

# DeSpec: Modeling the Windows Driver Environment<sup>\*</sup>

Tomas Matousek and Pavel Jezek

Charles University in Prague  
Department of Software Engineering  
Malostranske namesti 25,  
11800 Prague, Czech Republic  
{tomas.matousek, pavel.jezek}@mff.cuni.cz

**Abstract.** This paper introduces a new object-oriented specification and modeling language called DeSpec, targeting primarily model checking in the Windows NT kernel driver environment. It integrates the majority of Zing modeling language features and adds means for defining parameterized abstractions of the environment at varying levels of detail. The DeSpec language also enables capturing constraints imposed on drivers by the Windows kernel in a form of quantified temporal logic patterns – the easy-to-read templates of LTL formulae introduced by the Bandera toolset.

## 1 Introduction

In recent years, efforts to verify correctness of Windows kernel drivers [21] have emerged as it is crucial for stability of the whole operating system. Microsoft itself has developed several tools for driver verification including the latest Static Driver Verifier model checker. The key to successful application of the model checking approach in this area is a reasonable choice of the environment model. However, the environment models used in current tools are too (1) non-deterministic, degrading preciseness of the the model checker reports, and (2) oversimplified, losing the ability to check more specific kinds of properties of drivers. On the other hand, neither a formal or readable specification usable for documentation purposes is provided by these models. This paper targets these issues by introducing a new language for formal specification and modeling of kernel drivers and their environment.

Please note that due to space limitations the paper presents only a small excerpt of the language features. The full language specification, detailed elaboration of its features and also a large sample specification of the Windows environment can be found in [15].

---

<sup>\*</sup> This work was partially supported by the Czech Academy of Sciences project 1ET400300504 and the Grant Agency of the Czech Republic project GD201/05/H014.

## 1.1 Model Checking

The model checking technique is a formal verification method based on creation of a model emulating a software unit to be verified against a given property. This model should ideally retain those parts of the software that influences the property so that the verification is sound and complete with respect to the property. On the other hand, the model should be as simple as possible. The reason is that the model checker has to explore all the possible states of the system and the time and space requirements for the verification are growing exponentially with respect to the number of operations, threads and variables used in the model (the *state explosion problem* [16]).

Another problem arising from application of the model checking approach is a necessity to choose a suitable form in which the properties should be specified. Temporal logics (e. g. *Linear Temporal Logic* (LTL) [14]) are often used for this purpose since they can express changes of predicates validity in time. However, specifying properties of a real application by means of pure temporal logic has an important drawback. The specification will not easy to comprehend for the most of driver programmers and if a formula gets more complex neither for temporal logic experts.

This results in an effort to develop a higher-level language suitable for property specification. Properties expressed in this language are then translated back to the temporal logic that can be consumed by existing verification tools. Bandera *Temporal Logic Patterns* [9] are an example of such a language. They allow writing frequently used temporal logic formulae patterns in very simple plain English sentences used as templates, e. g. “P is absent between Q and R” is representing the  $\Box((Q \wedge \neg R \wedge \Diamond R) \Rightarrow [P U R])$  formula. Although it is not possible to express an arbitrary LTL formula by a pattern, the patterns are sufficient to specify widely used properties. Moreover, additional patterns can be added if needed.

Third issue of the model checking is the extraction of a program model from its source code. First, if a source code of the entire application is not available the *missing environment problem* [31] emerges. The model of the missing parts, i. e. the environment, needs to be specified by hand. Second, source code language itself can be a problem for the model extraction – it is definitely simpler to extract a model from Java or C# source code rather than from pure C source code. A C language extractor needs to understand constructs like pointers, dynamic arrays, unions, reinterpreting type casts, etc. Fortunately, even though these constructs in general do not allow the extractor to build sound nor complete model of the program, the results of the software verification are still valuable. However some error reports would be false and some of the true errors would remain undiscovered.

The choice of a target model checker is also important. The *Zing Model Checker* [1, 23], being developed by Microsoft Research group, is assumed as the target model checker in this paper. Therefore, the model extractor is expected to emit the resulting model in the Zing modeling language. The choice was made due to Zing’s rich modeling functionality and the state of its current develop-

ment (the preview implementation is available and works quite well). However, most ideas behind this work are not dependent on the target modeling language and can be applied to any other modeling language that provides at least classes, methods, exceptions, non-deterministic choices, and threads. Another modeling language meeting these criteria should be the new version of *Bandera Intermediate Representation* (BIR) – a modeling language of *Bogor* model checking framework [27].

