DARYA MELICHER and ANLUN XU, School of Computer Science, Carnegie Mellon University VALERIE ZHAO, University of Chicago

ALEX POTANIN, School of Engineering and Computer Science, Victoria University of Wellington JONATHAN ALDRICH, School of Computer Science, Carnegie Mellon University

Effect systems have been a subject of active research for nearly four decades, with the most notable practical example being checked exceptions in programming languages such as Java. While many exception systems support abstraction, aggregation, and hierarchy (e.g., via class declaration and subclassing mechanisms), it is rare to see such expressive power in more generic effect systems. We designed an effect system around the idea of protecting system resources and incorporated our effect system into the Wyvern programming language. Similar to type members, a Wyvern object can have effect members that can abstract lower-level effects, allow for aggregation, and have both lower and upper bounds, providing for a granular effect hierarchy. We argue that Wyvern's effects capture the right balance of expressiveness and power from the programming language design perspective. We present a full formalization of our effect-system design, showing that it allows reasoning about authority and attenuation. Our approach is evaluated through a security-related case study.

CCS Concepts: • Software and its engineering \rightarrow Object oriented languages; • Security and privacy \rightarrow Software security engineering;

Additional Key Words and Phrases: Effects, effect system, expressiveness, abstraction, language-based security

ACM Reference format:

Darya Melicher, Anlun Xu, Valerie Zhao, Alex Potanin, and Jonathan Aldrich. 2022. Bounded Abstract Effects. *ACM Trans. Program. Lang. Syst.* 44, 1, Article 5 (January 2022), 48 pages. https://doi.org/10.1145/3492427

1 INTRODUCTION

An effect system can be used to reason about the side effects of code, such as reads and writes to memory, exceptions, and I/O operations. Java's checked exceptions is a simple effect system that has found widespread use, and interest is growing in effect systems for reasoning about security [Turbak and Gifford 2008], memory effects [Lucassen and Gifford 1988], asynchronous event streams [Bračevac et al. 2018], and concurrency [Bocchino et al. 2009; Dolan et al. 2017].

Darya Melicher and Anlun Xu equal contribution.

Authors' addresses: D. Melicher, A. Xu, and J. Aldrich, School of Computer Science, Carnegie Mellon University, 5000 Forbes Ave. Pittsburgh, PA, 15213; emails: darya@cs.cmu.edu, anlunx@andrew.cmu.edu, jonathan.aldrich@cs.cmu.edu; V. Zhao, University of Chicago, 5801 S Ellis Ave, Chicago, IL, 60637, United States; email: vzhao@uchicago.edu; A. Potanin, School of Engineering and Computer Science, Victoria University of Wellington, P.O. Box 600, Wellington, New Zealand, Wellington, NZ, 6144; email: alex@ecs.vuw.ac.nz.



This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2022 Copyright held by the owner/author(s). 0164-0925/2022/01-ART5 \$15.00 https://doi.org/10.1145/3492427

5:2 D. Melicher et al.

Requirements for a scalable effect system. Unfortunately, effect systems have not been widely adopted other than checked exceptions in Java, a feature that is widely viewed as problematic [van Dooren and Steegmans 2005]. The root of the problem is that existing effect systems do not provide adequate support for scaling to programs that are larger and have complex structure. Any adequate solution must support *effect composition* and *effect abstraction*.

Composition and abstraction are keys to achieving scale in general. We define *effect composition* as the ability to define higher-level effects in terms of lower-level effects. For example, a database component might compose multiple lower-level effects such as file.Read and network.Access into a single, higher-level effect db.Query. Of course, we may want to ensure that clients do not depend on the particular implementation of the db.Query effect in case it changes. Thus, we define *effect abstraction* as the ability to *hide* the concrete definition of an effect from clients. In this case, db.Query becomes an *abstract effect*, similar in spirit to an abstract type [Mitchell and Plotkin 1988] and supported with an *effect member* construct that is directly analogous to abstract type members in Scala [Odersky and Zenger 2005]. Abstraction and composition can be used hierarchically. For example, the file.Read effect could abstract a lower-level system.FFI effect. Then, clients of a file should be able to reason about side effects in terms of file reads and writes, not in terms of the low-level calls that are made to the foreign function interface (FFI).

Effect polymorphism allows functions to be reused with different function arguments with different effects [Lucassen and Gifford 1988]. In systems at a larger scale, there are various possible effects, and each program component may cause different effects. With effect polymorphism, we can write general code that handles objects with different effects, thereby reducing the amount of replicated code. In practice, we have found that to make effects work well with modules, it is essential to extend effect polymorphism by assigning bounds to effect parameters. Therefore, we introduce bounded abstract effects, which allows programmers to define upper and lower bounds both on abstract effects and on polymorphic effect parameters.

We leverage *path-dependent effects*, that is, effects whose definitions depend on an object, as the foundation of our effect system. This adds expressiveness; for example, if we have two File objects, x and y, we can distinguish effects on one file from effects on the other: the effects x.Read and y.Read are distinct. Path-dependent effects are particularly important in the context of modules, in which two different modules may implement the same abstract effect in different ways. For example, it may be important to distinguish db1.Query from db2.Query if db1 is an interface to a database stored in the local file system whereas db2 is a database accessed over the network.

Design of the effect system in Wyvern. This article presents a novel and scalable effect-system design that supports effect abstraction and composition. The abstraction facility of our effect system is inspired by type members in languages such as Scala. Just as Scala objects may define type members, in our effect calculus, any object may define one or more *effect members*. An effect member defines a new effect in terms of the lower-level effects that are used to implement it. The set of lower-level effects may be empty in the base case or may include low-level effects that are hard-coded in the system. Type ascription can enable information hiding by concealing the definition of an effect member from the containing object's clients. In addition to completely concealing the definition of an effect, our calculus provides bounded abstraction, which exposes upper or lower bounds of the definition of an effect while still hiding the definition of it.

Just as Scala's type members can be used to encode parametric polymorphism over types, our effect members double as a way to provide effect polymorphism. Bounded effect polymorphism is also provided in our system, because abstract effect members can be bounded by upper or lower bounds. We follow numerous prior Scala formalisms in including polymorphism via this encoding rather than explicitly. This keeps the formal system simpler without giving up expressive power.

Finally, because effect members are defined on objects, our effects are *generative*, even dynamically [Dreyer et al. 2003]. This yields great expressivity: each object created at runtime defines a new effect for each effect member in that object so that, for example, we can separately track effects on different File objects, statically distinguishing the effects on one object from the effects on another.

Evaluation and Security Applications. A promising area of application for effects is software security. For example, in the setting of mobile code, Turbak and Gifford [2008] proposed that effects could be used to ensure that any untrusted code we download can access only the system resources it needs to do its tasks, thus following the principle of least privilege [Denning 1976]. We are not aware of prior work that explores this idea in depth.

In order to evaluate our design for effect abstraction, we have incorporated it into an effect system that tracks the use of system resources such as the file system, network, and keyboard. Our effect system is intended to help developers reason about which source code modules use these resources. Through the use of abstraction, we can "lift" low-level resources such as the file system into higher-level resources such as a logging facility or a database and enable application code to reason in terms of effects on those higher-level resources when appropriate. In fact, even the use of resources such as the file system is scaffolded as an abstraction on top of a primitive system.FFI effect that our system attaches to uses of the language's foreign function interface. A set of illustrative examples demonstrates the benefits of abstraction for effect aggregation as well as for information hiding and software evolution. Finally, we show how our effect system allows us to reason about the *authority* [Miller 2006] of code, that is, what effects a component can have, as well as the *attenuation* of that authority. We will demonstrate the application of our effect system to security via a case study. However, the security definitions are speculative due to the fact that we do not have formal proofs as validation, although they are based on ideas that were formalized and proved by Craig et al. [2018].

Our effect system is implemented in the context of Wyvern, a programming language designed for highly productive development of secure software systems. In this article, we give several concrete examples of how our effect-system design can be used in software production, all of which is functional Wyvern code that runs in the Wyvern regression test suite.

Outline and Contributions. Here, we describe the main contributions of our article, followed by a running example in the next section.

- The design of a novel effect system fulfilling the requirements outlined earlier. Our system is the first standalone language to bring together effect abstraction and composition with the effect member construct. Ours is also the first system to provide the programmer with a general form of bounded effect polymorphism and bounded effect abstraction, supporting upper and lower bounds that are other arbitrary effects (Section 3).
- The application of our effect system to a number of forms of security reasoning, illustrating its expressiveness and making the benefits described earlier concrete (Section 4).
- A precise, formal description of our effect system and proof of its soundness. Our formal system shows how to generalize and enrich earlier work on path-dependent effects by leveraging the type theory of DOT (Section 5).
- A formalization of authority using effects and of authority attenuation (Section 5.7).
- A feasibility demonstration via the implementation of our approach in the Wyvern programming language (Section 6).

The last sections in the article discuss related work and our conclusions.

5:4 D. Melicher et al.

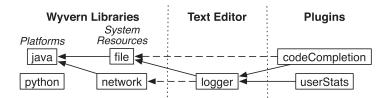


Fig. 1. The overall architecture of the text-editor application. Boxes represent modules and the arrows represent module imports. The solid arrows are imports that take place, and the dashed arrows represent potential imports that may or may not occur.

2 RUNNING EXAMPLE

Drawing inspiration from a recent report on security vulnerabilities in text editors [Azouri 2018], we use a text editor application as a running example to demonstrate the key features of our effect system design. The overall architecture of this application is shown in Figure 1. Each box in the diagram represents a module and the arrows represent module imports. For the purposes of our forthcoming examples, the solid arrows are imports that take place and the dashed arrows represent potential imports that may or may not occur.

The application is written using Wyvern's libraries, which contain modules representing system resources, such as the file system and network. These modules rely on access to native backend modules, such as <code>java</code> and <code>python</code>, which are Wyvern's Java and Python backends, respectively. In the text editor, we focus only on the <code>logger</code> module that implements the logging facility of the application. The text editor allows supplementing its core functionality with various third-party plugins. We assume that the application requires that all plugins and user-facing modules of the text editor itself update the log file with the user-observable actions they perform. The result produced by the logger module can be used by either the user of the text editor or telemetry, which helps the developer of the text editor to analyze the performance of the program. In our examples, we use two sample plugins: one that, as the user types in code, detects code patterns and offers to complete the code for them and another that analyzes the text editor's log file and provides insight into how the text-editor application is used.

The dashed vertical lines represent the conceptual boundaries between parts of the application that vary in the level of trust based on the security of the contained code. Modules in the Wyvern libraries are the most trusted since they provide functionality essential for all applications developed in Wyvern and were written with security in mind. Modules of the text editor application are less trusted since they are more likely to contain fallible code. Finally, the plugins are the least trusted since they are written by third parties and may be error-prone, vulnerable to exploitation, or outright malicious.

In the following, we will see how Wyvern's effect system can shed light on what effects each module in the system can have in terms of the effect abstractions provided by the modules it depends on—ensuring that security vulnerabilities that are caused by modules exceeding their authority are caught by effect checking.

3 WYVERN EFFECTS BASICS

Wyvern is a programming language that supports a first-class module system with abstract types. The type system of Wyvern is structural and supports path-dependent types similar to DOT. Moreover, Wyvern is designed as a capability-safe language, supporting a least-privilege approach to security and enabling architects to enforce a number of important design constraints [Melicher et al. 2017]. This paper focuses on a novel effect system based on the Wyvern programming language.

```
resource type Logger
effect ReadLog
effect UpdateLog
def readLog(): {this.ReadLog} String
def updateLog(newEntry: String): {this.UpdateLog} Unit

module def logger(f: File): Logger
effect ReadLog = {f.Read}
effect UpdateLog = {f.Append}
def readLog(): {ReadLog} String = f.read()
def updateLog(newEntry: String): {UpdateLog} Unit = f.append(newEntry)
```

Fig. 2. A type and a module implementing the logging facility in the text-editor application.

```
resource type File
effect Read
effect Write
effect Append

...
def read(): {this.Read} String
def write(s: String): {this.Write} Unit
def append(s: String): {this.Append} Unit
```

Fig. 3. The type of the file resource.

Consider the code in Figure 2 that shows a type and a module implementing the logging facility of the text-editor application. In the given implementation of the Logger type, the logger module accesses the log file. All modules of type Logger must have two methods: the readLog method that returns the content of the log file and the updateLog method that appends new entries to the log file. In addition, the Logger type declares two *abstract* effects, ReadLog and UpdateLog, that are produced by the corresponding methods. These effects are abstract because they are not given a definition in the Logger type. Thus, it is up to the module implementing the Logger type to define what they mean. The effect names are user defined, allowing the choice of meaningful names.

The logger module implements the Logger type. To access the file system, an object of type File (shown in Figure 3) is passed into logger as a parameter. The logger module's effect declarations are those of the Logger type, except that now they are *concrete*, that is, they have specific definitions. The ReadLog effect of the logger module is defined to be the Read effect of the File object. Accordingly, the readLog method, which produces the ReadLog effect, calls f's read method. Similarly, the UpdateLog effect of the logger module is defined to be f.Append. Accordingly, the updateLog method, which produces the UpdateLog effect, calls f's append method. In general, effects in a module or object definition must always be concrete, whereas effects in a type definition may be either abstract or concrete.

3.1 Path-dependent Effects

In Wyvern programming language, we introduce the mechanism of path-dependent effects. The basic mechanism is similar to the path-dependent types from the Dependent Object Types (DOT) calculus [Amin et al. 2014]. Therefore, in order to understand path-dependent effects, we need to

¹The keyword resource in the type definition indicates that the implementations of this type may have state and may access system resources; this is orthogonal to effect checking.

5:6 D. Melicher et al.

look at path-dependent types first. A path-dependent type describes a type definition that depends on a runtime value of a path to an object. For example, consider the following function that applies a key generator function to elements of a list of strings.

```
def map(g : Generator, s : List[String]): List[g.GeneratedKey] = s.map(g.generate)
```

The value g of Generator type provides a function g.generate that receives a String as input and produces a value of type g.GeneratedKey. The type g.GeneratedKey is dependent on the value g. Thus, the output type of the function map, List[g.GeneratedKey], is also dependent on the input g.

Now let's consider the scenario in which we want to track effects of computations. Assuming that the computation <code>g.generate</code> is effectful, we would want to annotate the function <code>map</code> with an effect label indicating that <code>map</code> has the effect that <code>g.generate</code> causes. One natural solution is to use the path-dependent effect:

```
def map(g : Generator, s : List[String]): {g.Generate} List[g.GeneratedKey] = s.map(g.generate)
```

The function map is annotated with an effect label g.Generate, which is dependent on the value g. This dependency is desirable because the type Generator allows multiple implementations that can potentially cause different types of effects. The label g.Generate will have different meanings if different values are passed into the map function as variable g.

In the Wyvern language, effects are members of objects.² Thus, we refer to them with the form variable.EffectName, where variable is a reference to the object defining the effect and EffectName is the name of the effect. For example, in the definition of the ReadLog effect of the logger module, f is the variable referring to a specific file and Read is the effect that the read method of f produces. This conveniently ties together the resource and the effects produced on it (which represent the operations performed on it), helping a software architect or a security analyst to reason about how resources are used by any particular module and its methods. For example, when analyzing the effects produced by logger's readLog method, a security analyst can quickly deduce that calling that method affects the file resource and, specifically, the file is read simply by looking at the Logger type and logger's effect definitions but not at the method's code. Furthermore, these properties can be automatically checked with an idiom of use: In addition to directly looking at the effect annotation of the method of the logger module, the security analyst may write client code that specifies the effect that the logger module is allowed to have. If the logger module accesses system resources outside of the specified effect set, then the compiler would automatically reject the program.

Because an effect includes a reference to an object instance, our effect system can distinguish reads and writes on different file instances. If the developer does not want this level of precision, it is still possible to declare effects at the module level (i.e., as members of a fileSystem module object instance) and to share the same Read and Write effects, for example, across all files in fileSystem.

The basic mechanisms of path dependence are borrowed from Scala and have been shown to scale well in practice. As in Scala, paths must be "stable" to ensure that assignment cannot change effects. To keep our formal system simple, all paths are simply immutable variables in this article. Our implementation extends this to allow paths involving mutable fields, following recent research [Rapoport and Lhoták 2019]. These mechanisms come from the Dependent Object Types (DOT) calculus [Amin et al. 2014], a type theory of Scala and related languages (including Wyvern). In our system, effects, instead of types, are declared as members of objects.

3.2 Effect Abstraction

The design of our effect system supports a novel form of *effect abstraction*, which is the ability to define higher-level effects in terms of lower-level effects and potentially to hide that definition

 $^{^2}$ Modules are an important special case of objects.

from clients of an abstraction. In the earlier logging example, through the use of abstraction, we "lifted" low-level resources such as the file system (i.e., the Read and Append effects of the file) into higher-level resources such as a logging facility (i.e., the ReadLog and UpdateLog effect of the logger) and enabled application code to reason in terms of effects on those higher-level resources when appropriate.

Effect abstraction has several concrete benefits. First, it can be used to distinguish different uses of a low-level effect. For example, <code>system.FFI</code> describes any access to system resources via calls through the foreign function interface (FFI), but modules that define file and network I/O can represent these calls as different effects, which enables higher-level modules to reason about file and network access separately. Second, multiple low-level effects can be aggregated into a single high-level effect to reduce effect specification overhead. For instance, the <code>db.Query</code> effect might include both <code>file.Read</code> and <code>network.Access</code> effects. Third, by keeping an effect abstract, we can hide its implementation from clients, which facilitates software evolution: code defining a high-level effect in terms of lower-level ones can be rewritten (or replaced) to use a different set of lower-level effects without affecting clients (more on this in Section 4.1).

3.3 Effect Bounds

Our effect system also gives the programmer the ability to define a subtyping hierarchy of effects via effect bounds. To define the hierarchy, the programmer gives the effect member an upper bound or a lower bound, hiding the definition of the effect from the client.

For example, consider the type BoundedLogger, which has the same method declarations and effect members as the type Logger in Figure 2 except that the ReadLog and UpdateLog effects are upper-bounded by the corresponding effects in the fileSystem module:

```
resource type BoundedLogger
effect ReadLog <= {fileSystem.Read}
effect UpdateLog <= {fileSystem.Append}
... // same as in the type Logger in Figure 2</pre>
```

Any object implementing type <code>BoundedLogger</code> may have an effect member <code>ReadLog</code>, which is at most <code>fileSystem.Read</code>. This allows programmers to compare the <code>ReadLog</code> effect with other effects while keeping its definition abstract. For instance, a library can provide two implementations of <code>BoundedLogger</code>, including an effectless logger in which the effects <code>ReadLog</code> and <code>UpdateLog</code> are empty sets and an effectful logger in which <code>ReadLog</code> and <code>UpdateLog</code> are defined as effects in the <code>fileSystem</code> module. The library's clients then can annotate the effects of both implementations with <code>fileSystem.Read</code> and <code>fileSystem.Append</code> according to the effect hierarchy without the need to know the exact implementation of the two instances.

An effect hierarchy can also be constructed using lower bounds. For example, consider the following type for I/O modules that supports writes:

```
type IO
effect Write >= {system.FFI}
def write(s: String): {this.Write} Unit
```

Since I/O is done using the FFI, the Write effect is *at least* the system.FFI effect. Similar to providing an upper bound on effects, this type does not specify the exact definition of the Write effect, and implementations of this type can define Write as an effect set with more effects than {system.FFI}.

The effect hierarchy achieved by bounding effect members is supported by the subtyping relations of our effect system (see Sections 5.5.1 and 5.5.2). If a type has an effect member with more strict bounds than another type, then the former type is a subtype of the latter type. For example,

5:8 D. Melicher et al.

when a logger with the effect member Read <= {fileSystem.Read} is expected, we can pass in a logger with Read = {} because the definition as an empty set is more strict than an upper bound.

The following two case studies demonstrate the expressiveness of the effect hierarchy:

3.4 Effect Aggregation

Wyvern's effect-system design allows for reducing the effect-annotation overhead by aggregating several effects into one. For example, if, to update the log file the logger module needed to first read the file and then write it back, the UpdateLog effect would consist of two effects: a file read and a file write. In other effect systems, this change may make effects more verbose since all of the methods that call the updateLog method would need to be annotated with the two effects. However, effect aggregation allows us to define the UpdateLog effect to be the two effects and then use UpdateLog to annotate the updateLog method and all methods that call it:

```
module def logger(f: File): Logger
effect UpdateLog = {f.Read, f.Write}
def updateLog(newEntry: String): {this.UpdateLog} Unit
...
```

This way, we need to use only one effect, UpdateLog, instead of two in method effect annotations, thus reducing the effect-annotation overhead. Because more code may add more effects, larger software systems might experience a snowballing of effects when method annotations have numerous effects in them.

3.5 Controlling FFI Effects

Wyvern programs access system resources via calls to other programming languages such as Java and Python, that is, through an FFI. To monitor and control the effects caused by FFI calls, we enforce that all functions from other programming languages, when called within Wyvern, are annotated with the system.FFI effect.

As was mentioned in Section 3.2, the system.ffl effect describes function calls though an FFI. Since every function call though FFI has this effect, the access to system resources via FFI is guaranteed to be monitored. system.ffl is the lowest-level effect in the effect system that can be used to build other higher-level effects. The programmer can lift system.ffl to higher-level effects and reason about those higher-level effects instead.

For example, Wyvern's import mechanism works by loading an object in a static field of a Java class, and the following code imports a field of a Java class that helps to implement file I/O:

```
import java:wyvern.stdlib.support.FileIO.file
```

The file object is itself of type FileIO. And FileIO has this method, among others:

```
public void writeStringIntoFile(String content, String filename) throws IOException { ... }
```

In Wyvern, there is a type wyvern.stdlib.support.FileIO as well as an object file (of that type) that gets added to the scope as a result of the import above. The type has the following member, corresponding to the method above:

```
def writeStringIntoFile(content:String, filename:String): { system.FFI } Unit
```

Here, the system.FFI effect was added to the signature because this is a function that was imported via the FFI. The Wyvern file library that uses the writeStringIntoFile function abstracts this system.FFI effect into a library-specific FileIO.Write effect.

4 USE CASES

In this section, we present a selected set of software development patterns that Wyvern's effect system helps facilitate.

ACM Transactions on Programming Languages and Systems, Vol. 44, No. 1, Article 5. Publication date: January 2022.

```
module def remoteLogger(net: Network): Logger
effect ReadLog = {net.Receive}

effect UpdateLog = {net.Send}
def readLog(): {ReadLog} String = net.receive()

def updateLog(newEntry: String): {UpdateLog} Unit = net.send(newEntry)
```

Fig. 4. An alternative implementation of the Logger type from Figure 2.

4.1 Information Hiding and Polymorphism

Introduced by Parnas in the early 1970s [Parnas 1971, 1972], information hiding is a key software development principle stating that, in a software application, implementation details of a particular software module should be hidden behind a stable interface. This principle promotes modularity in the software implementation and gives software developers more flexibility to modify the existing implementation of a module without affecting other modules. Our effect-system design facilitates the principle of information hiding.

