A high-level modular definition of the semantics of C♯

Egon Börger[a], Nicu G. Fruja[b], Vincenzo Gervasi[a], Robert F. Stärk[b],*

[a]*Dipartimento di Informatica, Università di Pisa, Via F. Buonarroti 2, I-56127 Pisa, Italy*
[b]*Computer Science Department, ETH Zürich, CH-8092 Zürich, Switzerland*

Abstract

We propose a structured mathematical definition of the semantics of C♯ programs to provide a platform-independent interpreter view of the language for the C♯ programmer, which can also be used for a precise analysis of the ECMA standard of the language and as a reference model for teaching. The definition takes care to reflect directly and faithfully—as much as possible without becoming inconsistent or incomplete—the descriptions in the C♯ standard to become comparable with the corresponding models for Java in Stärk et al. (Java and Java Virtual Machine—Definition, Verification, Validation, Springer, Berlin, 2001) and to provide for implementors the possibility to check their basic design decisions against an accurate high-level model. The model sheds light on some of the dark corners of C♯ and on some critical differences between the ECMA standard and the implementations of the language.
Crown Copyright © 2004 Published by Elsevier B.V. All rights reserved.

Keywords: Semantics of programming languages; Abstract State Machines; C♯; Java; NET

1. Introduction

In this paper the method developed in [36] for a rigorous definition and analysis of Java and its implementation on the Java Virtual Machine (JVM) is applied to formalize the semantics of the entire language C♯. We provide a succinct, purely mathematical (thus platform-independent) model, which reflects as much as possible the intuitions and design decisions

* Corresponding author.
 E-mail addresses: boerger@di.unipi.it (E. Börger), fruja@inf.ethz.ch (N.G. Fruja), gervasi@di.unipi.it (V. Gervasi), staerk@inf.ethz.ch (R.F. Stärk).

underlying the language as described in the ECMA standard [27] and in [25] and can be used as accurate and complete reference model by C♯ programmers, by implementors of the language and by students learning it. In Section 8 we point to some challenging applications of the model for proving interesting theorems about C♯ and its implementations.

The model clarifies a certain number of semantically relevant issues which are not handled by the ECMA standard, wherefore we also consulted the Microsoft Press books [3,30,31] and the documentation in [18,29,34,37,39]. A series of bugs and gaps in the ECMA standard for C♯ and in its implementation in .NET and incoherences between the two were detected during our attempt to build for the language a consistent and complete yet abstract *ground model* (in the sense described in [8]). Some of them are mentioned in this paper to shed light on some dark corners of C♯, for a complete discussion we refer the reader to the companion paper [20]. As a rule we adhere to an established scientific tradition for which one of the goals of defining the meaning of programs is to accurately specify the freedom the compiler writer has for the implementation. Nevertheless, we also want our model to support the practice of programming. Therefore, whenever we see for a language construct an incoherence or a to-be-closed gap between on the one side the view offered by the ECMA standard, which should support the understanding also by programmers, and on the other side the view current compilers seem to have, we give in our model a pragmatic preference to abstractly defining what the programmer is allowed to expect from the execution of his code in the current implementations of C♯ [29,34,39]. In each case we explicitly discuss the discovered discrepancy so that the parameters of the design decision become clear. To support the experimentation with the model a project has been started to refine the model developed here to .NET-executable AsmL code [19], similarly to the AsmGofer refinement developed by Schmid [32,33] for the Java and JVM models in [36].

To provide the programmer with a transparent view of the intricate interaction of various language features which depend on the run-time environment, our model comes as an *abstract interpreter*, which provides a simple way to reflect those run-time-related features encountered upon executing a given C♯ program. To exploit the flexibility the use of Abstract State Machines (ASMs) offers in high-level system modeling and to obtain the faithfulness and simplicity of abstract models the ASM method allows one to achieve, the interpreter takes the form of an ASM. [1] This allows us in particular to specify the static and the dynamic parts of the semantics separately, due to the ASM classification of abstract states into a static and a dynamic part. The *dynamic semantics* of the language is captured operationally by ASM rules which describe the run-time effect of program execution on the abstract state of the program, the *static semantics* comes as a mainly declarative description of the relevant syntactical and compile-time checked language features (like typing rules, rules for definite assignment and reachability, name resolution, method resolution for overloaded methods, etc.) and of pre-processing directives (like #define, #undef, #if, #else, #endif, etc.), which are mostly reflected in the attributed abstract syntax tree our model starts from.

To keep the size of the models small and to facilitate the understanding of clusters of language constructs in terms of local state transformations, similarly to the decomposition

---

[1] See Section 8 for more information on our choice of ASMs among the many frameworks in the literature to deal with language semantics.

of Java and the JVM in [36] we structure the C♯ programming language into *layered modules of orthogonal language features*, namely

- the imperative core, related to sequential control by while programs, built from statements and expressions over the simple types of C♯,
- classes, realizing procedural abstraction with global (module) variables and class initialization,
- object-orientation with class instances, instance methods, inheritance,
- exception handling,
- delegates together with events (including here for convenience also properties, indexers, attributes),
- concurrency (threads),
- so-called unsafe code with pointers and pointer arithmetic.

This yields a sequence of sublanguages $C\sharp_\mathcal{I}$, $C\sharp_\mathcal{C}$, $C\sharp_\mathcal{O}$, $C\sharp_\mathcal{E}$, $C\sharp_\mathcal{D}$, $C\sharp_\mathcal{T}$, $C\sharp_\mathcal{U}$ which altogether describe the entire language C♯. Each language L in the sequence extends its predecessor and for each one we build a submachine EXECCSHARP$_L$ which is a conservative (purely incremental) extension of its predecessor. The model EXECCSHARP for the entire language C♯ is a composition of all submachines.

EXECCSHARP ≡
  EXECCSHARP$_I$
  EXECCSHARP$_C$
  EXECCSHARP$_O$
  EXECCSHARP$_E$
  EXECCSHARP$_T$
  EXECCSHARP$_D$
  EXECCSHARP$_U$

This approach supports a systematic piecemeal introduction of the numerous language constructs in teaching C♯ (or similar programming languages).

To keep the definition of the models succinct, we avoid tedious and routine repetitions concerning language constructs which can be reduced in well-known ways to the core constructs in our models. Whenever instead of a direct formalization of a construct we use a syntactical translation to constructs dealt with in the core model, we have to justify that the translation is correct with respect to the semantics of the construct as intended by the standard. The ASM model we define provides a basis to rigorously formulate and mathematically prove the intended equivalence. [2] Since such a justification follows well-known patterns, it is skipped in this paper, but to remind the reader of the problem we usually mention it.

The handling of truly concurrent threads, not limited to interleaving or similar simple scheduling techniques, is closely related to the underlying memory model. Since the description of this memory model goes much beyond this paper, the submodel $C\sharp_\mathcal{T}$ and its further analysis is postponed to a separate paper [35].

---

[2] One has to define an extension of the core model by a direct formalization of the construct in question and then to prove that this model is equivalent to the core model modulo the syntactical translation of the construct.

By and large one can correctly understand an ASM as pseudo-code operating over abstract data (structures in the sense of logic). Therefore we skip a detailed definition of ASMs, which is available in textbook form in Chapter 2 of the AsmBook [16]. Since our paper is not a tutorial or manual on C♯, we restrict our explanations of language constructs to features a reader will appreciate who is already knowledgeable about the basic concepts of object-oriented programming. In a technical report [12] also the remaining details which are skipped in this paper are spelt out completely, together with further explanations and examples.

The paper is structured by the modularization we propose for the language description. The basic framework of our model is introduced in Section 2 together with the interpreter for the imperative kernel C♯$_\mathcal{I}$ of the language. Successively one more section is added for each model refinement to capture the related language extension. In general, each section has a first part where the static assumptions of the model are formulated, followed by a second part which contains the dynamics expressed by the ASM transition rules operating on the corresponding state components. In general, at each layer the interpreter consists of two submachines, one defining expression evaluation and one defining statement execution.

2. The imperative core C♯$_\mathcal{I}$

In this section we define the model for C♯$_\mathcal{I}$, which defines the basic machinery of the ASM model for C♯. It describes the semantics of the sequential imperative core of C♯ with to be executed statements (appearing in method bodies) and to be evaluated expressions (appearing in statements) built using predefined operators over simple types. The computations of this interpreter are supposed to start with an arbitrary but fixed C♯ program. We separate syntax and compile-time matters from run-time issues by assuming that the program is given as an attributed syntax tree (i.e. annotated abstract syntax tree resulting from parsing and elaboration), trying to achieve model simplicity also by assuming some useful syntactical simplifications which will be mentioned as we build the model. Before defining the transition rules for the dynamic semantics of C♯$_\mathcal{I}$, we formulate what has to be said about the static semantics.

2.1. Static semantics of C♯$_\mathcal{I}$

We view the grammar in Fig. 1, which defines expressions and statements of the sublanguage C♯$_\mathcal{I}$, as defining also the corresponding ASM domains *Exp* and *Stm*. To avoid lengthy repetitions we include here already the distinctions between checked and unchecked expressions and blocks, though they are semantically irrelevant in the submodel C♯$_\mathcal{I}$ and start to play a role only with C♯$_\mathcal{E}$. The set *Vexp* of variable expressions (lvalues) consists in this model of the local variables only and will be refined below. *Sexp* denotes the set of statement expressions than can be used either as (result yielding) expressions or as (result discarding) statements, such as an assignment to a variable expression using '=' or an assignment operator from the set *Aop* or '++' or '--'. *Lit* denotes the set of literals, similarly for *Type*, *Lab* and the set *Cexp* of constant expressions whose value is known at compile time. When referring to the set of sequences of elements from a set *Item* we write *Items*, e.g. *Sexps*

$$
\begin{aligned}
Exp \quad &::= Lit \mid Vexp \mid Uop\ Exp \mid Exp\ Bop\ Exp \mid Exp\ \text{‘?’}\ Exp\ \text{‘:’}\ Exp \\
&\quad \mid (\ Type\ \text{‘)’}\ ExpSexp \mid \text{‘(’}\ Exp\ \text{‘)’} \mid \text{‘checked’}\ \text{‘(’}\ Exp\ \text{‘)’} \\
&\quad \mid \text{‘unchecked’}\ \text{‘(’}\ Exp\ \text{‘)’} \\[4pt]
Vexp \quad &::= Loc \\[4pt]
Sexp \quad &::= Vexp\ \text{‘=’}\ Exp \mid Vexp\ Aop\ Exp \mid Vexp\ \text{‘++’} \mid Vexp\ \text{‘--’} \\[4pt]
Uop \quad &::= \text{‘+’} \mid \text{‘-’} \mid \text{‘!’} \mid \text{‘\textasciitilde’} \\[4pt]
Bop \quad &::= \text{‘*’} \mid \text{‘/’} \mid \text{‘\%’} \mid \text{‘+’} \mid \text{‘-’} \mid \text{‘<<’} \mid \text{‘>>’} \mid \text{‘<’} \mid \text{‘>’} \mid \text{‘<=’} \mid \text{‘>=’} \mid \text{‘==’} \\
&\quad \mid \text{‘!=’} \mid \text{‘\&’} \mid \text{‘\textasciicircum’} \mid \text{‘|’} \\[4pt]
Aop \quad &::= \text{‘*=’} \mid \text{‘/=’} \mid \text{‘\%=’} \mid \text{‘+=’} \mid \text{‘-=’} \mid \text{‘<<=’} \mid \text{‘>>=’} \mid \text{‘\&=’} \mid \text{‘\textasciicircum=’} \mid \text{‘|=’} \\[4pt]
Stm \quad &::= \text{‘;’} \mid Sexp\ \text{‘;’} \mid \text{‘break’}\ \text{‘;’} \mid \text{‘continue’}\ \text{‘;’} \mid \text{‘goto’}\ Lab\ \text{‘;’} \\
&\quad \mid \text{‘if’}\ \text{‘(’}\ Exp\ \text{‘)’}\ Stm\ \text{‘else’}\ Stm \\
&\quad \mid \text{‘while’}\ \text{‘(’}\ Exp\ \text{‘)’}\ Stm \mid \text{‘do’}\ Stm\ \text{‘while’}\ \text{‘(’}\ Exp\ \text{‘)’} \\
&\quad \mid \text{‘for’}\ \text{‘(’}\ [Sexps]\ \text{‘;’}\ [Exp]\ \text{‘;’}\ [Sexps]\ \text{‘)’}\ Stm \\
&\quad \mid \text{‘switch’}\ \text{‘(’}\ Exp\ \text{‘)’}\ \text{‘\{’}\ \{Case\ \{Case\}\ Bstm\ \{Bstm\}\}\ \text{‘\}’} \\
&\quad \mid \text{‘goto’}\ \text{‘case’}\ Cexp\ \text{‘;’} \mid \text{‘goto’}\ \text{‘default’}\ \text{‘;’} \\
&\quad \mid \text{‘checked’}\ Block \mid \text{‘unchecked’}\ Block \mid Block \\[4pt]
Sexps \quad &::= Sexp\ \{\text{‘,’}\ Sexp\} \\[4pt]
Case \quad &::= \text{‘case’}\ Cexp\ \text{‘:’} \mid \text{‘default’}\ \text{‘:’} \\[4pt]
Block \quad &::= \text{‘\{’}\ \{Bstm\}\ \text{‘\}’} \\[4pt]
Bstm \quad &::= Type\ Loc\ \text{‘;’} \mid \text{‘const’}\ Type\ Loc\ \text{‘=’}\ Cexp\ \text{‘;’} \mid Lab\ \text{‘:’}\ Stm \mid Stm
\end{aligned}
$$

Fig. 1. Grammar of expressions and statements in C♯$_\mathcal{I}$.
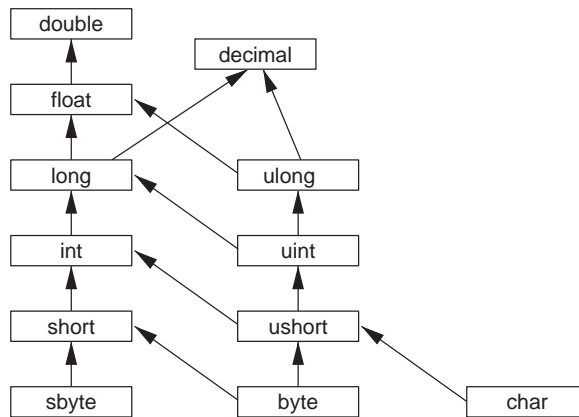
Fig. 2. The simple types of C♯$_\mathcal{I}$.

for the set of sequences of statement expressions. We usually write lower case letters *e* to denote elements of a set *E*, e.g. *lit* for elements of *Lit*.

The descriptions of implicit numeric conversions in [27, Section 13.1] and of binary numeric promotions in [27, Section 14.2.6] can be succinctly formulated as follows, using the type graph in Fig. 2 for the simple types of C♯, which are the types of C♯$_\mathcal{I}$ (for a classification of the types of C♯ see Fig. 4).

Definition 1 (*Implicit conversion [27, Section 13.1]*). We say that there exists an *implicit numeric conversion* from type A to B (written $A \prec B$) iff there exists a finite, non-empty path of arrows from A to B in the type graph in Fig. 2. We write $A \preccurlyeq B$ for $A \prec B$ or $A = B$. A type C is called an *upper bound* of A and B iff $A \preccurlyeq C$ and $B \preccurlyeq C$. A type C is the *least upper bound* of A and B iff

- C is an upper bound of A and B and
- $C \preccurlyeq D$ for each upper bound D of A and B.

We write $\sup(A, B)$ for the least upper bound of A and B if it exists.

We assume all the type constraints (on the operand and result values) and precedence conventions listed in [27] for the predefined (arithmetical, relational, bit and boolean logical) operators and the expression types. As usual each expression node *exp* in the attributed syntax tree has as attribute its compile-time type *type(exp)*.

About type conversions at compile-time we assume that type casts are inserted in the syntax tree if necessary. For example, if a binary numeric operator *bop* is applied to arguments in e_1 *bop* e_2, then the least upper bound T of the types of e_1 and e_2 must exist and the expression is transformed into $(T)e_1$ *bop* $(T)e_2$.

Definition 2 (*Binary numeric promotion [27, Section 14.2.6]*). The binary numeric promotion consists of applying the following rules:

- If the least upper bound of A and B exists, then
  - if $\sup(A, B) \preccurlyeq$ `int`, then A and B are converted to `int`,
  - otherwise, A and B are converted to $\sup(A, B)$.
- If the least upper bound of A and B does not exist, then a compile-time error occurs.

We also assume the syntactical constraints for statements listed in [27], e.g. the following ones for blocks (where the *scope of a local variable* (*local constant*) is defined as the block in which it is declared, the *scope of a label* is the block in which the label is declared, and a local variable is identified by its name *and* the position of its declaration, so that in particular local variables with the same name in disjoint blocks are considered as different):

- It is not allowed to refer to a local variable (local constant) in a textual position that precedes its declaration.
- It is not allowed to declare another local variable or local constant with the same name in the scope of a local variable (local constant).
- It is not allowed for two labels with the same name to have overlapping scopes.
- A `goto` *Lab* must be in the scope of a label with name *Lab*.
- Expressions in *constant declarations* are evaluated at compile-time.

To simplify the exposition of our model we assume some standard syntactical reductions as indicated in Table 1. The correctness of these replacements with respect to [27] can easily be checked on the basis of our semantics model for C♯.

