Automatically Verifying Typing Constraints for a Data Processing Language

Michael Backes[1,2], Cătălin Hriţcu[1,3], Thorsten Tarrach[1,4,5]

[1]Saarland University, Saarbrücken, Germany [2]MPI-SWS, Saarbrücken, Germany
[3]University of Pennsylvania, Philadelphia, USA
[4]Atomia AB, Västerås, Sweden [5]Troxo DOO, Niš, Serbia

Abstract. In this paper we present a new technique for automatically verifying typing constraints in the setting of a first-order data processing language with refinement types and dynamic type-tests. We achieve this by translating programs into a standard while language and then using a general-purpose verification tool. Our translation generates assertions in the while program that faithfully represent the sophisticated typing constraints in the original program. We use a generic verification condition generator together with an SMT solver to prove statically that these assertions succeed in all executions. We formalise our translation algorithm using an interactive theorem prover and provide a machine-checkable proof of its soundness. We provide a prototype implementation using Boogie and Z3 that can already be used to efficiently verify a large number of test programs.

1 Introduction

Dminor [7] is a first-order data processing language with *refinement types* (types qualified by Boolean expressions) and *dynamic type-tests* (Boolean expressions testing whether a value belongs to a type). The combination of refinement types and dynamic type-tests appears in a recent commercial language code-named M [2] and seems to be very useful in practice. However, the increased expressivity allowed by this combination makes statically type-checking programs very challenging.

In this paper we present a new technique for statically checking the typing constraints in Dminor programs by translating these programs into a standard while language. The sophisticated typing constraints in the original program are faithfully encoded as assertions in the generated program and we use a general-purpose verification tool to show statically that these assertions succeed in all executions. This opens up the possibility to take advantage of the huge amount of proven techniques and ongoing research done on general-purpose verification tools.

We have proved that if all assertions succeed in the translated program then the original Dminor program does not cause typing errors when executed. This proof was done in the Coq interactive theorem prover [5], based on a formalisation of our translation algorithm. We thus show formally that, for the language we are considering, a generic verification tool can check the same properties as a sophisticated type-checker. To the best of our knowledge, this is the first machine-checked proof of a translation to an intermediate verification language (IVL).

Finally, we provide a prototype implementation using Boogie and Z3 that can already be used to verify a large number of test programs and we report on an experimental evaluation against the original Dminor type-checker.

1.1 Related work

Biermann et al. [7] were the first to study the combination of refinement types and dynamic type-tests. They introduce a first-order functional language called Dminor, which captures the essence of M [2], but which is simple enough to express formally. They show that the combination of refinement types and dynamic type-tests is highly expressive; for instance intersection, union, negation, singleton, dependent sum, variant and algebraic types are all derivable in Dminor. This expressivity comes, however, at a cost: statically type-checking Dminor programs is very challenging, since the type information can be "hidden" deep inside refinements with arbitrary logical structure. For instance intersection types $T \& U$ are encoded in Dminor as refinement types $(x : \mathsf{Any\ where}\ (x\ \mathsf{in}\ T\ \&\&\ x\ \mathsf{in}\ U))$, where the refinement formula is the boolean conjunction of the results of two dynamic type-tests. Syntax-directed typing rules cannot deal with such "non-structural" types, so Biermann et al. [7] propose a solution based on semantic subtyping. They formulate a semantics in which types are interpreted as first-order logic formulae, subtyping is defined as a valid implication between the semantics of types and they use an SMT solver to discharge such logical formulae efficiently.

The idea of using an SMT solver for type-checking languages with refinement types is quite well established and was used in languages such as SAGE [17], F7 [6], Fine [30] and Dsolve [29]. Biermann et al. [7] show that, in the setting of a first-order language, the SMT solver can play a more central role: They successfully use the SMT solver to check semantic subtyping, not just the refinement constraints. However, while in Dminor [7] subtyping is semantic and checked by the SMT solver, type-checking is still specified by syntax-directed typing rules, and implemented by bidirectional typing rules. In the current work we show that we can achieve very similar results to those of the Dminor type-checker without relying on *any* typing rules, by using the logical semantics of Dminor types directly to generate assertions in a while program, and then verifying the while program using standard verification tools.

Relating type systems and software model-checkers is a topic that has received attention recently from the research community [15, 18, 25]. Our approach is different since we enforce typing constraints using a verification condition generator. Our implementation uses Boogie [20], the generic verification condition generation back-end used by the Verified C Compiler (VCC) [10] and Spec# [4].