For example, Figure 4 shows an alternative implementation of the Logger type from Figure 2. In this version, the log file is stored on some remote machine, and the network (instead of the file system) is used to perform operations on the log. Importantly, the Logger type contains no information about what resource should be used to implement the logging functionality. Thus, a module implementing the Logger type may use any resource or no resources at all (in which case Logger's effects could be defined as empty effects, i.e., {}). Yet the client modules that use a resource of type Logger, such as the two text editor's plugins, observe no difference in the logging functionality. The software architect may swap one logger version for the other at any time without affecting the modules using logger provided that the interface of the Logger type remains the same. Thus, using effect abstraction in the Logger type facilitates the principle of information hiding.

Information hiding is also facilitated by the bounded abstraction feature of our effect system. Consider the following type, which is a subtype of Logger and can be ascribed to the logger module defined in Figure 4:

```
resource type RemoteLogger
effect ReadLog <= {net.Receive}
effect UpdateLog <= {net.Send}
... // same as in the type Logger in Fig. 2</pre>
```

In contrast to Logger, RemoteLogger defines the ReadLog and UpdateLog effects as subeffects of {net.Receive} and {net.Send}, essentially hiding the definitions. Then, if RemoteLogger is used as the functor remoteLogger's return type, remoteLogger's two effect members may be defined using the network resource and remoteLogger's clients can use the net's effects to account for effects of remoteLogger's methods. However, effects in remoteLogger cannot be used to annotate methods that produce the lower-level net effects. Thus, the effect hierarchy allows programmers to annotate methods that have higher-level effects with lower-level effects but not the other way around.

Our design also supports effect polymorphism. For example, the following higher-order function can be used to invoke a function with an arbitrary effect:

```
def invokeTwice[effect E](f: Unit -> {E} Unit): Unit
  f()
  f()
invokeTwice[log.UpdateLog]( () -> log.updateLog("Updating log.") )
```

Here, <code>invokeTwice</code> is parameterized by an effect E. The <code>invokeTwice</code> function takes another function that has no arguments and produces no result but has effect E and <code>invokes</code> that function twice. We call <code>invokeTwice</code>, <code>instantiate</code> the effect parameter with <code>log.UpdateLog</code>, and <code>give invokeTwice</code> a function that updates the log file.

5:10 D. Melicher et al.

```
module def codeCompletion(log: Logger)
def findTemplate(wordSequence: String): {log.UpdateLog} String
log.updateLog("Searching for a matching template.") ...
log.updateLog("Found matching template.") ...

module def userStats(log: Logger)
def calculateUserStats(): {log.ReadLog, log.UpdateLog} String
log.updateLog("Starting to analyze the log content.")
analyzeLogContent(log.readLog()) ...
```

Fig. 5. Excerpts from the code-completion and user-statistics-analyzer plugins of the text-editor application.

Our implementation follows Scala's approach to type polymorphism. Internally, effect polymorphism is rewritten in terms of effect members so that <code>invokeTwice</code> takes an extra argument that has an effect member E.

The compiler rewrites invokeTwice using only effect members. In this rewriting, the invokeTwice function takes an extra parameter, an EffectHolder object, which holds the effect parameter E as an effect member. The desugared code would look like this:

```
type EffectHolder
  effect E

def invokeTwice(eh: EffectHolder, f: Unit -> {eh.E} Unit): Unit
  f()
  f()

let effectHolder: EffectHolder = new
  effect E = log.UpdateLog
in invokeTwice(effectHolder, () -> log.updateLog("Updating log."))
```

Note that this code creates an effectHolder object that instantiates effect E with log.UpdateLog. We also rely on path-dependent types [Amin et al. 2014]: the second parameter of invokeTwice can refer to the first parameter in order to describe the effect of the argument function f.

4.2 Controlling Operations Performed on Modules

Our effect system design allows software developers to control what operations are performed on system resources and other important modules. Consider the two plugins for the text editor. As we noted earlier, these plugins lie outside the trusted code base for the application because they were written by third parties and may contain bugs that could introduce vulnerabilities or could be malicious. To better maintain security of the text-editor application and minimize any potential damage from the plugins, developers of the text editor need to control what resources the plugins access and what operations they perform on those resources. The first part of this task, that is, controlling access to resources, is done via Wyvern's capability-based module system, which limits the plugins' access to resources [Melicher et al. 2017]. The second part of the task, that is, limiting what operations are performed on the resources the plugins have access to, is where Wyvern's effect system can help.

For example, Figure 5 shows some code of the two text editor plugins. Both plugins have access to the logger module, which is passed in as a functor parameter, but they use it differently. Both plugins must report the log.Update effect because they both invoke the method log.updateLog, but only the userStats plugin needs to perform more operations on logger rather than simply updating it. The codeCompletion module needs logger only to update the log file about the status of the search of an appropriate template in its findTemplate method. On the other hand, along with updating the log file, the userStats module reads the log file to analyze its content. Accordingly, codeCompletion's

```
module fileEffects: FileEffects

module fileEffects: FileEffects

effect Read = {system.FFI}

effect Write = {system.FFI}

effect Append = {system.FFI}

...

type FileEffects

effect Read >= {system.FFI}

effect Write >= {system.FFI}

effect Append >= {system.FFI}

...
```

(a) A pure module defining file effects.

(b) A type for the module containing file effects which provides lower bounds for them.

Fig. 6. Defining globally available file effects.

findTemplate method is annotated with the log.UpdateLog effect; therefore, it must call only logger's updateLog method. In contrast, userStats's calculateUserStats method is annotated with both the log.ReadLog and log.UpdateLog effects. Therefore, it may call either updateLog or readLog on the logger.

Wyvern's effect system ensures that the method bodies of findTemplate and calculateUserStats methods produce only the effects with which the methods are annotated (more details on this are in Section 5). Then, the developer can rely on the effect annotations in modules' interfaces to reason about the effects that methods may produce on resources. Thus, our effect-system design allows controlling what operations are performed on resources of an application and significantly simplifies the reasoning process during an analysis of the application security.

4.3 Effect Granularity and Visibility

Our approach offers library designers a choice of granularity of effects: effects that are defined per-object on the one hand and globally defined effects shared by many objects on the other hand. For instance, the example code shown so far has envisioned effects for File objects that are specific to each individual file, allowing fine-grained control of what code accesses what file. Using fine-grained effects, for example, we can verify that the logger accesses a single distinguished log file and no other files. This design has a cost, however—using fine-grained effects can result in verbose effect declarations.

In a different design for the file system libraries, we might define coarse-grained Read, Write, and Append effects in a globally accessible module. For example, Figure 6(a) shows a fileEffects module that defines these effects in terms of Wyvern's low-level foreign function access effect, system.FFI. We hide this concrete definition behind the FileEffects module type defined in Figure 6(b). FileEffects puts a lower bound of system.FFI on each of these effects, which has two purposes. First, the code implementing file reads and writes, which does so using the foreign function interface and therefore incurs the low-level system.FFI effect, can be annotated with higher-level effects, such as fileEffects.Read, since fileEffects.Read is known to subsume system.FFI. Second, file system client code has to assume that, in general, fileEffects.Read might include more than system.FFI, since it cannot see the true definition of Read in the fileEffects module due to the ascribed FileEffects signature, which hides this definition. Thus, client code must treat fileEffects.Read abstractly; it cannot treat it as merely being system.FFI or as any other effect implemented in terms of system.FFI.

Thus, Wyvern's design elegantly supports either local, fine-grained definitions of effects, such as reads on a particular file, or more coarse-grained effects, such as reads to any file in the file system. It is even possible to combine these designs; for example, we could define fileEffects.Read as shown earlier and then type File could declare a Read effect that is specific to that file, but yet is a sub-effect of fileEffects.Read. In this design, a method that takes a file argument f and reads from it could be annotated with a fine-grained f.Read effect while a caller of that method that accesses

5:12 D. Melicher et al.

multiple files could declare its effects with the more coarse-grained fileEffects.Read to keep its declaration succinct.

4.4 Authority Attenuation

A key principle to ensure security of a software system is the principle of least privilege [Denning 1976], which states that a software module must have privilege necessary only to implement its designated functionality and nothing else. In practice, this principle translates into protecting software modules representing system resources, such as the file system and network, and other important modules, such as those holding user data, from excessive access and abuse by other software modules. An important component of privilege is operations performed on a resource being accessed. In the field of software security, such operations represent *authority* over the accessed module [Miller 2006].³

Notably, Wyvern effects that describe operations performed on modules are a good medium for representing authority over modules. For example, the fact that the logger module's effects use only file's Read and Append effects in logger's effect definitions signifies that the only operations logger performs on the log file are the read and append operations, meaning that the only authority logger has over the log file is to read it and append to it.

Furthermore, our effect-system design allows expressing the notion of *authority attenuation*, which is a common software-security pattern [Murray 2008]. Authority attenuation happens when the original set of operations that can be performed on a resource is limited by an intermediary object [Miller 2006]. For example, consider the sequence of module dependencies from Figure 1 consisting of the file module, the logger module, and the codeCompletion module. There are several operations that can be performed on a file (at least the three shown in the File type in Figure 3), but logger performs only two of them (as was mentioned earlier and as can be seen from its effects' definitions in Figure 2). The codeCompletion module can access the logger module but not the file module. Thus, the only operations it can perform on file are those that logger can perform. Thus, the logger module attenuates codeCompletion's authority over the file module.

Therefore, our effect-system design helps developers in observing and establishing the authority-attenuating relationship between modules of a software application, which may be desired and beneficial during the design phase of a software application, a security audit, or an architecture review of a software application.

5 FORMALIZATION

As was mentioned earlier, Wyvern modules are first class and are, in fact, objects since they are only syntactic sugar on top of Wyvern's object-oriented core and can be translated into objects. The translation has been described in detail previously [Melicher et al. 2017]. Here, we provide only some intuition behind it. In this section, we start with describing the syntax of Wyvern's object-oriented core and then present an example of the module-to-object translation, followed by a description of Wyvern's static semantics and subtyping rules. Furthermore, we present the dynamic semantics and the type soundness theorems. Last, but not least, we provide the definitions on authority and discuss why they are useful for security analysis on programs written in Wyvern.

5.1 Object-Oriented Core Syntax

Figure 7 shows the syntax of Wyvern's object-oriented core. Wyvern expressions e include variables and the four basic object-oriented expressions: the new statement new $(x \Rightarrow \overline{d})$, a method

³Similar to the work by Maffeis et al. [2010], we widened the original definition of authority to be about being able to perform any operation on a module instead of being able to only modify it.

```
d ::= def m(x : \tau) : \{\varepsilon\} \tau = e
                                                                                                                        def m(x : \tau) : \{\varepsilon\} \tau
                                                                                                             ::=
                                                                                                         \sigma
::= x
                                               1
                                                        \operatorname{var} f : \tau = x
                                                                                                                1
                                                                                                                         var f : \tau
  1
          new(x \Rightarrow d)
                                                1
                                                        \mathsf{effect}\: g = \{\varepsilon\}
                                                                                                                         effect g
  ı
          e.m(e)
                                               ::=
                                                                                                                         effect g \ge \{\varepsilon\}
  1
          e.f
                                                        \{x \Rightarrow \overline{\sigma}\}
                                               ::=
                                                                                                                         effect q \leq \{\varepsilon\}
  1
          e.f = e
                                                       \emptyset \mid \Gamma, \ x : \tau
                                                                                                                         effect q = \{\varepsilon\}
```

Fig. 7. Wyvern's object-oriented core syntax.

```
let logger = new (x ⇒

def apply(f: File) : {} Logger

new (y ⇒

effect ReadLog = {f.Read}

effect UpdateLog = {f.Append}

def readLog(): {y.ReadLog} String = f.read()

def updateLog(newEntry: String): {y.UpdateLog} Unit = f.append(newEntry)))

in ... // calls logger.apply(...)
```

Fig. 8. A simplified translation of the logger module from Figure 2 into Wyvern's object-oriented core.

call e.m(e), a field access e.f, and a field assignment e.f = e. Objects are created by new statements that contain a variable x representing the current object along with a list of declarations. In our implementation, x defaults to this when no name is specified by the programmer. Declarations d come in three kinds: a method declaration $d \in m(x : \tau) : \{\varepsilon\} \ \tau = e$, a field $d \in \tau = x$, and an effect member effect $d \in \tau = x$. Method declarations are annotated with a set of effects $d \in x$. Object fields $d \in x$ are $d \in x$ and $d \in x$ are $d \in x$ and $d \in x$ are striction that simplifies our core language by ensuring that object initialization never has an effect. Although at first this may seem to be limiting, in fact, we do not limit the source language in this way. Side-effecting member initializations in the source language are translated to the core by wrapping the new object with a let expression that defines the variable to be used in the field initialization. For example, this code:

```
new
    var x: String = f.read()
can be internally rewritten as:
    let y = f.read()
    in new
    var x: String = y
```

Effects in method annotations and effect-member definitions effect $g = \{\epsilon\}$ are surrounded by curly braces to visually indicate that they are sets, and each effect in an effect set is defined to be a variable representing the object on which an effect is produced, followed by a dot and the effect name. For example, $\{file.Read, network.Access\}$ is an effect set that contains two effect labels from file and network modules. Abstract effects may be defined with an upper bound or a lower bound.

The type $\{x \Rightarrow \overline{\sigma}\}$ is the only possible form of object types τ . The variables $\overline{\sigma}$ in object types are a collection of declaration types, which include method signatures $\operatorname{def} \ m(x:\tau) : \{\varepsilon\} \ \tau$, field-declaration types var $f:\tau$, and the types of effect-member declarations and definitions. Similar to the difference between the modules and their types, effects in an object must always be defined (i.e., always be concrete), whereas effects in object types may or may not have definitions (i.e., be either abstract or concrete) and may have an upper or lower bound.

5.2 Modules-to-Objects Translation

Figure 8 presents a simplified translation of the logger module from Figure 2 into Wyvern's object-oriented core (for a full description of the translation mechanism, refer to Melicher et al. [2017]).

5:14 D. Melicher et al.

For our purposes, the functor becomes a regular method, called apply, that has the return type Logger and the same parameters as the module functor. The method's body is a new object containing all of the module declarations. The apply method is the only method of an outer object that is assigned to a variable whose name is the module's name. Later on in the code, when the logger module needs to be instantiated, the apply method is called with appropriate arguments passed in.

To aid this translation mechanism, we use the two relatively standard encodings:

```
let x = e in e' \equiv \text{new}(\_ \Rightarrow \text{def } f(x : \tau) : \tau' = e').f(e)
def m(\overline{x : \tau}) : \tau = e \equiv \text{def } m(x : (\tau_1 \times \tau_2 \times ... \times \tau_n)) : \tau = [x.n/x_n]e
```

The let expression is encoded as a method call on an object that contains that method, with the let variable being the method's parameter and the method body being the let's body. The multiparameter version of the method definition is encoded using indexing into the method parameters.

5.3 Well-formedness

Since Wyvern's effects are defined in terms of variables, before we type check expressions, we must make sure that effects and types are well formed. Wyvern well-formedness rules are mostly straightforward and are shown in Figure 9. The three judgments read that in the variable typing context Γ , the type τ , the declaration type σ , and the effect set ε are well formed, respectively.

An object type is well formed if all of its declaration types are well formed (WF-Type). A method-declaration type is well formed if the type of its parameter, its return type, and the effects in its effect annotation are well formed. A field-declaration type is well formed if its type is well formed. Since an effect-declaration type has no right-hand side, it is trivially well formed (WF-Def). An effect-definition type is well formed if the effect set in its right-hand side is well formed (WF-Effect2). Finally, a bounded effect declaration is well formed if the upper bound or lower bound on the right-hand side is well formed (WF-Effect3, WF-Effect4). An effect set is well formed if, for every effect it contains, the definition of the effect does not form a cycle, the variable in the first part of the effect is well typed, and the type of that variable contains either an effect-declaration or an effect-definition type, in which the effect name matches the effect name in the second part of the effect (WF-Effect). The typing judgments used by (WF-Effect) are defined in Section 5.4.

The $\Gamma \vdash safe(x.g, \varepsilon)$ judgment ensures that the definition of effect x.g does not contain a cycle. The effect set ε memorizes a set of effects that are defined by x.g. The judgment ensures that those effects do not appear in the definition of x.g, ensuring that there are no cycles in effect definitions.

In the rule (Safe-1), we want to ensure that the definition of the effect x.g does not contain a cycle. We first get the definition of x.g, which is ε' , from the typing context. Then, we look at the effects in the set $\{x.g\} \cup \varepsilon$ and check to see whether any of them appear in ε' , which makes sure that the definition of x.g does not introduce cycles. Then, by the judgment $safe(c.d, \{x.g\} \cup \varepsilon)$, we recursively check whether each effect in ε' could introduce a cycle. Rules Safe-2 and Safe-3 are identical except that the effect is bounded rather than fully specified. Rule Safe-4 expresses that a fully abstract type member cannot introduce a cycle.

5.4 Static Semantics

Wyvern's static semantics is presented in Figure 10. Expression type checking includes checking the effects that an expression may have, the set of which is denoted in a pair of curly braces between the colon and the type in the type annotation. Then, for expressions, the judgment reads that, in the variable typing context Γ , the expression e is a well-typed expression with the effect set ε and the type τ .

$$\frac{\forall \sigma \in \overline{\sigma}, \, \Gamma, \, x : \{x \Rightarrow \overline{\sigma}\} + \sigma \, wf}{\Gamma \vdash \{x \Rightarrow \overline{\sigma}\} \, wf} \quad \text{(WF-Type)}$$

$$\frac{\Gamma \vdash \sigma \, wf}{\Gamma \vdash \text{def} \, m(x : \tau_2) : \{\varepsilon\} \, \tau_1 \, wf} \quad \Gamma, \, x : \tau_2 \vdash \varepsilon \, wf}{\Gamma \vdash \text{def} \, m(x : \tau_2) : \{\varepsilon\} \, \tau_1 \, wf} \quad \text{(WF-Def)} \quad \frac{\Gamma \vdash \tau \, wf}{\Gamma \vdash \text{var} \, f : \tau \, wf} \quad \text{(WF-EFFECT2)}$$

$$\frac{\Gamma \vdash \varepsilon \, wf}{\Gamma \vdash \text{effect} \, g \, wf} \quad \text{(WF-EFFECT3)} \quad \frac{\Gamma \vdash \varepsilon \, wf}{\Gamma \vdash \text{effect} \, g \, g \, \{\varepsilon\} \, wf} \quad \text{(WF-EFFECT4)}$$

$$\frac{\Gamma \vdash \varepsilon \, wf}{\Gamma \vdash \text{effect} \, g \, g \, \{\varepsilon\} \, wf} \quad \text{(WF-EFFECT3)} \quad \frac{\Gamma \vdash \varepsilon \, wf}{\Gamma \vdash \text{effect} \, g \, g \, \{\varepsilon\} \, wf} \quad \text{(WF-EFFECT4)}$$

$$\frac{\forall i, j, \, x_i, g_j \in \varepsilon,}{\Gamma \vdash \text{effect} \, g \, g \, \{\varepsilon\} \, wf} \quad \text{(WF-EFFECT4)}$$

$$\frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \text{effect} \, g \, g \, \{\varepsilon\} \, wf} \quad \text{(WF-EFFECT4)}$$

$$\frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \text{effect} \, g \, g \, \{\varepsilon\} \, wf} \quad \text{(WF-EFFECT4)}$$

$$\frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \text{effect} \, g \, g \, \{\varepsilon\} \, wf} \quad \text{(WF-EFFECT4)}$$

$$\frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \text{effect} \, g \, g \, \{\varepsilon\} \, wf} \quad \text{(WF-EFFECT5)}$$

$$\frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \text{effect} \, g \, g \, \{\varepsilon\} \, wf} \quad \text{(WF-EFFECT5)}$$

$$\frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \text{(WF-EFFECT5)}$$

$$\frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \text{(WF-EFFECT5)}$$

$$\frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \text{(WF-EFFECT5)}$$

$$\frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \text{(WF-EFFECT5)}$$

$$\frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \text{(WF-EFFECT5)}$$

$$\frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \text{(WF-EFFECT5)}$$

$$\frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \text{(WF-EFFECT5)}$$

$$\frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \, wf} \quad \frac{\nabla \vdash \varepsilon \, wf}{\Gamma \vdash \varepsilon \,$$

Fig. 9. Wyvern well-formedness rules.

A variable trivially has no effects. A new expression also has no effects because of the fact that fields may be initialized using only variables. A new object is well typed if all of its declarations are well typed (T-Var).

A method call is well typed if the expression passed into the method as an argument is well typed, if the expression the method is called on is well typed, and if the expression's type contains a matching method-declaration type (T-Method). In addition, bearing the appropriate variable substitutions, the effect set annotating the method-declaration type must be well formed and the effect set ε in the method-call type must be a union of the effect sets of both expressions involved in the method call as well as the effect set of the method-declaration type.

5:16 D. Melicher et al.

Fig. 10. Wyvern static semantics.

Substitutions are carried out in this rule and others because the caller has a different point of view from the callee. For example, the argument type τ_2 may mention effect members of the receiver x, but the caller does not know about x; instead, the caller's name for the receiver is the expression e_1 . Thus, we substitute e_1 for x in τ_2 when checking whether e_2 has the right type. Similar substitutions are used elsewhere in the rule.

Definition 1 (Substitution). We write $[e/x]\tau$ to denote the substitution of e for the free occurrence of x in τ . Similarly, we use $[e/x]\varepsilon$ and $[e/x]\sigma$ to denote substitution of e for free variable x in effects and declaration types, respectively.

```
(1) [y/x]{z \Rightarrow \overline{\sigma}} = {z \Rightarrow \overline{[y/x]\sigma}}
```

- (2) $[y/x]\{x \Rightarrow \overline{\sigma}\} = \{x \Rightarrow \overline{\sigma}\}\$
- (3) $[y/x] \text{def } m(z:\tau_1): \{\varepsilon\} \ \tau_2 = \text{def } m(z:[y/x]\tau_1): \{[y/x]\varepsilon\} \ [y/x]\tau_2$
- (4) [y/x]def $m(x : \tau_1) : \{\varepsilon\} \ \tau_2 = \text{def } m(x : \tau_1) : \{\varepsilon\} \ \tau_2$
- (5) [y/x]var $f: \tau = \text{var } f: [y/x]\tau$
- (6) [y/x] effect q = effect q
- (7) [y/x] effect $g \le \{\varepsilon\}$ = effect $g \le \{[y/x]\varepsilon\}$
- (8) [y/x] effect $g \ge \{\varepsilon\}$ = effect $g \ge \{[y/x]\varepsilon\}$
- (9) [y/x] effect $q = \{\varepsilon\}$ = effect $q = \{[y/x]\varepsilon\}$
- (10) [y/x]{} = {} (Empty effect set)
- (11) $[y/x]\{x.g\} \cup \varepsilon = \{y.g\} \cup [y/x]\varepsilon$
- (12) $[y/x]\{z.q\} \cup \varepsilon = \{z.q\} \cup [y/x]\varepsilon$

Note that if e_1 is not a variable, we define the substitution $[e_1/x]\tau_2$ to have no effect, resulting in just τ_2 . Otherwise, the substitution could introduce invalid effects, such as x.f(y).E. In our implementation, the programmer is responsible for using let-binding variables where appropriate to avoid this problem. Any mistakes made simply result in a typing error.