2.1.1. Control-flow analysis

During the static program analysis where the compiler has to verify that the given program is well-typed, predicates *reachable* and *normal* with the following intended meaning are computed for statements, using the type information contained in the attributed syntax tree

Table 1
Standard syntactical reductions

exp_1 && exp_2	exp_1 ? exp_2 : false
exp_1 || exp_2	exp_1 ? true : exp_2
if (exp) stm	if (exp) stm else ;
++$vexp$	$vexp$ += 1
--$vexp$	$vexp$ -= 1
int x = 1, y, z = x * 2;	int x; x = 1; int y; int z; z = x * 2;
for (t loc = exp; tst; $step$) stm	{ t loc; for (loc = exp; tst; $step$) stm }

Table 2
Reachability rules for C♯$_\mathcal{I}$

s is a function body	$\Longrightarrow reachable(s)$
$reachable(;)$	$\Longrightarrow normal(;)$
$reachable(e;)$	$\Longrightarrow normal(e;)$
$reachable(\{\})$	$\Longrightarrow normal(\{\})$
$reachable(\{s\dots\})$	$\Longrightarrow reachable(s)$
$normal(s_i)$ in $\{\dots s_i\, s_{i+1}\dots\}$	$\Longrightarrow reachable(s_{i+1})$
$reachable(\texttt{goto}\,l;)$ in $\{\dots l{:}s\dots\}$	$\Longrightarrow reachable(l{:}s)$
$normal(s)$	$\Longrightarrow normal(\{\dots s\})$
$reachable(\texttt{if}\,(e)\,s_1\,\texttt{else}\,s_2) \land e \neq \texttt{false}$	$\Longrightarrow reachable(s_1)$
$reachable(\texttt{if}\,(e)\,s_1\,\texttt{else}\,s_2) \land e \neq \texttt{true}$	$\Longrightarrow reachable(s_2)$
$normal(s_1) \lor normal(s_2)$	$\Longrightarrow normal(\texttt{if}\,(e)\,s_1\,\texttt{else}\,s_2)$
$reachable(\texttt{while}\,(e)\,s) \land e \neq \texttt{false}$	$\Longrightarrow reachable(s)$
$reachable(\texttt{while}\,(e)\,s) \land e \neq \texttt{true}$	$\Longrightarrow normal(\texttt{while}\,(e)\,s)$
$reachable(\texttt{break};)$ in s	$\Longrightarrow normal(\texttt{while}\,(e)\,s)$

as the result of parsing and elaboration:

$$reachable(stm) \Longleftrightarrow stm \text{ can be reached}$$
$$normal(stm) \Longleftrightarrow stm \text{ can terminate normally}$$
$$\Longleftrightarrow \text{ the end point of } stm \text{ can be reached}$$

One of the language design goals was to guarantee the following two properties for programs to be accepted by the compiler:
- during the program execution, only *reachable* positions are reached,
- normal termination happens only in *normal* positions.

These two properties are obtained by checking two sufficient conditions via so-called reachability rules, which can be inductively defined for C♯$_\mathcal{I}$ in Table 2 (similarly for do, for, switch).[3] For constant boolean expressions in conditional and while statements we assume that they are replaced in the abstract syntax tree by true or false.

Unreachable statements indicate programming errors and therefore generate compile-time warnings. Function bodies that can terminate normally generated compile-time errors, since at run-time execution could fall off the bottom of the code array.

---

[3] We include these rules here to place the corresponding natural language specification in [27] on a firm ground for a mathematical proof of the above two properties as part of a type safety proof for C♯.
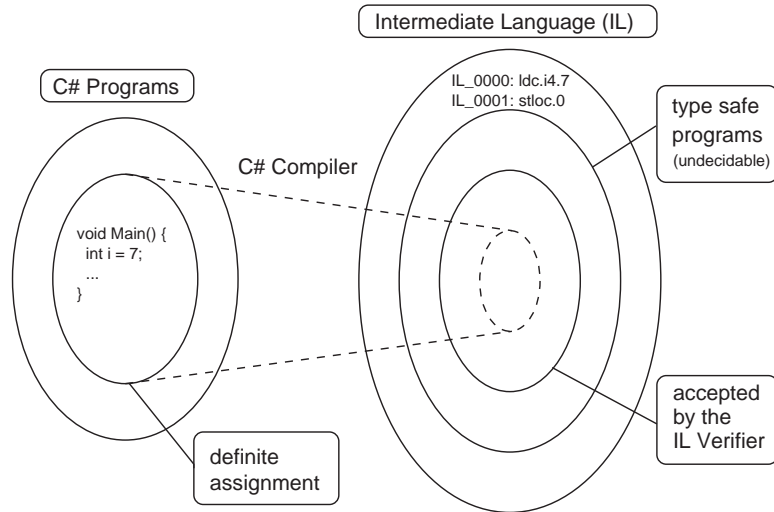
Fig. 3. Definite assignment and IL verification.

Another language design goal was to achieve the type safety of well-typed C♯ programs, i.e. that (a) variables at run-time contain values that are *compatible* with the declared types, and (b) expressions are evaluated at run-time to values that are *compatible* with their compile-time types. Among the desired consequences of the type safety of a program one has that at run-time its variables will never contain *undefined* values, that there are no *dangling* references, that the program cannot *corrupt* the memory, and that the dynamic *method lookup* always succeeds. Using the notation explained in the next section such invariants can be made precise and be proven to hold under appropriate assumptions. [4]

To guarantee the type safety the compiler checks a sufficient condition computing predicates *before*, *after* (for occurrences of statements and expressions in a function body) and *true*, *false* (for the two possible evaluation results of boolean expressions), which implement the so-called definite assignment rules to assure that a variable is *definitely assigned* before its value is used. The situation is illustrated in Fig. 3. Unfortunately the picture does not reflect reality. Microsoft has decided that in verified IL (intermediate language) code local variables are initialized by the run-time system with zero values. [5] Hence, also source code programs that do not fulfill the definite assignment constraints are accepted by the IL verifier.

A variable occurring in a position is called definitely assigned there, if on every execution path leading to that position (in the abstract syntax tree) a value is assigned to the variable.

---

[4] For example the following invariants can be proved to hold at run-time: (a) *before(pos)* ⊆ *Defined* where *Defined* = {$x \in Loc \mid mem(locals(x)) \neq Undef$}, (b) *after(pos)* ⊆ *Defined* if *values(pos)* = *Norm* or *values(pos)* ∈ *Value*. Specifically for boolean expressions holds *true(pos)* ⊆ *Defined* if *values(pos)* = *True*, the same for *false*. Such proofs can be carried out on the basis of the model developed in this paper, using the pattern developed in [36, Chapter 8] for proving that Java is type safe. For a different approach see [23].

[5] Maybe to simplify the job of the JIT verifiers, as one of our referees suggested.

Table 3
Definite assignment for statements

s is a function body	$before(s) = \emptyset$
`;`	$after(\texttt{;}) = before(\texttt{;})$
`exp;`	$before(exp\texttt{;}) = before(exp\texttt{;}), after(exp\texttt{;}) = after(exp)$
`break;`	$after(\texttt{break;}) = vars(\texttt{break;})$
`continue;`	$after(\texttt{continue;}) = vars(\texttt{continue;})$
`goto l;`	$after(\texttt{goto } l\texttt{;}) = vars(\texttt{goto } l\texttt{;})$
$stm = \{s_1 \ldots s_n\}$	$before(s_1) = before(stm), after(stm) = after(s_n),$
	$before(s_{i+1}) = after(s_i) \cap goto(s_{i+1})$ where
	$goto(l\texttt{:}s) =$
	$\quad \bigcap\{before(\texttt{goto } l\texttt{;}) \mid \texttt{goto } l\texttt{;} \text{ reachable in } stm\}$
	and $goto(s) = vars(s)$ if s is not a labeled statement
$stm = \texttt{if } (e)\ s_1 \texttt{ else } s_2$	$before(e) = before(stm), before(s_1) = true(e)$
	$before(s_2) = false(e), after(stm) = after(s_1) \cap after(s_2)$
$stm = \texttt{while } (e)\ s$	$before(e) = before(stm), before(s) = true(e),$
	$after(stm) = false(e) \cap break(s)$ where
	$break(s) = \bigcap\{before(\texttt{break;}) \mid \texttt{break;} \text{ reachable in } s\}$

Thus the intended meaning of the above predicates is as follows, where by "elaboration" of an *item* we mean "execution", if *item* is a statement, and "evaluation" if it is an expression:

$x \in before(item)$: x is definitely assigned *before* the elaboration of *item*

$x \in after(item)$: x is definitely assigned *after* normal elaboration of *item*

$x \in true(exp)$: x is definitely assigned *after exp* evaluates to *true*

$x \in false(exp)$: x is definitely assigned *after exp* evaluates to *false*

To provide a basis for a mathematical analysis, we turn the verbally stated definite assignment rules of [36, Section 12.3.3] into a precise set of equational constraints, where $vars(stm) = \{x \mid stm \text{ is in the scope of } x\}$.

Table 3 contains the constraints for the statements. Table 4 contains the equations for specific boolean expressions, which are imposed for the eager (short-circuit) evaluation of boolean expressions. Note that there is no equation in Table 4 for *after* sets since by definition $after(exp) = true(exp) \cap false(exp)$. If *exp* is a boolean expression which is not an instance of one of the expressions in Table 4, then the following are constraints for *exp*: $true(exp) = after(exp)$ and $false(exp) = after(exp)$.

Table 5 contains the equations for non-boolean expressions. In all other cases, if *exp* is an expression which has the *direct subexpressions* $e_1, e_2, \ldots, e_n$, then the left-to-right evaluation scheme yields

- $before(e_1) = before(exp)$,
- $before(e_{i+1}) = after(e_i)$ for $i \in [1 .. n - 1]$,
- $after(exp) = after(e_n)$.

Due to the goto statement the above constraints do not specify in a unique way the sets of variables that have to be considered as definitely assigned. Consider the following block (from [21]):

```
{ int i = 1 ; L : goto L ; }
```

Table 4
Definite assignment for boolean expressions

`true`	$true(\texttt{true}) = before(\texttt{true}), false(\texttt{true}) = vars(\texttt{true})$
`false`	$true(\texttt{false}) = vars(\texttt{false}), false(\texttt{false}) = before(\texttt{false})$
$exp = \;!e$	$before(e) = before(exp), true(exp) = false(e)$
	$false(exp) = true(e)$
$exp = (e_1 \;\&\&\; e_2)$	$before(e_1) = before(exp), before(e_2) = true(e_1),$
	$true(exp) = true(e_2), false(exp) = false(e_1) \cap false(e_2)$
$exp = (e_1 \;|\,|\; e_2)$	$before(e_1) = before(exp), before(e_2) = false(e_1),$
	$true(exp) = true(e_1) \cap true(e_2), false(exp) = false(e_2)$
$exp = (e_0 \;?\; e_1 : e_2)$	$before(e_0) = before(exp), before(e_1) = true(e_0)$
	$before(e_2) = false(e_0), true(exp) = true(e_1) \cap true(e_2)$
	$false(exp) = false(e_1) \cap false(e_2)$

Table 5
Definite assignment for arbitrary expressions

loc	$loc \in before(loc), after(loc) = before(loc)$
lit	$after(lit) = before(lit)$
$exp = (loc = e)$	$before(e) = before(exp), after(exp) = after(e) \cup \{loc\}$
$exp = (loc \; op\!= e)$	$loc \in before(exp), before(e) = before(exp)$
	$after(exp) = after(e)$
$exp = (e_0 \;?\; e_1 : e_2)$	$before(e_0) = before(exp), before(e_1) = true(e_0)$
	$before(e_2) = false(e_0), after(exp) = after(e_1) \cap after(e_2)$

Then the constraints of the definite assignment analysis are satisfied for both $before(\texttt{L: goto L;}) = \emptyset$ and $before(\texttt{L: goto L;}) = \{i\}$. Hence during the analysis the greatest sets of variables that satisfy the constraints for *before* and *after* have to be computed (cf. [21]). For blocks without goto statements, however, it can be proved from the above axioms that the *before* set determines the *after* set in a unique way.

2.2. Dynamic semantics for C♯$_\mathcal{I}$

The dynamic semantics for C♯$_\mathcal{I}$ describes the effect of statement execution and expression evaluation upon the program state, so that the transition rule for C♯$_\mathcal{I}$ (the same for its extensions) has the form

EXECCSHARP$_I$ ≡
    EXECCSHARPEXP$_I$
    EXECCSHARPSTM$_I$

The first subrule defines one execution step in the evaluation of expressions; the second subrule defines one step in the execution of statements.

To make the further model refinements possible via purely incremental extensions, our definition proceeds by walking through the attributed syntax tree and computing at each

node the effect of the program construct attached to the node. We formalize the walk by a cursor ▶, whose position in the tree—represented by a dynamic function *pos*: *Pos*—is updated using static tree functions, leading from a node in the tree down to its *first* child, from there to the *next* brother or *up* to the parent node (if any), as illustrated by the following self-explanatory example. The moves of *pos* contain implicitly the control-flow graph of C♯. A function *label*: *Pos* → *Label* decorates nodes with the information which identifies the grammar rule associated to the node. For the sake of notational succinctness we use concrete syntax of C♯ to describe the labels, thus avoiding the explicit introduction of auxiliary non-terminals the reader probably does not want to see. In the following example the four possible cursor positions are reachable from the root by following the tree functions *first*, *next* and *up*. The *label* of the root node is the auxiliary non-terminal *If*, identifying the grammar rule which produces in one step if (*exp*) *stm₁* else *stm₂*.

For updating the values of local variables in the memory we use two dynamic functions *locals*: *Loc* → *Adr* and *mem*: *Adr* → *SimpleValue* ∪ {*Undef*}, which assign to local variables memory addresses and store the values there. Since in C♯$_\mathcal{I}$ the values are of simple types, the equation *Value* = *SimpleValue* ∪ *Adr* holds, which will be refined in the extended models to include also references and structs. The uniquely identified local variables are modeled by stipulating *Loc* = *Identifier* × *Pos*, where *Pos* is the set of positions in the abstract syntax tree.

The indirection through memory addresses is not really needed in C♯$_\mathcal{I}$. In C♯$_\mathcal{I}$ one could assign values directly to local variables without storing them in an abstract memory. The addresses, however, are needed later for call-by-reference with ref and out parameters (one of the major differences between C♯ and Java from the modelling point of view).

Statements can terminate normally or abruptly, where in C♯$_\mathcal{I}$ the reasons of abruption are from the set *Abr* = *Break* | *Continue* | *Goto*(*Lab*), to be refined for the extended models. We use a dynamic function *values*: *Pos* → *Result* to store intermediate evaluation results from the set

$$Result = Value \cup Abr \cup \{Undef, Norm\}.$$

For the initial state we assume
- *mem*(*i*) = *Undef* for every *i* ∈ *Adr*,
- *pos* = root position of the attributed syntax tree,
- *locals*(*x*) ∈ *Adr* for every variable *x*. [6]

As intermediate *values* at a position *p* the cursor is at or is passing to, the computation may yield directly a simple value; at *AddressPos*itions as defined below it may yield an address;

---

[6] This amounts to assuming that the compiler chooses an address for each variable.

but it may also yield a *memValue* which has to be retrieved indirectly via the given address (where for C♯$_\mathcal{I}$ the memory value of a given type t at a given address *adr* is defined by *memValue*$(adr, t) = mem(adr)$; the parameter t will become relevant only in the refinement of *memValue* in C♯$_\mathcal{O}$ and C♯$_\mathcal{U}$). This is described by the following two macros:

YIELD(*val*, p) ≡
   values(p) := *val*
   pos := p

YIELDINDIRECT(*adr*, p) ≡
   if *AddressPos*(p) **then** YIELD(*adr*, p)
   else YIELD(*memValue*(*adr*, *type*(p)), p)

We will use the macros in the two forms YIELD(*val*) ≡ YIELD(*val*, *pos*) and YIELDUP(*val*) ≡ YIELD(*val*, *up*(*pos*)). Similarly we have two forms also for the second macro: YIELDINDIRECT(*adr*) and YIELDUPINDIRECT(*adr*).

Being in a context where an address and not a value is required can be defined as follows:

$$AddressPos(\alpha) \iff FirstChild(\alpha) \land$$
$$(label(up(\alpha)) \in \{++, --\} \lor label(up(\alpha)) \in Aop)$$
where $FirstChild(\alpha) \iff first(up(\alpha)) = \alpha$

To further reduce any notational overhead not needed by the human reader, in spelling out the rules below we identify positions with the occurrences of the syntactical constructs nodes are decorated with. This explains updates like *pos* := *exp* or *pos* := *stm*, which are used as shorthand for updating *pos* to the node labeled with the corresponding occurrence of *exp* respectively *stm*. [7] Furthermore, for a succinct formulation we use a macro *context*(*pos*) to describe the context of the currently to be handled expression or statement or intermediate result, which has to be matched against the syntactically possible cases (in the textual order of the rule) to select the appropriate computation step. If the elaboration of the subtree at the position *pos* has not yet started, then *context*(*pos*) is the construct encoded by the labels of *pos* and of its children. Otherwise, if *pos* carries already its result in *values*, *context*(*pos*) is the pseudo-construct encoded by the labels of the parent node of *pos* and of its children after replacing the already evaluated constructs by their *values* in the corresponding node. This explains notations like *uop* ▶ *val* to describe the *context* of *pos*, where *pos* is marked with the cursor (▶), resulting from the successful evaluation of the argument *exp* of the construct *uop exp* (encoded by *up*(*pos*) and its child *pos*), just before *uop* is applied to *val* to YIELDUP(*Apply*(*uop*, *val*)).

It thus remains to define the two submachines for expression evaluation and statement execution. This is done in a modular fashion, grouping behaviorally similar instructions into one parameterized instruction. [8]

---

[7] An identification of this kind, which is common in mathematics, has clearly to be resolved in an executable version of the model. See for example the formulation of the ASM model for Java in [36].

[8] The specializations can be regained instruction-wise by mere parameter expansion, a form of refinement that is easily proved to be correct.