## 1.2 Checking Windows Drivers

Windows kernel drivers are relatively small libraries usually written in the C language and running in a privileged mode that enables them to work directly with hardware. This introduces a high risk of damaging other parts of the kernel if a driver contains an error. Hence the correctness of drivers is crucial for security and stability of an operating system and drivers are common subject of software verification.

A driver can be seen as a component put into the environment of the kernel and other drivers. Since drivers usually communicate with each other only via calling kernel functions, this view can be simplified to two components – the driver being verified and the kernel. The verifier has to deal with the missing environment problem here as the source code of the Windows kernel is usually not available. Even if it was, it would be virtually impossible to extract and verify the kernel model as the whole kernel is very complex. Besides, drivers shouldn't depend on the exact behavior of the private parts of the kernel as they can change version to version.

The kernel should be viewed as a black box instead and the model extractor should only work with specification of the functionality provided to the drivers. However, such a specification is not currently available in a form that would be feasible to drive the model extractor – the only source of official documentation is the *Driver Development Kit* (DDK) [20] provided by Microsoft, where the rules the drivers should comply with are described in plain English and some important details are stated vaguely or even missing entirely.

Several tools that verify driver correctness have already been developed by Microsoft itself. These include the *Driver Verifier* [17] tool for run-time driver verification, the *PREfast* [18] static analysis tool based on local analysis of driver functions and finally the *Static Driver Verifier* (SDV) [19] (still in development) based on techniques of static analysis of the whole driver and model checking.

The Static Driver Verifier (SDV) is modeling the kernel environment in C language enriched with special functions and macros that handle non-determinism necessary for emulating various execution paths. The rules the drivers can be verified against are written in *Specification Language for Interface Checking* (SLIC) [2] created for this purpose. Expressing a rule in the SLIC language inheres in writing pieces of C pseudo-code and defining how the environment model should be instrumented by them. The resulting instrumented code is converted to an abstract Boolean program which is passed to the model checker. The very first Boolean program extracted from the instrumented code abstracts from all local

variables and replaces all conditions by non-deterministic choices. Error traces are then discovered by the model checker and confronted with the original program via symbolic execution. If an error trace describes the execution that is actually infeasible, the Boolean program is refined to be more specific with respect to the variables influencing the trace. The refined program is passed back to the model checker. This process of error search and model specialization repeats until there are no infeasible error traces found or a timeout elapses.

The environment model and the SLIC language allows safety properties to be checked with respect to operations performed sequentially on a single device object (an object representing a device in the driver). SLIC rules are limited to safety properties so it is not possible to encode all the rules defined in the DDK. The rules are specified separately from kernel environment which makes them less maintainable. Inability to model multi-threaded environment and simultaneous work on more device objects also prevents from verification of some race conditions commonly contained in faulty Windows drivers.

### 1.3 Paper Contribution

The aim of this work is to make it possible to specify and model the kernel environment in a formal yet comprehensible form, which could be used not only for precise documentation of the kernel API but above all as an input for a model extractor that produces verifiable concurrent models of the Windows drivers. For this purpose, the paper introduces a new specification and modeling language called *Driver Environment Specification Language* (abbreviated as *DeSpec*) [15]. As shown in [15], the language is able to capture a significant subset of the rules imposed on drivers by the DDK including those that are difficult or impossible to express in the SLIC language and hence currently not verifiable by the SDV.

The rest of the paper is laid out as follows. Section 2 briefly describes the Windows kernel environment from a point of view of the driver verification. Section 3 introduces the DeSpec language, explains its part on an example and describes how a model extractor should work with DeSpec specifications. Section 4 discusses related work and Section 5 concludes.

## 2 Windows Kernel Environment

The Windows kernel executive comprises of several components that manage various system resources – the managers [30]. The managers are providing services for the other parts of the executive and for drivers. The *I/O Manager*, the *Plug & Play Manager*, and the *Power Manager* are the ones that are most interesting for driver verification as they do the majority of communication with drivers. Note, this work is limited only to drivers following the *Windows Driver Model* (WDM) [25]. Such drivers have to implement Plug & Play and power management features.