ACM Transactions on Programming Languages and Systems, Vol. 44, No. 1, Article 5. Publication date: January 2022.

$$\frac{\varepsilon_1 \subseteq \varepsilon_2}{\Gamma + \varepsilon_1 <: \varepsilon_2} \text{ (Subeffect-Subset)}$$

$$\frac{\Gamma \vdash n : \{y \Rightarrow \overline{\sigma}\} \text{ effect } g \leqslant \{\varepsilon\} \in \sigma \quad \Gamma \vdash [n/y] \varepsilon \cup \varepsilon_1 <: \varepsilon_2}{\Gamma \vdash \varepsilon_1 \cup \{n.g\} <: \varepsilon_2} \text{ (Subeffect-Upperbound)}$$

$$\frac{\Gamma \vdash n : \{y \Rightarrow \overline{\sigma}\} \text{ effect } g \geqslant \{\varepsilon\} \in \sigma \quad \Gamma \vdash \varepsilon_1 <: [n/y] \varepsilon \cup \varepsilon_2}{\Gamma \vdash \varepsilon_1 <: \varepsilon_2 \cup \{n.g\}} \text{ (Subeffect-Lowerbound)}$$

$$\frac{\Gamma \vdash n : \{y \Rightarrow \overline{\sigma}\} \text{ effect } g = \{\varepsilon\} \in \sigma \quad \Gamma \vdash \varepsilon_1 <: [n/y] \varepsilon \cup \varepsilon_2}{\Gamma \vdash \varepsilon_1 <: \varepsilon_2 \cup \{n.g\}} \text{ (Subeffect-Def-1)}$$

$$\frac{\Gamma \vdash n : \{y \Rightarrow \overline{\sigma}\} \text{ effect } g = \{\varepsilon\} \in \sigma \quad \Gamma \vdash [n/y] \varepsilon \cup \varepsilon_1 <: \varepsilon_2}{\Gamma \vdash \varepsilon_1 \cup \{n.g\} <: \varepsilon_2} \text{ (Subeffect-Def-2)}$$

Fig. 11. Wyvern subeffecting rules.

An object field read is well typed if the expression on which the field is dereferenced is well typed and the expression's type contains a matching field-declaration type (T-Field). The effects of an object field type are those of the expression on which the field dereferencing is called.

A field assignment is well typed if the expression to which the field belongs is well typed and the expression's type has an appropriate field-declaration type, and if the expression in the right-hand side of the assignment is well typed (T-Assign). The effect set that a field assignment produces is the union of the effect sets produced by the two subexpressions.

Subsumption may happen only if the expression is well typed using the original type, the original type is a subtype of the new type, and when the effect set of the original set is a subeffect (discussed in Section 5.5.1) of the effect of the new type (T-Sub).

None of the object declarations produces effects. Thus, object-declaration type-checking rules do not include an effect set preceding the type annotation. For declarations, the judgment reads that, in the variable typing context Γ , the declaration d is a well-typed declaration with the declaration type σ .

When type-checking a method declaration, the effect set annotating the method must be well formed in the overall typing context extended with the method argument (DT-Def). Furthermore, the effect annotating the method must be a supereffect of the effect that the method body actually produced.

A field declaration is trivially well typed, and an effect declaration is well typed if the effect set that it is defined with is well formed in the given context.

5.5 Subtyping

5.5.1 Subeffecting Rules. As we already saw in the T-Sub and DT-Def rules earlier, and as we will see more in Section 5.5.2, to compare two sets of effects we use subeffecting rules, which are presented in Figure 11. If an effect is a subset of another effect, then the former effect is a subeffect of the latter (Subeffect-Subset). If an effect set contains an effect variable that is declared with an upper bound and the union of the rest of the effect set with the upper bound is a subeffect of another effect set, then the former effect set is a subeffect of the latter effect set (Subeffect-Lowerbound). If an effect set with the lower bound is a supereffect of another effect set, then the former effect set is a supereffect of the latter (Subeffect-Lowerbound). If an effect set contains an effect variable that has a definition, and the union of the rest of the effect set with the definition of the

5:18 D. Melicher et al.

 $\overline{size(\Gamma, \{\})} = 0 \quad \text{(SIZE-EMPTY)}$ $\frac{\Gamma \vdash x : \{\}\{y \Rightarrow \sigma\} \quad \text{effect } g \in \sigma}{size(\Gamma, x.g)} \quad \text{(SIZE-ABSTRACT)}$ $\overline{size(\Gamma, \overline{x.g})} = \sum_{x.g \in \overline{x.g}} size(\Gamma, x.g) \quad \text{(SIZE-LIST)}$ $\frac{\Gamma \vdash x : \{\}\{y \Rightarrow \sigma\} \quad \text{effect } g = \{\varepsilon\} \in \sigma}{size(\Gamma, x.g)} \quad \text{(SIZE-DEF)}$ $\frac{\Gamma \vdash x : \{\}\{y \Rightarrow \sigma\} \quad \text{effect } g \leqslant \{\varepsilon\} \in \sigma}{size(\Gamma, x.g)} \quad \text{(SIZE-DEF)}$ $\frac{\Gamma \vdash x : \{\}\{y \Rightarrow \sigma\} \quad \text{effect } g \leqslant \{\varepsilon\} \in \sigma}{size(\Gamma, x.g)} \quad \text{(SIZE-UPPERBOUND)}$ $\frac{\Gamma \vdash x : \{\}\{y \Rightarrow \sigma\} \quad \text{effect } g \geqslant \{\varepsilon\} \in \sigma}{size(\Gamma, x.g)} \quad \text{(SIZE-LOWERBOUND)}$

Fig. 12. Rules for determining the size of effect definitions.

variable is a supereffect of another effect set, then the former effect set is a supereffect of the latter (Subeffect-Def-1). Finally, if an effect set contains an effect variable that has a definition and the union of the rest of the effect set with the definition of the variable is a subeffect of another effect set, then the former effect set is a subeffect of the latter (Subeffect-Def-2).

We introduce the definition size as a tool for the later induction proofs. The value $size(\Gamma, \varepsilon)$ describes the depth of the effect definition tree. For example, if $x.g = \{y_1.g_1, y_2.g_2, y_3.g_3\}$ and effects $y_1.g_1$, $y_2.g_2$, $y_3.g_3$ are abstract effects, then there is only one level of effect definition in x.g. Thus, $size(\Gamma, x.g) = 1$. On the other hand, if $x.g = \{y.h\}$ and $y.h = \{z.i\}$, where $\{z.i\}$ is abstract in context, then $size(\Gamma, x.g) = 2$.

LEMMA 1. $size(\Gamma, \varepsilon)$ (Defined in Figure 12) is finite.

PROOF. By rules Safe-1, Safe-2, Safe-3, and Safe-4 in Figure 9, the size of an arbitrary effect x.g is bounded by the total number of effects in the context Γ .

Theorem 1. $\Gamma \vdash \varepsilon <: \varepsilon'$ is decidable.

PROOF. The proof is by induction on $size(\Gamma, \varepsilon \cup \varepsilon')$.

Base case: Since both effects are of size 0, the only applicable rule for subeffecting is Subeffect-Subset. The rule only checks if ε is a subset of ε' . Therefore, the relation is decidable.

Induction Step: Assume that the judgment $\Gamma \vdash \varepsilon <: \varepsilon'$ is derived from Subeffect-Upperbound. In the premise of this rule, we have that $\Gamma \vdash [n/y]\varepsilon \cup \varepsilon_1 <: \varepsilon_2$. Since we extract the definition of n.g to find ε , we have $size(\Gamma, [n/y]\varepsilon \cup \varepsilon_1 \cup \varepsilon_2) < size(\Gamma, \{n.g\} \cup \varepsilon_1 \cup \varepsilon_2)$. We can then use an induction hypothesis to show that the subeffecting judgment in the premise is decidable.

The inductive steps for rules Subeffect-Lowerbound, Subeffect-Def-1, and Subeffect-Def-2 have a similar structure.

5.5.2 Declarative Subtyping Rules. Wyvern subtyping rules are shown in Figure 13. Since to compare types we need to compare the effects in them using subeffecting, the subtyping relationship is checked in a particular variable typing context. The first four object-subtyping rules and the S-Refl2 rule are standard. In S-Depth, since effects may contain a reference to the current object,

$$\frac{\Gamma \vdash \tau <: \tau'}{\Gamma \vdash \tau <: \tau} \text{ (S-Refl1)} \qquad \frac{\Gamma \vdash \tau_1 <: \tau_2 \quad \Gamma \vdash \tau_2 <: \tau_3}{\Gamma \vdash \tau_1 <: \tau_3} \text{ (S-Trans)}$$

$$\frac{\{x \Rightarrow \sigma_i^{i \in 1...n}\} \text{ is a permutation of } \{x \Rightarrow \sigma_i^{i \in 1..n}\}}{\Gamma \vdash \{x \Rightarrow \sigma_i^{i \in 1..n}\}} \text{ (S-Perm)}$$

$$\frac{\forall i, \ \Gamma, \ x : \{x \Rightarrow \sigma_i^{i \in 1..n}\} \vdash \sigma_i <: \sigma_i'}{\Gamma \vdash \{x \Rightarrow \sigma_i^{i \in 1..n}\}} \text{ (S-Depth)}$$

$$\frac{\forall i, \ \Gamma, \ x : \{x \Rightarrow \sigma_i^{i \in 1..n}\} \vdash \sigma_i <: \sigma_i'}{\Gamma \vdash \{x \Rightarrow \sigma_i^{i \in 1..n}\}} \text{ (S-Depth)}$$

$$\frac{\Gamma \vdash \{x \Rightarrow \sigma_i^{i \in 1..n}\} \land (S-Width)}{\Gamma \vdash \{x \Rightarrow \sigma_i^{i \in 1..n}\}} \text{ (S-Width)}$$

$$\frac{\Gamma \vdash \sigma <: \sigma'}{\Gamma \vdash \sigma <: \sigma'}$$

$$\frac{\Gamma \vdash \tau_1' <: \tau_1 \quad \Gamma \vdash \tau_2 <: \tau_2' \quad \Gamma, x : \tau_1 \vdash \varepsilon_1 <: \varepsilon_2}{\Gamma \vdash \text{def } m(x : \tau_1) : \{\varepsilon_1\} \quad \tau_2 <: \text{def } m(x : \tau_1') : \{\varepsilon_2\} \quad \tau_2'} \text{ (S-Def)}$$

$$\frac{\Gamma \vdash \varepsilon <: \varepsilon'}{\Gamma \vdash \text{effect } g = \{\varepsilon\} <: \text{effect } g \leqslant \varepsilon'} \text{ (S-Effect-3)}$$

$$\frac{\Gamma \vdash \varepsilon <: \varepsilon'}{\Gamma \vdash \text{effect } g \leqslant \varepsilon <: \text{effect } g \leqslant \varepsilon'} \text{ (S-Effect-4)}$$

$$\frac{\Gamma \vdash \varepsilon' <: \varepsilon}{\Gamma \vdash \text{effect } g \leqslant \varepsilon <: \text{effect } g \leqslant \varepsilon'} \text{ (S-Effect-6)}$$

$$\frac{\Gamma \vdash \varepsilon' <: \varepsilon}{\Gamma \vdash \text{effect } g \leqslant \varepsilon'} \text{ (S-Effect-7)}$$

Fig. 13. Wyvern subtyping rules.

$$\frac{ \left[\Gamma \vdash \tau <: \tau' \right] }{ \exists \text{ an injection } p: \{1...n\} \mapsto \{1...m\}, \quad \forall i \in 1...n, \ \Gamma, \ x: \{x \Rightarrow \sigma_i^{i \in 1..m}\} \vdash \sigma_{p(i)} <: \sigma_i' }{ \Gamma \vdash \{x \Rightarrow \sigma_i^{i \in 1...m}\} <: \Gamma \vdash \{x \Rightarrow \sigma_i^{i \in 1...m}\} }$$
 (S-ALG)

Fig. 14. Algorithmic subtyping.

to check the subtyping relationship between two type declarations, we extend the current typing context with the current object. Method-declaration typing is contravariant in the argument types and covariant in the return type. Furthermore, there must be a covariant-like relationship between the effect sets in the method annotations on the two method declarations: the effect set of the subtype method declaration must be a subeffect of the effect set of the supertype method declaration (S-Def). An effect definition or an effect declaration with a bound is trivially a subtype of an effect declaration (S-Effect-1, S-Effect-2, S-Effect-5). An effect definition is a subtype of an effect declaration with the upper bound if the definition is a subeffect of the upper bound (S-Effect-3). Similarly, an effect definition is a subtype of an effect declaration with the lower bound is a subtype of the effect declaration with another upper bound if the former upper bound is a subtype of the effect declaration with another lower bound if the former upper bound is a subtype of the effect declaration with another lower bound if the former upper bound is a subtype of the latter upper bound (S-Effect-7).

5.5.3 Algorithmic Subtyping Rules. The S-Alg rule encodes the S-Refl1, S-Perm, S-Depth, and S-Width rule using an injective function *p*. Different forms of injective function *p* make S-Alg encode

5:20 D. Melicher et al.

```
\operatorname{def} m(x : \tau) : \{\varepsilon\} \tau
                                                                                                                                                 declaration types
                                                                                         \sigma
                                                                                                         var f : \tau
                                                                                                         effect q
n
      ::=
               x \mid l
                                                                  names
                                                                                                         effect g \ge \{\varepsilon\}
      ::=
               n
                                                            expressions
               new(x \Rightarrow d)
                                                                                                         effect g \leq \{\varepsilon\}
       ı
       1
               e.m(e)
                                                                                                         effect q = \{\varepsilon\}
                                                                                         Γ
                                                                                                         \emptyset \mid \Gamma, x : \tau
                                                                                                                                              var. typing context
       1
               e.f
                                                                                                         \emptyset \mid \mu, l \mapsto \{x \Rightarrow d\}
       ı
               e.f = e
                                                                                         μ
                                                                                                ::=
      ::=
              \overline{n.g}
                                                                   effects
                                                                                         Σ
                                                                                                ::=
                                                                                                         \emptyset \mid \Sigma, l : \tau
                                                                                                                                             store typing context
ε
                                                                                         Е
      ::=
               \mathsf{def}\ m(x:\tau):\{\varepsilon\}\ \tau=e
                                                           declarations
                                                                                                         []
                                                                                                                                               evaluation context
                                                                                                         E.m(e)
       var f : \tau = n
       1
               effect g = \{\varepsilon\}
                                                                                                         l.m(E)
              \{x \Rightarrow \overline{\sigma}\}
                                                                                                         E.f
      ::=
                                                            object type
                                                                                                         E.f = e
                                                                                                         l.f = E
```

Fig. 15. Wyvern's object-oriented core syntax with dynamic forms.

one of the three original rules. For example, if m=n, then $\overline{\sigma_{p(i)}}$ is a permutation of $\overline{\sigma_i}$, which makes the rule S-Alg equivalent to S-Perm. And if m>n, then S-Alg subsumes S-Width using correct function p. Moreover, S-Alg will encode S-Depth if m=n and $\forall i\in 1\dots n, p(i)=i$. The subtyping rules of declaration types are identical to the declarative subtyping. We prove that S-Trans rules are admissible in Theorem 2, which is proved in Appendix A. Therefore, we have shown that there is a set of algorithmic subtyping rules for object types and declaration types that is complete and syntax- directed, which means that subtyping is decidable.

```
Theorem 2. (Transitivity of Algorithmic Subtyping). If \Gamma \vdash \tau_1 <: \tau_2 \text{ and } \Gamma \vdash \tau_2 <: \tau_3, \text{ then } \Gamma \vdash \tau_1 <: \tau_3. If \Gamma \vdash \sigma_1 <: \sigma_2 \text{ and } \Gamma \vdash \sigma_2 <: \sigma_3, \text{ then } \Gamma \vdash \sigma_1 <: \sigma_3.
```

5.6 Dynamic Semantics and Type Soundness

- 5.6.1 Object-Oriented Core Syntax. We present the dynamic semantics of our calculus, which uses a small-step style of operational semantics with evaluation contexts. Figure 15 shows the version of the syntax of Wyvern's object-oriented core that includes dynamic semantics. The definitions are a strict extension of those in Figure 7. Expressions now include locations l, which are the result of object allocation. All runtime values in the language are locations. Thus, the name-dependent part of an effect can now be either a variable or the location that the variable is substituted with at runtime. To support locations, we also define a store μ and its typing context Σ . Finally, to make the dynamics more compact, we use an evaluation context E.
- 5.6.2 Changes in Static Semantics. Type checking a location (T-Loc) and a field declaration (DT-Var) is straightforward; we also need to ensure that the store is well formed and contains objects that respect their types. In (T-Store), we check that for all locations l in the store, declarations d_i are well typed in a context that contains self-variable x. The well-formedness rules and the definition of substitution stay mostly unchanged except that we can use location l as a name and substitute a location for a variable in substitutions.
- 5.6.3 Dynamic Semantics. The dynamic semantics that we use for Wyvern's effect system is shown in Figure 17 and is similar to the one described in prior work [Melicher et al. 2017]. In comparison with the prior work, this version of Wyvern's dynamic semantics has fewer rules and the E-Method rule is simplified.

$$\begin{array}{c} \Gamma \mid \Sigma \vdash e : \{\varepsilon\} \; \tau \\ \\ \dots \\ \hline \Gamma \mid \Sigma \vdash l : \{\} \; \tau \end{array} \; \text{(T-Loc)} \\ \\ \Gamma \mid \Sigma \vdash d : \sigma \\ \\ \dots \\ \hline \Gamma \mid \Sigma \vdash n : \{\} \; \tau \\ \\ \dots \\ \hline \Gamma \mid \Sigma \vdash \text{var} \; f : \tau = n \; : \; \text{var} \; f : \tau \end{array} \; \text{(DT-Var)} \\ \\ \mu \colon \Sigma \\ \\ \forall l \mapsto \{x \Rightarrow \overline{d}\} \in \mu, \; l : \{x \Rightarrow \overline{\sigma}\} \in \Sigma \\ \forall i, \; d_i \in \overline{d}, \; \sigma_i \in \overline{\sigma}, \\ x \colon \{x \Rightarrow \overline{\sigma}\} \mid \Sigma \vdash d_i : \sigma_i \\ \hline \mu \colon \Sigma \end{array} \; \text{(T-Store)}$$

Fig. 16. Wyvern static semantics affected by dynamic semantics.

Fig. 17. Wyvern's dynamic semantics.

The judgment $\langle e \mid \mu \rangle \rightarrow \langle e' \mid \mu' \rangle$ can be read as follows: given the store μ , the expression e evaluates to the expression e' and the store becomes μ' . The E-Congruence rule handles all nonterminal forms. To create a new object (E-New), we select a fresh location in the store and assign the object's definition to it. Provided that there is an appropriate method definition in the object on which a method is called, the method call is reduced to the method's body (E-Method). In the method's body, the locations representing the method argument and the object on which the method is called are substituted for corresponding variables. An object field is reduced to the value held in it (E-Field). When an object field's value changes (E-Assign), appropriate substitutions are made in the object's declaration set and the store.

5.6.4 Type Soundness. We prove the soundness of the effect system presented earlier using the standard combination of progress and preservation theorems. Proof of these theorems can be found in Appendix B. Because of the "bad bounds" problem, the soundness theorem for DOT-like calculi has historically been complex and its proofs tend to be difficult [Amin et al. 2014]. The soundness proof of our system is surprisingly simple due to several properties of our system that are different from regular DOT calculi. First, our system does not contain intersection types, a feature that

5:22 D. Melicher et al.

interacts poorly with the narrowing property. Moreover, our effect members cannot have both an upper bound and a lower bound, which fundamentally avoids the "bad bounds" situation in which the lower bound is not a subtype of the upper bound. Finally, the use of effect members is very restricted when compared with type members in DOT calculi because effect members can appear only in effect annotations or effect bounds. Thus, the type structure of the application does not depend on effect members—only effects, which are at the "leaves" of the type and are dependent on effect members.

```
Theorem 3 (Preservation). If \Gamma \mid \Sigma \vdash e : \{\varepsilon\} \ \tau, \mu : \Sigma, and \langle e \mid \mu \rangle \longrightarrow \langle e' \mid \mu' \rangle, then \exists \Sigma' \supseteq \Sigma, \mu' : \Sigma', \exists \varepsilon', such that \Gamma \vdash \varepsilon' <: \varepsilon, and \Gamma \mid \Sigma' \vdash e' : \{\varepsilon'\} \ \tau.
```

Theorem 4 (Progress). If $\emptyset \mid \Sigma \vdash e : \{\varepsilon\} \ \tau$ (i.e., e is a closed, well-typed expression), then either

- (1) e is a value (i.e., a location) or
- (2) $\forall \mu$ such that $\mu : \Sigma, \exists e', \mu'$ such that $\langle e \mid \mu \rangle \longrightarrow \langle e' \mid \mu' \rangle$.

5.7 Authority

Our definition of authority is based on prior research [Drossopoulou et al. 2016; Miller 2006] and says that authority is the ability to operate on resources. Using the extra information that effect members and annotations provide, we can now talk about authority of modules (and objects) in an application.

5.7.1 Authority Safety. We define an authority-safe programming language as one that provides a way for a software developer to specify and limit modules' (or objects') authority over other modules (or objects) using a set of well-defined rules. Through examples in Sections 2 through 4, we illustrated how a software developer could use effect annotations to specify and control modules' authority. Our formal system ensures that the program behavior adheres to the rules specified by the software developer. Specifically, Wyvern's static semantics (Section 5.4) checks that effect annotations correspond to the effects produced by each method body, and Theorem 3 guarantees that effects produced during execution adhere to the effect annotations in the program, because preservation states that the effect of a program decreases only as a program executes and never increases (the effect may decrease because it is conditional and the guarding condition is not fulfilled or because the effect takes place and the remainder of the program does not have that effect). Then, since we proved the type soundness of Wyvern's effect system, we proved Wyvern authority safe.

5.7.2 Authority Provided by an Object. A basic notion in the authority analysis of an application is the notion of what authority an object provides, which we define next.

Definition 2 (Authority Provided by an Object). The authority provided by an object is a set of effects that a client can produce using that object's public interface.

This definition is "outward facing" in a sense that it helps reasoning about the authority that the client object can gain by using the object in question. We chose such a definition because it seems to be more useful in a security analysis. For example, if an application's programming interface allows plugins to access a specific module (e.g., the logger module described in Figure 2), it is useful to be able to determine what effects a plugin could produce by using that module, accessing its public fields, and calling public methods on it.