2.2.1. Expression evaluation rules

We are now ready to define the machine EXECCSHARPEXP$_I$ in a compositional way, namely proceeding expression-wise: for each syntactical form of expressions there is a set of rules covering each intermediate phase of their evaluation. The machine passes control from unevaluated expressions to the appropriate subexpressions until an atom (a literal or a local variable) is reached. It can continue its computation only as long as no operator exception occurs, as a consequence it does not distinguish between checked and unchecked expression evaluation—the extension by rules to handle exceptions is defined in the model extension C♯$_\mathcal{E}$. The expressions for numeric casts will be refined in C♯$_\mathcal{O}$ and in C♯$_\mathcal{E}$. The macro WRITEMEM(adr, t, val) denotes here $mem(adr) := val$; it will be refined in the model for C♯$_\mathcal{O}$.

EXECCSHARPEXP$_I$ ≡ **match** $context(pos)$

　$lit \to$ YIELD($ValueOfLiteral(lit)$)
　$loc \to$ YIELDINDIRECT($locals(loc)$)

　$uop\ exp\quad \to pos := exp$
　$uop\ ▶val \to\quad$ **if** $\neg UopException(uop, val)$ **then** YIELDUP($Apply(uop, val)$)

　$exp_1\ bop\ exp_2\quad \to pos := exp_1$
　$▶val\ bop\ exp\quad \to pos := exp$
　$val_1\ bop\ ▶val_2 \to\ $ **if** $\neg BopException(bop, val_1, val_2)$ **then**
　　　　　　　　　　　YIELDUP($Apply(bop, val_1, val_2)$)

　$exp_0\ ?\ exp_1\ :\ exp_2 \to pos := exp_0$
　$▶val\ ?\ exp_1\ :\ exp_2 \to\quad$ **if** val **then** $pos := exp_1$ **else** $pos := exp_2$
　$True\ ?\ ▶val\ :\ exp \to$ YIELDUP(val)
　$False\ ?\ exp\ :\ ▶val \to$ YIELDUP(val)

　$loc = exp\quad \to pos := exp$
　$loc = ▶val \to \{$WRITEMEM($locals(loc), type(loc), val$), YIELDUP($val$)$\}$

　$(type)\ exp\quad \to pos := exp$
　$(type)\ ▶val \to$ **if** $type(pos) \in NumericType \wedge type \in NumericType$ **then**
　　　　　　　　　if $\neg UopException(type, val)$ **then**
　　　　　　　　　　YIELDUP($Convert(type, val)$)

　$vexp\ op=\ exp\quad \to pos := vexp$
　$▶adr\ op=\ exp \to pos := exp$
　$adr\ op= ▶val \to$ **let** $t = type(up(pos))$ **and** $v = memValue(adr, t)$ **in**
　　　　　　　　　if $\neg BopException(op, v, val)$ **then**
　　　　　　　　　　let $w = Apply(op, v, val)$ **and** $r = Convert(t, w)$ **in**
　　　　　　　　　　　$\{$WRITEMEM(adr, t, r), YIELDUP(r)$\}$

　$vexp\ op\quad \to pos := vexp$// for postfix operators $op \in \{++, --\}$
　$▶adr\ op \to$ **let** $old = memValue(adr, type(pos))$ **in**
　　　　　　　if $\neg UopException(op, old)$ **then**
　　　　　　　　WRITEMEM($adr, type(up(pos))$
　　　　　　　　$Apply(op, old)$)
　　　　　　　　YIELDUP(old)

$$
\begin{aligned}
\texttt{checked}(exp) &\rightarrow pos := exp \\
\texttt{checked}(\blacktriangleright val) &\rightarrow \text{YIELDUP}(val) \\
\texttt{unchecked}(exp) &\rightarrow pos := exp \\
\texttt{unchecked}(\blacktriangleright val) &\rightarrow \text{YIELDUP}(val)
\end{aligned}
$$

Being in a checked context is used to define whether operators throw an overflow exception (in which case a rule will be added in the model for $C\sharp_{\mathcal{E}}$). The general rule is that operators for the type `decimal` always throw overflow exceptions whereas operators for integral types only throw overflow exceptions in a checked context except for the division by zero. By default every position is unchecked, unless explicitly declared otherwise.

$Checked(\alpha) \iff label(\alpha) = Checked \vee$
   $(label(\alpha) \neq Unchecked \wedge up(\alpha) \neq Undef \wedge Checked(up(\alpha)))$

$UopException(uop, val) \iff Checked(pos) \wedge Overflow(uop, val)$

$BopException(bop, val_1, val_2) \iff$
   $DivisonByZero(bop, val_2) \vee DecimalOverflow(bop, val_1, val_2) \vee$
   $(Checked(pos) \wedge Overflow(bop, val_1, val_2))$

2.2.2. Statement execution rules

The machine EXECCSHARPSTM$_I$ is defined below statement-wise. It transfers control from structured statements to the appropriate substatements, until the current statement has been computed normally or abrupts the computation. Abruptions trigger the control to propagate through all the enclosing statements up to the target labeled statement. The concept of propagation is defined for $C\sharp_{\mathcal{I}}$ in such a way that in the refined models it is easily extended to abruptions due to return from procedures or to exceptions. [9] In case of a new execution of the body of a while statement, the previously computed intermediate results have to be cleared. [10] For the sake of brevity we skip the analogous transition rules for statements `do`, `for`, `switch`, `goto case`, `goto default`. Since we formulate the model for the human reader, we also use the ...-notation, for example in the rules for abruption or for sequences of block statements. This avoids having to fuss with an explicit formulation of the context, typically determined by a walk through a list. This simplification, which is tailored for the human reader, can easily be resolved for an executable model version without increasing the number of rules.

EXECCSHARPSTM$_I$ $\equiv$ **match** $context(pos)$
   $;\ \rightarrow$ YIELD$(Norm)$

   $exp;\ \rightarrow pos := exp$
   $\blacktriangleright val;\ \rightarrow$ YIELDUP$(Norm)$

---

[9] For $C\sharp_{\mathcal{I}}$ alone it would be simpler to transfer control directly by updating *pos* to the value of a corresponding static function.

[10] CLEARVALUES is needed in the present rule formulation, which is close to an executable format. In a more abstract SOS-style, as used for the Java model in [36], it would not be necessary because there the intermediate *values* can be written into a dynamic function for the still to be executed rest program.

```
break;       → YIELD(Break)
continue;    → YIELD(Continue)
goto lab;    → YIELD(Goto(lab))
```

if (exp) stm_1 else stm_2 $\rightarrow pos := exp$
if $(\blacktriangleright val)$ stm_1 else stm_2 $\rightarrow$ **if** val **then** $pos := stm_1$ **else** $pos := stm_2$
if $(True)$ $\blacktriangleright Norm$ else stm $\rightarrow$ YIELDUP$(Norm)$
if $(False)$ stm else $\blacktriangleright Norm$ $\rightarrow$ YIELDUP$(Norm)$

while (exp) stm $\rightarrow pos := exp$
while $(\blacktriangleright val)$ stm $\rightarrow$ **if** val **then** $pos := stm$
                                  else YIELDUP$(Norm)$
while $(True)$ $\blacktriangleright Norm$ $\rightarrow \{pos := up(pos),$ CLEARVALUES$(up(pos))\}$
while $(True)$ $\blacktriangleright Break$ $\rightarrow$ YIELDUP$(Norm)$
while $(True)$ $\blacktriangleright Continue$ $\rightarrow \{pos := up(pos),$ CLEARVALUES$(up(pos))\}$
while $(True)$ $\blacktriangleright abr$ $\rightarrow$ YIELDUP(abr)

type loc; $\rightarrow$ YIELD$(Norm)$

lab: stm $\rightarrow pos := stm$
lab: $\blacktriangleright Norm$ $\rightarrow$ YIELDUP$(Norm)$

```
checked block        → pos := block
checked ▶Norm        → YIELDUP(Norm)
unchecked block      → pos := block
unchecked ▶Norm      → YIELDUP(Norm)
```

$\ldots \blacktriangleright abr \ldots \rightarrow$ **if** $up(pos) \neq Undef \wedge PropagatesAbr(up(pos))$ **then**
                  YIELDUP(abr)

$\{\ \}$ $\rightarrow$ YIELD$(Norm)$
$\{stm \ldots \}$ $\rightarrow pos := stm$
$\{\ldots \blacktriangleright Norm\}$ $\rightarrow$ YIELDUP$(Norm)$
$\{\ldots \blacktriangleright Norm \; stm \ldots \}$ $\rightarrow pos := stm$
$\{\ldots \blacktriangleright Goto(l) \ldots \}$ $\rightarrow$ **let** $\alpha = GotoTarget(first(up(pos)), l)$
                         if $\alpha \neq Undef$ **then**
                             $\{pos := \alpha,$ CLEARVALUES$(up(pos))\}$
                         else YIELDUP$(Goto(l))$
$\{\ldots \blacktriangleright abr \ldots \}$ $\rightarrow$ YIELDUP(abr)

In $C\sharp_{\mathcal{I}}$ abruptions are propagated upwards except at the following statements:

$PropagatesAbr(\alpha) \iff label(\alpha) \notin \{Block, While, Do, For, Switch\}$

To compute the target of a label in a list of block statements we define:

$GotoTarget(\alpha, l) =$
  if $label(\alpha) = Lab(l)$ **then** α
  elseif $next(\alpha) = Undef$ **then** $Undef$
  else $GotoTarget(next(\alpha), l)$

The auxiliary macro CLEARVALUES(α) to clear all values in the subtree at position α can be defined by recursion as follows, proceeding from top to bottom and from left to right: [11]

CLEARVALUES(α) ≡
   values(α) := *Undef*
   if *first*(α) $\neq$ *Undef* **then** CLEARVALUESSEQ(*first*(α))

CLEARVALUESSEQ(α) ≡
   CLEARVALUES(α)
   if *next*(α) $\neq$ *Undef* **then** CLEARVALUESSEQ(*next*(α))

3. C♯$_\mathcal{C}$: refining C♯$_\mathcal{I}$ by static class features

In this section we refine the imperative core C♯$_\mathcal{I}$ to C♯$_\mathcal{C}$ by adding classes (modules) concentrating upon their static features (static fields, methods, constructors), including their initialization and the parameter mechanism that provides value, `ref` and `out` parameters. For such a refinement we (a) extend the ASM universes and functions, or introduce new ones, to reflect the grammar extensions for expressions and statements, (b) add the appropriate constraints needed for the static analysis of the new items (type constraints, definite assignment rules), (c) extend some of the macros, e.g. *PropagatesAbr*(α), to make them work also for the newly occurring cases, (d) add rules which define the semantics of the new instructions that operate over the new domains.

3.1. Static semantics of C♯$_\mathcal{I}$

In C♯$_\mathcal{C}$ a program is a set of compilation units, each coming with "using directives" and declarations of names spaces (including a global namespace) and types (for classes and interfaces [12]) in the global namespace. For simplicity of exposition, we disregard "using" directives and nested namespaces by assuming everywhere the adoption of (equivalent) fully qualified names. The precise syntax of classes and their static members, the rules for the accessibility of types and members via the access modifiers (public, internal, protected, private) and illustrating examples are spelt out in [12]. We define here the extension of the grammars for *Vexp*, *Sexp*, *Stm* and thereby of the corresponding ASM domains, which reflects the introduction of sets of *Class*es with static *Field*s and static *Meth*ods in C♯$_\mathcal{C}$. The new set *Arg* of arguments appearing here reflects that besides value parameters also `ref` and `out` parameters can be used.

   Vexp ::= … | *Field* | *Class* '.' *Field*
   Sexp ::= … | *Meth* ([*Args*]) | *Class* '.' *Meth* ([*Args*])

---

[11] Intuitively it should be clear that the execution of this recursive ASM provides simultaneously—in one step—the set of all updates of all its recursive calls, as is needed here for the clearing purpose; see [10] for a precise definition.

[12] Note that struct and enum types and delegates are introduced by further refinement steps below.

Arg   ::= *Exp* | 'ref' *Vexp* | 'out' *Vexp*

Args ::= *Arg* {',' *Arg*}

Stm   ::= … | 'return' *Exp* ';' | 'return' ';'

The type constraints for the new expressions and the return statement are spelt out in [12]. The difference between ref and out parameters at function calls and in function bodies is reflected by including as *AddressPos*itions all nodes whose parent node is labeled by ref or out and by adding the following definite assignment constraints:

- ref arguments must be definitely assigned *before* the function call.
- out arguments are definitely assigned *after* the function call.
- ref parameters and value parameters of a function are definitely assigned at the beginning of the function body.
- out parameters must be definitely assigned when the function returns.

Therefore the definite assignment constraints for expressions are extended by the following constraints for general argument expressions in function calls and for ref and out argument expressions:

- For *exp* = *M*(*args*):
  - *before*(*args*) = *before*(*exp*)
  - *RefParams*(*args*) ⊆ *after*(*args*)
  - *after*(*exp*) = *after*(*args*) ∪ *OutParams*(*args*)
- For *exp* = (ref *e*) or *exp* = (out *e*):
  - *before*(*e*) = *before*(*exp*)
  - *after*(*exp*) = *after*(*e*)

The definite assignment constraints for statements are extended for function bodies and return statements as follows:

- If *s* is the body of *M*, then
  - *before*(*s*) = *ValueParams*(*M*) ∪ *RefParams*(*M*)
- If *stm* = return; is in *M*, then
  - *OutParams*(*M*) ⊆ *before*(*stm*)
  - *after*(*stm*) = *vars*(*stm*)
- If *stm* = return *e*; is in *M*, then
  - *before*(*e*) = *before*(*stm*)
  - *OutParams*(*M*) ⊆ *after*(*e*)
  - *after*(*stm*) = *vars*(*stm*)

The presence of to-be-initialized classes and of method calls is reflected by the introduction of new universes to denote methods, the initialization status of a type (which will be refined below by exceptions) and the sequence of still active method calls (frame stack):

$$Meth = Type \times Msig$$
$$TypeState = Linked \mid InProgress \mid Initialized$$
$$Frame = Meth \times Pos \times Locals \times Values$$
$$\text{where } Values = (Pos \rightarrow Result) \text{ and } Locals = (Loc \rightarrow Adr)$$

A method signature *Msig* consists of the name of a method plus the sequence of types of the arguments of the method. A method is uniquely determined by the type in which it is declared and its signature. The reasons for abruptions are extended by

method return:

$$Abr = \ldots \mid Return \mid Return(Value)$$

3.2. Dynamic semantics of $C\sharp_{\mathcal{I}}$

To dynamically handle the addresses of static fields (global or class variables), the initialization state of types, the current method, the execution stack and the (initially) to be initialized type we use the following new dynamic functions:

$$globals : Type \times Field \rightarrow Adr \qquad frames : List(Frame)$$
$$typeState : Type \rightarrow TypeState \qquad meth : Meth$$

We extend the stipulations for the initial state as follows:

- $typeState(c) = Linked$ for each class c
- $meth = EntryPoint::\texttt{Main()}$ [*EntryPoint* is the main class]
- $pos = body(meth)$ [The root position of the body]
- $locals = values = \emptyset$ and $frames = []$

The submachine EXECCSHARP$_C$ extends the machine EXECCSHARP$_I$ for $C\sharp_{\mathcal{I}}$ by additional rules for the evaluation of the new expressions and for the execution of return statements. In the same way the further refinements in the sections below consist in the parallel addition of appropriate submachines.

EXECCSHARP$_C$ ≡
    EXECCSHARP$_I$
    EXECCSHARPEXP$_C$
    EXECCSHARPSTM$_C$

3.2.1. Expression evaluation rules

The rules for class field evaluation in the submachine EXECCSHARPEXP$_C$ are analogous to those for the evaluation of local variables in EXECCSHARPEXP$_I$, except for using *globals* instead of *locals* and for the additional clause for class initialization. The rules for method calls use the macro INVOKESTATIC defined below and reflect that the arguments are evaluated from left to right. [13]

EXECCSHARPEXP$_C$ ≡ **match** *context*(*pos*)
    $c.f \rightarrow$ **if** *Initialized*(*c*) **then** YIELDINDIRECT(*globals*(*c*::*f*))
            else INITIALIZE(*c*)

    $c.f = exp \rightarrow pos := exp$
    $c.f = {}^{\blacktriangleright}val \rightarrow$ **if** *Initialized*(*c*) **then**
                    WRITEMEM(*globals*(*c*::*f*), *type*(*c*::*f*), *val*)
                    YIELDUP(*val*)
                else INITIALIZE(*c*)

---

[13] These are the rules to be modified in case one wants to specify another evaluation order for expressions, involving the use of the ASM **choose** construct if some non-deterministic choice has to be formulated. For a discussion of such model variations we refer to [40] where an ASM model is developed which can be instantiated to capture the different expression evaluation strategies in Ada95, C, C++, Java, C♯ and Fortran.

$c.m\,(\,args\,)\quad \rightarrow pos := (\,args\,)$
$c.m^{\blacktriangleright}(\,vals\,) \rightarrow \text{INVOKESTATIC}(c::m,\,vals)$

`ref` $vexp\;\; \rightarrow pos := vexp$
`ref` $^{\blacktriangleright}adr \rightarrow \text{YIELDUP}(adr)$

`out` $vexp\;\; \rightarrow pos := vexp$
`out` $^{\blacktriangleright}adr \rightarrow \text{YIELDUP}(adr)$

$(\,)\qquad\qquad\qquad\quad\; \rightarrow \text{YIELD}([])$
$(\,arg\,,\,\dots\,)\qquad\qquad \rightarrow pos := arg$
$(\,val_1\,,\,\dots\,,^{\blacktriangleright}val_n\,) \rightarrow \text{YIELDUP}([val_1,\,\dots,\,val_n])$
$(\,\dots\,^{\blacktriangleright}val\,,arg\dots\,) \rightarrow pos := arg$

The macro INVOKESTATIC invokes the method if the initialization of its class is not triggered, otherwise it initializes the class. The initialization of a class (or struct, see Section 4) is not triggered if the class is already initialized.[14] For methods which are not declared external, INVOKEMETHOD updates the frame stack and the current frame in the expected way, allocating via INITLOCALS for every local variable or value parameter a new address and copying every value argument there. Since we will also have to deal with external methods— whose declaration includes an `extern` modifier and which may be implemented using a language other than C♯—we provide here for their invocation a submachine INVOKEEXTERN, to be defined separately depending on the class of external (e.g. library) methods.[15] The predicate *StaticCtor* recognizes static class constructors; their implicit call interrupts the member access at *pos*, to later return to the evaluation of *pos* instead of *up*(*pos*). We separate the current frame—consisting of *meth*, *pos*, *locals* and *values*—from the stack of such frames to notationally smooth the transition from C♯$_\mathcal{I}$ to C♯$_\mathcal{C}$.