There is previous work on integrating verification tools such as Boogie [8] and Why [13] with proof assistants, for the purpose of manually aiding the verification process or proving the correctness of background theories with respect to natural models. However, we are not aware of other machine-checked correctness proofs for translations from surface programming languages into IVLs, even for a language as simple as the one described in this paper. A translation from Java bytecode into Boogie was proved correct in the Mobius project [1, 19], but we are not aware of any mechanized formalisation of this proof.

1.2 Overview

In §2 we provide a brief review of Dminor and in §3 we give a short introduction to our intermediate verification language. §4 and §5 describe our translation algorithm and its implementation. In §6 we compare our work to the Dminor type-checker [7]. Finally, in §7 we conclude and discuss some future work. Further details, our implementation, our Coq formalisation and proofs are all available online at:
`http://www.infsec.cs.uni-saarland.de/projects/dverify`.

2 Review of Dminor (Data Processing Language)

Dminor is a first-order functional language for data processing. We will briefly review this language below; full details are found in the paper by Bierman et al. [7].

Values in Dminor can be simple values (integers, strings, Booleans or null), collections (multi-sets of values) or entities (records). Dminor types include the top type (Any), scalar types (Integer, Text, Logical), collection types ($T*$) and entity types ($\{\ell\colon T\}$). More interestingly, Dminor has *refinement types*: the refinement type $(x : T \text{ where } e)$ consists of the values x of type T satisfying the Boolean expression e.

Syntax of Dminor Expressions:

$e ::=$	Dminor expression
$x \mid k$	variables and scalar constants
$\oplus(e_1, \ldots, e_n)$	primitive operator application
$e_1?e_2 : e_3$	conditional (if-then-else)
$\text{let } x = e_1 \text{ in } e_2$	let-expression (x bound in e_2)
$e \text{ in } T$	dynamic type-test
$\{\ell_i \Rightarrow e_i \ ^{i\in 1..n}\}$	entity (record with n fields $\ell_i \ldots \ell_n$)
$e.\ell$	selects field ℓ of entity e
$\{v_1, \ldots, v_n\}$	collection (multiset)
$e_1 :: e_2$	adding element e_1 to collection e_2
$\text{from } x \text{ in } e_1 \text{ let } y = e_2 \text{ accumulate } e_3$	collection iteration (x, y bound in e_3)
$f(e_1, \ldots, e_n)$	function application

Refinement types can be used to express pre- and postconditions of functions, as shown in the type of removeNulls below, where the postcondition states that the resulting collection has at most as many elements as the original collection.

Refinement type used to encode a pre- and postconditions

$e.\text{Count} \triangleq \text{from } x \text{ in } e \text{ let } y = 0 \text{ accumulate } y + 1$

$\text{NullableInteger} \triangleq x : \text{Any where } (x \text{ in Integer} \mid\mid x == \text{null})$

$\text{removeNulls}(c : \text{NullableInteger}*) : (x : \text{Integer}* \text{ where } x.\text{Count} \leq c.\text{Count})\ \{$
  $\text{from } x \text{ in } c \text{ let } y = \{\} \text{ accumulate } ((x \neq \text{null})?(x :: y) : y)$
$\}$

The dynamic type-test expression e is T pattern-matches the result of expression e against the type T; it returns **true** if e has type T and **false** otherwise. While dynamic type-test are useful on their own in a data processing language (e.g. for pattern-matching an XML document against a schema represented as a type [14]), they can also be used inside refinement types, which greatly increases the expressivity of Dminor (e.g. it allows encoding union, intersection, negation types, etc., as seen in the example above, where NullableInteger is an encoded union type).

Bierman et al. [7] define a big-step operational semantics for Dminor, in which evaluating an expression can return either a value or "error". An error can for instance arise if a non-existing field is selected from an entity. In Dminor such errors are avoided by the type system, but in this work we rule them out using standard verification tools. The type system by Bierman et al. uses semantic subtyping: they formulate a logical semantics (denotational) in which types are interpreted as first-order logic formulae and subtyping is defined as the valid implication between such formulae. More precisely, they define a function $\mathbf{F}[\![T]\!](v)$ that returns a first-order logic formula testing if the value v belongs to a type T. Since (pure) expressions can appear inside refinement types, $\mathbf{F}[\![T]\!]$ is defined by mutual recursion together with two other functions: $\mathbf{R}[\![e]\!]$ returns a first-order logic term denoting the result[1] of an expression e; and $\mathbf{W}[\![T]\!](v)$ a formula that tests if checking whether v is in type T causes an execution error. The reason for the existence of **W** is that **F** is total and has to return a boolean even when evaluating the expression inside a refinement type causes an error. Our translation makes use of the functions **F** and **W** to faithfully encode the typing constraints in Dminor as assertions in the generated while program.