Formally, we represent the authority provided by an object as a set of *auth* rules, shown in Figure 18. An object's authority (AUTH-OBJECT) is the authority of the object's declarations. The authority of a method (AUTH-DEF) is the effects that the method produces during execution and the authority of objects that the method can release to the caller—including authority in covariant

```
auth(new(x \Rightarrow \overline{d}))
                                                   auth(\text{new}(x \Rightarrow \overline{d})) = \bigcup_{d \in \overline{d}} auth(d) \text{ (Auth-Object)}
auth(d)
                               auth(\text{def } m(x:\tau_1):\{\varepsilon\} \ \tau_2=e)=\varepsilon \cup contra-auth(\tau_1) \cup auth(\tau_2) \ (\text{Auth-Def})
                                                           auth(var f : \tau = n) = auth(\tau) (Auth-Var)
                                                           auth(effect q = \{\varepsilon\}) = \emptyset (Auth-Effect)
auth(\tau)
                contra-auth(\tau)
          \frac{\tau = \{x \Rightarrow \overline{\sigma}\}}{auth(\tau) = \bigcup_{\sigma \in \overline{\sigma}} auth(\sigma)} \text{ (Auth-Type)} \quad \frac{\tau = \{x \Rightarrow \overline{\sigma}\}}{contra-auth(\tau) = \bigcup_{\sigma \in \overline{\sigma}} contra-auth(\sigma)} \text{ (Contra-Auth-Type)}
auth(\sigma)
                contra-auth(\sigma)
                              auth(\text{def } m(x:\tau_1):\{\varepsilon\} \ \tau_2) = \varepsilon \cup contra-auth(\tau_1) \cup auth(\tau_2) \ (\text{Auth-DefType})
                                                          auth(\text{var } f: \tau) = auth(\tau) \text{ (AUTH-VARTYPE)}
                                                          auth(effect q...) = \emptyset (Auth-EffectType)
                    contra-auth(def\ m(x:\tau_1):\{\varepsilon\}\ 	au_2)=auth(	au_1)\cup contra-auth(	au_2) (Contra-Auth-DefType)
                                         contra-auth(var\ f:\tau) = contra-auth(\tau)\ (Contra-Auth-VarType)
                                              contra-auth(effect g...) = \emptyset(ContraAuth-EffectType)
```

Fig. 18. Rules defining authority of an object.

occurrences ($auth(au_2)$) or contravariant occurrences (handled with a separate judgment, contra-auth). The reason for including the latter authority component is that whenever the method is called, an object of the return type is returned to and may be operated on by the caller, thus increasing the caller's authority. For the same reason, authority of an object's field (Auth-Var) is the authority of objects of the field's type. An effect declaration carries no authority by itself (Auth-Effect) because authority is about what effects the object can have, not what effects it can describe.

Defining the authority of a type is important because narrowing the interface of an object can be used to provide less authority to clients. The authority provided by objects of a particular type (AUTH-TYPE) is the authority of the type's declarations. Authority of a method-declaration type (AUTH-DEFTYPE) and a field-declaration type (AUTH-VARTYPE) is similar to the authority of corresponding declarations in an object. An effect-declaration type produces no authority regardless of whether the effect is abstract, concrete, or bounded (AUTH-EFFECTTYPE).

The *contra-auth* judgment treats method arguments contravariantly in the authority analysis. It is not necessary to capture the effects of method arguments directly, as we would if we just used the *auth* rule on the method argument type. The method arguments represent the authority of the caller, not of the object; thus, this object does not have authority to them. On the other hand, if the method argument is an object, it may itself have methods. If those methods are invoked by the original method, it may pass objects containing authority to them. Thus, when *contra-auth* analyzes an argument object, it contravariantly invokes *auth* on its arguments. The rules for *contra-auth* thus mirror those for *auth* with the exception that *contra-auth* does not capture the effect of the method itself because that method is from the original caller, not the callee—the callee may not even invoke it! This form of analysis follows the general rule of contravariant function types, but it also more specifically echos the capability-based, effect-bounding analysis of Craig et al. [2018].

As an example of how these rules can be applied in practice, we introduce the following definitions:

5:24 D. Melicher et al.

```
1
   resource type FileUser
    def useFile(f: File): {} Unit
2
   resource type FileHolder
     def accessFile(m: FileUser): {} Unit
5
   module def fileHolder(f: File): {} FileHolder
      var myFile: File = f // private variable
8
      def accessFile(m: FileUser): {} Unit
9
       m.useFile(myFile)
10
11
   module def example(fileHolder: FileHolder)
12
     var exfiltratedFile : Option[File] = None()
13
     var fileSetter: FileUser = new
14
       def useFile(f: File): {} Unit
15
          exfiltratedFile = Some(f)
16
      fileHolder.accessFile(fileSetter)
17
      exfiltratedFile.get().write("I have access to this file!")
19
   // Top-level code
20
   require fileSystem
21
   val f = fileSystem.fileFor(/* path to file */)
   val fileHolder = fileHolder(f)
   example(fileHolder)
```

The FileUser type contains a method useFile that receives a File and returns a unit value without causing any effect. Then, we define the fileHolder module of type fileHolder. The module contains a variable myFile, which is a file that is hidden from the user of the module, since the variable does not appear in the type fileHolder. The module also defines a method accessFile, which receives a FileUser m and simply calls the method m.useFile on myFile.

We define the module example on line 12. The variable exfiltratedFile is initially defined as an empty optional value. The object fileSetter has type FileUser and its method sets the variable exfiltratedFile to the argument f. Then, on line 17, fileSetter is passed to method accessFile. This method call sets exfiltratedFile to fileHolder.myFile, which is a private variable. The example module therefore gains the authority of writing to a file from fileHolder

We need to account for the effect fileSystem.Write in the authority of FileHolder. By (Auth-Def), the authority of FileHolder contains auth(Unit) and contra-auth(FileUser). The former authority is empty, while by (ContraAuth-DefType), the latter authority contains auth(File), which includes the file writing effect. Therefore, the authority auth(FileHolder) contains fileSystem.Write, even though it does not directly provide access to a File.

5.7.3 Authority Attenuation. Introduced in Mark Miller's dissertation [Miller 2006], the notion of authority attenuation can be described as follows. If a module (or an object) accesses a resource and produces less than the total possible set of operations on that resource, we say that the module (or object) attenuates the resource. For example, consider the modules in Figures 2, 3, and 5. We observe that while the file module can have a number of effects (Read, Write, Append, etc.), the logger module produces only two of file's effects (Read and Append). Then, any module that uses logger and does not have access to the file module (e.g., the codeCompletion plugin module) can produce on file at most the two effects that logger can produce. Thus, the logger module attenuates the file module by giving access to only a subset of file's effects.

$$tLookup(\Gamma, \tau, \overline{x.g}) = \bigcup_{x.g \in \overline{x.g}} tLookup(\Gamma, \tau, x.g) \text{ (TLookup)}$$

$$\frac{\Gamma \vdash x : \{\} \ \tau}{tLookup(\Gamma, \tau, x.g)} = \bigcup_{x.g \in \overline{x.g}} tLookup(\Gamma, \tau, x.g) \text{ (TLookup)}$$

$$\frac{\Gamma \vdash x : \{\} \ \tau}{tLookup(\Gamma, \tau, x.g) = \tau.g} \text{ (TLookup-Stop-1)}$$

$$\frac{\Gamma \vdash x : \{\} \ \tau' \quad \tau' \neq \tau \quad \tau' = \{y \Rightarrow \overline{\sigma}\} \quad \text{(effect } g \in \overline{\sigma} \lor \text{effect } g \geqslant \{\varepsilon\} \in \overline{\sigma})}{tLookup(\Gamma, \tau, x.g) = \tau'.g} \text{ (TLookup-Stop-2)}$$

$$\frac{\Gamma \vdash x : \{\} \ \{y \Rightarrow \overline{\sigma}\} \quad \tau \neq \{y \Rightarrow \overline{\sigma}\} \quad \text{(effect } g = \{\varepsilon\} \in \overline{\sigma} \lor \text{effect } g \leqslant \{\varepsilon\} \in \overline{\sigma})}{tLookup(\Gamma, \tau, x.g) = tLookup(\Gamma, \tau, [x/y]\varepsilon)} \text{ (TLookup-Recurse)}$$

Fig. 19. Wyvern effect-lookup rules that target a specific type.

To aid a security analyst in a formal security analysis of an application, we formalized the notion of authority attenuation. Importantly, our definition of authority attenuation is static. We examine only an object's code and do not know which specific objects the object uses at runtime. Instead, we can talk about objects of a specific *type* that the object uses. Our definition of authority attenuation benefits from this since we can talk about *groups of objects*, any object in which is attenuated. For example, using our static definition of authority attenuation, instead of knowing that the logger module attenuates the file module (which is of type File), we know that logger attenuates *all* objects of type File.

In essence, our formal definition says that if we let F_1 be the set of effects that represents an object's authority and let F_2 be the set of effects that represents authority of objects of a specific type, then, if F_1 and F_2 share at least one effect and there is at least one effect that is in F_2 but not in F_1 , we say that the object attenuates objects of that type. For example, if we let F_1 be the set of effects that represents the logger's authority and let F_2 be the set of effects that represents authority of objects of type File, then, if F_1 and F_2 share at least one effect and there is at least one effect that is in F_2 but not in F_1 , we say that logger attenuates objects of type File. Formally, we write these conditions as follows.

Definition 3 (Authority Attenuation⁴). An object o attenuates objects of type τ , if $F_1 = tLookup(\Gamma, \tau, auth(o)), F_2 = tLookup(\Gamma, \tau, auth(\tau)), F_1 \cap F_2 \neq \emptyset$, and $F_2 \setminus F_1 \neq \emptyset$.

First, using the *auth* rules shown in Figure 18, we find authority of object o and of objects of type τ . Then, we use the *tLookup* rules, shown in Figure 19, to "normalize" the two effect sets, making it possible to compare them. Finally, we compare the two effect sets.

The *tLookup* rules support the static nature of our definition of authority attenuation. They resolve effects to lower-level effects by "searching" for effects of an object of a particular *type* and stopping when an object of that type is found. When we apply *tLookup* to a set of effects, we apply *tLookup* to each effect in that set (TLOOKUP). If the type that we are looking for is the type of the current object (TLOOKUP-STOP-1), we return the "normalized" form of the effect, which differs from the original form in that we substitute the variable name with the type name. If we encounter an abstract effect (TLOOKUP-STOP-2), we return the "normalized" form of that effect that uses the type of the current object. Otherwise, the effect is concrete, and we proceed by examining the effect's definition (TLOOKUP-RECURSE).

⁴It is possible to create a more general formal definition of authority attenuation by, instead of considering one object that attenuates objects of a specific type, considering objects of one type that attenuate objects of another type.

5:26 D. Melicher et al.

```
resource type Plugin
effect Run
def getName(): {} String
def run(): {this.Run} Unit
```

Fig. 20. The Plugin type that each text editor's plugin must implement.

As an example, let us apply our definition of authority attenuation to the logger module and the objects of type File (e.g., the file module) from Figures 2 and 3, respectively.

```
auth(\operatorname{logger}) = auth(\operatorname{effect} \ \operatorname{ReadLog} = \{f.\operatorname{Read}\}) \cup auth(\operatorname{effect} \ \operatorname{UpdateLog} = \{f.\operatorname{Append}\}) \cup \\ auth(\operatorname{def} \ \operatorname{readLog})\colon \{\operatorname{this}.\operatorname{ReadLog}\} \ \operatorname{String} = f.\operatorname{Read}()) \cup \\ auth(\operatorname{def} \ \operatorname{updateLog}(\operatorname{newEntry}\colon \operatorname{String})\colon \{\operatorname{this}.\operatorname{UpdateLog}\} \ \operatorname{Unit} = f.\operatorname{append}(\operatorname{newEntry})) \\ = \emptyset \cup \emptyset \cup \{\operatorname{this}.\operatorname{ReadLog} \cup auth(\operatorname{String})\} \cup \{\operatorname{this}.\operatorname{UpdateLog} \cup contra-auth(\operatorname{String})\} \\ = \{\operatorname{this}.\operatorname{ReadLog}, \ \operatorname{this}.\operatorname{UpdateLog}\} \\ tLookup(\Gamma, File, auth(\operatorname{logger})) \\ = tLookup(\Gamma, File, \{\operatorname{this}.\operatorname{ReadLog}, \ \operatorname{this}.\operatorname{UpdateLog}\}) \\ = tLookup(\Gamma, File, \ \operatorname{this}.\operatorname{ReadLog}) \cup tLookup(\Gamma, File, \ \operatorname{this}.\operatorname{UpdateLog}) \\ = tLookup(\Gamma, \operatorname{File}, \ \operatorname{f.Read}) \cup tLookup(\Gamma, \operatorname{File}, \ \operatorname{f.Append}) \\ = \{\operatorname{File}.\operatorname{Read}, \ \operatorname{File}.\operatorname{Append}\}
```

Using the *auth* and *tLookup* rules on the logger object, we find that logger's authority is $F_1 = \{File.Read, File.Append\}$. Similarly, we find that the authority of objects of type File is $F_2 = \{File.Read, File.Write, File.Append, \ldots\}$. Then, comparing the two sets, we have that $F_1 \cap F_2 = \{File.Read, File.Append\} \neq \emptyset$ and $F_2 \setminus F_1 = \{File.Write, \ldots\} \neq \emptyset$. Thus, by our definition, the logger module attenuates modules of type File.

Definition 4 (Authority Attenuation (more generally)). Objects of type τ_1 attenuate objects of type τ_2 , if

```
(1) F_1 = tLookup(\Gamma, \tau, auth(\tau_1)), F_2 = tLookup(\Gamma, \tau, auth(\tau_2)),
```

- (2) $F_1 \cap F_2 \neq \emptyset$, and
- (3) $F_2 \setminus F_1 \neq \emptyset$.

This definition essentially says that if we let F_1 be the set of effects that represents authority of objects of one type and let F_2 be the set of effects that represents authority of objects of another type, then, if F_1 and F_2 share at least one effect and there is at least one effect that is in F_2 but not in F_1 , we say that objects of the former type attenuate objects of the latter type. Definition 3 is more general than Definition 2 because Definition 3 defines authority attenuation for all of the objects of a particular type rather than for a single object.

6 CASE STUDY: AN EXTENSIBLE TEXT-EDITOR APPLICATION

Effect checking was implemented as part of Wyvern.⁵ To evaluate our effect-system design, we used Wyvern to create an extensible text-editor application and plugins for it, similar to the running example described earlier.⁶

6.1 Application Description

The text-editor application provides basic text-editing functionality. When started, the text-editor window has a text area where the user may enter or edit text. The title bar shows the path to the currently opened document or "Untitled" if the document has not been saved yet. The menu bar

⁵https://github.com/wyvernlang/wyvern.

 $^{^6} https://github.com/wyvernlang/wyvern/blob/master/examples/text-editor/README.md.\\$

,					
	Definitions	Annotations			
	Delimitions	and parameters			
Text editor	1.2	1.1			
Plugins	3.6	0.8			
Overall	1.3	1.0			

Table 1. The Average Number (Geometric Mean) of Effects Per an Effect Set

has three options: it allows users to perform operations on files, perform editing operations on the text in the text area, or run plugins. The main application module is called textEditor and is of type TextEditor.

We implemented the following three plugins:

- darkTheme sets the theme of the text editor to have a dark background and light text,
- ullet questionnaireCreator extracts questions from the currently opened document and creates a questionnaire in a separate file, and
- wordCount counts the number of words in the currently opened document and displays that number to the user in a pop-up window.

All plugins must implement the Plugin type, shown in Figure 20.

6.2 Observations and Discussion

During the implementation of the text-editor application, we made several observations that stem from the way we designed Wyvern's effect systems, which we present and discuss next.

6.2.1 Effect Aggregation. Table 1 shows the average number (geometric mean⁷) of effects in each effect set used in the implementation. This aspect speaks to the amount of boilerplate code that the effect-aggregation feature of our effect-system design eliminates.

The average number of effects in the effect-definition sets is much lower for the text editor than for the plugins, which signals that effects declared in the text editor are usually composed of fewer effects than those declared in plugins. There are at least two reasons for that. The main reason is that a text editor's methods frequently use only one resource each and perform only one operation on it whereas, in a plugin, the run method tends to use all of the resources that the plugin has access to. Another minor reason is that the textEditor module defines an effect, called SaveFile, whose definition consists of four effects, which is then used as a shorthand in defining two out of seven textEditor's effects.

In contrast with effect definitions, the difference between the average numbers of effects in effect-annotation sets in the text editor and the plugins is insignificant, and the numbers are low. For the text-editor application, the reason is that the same <code>SaveFile</code> is used to annotate 5 out of 15 (i.e., one third of) <code>textEditor</code>'s methods. In addition, three more <code>textEditor</code>'s methods have empty effect annotations. For the plugins, the reason is that there is only one method (the <code>run</code> method) that has an effect annotation with an effect (<code>Plugin</code>'s <code>Run</code> effect) in it, and the rest of the methods have empty effect annotations.

Overall, these observations imply that the effect-aggregation feature has its merits and indeed serves to reduce the amount of effect-related code.

6.2.2 Effect-Annotation Overhead. Table 2 presents a higher-level picture of the effect-annotation overhead. Overall, the effect-annotation overhead comes from three sources: effect

 $^{^7}$ To handle zeros in the data, we added one to each value, calculated the mean, and then subtracted one from the result.

5:28 D. Melicher et al.

	LoC	Effect declarations		Effect annotations		Effect parameters		Total	
Text editor	250	38	(15%)	56	(22%)	3	(1%)	97	(39%)
Plugins	110	3	(3%)	12	(11%)	6*	(5%)	21	(19%)
Total	360	41	(11%)	68	(19%)	9	(3%)	118	(33%)

Table 2. The Effect-Annotation Overhead in the Text-Editor Application and Its Plugins

For the plugins, the number of lines that contain effect parameters, marked with an asterisk, includes the lines where plugins are instantiated that are located in the text editor's code.

declarations, effect annotations on methods, and effect parameters. The important distinction among these types of annotation overhead is that effect declarations require adding new lines of code to the implementation, whereas effect annotations and effect parameters change the lines of code that would still exist in unannotated code.

Incorporating effects into the text editor's code base led to an 11% increase in its size and affected 22% of (the enlarged version of) it. Thus, overall, 33% of code was affected by the inclusion of the effect information. The overhead is lower for the plugins than for the text-editor application itself. The reason for this difference is that the text-editor application accesses and operates on more resources than any one plugin does, and the application defines the effects that the plugins may have. In contrast, all three plugins define and use only one effect, Run, that involves using resources (and one more empty effect). In addition, none of the plugins introduces new effects that would be used only by the plugin itself. Based on this pattern, we do not expect effect annotations to be a deterrent to implementing plugins, and we expect the ratio of affected lines to go down as more plugins are added.

We did see one source of verbosity from effects. Since the TextEditor type is agnostic to its plugins, it has an abstract effect member Run that represents the effects that its plugins have. When the text editor runs a plugin, it incurs the Run effect. This saves annotation from within the TextEditor, but when we instantiate a TextEditor to be run with a specific set of plugins, we must parameterize it in a way that binds the Run effect to the set of all the effects of all of its plugins. In our example, this is 7 separate effects, and there might be even more if there were additional plugins. This verbose parameterized type also appears in the module header of the textEditor implementation. While this creates a couple of very long lines of code (e.g., 192 characters in the type instantiation in main), we view it as a good trade-off given that this large set of effects is encapsulated by the abstract Run effect everywhere else.

6.2.3 Information Hiding and Polymorphism. There are three examples of information hiding and polymorphism that we observed in the text-editor application. The first example is in the plugin modules. Different plugins naturally have different effects; for example, the darkTheme and questionnaireCreator plugins both use the logging module, but otherwise have disjoint effects. We defined an abstract effect Run in the Plugin module, which is then given a different concrete definition in each concrete plugin implementation.

The second example is in the logging module. Currently, the logger module is implemented using the file system and stores the log file locally in a file. In the future, the text editor can be made distributed, and the logging module could maintain a log file which is stored somewhere else on the network (e.g., as was suggested in Section 4.1). Due to effect abstraction, this change is easily accommodated in the current version of the text editor's code. As long as the new, distributed logger implements the Logger type, the modules that use logger are not affected by the substitution.

The third example uses the effect-hierarchy feature of our effect system. Similar to the RemoteLogger example in Section 4.1, we use an effect hierarchy to hide the definitions of the

UI-related effects. Namely, we define the UIEffects type, which specifies a lower bound for each UI-related effect, and then ascribe this type to the uiEffects pure module that defines those effects.

In addition, our design can accommodate one more possible change. Currently, the logger module appends to the log file only, thus producing the Append effect on the logFile module, which is reflected in the definition of logger's Update effect. Alternatively, logger could write to the log file, aggregating the information that has been already logged, for example, substituting "X action occurred. X action occurred." with "X action occurred 2 times." In such a case, logger would produce the Write effect on the logFile module, and the definition of logger's Update effect would change accordingly. Due to effect abstraction, there would be no difference for the modules that use the logger module, which would still produce logger's Update effect.

6.2.4 Controlling Operations Performed on Modules. In the text-editor application, plugins may be written by some third party; thus, their code is untrusted. To verify that plugins make no illegal calls, we need to check the effect annotations on the plugins and analyze the legitimacy of the effects produced on each text-editor's module that plugins access. For example, the questionnaireCreator plugin produced the Update effect on logger, the Append and Write effects on fileSystem, and the Read effect on textArea. These effects are congruent with questionnaireCreator's expected functionality: the plugin produces the Update effect on logger to update the log file, the Write and Append effects on fileSystem to create a new file containing the resulting questionnaire and to append to it when a question is encountered in the original text, and the Read effect on textArea to read in the current version of the opened document. If questionnaireCreator had any more effects, those effects would have been unauthorized. Thus, all of the effects that the plugin produces are legitimate. We performed a similar verification on the other plugins and determined that they are given access to the minimal number of text editor's modules and, on those, they perform only the necessary operations.

In additional to relying on a security analyst to manually inspect the effect of a module. We can define a type whose effect signature bounds what the plugin can do-limiting all access to resources, not just files but also system.FFI. Then, we can assign the plugin to a variable of that type. If the assignment succeeds, we are guaranteed that the plugin obeys the effects of the type we defined.

Moreover, our effect system allows expressing the intent that a method may not produce any effects, which is then enforced by Wyvern's effect system. During the implementation of the text editor, we used this feature when defining the Plugin type (Figure 20). We added a method, called getName, that returns the plugin's name so that the plugin can be added to the text editor's user menu. All that getName needs to do is to return a String with the plugin's name and, thus, the method must produce no effects. To enforce this restriction, we annotated the method with {}, that is, an empty effect set, which achieved the desired behavior.

- 6.2.5 Designating Important Resources Using Globally Available Effects. While implementing the text-editor application, to define the effects of the UI-related modules, we made a design choice to use globally available effects (discussed in detail in Section 4.3). We defined the UI-related globally available effects in the pure module called <code>uieffects</code> and used them throughout the code, designating the importance of tracking effects on the UI and making it more obvious where those effects are produced in the code.
- 6.2.6 Authority Attenuation. Section 5.7.3 describes how our effect system allows formalizing authority attenuation. In practice, as suggested in Section 4.4, since method effect annotations expose the information about how resources are used, a software developer is able to identify occurrences of authority attenuation by looking at modules' interfaces.