INVOKESTATIC($c::m$, *vals*) $\equiv$
   if not *TriggerInit*(c) **then** INVOKEMETHOD($c::m$, *vals*)
   else INITIALIZE(c)
   where *TriggerInit*(c) $= \neg$*Initialized*(c)

INVOKEMETHOD($c::m$, *vals*) $\equiv$
   if `extern` $\in$ *modifiers*($c::m$) **then** INVOKEEXTERN($c::m$, *vals*)
   else let $p =$ **if** *StaticCtor*($c::m$) **then** *pos* **else** *up*(*pos*) **in**
     frames $:= push(frames, (meth, p, locals, values))$
     meth $:= c::m$
     pos $:= body(c::m)$
     values $:= \emptyset$
     INITLOCALS($c::m$, *vals*)

---

[14] As analyzed in [20], it is also not triggered if the class is marked with the implementation flag *beforefieldinit*, indicating that the reference of the static method does not trigger the class (or struct) initialization. If one wants to model this flag the definition has to be refined to *TriggerInit*(c) $= \neg$*Initialized*(c) $\wedge \neg$*beforefieldinit*(c) and furthermore in Section 5.

[15] For an illustration of this use of external methods see below the model for delegates.

The following definition for the initialization of local variables reflects that C♯ permits to pass function call parameters by value, as Java does, but also by `reference`. Also `out` parameters are allowed, treated as `ref` parameters except that they need not be definitely assigned until the function call returns. [16]

In the following definition, all (also simultaneous) applications of the external function *new* during the computation of the ASM are supposed to provide pairwise different fresh elements from the underlying domain *Adr*. See [22,16, 2.4.4] for a justification of this assumption. See also the end of Section 4 where we provide an abstract specification of the needed memory allocation to addresses of references and objects of struct type and to their instance fields. *paramIndex*(c::m, x) yields the index of the formal parameter x in the signature of c::m.

INITLOCALS(c::m, *vals*) ≡
  forall $x \in$ *LocalVars*(c::m) **do**   // addresses for local variables
    locals(x) := *new*(*Adr*, *type*(x))
  forall $x \in$ *ValueParams*(c::m) **do**   // copy value arguments
    let *adr* = *new*(*Adr*, *type*(x)) **in**
      locals(x) := *adr*
      WRITEMEM(*adr*, *type*(x), *vals*(*paramIndex*(c::m, x)))
  forall $x \in$ *RefParams*(c::m) ∪ *OutParams*(c::m) **do**   // ref, out arguments
    locals(x) := *vals*(*paramIndex*(c::m, x))

3.2.2. Statement execution rules

The rules for method return in the submachine EXECCSHARPSTM$_C$ trigger an abruption upon returning from a method, resulting (via the propagation of this abruption to the beginning of the method body where it occurred) in the execution of the machine EXITMETHOD. The rule to YIELDUP(*Norm*) does not capture falling off the method body, but yields up the result of the normal execution of the invocation of a method with void return type in an expression statement.

EXECCSHARPSTM$_C$ ≡ **match** *context*(*pos*)
  `return` *exp*;   → *pos* := *exp*
  `return` ▶*val*; → YIELDUP(*Return*(*val*))
  `return`;   → YIELD(*Return*)

  Return   → **if** *pos* = *body*(*meth*) ∧ ¬*Empty*(*frames*) **then**
               EXITMETHOD(*Norm*)
  Return(*val*) → **if** *pos* = *body*(*meth*) ∧ ¬*Empty*(*frames*) **then**
               EXITMETHOD(*val*)

  ▶*Norm*; → YIELDUP(*Norm*)

The machine EXITMETHOD restores the frame of the invoker and passes the result value (if any). Upon normal return from a static constructor it also updates the *typeState* of the

---

[16] To reflect different parameter passing mechanisms as encountered in other programming languages, due to the modular character of our model essentially only the above submachine INITLOCALS would have to be appropriately modified.

relevant class as *Initialized*. We also add a rule FREELOCALS to free the memory used for local variables and value parameters, using an abstract notion FREEMEMORY of how addresses of local variables and value parameters are actually de-allocated. [17]

EXITMETHOD(*result*) ≡
    let (*oldMeth*, *oldPos*, *oldLocals*, *oldValues*) = *top*(*frames*) **in**
        meth   := *oldMeth*
        pos   := *oldPos*
        locals := *oldLocals*
        frames := *pop*(*frames*)
        if *StaticCtor*(*meth*) ∧ *result* = *Norm* **then**
           typeState(*type*(*meth*)) := *Initialized*
           values := *oldValues*
        else
           values := *oldValues* ⊕ {*oldPos* ↦ *result*}
        FREELOCALS

FREELOCALS ≡
    forall *x* ∈ *LocalVars*(*meth*) ∪ *ValueParams*(*meth*) **do**
        FREEMEMORY(*locals*(*x*), *type*(*x*))

Following [27, Sections 17.11,17.4.5.1,10.11,10.4.5.1] a type *c* is considered as initialized if its static constructor has been invoked (see the update of *typeState*(*c*) to *InProgress* in INITIALIZE below) or has terminated normally (see the update of *typeState*(*c*) to *Initialized* in EXITMETHOD above). We therefore define:

$$Initialized(c) \iff typeState(c) = Initialized \lor typeState(c) = InProgress$$

To initialize a class its static constructor is invoked (.cctor = class constructor). This macro will be further refined in $C\sharp_{\mathcal{E}}$ to account for exceptions during an initialization.

INITIALIZE(*c*) ≡
    if *typeState*(*c*) = *Linked* **then**
        typeState(*c*) := *InProgress*
        forall *f* ∈ *staticFields*(*c*) **do**
           let *t* = *type*(*c*::*f*) **in** WRITEMEM(*globals*(*c*::*f*), *t*, *defaultValue*(*t*))
        INVOKEMETHOD(*c*::.cctor, [])

Note that in C♯ the initialization of a class does not trigger the initialization of its direct base class, differing on this point from Java where the rule for calling static constructors (see [36, Fig. 4.5]) triggers the initialization of the superclass in case the superclass is not yet initialized.

    With respect to the execution of initializers of static class fields the ECMA standard [27, Section 17.4.5.1] says that the static field initializers of a class correspond to a sequence of assignments that are executed in the textual order in which they appear in the

---

[17] Under the assumption of a potentially infinite supply of addresses, which is often made when describing the semantics of a programming language, one can dispense with FREELOCALS.

class declaration. If a static constructor exists in the class, execution of the static field initializers occurs immediately prior to executing that static constructor. Otherwise, the static field initializers are executed at an *implementation-dependent* time prior to the first use of a static field of that class. Our model expresses the decision taken by Microsoft's current C♯ compiler, which in the second case creates a static constructor. If one wants to reflect also the non-determinism suggested by the ECMA formulation, one can formalize the implementation-dependent external control by a monitored function *typeToBeInitialized* (which by the way can also be used for the classes and structs classified by an implementation flag as *beforefieldinit* type). The C♯ interpreter then takes the following form: [18]

if *typeToBeInitialized* $\neq$ *Undef* **then**
    INITIALIZE(*typeToBeInitialized*)
else EXECCSHARP

4. Refinement C♯$_{\mathcal{O}}$ of C♯$_{\mathcal{C}}$ by object related features

In this section we refine the static class features of C♯$_{\mathcal{C}}$ by adding objects (for class instances, comprising arrays and structs) together with *instance* fields, methods and constructors [19] as well as inheritance (including overriding and overloading of methods). Accordingly we extend the ASM universes and functions of C♯$_{\mathcal{C}}$ to reflect the new expressions and statements together with the appropriate constraints and new rules, using appropriate refinements of some of the macros to define the semantics of the new instructions of C♯$_{\mathcal{O}}$. For the detailed definition of the syntax of (members of) classes, interfaces, structs, etc., and of the constraints for the new modifiers (`abstract`, `sealed`, `readonly`, `volatile`, `virtual`, `override`) together with illustrating examples, we refer the reader to [12].

4.1. Static semantics of C♯$_{\mathcal{O}}$

The first extension concerns the sets *Exp*, *Vexp*, *Sexp* where the new reference and array types appear. *Rank* serves to denote the dimension of array types; *NonArrayType* stands for value types, classes and interfaces and will be extended in C♯$_{\mathcal{D}}$ to comprise also delegates. Value types represent a feature that distinguishes C♯ from Java. In C♯$_{\mathcal{I}}$ we have cast expressions $(t)\,exp$ where the type t and the type of *exp* are numeric types. Here, we extend the grammar to $(t)\,exp$ where t and the type of *exp* can be any type. A *RefExp* is an expression

---

[18] This is discussed in detail in [20]. The reader finds there also a detection of further class initialization features that are missing in the ECMA specification, related to the definition of when a static class constructor has to be executed and to the initialization of structs.

[19] Destructors or finalizers which relate to garbage collection are not modeled here.
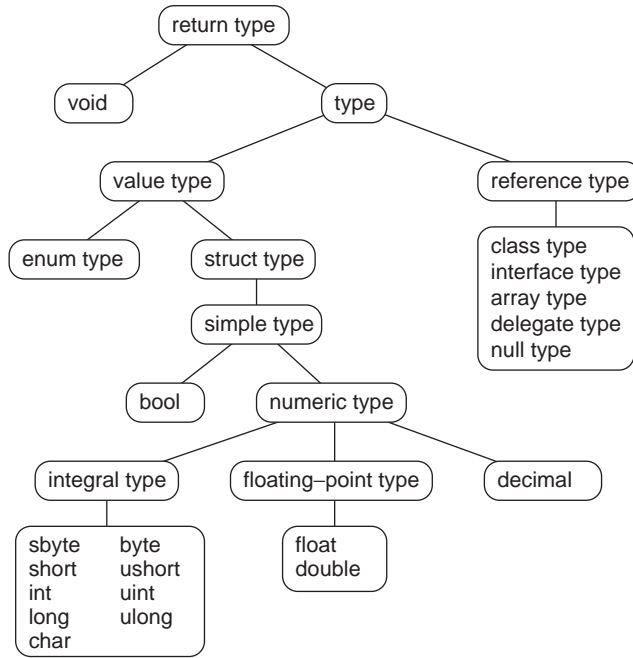
Fig. 4. The classification of types of C♯.

of a reference type and an *ArrayExp* is an expression of an array type.

$$
\begin{aligned}
Exp \ &::= \ldots \mid \text{`null'} \mid \text{`this'} \mid \text{`typeof'} \text{ `('} \textit{RetType} \text{ `)'} \mid Exp \text{ `is'} \textit{Type} \\
&\mid \ Exp \text{ `as'} \textit{RefType} \mid \text{`('} \textit{Type} \text{ `)'} Exp \\
&\mid \ \text{`new'} \textit{NonArrayType} \text{ `['} \textit{Exps} \text{ `]'} \{\textit{Rank}\} [\textit{ArrayInitializer}] \\[4pt]
Vexp \ &::= \ldots \mid Vexp \text{ `.'} \textit{Field} \mid \textit{RefExp} \text{ `.'} \textit{Field} \mid \text{`base'} \text{ `.'} \textit{Field} \\
&\mid \ \textit{ArrayExp} \text{ `['} \textit{Exps} \text{ `]'} \\[4pt]
Sexp \ &::= \ldots \mid \text{`new'} \textit{Type} ( [\textit{Args}] ) \mid Exp \text{ `.'} \textit{Meth} ([\textit{Args}]) \\
&\mid \ \text{`base'} \text{ `.'} \textit{Meth} ([\textit{Args}]) \\[4pt]
Exps \ &::= \ Exp \{\text{`,'} Exp\} \\[4pt]
Rank \ &::= \ \text{`['} \{ \text{`,'} \} \text{`]'}
\end{aligned}
$$

A `this` in an instance constructor or instance method of a struct is considered to be a *Vexp*. When a `this` occurs in a class it is not a *Vexp*.

The extended type classification where simple types become aliases for struct types is re-assumed by Fig. 4. We refer the reader to [12] for the detailed list of new type constraints. Also the constraints for overriding and overloading of methods and the resolution of overloaded methods at compile-time are spelt out there.

The subtype relation (i.e. the standard implicit conversion) is based on the inheritance relation—defined as a finite tree with root `object`—together with the "implements" relation between classes and interfaces. It is defined as follows (and should not be confused

with the classification of types in Fig. 4):

- T any type $\Longrightarrow T \preccurlyeq$ object and $T \preccurlyeq T$,
- class S derived from $T \Longrightarrow S \preccurlyeq T$,
- class, interface or struct S implements interface $T \Longrightarrow S \preccurlyeq T$,
- T array type $\Longrightarrow T \preccurlyeq$ System.Array,
- T delegate type $\Longrightarrow T \preccurlyeq$ System.Delegate,
- T value type $\Longrightarrow T \preccurlyeq$ System.ValueType,
- T array or delegate type $\Longrightarrow T \preccurlyeq$ System.ICloneable,
- T reference type $\Longrightarrow \Lambda \preccurlyeq T$, [$\Lambda$ is the null type],
- S and T reference types, $S \preccurlyeq T \Longrightarrow S[R_1] \cdots [R_k] \preccurlyeq T[R_1] \cdots [R_k]$.

Note that types of one category in Fig. 4 can be subtypes of another (disjoint) category. For example, if a struct type S implements the interface I, then (the value type) S is considered to be a subtype of (the reference type) I.

We list here the additional definite assignment rules for local variables of struct type:

- If p is a local variable of a struct type S, then $p.f$ is considered as a local variable for each instance field f of S.
- A local variable p of struct type S is definitely assigned $\Longleftrightarrow$
  $p.f$ is definitely assigned for each instance field f of S.

We assume that as a result of field and method resolution the attributed syntax tree has exact information. Each field access has the form $T::f$ where f is a field declared in the type T. Each method call has the form $T::m(args)$ where m is the signature of a method declared in type T. Moreover, at compile-time certain expressions are reduced to basic expressions as follows.

For the base access of fields and methods we have:

- base.f in class C is replaced by this.$B::f$, where B is the first base class of C where the field f is declared.
- base.$m(args)$ in class C is replaced by this.$B::M(args)$, where $B::M$ is the method signature of the method selected by the compiler (the set of applicable methods is constructed starting in the *base class* of C). This selection algorithm is described in [12], formalizing the conditions stated in [27, Section 14.4.2/3].

For instance field access and class instance creation we have:

- If f is a field, then f is replaced by this.$T::f$, where f is declared in T.
- Let T be a class type. Then the instance creation new $T::M(args)$ is replaced by new $T . T::M(args)$.

Hence we split an instance creation expression into a creation part and an invocation of an instance constructor. To make the splitting correctly reflect the intended meaning of new $T::M(args)$, we assume in our model that class instance constructors return the value of this.

For instance constructors of structs one has to reflect that in addition they need an address for this. Also for constructors of structs we assume that they return the value of this. Let S be a struct type. Then the following transformations are applied:

- An assignment $vexp =$ new $S::M(args)$ is replaced by $vexp . S::M(args)$. This reflects that such a new triggers no object creation or memory allocation since structs get their memory allocated at declaration time.

- Other occurrences of new S::M(*args*) are replaced by x . S::M(*args*), where x is a new temporary local variable of type S.

For automatic boxing we have:

- An assignment *vexp* = *exp* is replaced by *vexp* = (T)*exp* if *type*(*exp*) is a value type, $T = type(vexp)$ and T is a reference type. In this case we must have $type(exp) \preccurlyeq T$.
- An argument *arg* is replaced by (T)*arg* if *type*(*arg*) is a value type, the selected method expects an argument of type T and T is a reference type. In this case we must have $type(arg) \preccurlyeq T$.

4.2. Dynamic semantics for C♯$_\mathcal{O}$

We are now ready to describe the extension of the dynamic state for the model of C♯$_\mathcal{O}$. The domain of values is extended to contain also references (assuming $Ref \cap Adr = \emptyset$ and $null \in Ref$) and struct values: $Value = SimpleValue \cup Adr \cup Ref \cup Struct$. The set $Struct$ of struct values can be defined as the set of mappings from $StructType$::$Field$ to $Value$. The value of an instance field of a value of struct type T can then be extracted by applying the map to the field name, i.e. $structField(val, T, f)$.

Two dynamic functions keep track of the $runTimeType$: $Ref \rightarrow Type$ of references and of the type object $typeObj$: $RetType \rightarrow Ref$ of a given type, where $RetType ::= Type \mid$ 'void'. The memory function is extended to store also references: mem: $Adr \rightarrow SimpleValue \cup Ref \cup \{Undef\}$. For boxing we need a dynamic function $valueAdr$: $Ref \rightarrow Adr$ to record the address of a value in a box. If $runTimeType(ref)$ is a *value type t*, then $valueAdr(ref)$ is the address of the struct value of type t stored in the box. The static function $instanceFields$: $Type \rightarrow Powerset(Type::Field)$ yields the set of instance fields of any given type t; if t is a class type, it includes the fields declared in base classes of t. We use the common programming notation $Type$::$Field$ instead of the set-theoretic product set notation. We abstract from the implementation-dependent layout of structs and objects and use a function $fieldAdr$: $(Adr \cup Ref) \times Type::Field \rightarrow Adr$ to record addresses of fields. This function satisfies the following properties:

- If t is a *struct type*, then $fieldAdr(adr, t::f)$ is the address of field f of a value of type t stored in *mem* at address *adr*.
- A value of struct type t at address *adr* occupies the following addresses in *mem*:

$$\{fieldAdr(adr, f) \mid f \in instanceFields(t)\}$$

- If $runTimeType(ref)$ is a *class type*, then $fieldAdr(ref, t::f)$ is the address of field $t::f$ of the object referenced by *ref*.
- An object of class c is represented by a reference *ref* with the property $runTimeType(ref) = c$ and occupies the following addresses in *mem*:

$$\{fieldAdr(ref, f) \mid f \in instanceFields(c)\}$$

A function $elemAdr$: $Ref \times \mathbb{N}^* \rightarrow Adr$ records addresses of array elements. The this reference is treated as first parameter and is passed by value. Therefore $paramIndex(c::m,$ this$) = 0$ and this is an element of $ValueParams(c::m)$.