The I/O Manager loads and unloads drivers and issues I/O requests on them. The drivers are directly controlled by the I/O Manager, which issues I/O requests

in form of *I/O Request Packets* (IRPs). If a driver can complete the request it fills in a place in the packet reserved for output parameters and passes the packet back to the manager. If it doesn't implement the required functionality it can pass the request to an optional lower level driver – a driver hierarchy is then formed. The other managers issues their requests and notifications to the drivers through the I/O Manager. For example, the Plug & Play Manager keeps track of the device state transitions (device removal, stopping, starting, etc.) and the Power Manager monitors the power state of the machine (whether it is going to sleep, awaking, etc.). Both managers notify the driver appropriately by sending it the respective IRPs.

Each driver has to respond correctly to an arbitrary request and content of the packet. It can return a result indicating an error, but it must never crash or damage other parts of the kernel. The driver cannot make any assumptions about drivers above or below it in the hierarchy. This requirement allows the verification environment to isolate the driver and test it on arbitrary inputs and outputs from the I/O Manager or higher/lower level drivers.

Therefore, each driver is verified separately. In the time of verification, other drivers are seen as a part of the environment and details of their interior are not visible to the driver being verified as they interact with it only via the I/O Manager. Due to the clear separation of the driver being verified and the environment, it is possible to model the driver using its source code and the specification of the interfaces provided and required by the environment. Neither the kernel's nor third party driver source code is necessary for the verification.

### 3 Driver Environment Specification Language

#### 3.1 Overview

The Driver Environment Specification Language (DeSpec) is an object-oriented specification and modeling language incorporating the majority of features of the Zing modeling language [1] combined with design-by-contract elements with syntax inspired by Spec# language [3], and Bandera Temporal Logic Patterns [9]. It is designed to provide necessary information to the model extraction process leading to the creation of the Zing model of the driver.

It allows modeling the behavior of kernel functions and specifying constraints and rules that drivers should obey when calling these functions. It also enables modeling of the I/O Manager's behavior to drivers. In the DeSpec language, models and abstractions can be defined in various levels of detail, which, as one of the solutions fighting against the state space explosion problem, enables the model extractor to infer the smallest available model sufficient for the verification of a particular rule.

The basic elements of the driver and its environment that the specification should be able to capture are functions, global variables and data structures used to exchange information between the driver and the environment.

The specifications have to map constructs of the C language to their equivalents in the Zing language. Pointers, function pointers, unions and other constructs that are not directly expressible in the Zing language have to be modeled in a special way using only the means available in the Zing language. Apparently, constructs exploiting memory layout, such as reinterpreting casts or unions, cannot be modeled in a feasible way. Fortunately, a well written driver should be as platform independent as possible and thus these constructs should be used rarely.

### 3.2 Structure of Specifications

The DeSpec language is similar to the C# language in its syntactical structure. Each source file contains a list of declarations grouped to namespaces. Declarations include classes, integer enumerations, integer ranges, method delegates and method groups. A class declaration comprises of its members. Apart from fields and methods, which are common for standard object-oriented languages, DeSpec classes can also contain rules. A rule specifies constraints on fields and methods by the means of temporal logic patterns. This section briefly describes DeSpec namespaces, classes and rules.

**Namespaces** A namespace defines a scope for abstractions of kernel functions and structures. When the model extractor searches for an abstraction of a kernel function or structure used in the driver's source code it looks up a single namespace only. The choice of the namespace to be searched by the extractor depends on the constraints the user wants to verify. The default (global) namespace contains a minimal model for the kernel functions and structures. Other namespaces usually *refine* the default model – making it more complex to enable verification of the specific constraint. Constraints are also stated in declarations within the specification, e. g. in a form of rules. By choosing the constraint to verify, the containing namespace is designated for being searched by the extractor. The ability to differentiate specifications by level of details is important for reducing the size of the resulting model.

**Classes** Although Windows kernel is written in the C programming language its design is object oriented. Usually, a structure representing an object within the kernel (e. g. a semaphore, mutex or device) is provided along with functions working with it. These functions in fact behave like methods of the structure (object) as they all take a pointer to the structure as one of their parameters (actually implementing the “this” reference). A notion of inheritance is also present on several places. Inheritance is used for sharing data among structures representing different yet related objects. The sharing technically inheres in declaring common initial fields in the related structures.