5:30 D. Melicher et al.

In the text-editor application, by examining only module interfaces, we were able to determine that the logger module attenuates the logFile object. While logFile has three effects: Read, Write, and Append, logger produces only the Append effect. This means that the logger module allows for only a limited set of effects to be produced on logFile, thus, attenuating it. Considering the structured nature of module interfaces, we believe that it is feasible to automate this discovery process.

- 6.2.7 Effect Hierarchy. In the text editor, we used the effect-hierarchy feature of our effect system twice. The first use case is in hiding the definitions of the UI-related effects, which is discussed in Section 6.2.3. The second use case is in refining the Plugin type (see Figure 20) specifically for the text editor's theme plugins. We created a new type, called ThemePlugin, which is identical to the original Plugin type except for its Run effect being lower-bounded by the UI effects necessary for updating the way the UI looks. We then used the ThemePlugin type when adding theme-related plugins to the text editor.
- 6.2.8 Controlling FFI Effects. The implementation of the text editor requires calling into several Java methods through an FFI. Our effect system expects all Wyvern methods calling into the Java FFI to be annotated with the system.FFI effect. Therefore, whenever a Wyvern method calling into the Java FFI was not annotated with a proper effect, the compiler rejected the program. Moreover, to have a fine-grained control of different kinds of FFI effects, based on system.FFI, we defined higher-level UI effects, such as ShowDialog and ReadTextArea, in the uiEffects module. This way, we can reason about the effects of Wyvern methods that call into different Java methods separately.
- 6.2.9 Additional Validation of the Wyvern Effect System . Fish et al. [2020] describe the application of our effect system to the Wyvern standard library. The article presents a case study of a small standard I/O library seeking to use the effect system of Wyvern for tracking the secure use of resources. The study suggests that the effect system of Wyvern is indeed practicable and useful and, thus, potentially promising for inclusion in other future language designs.

7 DISCUSSION: LIMITATIONS AND BENEFITS OF THE DESIGN

Our contributions center around abstraction, but it is instructive to consider how fundamental this is to our design. Our formal system has no "built-in" effects, which means that the soundness theorem does not actually say anything about whether the annotated effects describe any technical aspects of execution behavior. Instead, soundness means that declared effects are respected: as a program executes, the overall effect of the program never increases. Our practical implementation does more to connect effects to program semantics than the formal system, but only barely: it builds in only one effect, representing foreign function interface use.

On the one hand, this can be considered a weakness. It means that programmers must take care that when they use the foreign function interface or do something else in code that they want to track in the effect system, they label the relevant low-level methods with appropriate abstract effects. When using effects for security purposes, the security of the system rests not just on the effect system but also on the way that it is used. Code reviews or careful study by security analysts—doing reasoning of the form illustrated in the case study described earlier—is likely necessary to accomplish this. We believe, however, that this is not so much of a limitation as it may initially seem: after all, the security of any system does not rest merely on types or theorems about the system but also on expert analysis connecting those types and theorems to security characteristics that people care about in the real world.

The strength of our approach is the flip side of this limitation. In particular, programmers can use abstraction to reason about effects which are internally defined in terms of system.FFI (or perhaps nothing at all) but abstract to clients. In essence, what our approach does is allow security

analysis to focus primarily on the lowest-level parts of the system, for example, where the FFI is used—places in code that typically must be treated with great care in any approach. Our effect system then allows programmers to leverage arbitrary effect abstractions expressed at this low level, track them throughout the application, and build higher-level effect abstractions. Although our effect system lacks the precise semantics of effects, which is provided by denotational effect systems, our preservation theorem implies that abstract effects unrelated by subtyping cannot be mixed, which still provides powerful tools for reasoning about the effectful behavior of a program.

8 RELATED WORK

This article is derived in part from the Ph.D. and Master's theses of the two first authors [Melicher 2020; Xu 2020]. The Ph.D. thesis includes a version of the formal system that includes type members but no lower or upper bounds on them; it also includes the case study. The definition of authority in this article also comes from the Ph.D. thesis but is extended now to consider authority over objects that can escape by being passed to the funarg of a higher-order function. The master's thesis includes the version of the formal system with lower and upper bounds.

Origins of Effect Systems. Effect Systems were originally proposed by Lucassen [1987] to track reads and writes to memory Then, Lucassen and Gifford [1988] extended this effect system to support polymorphism. Effects have since been used for a wide variety of purposes, including exceptions in Java [Kiniry 2006] and asynchronous event handling [Bračevac et al. 2018]. Turbak and Gifford [2008] previously proposed effects as a mechanism for reasoning about security, which is the main application that we discuss.

Denotational vs. Descriptive Effects. Filinski [2010] makes a distinction between two strands of work on effects. A *denotational* approach, which includes algebraic effects, defines the semantics of computational effects based on primitives. A *descriptive* approach (e.g., Java's checked exceptions) takes effects that are already built into the language—such as reading and writing state or exceptions—and provides a way to restrict them. Although descriptive effect systems are capable of controlling side effects, one of the limitations is that it does not have semantic meanings that precisely describe the semantics of a program. In this terminology, our approach is descriptive rather than denotational.

Prior Work on Bounded Effect Polymorphism. A limited form of bounded effect polymorphism was explored by Trifonov and Shao [1999], who bound effect parameters by the resources that they may act on. However, the bound cannot be another arbitrary effect, as in our system. Long et al. [2015] use a form of bounded effect polymorphism internally but do not expose it to users of their system.

Brachthäuser et al. [2020a] present a system with contextual effect polymorphism, which is different from parametric effect polymorphism because contextual effect polymorphism is not explicit; rather, it implicitly arises from the calling context. They did not explore bounded effect polymorphism in their article.

Defining Application-Specific Effects. Marino and Millstein [2009] discuss an effect system in which application-specific effects can be defined. One of their examples is system calls that can block, but their design does not provide the benefit of a semantic tie-in to the foreign function interface, as ours does.

Capability-Based Module System. Melicher et al. [2017] introduced a capability-based module system that allows authority control. The idea of authority attenuation in our article is realized by using the capability-safe property of the module system in their article. However, their article does

5:32 D. Melicher et al.

not provide an effect system; therefore, it does not enjoy the benefits of the design of our effect system, such as effect abstraction and effect bounds. Furthermore, we are able to define authority and authority attenuation using effects, which was not possible in the earlier work. Notably, the earlier work defined an object's "authority" as the set of objects it has access to. Following Miller [2006] and others, we now prefer the term "permissions" for this and define "authority" as the ability to operate on resources.

Path-Dependent Effects. JML's data groups [Leino et al. 2002] have some superficial similarities to Wyvern's effect members. Data groups are identifiers bound in a type that refer to a collection of fields and other data groups. They allow a form of abstract reasoning in that clients can reason about reads and writes to the relevant state without knowing the underlying definitions. Data groups are designed specifically to capture the modification of state, and it is not obvious how to generalize them to other forms of effects.

The closest prior work on path-dependent effects, by Greenhouse and Boyland [1999], allows programmers to declare regions as members of types. This supports a form of path dependency in read and write effects on regions. Our formalism expresses path-dependent effects based on the type theory of DOT [Amin et al. 2014], which we find to be cleaner and easier to extend with the unique bounded abstraction features of our system. The type members of Amin et al.s can be left abstract or refined by upper or lower bounds, and were a direct inspiration for our work on bounded abstract effects.

Subeffecting. Rytz et al. [2012] supports more flexibility via an extensible framework for effects. Users can plug in their own domain of effects, specifying an effect lattice representing subeffecting relationships. Each plugin is monolithic. In contrast, our effect members allow new effects to be incrementally added and related to existing effects using declared subeffect bounds.

Algebraic Effects, Generativity, and Abstraction. Algebraic effects and handlers [Plotkin and Power 2003; Plotkin and Pretnar 2009] are a way of implementing certain kinds of side effects and control effects, such as exceptions and mutable state in an otherwise purely functional setting. As described earlier, algebraic effects fall into the "denotational" rather than "descriptive" family of effects work; these lines of work are quite divergent, and it is often unclear how to translate technical ideas from one setting to the other. However, certain articles explore parallels to our work despite the major contextual differences.

Biernacki et al. [2020] discuss the use of generative effects in the setting of algebraic effects and handlers and provide a calculus that supports instances of algebraic effects. Their system does not provide mechanisms for existential abstraction of generative effects and does not support path dependency or bounds, like our system. Bračevac et al. [2018] use generative effects to support asynchronous, event-based reactive programs. However, their generativity is at a per-module level, whereas our work supports per-object generativity.

Zhang and Myers [2019] describe a design for algebraic effects that preserves abstraction in the sense of parametric functions: if a function does not statically know about an algebraic effect, that effect tunnels through that function. This is different from our form of abstraction, in which the definition of an effect is hidden from clients.

Koka [Leijen 2014] provides a static effect system that is capable of tracking external side effects such as mutable reference cells, but their system does not support effect abstraction and lacks the expressiveness provided by path-dependent types. Koka provides a hierarchy of effects inspired by Haskell's standard monads. However, its language does not provide the mechanism that allows the programmer to construct a hierarchy of abstract effects or to compose abstract effects whose operations are unknown.

Biernacki et al. [2019] discuss how to abstract algebraic effects using existentials. Our effect members function similarly to existentials but provide more expressiveness because of generativity of path-dependent types and the ability to bound effect members by other effects.

JEff [Inostroza and van der Storm 2018] explores integrating algebraic effects with object-oriented programming languages, mainly focusing on effect handling. In the effect system of JEff, all methods declarations need to be annotated with a concrete effect type. Therefore, JEff does not provide effect abstraction or effect polymorphism.

Bounded Abstract Algebraic Effects in the Effekt Library. The Effekt library [Brachthäuser et al. 2020b] explores algebraic effects as a library of the Scala programming language. Since their effect system is built on type members of Scala, it bears some similarities to our system. In particular, Effekt supports a form of effect composition as well as a form of bounded effect abstraction.

Compared with Effekt, our work contributes a formal definition of bounded abstract effects and the confidence that comes from the effect soundness theorem that we proved. While the core of Scala's type system has been formalized and proven sound, the related mechanisms for effects had not previously been formalized nor theorems specifically about effect soundness proven. These formalizations and theorems share some mechanisms with Scala, but also have unique aspects (e.g., preservation of effects and effect composition). We also formalize authority and authority attenuation.

More subtly, the fact that Effekt captures algebraic (i.e., denotational) effects places restrictions on how abstraction can be used. The form of abstraction in Effekt is indeed useful for algebraic effects, but it is unsuitable for the security applications we explore in the context of descriptive effects. The key issue is that a denotational effect system relates two program points: the place where an effect operation is used and the place where the effect operation is handled. Although effects can be abstracted in Effekt, this abstraction applies only between the effect and the handler that implements the effect. For example, you can define a handler for log operations that works by writing the log to a file; between that handler and the uses of the log operation, abstraction can be used to show only an abstract log effect and not a file write effect. However, outside the handler, the file write effect will be shown; what effects are used in the implementation of the logger cannot be hidden from the surrounding context. Hiding the low-level effects used in the implementation of the logger or of other similar modules is, of course, exactly the point of our system—we are able to hide this from the surrounding context exactly because our effects are descriptive, not denotational, and they do not have handlers.8 Our applications to security and our case study use exactly the kind of abstraction that is supported in our descriptive effect system but which cannot be (modularly) supported in an algebraic/denotational effect system such as Effekt.

A final difference between our work and Effekt is a practical one: Effekt is implemented as a library, not a language extension. Thus, it requires mechanisms both to support the runtime semantics and the static effect-checking semantics of algebraic effects. Effekt uses monads for both of these purposes, which has an impact on developers and the way they write their programs. Using monads is likely acceptable in small portions of the program where control effects are desired; thus, monads may be a good solution in a system focused on denotational effects. However, a system focused on descriptive effects needs to reason about the effects program-wide: this means that most of the program would be written with monads. Many developers feel that monads make code more awkward, and may be unwilling to use monads at this large scale. Overall, Effekt's strategy clearly facilitates adoption because the language need not be changed, but there are trade-offs in usability,

⁸One could argue that such abstraction could be supported in Effekt by placing all such handlers around the top-level program; we regard this as anti-modular because it requires a global restructuring of the program—that is to say, our solution adds expressiveness in the context of Felleisen's work [Felleisen 1991].

5:34 D. Melicher et al.

especially at large scale. Our work shows how to design bounded effect abstraction as a language extension, providing a cleaner interface to developers.

Coeffect systems. Recently, coeffects have been proposed as a dual construct to effects [Petricek et al. 2014]. The relationship between effects and coeffects has been explored mostly in a denotational setting (e.g., see Gaboardi et al. [2016]), and it is less clear how they relate in the descriptive setting where we work. Broadly, coeffects use "input resources" to limit what a piece of code can do, whereas effects are an output that describes what it actually does. Many phenomena can be modeled either as effects or coeffects: for example, exceptions are normally viewed as an effect, but they can also be phrased as coeffects in which the exception handler is the input resource. We follow both the surface-level design and language formalism from the effect literature, which is still the dominant line of research; exploring connections to coeffects is left to future work. We observe, however, that since our effects are path dependent, they can refer to "resources" such as files that are passed in. Thus, our system features some kinds of expressiveness that are more typical of coeffect systems.

Authority Attenuation. Although a number of works on authority safety mention and explain authority attenuation (e.g., Mettler et al. [2010]; Miller [2006]), the only work on formalizing authority attenuation that we are aware of is a workshop presentation by Loh and Drossopoulou [2017]. In the presentation, the authors used Hoare triples and invariants to show how authority can be attenuated in a restricted document object model (DOM) tree. In contrast, our approach to authority attenuation uses effect abstraction and is more general, for example, allowing reasoning in contexts other than the DOM.

9 CONCLUSION

This article presented an effect member mechanism to support effect abstraction: the ability to define higher-level effects in terms of lower-level effects, to hide that definition from clients of an abstraction and to reveal partial information about an abstract effect through effect bounds. A set of illustrative examples as well as a framework/plugin case study demonstrates the expressiveness of effect abstraction, including the ability to support information hiding, the ability to characterize effects on application- or library-specific higher-level resources, and the ability to reason rigorously about the authority of untrusted code in a security context. We formally proved the soundness of our effect system and showed that it is able to formally model authority attenuation for the first time. Overall, these contributions lay a solid foundation for effect systems that can scale up and can deal with the complexities of real-world code.

APPENDIX

A PROOF OF THEOREM 2 (TRANSITIVITY OF SUBTYPING)

Lemma 2 (Narrowing For Subeffecting). *If* Γ , $x : \tau \vdash \varepsilon_1 <: \varepsilon_2$, and $\Gamma \vdash \tau' <: \tau$, then Γ , $x : \tau' \vdash \varepsilon_1 <: \varepsilon_2$.

PROOF. The proof is by structural induction on the rule to derive $\Gamma, x : \tau \vdash \varepsilon_1 <: \varepsilon_2$.

- (1) Subeffect-Subset. Since the premise does not rely on the context, this case is trivially true.
- (2) Subeffect-Upperbound. If the type of n is not changed, then we can apply the same rule to derive $\Gamma, x: \tau' \vdash \varepsilon_1 \cup \{n.g\} <: \varepsilon_2$. If the type of n is replaced by τ' , then we have that effect $g \le \varepsilon' \in \sigma$, where $\Gamma, n: \tau' \vdash \varepsilon' <: \varepsilon$. By IH, we have that $\Gamma, n: \tau' \vdash [n/y]\varepsilon \cup \varepsilon_1 <: \varepsilon_2$. By transitivity of subeffecting, we have that $\Gamma, n: \tau' \vdash [n/y]\varepsilon' \cup \varepsilon_1 <: \varepsilon_2$. Then, we can apply Subeffect-Upperbound again to derive $\Gamma, x: \tau' \vdash \varepsilon_1 \cup \{n.g\} <: \varepsilon_2$.
- (3) Subeffect-Lowerbound. This case is similar to Subeffect-Upperbound.

- (4) Subeffect-Def-1. Since the declaration type effect $g = \{\varepsilon\}$ is not changed, the result follows directly by induction hypothesis.
- (5) Subeffect-Def-2. Since the declaration type effect $g = \{\varepsilon\}$ is not changed, the result follows directly by induction hypothesis.

Lemma 3 (Narrowing For Subtyping). If $\Gamma, x : \tau \vdash \tau_1 <: \tau_2$, and $\Gamma \vdash \tau' <: \tau$, then $\Gamma, x : \tau' \vdash \tau_1 <: \tau_2$.

If $\Gamma, x : \tau \vdash \sigma_1 <: \sigma_2$, and $\Gamma \vdash \tau' <: \tau$, then $\Gamma, x : \tau' \vdash \sigma_1 <: \sigma_1$.

Proof. We induct on the number of S-Alg used to derive the typing judgment in the premise of the statement.

- BC S-Alg is not used; thus, we have that $\Gamma, x : \tau + \sigma_1 <: \sigma_2$ derived by S-Refl2 or one of the S-Effect rules. The proof is trivial if we apply Lemma 2.
- IS1 Assume that we used S-Alg n times to derive $\Gamma, x: \tau \vdash \{y \Rightarrow \sigma_i^{i \in 1...m}\} <: \Gamma \vdash \{y \Rightarrow \sigma_i^{i \in 1...n}\}$. Then, for each subtyping judgment in the premise of S-Alg, we can apply an induction hypothesis to derive $\Gamma, x: \tau', y: \{y \Rightarrow \sigma_i^{i \in 1...m}\} \vdash \sigma_{p(i)} <: \sigma_i'$. Then, by applying S-Alg, we have that $\Gamma, x: \tau' \vdash \{y \Rightarrow \sigma_i^{i \in 1...m}\} <: \Gamma \vdash \{y \Rightarrow \sigma_i^{i \in 1...m}\}$.
- IS2 Assume that we used S-Alg n times to derive $\Gamma, y: \tau \vdash \text{def } m(x:\tau_1): \{\varepsilon_1\} \ \tau_2 <: \text{def } m(x:\tau_1'): \{\varepsilon_2\} \ \tau_2'$, by inversion on S-Def; we have that $\Gamma, y: \tau \vdash \tau_1' <: \tau_1, \ \Gamma, y: \tau \vdash \tau_2 <: \tau_2'$ and $\Gamma, y: \tau, x: \tau_1 \vdash \varepsilon_1 <: \varepsilon_2$. Then, by induction hypothesis and Lemma 2, we have that $\Gamma, y: \tau' \vdash \tau_1' <: \tau_1, \ \Gamma, y: \tau' \vdash \tau_2 <: \tau_2' \ \text{and} \ \Gamma, y: \tau', x: \tau_1 \vdash \varepsilon_1 <: \varepsilon_2$. Then, we use S-Def to derive $\Gamma, y: \tau' \vdash \text{def } m(x:\tau_1): \{\varepsilon_1\} \ \tau_2 <: \text{def } m(x:\tau_1'): \{\varepsilon_2\} \ \tau_2'$.

A.0.1 Proof of Theorem 2. If $\Gamma \vdash \tau_1 <: \tau_2$ and $\Gamma \vdash \tau_2 <: \tau_3$, then $\Gamma \vdash \tau_1 <: \tau_3$. If $\Gamma \vdash \sigma_1 <: \sigma_2$ and $\Gamma \vdash \sigma_2 <: \sigma_3$, then $\Gamma \vdash \sigma_1 <: \sigma_3$.

PROOF. We induct on the the number of S-Alg used to derive the two judgments in the premise of the first statement, $\Gamma \vdash \tau_1 <: \tau_2$ and $\Gamma \vdash \tau_2 <: \tau_3$, or the two judgments in the premise of the second statement, $\Gamma \vdash \sigma_1 <: \sigma_2$ and $\Gamma \vdash \sigma_2 <: \sigma_3$.

- BC The S-Alg is not used; thus, we have that $\Gamma \vdash \sigma_1 <: \sigma_2$ and $\Gamma \vdash \sigma_2 <: \sigma_3$ by S-Refl2 or one of S-Effect. By Lemma 7 transitivity of subeffecting, it is easy to see that $\Gamma \vdash \sigma_1 <: \sigma_3$.
- IS1 Assume that we used S-Alg n times to derive $\Gamma \vdash \{x \Rightarrow \sigma_i^{i \in 1...m}\}$ <: $\{x \Rightarrow \sigma_i'^{i \in 1...n}\}$ and $\Gamma \vdash \{x \Rightarrow \sigma_i'^{i \in 1...n}\}$ <: $\{x \Rightarrow \sigma_i'^{i \in 1...n}\}$ Sy inversion of S-Alg, there is an injection $p:\{1..n\} \mapsto \{1..m\}$ such that $\forall i \in 1..n, \ \Gamma, x: \{x \Rightarrow \sigma_i^{i \in 1..m}\} \vdash \sigma_{p(i)} <: \sigma_i'$. There is another injection $q:\{1..k\} \mapsto \{1..n\}$ such that $\forall i \in 1..k, \ \Gamma, x: \{x \Rightarrow \sigma_i'^{i \in 1..m}\} \vdash \sigma_{q(i)}' <: \sigma_i''$. Thus, for each $i \in 1..k$, we have two judgments:

$$\begin{split} & \Gamma, x : \left\{ x \Rightarrow \sigma_i^{i \in 1..m} \right\} \vdash \sigma_{p(q(i))} <: \sigma_{q(i)}' \\ & \Gamma, x : \left\{ x \Rightarrow {\sigma'}_i^{i \in 1..n} \right\} \vdash \sigma_{q(i)}' <: \sigma_i'' \end{split}$$

By Lemma 3, we can write the second judgment as $\Gamma, x: \{x \Rightarrow \sigma_i^{i \in 1..m}\} \vdash \sigma_{q(i)}' <: \sigma_i''$. By IH, for all $i \in 1..k$, $\Gamma, x: \{x \Rightarrow \sigma_i''^{i \in 1..k}\} \vdash \sigma_{p(q(i))} <: \sigma_i''$. Since the function $p \circ q$ is a bijection from $\{1..k\} \mapsto \{1..n\}$, we can use the rule S-Alg again to derive $\Gamma \vdash \{x \Rightarrow \sigma_i^{i \in 1...m}\} <: \{x \Rightarrow \sigma_i''^{i \in 1...k}\}$

IS2 Assume that we used S-Alg n times to derive $\Gamma \vdash \text{def } m(x:\tau_1):\{\varepsilon_1\}\ \tau_1' <: \text{def } m(x:\tau_2):\{\varepsilon_2\}\ \tau_2' \text{ and } \Gamma \vdash \text{def } m(x:\tau_2):\{\varepsilon_2\}\ \tau_2' <: \text{def } m(x:\tau_3):\{\varepsilon_3\}\ \tau_3'.$ By inverse on S-Def, we have that $\Gamma \vdash \tau_2 <: \tau_1$, $\Gamma \vdash \tau_3 <: \tau_2$, $\Gamma \vdash \tau_1' <: \tau_2' \text{ and } \Gamma \vdash \tau_2' <: \tau_3'.$ By IH, we have that $\Gamma \vdash \tau_1' <: \tau_3'$

5:36 D. Melicher et al.

and $\Gamma \vdash \tau_3 <: \tau_1$. We have that $\Gamma \vdash \varepsilon_1 <: \varepsilon_3$ by transitivity of subeffects. Hence, we can use S-Def again to derive $\Gamma \vdash \text{def } m(x : \tau_1) : \{\varepsilon_1\} \ \tau_1' <: \text{def } m(x : \tau_3) : \{\varepsilon_3\} \ \tau_3'$.