For the refinement of the EXECCSHARP$_C$ transition rules it suffices to add the new rule EXECCSHARPEXP$_O$ for evaluating the new expressions, since C♯$_O$ introduces no new statements.

EXECCSHARP$_O$ ≡
  EXECCSHARP$_C$
  EXECCSHARPEXP$_O$

For better readability we organize the numerous EXECCSHARPEXP$_O$ rules for each of the new expressions into parallel submachines each of which collects the rules for expressions which belong to the same category (for type testing and casting, for fields, for arrays). As analyzed in [20], the invocation of an instance constructor of a class may trigger the class initialization.

EXECCSHARPEXP$_O$ ≡ **match** *context*(*pos*)
  null → YIELD(*null*)
  this → YIELDINDIRECT(*locals*(this))
  TESTCASTEXP$_O$
  FIELDEXP$_O$
  new *c* → **let** *ref* = *new*(*Ref*, *c*) **in**
           runTimeType(*ref*) := *c*
           forall *f* ∈ *instanceFields*(*c*) **do**
             let *adr* = *fieldAdr*(*ref*, *f*) **and** *t* = *type*(*f*) **in**
               WRITEMEM(*adr*, *t*, *defaultValue*(*t*))
           YIELD(*ref*)
  exp . *T* ::*M* (*args*) → *pos* := *exp*
  ▶*val* . *T* ::*M* (*args*) → **if** *StructValueInvocation*(*up*(*pos*)) **then**
                 // create home for struct value
                 let *adr* = *new*(*Adr*, *type*(*pos*)) **in**
                   WRITEMEM(*adr*, *type*(*pos*), *val*)
                   values(*pos*) := *adr*
               pos := (*args*)
  val . *T* ::*M* ▶ (*vals*) → **if** *InstanceCtor*(*M*) ∧ *TriggerInit*(*T*) **then**
              INITIALIZE(*T*)
           elseif *val* ≠ *null* **then**
             INVOKEINSTANCE(*T* ::*M*, *val*, *vals*)
  ARRAYEXP$_O$

A struct value invocation is a method invocation on a struct value which is not stored in a variable. For such struct values the above rule creates a temporary storage area (called "home") to be passed in the invocation as value of this.

StructValueInvocation(*exp* . *T* ::*M* (*args*)) ⟺
  type(*exp*) ∈ *StructType* ∧ *exp* ∉ *Vexp*

The rules for casting in TESTCASTEXP$_O$ use the new macro YIELDUPBOX defined below. Note that in expressions '*exp* is *t*' and (*t*)*exp* the type *t* can be any type, whereas in '*exp* as *t*' the type *t* must be a reference type. The type of '*exp* is *t*' is bool, while the

type of $(t)\,exp$ and 'exp as t' is t.

$\text{TESTCASTEXP}_O \equiv$
  $\texttt{typeof}(t) \rightarrow \text{YIELD}(typeObj(t))$

  exp is $t \quad \rightarrow pos := exp$
  ▶val is $t \rightarrow$ **if** $type(pos) \in ValueType$ **then**
             $\text{YIELDUP}(type(pos) \preccurlyeq t)$ // compile-time property
          else
             $\text{YIELDUP}(val \neq null \wedge runTimeType(val) \preccurlyeq t)$

  exp as $t \quad \rightarrow pos := exp$
  ▶val as $t \rightarrow$ **if** $type(pos) \in ValueType$ **then**
             $\text{YIELDUPBOX}(type(pos), val)$ // box a copy of the value
          elseif $(val \neq null \wedge runTimeType(val) \preccurlyeq t)$ **then**
             $\text{YIELDUP}(val)$ // pass reference through
          else $\text{YIELDUP}(null)$ // convert to null reference

  $(t)\,exp \quad \rightarrow pos := exp$
  (t)▶$val \rightarrow$ **if** $type(pos) \in ValueType$ **then**
          // compile-time identity
          if $t = type(pos)$ **then** $\text{YIELDUP}(val)$
          // box value
          if $t \in RefType$ **then** $\text{YIELDUPBOX}(type(pos), val)$
       if $type(pos) \in RefType$ **then**
          if $t \in RefType \wedge (val = null \vee runTimeType(val) \preccurlyeq t)$ **then**
            $\text{YIELDUP}(val)$ // pass reference through
          if $t \in ValueType \wedge val \neq null \wedge t = runTimeType(val)$ **then**
            // un-box a copy of the value
            $\text{YIELDUP}(memValue(valueAdr(val), t))$

The rules for instance field access and assignment in FIELDEXP_O are analogous to those for class variables, adding the evaluation of the instance, using *fieldAdr* instead of *globals*, and instead of WRITEMEM the macro SETFIELD defined below. The second rule for instance field access has to distinguish two cases, depending upon the statically known instance type. Since this type information is already known at the time of static analysis, it could be resolved by introducing two separate constructs for field access, as one of our referees observed pointing also out that in fact the CLI has a single, overloaded instruction for field access with overloading to be resolved by the JIT. However from the modeling point of view, having two separate constructs for field access would lead to essentially the same two rules we have formulated here, except for having as rule guard a matching condition for the two constructs instead of the type test. We use $type(exp\,.\,t::f) = type(t::f)$.

$\text{FIELDEXP}_O \equiv$
  $exp\,.\,t::f \quad \rightarrow pos := exp$
  ▶$val\,.\,t::f \rightarrow$ **if** $type(pos) \in ValueType \wedge val \notin Adr$ **then**
            $\text{YIELDUP}(structField(val, type(pos), t::f))$
         elseif $val \neq null$ **then**
            $\text{YIELDUPINDIRECT}(fieldAdr(val, t::f))$

$$exp_1 . t :: f = exp_2 \quad \rightarrow pos := exp_1$$
$$\blacktriangleright val . t :: f = exp \quad \rightarrow pos := exp$$
$$val_1 . t :: f = \blacktriangleright val_2 \rightarrow \textbf{if } val_1 \neq null \textbf{ then}$$
$$\text{SETFIELD}(val_1, t :: f, val_2)$$
$$\text{YIELDUP}(val_2)$$

$C\sharp_O$ supports single dimensional as well as multi-dimensional arrays. Array types are read from right to left. For example, $\texttt{int[][,]}$ is the type of single-dimensional arrays of two-dimensional arrays with elements of type $\texttt{int}$. By $dim(n)$ we denote a sequence of $n - 1$ commas, hence $T[\,dim(n)\,]$ is the type of n-dimensional arrays with elements of type T. The length of the ith dimension of an n-dimensional array represented by a reference ref is stored as the value of $dimLength(ref, i)$. Note that the rules for using array indexing expressions as rvalue respectively as lvalue appear together as subgroups of ARRAYEXP_O, separated by pattern matching.

$\text{ARRAYEXP}_O \equiv$

$\quad$ new $T[\,exp_1, \ldots, exp_n\,][\,R_1\,] \cdots [\,R_k\,] \rightarrow pos := exp_1$

$\quad$ new $T[\,l_1, \ldots, \blacktriangleright l_i, exp_{i+1}, \ldots, exp_n\,][\,R_1\,] \cdots [\,R_k\,] \rightarrow pos := exp_{i+1}$

$\quad$ new $T[\,l_1, \ldots, \blacktriangleright l_n\,][\,R_1\,] \cdots [\,R_k\,] \rightarrow$

$\qquad$ **if** $\forall i \in [1 .. n]\,(0 \leqslant l_i)$ **then**

$\qquad\quad$ **let** $S = T[\,R_1\,] \cdots [\,R_k\,]$ **in**

$\qquad\quad$ **let** $ref = new(Ref, [l_1, \ldots, l_n], S)$ **in**

$\qquad\qquad$ $runTimeType(ref) := T[\,dim(n)\,][\,R_1\,] \cdots [\,R_k\,]$

$\qquad\qquad$ **forall** $i \in [1 .. n]$ **do** $dimLength(ref, i - 1) := l_i$

$\qquad\qquad$ **forall** $\alpha \in [0 .. l_1 - 1] \times \cdots \times [0 .. l_n - 1]$ **do**

$\qquad\qquad\quad$ $\text{WRITEMEM}(elemAdr(ref, \alpha), S, defaultValue(S))$

$\qquad\qquad$ $\text{YIELDUP}(ref)$

$\quad exp_0[\,exp_1, \ldots, exp_n\,] \rightarrow pos := exp_0$

$\quad \blacktriangleright ref[\,exp_1, \ldots, exp_n\,] \rightarrow pos := exp_1$

$\quad ref[\,i_1, \ldots, \blacktriangleright i_k, exp_{k+1}, \ldots, exp_n\,] \rightarrow pos := exp_{k+1}$

$\quad ref[\,i_1, \ldots, \blacktriangleright i_n\,] \rightarrow$

$\qquad$ **if** $ref \neq null \wedge \forall k \in [1 .. n]\,(0 \leqslant i_k < dimLength(ref, k - 1)) \wedge$

$\qquad\quad (RefOrOutArg(up(pos)) \wedge type(up(pos)) \in RefType \rightarrow$

$\qquad\qquad elementType(runTimeType(ref)) = type(up(pos)))$

$\qquad$ **then**

$\qquad\quad$ $\text{YIELDUPINDIRECT}(elemAdr(ref, (i_1, \ldots, i_n)))$

$\quad exp_0[\,exp_1, \ldots, exp_n\,] = exp \rightarrow pos := exp_0$

$\quad \blacktriangleright ref[\,exp_1, \ldots, exp_n\,] = exp \rightarrow pos := exp_1$

$\quad ref[\,i_1, \ldots, \blacktriangleright i_k, exp_{k+1}, \ldots, exp_n\,] = exp \rightarrow pos := exp_{k+1}$

$\quad ref[\,i_1, \ldots, \blacktriangleright i_n\,] = exp \rightarrow pos := exp$

$\quad ref[\,i_1, \ldots, i_n\,] = \blacktriangleright val \rightarrow$

$\qquad$ **let** $T = elementType(runTimeType(ref))$ **in**

$\qquad\quad$ **if** $ref \neq null \wedge \forall k \in [1 .. n]\,(0 \leqslant i_k < dimLength(ref, k - 1)) \wedge$

$\qquad\qquad (type(pos) \in RefType \rightarrow runTimeType(val) \leqslant T)$

then
    WRITEMEM(*elemAdr*(*ref*, ($i_1, \ldots, i_n$)), *T*, *val*)
    YIELDUP(*val*)

4.2.1. Refinement of macros

Invocation of instance methods splits into virtual and non-virtual calls. The function *lookup* yields the class where the given method specification is defined in the class hierarchy, depending on the run-time type of the given reference.

INVOKEINSTANCE(*T*::*M*, *val*, *vals*) ≡
    if *callKind*(*T*::*M*) = *Virtual* **then**   // indirect call, *val* ∈ *Ref*
        let *S* = *lookup*(*runTimeType*(*val*), *T*::*M*) **in**
        let *this* = **if** *S* ∈ *StructType* **then** *valueAdr*(*val*) **else** *val* **in**
            INVOKEMETHOD(*S*::*M*, [*this*] · *vals*)
    if *callKind*(*T*::*M*) = *NonVirtual* **then**   // direct call, *val* ∈ *Adr* ∪ *Ref*
        INVOKEMETHOD(*T*::*M*, [*val*] · *vals*)

In $C\sharp_{\mathcal{O}}$ the notion of reading from the memory is refined by extending the simple equation *memValue*(*adr*, *t*) = *mem*(*adr*) of $C\sharp_{\mathcal{I}}$ to fit also reference and struct types. This is done by the following simultaneous recursive definition of *memValue* and *getField* along the given struct type.

memValue(*adr*, *t*) =
    if *t* ∈ *SimpleType* ∪ *RefType* **then** *mem*(*adr*)
    elseif *t* ∈ *StructType* **then** {*f* ↦ *getField*(*adr*, *f*) | *f* ∈ *instanceFields*(*t*)}

getField(*adr*, *t*::*f*) = *memValue*(*fieldAdr*(*adr*, *t*::*f*), *type*(*t*::*f*))

Writing to memory is refined recursively together with SETFIELD along the given struct type:

WRITEMEM(*adr*, *t*, *val*) ≡
    if *t* ∈ *SimpleType* ∪ *RefType* **then** *mem*(*adr*) := *val*
    elseif *t* ∈ *StructType* **then**
        forall *f* ∈ *instanceFields*(*t*) **do** SETFIELD(*adr*, *f*, *val*(*f*))

SETFIELD(*adr*, *t*::*f*, *val*) ≡ WRITEMEM(*fieldAdr*(*adr*, *t*::*f*), *type*(*t*::*f*), *val*)

The notion of *AddressPos* from $C\sharp_{\mathcal{C}}$ is refined to include also lvalue nodes of *StructType*, so that address positions are of the following form: `ref` □, `out` □, □`++`, □`--`, □ *op*= *exp*, □`.`*f*, □`.`*m*`(`*args*`)`.

AddressPos(α) ⟺ *FirstChild*(α) ∧
    label(*up*(α)) ∈ {`ref`, `out`, `++`, `--`} ∨ *label*(*up*(α)) ∈ *Aop* ∨
    (*label*(*up*(α)) = '.' ∧ α ∈ *Vexp* ∧ *type*(α) ∈ *StructType*)

YIELDUPBOX creates a box for a given value of a given type and returns its reference. The run-time type of a reference to a boxed value of struct type *t* is defined to be (the value type) *t* of the value. There is no need to introduce special reference types for boxed values. If *type*(*exp*) is a value type that implements the interface *I*, then *type*(*exp*)⩽*I* and the value

can be boxed using '$(I)\,exp$' or 'exp as i'.

YIELDUPBOX$(t, val) \equiv$ **let** $ref = new(Ref)$ **and** $adr = new(Adr, t)$ **in**
  $runTimeType(ref) := t$
  $valueAdr(ref) := adr$
  WRITEMEM(adr, t, val)
  YIELDUP(ref)

The struct value is copied in both cases, when it is boxed and when it is un-boxed.

4.2.2. ASM function new

We now justify in the context of the basic parallel execution mechanism of ASM rules the sequentiality which is used in the following macros:

let $adr = new(Adr, T)$ **in** P
let $ref = new(Ref, T)$ **in** P
let $ref = new(Ref, [l_1, \ldots, l_n], T)$ **in** P

In the context of the machine EXECCSHARP this comes up to specify an abstract memory management. In fact **let** $adr = new(Adr, T)$ **in** P stands for the sequential execution of a new address allocation (which uses the ASM construct **import** to provide a completely fresh element) followed by P:

let $adr = new(Adr, T)$ **in** $P \equiv$ (**import** adr **do** ALLOCADR(adr, T)) **seq** P

where the operator **seq** for sequential execution of two ASMs M, N is to be understood as defined for turbo ASMs in [14] (alternatively see [16, Chapter 4.1]), namely as binding into one overall ASM step the two steps of first executing M in the given state and then N in the resulting state. Similarly **let** $ref = new(Ref, T)$ **in** P stands for the sequential execution of address allocation for all instance fields of a given type followed by P:

let $ref = new(Ref, T)$ **in** $P \equiv$
  import ref **do**
    $Ref(ref) := True$
    ALLOCFIELDS$(ref, instanceFields(T))$
  seq P

Similarly, we define the address allocation for elements of an n-dimensional array:

let $ref = new(Ref, [l_1, \ldots, l_n], T)$ **in** $P \equiv$
  import ref **do**
    $Ref(ref) := True$
    forall $\alpha \in [0 \mathrel{..} l_1 - 1] \times \cdots \times [0 \mathrel{..} l_n - 1]$ **do**
      import adr **do**
        $elemAdr(ref, \alpha) := adr$
        ALLOCADR(adr, T)
  seq P

The two macros for allocation of addresses and fields can be recursively defined as follows, relying again upon the definition of recursive turbo ASMs in [10] (or see alternatively [16, Chapter 4.1.2]):

ALLOCADR(adr, T) ≡
   $Adr(adr) := True$
   if $T \in StructType$ **then** ALLOCFIELDS($adr, instanceFields(T)$)

ALLOCFIELDS(x, fs) ≡
   forall $f \in fs$ **import** adr **do**
     $fieldAdr(x, f) := adr$
     ALLOCADR($adr, type(f)$)

5. Refinement C♯$_{\mathcal{E}}$ of C♯$_{\mathcal{O}}$ by exception handling

In this section we extend C♯$_{\mathcal{O}}$ with the exception handling mechanism of C♯, which separates normal program code from exception handling code. To this purpose exceptions are represented as objects of predefined system exception classes or of user-defined application exception classes. Once created (thrown), these objects trigger an abruption of the normal program execution to catch the exception—in case it is compatible with one of the exception classes appearing in the program in an enclosing try–catch–finally statement. Optional finally statements are guaranteed to be executed independently of whether the try statement completes normally or is abrupted.