These observations justify introduction of classes as main elements of the specifications – the kernel structures provided to drivers are modeled in DeSpec

as classes. The functions bound to these structures are the respective class instance methods. Functions not bound to any instance are mapped to static methods. The formal argument referring to the instance the method is working on is specified by the *instance* keyword. The method (whether static or instance) abstracting a kernel function has to have the same name as the kernel function and no other method (even in another class) can have the same name. This rule allows the model extractor to find a specification of a function whose call has been observed in the source code. An example of a class specification follows:

*Example 1.*

```
class DEVICE_OBJECT
{
    NTSTATUS IoAttachDevice(instance,_,out DEVICE_OBJECT attachedTo)
        requires !Driver.IsLowest;
    {
        NTSTATUS status = choose
        {
            NTSTATUS.STATUS_SUCCESS,
            NTSTATUS.STATUS_INVALID_PARAMETER,
            NTSTATUS.STATUS_OBJECT_TYPE_MISMATCH,
            NTSTATUS.STATUS_OBJECT_NAME_INVALID,
            NTSTATUS.STATUS_INSUFFICIENT_RESOURCES
        };
        attachedTo = IsSuccessful(status) ? Driver.LowerDevice : null;
        return status;
    }

    DEVICE_OBJECT IoAttachDeviceToDeviceStack(instance,_)
        requires !Driver.IsLowest;
    {
        return (choose(bool)) ? Driver.LowerDevice : null;
    }

    void IoDetachDevice(instance);

    /* more members follow */
}
```

In Example 1, the `DEVICE_OBJECT` class abstracts the structure of the same name. Instances of the structure represent devices that drivers are working with. Both *IoAttachDevice* and *IoAttachDeviceToDeviceStack* kernel functions attach the device object to the top of the device objects chain. The immediate lower device object, where the instance is attached to, is returned in the *attachedTo*

output argument and in the return value, respectively. The *IoDetachDevice* simply detaches the immediate higher level device from this device object instance<sup>1</sup>.

The signature of a method abstracting a kernel function defines how parameters of the function are treated within the specification. The *placeholder* token (a single underscore) is used for arguments that are not important for the specification. The models of *IoAttach*- functions do not care about the second parameter. When a specification refers to the *IoAttachDevice* method, only one argument is stated in the list of actual arguments. The instance argument is picked from the argument list out before the method to denote the target instance using the dot notation. The arguments on the placeholders positions are also omitted in the actual argument list. Methods declared in Example 1 are referred to as follows:

```
device.IoAttachDevice(out lower_device)
device.IoAttachDeviceToDeviceStack()
lower_device.IoDetachDevice()
```

The *out* keyword specifies that the argument is an output argument and has to be assigned within the method's body. The output argument is mapped to the C language by an additional level of indirection. The C type of the argument is thus `DEVICE_OBJECT**`. The *ref* keyword is also supported for marking in-out arguments.

A possibly empty list of preconditions and postconditions follows the signature. The syntax is similar to the one used in the Spec# language – the conditions are introduced by *requires* and *ensures* keywords, respectively. The condition is a Boolean expression with some limitations on the terms that form it. The conditions stated in Example 1 require the lowest level driver not to call the *IoAttach*- functions. Pre- and postconditions are translated to assertions when the Zing model of the method is generated.

The body defines a model of the method's behavior using Zing syntax enriched with additional constructs that are translated to the Zing when the resulting model is generated. In Example 1, extended forms of the Zing's *choose* operator are used. Type `NTSTATUS` is an integer enumeration abstracting the kernel type of the same name. The operator *IsSuccessful* determines whether a value is a successful value of its type as recognized by the kernel.

The body can also be omitted at all if the real function whose model is embodied to the method does nothing that influences the driver at the current level of abstraction and only its calls are significant. However, if the real function returns some values to the caller (via a return value or output parameters), throws any exceptions or if it changes the state of the object or any global state the specification method should have a body that models these operations.