3 Bemol (Intermediate Verification Language)

We define a simple intermediate verification language (IVL) we call Bemol. Bemol is much simplified compared to a generic IVL: the number of language constructs has been reduced and some Dminor-specific constructs that would normally be encoded were added as primitives. We use Bemol to simplify the presentation, the formalisation of our translation and the soundness proof. In our implementation we use Boogie [3, 11, 20] as the IVL and we encode all Bemol constructs that do not have a direct correspondent in Boogie.

3.1 Syntax and Informal Semantics

Bemol is a while language with collections, records, asserts, mutually recursive procedures, variable scoping and evaluation of logical formulae. The syntax of Bemol is separated into two distinct classes: expressions e, which are side-effect free, and commands c, which have side-effects.

Our expressions allow basic operations on values, most of which directly correspond to the operations in Dminor. Also the available primitive operators $\oplus$ are the

[1] Bierman et al. [7] show that $\mathbf{R}[\![e]\!]$ coincides with the big-step operational semantics on pure expressions – i.e., expressions without side-effects such as non-determinism (accumulate) and non-termination (recursive functions).

same as in Dminor. The only significant difference is the expression formula f which "magically" evaluates the logical formula f and returns a boolean encoding the validity or invalidity of the formula – such a construct is standard in most IVLs. We use the notation $[\![e]\!]_{st}$ for the evaluation of expression e under state st. In case of a typing error (such as selecting a non-existing field from an entity) $\perp$ is returned.

Syntax of Bemol Expressions:

$e ::=$	Bemol expression
x	variable
v	Dminor value (scalar, collection or entity)
$\oplus(e_1, \ldots, e_n)$	primitive Dminor operator application
$e.\ell$	selects field ℓ of entity e
$e_1[\ell := e_2]$	updates field ℓ in entity e_1 with e_2 (produces new entity)
$e_1 :: e_2$	adds element e_1 to collection e_2 (produces new collection)
$e_1 \setminus \{e_2\}$	removes one instance of e_2 from e_1 (produces new collection)
is_empty e	returns **true** if e is the empty collection; **false** otherwise
formula f	returns **true** if formula f is valid in the current state

Syntax of Bemol Commands:

$c ::=$	Bemol command
skip	does nothing
$c_1;\ c_2$	executes c_1 and then c_2
$x\ :=\ e$	assigns the result of e to x
if e then c_1 else c_2	conditional
while e inv a do c end	while loop with invariant a
assert f	expects that formula f holds, causes error otherwise
$x\ :=$ pick e	puts an element of e in x (non-deterministic)
call P	calls the procedure P
backup x in c	backs up the current state

Bemol commands manipulate the current global state, which is a total function that maps variables to values. The invariant specified in the while command does not affect evaluation; its only goal is to aid the verification condition generator. The pick command chooses non-deterministically an element from collection e and assigns its value to variable x. The call P command transfers control to procedure P, which will also operate on the same global state. The backup x in c command backs up the current state, executes c and once this is finished restores all variables to their former value except for x. This is useful for simulating a call-stack for procedures, and we also introduce it during the translation to simplify the soundness proof. A similar technique is employed by Nipkow [26] for representing local variables. We use this in our encoding for procedure calls below. The encoding uses an entity to pass multiple arguments.

Encoding of procedure calls

$x \;:=\; \mathsf{call}\ \mathsf{P}(e_1, \ldots, e_n) \triangleq$
  $\mathsf{backup}\ x\ \mathsf{in}\ ($
    $\mathsf{arg} \;:=\; \{\}; \mathsf{arg} \;:=\; \mathsf{arg}[\ell_1 := e_1]; \ldots; \mathsf{arg} \;:=\; \mathsf{arg}[\ell_n := e_n]; \mathsf{call}\ \mathsf{P}; x \;:=\; \mathsf{ret})$

$\mathsf{procedure}\ \mathsf{P}(x_1, \ldots, x_n)\{\ c\ \} \triangleq \mathsf{proc}\ \mathsf{P}\ \{\ x_1 \;:=\; \mathsf{arg}.\ell_1; \ldots; x_n \;:=\; \mathsf{arg}.\ell_n;\ c\ \}$

3.2 Operational semantics

We define the big-step semantics of Bemol as a relation $st_{init} \xrightarrow{c} r$, where r can be either a final state st_{final} or $\mathbf{Error}$. The only command that can cause an $\mathbf{Error}$ is the assert command; all the other commands simply "bubble up" the errors produced by failed assertions. If an expression evaluates to $\bot$ it will lead to the divergence of the command that contains it, but this does not cause an error.[2]

3.3 Hoare logics and verification condition generation

We define a Hoare logics for our commands, based on the Software Foundations lecture notes [27] and the ideas of Nipkow [26].