5.1. Static semantics of C♯$_{\mathcal{E}}$

For the refinement of EXECCSHARP$_O$ by exceptions, as in the previous section it suffices to add the new rules for exception handling and to extend the static semantics. The set of statements is extended by throw and try–catch–finally statements satisfying the following constraints:

       Stm ::= ... | 'throw' Exp ';' | 'throw' ';'
                | 'try' $Block$ {$Catch$} ['catch' $Block$] ['finally' $Block$]

       $Catch$ ::= 'catch' '(' $ClassType$ [Loc] ')' $Block$

- every try–catch–finally statement contains at least one *catch clause*, *general catch clause* (i.e. of form catch *block*), or *finally block*,
- no return statements are allowed in finally blocks,
- a break, continue, or goto statement is not allowed to jump out of a finally block,
- a throw statement without expression is only allowed in catch blocks,
- the exception classes in a *Catch* clause appear there in a non-decreasing type order, more precisely, for every try–catch statement of the form.

      try *block* catch (E_1 x_1) $block_1$ ... catch (E_n x_n) $block_n$

the following holds $i < j \implies E_j \not\preccurlyeq E_i$ (and $E_i \preccurlyeq$ System.Exception).

In our model the sets of abruptions and type states have to be extended by exceptions. Due to the presence of `throw` statements without expression, a stack of references is needed to record exceptions which are to be re-thrown.

$$Abr = \ldots \mid Exc(Ref), \; TypeState = \ldots \mid Exc(Ref), \; excStack{:}\,List(Ref)$$

To simplify the exposition we assume that general catch clauses '`catch` *block*' are replaced at compile-time by '`catch (Object` x `)` *block*' with a new variable x. We also reduce try–catch–finally statements to try–catch and try–finally statements as follows. Both reductions can easily be shown to correctly express the ECMA specification:

```
try TryBlock
catch ( E₁ x₁ ) CatchBlock₁
   ⋮
catch ( Eₙ xₙ ) CatchBlockₙ
finally FinallyBlock
```
$\implies$
```
try {
  try TryBlock
  catch ( E₁ x₁ ) CatchBlock₁
     ⋮
  catch ( Eₙ xₙ ) CatchBlockₙ
} finally FinallyBlock
```

If a static constructor throws an exception, and no catch clauses exists to catch it, then this exception is wrapped into a `TypeInitializationException` by translating `static` T`()` `{` *BlockStatements* `}` into

```
static T ( ) {
  try { BlockStatements }
  catch (Exception e) {
    throw new TypeInitializationException(T,e);
  }
}
```

The reachability rules and the definite assignment constraints for a try–catch $stm \equiv$ `try` *tryBlock* `catch (` $E_1\, x_1$ `)` *catchBlock*$_1$ `...catch (` $E_n\, x_n$ `)` *catchBlock*$_n$ are:

- If *reachable*(*stm*), then *reachable*(*tryBlock*) and *reachable*(*catchBlock*$_i$) for every $i \in [1..n]$.
- If *normal*(*tryBlock*) or *normal*(*catchBlock*) for at least one $i \in [1..n]$, then *normal*(*stm*).
- *before*(*tryBlock*) = *before*(*stm*).
- *before*(*catchBlock*$_i$) = *before*(*stm*) $\cup \{x_i\}$ for every $i \in [1..n]$.
- *after*(*stm*) = *after*(*tryBlock*) $\cap \bigcap_{i=1}^{n}$ *after*(*catchBlock*$_i$).

For a statement *stm* of the form `try` *tryBlock* `finally` *finallyBlock* the rules and constraints are:

- If *reachable*(*stm*), then *reachable*(*tryBlock*) and *reachable*(*finallyBlock*).
- If *normal*(*tryBlock*) and *normal*(*finallyBlock*), then *normal*(*stm*).
- *before*(*tryBlock*) = *before*(*stm*).
- *before*(*finallyBlock*) = *before*(*stm*).
- *after*(*stm*) = *after*(*tryBlock*) $\cup$ *after*(*finallyBlock*).

5.2. Dynamic semantics for C$\sharp_{\mathcal{E}}$

The transition rules for EXECCSHARP$_E$ are defined by adding two submachines to EXECCSHARP$_O$. The first one provides the rules for handling the exceptions which may occur during the evaluation of expressions, the second one describes the meaning of the new throw and try–catch–finally statements.

EXECCSHARP$_E$ ≡
  EXECCSHARP$_O$
  EXECCSHARPEXP$_E$
  EXECCSHARPSTM$_E$

5.2.1. Expression evaluation rules

EXECCSHARPEXP$_E$ contains rules for each of the numerous forms of run-time exceptions defined in the subclasses of `System.Exception`. We give here seven characteristic examples and group them for the ease of presentation into parallel submachines by the form of expression they are related to, namely for arithmetical exceptions and for those related to cast expressions, reference expressions or array expressions. The notion of FAILUP we use is supposed to execute the code `throw new E()` at the parent position, so that we define the macro by invoking an internal method `Throw`E with that effect for each of the exception classes E used as parameter of FAILUP.

EXECCSHARPEXP$_E$ ≡ **match** *context*(*pos*)
  uop ▶*val* → **if** *Checked*(*pos*) ∧ *Overflow*(*uop*, *val*) **then**
           FAILUP(`OverflowException`)
  val$_1$ *bop* ▶*val*$_2$ → **if** *DivisionByZero*(*bop*, *val*$_2$) **then**
            FAILUP(`DivideByZeroException`)
          elseif *DecimalOverflow*(*bop*, *val*$_1$, *val*$_2$)∨
           (*Checked*(*pos*) ∧ *Overflow*(*bop*, *val*$_1$, *val*$_2$))
          then FAILUP(`OverflowException`)
  CASTEXCEPTIONS
  NULLREFEXCEPTIONS
  ARRAYEXCEPTIONS

FAILUP(E) ≡ INVOKEMETHOD(*ExcSupport*::Throw$E$, [ ])

CASTEXCEPTIONS ≡ **match** *context*(*pos*)
  (*t*)▶*val* →
    if *type*(*pos*) ∈ *RefType* **then**
      if *t* ∈ *RefType* ∧ *val* ≠ *Null* ∧ *runTimeType*(*val*) ≰*t* **then**
        FAILUP(`InvalidCastException`)
      if *t* ∈ *ValueType* **then**   // attempt to unbox
        if *val* = *Null* **then** FAILUP(`NullReferenceException`)
        elseif *t* ≠ *runTimeType*(*val*) **then**
          FAILUP(`InvalidCastException`)
      if *type*(*pos*) ∈ *NumericType* ∧ *t* ∈ *NumericType* ∧ *Checked*(*pos*) ∧

$$Overflow(t, val)$$
then FAILUP(OverflowException)

The semantics of assignments as defined by the ECMA standard and formalized by the rule
NULLREFEXCEPTIONS is violated by a compiler optimization in [39] related to the timing
of the *Null* check for certain expressions, see the analysis in [20].

NULLREFEXCEPTIONS ≡ **match** *context(pos)*
$\quad$ ▶*ref* . *t* :: *f* $\qquad$ → **if** *ref* = *Null* **then**
$\qquad\qquad\qquad\qquad\qquad$ FAILUP(NullReferenceException)
$\quad$ *ref* . *t* :: *f* = ▶*val* $\quad$ → **if** *ref* = *Null* **then**
$\qquad\qquad\qquad\qquad\qquad$ FAILUP(NullReferenceException)
$\quad$ *ref* . *T* :: *M* (▶*vals*) $\;$ → **if** *ref* = *Null* **then**
$\qquad\qquad\qquad\qquad\qquad$ FAILUP(NullReferenceException)

If the address of an array element is passed as a `ref` or `out` argument to a method,
then the run-time element type of the array must be *equal* to the parameter type that the
method expects. If an object is assigned to an array element, then the type of the object
must be a *subtype* of run-time element type of the array (array covariance problem). In
both cases, if the condition is not satisfied, an `ArrayTypeMismatchException` is
thrown.

ARRAYEXCEPTIONS ≡ **match** *context(pos)*
$\quad$ new $T [l_1 , \ldots , $▶$l_n] [R_1] \cdots [R_k] \to$
$\qquad$ **if** $\exists i \in [1 \ldots n] (l_i < 0)$ **then** FAILUP(OverflowException)
$\quad$ *ref* $[i_1 , \ldots , $▶$i_n] \to$
$\qquad$ **if** *ref* = *Null* **then** FAILUP(NullReferenceException)
$\qquad$ **elseif** $\exists k \in [1 \ldots n] (i_k < 0 \vee dimLength(ref , k - 1) \leqslant i_k)$ **then**
$\qquad\quad$ FAILUP(IndexOutOfRangeException)
$\qquad$ **elseif** *RefOrOutArg(up(pos))* ∧ *type(up(pos))* ∈ *RefType* ∧
$\qquad\quad$ *elementType(runTimeType(ref))* ≠ *type(up(pos))*
$\qquad$ **then** FAILUP(ArrayTypeMismatchException)
$\quad$ *ref* $[i_1 , \ldots , i_n] = $▶*val* →
$\qquad$ **if** *ref* = *Null* **then** FAILUP(NullReferenceException)
$\qquad$ **elseif** $\exists k \in [1 \ldots n] (i_k < 0 \vee dimLength(ref , k - 1) \leqslant i_k)$ **then**
$\qquad\quad$ FAILUP(IndexOutOfRangeException)
$\qquad$ **elseif** *type(pos)* ∈ *RefType* ∧
$\qquad$ *runTimeType(val)* ⪯̸ *elementType(runTimeType(ref))* **then**
$\qquad\quad$ FAILUP(ArrayTypeMismatchException)

5.2.2. Statement execution rules

The statement execution submachine splits naturally into submachines for throw, try–
catch, try–finally statements and a rule for the propagation of an exception (from the
root position of a method body) to the method caller. To support a correct understand-
ing of the exception messages that are printed to the console we include into the rule for
throw statements the initialization of stack traces. The initialization of stack traces in Java

and C♯ is different. In Java, the stack trace is initialized with the complete trace up to the main function once and for all when the exception object is created. In C♯ the stack trace is initialized with the empty trace each time when the exception object is thrown with `throw` and then augmented whenever the exception is propagated to a parent frame. The semantics of the parameterless `throw;` instruction is explained in terms of the exception Stack *excStack*. When an exception is caught, it is pushed on top of the exception stack (which as explained above is needed to record exceptions which are to be rethrown). Whenever a catch block terminates (normally or abruptly) the topmost element of the exception stack is deleted. No special rules are needed for general catch clauses '`catch` *block*' in try–catch statements, due to their compile-time transformation mentioned above. The completeness of the try–finally rules is due to the constraints listed above, which restrict the possibilities for exiting a finally block to normal completion or triggering an exception.

$\text{ExecCsharpStm}_E \equiv \textbf{match } \textit{context}(pos)$

   `throw` *exp;* → *pos := exp*

   `throw` ▶*ref;* → **if** *ref = Null* **then** $\text{FailUp}(\texttt{NullReferenceException})$
                 else $\text{InitStackTrace}(ref, meth)$
                     $\text{YieldUp}(Exc(ref))$

   `throw;` → $\text{Yield}(Exc(top(excStack)))$

   `try` *block* `catch` `(` *E x* `)` *stm* ... → *pos := block*
   `try` ▶*Norm* `catch` `(` *E x* `)` *stm* ... → $\text{YieldUp}(Norm)$
   `try` ▶*Exc(ref)* `catch(`*E₁ x₁*`)` *stm₁* ... `catch(`*Eₙ xₙ*`)` *stmₙ* →
     if $\exists i \in [1..n]\ runTimeType(ref) \preccurlyeq E_i$ **then**
       let $j = \min\{i \in [1..n] \mid runTimeType(ref) \preccurlyeq E_i\}$ **in**
         $pos := stm_j$
         $excStack := push(ref, excStack)$
         $\text{WriteMem}(locals(x_j), \texttt{object}, ref)$
     else $\text{YieldUp}(Exc(ref))$
   `try` ▶*abr* `catch(`*E₁ x₁*`)` *stm₁* ... `catch(`*Eₙ xₙ*`)` *stmₙ* → $\text{YieldUp}(abr)$
   `try` *Exc(ref)* ... `catch(`...`)` ▶*res* ... →
     $\{excStack := pop(excStack),\ \text{YieldUp}(res)\}$

   `try` *tryBlock* `finally` *finallyBlock* → *pos := tryBlock*
   `try` ▶*res* `finally` *finallyBlock* → *pos := finallyBlock*
   `try` *res* `finally` ▶*Norm* → $\text{YieldUp}(res)$
   `try` *res* `finally` ▶*Exc(ref)* → $\text{YieldUp}(Exc(ref))$

   Exc(ref) → **if** *pos = body(meth)* $\wedge \neg Empty(frames)$ **then**
     if *StaticCtor(meth)* **then** $typeState(type(meth)) := Exc(ref)$
     else $\text{AppendStackTrace}(ref, meth(top(frames)))$
     $\text{ExitMethod}(Exc(ref))$

In case an exception happened in the static constructor of a type, its type state is set to that exception to prevent its re-initialization and instead to re-throw the old exception object. The refinement of the macro Initialize defined in C♯$_C$ re-throws the exception

object of a type which had an exception in the static constructor, thus preventing its re-initialization. [20]

$\text{INITIALIZE}(c) \equiv$

  . . .
  if $typeState(c) = Exc(ref)$ **then** $\text{YIELD}(Exc(ref))$

6. Refinement C♯$_\mathcal{D}$ of C♯$_\mathcal{E}$ by delegates

In this section we extend C♯$_\mathcal{E}$ by features which distinguish C♯ from other languages, e.g. Java. We start with delegates and then add further constructs whose semantics can be defined mainly by reducing them via syntactical translations to the language model developed so far: properties, indexers, overloaded operators, enumerators with the `foreach` statement, the `using` statement, events and attributes. We use the model developed so far as ground model (in the sense of [8]) for C♯, thus providing a basis to justify the correctness (with respect to the ECMA standard [27]) of the "semantics of syntactic sugar" introduced in this section to define the semantics for delegates, properties, etc.

6.1. Delegates

Delegate types in C♯ are reference types that encapsulate a static or instance method with a specific signature, with the intention of having delegates playing the role of type-safe function pointers. A delegate type D is declared as follows:

```
delegate T D(S₁ x₁ , . . . , Sₙ xₙ);
```

It represents the type of methods that take n arguments of type $S_1, \ldots, S_n$ and have return type T. Delegate types appear as subtypes of `System.Delegate` and provide in particular the *callback* functionality and asynchronous event handling. More precisely, the characteristic ability of delegates is to call a list of multiple methods sequentially. This feature is realized by means of an *invocationList*: $Ref \rightarrow Delegate^* \cup \{Undef\}$ with which each delegate instance is equipped upon its creation. Each such list is a per instance immutable, non-empty, ordered list of static methods or pairs of target objects and instance methods. Upon invocation of a delegate instance with arguments *args*, the methods of its invocation list are called one after the other with these arguments *args*, returning to the caller of the delegate either the *return value* of the last list element or the first *exception* a list element has thrown during its execution, preventing the remaining list elements from being invoked.

Therefore we introduce a new universe $Delegate = Meth \cup (Ref \times Meth)$. To express the creation and use of new delegate expressions the sets *Exp*, *Sexp* are extended by additional

---

[20] For modeling the implementation flag *beforefieldinit* mentioned above this implies, as observed in [20], to refine also the predicate *TriggerInit*, used for invoking static or instance methods, namely to guarantee for a class in exception state its initialization even if the class is marked *beforefieldinit*: $TriggerInit(c) = (\neg Initialized(c) \land \neg beforefieldinit(c)) \lor typeState(c) = Exc(ref)$.

grammar rules as follows, using a new set *Dexp* of delegate expressions:

$$Sexp ::= \ldots | Exp \ (\ [Args] \ )$$
$$Exp ::= \ldots | \text{'new'} \ DelegateType \ \text{'('} \ Dexp \ \text{')'}$$
$$Dexp ::= Meth \ | \ Type \ \text{'.'} \ Meth \ | \ Exp \ \text{'.'} \ Meth \ | \ Exp$$

A method $T{::}M$ is called *compatible* with the delegate type D iff $T{::}M$ and D have the same return type and the same number of parameters with the same parameter types (including `ref`, `out`, `params` modifiers). The type constraints on the new expressions are spelt out in [12].

We use the model EXECCSHARPSTM$_I$, which includes a description of the `for` statement of C♯$_\mathcal{I}$, to express the sequentiality of the execution of delegate invocation list elements. In fact the above delegate declaration can be translated for $T \neq$ `void` in the following class:

```
sealed class D : System.Delegate {
  public T Invoke( S₁ x₁ , ... , Sₙ xₙ ) {
    T result;
    for (int i = 0; i < this._length() ; i++)
      result = this._invoke(i,x₁, ... ,xₙ);
    return result;
  }
  private extern int _length();
  private extern T _invoke(int i , S₁ x₁ , ... , Sₙ xₙ );
}
```

A delegate invocation expression $exp(args)$ can be syntactically translated into a normal method call exp . $D{::}\text{Invoke}(args)$ where D is the type of exp.[21] It then suffices to refine the ASM rule INVOKEEXTERN defined in the model EXECCSHARPEXP$_C$ to describe the meaning of the method $D{::}\_\text{invoke}$, which is to invoke the ith element of the invocation list on the given arguments, and analogously of `_length`.

INVOKEEXTERN($T{::}M$, *vals*) ≡
  if $T \in DelegateType$ **then**
    if $name(M) = \_\text{length}$ **then**
      DELEGATELENGTH(*vals*|(0))
    if $name(M) = \_\text{invoke}$ **then**
      INVOKEDELEGATE(*vals*(0), *vals*(1), *drop*(*vals*, 2))

DELEGATELENGTH(*ref*) ≡
  YIELDUP(*length*(*invocationList*(*ref*)))

INVOKEDELEGATE(*ref*, *i*, *vals*) ≡
  match *invocationList*(*ref*)(*i*)

---

[21] In [27, Section 10.4.7] the members of a delegate are defined to be the members inherited from the class `System.Delegate`. However neither .NET nor Rotor nor Mono do respect this stipulation since they add further methods to those inherited. One such example is the method `_invoke` we use here.

$$T::M \qquad\qquad \rightarrow \text{INVOKESTATIC}(T::M, vals)$$
$$(target, T::M) \rightarrow \text{INVOKEINSTANCE}(T::M, target, vals)$$

Since there are no new statements appearing in $C\sharp_{\mathcal{D}}$, the addition of the rule EXECCSHARP_D to EXECCSHARP_E consists in the following ASM subrule EXECCSHARPEXP_D, which defines the meaning of delegate instance creation. For a detailed analysis of the discrepancy we exhibit here between the ECMA standard and the .NET implementation see [20].