Since classes are abstractions of the structures, they should contain fields corresponding to the fields of the structures. Only those fields that are documented should be included and the methods should set their values as described in the

---

<sup>1</sup> Note the reverse roles of the device objects – the higher level device object is attaching but the lower level device object is detaching.

documentation. Usually, additional fields are necessary for storing auxiliary data used only for the sole purpose of modeling. Such field does not correspond to any field visible for drivers and is marked by the *synthetic* keyword. Similarly, *synthetic methods* and also *synthetic classes* can be defined as well. In general, DeSpec distinguishes synthetic language elements from non-synthetic ones. Note that all elements used in the first example are non-synthetic. Synthetic classes contain no abstractions, particularly no kernel function is mapped to a method of a synthetic class. Example of a class containing synthetic attributes follows:

*Example 2.*

```
static class Driver
{
    synthetic DEVICE_OBJECT LowerDevice = new DEVICE_OBJECT;

    [ModelParam]
    synthetic const bool IsLowest = false;

    /* more members follow */
}
```

In Example 2, two synthetic fields are defined in the static class. The first one, *LowerDevice*, is used as a dummy device object that all devices of the current driver are attached to. The model can abstract from the precise device objects chain because the drivers shouldn't care about what drivers are layered beneath them in the chain. Similar simplifications are necessary to reduce the size of the generated model.

The second field named *IsLowest* is a literal constant field defining whether or not the driver is the lowest level driver in the driver chain. The field is annotated by the *ModelParam* attribute, which means that its initial value should be set by the user prior to the model extraction. Model parameterization is utilized when the model depends on a property that is difficult to deduce automatically from the driver's source code. It can be also used for model size tuning.

**Rules** Another member that can be present in the class is a *rule*. The rule is a list of quantified temporal logic patterns [9] with pattern parameters filled with Boolean expressions.

*Example 3.*

```
class DEVICE_OBJECT
{
    /* method declarations from Example 1 omitted */

    static rule
        forall(DEVICE_OBJECT device)
```

```

    {
      _ .IoAttachDevice(out device)::succeeded ||
      (device == _ .IoAttachDeviceToDeviceStack()) && device!=null
    }
    corresponds to
    {
      device.IoDetachDevice()
    }
    globally;
}

```

The rule in Example 3 is a single pattern, however, in general, a rule is a list of quantified temporal logic patterns separated by commas and ending by a semicolon. The rule presented has the following meaning: “Each successfully attached device is eventually detached and each device that is detached has previously been successfully attached.” Rest of the section explains the patterns in more detail.

Each temporal logic pattern is formed by pattern keywords and pattern expressions. The pattern used in Example 3 can be generalized to  $\{P\}$  *corresponds to*  $\{Q\}$  *globally*, where  $P$  and  $Q$  are Boolean expressions. Each pattern can be split into two parts: the property and the scope. In this case, the property is  $\{P\}$  *corresponds to*  $\{Q\}$  and the scope is *globally*. A list of available pattern properties follows:

1.  $\{Q\}$  *is universal*
2.  $\{Q\}$  *is absent*
3.  $\{Q\}$  *exists*
4.  $\{Q\}$  *precedes*  $\{R\}$
5.  $\{Q\}$  *leads to*  $\{R\}$
6.  $\{R\}$  *responds to*  $\{Q\}$
7.  $\{Q\}$  *corresponds to*  $\{R\}$

The properties 1 to 6 are defined in [9]. The properties 5 and 6 are equivalent and it depends on the situation which one is more appropriate to use. The property 7 is equivalent to a conjunction of properties 5 and 4, i. e. to  $\{Q\}$  *leads to*  $\{R\} \wedge \{Q\}$  *precedes*  $\{R\}$ . It has been introduced to the language since the combination is frequently used in the kernel environment and it would be inconvenient to write the two patterns separately. The available scopes are:

1. *globally*
2. *before*  $\{S\}$
3. *after*  $\{P\}$
4. *after*  $\{P\}$  *until*  $\{S\}$
5. *between*  $\{P\}$  *and*  $\{S\}$

The meaning of each property and scope is obvious. Detailed definitions can be found in [9] along with the equivalent LTL formulae. The LTL formulae for

$Q$ -corresponds-to- $R$  pattern with the global scope is

$$\Box(Q \Rightarrow \Diamond R) \wedge [R \ W \ Q]^2.$$

Temporal patterns can be quantified over value types or reference types. Patterns of instance rules are implicitly quantified by a variable of the declaring type. Instance rules can refer to that variable by using *this* keyword. This keyword can be omitted when referring to the instance members of the type. Unlike Bandera [8], DeSpec allows to quantify over value types (i. e. integers, Boolean, enumerations). Zing symbolic value types can be used for the implementation. The reference type quantification may be implemented in the same way as in the Bandera, however more scalable implementation would be possible using Zing symbolic reference types, which should be available in the next version of the Zing.