Definition 1 (Hoare triple). *We say that a Hoare triple* $\models \{P\}\, c\, \{Q\}$ *holds semantically if and only if* $\forall st\ r.\ st \xrightarrow{c} r \implies \forall z.\ P\ z\ st = \mathbf{true} \implies \exists st'.\ r = st' \wedge Q\ z\ st' = \mathbf{true}$.

By requiring that the result of the command is not $\mathbf{Error}$ but an actual state st' we ensure that correct programs do not cause assertions to fail. The meta-variable z is an addition to the traditional Hoare triple and models auxiliary variables, which need to be made explicit in the presence of recursive procedures. Our treatment of auxiliary variables and procedures follows the one of Nipkow [26], who formalises an idea by Morris [24] and Kleymann [16].

The Hoare rules for the standard commands are the same as in Nipkow's work [26], we just list the rules for the constructs that are new to Bemol.

Selected Hoare Rules for Bemol

(Hoare Assert)

$$\frac{}{C \models \{Q \wedge a\}\ \mathsf{assert}\ a\ \{Q\}}$$

(Hoare Pick)

$$\frac{}{C \models \{\lambda z\ st.\ \forall v \in [\![e]\!]_{st}, P\{v/x\}\ z\ st\}\ x\ :=\ \mathsf{pick}\ e\ \{P\}}$$

(Hoare Backup)

$$\frac{\forall st'.\ C \models \{\lambda z\ st.\ P\ z\ st \wedge st' = st\}\, c\, \{\lambda z\ st.\ Q\{st\ x/x\}\ z\ st'\}}{C \models \{P\}\ \mathsf{backup}\ x\ \mathsf{in}\ c\ \{Q\}}$$

[2] Since we only reason about partial-correctness, diverging programs are considered correct. This makes the assumptions on our encoding of Bemol into Boogie be minimal: we only assume that the asserts and successful evaluations of the other commands are properly encoded in Boogie. In §4 we prove that we add enough asserts to capture all errors in the original Dminor program, even under these conservative assumptions we make in the Bemol semantics.

The backup x in c command requires that the Hoare triple for c has the same state for the pre- and postcondition, except for variable x which is updated. We "transfer" the state from the pre- to the postcondition by quantifying over a new state st' that we require to be equal to the state in the precondition.

For our semantics of the Hoare triples it is possible to define a weakest precondition, but not a strongest postcondition function. This is because if c evaluates to $\mathbf{Error}$ no postcondition is strong enough to make the triple valid. Corresponding to the Hoare rules, we define a verification condition generator (VCgen $c\,Q$), which takes a command c and a postcondition Q as arguments and generates a precondition. We have proved that this is sound, however, the VCgen is not guaranteed to return the weakest precondition, because the user-provided loop invariants are not necessarily the best. The soundness proof of the VCgen crucially relies on the soundness of the Hoare logic rules above.

More importantly for our application, we have proved as a corollary that the programs deemed correct by our VCgen do not cause errors when executed.

Theorem 1 (Soundness of VCgen).
If VCgen $c\,Q$ *returns a valid formula, then* $\nexists st.\ st \xrightarrow{c} \mathbf{Error}$.

4 Translation from Dminor to Bemol

Our translation algorithm is a function $\langle\!\langle\, e \,\rangle\!\rangle_x$ that takes a Dminor program and a variable name x as input and outputs a Bemol program. The variable x is where the generated Bemol program should store the result after it executes. We will introduce the translation using two examples.

In the examples below we consider out to be the variable where the result is put. In Example 1 we show how the removeNulls example from §2 is translated to a while loop that picks and removes elements from the collection until it is empty.