$\text{EXECCSHARPEXP}_D \equiv$ **match** *context*(*pos*)
  new D ($T::M$) $\rightarrow$
    let $d = new(Ref, D)$ **in**
      runTimeType(*d*) := *D*
      invocationList(*d*) := [*T*::*M*]
      $\text{YIELD}(d)$
  new D (*exp* . $T::M$) $\rightarrow$ *pos* := *exp*
  new D (►*ref* . $T::M$) $\rightarrow$
    if *ref* = *Null* **then** $\text{FAILUP}(\texttt{NullReferenceException})$
    else let $d = new(Ref, D)$ **in**
      runTimeType(*d*) := *D*
      invocationList(*d*) := [(*ref*, *T*::*M*)]
      $\text{YIELDUP}(d)$
  new D (*exp*) $\rightarrow$ *pos* := *exp*
  new D (►*ref*) $\rightarrow$
    if *ref* = *Null* **then** $\text{FAILUP}(\texttt{NullReferenceException})$
    else let $d = new(Ref, D)$ **in**
      runTimeType(*d*) := *D*
      invocationList(*d*) := *invocationList*(*ref*)   // ECMA §14.5.10.3
      // Microsoft .NET Framework:
      // *invocationList*(*d*) := [(*ref*, *D*::Invoke)]
      $\text{YIELDUP}(d)$

To be complete, one should add some rules which reflect the special character of delegate invocation lists. As usual for lists, two invocation lists are *equal* (==) iff they have the same length and the elements of the lists are pairwise equal; they can be *combined* (concatenated with '+') and sublists determined by a particular prefix and suffix condition can be *removed* from them (with '-'). To describe this specialization of list operations in our model it suffices to refine the macro INVOKEEXTERN by new rules for these operators `operator+`, `operator-`, `operator==`.

$\text{INVOKEEXTERN}(T::M, vals) \equiv$

  ...
  if $T \in DelegateType$ **then**
    if *name*(*M*) = `operator+` **then**
      $\text{DELEGATECOMBINE}(T, vals(0), vals(1))$
    if *name*(*M*) = `operator-` **then**
      $\text{DELEGATEREMOVE}(T, vals(0), vals(1))$

if $name(M) =$ operator== **then**
    DELEGATEEQUAL($vals(0)$, $vals(1)$)

Since invocation lists are considered to be immutable, combination and removal return *new* delegate instances unless one of the arguments is null. The null reference represents a delegate instance with an empty invocation list.

DELEGATECOMBINE(D, r_1, r_2) ≡
  if $r_1 = Null$ **then** YIELDUP(r_2)
  elseif $r_2 = Null$ **then** YIELDUP(r_1)
  else let $d = new(Ref, D)$ **in**
    $runTimeType(d) := D$
    $invocationList(d) := invocationList(r_1) \cdot invocationList(r_2)$
    YIELDUP(d)

DELEGATEREMOVE(D, r_1, r_2) ≡
  if $r_1 = Null$ **then** YIELDUP($Null$)
  elseif $r_2 = Null$ **then** YIELDUP(r_1)
  else let $l_1 = invocationList(r_1)$ **and** $l_2 = invocationList(r_2)$ **in**
    if $l_1 = l_2$ **then** YIELDUP($Null$)
    elseif $Subword(l_2, l_1)$ **then let** $d = new(Ref, D)$ **in**
      $runTimeType(d) := D$
      $invocationList(d) := prefix(l_2, l_1) \cdot suffix(l_2, l_1)$
      YIELDUP(d)
    else YIELDUP(r_1)

The notions of *prefix* and *suffix* are defined here in terms of the *last* occurrence of a subword: $prefix(u, v)$ is the part of v before the last occurrence of u in v and $suffix(u, v)$ the part of v after the last occurrence of u in v.

DELEGATEEQUAL(r_1, r_2) ≡
  if $r_1 = Null \lor r_2 = Null$ **then** YIELDUP($r_1 = r_2$)
  else let $l_1 = invocationList(r_1)$ **and** $l_2 = invocationList(r_2)$ **in**
    YIELDUP($length(l_1) = length(l_2) \land \forall i < length(l_1)\,(l_1(i) = l_2(i))$)

6.2. Properties, events and further features in C$\sharp_{\mathcal{D}}$

In this section we add further language features of C$\sharp$ whose semantics can be easily defined in terms of the model developed so far, essentially by simple syntactical reductions which one can easily justify to formalize correctly the explanations in [27].

6.2.1. Properties
Collections of a read and/or a write method for attributes of a class or struct are called *properties* in C$\sharp$ and declared in the following form (we skip the modifiers):

    Type Identifier '{' ['get' *Block*] ['set' *Block*] '}'

By definition a *read–write* property has a `get` and a `set` accessor, a *read-only* property has only a `get` accessor, a *write-only* property has only a `set` accessor. The identifier of a property P of type T can be used like a field identifier,[22] except that it cannot be passed as `ref` or `out` argument. Furthermore, it is required that the body of a `get` accessor is the body of a method with return type T, that a `set` accessor has a value parameter named `value` of type T and that its body is the body of a `void` method. Using the signatures T `get_P();` and `void set_P(`T `value);`, which are reserved for get and set accessors, the intended semantics of properties is reduced to the semantics of methods, using the following syntactical reductions:

$$
\begin{array}{l}
T\ P\ \{ \\
\quad \texttt{get}\ \{\ getAccessor\ \} \\
\quad \texttt{set}\ \{\ setAccessor\ \} \\
\}
\end{array}
\quad \Longrightarrow \quad
\begin{array}{l}
T\ \texttt{get\_}P()\ \{\ getAccessor\ \} \\
\texttt{void set\_}P(T\ \texttt{value})\ \{ \\
\quad setAccessor \\
\}
\end{array}
$$

$$
exp.P \Longrightarrow exp.\texttt{get\_}P() \qquad exp_1.P = exp_2; \Longrightarrow exp_1.\texttt{set\_}P(exp_2);
$$

The above translation comprises also expressions of the form $exp_1.P\ op= exp_2$, because we can assume that these compound assignments are compiled to $\langle x = exp_1,\ y = x.\texttt{get\_P()}\ op\ exp_2,\ x.\texttt{set\_P(}y\texttt{)},\ y\rangle$ with fresh local variables x, y, using as auxiliary operator the comma operator familiar from C/C++. This necessitates auxiliary rules for going through sequences of expressions of the following form:

$$
\begin{array}{ll}
\langle exp, \ldots \rangle & \to pos := exp \\
\langle val_1, \ldots, \blacktriangleright val_n \rangle & \to \textsc{YieldUp}(val_n) \\
\langle \ldots \blacktriangleright val, exp \ldots \rangle & \to pos := exp
\end{array}
$$

6.2.2. Indexers

Indexers can be used like array elements except that they cannot contain `ref` or `out` parameters and their elements cannot be passed as `ref` or `out` arguments. They are declared in a class or struct type as follows (we skip the modifiers):

$$
\textit{Type}\ \text{`this'}\ \text{`['}\ [\textit{Params}]\ \text{`]'}\ \text{`\{'}\ [\text{`get'}\ \textit{Block}]\ [\text{`set'}\ \textit{Block}]\ \text{`\}'}
$$

Analogously to the constraints for properties, for an indexer of type T with parameters p, the body of a `get` accessor is the body of a method with parameters p and return type T, the body of a `set` accessor is the body of a `void` method with parameters p and an implicit value parameter named `value` of type T. A base class indexer can be accessed by `base[`$exps$`]`. Using the signatures T `get_Item(`$params$`)` and `void set_Item(`$params$`,`T `value)`, which are reserved for get and set accessors, the intended semantics of indexers is reduced to the semantics of arrays and methods via the following compile-time translation (and corresponding operator expression translation as explained

---

[22] Without knowing whether it is accessed directly or whether an accessor method is being called.

for properties):

$$
\begin{array}{ll}
T \text{ this}[params] \{ & T \text{ get\_Item}(params) \{ getAccessor \} \\
\quad \text{get} \{ getAccessor \} & \text{void set\_Item}(params, T \text{ value}) \{ \\
\quad \text{set} \{ setAccessor \} \Longrightarrow & \quad setAccessor \\
\} & \}
\end{array}
$$

Events. Events can be declared in C♯ like fields as follows: in the form 'event' *DelegateType Identifier* ';' (we omit the modifiers), or like properties, in the form

    'event' *DelegateType Identifier* '{' 'add' *Block* 'remove' *Block* '}'.

Outside the type that contains the declaration, an event X can only be used as the left-hand operand of += and -= in expressions X += *exp* and X -= *exp* of type void; within the type that contains the declaration, field-like events can be used like fields of delegate types. The accessors of property-like events have to be bodies of void methods with an implicit parameter value of *DelegateType*.

    The semantics of events in C♯ follows the *Publish/Subscribe* pattern. A class publishes an event it can raise, so that any number of classes can subscribe to that event. When the event is actually raised, each subscriber is notified that the event has occurred, namely by calling a delegate whose invocation list is executed with the sender object and the event data as its arguments. This idea is realized as follows. The *event sender* class that raises an event named X has the member event X_EventHandler X; where the delegate type X_EventHandler for the event is declared as follows (with two arguments, the first one for the publisher and the second one for the event information object, which must be derived from the class EventArgs):

    delegate void X_EventHandler(object *sender*, X_EventArgs *e*);

To consume the event, the *event receiver* declares an event-handling method Receive_X with the same signature as the event delegate:

    void Receive_X(object *sender*, X_EventArgs *e*) { ... }

To *register* the event handler, the event receiver has to add the Receive_X method to the event X of the event sender object:

    X += new X_EventHandler(this.Receive_X);

The event sender *raises* the event by invoking the invocation list of X with the sender object and the event data, e.g.

    void On_X(X_EventArgs *e*) { if (X != null) X(this,*e*); }

It suffices to assign a meaning to the signatures void add_X(D value) and void remove_X(D value), which are reserved for every event X of delegate type D. This is done by the following translation of field-like events, anticipating the lock statement

of C♯$_\mathcal{T}$ which is explained in [35]. [23]

```
class C {
  private D _ _X;
  void add_X(D value) {
    lock (this) { _ _X = _ _X + value; }
  }
  void remove_X(D value) {
    lock (this) { _ _X = _ _X - value; }
  }
}
```

6.2.3. Further constructs

Similar syntactical reductions to those given above can be used to define the semantics of overloaded standard mathematical operators and user-defined conversions, of enumeration related statements 'foreach (*T x* in *exp*) *stm*', of using statements 'using (*resource*) *stm*', of parameter arrays and of attributes.

7. Refinement C♯$_\mathcal{U}$ by pointers in unsafe code

In this section we add the features C♯ offers for using pointers (coming with address-of and dereferencing operators '&', '*', '->' together with pointer arithmetic) to directly work on memory addresses, bypassing the type checking by the compiler—hence the name 'unsafe' code blocks. Java has no such unsafe extension. The extension includes a mechanism called *pinning* of objects to prevent the runtime during the execution of a 'fixed' statement to manage via the garbage collector memory one wants to address directly. Code for which (de-) allocation is not controlled by the runtime is called *unmanaged*. As an alternative to pinning, data of unmanaged type can also be 'stackalloc'ated, instead of using the heap.

The refinement consists, besides some new rules, mainly in a definition of the *mem*ory function in terms of byte sequences. This is a typical data refinement, using an encoding of simple types and a corresponding refinement of the function *structField*.

7.1. Signature refinement for C♯$_\mathcal{U}$

We refine *Type* by adding pointer types to value and reference types:

$$Type \quad ::= ValueType \mid RefType \mid PointerType$$
$$PointerType ::= UnmanagedType \text{ '*'} \mid \text{'void' '*'}$$

where *unmanaged* types are types which are not managed and *managed* types are recursively defined as (a) reference types or (b) struct types that contain a field of a managed type

---

[23] If one prefers to be independent of the thread model C♯$_\mathcal{T}$, one can consider lock statements lock (*exp*) *stm* translated for single-thread execution by { Object *o* = *exp*; *stm* } (with a fresh variable *o*), which is then refined in C♯$_\mathcal{T}$ for the multiple thread execution model.

Table 6
Type constraints for unsafe expressions

Expression	Constraints	Expression type
`sizeof(`t`)`	t unmanaged type	`int`
$*e$	$type(e) = T*, T \neq$ `void`	T
$\&v$	v a fixed variable	$T*$, where $T = type(v)$
$e\ \texttt{->}\ m$	$type(e) = T*, T \neq$ `void`	$type(T::m)$
$e[i]$	$type(e) = T*, T \neq$ `void`, $type(i)$ integral	T

or a pointer to a managed type. The subtype relation is extended to pointer types such that $\Lambda \leqslant T* \leqslant$ `void*`. *Exp* and *Vexp* are extended by address-of and dereferencing expressions and expressions to denote the values of a new function indicating the '`sizeof`' unmanaged types. *Stm* is extended to reflect '`unsafe`' code blocks, '`fixed`' statements and '`stackalloc`'ation of arrays. '`unsafe`' can also appear as modifier for classes, structs, interfaces, delegates as well as for fields, methods, properties, indexers, operators, events, constructors, destructors.

$$Exp \quad ::= \ldots \mid \text{`\&'} \, Vexp \mid Exp \text{ `->' } Meth \ (\ [Args]\ ) \mid Exp \text{ `->' } Field$$
$$\mid \text{ `sizeof' `(' } UnmanagedType \text{ `)'}$$

$$Vexp \quad ::= \ldots \mid \text{ `*' } Exp$$

$$Stm \quad ::= \ldots \mid \text{ `unsafe' } Block \mid \text{ `fixed' `(' } PointerType \, Loc = Exp \text{ `)' } Stm$$

$$Bstm \quad ::= \ldots \mid PointerType \, Loc \text{ `='}$$
$$\text{ `stackalloc' } UnmanagedType \text{ `[' } Exp \text{ `]' `;'}$$

In the following expressions, the basic arithmetical operators are used for pointer increment and decrement, pointer addition and subtraction, pointer comparison, and pointer conversion (where p and q are of a pointer type, i is of integer type):

- $\texttt{++}p, \texttt{--}p, p\texttt{++}, p\texttt{--}, p + i, i + p, p - i, p - q, p \texttt{ == } q, p \texttt{ != } q, p \texttt{ < } q, p \texttt{ <= } q, p \texttt{ > } q, p \texttt{ >= } q$
- $(T*)i, (T*)p, (\texttt{int})p, (\texttt{uint})p, (\texttt{long})p, (\texttt{ulong})p$

On the types of the new expressions the constraints in Table 6 are imposed. We assume the dereferencing and member access operator $e\texttt{->}m$ to be translated to $(*e).m$, similarly $e[i]$ to $*(e + i)$.

For statements the following type constraints in Table 7 are assumed. A variable is called *moveable* (by the garbage collector) iff it is not a fixed variable. Fixed variables are (by recursive definition): local variables, value parameters, $*exp$ for exp of pointer type, and instance field expressions $v.f$ if v is a fixed variable of struct type T and f is an instance field of T.

The local variable p in the fixed statement is called a *pinned* local variable. A pinned local variable is a read-only variable. It is not allowed to assign a new value to it in the body of the fixed statement.

The principal refinement in the ASM extension EXECCSHARP$_U$ for C♯$_\mathcal{U}$ is that of the *mem*ory together with its operators, where the set of *SimpleValue*s is replaced by *Byte*s

Table 7
Type constraints for unsafe statements

Statement	Constraints
$T*$ p = stackalloc $T[exp]$;	$type(exp)$ = int, T unmanaged
fixed (char* $p = exp$) stm	$type(exp)$ = string, p read-only in stm
fixed ($T*$ $p = exp$) stm	$type(exp)$ = $T[R]$, T unmanaged,
	p read-only in stm
fixed ($T*$ $p = \&vexp$) stm	$type(vexp)$ = T, T unmanaged,
	$vexp$ a moveable variable,
	p read-only in stm

(8-bit strings), using non-negative integers as memory addresses ($Adr = \mathbb{N}$):

$$mem: Adr \rightarrow Byte \cup Ref \cup \{Undef\}$$

The partial functions to *encode* (resp. *decode*) values of a given simple type T by byte sequences, of a length (number of bytes) depending on $sizeOf(T)$, satisfy for values *val* the equations

$$decode(T, encode(val)) = val \quad \text{and} \quad length(encode(val)) = sizeOf(T).$$

For every pointer type $T*$ holds $sizeOf(T*) = sizeOf(\text{void*})$.

A function $fieldOffset: UnmanagedStructType \times Field \rightarrow \mathbb{N}$ is used to describe the layout of unmanaged structs. It has to satisfy the following constraint for every unmanaged struct type T and instance field f of T (overlapping fields are allowed in C$\sharp_\mathcal{U}$):

$$fieldOffset(T, f) + sizeOf(type(f)) \leqslant sizeOf(T)$$

We assume that if *adr* is an address allocated using $new(Adr, T)$ for struct type T, then for every instance field f of T the equation $fieldAdr(adr, f) = adr + fieldOffset(T, f)$ holds.