Method event operator	Effect
<code>M(args)::entered</code>	Returns <i>true</i> after $M$ is entered with specified arguments until $M$ returns. Return value is ignored.
<code>M(args)</code> <code>M(args)::returned</code>	Returns <i>true</i> after $M$ is called with specified arguments and with any return value until $M$ is entered again with the same arguments.
<code>M(args)::succeeded</code>	Returns <i>true</i> after $M$ is called with specified arguments and with a successful return value until $M$ is entered again with the same arguments.
<code>M(args)::failed</code>	Returns <i>true</i> after $M$ is called with specified arguments and with an unsuccessful return value until $M$ is entered again with the same arguments.
<code>expr === M(args)</code> <code>expr !== M(args)</code>	Returns <i>true</i> after $M$ is called with specified arguments and with a return value (un)equal to the value of the <i>expr</i> until $M$ is entered again with the same arguments.

**Table 1.** Source code event operators for methods.  $M$  stands for non-synthetic methods,  $args$  stands for a list of arguments (possibly empty), and  $expr$  denotes an expression.

Boolean expressions comprising pattern parameters should refer to so called *source code events* via *source code event operators*. The source code event inheres in executing a particular piece of code. DeSpec allows to specify events corresponding to function calls and operations on fields (read and write) within the

<sup>2</sup>  $\Box$  is the universal time quantifier (always in the future),  $\Diamond$  is the existential time quantifier (sometime in the future),  $[\varphi \ W \ \psi]$  is the weak until operator (either  $\psi$  never holds and  $\varphi$  holds always, or  $\psi$  holds sometime in the future and  $\varphi$  holds until that moment).

Field event operator	Effect
<code>F::read</code>	Returns <i>true</i> after <i>F</i> is read from.
<code>F::written</code>	Returns <i>true</i> after <i>F</i> is written to.
<code>expr == F::get</code> <code>expr != F::get</code>	Returns <i>true</i> after <i>F</i> is read from and the value read is (not) equal to the value of the <i>expr</i> .
<code>expr == F::set</code> <code>expr != F::set</code>	Returns <i>true</i> after <i>F</i> is written to and the value written is (not) equal to the value of the <i>expr</i> .

**Table 2.** Source code event operators for fields. *F* stands for a non-synthetic field and *expr* denotes an expression.

driver’s source code. Hence, source code event operators are applicable on non-synthetic methods and fields. Available source code event operators are listed in Table 1 and Table 2.

In Example 3, the source code event defined by the `method(args)::succeeded` operator establishes a watchdog for successful returns from the kernel function `IoAttachDevice` such that first two arguments were arbitrary when the function was called, and the third argument can be unified with the *device* quantification variable after the function successfully returns.

Each source code event operator is replaced by the corresponding predicate for the purpose of rule verification. The use of the source code event operator inside a pattern expression implies adding a global state variable to the resulting Zing model and instrumentation of the model with pieces of Zing code that make transitions of the state. The value of the operator state variable determines the value of the LTL formula predicate. Although Zing doesn’t support LTL verification directly, it is possible to use run-time verification algorithm proposed by [11].

### 3.3 DeSpec Driven Model Extraction

Inputs to the model extraction process are the source code of the driver being verified, kernel header files, and the specifications of kernel functions and data structures written in DeSpec. At the beginning, the user should select a set of constraints that he or she wants to verify.

The user also chooses the *top-level model* to be used for the verification. This model is also written in DeSpec as a class implementing the predefined methods. Its task is to emulate the kernel’s behavior to the driver including driver loading and initialization and issuing I/O requests (IRPs). Default top-level model is the most complex one. It emulates multiprocessor environment, multiple device objects, and concurrent IRPs. However, for a verification of some rules a simpler model may be sufficient. DeSpec allows to write and use such model. The choice of the simpler model may radically reduce the size of state space and make the verification faster and sometimes even allow the verification to be completed in realistic time. However, some errors may remain undiscovered.