Example 1: Accumulate filtering null **values**

```
removeNulls(c : NullableInteger*) :
  (x : (Integer*) where (x.Count ≤ c.Count))
{
  from x in c let y = {}
    accumulate ((x ≠ null)?(x :: y) : y)
}
```

```
procedure removeNulls(c) {
  assert F⟦NullableInteger*⟧(c);
  y := {};
  c' := c;
  while !is_empty c' inv i(c, c', y) do
    x := pick c';
    c' := c'\{x};
    if x ≠ null then
      y := x :: y
    else
      y := y
  end;
  ret := y;
  assert (F ⟦x : Integer*
    where y.Count ≤ c.Count) ret⟧
}
```

$i(c, c', y) \triangleq \mathbf{F}[\![y : (\mathsf{Integer}*) \text{ where } (c'.\mathsf{Count} + y.\mathsf{Count} \leq c.\mathsf{Count})]\!]\ y$

The loop invariant specifies that the sum of the number of elements in the the intermediate collection c' and the resulting collection y is less or equal than the number of elements in the original collection c. It is not sufficient for the invariant to just reason over y and c as this would be too weak. In this case the invariant is provided by hand on the generated code because this loop invariant is not expressible as a Dminor type. Loop invariant inference on the Dminor side is in general deemed to fail for global properties of collections. Our implementation successfully verifies this example with the provided invariant and in the future we hope to infer such invariants automatically.

As seen in Example 2 for type-tests we first use an assert to check that the type-test does not cause a typing-error and then perform the actual type-test which returns a Logical. Note that **F** is total and would also return a value on a wrongly typed argument.

Example 2: Type-test

x in (y : Integer where $y > 5$)	assert ($\neg(\mathbf{W}[\![y$: Integer where $y > 5]\!]\ x)$);
	out := formula ($\mathbf{F}[\![y$: Integer where $y > 5]\!]\ x$)

For illustration, we expand $\mathbf{W}$ and $\mathbf{F}$ in the example above; please see the paper by Bierman et al. [7] for the precise definition of these functions of the logical semantics.

$$
\begin{aligned}
&\mathbf{W}[\![y : \text{Integer where } y > 5]\!]\ x \\
\models\ &\mathbf{W}[\![\text{Integer}]\!]\ x \vee \text{let } y = x \text{ in } \neg(\mathbf{R}[\![y > 5]\!] = \mathbf{Return}(\mathbf{false}) \\
&\qquad \vee \mathbf{R}[\![y > 5]\!] = \mathbf{Return}(\mathbf{true})) \\
\models\ &\mathbf{false} \vee \neg((\text{if } \mathbf{F}[\![\text{Integer}]\!]\ x \text{ then } \mathbf{Return}(x > 5) \text{ else } \mathbf{Error}) = \mathbf{Return}(\mathbf{false}) \\
&\qquad \vee (\text{if } \mathbf{F}[\![\text{Integer}]\!]\ x \text{ then } \mathbf{Return}(x > 5) \text{ else } \mathbf{Error}) = \mathbf{Return}(\mathbf{true})) \\
\models\ &\neg\mathbf{F}[\![\text{Integer}]\!]\ x \models \neg(\text{In\_Integer } x)
\end{aligned}
$$

$$
\begin{aligned}
&\mathbf{F}[\![y : \text{Integer where } y > 5]\!]\ x \\
\models\ &\mathbf{F}[\![\text{Integer}]\!]\ x \wedge \text{let } y = x \text{ in } \mathbf{R}[\![y > 5]\!] = \mathbf{Return}(\mathbf{true}) \\
\models\ &\text{In\_Integer } x \wedge (\text{if In\_Integer } x \text{ then } \mathbf{Return}(x > 5) \text{ else } \mathbf{Error}) = \mathbf{Return}(\mathbf{true}) \\
\models\ &\text{In\_Integer } x \wedge x > 5
\end{aligned}
$$

In case x is not an integer the formula In_Integer $x \wedge x > 5$ is logically equivalent to $\mathbf{false}$. Our translation asserts that x is an integer before calling formula in order to match the semantics of Dminor, in which $x > 5$ causes an error when x is not an integer.

4.1 Soundness

We have proved in Coq that if a Dminor program e can raise an error, then the translated program $\langle\!\langle\, e \,\rangle\!\rangle_x$ can evaluate to an error in Bemol. The contrapositive of this is: if the translated program cannot evaluate to an error, then the original Dminor program cannot evaluate to an error either. We have proved this theorem in Coq by induction over the big-step semantics of Dminor $\Downarrow$.

Theorem 2 (Soundness of translation). *If* $e \Downarrow \mathbf{Error}$ *then* $\forall st.\ st \xrightarrow{\langle\!\langle e \rangle\!\rangle_x} \mathbf{Error}$.

As an immediate consequence of Theorem 1 and Theorem 2 we obtain the soundness of our whole technique.

Corollary 1 (Soundness). *If* VCgen $\langle\!\langle\, e \,\rangle\!\rangle_x$ true *is a valid formula, then* $e \not\Downarrow \mathbf{Error}$.

