A. Schmidt. *Denotational Semantics: A Methodology for Language Development*. Allyn and Bacon, Inc., Newton, MA, 1986.

[21] M.G.J. van den Brand, A. van Deursen, J. Heering, H.A. de Jong, M. de Jonge, T. Kuipers, P. Klint, L. Moonen, P.A. Olivier, J. Scheerder, J.J. Vinju, E. Visser, and J. Visser. The ASF+SDF meta-environment: A component-based language development environment. In R. Wilhelm, editor, *Compiler Construction*, volume 2027 of *LLNCS*, pages 365–ff. Springer-Verlag, 2001.

[22] R.T. Yeh. An alternate paradigm for software evolution. In P.A. Ng and R.T. Yeh, editors, *Modern Software Engineering: Foundations and Current Perspectives*, New York, NY, 1990. Van Nostrand Reinhold.