Once the top-level model is chosen, the model extractor generates Zing model of the driver (using its C source code and kernel headers) and combines it with the environment model. Since the resulting model is too large to be verified, the slicing [13, 10] should take place retaining only those parts transitively referred to by the top-level model and the constraints being verified. As a final result, a Zing model of the driver and the related kernel functions and structures are output.

## 4 Related Work

This work incorporates or relies on ideas and approaches of the model checking [6, 16], model extraction [7, 29], temporal logics [26, 14, 5], source code static analysis and slicing [13, 10], and Windows kernel driver environment [30, 25].

In particular, the Zing model checker [1], Bandera toolset (especially the Bogor model checking framework [27, 28]), *Java Path Finder* [24], and *SPIN model checker* [12, 4] are related tools devoted to the model checking.

The *SLAM project* [22] is addressing the static analysis and verification of the C programs, especially the Windows kernel drivers. The beta version of Microsoft Static Driver Verifier (SDV) tool [19] has been recently released as a result of efforts in this area. Since this paper targets on Windows kernel drivers verification, the SDV is the closest related work. The way how rules are specified in this tool limits its verification power to safety properties. The environment model used by SDV is single-threaded, preventing verification of some race conditions, and quite non-deterministic, introducing additional false reports. It neither provides a specification of the kernel functions that might be used as a documentation. On the other hand, SDV is a functional tool whose application in practice already led to discovering several errors in Microsoft's own drivers.

Finding errors in drivers is not limited to the model checking technique. Microsoft *PREfast* tool for drivers [18] performs static analysis of the source code and searches for common error patterns. It can, for example, find memory leaks incurred by missing function calls, dereferences of null pointers, buffer overruns, kernel functions called on incorrect IRQL level, and so on. The analysis is function scoped and hence it introduces false negatives and also restricts a set of errors it is able to detect.

The Windows operating system also enables to check how drivers work in stress conditions such as lack of memory, missing resources, lost packets, etc. In cooperation with the kernel, *Driver Verifier* tool [17] emulates such conditions and runs tests on the specified driver. The tool is able to detect many errors but it doesn't do any static verification so many execution paths remain unchecked.

## 5 Conclusion and Future Work

This paper introduces the DeSpec language – a new specification and modeling language designed to enable writing modular, readable, and well arranged specifications of the Windows kernel driver environment as well as formally, yet still

comprehensibly, capture rules imposed on drivers by the kernel and documented in plain English in DDK.

Expressiveness and suitability of the language are demonstrated on a part of the kernel functionality in [15]. This work also shows that the available documentation of the kernel environment [20] is not sufficient for its formal specification without a deeper understanding of the Windows kernel.

As the DeSpec language is intended to be utilized by model checking tools, it addresses the main issue of this verification method – the state explosion problem. The abstractions may vary in the level of detail chosen according to the properties being verified. Complexity of the model can be further tuned by the user specified model parameters. By setting these parameters, the user can influence how complex the extracted model will be and what may it neglect. The user may also select a subset of tested driver functionality by choosing an appropriate top-level model.

The possibility of verifying LTL formulae with finite trace semantics using assertions only (see [11]) arises a question whether the use of temporal rules brings something new beyond the use of explicit assertions. Although many rules may be equivalently verified manually, i. e. by adding assertions (or method contracts) on the right places in the functions' model code, the use of rules has some advantages. Several advantages are implied by the locality. If entire “business logic” of the rule is written on a single place it is easier maintainable, more readable, and the verification of the rule can be easier (un)selected for verification. Besides, when the rule is more complex it wouldn't be easy to manually keep track of all operations in the code that influences the verified property. On the other hand, some rules are too complicated to write or comprehend that it is better to implement them manually by explicit assertions.

The ideas proposed by this paper are currently being implemented. The implementation comprises of the DeSpec language analyzer and a model extractor consuming C source code and producing a Zing model driven by DeSpec specifications.