5 Implementation

Our implementation is called DVerify and translates a Dminor program into a Boogie program. DVerify is written in F# 2.0 [22] and consists of more than 1200 lines of code, as well as a 700 line axiomatisation that defines the Dminor types and functions in Boogie. The Boogie tool then takes the translated Boogie program as input and outputs either an error message that describes points in the program where certain postconditions or assertions may not hold [21] or otherwise prints a message indicating that the program has been verified successfully.

5.1 High-level overview

The heart of our translation algorithm consists of a recursive function that goes over a Dminor expression and translates it into Boogie code. This function is called once per Dminor function and produces a Boogie procedure. Types in Dminor are translated into Boogie function symbols returning bool, using another recursive function in our implementation.

The while loops produced by the translation use the type annotation of accumulate in the source program to generate an invariant for our while loop. In the future we intend to infer such loop invariants automatically using the Boogie infrastructure for this task. The Dminor language as implemented by the type-checker allows for one more construct to define a loop, a from-where-select as in LINQ [23]. In theory from-where-select can be encoded using accumulate, but in the Dminor implementation it is considered primitive in the interest of efficiency and to reduce the type annotation burden. Since from-where-select does not carry a type annotation, we have to find one during translation. For that we use a modified version of the type-synthesis from Dminor that does not call the type-checking algorithm and therefore never fails to synthesise a type for an expression.

We use the Dminor implementation as a library so that we do not have to reimplement existing functionality. This is mainly the parser for Dminor files, the purity checking and a weak form of type-synthesis for from-where-select.

5.2 Axiomatisation

A big part of the implementation is the axiomatisation of Dminor values and functions in Boogie. This is necessary because Boogie as such understands only two sorts, bool and int, whereas Dminor and Bemol have a number of primitive and composite values, such as collections and entities. Our axiomatisation is similar to the axiomatisation the Dminor type-checker feeds to Z3 [7]. In Dminor this axiomatisation is written in SMT-LIB 1.2 syntax [28] and directly fed to Z3 with the proof obligation. Our axiomatisation is in the Boogie language and Boogie translates it to Simplify syntax [12] and feeds it to Z3 along with the verification conditions it generates. Dminor makes heavy usage of the theories Z3 offers, such as extensional arrays and datatypes for example. We use the weak arrays provided by Boogie by default and encode datatypes by hand.

Table 1 Precision comparison

	well-typed	ill-typed
Test suite	76	33
DMinor accepts	66	0
DVerify accepts	62	2 (correct programs)

Chart 1 Speed comparison (average times for 66 well-typed test programs)

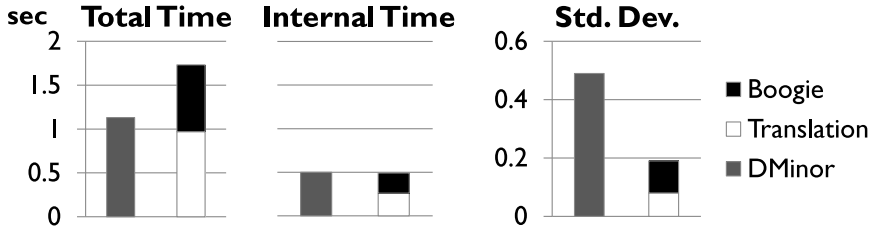

6 Comparison between Dminor and DVerify

We have tested our implementation against Dminor 0.1.1 from September 2010. Microsoft Research gave us access to their Dminor test suite that contains 109 sample programs. Out of these 109 tests 76 are well-typed Dminor programs and 33 are ill-typed. Out of the 76 well-typed programs the Dminor type-checker cannot verify 10 tests because of incompleteness.

As shown in Table 1, from the 66 cases on which Dminor succeeds, DVerify manages to verify 62 as correct. Out of the 33 that Dminor rejects, DVerify rejects 31. The other two are correct operationally, but are ill-typed with respect to the (inherently incomplete) Dminor type system. Overall this means that DVerify succeeds on 94% of the cases Dminor succeeds on and is able to verify two correct programs Dminor cannot verify. For the 4 correct programs that DVerify cannot verify the most common problem is that type-synthesis generates too complicated loop invariants and Z3 cannot handle the resulting proof obligations.