IS3 By transitivity of subeffecting, other cases for $\Gamma \vdash \sigma_1 <: \sigma_3$ are trivial.

B PROOF OF THE TYPE SOUNDNESS THEOREMS

B.1 Lemmas

Proof. Straightforward induction on typing derivations.

LEMMA 4 (WEAKENING). If $\Gamma \mid \varnothing \vdash e : \{\varepsilon\} \tau$ and $x \notin dom(\Gamma)$, then Γ , $x : \tau' \mid \varnothing \vdash e : \{\varepsilon\} \tau$, and the latter derivation has the same depth as the former.

Proof. Straightforward induction on typing derivations.

Lemma 5 (Reverse of Subeffecting-Lowerbound). If $\Gamma \vdash \varepsilon_1 <: \varepsilon_2 \cup \{x.g\}$, $\Gamma \vdash x : \{y \Rightarrow \sigma\}$, and effect $g \leqslant \varepsilon \in \sigma$, then $\Gamma \vdash \varepsilon_1 <: \varepsilon_2 \cup [x/y]\varepsilon$.

PROOF. We prove this by induction on $size(\varepsilon_1 \cup \varepsilon_2 \cup \{x.g\})$, which is defined in Figure 12. BC If $size(\varepsilon_1 \cup \varepsilon_2 \cup \{x.g\}) = 0$, then x.g cannot have a definition. This case is vacuously true. IS We case on the rule used to derive $\Gamma \vdash \varepsilon_1 <: \varepsilon_2 \cup \{x.g\}$:

- (a) $\Gamma \vdash \varepsilon_1 <: \varepsilon_2 \cup \{x.g\}$ is derived by Subeffect-Subset: If $x.g \notin \varepsilon_1$, then we can use Subeffect-Subset to show that $\Gamma \vdash \varepsilon_1 <: \varepsilon_2 \cup [x/y]\varepsilon$. If $x.g \in \varepsilon_1$, then $\varepsilon_1 = \varepsilon_1' \cup \{x.g\}$, where $\varepsilon_1' \subseteq \varepsilon_2$. Thus, we can use Subeffect-Def-1 to show that $\Gamma \vdash \varepsilon_1' \cup \{x.g\} <: \varepsilon_2 \cup [x/y]\varepsilon$.
- (b) $\Gamma \vdash \varepsilon_1 <: \varepsilon_2 \cup \{x.g\}$ is derived by Subeffect-Upperbound: Then, we have that $\varepsilon_1 = \varepsilon_1' \cup \{z.h\}$, $\Gamma \vdash z : \{y' \Rightarrow \sigma\}$, effect $h = \{\varepsilon'\} \in \sigma$ and that $\Gamma \vdash \varepsilon_1' \cup [z/y']\varepsilon' <: \varepsilon_2 \cup \{x.g\}$. By IH, we have that $\Gamma \vdash \varepsilon_1' \cup [z/y']\varepsilon' <: \varepsilon_2 \cup [x/y]\varepsilon$. Using Subeffect-Upperbound, we have that $\Gamma \vdash \varepsilon_1' \cup \{z.h\} <: \varepsilon_2 \cup [x/y]\varepsilon$.
- (c) $\Gamma \vdash \varepsilon_1 <: \varepsilon_2 \cup \{x.g\}$ is derived by Subeffect-Def-1: If Subeffect-Def-1 uses the effect x.g, then we immediately have that $\Gamma \vdash \varepsilon_1 <: \varepsilon_2 \cup [x/y]\varepsilon$. Otherwise, if Subeffect-Def-1 does not use x.g, then we have that $\varepsilon_2 = \varepsilon_2' \cup \{z.h\}$, $\Gamma \vdash z : \{y' \Rightarrow \sigma\}$, effect $h = \{\varepsilon'\} \in \sigma$, and $\Gamma \vdash \varepsilon_1 <: \varepsilon_2' \cup [z/y']\varepsilon' \cup \{x.y\}$. By IH, we have that $\Gamma \vdash \varepsilon_1 <: \varepsilon_2' \cup [z/y']\varepsilon' \cup [x/y]\varepsilon$. Using Subeffect-Def-1, we have that $\Gamma \vdash \varepsilon_1 <: \varepsilon_2 \cup [x/y]\varepsilon$.
- (d) $\Gamma \vdash \varepsilon_1 <: \varepsilon_2 \cup \{x.g\}$ is derived by Subeffect-Def-2: This case is similar to (b).

Lemma 6 (Reverse of Subeffecting-Def-2). If $\Gamma \vdash \varepsilon_1 \cup \{x.g\} <: \varepsilon_2$, $\Gamma \vdash x : \{y \Rightarrow \sigma\}$ and effect $g = \{\varepsilon\} \in \sigma$, then $\Gamma \vdash \varepsilon_1 \cup [x/y]\varepsilon <: \varepsilon_2$.

PROOF. We prove this by induction on $size(\varepsilon_1 \cup \varepsilon_2 \cup \{x.g\})$, which is defined in Figure 12. BC If $size(\varepsilon_1 \cup \varepsilon_2 \cup \{x.g\}) = 0$, then x.g cannot have a definition. This case is vacuously true. IS We case on the rule used to derive $\Gamma \vdash \varepsilon_1 \cup \{x.g\} <: \varepsilon_2$:

- (a) $\Gamma \vdash \varepsilon_1 \cup \{x.g\} <: \varepsilon_2$ is derived by Subeffect-Subset: Then, $x.g \in \varepsilon_2$. Thus, we can use Subeffect-Def-1 to derive $\Gamma \vdash \varepsilon_1 \cup [x/y]\varepsilon <: \varepsilon_2$.
- (b) $\Gamma \vdash \varepsilon_1 \cup \{x.g\} <: \varepsilon_2$ is derived by Subeffect-Upperbound: If the Subeffect-Upperbound rule uses the effect x.g, then by the premise of Subeffect-Upperbound, we have that $\Gamma \vdash \varepsilon_1 \cup [x/y]\varepsilon <: \varepsilon_2$. If the Subeffect-Upperbound rule does not use the effect x.g, then we have that $\varepsilon_1 = \varepsilon_1' \cup \{z.h\}, \Gamma \vdash z : \{y' \Rightarrow \sigma\}$, effect $h \le \varepsilon' \in \sigma$, and $\Gamma \vdash \varepsilon_1' \cup [z/y']\varepsilon' \cup \{x.g\} <: \varepsilon_2$. By IH, we have that $\Gamma \vdash \varepsilon_1' \cup [z/y']\varepsilon' \cup [x/y]\varepsilon <: \varepsilon_2$. Using Subeffect-Upperbound, we derive $\Gamma \vdash \varepsilon_1 \cup [x/y]\varepsilon <: \varepsilon_2$.

(c) $\Gamma \vdash \varepsilon_1 \cup \{x.g\} <: \varepsilon_2$ is derived by Subeffect-Def-1: Then, we have that $\varepsilon_2 = \varepsilon_2' \cup \{z.h\}$, $\Gamma \vdash z : \{y' \Rightarrow \sigma\}$, effect $h = \{\varepsilon'\} \in \sigma$, and $\Gamma \vdash \varepsilon_1 \cup \{x.g\} <: \varepsilon_2' \cup [z/y']\varepsilon'$. By IH, we have that $\Gamma \vdash \varepsilon_1 \cup [x/y]\varepsilon <: \varepsilon_2' \cup [z/y']\varepsilon'$. Using Subeffect-Def-1, we have that $\Gamma \vdash \varepsilon_1 \cup [x/y]\varepsilon <: \varepsilon_2 \cup \{z.h\}$.

(d) $\Gamma \vdash \varepsilon_1 \cup \{x,g\} <: \varepsilon_2$ is derived by Subeffect-Def-2: This case is similar to (b).

Lemma 7 (Transitivity In Subeffecting). *If* $\Gamma \vdash \varepsilon_1 <: \varepsilon_2 \text{ and } \Gamma \vdash \varepsilon_2 <: \varepsilon_3, \text{ then } \Gamma \vdash \varepsilon_1 <: \varepsilon_3.$

PROOF. We prove this using structural induction on $size(\Gamma, \varepsilon_1 \cup \varepsilon_2 \cup \varepsilon_3)$, which is defined in Figure 12.

- BC Let $size(\Gamma, \varepsilon_1 \cup \varepsilon_2 \cup \varepsilon_3) = 0$. The judgments $\Gamma + \varepsilon_1 <: \varepsilon_2$ and $\Gamma + \varepsilon_2 <: \varepsilon_3$ are derived from Subeffect-Subset. Thus, we have transitivity immediately.
- IS Let $N \ge 0$, assume that $\forall \varepsilon_1, \varepsilon_2, \varepsilon_3$ with $size(\Gamma, \varepsilon_1 \cup \varepsilon_2 \cup \varepsilon_3) \le N$; if $\varepsilon_1 <: \varepsilon_2$ and $\varepsilon_2 <: \varepsilon_3$, then $\varepsilon_1 <: \varepsilon_3$. Let $\Gamma \vdash \varepsilon_1 <: \varepsilon_2$ and $\Gamma \vdash \varepsilon_2 <: \varepsilon_3$ and $size(\Gamma, \varepsilon_1 \cup \varepsilon_2 \cup \varepsilon_3) = N + 1$. We want to show $\varepsilon_1 <: \varepsilon_3$. We case on the rules used to derive $\Gamma \vdash \varepsilon_1 <: \varepsilon_2$ and $\Gamma \vdash \varepsilon_2 <: \varepsilon_3$.
 - (a) $\Gamma \vdash \varepsilon_1 <: \varepsilon_2$ by Subeffect-Subset.
 - (i) $\Gamma \vdash \varepsilon_2 <: \varepsilon_3$ by Subeffect-Subset. Transitivity in this case is trivial.
 - (ii) $\Gamma \vdash \varepsilon_2 <: \varepsilon_3$ by Subeffect-Upperbound. Let $\varepsilon_2 = \varepsilon_2' \cup \{x.g\}$. By Subeffect-Upperbound, we have that $\Gamma \vdash x : \{y \Rightarrow \sigma\}$ effect $g \le \varepsilon \in \sigma$ and $\varepsilon_2' \cup [x/y]\varepsilon <: \varepsilon_3$ There are two cases:
 - (A) If $\{x.g\} \notin \varepsilon_1$, then $\varepsilon_1 \subseteq \varepsilon_2'$. Therefore, $\Gamma \vdash \varepsilon_1 <: \varepsilon_2' \cup [x/y]\varepsilon$. By induction hypothesis, we have that $\Gamma \vdash \varepsilon_1 <: \varepsilon_3$.
 - (B) If $\{x.g\} \in \varepsilon_1$, then $\varepsilon_1 = \varepsilon_1' \cup \{x.g\}$ and $\varepsilon_1' \subseteq \varepsilon_2'$. Thus, we have that $\Gamma \vdash \varepsilon_1' \cup [x/y]\varepsilon <: \varepsilon_2' \cup [x/y]\varepsilon$ by Subeffect-Subset. By IH, we have that $\varepsilon_1' \cup [x/y]\varepsilon <: \varepsilon_3$. Then, we use Subeffect-Upperbound to derive $\varepsilon_1' \cup \{x.g\} <: \varepsilon_3$.
 - (iii) $\Gamma \vdash \varepsilon_2 <: \varepsilon_3$ by Subeffect-Def-1. Let $\varepsilon_3 = \varepsilon_3' \cup \{x.g\}$. We have that $\Gamma \vdash x : \{y \Rightarrow \sigma\}$, effect $g = \{\varepsilon\}$, and $\Gamma \vdash \varepsilon_2 <: \varepsilon_3' \cup [x/y]\varepsilon$. By IH, we have that $\Gamma \vdash \varepsilon_1 <: \varepsilon_3' \cup [x/y]\varepsilon$. Then, we can use Subeffect-Def-1 again to derive $\Gamma \vdash \varepsilon_1 <: \varepsilon_3$.
 - (iv) $\Gamma \vdash \varepsilon_2 <: \varepsilon_3$ by Subeffect-Def-2. The proof is identical to ii.
 - (b) $\Gamma \vdash \varepsilon_1 <: \varepsilon_2$ by Subeffect-Upperbound. Thus, we have that $\varepsilon_1 = \varepsilon_1' \cup \{x.g\}$. $\Gamma \vdash x : \{y \Rightarrow \sigma\}$, effect $g = \{\varepsilon\}$, and $\Gamma \vdash \varepsilon_1' \cup [x/y]\varepsilon <: \varepsilon_2$. Using IH, we have that $\Gamma \vdash \varepsilon_1' \cup [x/y]\varepsilon <: \varepsilon_3$. Using Subeffect-Upperbound again, we have that $\Gamma \vdash \varepsilon_1 <: \varepsilon_3$.
 - (c) $\Gamma \vdash \varepsilon_1 <: \varepsilon_2$ by Subeffect-Def-1. Therefore, we let $\varepsilon_2 = \varepsilon_2' \cup \{x.g\}$, $\Gamma \vdash x : \{y \Rightarrow \sigma\}$, and $effect g = \{\varepsilon\} \in \sigma$. By premise of Subeffect-Def-1, we have that $\Gamma \vdash \varepsilon_1 <: [x/y]\varepsilon \cup \varepsilon_2'$. Since $\Gamma \vdash \varepsilon_2 <: \varepsilon_3$, we have that $\Gamma \vdash \varepsilon_2' \cup \{x.g\} <: \varepsilon_3$.
 - (i) $\Gamma \vdash \varepsilon_2' \cup \{x.g\} <: \varepsilon_3$ by Subeffect-Subset. Then, we have that $\varepsilon_3 = \varepsilon_3' \cup \{x.g\}$ and $\varepsilon_2' \subseteq \varepsilon_3'$. Therefore, we have that $\varepsilon_2' \cup [x/y]\varepsilon \subseteq \varepsilon_3' \cup [x/y]\varepsilon$. Therefore, $\Gamma \vdash \varepsilon_2' \cup [x/y]\varepsilon <: \varepsilon_3' \cup [x/y]\varepsilon$. By IIH, we have that $\Gamma \vdash \varepsilon_1 <: \varepsilon_3' \cup [x/y]\varepsilon$. By Subeffect-Def-1,we have that $\Gamma \vdash \varepsilon_1 <: \varepsilon_3' \cup \{x.g\} = \varepsilon_3$.
 - (ii) $\Gamma \vdash \varepsilon_2 <: \varepsilon_3$ by Subeffect-Upperbound. Since the effect $\{x.g\}$ is used by Subeffect-Def-1, it is not used by the rule Subeffect-Upperbound. Let $\varepsilon_2 = \varepsilon_2'' \cup \{x.g\} \cup \{z.h\}$. By Subeffect-Def-1, we have that

5:38 D. Melicher et al.

 $\begin{array}{l} \Gamma \vdash \varepsilon_1 <: \varepsilon_2'' \cup [x/y] \varepsilon \cup \{z.h\}. \ \text{By Subeffect-Upperbound, we have that} \ \Gamma \vdash z : \{y' \Rightarrow \sigma'\}, \\ \text{effect} \quad h \leqslant \varepsilon' \in \sigma', \ \text{and} \ \Gamma \vdash \varepsilon_2'' \cup \{x.g\} \cup [z/y'] \varepsilon' <: \varepsilon_3. \ \text{By Lemma 5 and} \\ \Gamma \vdash \varepsilon_1 <: \varepsilon_2'' \cup [x/y] \varepsilon \cup \{z.h\}, \ \text{we have that} \ \Gamma \vdash \varepsilon_1 <: \varepsilon_2'' \cup [x.y] \varepsilon \cup [z/y'] \varepsilon'. \ \text{By Lemma} \\ 6, \ \text{and} \ \Gamma \vdash \varepsilon_2'' \cup \{x.g\} \cup [z/y'] \varepsilon' <: \varepsilon_3, \ \text{we have that} \ \Gamma \vdash \varepsilon_2'' \cup [x/y] \varepsilon \cup [z/y'] \varepsilon' <: \varepsilon_3. \\ \text{Therefore, we use IH to derive} \ \Gamma \vdash \varepsilon_1 <: \varepsilon_3. \end{array}$

- (iii) $\Gamma \vdash \varepsilon_2 <: \varepsilon_3$ by Subeffect-Def-1. Therefore, let $\varepsilon_3 = \varepsilon_3' \cup \{z.h\}$, $\Gamma \vdash z : \{y \Rightarrow \sigma'\}$, and $effect\ h = \{\varepsilon'\} \in \sigma'$. We have that $\Gamma \vdash \varepsilon_2 <: \varepsilon_3' \cup \{z.h\}$. By premise of Subeffect-Def-1, we have that $\Gamma \vdash \varepsilon_2 <: [z/y]\varepsilon' \cup \varepsilon_3'$. By IH, we have that $\Gamma \vdash \varepsilon_1 <: [z/y]\varepsilon' \cup \varepsilon_3'$. Using Subeffect-Def-1, we derive that $\Gamma \vdash \varepsilon_1 <: \varepsilon_3$.
- (iv) $\Gamma \vdash \varepsilon_2 <: \varepsilon_3$ by Subeffect-Def-2. This case is identical to c (ii).
- (d) $\Gamma \vdash \varepsilon_1 <: \varepsilon_2$ by Subeffect-Def-2. This case is identical to (b).
- (e) $\Gamma \vdash \varepsilon_1 <: \varepsilon_2$ by Subeffect-Lowerbound. This case is identical to (c).

Lemma 8 (Substitution In Types). If Γ , $z:\tau\vdash\tau_1<:\tau_2$ and $\Gamma\mid\Sigma\vdash l:\{\}$ $[l/z]\tau$, then $\Gamma\vdash[l/z]\tau_1<:[l/z]\tau_2$. Furthermore, if Γ , $z:\tau\vdash\sigma_1<:\sigma_2$ and $\Gamma\mid\Sigma\vdash l:\{\}$ $[l/z]\tau$, then $\Gamma\vdash[l/z]\sigma_1<:[l/z]\sigma_2$.

PROOF. The proof is by simultaneous induction on a derivation of Γ , $z: \tau \vdash \tau_1 <: \tau_2$ and Γ , $z: \tau \vdash \sigma_1 <: \sigma_2$. For a given derivation, we proceed by cases on the final typing rule used in the derivation:

Case S-Refl.1: $\tau_1 = \tau_2$, and the desired result is immediate.

<u>Case S-Trans</u>: By inversion on S-Trans, we get that Γ , $z:\tau \vdash \tau_1 <: \tau_2$ and Γ , $z:\tau \vdash \tau_2 <: \tau_3$. By the induction hypothesis, $\Gamma \vdash [l/z]\tau_1 <: [l/z]\tau_2$ and $\Gamma \vdash [l/z]\tau_2 <: [l/z]\tau_3$. Then, by S-Trans, $\Gamma \vdash [l/z]\tau_1 <: [l/z]\tau_3$.

<u>Case S-Perm</u>: $\tau_1 = \{x \Rightarrow \sigma_i^{i \in 1..n}\}$ and $\tau_2 = \{x \Rightarrow \sigma_i'^{i \in 1..n}\}$. Substitution preserves the permutation relations; thus, $[l/z]\{x \Rightarrow \sigma_i^{i \in 1..n}\}$ is a permutation of $[l/z]\{x \Rightarrow \sigma_i'^{i \in 1..n}\}$. Then, by S-Perm, $\Gamma \vdash [l/z]\{x \Rightarrow \sigma_i^{i \in 1..n}\} <: [l/z]\{x \Rightarrow \sigma_i'^{i \in 1..n}\}$.

 $\underline{\textit{Case S-Width}}: \tau_1 = \{x \Rightarrow \sigma_i^{i \in 1..n+k}\} \text{ and } \tau_2 = \{x \Rightarrow \sigma_i^{i \in 1..n}\}, \text{ and the desired result is immediate.}$

<u>Case S-Depth</u>: $\tau_1 = \{x \Rightarrow \sigma_i^{i \in 1...n}\}$ and $\tau_2 = \{x \Rightarrow {\sigma'}_i^{i \in 1...n}\}$. By inversion on S-Depth, we get that $\forall i, \ \Gamma, \ x : \{x \Rightarrow \sigma_i^{i \in 1...n}\}, \ z : \tau \vdash \sigma_i <: \sigma'_i$. By the induction hypothesis, $\forall i, \ \Gamma, \ x : \{x \Rightarrow \sigma_i^{i \in 1...n}\}$ $\vdash [l/z]\sigma_i <: [l/z]\sigma_i'$. Then, by S-Depth, $\Gamma \vdash [l/z]\{x \Rightarrow \sigma_i^{i \in 1...n}\} <: [l/z]\{x \Rightarrow \sigma_i'^{i \in 1...n}\}$.

Case S-Refl2: $\sigma_1 = \sigma_2$, and the desired result is immediate.

<u>Case S-Def:</u> $\sigma_1 = \text{def } m(x:\tau_1): \{\varepsilon_1\} \ \tau_2 \text{ and } \sigma_2 = \text{def } m(x:\tau_1'): \{\varepsilon_2\} \ \tau_2'.$ By inversion on S-Def, we get that Γ , $z:\tau \vdash \tau_1' <: \tau_1$, Γ , $z:\tau \vdash \tau_2 <: \tau_2'$, Γ , $z:\tau \vdash \varepsilon_1 <: \varepsilon_2$. By the induction hypothesis, $\Gamma \vdash [l/z]\tau_1' <: [l/z]\tau_1$ and $\Gamma \vdash [l/z]\tau_2 <: [l/z]\tau_2'$. By Lemma 9, $\Gamma \vdash [l/z]\varepsilon_1 <: [l/z]\varepsilon_2$. Then, by S-Def, $\Gamma \vdash [l/z](\text{def } m(x:\tau_1): \{\varepsilon_1\} \ \tau_2) <: [l/z](\text{def } m(x:\tau_1'): \{\varepsilon_2\} \ \tau_2').$

Case S-Effect: $\sigma_1 = \text{effect } q = \{\varepsilon\}$ and $\sigma_2 = \text{effect } q$, and the desired result is immediate.

Thus, substituting terms in types preserves the subtyping relationship.

Lemma 9 (Substitution in Expressions and Effects). If Γ , $z:\tau'\mid \Sigma \vdash e:\{\varepsilon\}$ τ and $\Gamma\mid \Sigma \vdash l:\{\}$ $[l/z]\tau'$, then $\Gamma\mid \Sigma \vdash [l/z]e:\{[l/z]\varepsilon\}$ $[l/z]\tau$.