## References

1. Andrews, T., Qadeer, S., Rajamani, S. K., Rehof, J., Xie, Y: Zing: A model checker for concurrent software, Technical report, Microsoft Research, 2004.
2. Ball T., Rajamani, S. K.: SLIC: a Specification Language for Interface Checking, Technical Report, MSR-TR-2001-21, Microsoft Research, 2002
3. Barnett, M., Leino, K. R. M., Schulte, W.: The Spec# Programming System - An Overview, Microsoft Research, 2004
4. Bell Labs: SPIN model checker, <http://spinroot.com>
5. Clarke, E. M., Emerson, E. A., Sistla, A. P.: Automatic verification of finite-state concurrent systems using temporal logic specifications, ACM Transactions on Programming Languages & Systems, 244-263, 1986
6. Clarke, E. M., Grumberg, O., Peled, D. A.: Model Checking, MIT Press, 2000.
7. Corbett, J. C., Dwyer, M. B., Hatcliff, J., Laubach, S., Pasareanu, C. S., Robby, Zheng, H.: Bandera: Extracting Finite-state Models from Java Source Code, proceedings of the International Conference on Software Engineering (ICSE), 2000

8. Corbett, J. C., Dwyer, M. B., Hatcliff, J., Robby: Expressing Checkable Properties of Dynamic Systems: The Bandera Specification Language, 2001
9. Dwyer, M. B., Avrunin, G. S., Corbett, J. C.: Patterns in property specifications for finite-state verification, in Proceedings of the 21st international Conference on Software Engineering, 411-420, 1999
10. Dwyer, M. B., Hatcliff J.: Slicing Software for Model Construction, Journal of High-order and Symbolic Computation, 2000
11. Giannakopoulou D., Havelund K.: Runtime Analysis of Linear Temporal Logic Specifications, RIACS Technical Report 01.21, 2001
12. Holzmann, G. J.: The SPIN Model Checker: Primer and Reference Manual, Addison-Wesley Professional, 2003
13. Krinke, J.: Advanced Slicing of Sequential and Concurrent Programs, PhD thesis, Fakultt Fr Mathematik und Informatik, Universitt Passau, 2003
14. Lamport, L.: "Sometime" is sometimes "not never" - on the temporal logic of programs, in Proceedings of 7th ACM Symposium on Principles of Programming Languages, pages 174-185, 1980.
15. Matousek, T.: Model of the Windows Driver Environment, Master Thesis at Department of Software Engineering, Charles University in Prague, 2005, <http://nenya.ms.mff.cuni.cz/publications/Matousek-thesis.pdf>
16. McMillan, K. L.: Symbolic model checking - an approach to the state explosion problem, PhD thesis, SCS, Carnegie Mellon University, 1992
17. Microsoft: Driver Verifier, <http://www.microsoft.com/whdc/DevTools/tools/DrvVerifier.mspx>
18. Microsoft: PREfast, <http://www.microsoft.com/whdc/devtools/tools/PREfast.mspx>
19. Microsoft: Static Driver Verifier - Finding Driver Bugs at Compile-Time, WHDC, <http://www.microsoft.com/whdc/devtools/tools/sdv.mspx>
20. Microsoft: Windows Driver Development Kit, WHDC, <http://www.microsoft.com/whdc/devtools/ddk/default.mspx>
21. Microsoft: Windows Driver Foundation, WHDC, <http://www.microsoft.com/whdc/driver/wdf/default.mspx>
22. Microsoft Research: SLAM Project, <http://research.microsoft.com/slam>
23. Microsoft Research: Zing Model Checker, <http://research.microsoft.com/zing>
24. NASA Intelligent Systems Division: Java Path Finder, <http://ase.arc.nasa.gov/havelund/jpf.html>
25. Oney, W.: Programming the Microsoft Windows Driver Model, 1999, Microsoft Press
26. Pnueli, A.: The temporal logic of programs, in 18th IEEE Symposium on Foundation of Computer Science, pages 46-57, 1977.
27. Robby, Dwyer, M. B., Hatcliff, J.: Bogor: An Extensible and Highly Modular Software Model Checking Framework, SIGSOFT Softw. Eng. Notes 28, 5, 267-276, 2003
28. Robby, Dwyer, M. B., Hatcliff, J.: Bogor, <http://bogor.projects.cis.ksu.edu>
29. SAnToS laboratory: Bandera project, <http://bandera.projects.cis.ksu.edu>
30. Solomon, D. A., Russinovich, M. E.: Inside Microsoft Windows 2000 Third Edition, Microsoft Press, 2000
31. Tkachuk, O., Dwyer, M. B., Pasareanu, C.: Automated Environment Generation for Software Model Checking, Proceedings of ASE 2003.