In order to compare efficiency, we first measured the overall wall-clock time that is needed by the two tools, which includes the time the operating system requires to start the process. Because we are dealing with a large number of small test files and both tools are managed .NET assemblies, initialisation dominates the total running times of both tools. Since initialisation is a constant factor that becomes negligible on bigger examples, we also measured the time excluding initialisation and parsing, which we call "internal time". Chart 1 shows both times (averaged over the 66 well-typed samples accepted by Dminor) on a 2.1 GHz laptop. The internal time is 0.5s on average for both Dminor and DVerify, which means that both tools are very efficient and that our combi-

Table 2 Qualitative comparison of Dminor and DVerify

Area	Our verification approach (DVerify)	Type-checking approach (DMinor)
Verification cond. generation	Weakest precondition	Bidirectional type-checking (type synthesis ≈ strongest postcondition)
Formulae discharged	One per postcondition/assertion (larger, but less obligations)	One per subtyping test (smaller, but more obligations)
Backend	Boogie + SMT-Solver (Z3)	SMT-Solver (Z3)
Loop invariants	In principle Boogie could infer some (even for accumulates) (even for global properties)	For from-where-select (but not for accumulates) (but not for global properties)
Error reporting	Abstract trace	Counterexample
Speed	similar	similar
Precision (practise)	similar	similar
Completeness (theory)	possibly better	possibly worse (type system)
Theories	equality, integers, datatypes, weak arrays	equality, integers, extensional arrays and native encoding of datatypes

nation of a translation and an off-the-shelf verification condition generator matches the average speed of a well-optimised type-checker on its own test suite.

One should keep in mind that all examples in this test suite are relatively small, the biggest one consisting of 90 lines. With bigger examples we expect DVerify to have a speed advantage over Dminor. This is for once due to the much lower standard deviation of DVerify, which indicates a better predictability. It is also reflected in the maximum internal time, where Dminor (3.97s) does three times worse than DVerify (1.35s). The qualitative comparison in Table 2 visualises a summary of this discussion.

7 Conclusion

In this paper we have presented a new technique for statically checking the typing constraints in Dminor programs by translating these programs into a standard while language and then using a general-purpose verification tool.

Future Work

Using a general verification tool for checking the types of Dminor programs should allow us to increase the expressivity of the Dminor language more easily. For example, adding support for *mutable state* would be easy in DVerify: Bemol already supports state, moreover Boogie is used mainly for imperative programming languages [9]. An interesting consequence is that it should be easier to support strong updates in DVerify (i.e. updates that change the type of variables), which is usually quite hard to achieve with a type-checker.

Another very interesting extension is *inferring loop invariants*. Dminor requires that each `accumulate` expression is annotated with a type for the accumulator which constitutes the invariant of the loop, whereas Boogie has build-in support for abstract interpretation for automatically inferring such invariants [3]. While the invariant inference support in Boogie seems currently very much focused on integer domains, it seems possible to extend it to include support for our Dminor types.

Finally, we would like to improve the *error reporting* capabilities of DVerify. When an assertion fails, Boogie produces an abstract execution trace that outlines which branches were taken to reach the failing assertion and where that assertion is located in the code [21]. In the future we would like to map this trace back to a Dminor trace, and produce errors in terms of the original Dminor program. More interestingly, we would also like to map the partial model produced by the SMT solver when it fails to prove a proof obligation, back as a potential counterexample assignment that maps variables to Bemol values. The Dminor type-checker [7] already implements this and it works quite well, but in Dminor this is implemented from scratch. For DVerify the IVL toolset could in principle provide more support for mapping back the models produced by the SMT solver to something that the end-user can understand.

Acknowledgements We thank Andrew D. Gordon for his helpful comments and the BOOGIE 2011 reviewers for their very useful feedback. Microsoft Research made our work much easier by making the Dminor source code available to us. Cătălin Hriţcu was supported by a fellowship from Microsoft Research and the International Max Planck Research School for Computer Science.

References

1. Bytecode level specification language and program logic. Mobius Project, Deliverable D3.1, 2006.
2. The Microsoft code name "M" Modeling Language Specification, October 2009. `http://msdn.microsoft.com/en-us/library/dd548667.aspx`.
3. M. Barnett, B.-Y. E. Chang, R. DeLine, B. Jacobs, and K. R. M. Leino. Boogie: A modular reusable verifier for object-oriented programs. In *4th International Symposium on Formal Methods for Components and Objects (FMCO)*, Lecture Notes in Computer Science, pages 364–387. Springer, 2005.
4. M. Barnett, K. Leino, and W. Schulte. The Spec# programming system: An overview. In *Workshop on Construction and Analysis of Safe, Secure and Interoperable Smart devices (CASSIS)*, pages 49–69. Springer, 2005.
5. B. Barras, S. Boutin, C. Cornes, J. Courant, Y. Coscoy, D. Delahaye, D. de Rauglaudre, J. Filliâtre, E. Giménez, H. Herbelin, et al. The Coq proof assistant reference manual, version 8.2. INRIA, 2009.
6. J. Bengtson, K. Bhargavan, C. Fournet, A. D. Gordon, and S. Maffeis. Refinement types for secure implementations. *ACM TOPLAS*, 33(2):8, 2011.
7. G. M. Bierman, A. D. Gordon, C. Hriţcu, and D. Langworthy. Semantic subtyping with an SMT solver. In *15th ACM SIGPLAN International Conference on Functional programming (ICFP 2010)*, pages 105–116. ACM Press, 2010.
8. S. Böhme, K. R. M. Leino, and B. Wolff. HOL-Boogie - an interactive prover for the Boogie program-verifier. In *21st International Conference on Theorem Proving in Higher Order Logics (TPHOLs)*, pages 150–166. Springer, 2008.