5:39

And if $\Gamma, z : \tau' \mid \Sigma \vdash \varepsilon_1 <: \varepsilon_2 \text{ and } \Gamma \mid \Sigma \vdash l : \{\}[l/z]\tau, \text{ then } \Gamma \mid \Sigma \vdash [l/z]\varepsilon_1 <: [l/z]\varepsilon_2.$

And if Γ , $z : \tau' \mid \Sigma \vdash d : \sigma$ and $\Gamma \mid \Sigma \vdash l : \{\} [l/z]\tau'$, then $\Gamma \mid \Sigma \vdash [l/z]d : [l/z]\sigma$.

Furthermore, if $\Gamma, z : \tau' \mid \Sigma \vdash \varepsilon \ w f$, then $\Gamma \mid \Sigma \vdash [l/z] \varepsilon \ w f$.

PROOF. The proof is by simultaneous induction on a derivation of Γ , $z:\tau'\mid \Sigma\vdash e:\{\varepsilon\}$ τ , Γ , $z:\tau'\mid \Sigma\vdash d:\sigma$, Γ , $z:\tau'\mid \Sigma\vdash \varepsilon_1$, and Γ , $z:\tau'\mid \Sigma\vdash \varepsilon wf$. For a given derivation, we proceed by cases on the final typing rule used in the derivation:

<u>Case T-Var</u>: e = x, and by inversion on T-Var, we get that $x : \tau \in (\Gamma, z : \tau')$. There are two subcases to consider depending on whether x is z or another variable. If x = z, then $\lfloor l/z \rfloor x = l$ and $\tau = \tau'$. The required result is then $\Gamma \mid \Sigma \vdash l : \{\} \lfloor l/z \rfloor \tau'$, which is among the assumptions of the lemma. Otherwise, $\lfloor l/z \rfloor x = x$, and the desired result is immediate.

<u>Case T-Field:</u> $e = e_1.f$, and by inversion on T-Field, we get that Γ , $z : \tau' \mid \Sigma \vdash e_1 : \{\varepsilon\} \{x \Rightarrow \overline{\sigma}\}$ and var $f : \tau \in \overline{\sigma}$. By the induction hypothesis, $\Gamma \mid \Sigma \vdash [l/z]e_1 : \{[l/z]\varepsilon\} [l/z]\{x \Rightarrow \overline{\sigma}\}$ and var $f : [l/z]\tau \in [l/z]\overline{\sigma}$. Then, by T-Field, $\Gamma \mid \Sigma \vdash ([l/z]e_1).f : \{[l/z]\varepsilon\} [l/z]\tau$, that is, $\Gamma \mid \Sigma \vdash [l/z](e_1.f) : \{[l/z]\varepsilon\} [l/z]\tau$.

Case T-Loc: e = l, [l/z]l = l, and the desired result is immediate.

5:40 D. Melicher et al.

<u>Case T-Sub</u>: $e = e_1$ and, by inversion on T-Sub, we get that $\Gamma, z : \tau' \mid \Sigma \vdash e_1 : \{\varepsilon_1\}\tau_1, \Gamma, z : \tau' \mid \Sigma \vdash \tau_1 <: \tau_2$, and $\Gamma, z : \tau' \mid \Sigma \vdash \varepsilon_1 <: \varepsilon_2$. By induction hypothesis, we have that $\Gamma \mid \Sigma \vdash [l/z]e_1 : \{[l/z]\varepsilon_1\}[l/z]\tau_1, \Gamma \mid \Sigma \vdash [l/z]\tau_1 <: [l/z]\tau_2$, and $\Gamma \mid \Sigma \vdash [l/z]\varepsilon_1 <: [l/z]\varepsilon_2$. Then, by T-sub, $\Gamma \mid \Sigma \vdash [l/z]e_1 : \{[l/z]\varepsilon_2\}[l/z]\tau_2$.

<u>Case DT-Def:</u> By inversion, we have that $\Gamma, z : \tau, x : \tau_1 \mid \Sigma \vdash e : \{\varepsilon'\} \tau_2, \Gamma, z : \tau, x : \tau_1 \mid \Sigma \vdash \varepsilon \text{ wf},$ and $\Gamma, z : \tau \mid \Sigma \vdash \varepsilon' <: \varepsilon$. By IH, we have that $\Gamma, x : [l/z]\tau_1 \mid \Sigma \vdash [l/z]e : \{[l/z]\varepsilon'\} [l/z]\tau_2,$ $\Gamma, x : [l/z]\tau_1 \mid \Sigma \vdash [l/z]\varepsilon \text{ wf},$ and $\Gamma \mid \Sigma \vdash [l/z]\varepsilon' <: [l/z]\varepsilon$. By DT-Def, we have that $\Gamma \mid \Sigma \vdash \text{def } m(x : [l/z]\tau_1) : \{[l/z]\varepsilon\} [l/z]\tau_2 = [l/z]e : \text{def } m(x : [l/z]\tau_1) : \{[l/z]\varepsilon\} [l/z]\tau_2.$

Case DT-VAR: $d = \text{var } f : \tau = n$, and, by definition of n, there are two subcases:

<u>Subcase n is x:</u> In this case, $d = \text{var } f : \tau = x$ and, by inversion on DT-VAR, we get that Γ , $z : \tau' \mid \Sigma \vdash x : \{\}$ τ . There are two subcases to consider depending on whether x is z or another variable. If x = z, then by the induction hypothesis, $\Gamma \mid \Sigma \vdash [l/z]x : \{\}$ $[l/z]\tau$, which yields $\Gamma \mid \Sigma \vdash l : \{\}$ $[l/z]\tau$ and $\tau = \tau'$. Thus, $\Gamma \mid \Sigma \vdash \text{var } f : [l/z]\tau = l : \text{var } f : [l/z]\tau$, that is, $\Gamma \mid \Sigma \vdash [l/z](\text{var } f : \tau = l) : [l/z](\text{var } f : \tau)$, as required. If $x \neq z$, then $\Gamma \mid \Sigma \vdash [l/z]x : \{\}$ $[l/z]\tau$ yields $\Gamma \mid \Sigma \vdash x : \{\}$ $[l/z]\tau$. Thus, $\Gamma \mid \Sigma \vdash \text{var } f : [l/z]\tau = x : \text{var } f : [l/z]\tau$, that is, $\Gamma \mid \Sigma \vdash [l/z]$ (var $f : \tau = x$) : $[l/z](\text{var } f : \tau)$, as required.

<u>Subcase n is l:</u> In this case, $d = \text{var } f : \tau = l$, that is, the field is resolved to a location l. This is not affected by the substitution, and the desired result is immediate.

<u>Case DT-Effect</u>: By IH, we have that $\Gamma \mid \Sigma \vdash [l/z]\varepsilon$ wf. We use DT-Effect to derive $\Gamma \mid \Sigma \vdash \text{effect } q = \{[l/z]\varepsilon\} : \text{effect } q = \{[l/z]\varepsilon\}.$

<u>Case Subeffect-Subset:</u> By inversion, we have that $\varepsilon_1 \subseteq \varepsilon_2$. Thus, $[l/z]\varepsilon_1 \subseteq [l/z]\varepsilon_2$. By Subeffect-Subset, we have that $\Gamma \mid \Sigma \vdash [l/z]\varepsilon_1 <: [l/z]\varepsilon_2$.

<u>Case Subeffect-Upperbound:</u> By inversion, we have that $\varepsilon_1 = \varepsilon_1' \cup \{x.g\}$, $\Gamma, z : \tau \mid \Sigma \vdash x : \{y \Rightarrow \sigma\}$, $effect g \leq \{\varepsilon\} \in \sigma$, and $\Gamma, z : \tau \mid \Sigma \vdash \varepsilon_1' \cup [x/y]\varepsilon <: \varepsilon_2$. By IH, we have that $\Gamma \mid \Sigma \vdash [l/z]\varepsilon_1' \cup [l/z][x/y]\varepsilon <: [l/z]\varepsilon_2$. Since y is a free variable, we select y such that $x \neq y$ and $y \neq z$. We case on if z = x:

- (1) If $z \neq x$, then we can swap the order of the substitutions on $\varepsilon \Gamma \mid \Sigma \vdash [l/z]\varepsilon'_1 \cup [x/y][l/z]\varepsilon <: [l/z]\varepsilon_2$. Using a substitution lemma for typing on $\Gamma, z : \tau \mid \Sigma \vdash x : \{y \Rightarrow \sigma\}$, we have that $\Gamma \mid \Sigma \vdash x : \{y \Rightarrow [l/z]\sigma\}$, effect $g \leq [l/z]\varepsilon \in [l/z]\sigma$. Using Subeffect-Upperbound, we have that $\Gamma \mid \Sigma \vdash [l/z]\varepsilon'_1 \cup \{x.g\} <: [l/z]\varepsilon_2$, which is equivalent to $\Gamma \mid \Sigma \vdash [l/z]\varepsilon_1 <: [l/z]\varepsilon_2$.
- (2) If z = x, then we have that $\Gamma \mid \Sigma \vdash [l/z]\varepsilon'_1 \cup [l/x, y]\varepsilon <: [l/z]\varepsilon_2$, which is equivalent to $\Gamma \mid \Sigma \vdash [l/z]\varepsilon'_1 \cup [l/y][l/z]\varepsilon <: [l/z]\varepsilon_2$. We case on the derivation of $\Gamma, z : \tau \mid \Sigma \vdash z : \{y \Rightarrow \sigma\}$. (a) (T-Var)

$$\frac{z:\tau\in\Gamma,z:\tau}{\Gamma,z:\tau\mid\Sigma\vdash z:\tau}$$

Thus, $\tau = \{y \Rightarrow \sigma\}$. By our assumption, we have that $\Gamma \mid \Sigma \vdash l : \{y \Rightarrow [l/z]\sigma\}$. Since effect $g \leqslant \varepsilon \in \sigma$, we have that effect $g \leqslant [l/z]\varepsilon \in [l/z]\sigma$. Therefore, we can use Subeffect-Upperbound on $\{l.g\}$ to derive $\Gamma \mid \Sigma \vdash [l/z]\varepsilon'_1 \cup \{l.g\} <: [l/z]\varepsilon_2$, which is equivalent to $\Gamma \mid \Sigma \vdash [l/z]\varepsilon_1 <: [l/z]\varepsilon_2$.

(b) (T-Sub)

$$\frac{\Gamma, z : \tau \mid \Sigma \vdash z : \tau_1 \quad \Gamma, z : \tau \mid \Sigma \vdash \tau_1 <: \{y \Rightarrow \sigma\}}{\Gamma, z : \tau \mid \Sigma \vdash z : \{y \Rightarrow \sigma\}}$$

Note that we introduced a new type τ_1 that z can be ascribed to. The judgment $\Gamma, z: \tau \mid \Sigma \vdash z: \tau_1$ can be derived by T-Sub, which introduces a new type τ_2 such that $\Gamma, z: \tau \mid \Sigma \vdash \tau_2 <: \tau_1$, or T-Var, which shows that $\tau_1 = \tau$. Therefore, if we follow the derivation tree, we get a chain relation $\Gamma, z: \tau \mid \Sigma \vdash \tau_1 <: \{y \Rightarrow \sigma\}, \Gamma, z: \tau \mid \Sigma \vdash \tau_2 <: \tau_1, \ldots, \Gamma, z: \tau \mid \Sigma \vdash \tau <: \tau_n$. We can apply IH on these judgments so that we have a chain $\Gamma \mid \Sigma \vdash [l/z]\tau_1 <: \{y \Rightarrow [l/z]\sigma\}, \Gamma \mid \Sigma \vdash [l/z]\tau_2 <: [l/z]\tau_1 \ldots, \Gamma \mid \Sigma \vdash [l/z]\tau <: [l/z]\tau_n$. By transitivity of subtyping, we have that $\Gamma \mid \Sigma \vdash [l/z]\tau <: \{y \Rightarrow [l/z]\sigma\}$. Thus, we have that $\Gamma \mid \Sigma \vdash l: \{y \Rightarrow [l/z]\sigma\}$. The rest of the proof is similar to case (a).

<u>Case Subeffect-Def-1:</u> By inversion, we have that $\varepsilon_2 = \varepsilon_2' \cup \{x.g\}$, $\Gamma, z : \tau \mid \Sigma \vdash x : \{y \Rightarrow \sigma\}$, effect $g = \{\varepsilon\} \in \sigma$, and $\Gamma, z : \tau \mid \Sigma \vdash \varepsilon_1 <: \varepsilon_2' \cup [x/y]\varepsilon$. By IH, we have that $\Gamma \mid \Sigma \vdash [l/z]\varepsilon_1 <: [l/z]\varepsilon_2' \cup [l/z][x/y]\varepsilon$. Since y is a free variable, we can select y such that $y \neq x$ and $y \neq z$. We case on if x = z:

- (1) If $z \neq x$, then $\Gamma \mid \Sigma \vdash [l/z]\varepsilon_1 <: [l/z]\varepsilon_2' \cup [x/y][l/z]\varepsilon$. By substitution lemma for typing, we have that $\Gamma \mid \Sigma \vdash x : \{y \Rightarrow [l/z]\sigma\}$, effect $g = [l/z]\varepsilon \in [l/z]\sigma$. Using Subeffect-Def-1, we have that $\Gamma \mid \Sigma \vdash [l/z]\varepsilon_1 <: [l/z]\varepsilon_2' \cup \{x.g\}$, which is equivalent to $\Gamma \mid \Sigma \vdash [l/z]\varepsilon_1 <: [l/z]\varepsilon_2$.
- (2) If z=x, then we have that $\Gamma\mid \Sigma\vdash [l/z]\varepsilon_1<:[l/z]\varepsilon_2'\cup [l/x,y]\varepsilon$, which is equivalent to $\Gamma\mid \Sigma\vdash [l/z]\varepsilon_1<:[l/z]\varepsilon_2'\cup [l/y][l/z]\varepsilon$.

We case on the derivation of $\Gamma, z : \tau \mid \Sigma \vdash z : \{y \Rightarrow \sigma\}$.

(a) (T-Var)

$$\frac{z:\tau\in\Gamma,z:\tau}{\Gamma,z:\tau\mid\Sigma\vdash z:\tau}$$

Thus, $\tau = \{y \Rightarrow \sigma\}$. By our assumption, we have that $\Gamma \mid \Sigma \vdash l : \{y \Rightarrow [l/z]\sigma\}$. Since effect $g = \{\varepsilon\} \in \sigma$, we have that effect $g = \{[l/z]\varepsilon\} \in [l/z]\sigma$. Therefore, we can use Subeffect-Def-1 on $\{l.g\}$ to derive $\Gamma \mid \Sigma \vdash [l/z]\varepsilon_1 <: [l/z]\varepsilon_2' \cup \{l.g\}$, which is equivalent to $\Gamma \mid \Sigma \vdash [l/z]\varepsilon_1 <: [l/z]\varepsilon_2$.

(b) (T-Sub)

$$\frac{\Gamma, z : \tau \mid \Sigma \vdash z : \tau_1 \quad \Gamma, z : \tau \mid \Sigma \vdash \tau_1 <: \{y \Rightarrow \sigma\}}{\Gamma, z : \tau \mid \Sigma \vdash z : \{y \Rightarrow \sigma\}}$$

Note that we introduced a new type τ_1 that z can be ascribed to. The judgment $\Gamma, z: \tau \mid \Sigma \vdash z: \tau_1$ can be derived by T-Sub, which introduces a new type τ_2 such that $\Gamma, z: \tau \mid \Sigma \vdash \tau_2 <: \tau_1$, or T-Var, which shows that $\tau_1 = \tau$. Therefore, if we follow the derivation tree, we get a chain relation $\Gamma, z: \tau \mid \Sigma \vdash \tau_1 <: \{y \Rightarrow \sigma\}, \Gamma, z: \tau \mid \Sigma \vdash \tau_2 <: \tau_1, \ldots, \Gamma, z: \tau \mid \Sigma \vdash \tau <: \tau_n$. We can apply IH on these judgments so that we have a chain $\Gamma \mid \Sigma \vdash [l/z]\tau_1 <: \{y \Rightarrow [l/z]\sigma\}, \Gamma \mid \Sigma \vdash [l/z]\tau_2 <: [l/z]\tau_1, \ldots, \Gamma \mid \Sigma \vdash [l/z]\tau <: [l/z]\tau_n$. By transitivity of subtyping, we have that $\Gamma \mid \Sigma \vdash [l/z]\tau <: \{y \Rightarrow [l/z]\sigma\}$. Thus, we have that $\Gamma \mid \Sigma \vdash l: \{y \Rightarrow [l/z]\sigma\}$. The rest of the proof is similar to case (a).

5:42 D. Melicher et al.

Case Subeffect-Def-2: This case is identical to Case Subeffect-Upperbound.

Case Subeffect-Lowerbound: This case is identical to Case Subeffect-Def-1.

<u>Case WF-Effect</u>: Let $n_i.g_j \in \varepsilon$ be arbitrary. By inversion, we have that $\Gamma, z : \tau \mid \Sigma \vdash n_i : \{\}\{y_i \Rightarrow \overline{\sigma_i}\}$. and the effect declaration of g_j is in $\overline{\sigma_i}$. By IH, we have that $\Gamma \mid \Sigma \vdash [l/z]n_i : \{\}\{y_i \Rightarrow [l/z]\overline{\sigma_i}\}$ and the effect declaration of g_j is in $\overline{\sigma_i}$. Thus, we have that $[l/z]\varepsilon \ wf$ by WF-Effect.

Thus, substituting terms in a well-typed expression preserves the typing.

B.2 Proof of Theorem 3 (Preservation)

If $\Gamma \mid \Sigma \vdash e : \{\varepsilon\} \ \tau$, $\mu : \Sigma$, and $\langle e \mid \mu \rangle \longrightarrow \langle e' \mid \mu' \rangle$, then $\exists \Sigma' \supseteq \Sigma$, $\mu' : \Sigma'$, $\exists \varepsilon'$, such that $\Gamma \vdash \varepsilon' <: \varepsilon$ and $\Gamma \mid \Sigma' \vdash e' : \{\varepsilon'\} \ \tau$.

PROOF. The proof is by induction on a derivation of $\Gamma \mid \Sigma \vdash e : \{\varepsilon\} \mid \tau$. At each step of the induction, we assume that the desired property holds for all subderivations and proceed by case analysis on the final rule in the derivation. Since we assumed that $\langle e \mid \mu \rangle \longrightarrow \langle e' \mid \mu' \rangle$ and there are no evaluation rules corresponding to variables or locations, the cases when e is a variable (T-VAR) or a location (T-Loc) cannot arise. For the other cases, we argue as follows:

<u>Case T-New:</u> $e = \text{new}(x \Rightarrow \overline{d})$ and, by inversion on T-New, we get that $\forall i, d_i \in \overline{d}, \sigma_i \in \overline{\sigma}, \Gamma, x : \{x \Rightarrow \overline{\sigma}\} \mid \Sigma \vdash d_i : \sigma_i$. The store changes from μ to $\mu' = \mu, l \mapsto \{x \Rightarrow \overline{d}\}$, that is, the new store is the old store augmented with a new mapping for the location l, which was not in the old store $(l \notin dom(\mu))$. From the premise of the theorem, we know that $\mu : \Sigma$, and by the induction hypothesis, all expressions of Γ are properly allocated in Σ . Then, by T-Store, we have that $\mu, l \mapsto \{x \Rightarrow \overline{d}\} : \Sigma, l : \{x \Rightarrow \overline{\sigma}\}$, which implies that $\Sigma' = \Sigma, l : \{x \Rightarrow \overline{\sigma}\}$. Finally, by T-Loc, $\Gamma \mid \Sigma \vdash l : \{\} \{x \Rightarrow \overline{\sigma}\}$ and $\varepsilon' = \varnothing = \varepsilon$. Thus, the right-hand side is well typed.

<u>Case T-METHOD</u>: $e = e_1.m(e_2)$ and, by the definition of the evaluation relation, there are two subcases:

<u>Subcase E-Congruence:</u> In this case, either $\langle e_1 \mid \mu \rangle \longrightarrow \langle e'_1 \mid \mu' \rangle$ or e_1 is a value and $\langle e_2 \mid \mu \rangle \longrightarrow \langle e'_2 \mid \mu' \rangle$. Then, the result follows from the induction hypothesis and T-METHOD.

<u>Subcase E-Method</u>: In this case, both e_1 and e_2 are values, namely, locations l_1 and l_2 , respectively. Then, by inversion on T-Method, we get that $\Gamma \mid \Sigma \vdash e_1 : \{\varepsilon_1\} \mid \{x \Rightarrow \overline{\sigma}\}$, def $m(y : \tau_2) : \{\varepsilon_3\} \mid \tau_1 \in \overline{\sigma}$, $\Gamma \mid \Sigma \vdash [e_1/x][e_2/y]\varepsilon_3 \mid f$, $\Gamma \mid \Sigma \vdash e_2 : \{\varepsilon_2\} \mid [e_1/x]\tau_2$, and $\varepsilon = \varepsilon_1 \cup \varepsilon_2 \cup [e_1/x][e_2/y]\varepsilon_3$. The store μ does not change and, since T-Store has been applied throughout, the store is well typed. Thus, $\Gamma \mid \Sigma \vdash \text{def } m(x : \tau_1) : \{\varepsilon\} \mid \tau_2 = e : \text{def } m(x : \tau_1) : \{\varepsilon\} \mid \tau_2$. Then, by inversion on DT-Def, we know that Γ , $x : \tau_1 \mid \Sigma \vdash e : \{\varepsilon'\} \mid \tau_2$ and Γ , $x : \tau_1 \mid \Sigma \vdash \varepsilon' < : \varepsilon$. Finally, by the subsumption lemma, substituting locations for variables in e preserves its type. Therefore, the right-hand side is well typed.

<u>Case T-FIELD:</u> $e = e_1.f$ and, by the definition of the evaluation relation, there are two subcases: <u>Subcase E-Congruence:</u> In this case, $\langle e_1 \mid \mu \rangle \longrightarrow \langle e_1' \mid \mu' \rangle$, and the result follows from the induction hypothesis and T-FIELD.

<u>Subcase E-Field:</u> In this case, e_1 is a value, that is, a location l. Then, by inversion on T-Field, we have that $\Gamma \mid \Sigma \vdash l : \{\varepsilon\} \{x \Rightarrow \overline{\sigma}\}$, where $\varepsilon = \emptyset$, and var $f : \tau \in \overline{\sigma}$. The

store μ does not change and, since T-Store has been applied throughout, the store is well typed. Thus, $\Gamma \mid \Sigma \vdash \text{var } f : \tau = l_1 : \text{var } f : \tau$. Then, by inversion on DT-VARL, we know that $\Gamma \mid \Sigma \vdash l_1 : \{\} \tau$ and $\varepsilon' = \emptyset = \varepsilon$, and the right-hand side is well typed.

<u>Case T-Assign:</u> $e = (e_1.f = e_2)$ and, by the definition of the evaluation relation, there are two subcases:

<u>Subcase E-Congruence:</u> In this case, either $\langle e_1 \mid \mu \rangle \longrightarrow \langle e_1' \mid \mu' \rangle$ or e_1 is a value and $\langle e_2 \mid \mu \rangle \longrightarrow \langle e_2' \mid \mu' \rangle$. Then, the result follows from the induction hypothesis and T-Assign.