To determine the layout of arrays with unmanaged element type we stipulate the following refinement of the function *elemAdr* which reflects that array elements are stored such that the indices of the right most dimension are increased first, then the next left dimension, and so on. For $runTimeType(ref) = T[dim(n)]$, where T is an unmanaged type and $l_i = dimLength(ref, i-1)$ for $i \in [1..n]$, we assume the following:

$elemAdr(ref, [i_1, i_2, \ldots, i_n]) =$
    $elemAdr(ref, [0, \ldots, 0]) + (\ldots (i_1 \cdot l_2 + i_2) \cdot l_3 + \ldots + i_n) \cdot sizeOf(T)$

7.2. Transition rule refinement for unsafe code

Besides the rules below which define the semantics of the new expressions and statements, to be added to EXECCSHARP$_D$, we have to data refine the notions of reading from and writing to memory for values of unmanaged type.

$memValue(adr, t) =$
    if $t \in RefType$ **then** $mem(adr)$

    elseif $t \in UnmanagedType$ **then**
      $[mem(adr + i) \mid i \in [0 \, .. \, sizeOf(t) - 1]\,]$
    elseif $t \in StructType$ **then**
      $\{f \mapsto getField(adr, f) \mid f \in instanceFields(t)\}$

$getField(adr, t::f) = memValue(fieldAdr(adr, t::f), type(t::f))$

WRITEMEM$(adr, t, val) \equiv$
  if $t \in RefType$ **then** $mem(adr) := val$
  elseif $t \in UnmanagedType$ **then**
    forall $i \in [0 \, .. \, sizeOf(t) - 1]$ **do** $mem(adr + i) := val(i)$
  elseif $t \in StructType$ **then**
    forall $f \in instanceFields(t)$ **do** SETFIELD$(adr, f, val(f))$

SETFIELD$(adr, t::f, val) \equiv$ WRITEMEM$(fieldAdr(adr, t::f), type(t::f), val)$

Values of unmanaged struct types are directly represented as sequences of bytes. Hence, the function *structField* has to be refined to extract a subsequence in case of unmanaged struct types:

$structField(val, T, f) =$
  if $T \in ManagedType$ **then** $val(f)$
  else let $n = fieldOffset(T, f)$ **in** $[val(i) \mid n \leqslant i < n + sizeOf(type(f))]$

In the rules for EXECCSHARPEXP$_U$ we have & $\square$ as additional address position. We follow the implementation in Rotor and .NET in formulating the *Null* check to prevent writing to null addresses; the ECMA standard describes this check as optional.

EXECCSHARPEXP$_U \equiv$ **match** $context(pos)$
  sizeof(T) $\to$ YIELD$(sizeOf(T))$

  &exp $\to pos := exp$
  &$^{\blacktriangleright}adr \to$ YIELDUP(adr)

  *exp $\to pos := exp$
  *$^{\blacktriangleright}adr \to$ **if** $adr = Null$ **then**   // null pointer check optional
            FAILUP(NullReferenceException)
         else YIELDUPINDIRECT(adr)

  *$exp_1 = exp_2$ $\to pos := exp_1$
  *$^{\blacktriangleright}adr = exp_2 \to pos := exp_2$
  *$adr = {}^{\blacktriangleright}val$ $\to$ **if** $adr = Null$ **then**   // null pointer check optional
               FAILUP(NullReferenceException)
            else
              WRITEMEM$(adr, type(pos), val)$
              YIELDUP(val)

The rules for pointer arithmetic can be summarized as follows:

$Apply(+(T\,*, \texttt{int}), adr, i) = adr + i \cdot sizeOf(T)$
$Apply(+(\texttt{int}, T\,*), i, adr) = adr + i \cdot sizeOf(T)$
$Apply(-(T\,*, T\,*), adr_1, adr_2) = (adr_1 - adr_2)/sizeOf(T)$

$Convert(T*, adr) = adr = Convert(S, adr)$
  for $S \in \{$int, uint, long, ulong$\}$
$Convert(T*, i) = i$

In the execution of the stackalloc statement we assume that $new(adr, T, i)$ allocates i consecutive chunks of addresses of size $sizeOf(T)$ which are later de-allocated on method exit in FREELOCALS.

EXECCSHARPSTM$_U$ ≡ **match** *context*(*pos*)
  unsafe *block*   → *pos* := *block*
  unsafe ▶*Norm* → YIELDUP(*Norm*)

  $T* loc$ = stackalloc $T[exp]$; → *pos* := *exp*
  $T* loc$ = stackalloc $T[▶i]$; → **let** $adr = new(Adr, T, i)$ **in**
                                                    WRITEMEM(*locals*(*loc*), $T*$, *adr*)
                                                    YIELDUP(*Norm*)

The run-time execution of fixed statements can be explained by syntactical transformations:

Statement	Run-time execution
fixed (char* $p = exp$) *stm*	{ char* p; $p = Cstring(exp)$; *stm* }
fixed ($T* p = exp$) *stm*	{ $T* p$; $p = \&exp[0]$; *stm* }
fixed ($T* p = \&vexp$) *stm*	{ $T* p$; $p = \&vexp$; *stm* }

In the first case, it is assumed that $Cstring(s)$ is an internal function that returns the address of the first element of a C-style null-terminated character array representation of the string s. How it is related to the original representation of the string is not specified in [27].

8. Related work and conclusion

One of our referees would like to see a critical assessment of the ASM method we used for this work and a comparison to alternative approaches. Some justification of the kind from the perspective of semantic methods for programming languages has been given in [6, Section 4], containing concrete illustrations of and references to the numerous and earlier competing proposals. This was at a time when ASMs were applied for the first time to (successfully) specify an industrial language standard, namely the ISO Prolog standard [11]. A decade later, a broader comparison of the then well-developed ASM method with respect to other system design and analysis frameworks has been provided in [5,7]. However, a systematic, comprehensive and authoritative evaluation of the multitude of system design and analysis proposals in the literature remains a highly desirable and challenging task to be accomplished, even if limited to the use of the major so-called formal methods for the development and investigation of programming languages and their implementations. From the perspective of practical system design and analysis some comparative studies of

this kind have been published, see e.g. [1,17,28] (the interested reader may also consult the corresponding ASM-based work in [4,13,26]). For work centered around Java and the JVM the reader finds in [2] a collection of formal-method approaches to language specification and analysis; [24] contains an excellent, detailed and at the time complete review of the huge literature on the subject (including an evaluation of the ASM-based Java/JVM investigations), with a focus on safety issues and their impact on smart cards. We cannot perform here a similar analysis for work on C♯ or other major programming languages. This explains why the references in this paper stick to C♯ documentation from ECMA and Microsoft and to some ASM work we have built upon directly.

For the work presented in this paper we set ourselves a more modest though not completely trivial major goal, namely to test whether the method developed in [36] for the definition and a proven to be correct implementation of a real-life programming language like Java scales naturally to the somewhat richer and more complex C♯. It is up to the reader to judge whether this ASM reuse case study for a real-life complex model succeeded. For the formalization of other programming languages something can also be learnt directly from the formalization of the semantics of C♯ worked out here. For example, how to "divide and conquer" the static and the dynamic semantics of a language, how to separate the description of conceptually independent programming constructs by dividing them into sublanguages, how to unify and streamline the formalization of similar constructs by appropriate parameterizations (which means abstractions), how to model and evaluate variations of specific features (e.g. expression evaluation, parameter passing mechanism, class initialization, etc.) by varying macros, rules and/or domains together with their operations, how to extend within a single framework the model for a language core by a form of bootstrapping (including in particular syntactical translations) to a model for the entire language, etc.

There are several by-products of the work presented here. Through the ASM-model-oriented analysis of the ECMA standard for C♯ we found several bugs and gaps in the formulation of the standard and in its .NET implementation as well as some incoherences between the two, as documented in detail in [20] in terms of our ASM model for C♯. Another by-product of the high-level modular interpreter defined here is the support it provides to teachers of C♯, in particular if they want to shed light on certain subtle language features which are not clarified by the ECMA documents. In the forthcoming paper [15] we are going to work out a concrete comparison of the two models for C♯ and for Java, which will allow us to formulate in a precise technical manner where and in which respect the two languages differ among each other and from other programming languages—methodologically, semantically and pragmatically. As a specific part of this reuse-case-study the second author is investigating how the main new features of C♯ 2.0 can be modeled by appropriate extensions of the ASM model developed here for C♯, in particular generic types (parametric polymorphism), anonymous methods and iterators. Last but not least, with our C♯ model and its extension to threads in [35] we have laid the ground for a mathematical analysis and possibly mechanical verification of interesting properties of the language and its implementation, like type safety,[24] compiler correctness and completeness, correctness of

---

[24] For a fragment of Microsoft's Intermediate Language, which is executed by Microsoft's Common Language Runtime, a type safety proof has been given in [23], based upon Syme's method [38] for writing functional specifications which can be subject to theorem proving in HOL.

(a mathematical model to be developed for) garbage collection, security, etc. For the correctness of the definite assignment analysis performed by a C♯ compiler, we may refer the interested reader to [21]. We hope somebody will feel challenged to use our model for precisely formulating and proving such theorems for C♯ and to build a corresponding model for Microsoft's Common Language Runtime together with a compilation scheme from C♯ to IL code, applying to our model the powerful ASM refinement technique [9] along the lines shown in the ASM-based Java/JVM study in [36].

The questions asked by our referees lead us to mention another practical and industrially viable use that can be made of a modeling and analysis activity as the one reported in this paper, except if the extreme time pressure usually imposed on developers to produce executable code from incomplete verbal specifications (mostly formulated in natural language) prevents them from at least once trying out a more reliable option. Here is a concrete example what could have been done. On September 27, 2000, the penultimate day of his sabbatical stay with Microsoft Research in Redmond, in a seminar talk to representatives of the C♯ development team, the first author suggested to use the method, at the time formulated and presented in terms of Java/JVM for publication in what became the Jbook [36], for the following five fundamental activities in relation to the at-the-time ongoing development of what became known as the C♯ language with the underlying CLR virtual machine:

- defining an ASM model as *executable specification of critical language constructs* or layers (if not of the entire language) and of the mapping to IL code,
- *generating test cases* for the implementing code from the ASM model,
- using the ASM model as *oracle for test evaluations* and for comparing model test runs with code test runs,
- using the ASM model as *internal documentation* for future language extensions and for relating other .NET languages to C♯, in particular those which are equipped already with an ASM model of their semantics,
- using the ASM model as *basis for writing innovative handbooks* for users and for maintenance professionals, where the innovative character derives from being (a) accurate yet simple and easy to understand, (b) complete and detailed yet succinct.

Acknowledgements

References

[1] J.-R. Abrial, E. Börger, H. Langmaack, The steam boiler case study: competition of formal program specification and development methods, in: J.-R. Abrial, E. Börger, H. Langmaack (Eds.), Formal Methods for Industrial Applications. Specifying and Programming the Steam-Boiler Control, Lecture Notes in Computer Science, Vol. 1165, Springer, Berlin, 1996, pp. 1–12.
[2] J. Alves-Foss, Formal Syntax and Semantics of Java, Lecture Notes in Computer Science, Vol. 1523, Springer, Berlin, 1998.

[3] T. Archer, A. Whitechapel, Inside C♯, Microsoft Press, 2002.

[4] C. Beierle, E. Börger, I. Durdanović, U. Glässer, E. Riccobene, Refining abstract machine specifications of the steam boiler control to well documented executable code, in: J.-R. Abrial, E. Börger, H. Langmaack (Eds.), Formal Methods for Industrial Applications. Specifying and Programming the Steam-Boiler Control, Lecture Notes in Computer Science, Vol. 1165, Springer, Berlin, 1996, pp. 62–78.

[5] E. Börger, Abstract state machines: a unifying view of models of computation and of system design frameworks, Ann. Pure Appl. Logic (2005).

[6] E. Börger, A logical operational semantics for full Prolog. Part I: selection core and control, in: E. Börger, H. Kleine Büning, M.M. Richter, W. Schönfeld (Eds.), CSL'89. 3rd Workshop on Computer Science Logic, Lecture Notes in Computer Science, Vol. 440, Springer, Berlin, 1990, pp. 36–64.

[7] E. Börger, High-level system design and analysis using abstract state machines, in: D. Hutter, W. Stephan, P. Traverso, M. Ullmann (Eds.), Current Trends in Applied Formal Methods (FM-Trends 98), Lecture Notes in Computer Science, Vol. 1641, Springer, Berlin, 1999, pp. 1–43.

[8] E. Börger, The ASM ground model method as a foundation of requirements engineering, in: N. Dershowitz (Ed.), Manna-Sympo., Lecture Notes in Computer Science, Vol. 2772, Springer, Berlin, 2003.

[9] E. Börger, The ASM refinement method, Formal Aspects Comput. 15 (2003) 237–257.

[10] E. Börger, T. Bolognesi, Remarks on turbo ASMs for computing functional equations and recursion schemes, in: E. Börger, A. Gargantini, E. Riccobene (Eds.), Abstract State Machines 2003—Advances in Theory and Applications, Lecture Notes in Computer Science, Vol. 2589, Springer, Berlin, 2003, pp. 218–228.

[11] E. Börger, K. Dässler, Prolog: DIN papers for discussion, ISO/IEC JTCI SC22 WG17 Prolog Standardization Document 58, National Physical Laboratory, Middlesex, England, 1990.

[12] E. Börger, N.G. Fruja, V. Gervasi, R. Stärk, A complete formal definition of the semantics of C♯, Technical Report, 2004, in preparation.

[13] E. Börger, L. Mearelli, Integrating ASMs into the software development life cycle, J. Univ. Comput. Sci. 3 (5) (1997) 603–665.

[14] E. Börger, J. Schmid, Composition and submachine concepts for sequential ASMs, in: P. Clote, H. Schwichtenberg (Eds.), Computer Science Logic (Proc. CSL, 2000), Lecture Notes in Computer Science, Vol. 1862, Springer, Berlin, 2000, pp. 41–60.

[15] E. Börger, R.F. Stärk, Exploiting abstraction for specification reuse. The Java/C♯ case study, in: F.S. de Boer, M.M. Bonsangue, S. Graf, W.-P. de Roever (Eds.), Proc. FMCO'03, Lecture Notes in Computer Science, Springer, Berlin, 2004.

[16] E. Börger, R.F. Stärk, Abstract State Machines. A Method for High-Level System Design and Analysis, Springer, Berlin, 2003.

[17] M. Broy, S. Merz, K. Spies, Formal Systems Specification—The RPC-Memory Specification Case Study, Lecture Notes in Computer Science, Vol. 1169, Springer, Berlin, 1996.

[18] Common Language Infrastructure, Standard ECMA-335, 2003, http://www.ecma-international.org.

[19] Foundations of Software Engineering Group, Microsoft Research, AsmL, 2001, web pages at http://research.microsoft.com/foundations/AsmL/.

[20] N.G. Fruja, Specification and implementation problems for C♯, in: B. Thalheim, W. Zimmermann (Eds.), Abstract State Machines 2004, Lecture Notes in Computer Science, Springer, Berlin, 2004.

[21] N.G. Fruja, The correctness of the definite assignment analysis in C♯, J. Object Technology 3 (9) (2004).

[22] N.G. Fruja, R.F. Stärk, The hidden computation steps of turbo abstract state machines, in: E. Börger, A. Gargantini, E. Riccobene (Eds.), Abstract State Machines 2003—Advances in Theory and Applications, Lecture Notes in Computer Science, Vol. 2589, Springer, Berlin, 2003, pp. 244–262.

[23] A.D. Gordon, D. Syme, Typing a multi-language intermediate code, in: Proc. 28th Ann. ACM SIGPLAN-SIGACT Symp. on Principles of Programming Languages, London, 2001, pp. 248–260.

[24] P. Hartel, L. Moreau, Formalizing the safety of Java, the Java Virtual Machine and Java Card, ACM Comput. Surv. 33 (4) (2001) 517–558.

[25] A. Hejlsberg, S. Wiltamuth, P. Golde, C♯ Language Specification, Addison-Wesley, Reading, MA, 2003.

[26] J. Huggins, Broy–Lamport specification problem: a Gurevich abstract state machine solution, Tech. Report CSE-TR-320-96, EECS Department, University of Michigan, 1996.

[27] C♯ Language Specification, Standard ECMA-334, 2001, http://www.ecma-international.org.

[28] C. Lewerentz, T. Lindner, Formal development of reactive systems, Case Study "Production Cell", Lecture Notes in Computer Science, Vol. 891, Springer, Berlin, 1995.

[29] Mono compiler for C♯, http://www.go-mono.com/c-sharp.html.

[30] J. Prosise, Programming Microsoft .NET, Microsoft Press, 2002.

[31] J. Richter, Applied Microsoft .NET Framework Programming, Microsoft Press, 2002.

[32] J. Schmid, Refinement and implementation techniques for abstract state machines, Ph.D. Thesis, University of Ulm, Germany, 2002.

[33] J. Schmid, Executing ASM specifications with AsmGofer, web pages at http://www.tydo.de/AsmGofer.

[34] SSCLI (Rotor) web site, http://www.sscli.net

[35] R.F. Stärk, E. Börger, An ASM specification of C♯ threads and the .NET memory model, in: B. Thalheim, W. Zimmermann (Eds.), Abstract State Machines 2004, Lecture Notes in Computer Science, Springer, Berlin, 2004.

[36] R.F. Stärk, J. Schmid, E. Börger, Java and the Java Virtual Machine—Definition, Verification, Validation, Springer, Berlin, 2001.

[37] D. Stutz, T. Neward, G. Shilling, Shared Source CLI Essentials, O'Reilly, Sebastopol, 2003.

[38] D. Syme, Declarative theorem proving for operational semantics, Ph.D. Thesis, University of Cambridge, 1998.

[39] Visual Studio .NET 2003, http://msdn.microsoft.com/vstudio/

[40] W. Zimmermann, A. Dold, A framework for modeling the semantics of expression evaluation with abstract state machines, in: E. Börger, A. Gargantini, E. Riccobene (Eds.), Abstract State Machines 2003—Advances in Theory and Applications, Lecture Notes in Computer Science, Vol. 2589, Springer, Berlin, 2003, pp. 391–406.