9. E. Cohen, M. Moskal, S. Tobies, and W. Schulte. A precise yet efficient memory model for C. *Electronic Notes in Theoretical Computer Science*, 254:85–103, 2009.
10. M. Dahlweid, M. Moskal, T. Santen, S. Tobies, and W. Schulte. VCC: Contract-based modular verification of concurrent C. In *31st International Conference on Software Engineering (ICSE)*, pages 429–430. IEEE, 2009.
11. R. DeLine and K. Leino. BoogiePL: A typed procedural language for checking object-oriented programs. Technical Report MSR-TR-2005-70, Microsoft Research, 2005.
12. D. Detlefs, G. Nelson, and J. Saxe. Simplify: A theorem prover for program checking. *Journal of the ACM (JACM)*, 52(3):473, 2005.
13. J.-C. Filliâtre and C. Marché. The Why/Krakatoa/Caduceus platform for deductive program verification. In *19th International Conference on Computer Aided Verification (CAV)*, pages 173–177. Springer, 2007.
14. H. Hosoya and B. Pierce. XDuce: A statically typed XML processing language. *ACM Transactions on Internet Technology*, 3(2):117–148, 2003.
15. R. Jhala, R. Majumdar, and A. Rybalchenko. HMC: Verifying functional programs using abstract interpreters. Accepted at CAV, 2011. To appear.
16. T. Kleymann. Hoare logic and auxiliary variables. *Formal Aspects of Computing*, 11(5):541–566, 1999.
17. K. Knowles, A. Tomb, J. Gronski, S. Freund, and C. Flanagan. SAGE: Unified hybrid checking for first-class types, general refinement types and `Dynamic`. Technical report, UCSC, 2007.
18. N. Kobayashi and C.-H. L. Ong. A type system equivalent to the modal mu-calculus model checking of higher-order recursion schemes. In *24th Annual IEEE Symposium on Logic in Computer Science (LICS)*, pages 179–188. IEEE Computer Society, 2009.
19. H. Lehner and P. Müller. Formal translation of bytecode into BoogiePL. *Electronic Notes in Theoretical Computer Science*, 190(1):35–50, 2007.
20. K. R. M. Leino. This is Boogie 2. TechReport, 2008.
21. K. R. M. Leino, T. Millstein, and J. Saxe. Generating error traces from verification-condition counterexamples. *Science of Computer Programming*, 55(1-3):209–226, 2005.
22. C. Marinos. An Introduction to Functional Programming for .NET Developers. *MSDN Magazine*, April 2010.
23. E. Meijer, B. Beckman, and G. M. Bierman. LINQ: reconciling object, relations and XML in the .NET framework. In *ACM SIGMOD International Conference on Management of Data (SIGMOD)*, page 706. ACM, 2006.
24. J. Morris. Comments on "procedures and parameters". Undated and unpublished.
25. M. Naik and J. Palsberg. A type system equivalent to a model checker. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 30(5):29, 2008.
26. T. Nipkow. Hoare Logics in Isabelle/HOL. In *Proof and System-Reliability*, pages 341–367. Kluwer, 2002.
27. B. Pierce, C. Casinghino, M. Greenberg, V. Sjöberg, and B. Yorgey. *Software Foundations*. `http://www.cis.upenn.edu/~bcpierce/sf/`, 2010.
28. S. Ranise and C. Tinelli. The satisfiability modulo theories library (SMT-LIB). `www.SMT-LIB.org`, 2006.
29. P. M. Rondon, M. Kawaguchi, and R. Jhala. Liquid types. In *ACM SIGPLAN 2008 Conference on Programming Language Design and Implementation (PLDI)*, pages 159–169, 2008.
30. N. Swamy, J. Chen, and R. Chugh. Enforcing stateful authorization and information flow policies in Fine. In *Proc. 19th European Symposium on Programming (ESOP 2010)*, pages 529–549, 2010.