<u>Subcase E-Assign:</u> In this case, both e_1 and e_2 are values, namely, locations l_1 and l_2 , respectively. Then, by inversion on T-Assign, we get that $\Gamma \mid \Sigma \vdash l_1 : \{\varepsilon_1\} \mid \{x \Rightarrow \overline{\sigma}\}$, var $f: \tau \in \overline{\sigma}$, $\Gamma \mid \Sigma \vdash l_2 : \{\varepsilon_2\} \mid \tau$, and $\varepsilon = \varepsilon_1 = \varepsilon_2 = \emptyset$. The store changes as follows: $\mu' = [l_1 \mapsto \{x \Rightarrow \overline{d}'\} / l_1 \mapsto \{x \Rightarrow \overline{d}\}] \mu$, where $\overline{d}' = [\text{var } f: \tau = l_2/\text{var } f: \tau = l] \overline{d}$. However, since T-Store has been applied throughout and the substituted location has the type expected by T-Store, the new store is well typed (as well as the old store). Thus, $\Gamma \mid \Sigma \vdash \text{var } f: \tau = l_2: \text{var } f: \tau$. Then, by inversion on DT-Varl, we know that $\Gamma \mid \Sigma \vdash l_2: \{\} \mid \tau \text{ and } \varepsilon' = \emptyset$, and the right-hand side is well typed.

Case T-SuB: The result follows directly from the induction hypothesis.

Thus, the program written in this language is always well typed.

B.3 Proof of Theorem 4 (Progress)

If $\emptyset \mid \Sigma \vdash e : \{\varepsilon\} \mid \tau$ (i.e., *e* is a closed, well-typed expression), then either

- (1) e is a value (i.e., a location) or
- (2) $\forall \mu$ such that $\mu : \Sigma, \exists e', \mu'$ such that $\langle e \mid \mu \rangle \longrightarrow \langle e' \mid \mu' \rangle$.

PROOF. The proof is by induction on the derivation of $\Gamma \mid \Sigma \vdash e : \{\varepsilon\} \tau$, with a case analysis on the last typing rule used. The case in which e is a variable (T-VAR) cannot occur and the case in which e is a location (T-Loc) is immediate since, in that case, e is a value. For the other cases, we argue as follows:

<u>Case T-New:</u> $e = \text{new}(x \Rightarrow \overline{d})$, and by E-New, e can make a step of evaluation if the new expression is closed and there is a location available that is not in the current store μ . From the premise of the theorem, we know that the expression is closed and that there are infinitely many available new locations. Therefore, e indeed can take a step and become a value (i.e., a location l). Then, the new store μ' is $\mu, l \mapsto \{x \Rightarrow \overline{d}\}$, and all of the declarations in \overline{d} are mapped in the new store.

<u>Case T-METHOD:</u> $e = e_1.m(e_2)$ and, by the induction hypothesis applied to $\Gamma \mid \Sigma \vdash e_1 : \{\varepsilon_1\}$ $\{x \Rightarrow \overline{\sigma}\}$, either e_1 is a value or else it can make a step of evaluation. Similarly, by the induction hypothesis applied to $\Gamma \mid \Sigma \vdash e_2 : \{\varepsilon_2\} [e_1/x]\tau_2$, either e_2 is a value or else it can make a step of evaluation. Then, there are two subcases:

Subcase $\langle e_1 \mid \mu \rangle \longrightarrow \langle e_1' \mid \mu' \rangle$ or e_1 is a value and $\langle e_2 \mid \mu \rangle \longrightarrow \langle e_2' \mid \mu' \rangle$: If e_1 can take a step or if e_1 is a value and e_2 can take a step, then rule E-Congruence applies to e, and e can take a step.

Subcase e_1 and e_2 are values: If both e_1 and e_2 are values, that is, they are locations l_1 and l_2 , respectively, then by inversion on T-Method, we have that $\Gamma \mid \Sigma \vdash l_1 : \{\varepsilon_1\} \{x \Rightarrow \overline{\sigma}\}$ and def $m(y : \tau_2) : \{\varepsilon_3\} \tau_1 \in \overline{\sigma}$. By inversion on T-Loc, we know that the store contains an appropriate

5:44 D. Melicher et al.

mapping for the location l_1 and, since T-Store has been applied throughout, the store is well typed and $l_1 \mapsto \{x \Rightarrow \overline{d}\} \in \mu$ with def $m(y : \tau_1) : \{\varepsilon_3\} \ \tau_2 = e \in \overline{d}$. Therefore, the rule E-Method applies to e, e can take a step, and $\mu' = \mu$.

<u>Case T-FIELD:</u> $e = e_1 \cdot f$ and, by the induction hypothesis, either e_1 can make a step of evaluation or it is a value. Then, there are two subcases:

<u>Subcase</u> $\langle e_1 \mid \mu \rangle \longrightarrow \langle e_1' \mid \mu' \rangle$: If e_1 can take a step, then rule E-Congruence applies to e, and e can take a step.

Subcase e_1 is a value: If e_1 is a value, that is, a location l, then by inversion on T-FIELD, we have $\Gamma \mid \overline{\Sigma} \vdash l : \{\varepsilon\} \mid \{x \Rightarrow \overline{\sigma}\}$ and var $f : \tau \in \overline{\sigma}$. By inversion on T-Loc, we know that the store contains an appropriate mapping for the location l and, since T-Store has been applied throughout, the store is well typed and $l \mapsto \{x \Rightarrow \overline{d}\} \in \mu$ with var $f : \tau = l_1 \in \overline{d}$. Therefore, the rule E-FIELD applies to e, e can take a step, and $\mu' = \mu$.

<u>Case T-Assign:</u> $e = (e_1.f = e_2)$ and, by the induction hypothesis, either e_1 is a value or else it can make a step of evaluation and, likewise, e_2 . Then, there are two subcases:

Subcase $\langle e_1 \mid \mu \rangle \longrightarrow \langle e_1' \mid \mu' \rangle$ or e_1 is a value and $\langle e_2 \mid \mu \rangle \longrightarrow \langle e_2' \mid \mu' \rangle$: If e_1 can take a step or if e_1 is a value and e_2 can take a step, then rule E-Congruence applies to e and e can take a step.

<u>Subcase e_1 and e_2 are values</u>: If both e_1 and e_2 are values, that is, they are locations l_1 and l_2 , respectively, then by inversion on T-Assign, we have that $\Gamma \mid \Sigma \vdash l_1 : \{\varepsilon_1\} \{x \Rightarrow \overline{\sigma}\}$, var $f : \tau \in \overline{\sigma}$, and $\Gamma \mid \Sigma \vdash l_2 : \{\varepsilon_2\} \tau$. By inversion on T-Loc, we know that the store contains an appropriate mapping for the locations l_1 and l_2 and, since T-Store has been applied throughout, the store is well typed and $l_1 \mapsto \{x \Rightarrow \overline{d}\} \in \mu$ with var $f : \tau = l \in \overline{d}$. A new well-typed store can be created as follows: $\mu' = [l_1 \mapsto \{x \Rightarrow \overline{d}'\}/l_1 \mapsto \{x \Rightarrow \overline$

Case T-SuB: The result follows directly from the induction hypothesis.

Thus, the program written in this language never gets stuck.

C CASE STUDIES ON EFFECT BOUNDS

C.1 Controlling Access to UI Objects

This main idea of the work of Gordon et al. [2013] is to control the access of user interface (UI) framework methods so that unsafe UI methods can be called only by the UI thread. There are three different method annotations, @SafeEffect, @UIEffect, and @PolyUIEffect, where

- (1) @SafeEffect annotates methods that are safe to run on any thread,
- (2) @UIEffect annotates methods that are callable only on a UI thread, and
- (3) @PolyUIEffect annotates methods whose effects are polymorphic over the receiver type's effect parameter.

In Wyvern, we can model @UIEffect as a member of the UI module, for example:

```
type UILibrary
  effect UIEffect >= {system.FFI}
  def unsafeUIMethod1(): {this.UIEffect} Unit
  def unsafeUIMethod2(): {this.UIEffect} Unit
  ...
```

Bounded Abstract Effects 5:45

This way, any client code of a UI library that calls UI methods will have the uilibrary.UIEffect effect.

An interface could be used for UI-effectful or UI-safe work. To accommodate such flexibility, Java_{UI} introduced the <code>@PolyUIType</code> annotation. For example, a <code>Runnable</code> interface that can be UI-safe or UI-unsafe is declared as

```
@PolyUIType public interface Runnable {
    @PolyUIEffect void Run();
}
```

Whether the method Run() will have a UI effect depends on an annotation when the type is instantiated. For example:

```
@Safe Runnable s = ....;
s.run(); // is UI safe
@UI Runnable s = .....;
s.run(); // has UI effect
```

In Wyvern, such a polymorphic interface can be created by defining the interface with a bounded effect member:

```
type Runnable
  effect Run <= {uiLibrary.UIEffect}
  def run(): {this.Run} Unit</pre>
```

This type ensures that the run method is safe to be called on the UI thread. Moreover, if an instance of Runnable does not have UIEffect, it can be ascribed with the type SafeRunnable, which is a subtype of Runnable:

```
type SafeRunnable
  effect Run = {}
  def run(): {this.Run} Unit
```

This indicates that run is safe to be called on any thread.

C.2 Controlling Mutable States Using Abstract Regions

Greenhouse and Boyland [1999] proposed a region-based effect system that describes how state may be accessed during the execution of some program component in object-oriented programming languages. One example of the usage of regions is as follows:

```
class Point {
  public region Position;
  private int x in Position;
  private int y in Position;
  public scale(int sc) reads nothing writes Position {
    x *= sc;
    y *= sc;
  }
}
```

The two variables x and y are declared inside a region Position. For each region, there can be two possible effects: read and write. The scale method has the effect of writing on the region this Position.

To achieve access control on regions in Wyvern, we need to keep track of the read and write effect on each variable in a region. We declare the resource type Var representing a variable wrapper.

5:46 D. Melicher et al.

```
resource type Var[T]
  effect Read
  effect Write
  def set (x: T): {this.Write} Unit
  def get (): {this.Read} T
```

Since the set and get methods are annotated with the corresponding effects and there is no exposed access to the variable that holds the value, the two methods protect the access to the variable inside the type Var. To avoid code boilerplate, this wrapper type can be added as a language extension. The earlier Point example can be rewritten in Wyvern as:

```
resource type Point
  val x: Var[Int]
  val y: Var[Int]
  effect Read >= {this.x.Read, this.y.Read}
  effect Write >= {this.x.Write, this.y.Write}
  def scale(sc: Int): {this.Write} Unit
```

Note that the Write effect of Point composes the write effects of x and y into a single higher-level effect, which is analogous to Position in Greenhouse and Boyland [1999].

We can also extend the type Point to 3DPoint in the following way:

```
resource type 3DPoint
  val x: Var[Int]
  val y: Var[Int]
  val z: Var[Int]
  effect Read = {this.x.Read, this.y.Read, this.z.Read}
  effect Write = {this.x.Write, this.y.Write, this.z.Write}
  def scale(sc: Int): {this.Write} Unit
```

Since the effect Read and Write in the type Point is declared with a lower bound, the type 3DPoint is a subtype of Point.

REFERENCES

Nada Amin, Tiark Rompf, and Martin Odersky. 2014. Foundations of path-dependent types. In Proceedings of the 2014 ACM International Conference on Object Oriented Programming Systems Languages and Applications (OOPSLA'14). ACM, New York, NY, USA, 233–249. https://doi.org/10.1145/2660193.2660216

Dor Azouri. 2018. Abusing text editors with third-party plugins. Retrieved November 17, 2021 from https://go.safebreach.com/rs/535-IXZ-934/images/Abusing_Text_Editors.pdf.

Dariusz Biernacki, Maciej Piròg, Piotr Polesiuk, and Filip Sieczkowski. 2019. Abstracting algebraic effects. *Proc. ACM Program. Lang.* 3, Article 28 (2019). DOI: https://doi.org/10.1145/3290319

Dariusz Biernacki, Maciej Piróg, Piotr Polesiuk, and Filip Sieczkowski. 2020. Binders by day, labels by night: Effect instances via lexically scoped handlers. *Proc. ACM Program. Lang.* 4, POPL, Article 48 (Dec. 2020), 29 pages. https://doi.org/10.1145/3371116

Robert L. Bocchino, Jr., Vikram S. Adve, Danny Dig, Sarita V. Adve, Stephen Heumann, Rakesh Komuravelli, Jeffrey Overbey, Patrick Simmons, Hyojin Sung, and Mohsen Vakilian. 2009. A type and effect system for deterministic parallel Java. In Object Oriented Programming Systems Languages and Applications. 20. https://doi.org/10.1145/1640089.1640097

Jonathan Immanuel Brachthäuser, Philipp Schuster, and Klaus Ostermann. 2020a. Effects as capabilities: Effect handlers and lightweight effect polymorphism. *Proc. ACM Program. Lang.* 4, OOPSLA, Article 126 (Nov. 2020), 30 pages. https://doi.org/10.1145/3428194

Jonathan Immanuel Brachthäuser, Philipp Schuster, and Klaus Ostermann. 2020b. Effekt: Capability-passing style for typeand effect-safe, extensible effect handlers in Scala. *Journal of Functional Programming* 30 (2020), e8. https://doi.org/10. 1017/S0956796820000027

Oliver Bračevac, Nada Amin, Guido Salvaneschi, Sebastian Erdweg, Patrick Eugster, and Mira Mezini. 2018. Versatile event correlation with algebraic effects. *Proceedings of the ACM on Programming Languages* 2, ICFP, Article 67 (2018), 31 pages. https://doi.org/10.1145/3236762

- Aaron Craig, Alex Potanin, Lindsay Groves, and Jonathan Aldrich. 2018. Capabilities: Effects for free. In Formal Methods and Software Engineering—20th International Conference on Formal Engineering Methods (ICFEM'18), Gold Coast, QLD, Australia, November 12–16, 2018, Proceedings (Lecture Notes in Computer Science), Jing Sun and Meng Sun (Eds.), Vol. 11232. Springer, Berlin, 231–247. https://doi.org/10.1007/978-3-030-02450-5_14
- Peter J. Denning. 1976. Fault tolerant operating systems. ACM Comput. Surv. 8, 4 (Dec. 1976), 359–389. https://doi.org/10. 1145/356678.356680
- Stephen Dolan, Spiros Eliopoulos, Daniel Hillerström, Anil Madhavapeddy, K. C. Sivaramakrishnan, and Leo White. 2018. Concurrent system programming with effect handlers. In *Trends in Functional Programming*, Wang, and Owens, Scott.
- Derek Dreyer, Karl Crary, and Robert Harper. 2003. A type system for higher-order modules. SIGPLAN Not. 38, 1 (Jan. 2003), 236-249. https://doi.org/10.1145/640128.604151
- Sophia Drossopoulou, James Noble, Mark S. Miller, and Toby Murray. 2016. Permission and authority revisited towards a formalisation. In *Proceedings of the 18th Workshop on Formal Techniques for Java-like Programs*. Association for Computing Machinery. DOI: 10.1145/2955811.2955821
- Matthias Felleisen. 1991. On the expressive power of programming languages. Science of Computer Programming 17, 1 (1991), 35–75. https://doi.org/10.1016/0167-6423(91)90036-W
- Andrzej Filinski. 2010. Monads in action. In Proceedings of the 37th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'10). ACM, New York, NY, 483–494. https://doi.org/10.1145/1706299.1706354
- Jennifer Fish, Darya Melicher, and Jonathan Aldrich. 2020. A case study in language-based security: Building an I/O library for Wyvern. In *Proceedings of the 2020 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software*. Association for Computing Machinery. DOI: 10.1145/3426428.3426913
- Marco Gaboardi, Shin-ya Katsumata, Dominic Orchard, Flavien Breuvart, and Tarmo Uustalu. 2016. Combining effects and coeffects via grading. In Proceedings of the 21st ACM SIGPLAN International Conference on Functional Programming. Association for Computing Machinery, 476–489. DOI: https://doi.org/10.1145/2951913.2951939
- Colin S. Gordon, Werner Dietl, Michael D. Ernst, and Dan Grossman. 2013. JavaUI: Effects for controlling UI object access. In ECOOP 2013 Object-Oriented Programming, Lecture Notes in Computer Science, Giuseppe Castagna (Ed.). Springer, Berlin, 179–204.
- Aaron Greenhouse and John Boyland. 1999. An object-oriented effects system. In ECOOP.
- Pablo Inostroza and Tijs van der Storm. 2018. JEff: Objects for effect. In Proceedings of the 2018 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software (Onward! 2018). ACM, New York, NY, 111–124. https://doi.org/10.1145/3276954.3276955
- Joseph R. Kiniry. bibinfoyear2006. Advanced Topics in Exception Handling Techniques. Springer-Verlag, Chapter Exceptions in Java and Eiffel: Two Extremes in Exception Design and Application. http://dl.acm.org/citation.cfm?id=2124243. 2124264
- Daan Leijen. 2014. Koka: Programming with Row Polymorphic Effect Types. In *Electronic Proceedings in Theoretical Computer Science*, vol. 153, 10.4204/EPTCS.153.8
- K. Rustan M. Leino, Arnd Poetzsch-Heffter, and Yunhong Zhou. 2002. Using data groups to specify and check side effects. In Conference on Programming Language Design and Implementation. 12. https://doi.org/10.1145/512529.512559
- Shu-Peng Loh and Sophia Drossopoulou. 2017. Specifying Attenuation. Retrieved November 17, 2021 from https://2017. splashcon.org/event/ocap-2017-specifying-attenuation.
- Yuheng Long, Yu David Liu, and Hridesh Rajan. 2015. Intensional effect polymorphism. In 29th European Conference on Object-Oriented Programming, ECOOP 2015, July 5–10, 2015, Prague, Czech Republic (LIPIcs), John Tang Boyland (Ed.), Vol. 37. Schloss Dagstuhl Leibniz-Zentrum für Informatik, 346–370. https://doi.org/10.4230/LIPIcs.ECOOP.2015.346
- John M. Lucassen. 1987. Types and Effects towards the Integration of Functional and Imperative Programming. Ph.D. Dissertation. Massachusetts Institute of Technology. Retrieved from https://www.proquest.com/dissertations-theses/types-effects-towards-integration-functional/docview/303629227/se-2?accountid=9902.
- John M. Lucassen and David K. Gifford. 1988. Polymorphic effect systems. In Symposium on Principles of Programming Languages. 11. https://doi.org/10.1145/73560.73564
- Sergio Maffeis, John C. Mitchell, and Ankur Taly. 2010. Object capabilities and isolation of untrusted web applications. In *IEEE Symposium on Security and Privacy*. 16.
- Daniel Marino and Todd Millstein. 2009. A generic type-and-effect system. In *International Workshop on Types in Language Design and Implementation*. 12. https://doi.org/10.1145/1481861.1481868
- Darya Melicher. 2020. Controlling Module Authority Using Programming Language Design. Ph.D. Dissertation. Carnegie Mellon University.
- Darya Melicher, Yangqingwei Shi, Alex Potanin, and Jonathan Aldrich. 2017. A capability-based module system for authority control. In 31st European Conference on Object-Oriented Programming, (ECOOP'17), Peter Müller, Vol. 74. Schloss Dagstuhl Leibniz-Zentrum für Informatik, 20:1–20:27. DOI: 10.4230/LIPIcs.ECOOP.2017.20
- Adrian Mettler, David Wagner, and Tyler Close. 2010. Joe-E: A security-oriented subset of Java. In *Network and Distributed System Security Symposium*.
 - ACM Transactions on Programming Languages and Systems, Vol. 44, No. 1, Article 5. Publication date: January 2022.

5:48 D. Melicher et al.

Mark S. Miller. 2006. Robust Composition: Towards a Unified Approach to Access Control and Concurrency Control. Ph.D. Dissertation. Johns Hopkins University.

John C. Mitchell and Gordon D. Plotkin. 1988. Abstract types have existential type. ACM Trans. Program. Lang. Syst. 10, 3 (July 1988), 470–502. https://doi.org/10.1145/44501.45065

Toby Murray. 2008. Analysing object-capability security. In Joint Workshop on Foundations of Computer Security, Automated Reasoning for Security Protocol Analysis and Issues in the Theory of Security.

Martin Odersky and Matthias Zenger. 2005. Scalable component abstractions. SIGPLAN Not. 40, 10 (Oct. 2005), 41–57. https://doi.org/10.1145/1103845.1094815

David L. Parnas. 1971. Information distribution aspects of design methodology. IFIPS Congress 71, 339-344.

David L. Parnas. 1972. On the criteria to be used in decomposing systems into modules. Commun. ACM 15, 12 (1972), 1053–1058. https://doi.org/10.1145/361598.361623

Tomas Petricek, Dominic Orchard, and Alan Mycroft. 2014. Coeffects: A calculus of context-dependent computation. In *Proceedings of the 19th ACM SIGPLAN International Conference on Functional Programming*. Association for Computing Machinery, 123–135. DOI: 10.1145/2628136.2628160

Gordon Plotkin and John Power. 2003. Algebraic operations and generic effects. *Applied Categorical Structures* 11, 1 (2003), 69–94. https://doi.org/10.1023/A:1023064908962

Gordon Plotkin and Matija Pretnar. 2009. Handlers of algebraic effects. In *Programming Languages and Systems*, Giuseppe Castagna (Eds.). Springer Berlin Heidelberg, 80–94.

Marianna Rapoport and Ondřej Lhoták. 2019. A path to DOT: Formalizing fully path-dependent types. *Proc. ACM Program. Lang.* 3, OOPSLA, Article 145 (Oct. 2019), 29 pages. https://doi.org/10.1145/3360571

Lukas Rytz, Martin Odersky, and Philipp Haller. 2012. Lightweight polymorphic effects. In European Conference on Object-Oriented Programming. 25. https://doi.org/10.1007/978-3-642-31057-7_13

Valery Trifonov and Zhong Shao. 1999. Safe and principled language interoperation. In Programming Languages and Systems, S. Doaitse Swierstra (Eds.). Springer Berlin Heidelberg, 128–146.

Franklyn A. Turbak and David K. Gifford. 2008. Design Concepts in Programming Languages. The MIT Press, Cambridge, MA.

Marko van Dooren and Eric Steegmans. 2005. Combining the robustness of checked exceptions with the flexibility of unchecked exceptions using anchored exception declarations. SIGPLAN Not. 40, 10 (Oct. 2005), 455–471. https://doi.org/10.1145/1103845.1094847

Anlun Xu. 2020. Extending Abstract Effects with Bounds and Algebraic Handlers. Master's Thesis. Carnegie Mellon University.

Yizhou Zhang and Andrew C. Myers. 2019. Abstraction-safe effect handlers via tunneling. Proceedings of the ACM on Programming Languages 3, POPL, Article 5 (2019), 29 pages. https://doi.org/10.1145/3290318

Received November 2020; revised September 2021; accepted October 2021