MIR: Automated Quantifiable Privilege Reduction
Against Dynamic Library Compromise in JavaScript

Nikos Vasilakis   Cristian-Alexandru Staicu†   Greg Ntousakis°   Konstantinos Kallas‡

Ben Karel◇   André DeHon‡   Michael Pradel††

MIT, CSAIL   †*TU Darmstadt & CISPA*   °*TU Crete*   ‡*University of Pennsylvania*   ◇*Aarno Labs*   ††*University of Stuttgart*

Abstract

Third-party libraries ease the development of large-scale software systems. However, they often execute with significantly more privilege than needed to complete their task. This additional privilege is often exploited at runtime via dynamic compromise, even when these libraries are not actively malicious. MIR addresses this problem by introducing a fine-grained read-write-execute (RWX) permission model at the boundaries of libraries. Every field of an imported library is governed by a set of permissions, which developers can express when importing libraries. To enforce these permissions during program execution, MIR transforms libraries and their context to add runtime checks. As permissions can overwhelm developers, MIR's permission inference generates default permissions by analyzing how libraries are used by their consumers. Applied to 50 popular libraries, MIR's prototype for JavaScript demonstrates that the RWX permission model combines simplicity with power: it is simple enough to automatically infer 99.33% of required permissions, it is expressive enough to defend against 16 real threats, it is efficient enough to be usable in practice (1.93% overhead), and it enables a novel quantification of privilege reduction.

1   Introduction

Modern software development relies heavily on third-party libraries.[1] Such reliance has led to an explosion of attacks [14, 29, 33, 34, 56, 68]: overprivileged code in imported libraries provides an attack vector that is exploitable long after libraries reach their end-users. Even when libraries are created and authored with the best possible intentions—*i.e.*, are not actively malicious—their privilege can be exploited at runtime to compromise the entire application—or worse, the broader system on which the application is executing.

Such *dynamic compromise*—as opposed to compromise achieved earlier in the supply chain, for actively-malicious
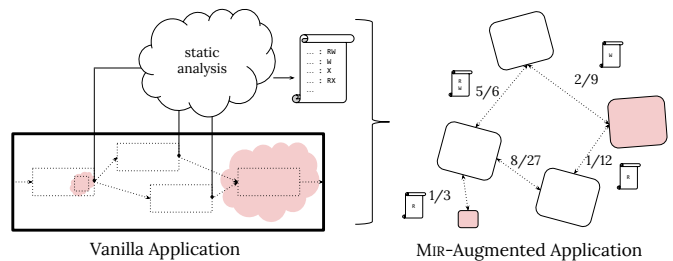
---

[1]This paper uses the terms library, module, and package interchangeably. It also uses `import` everywhere, for both Node's `require` and ES6's `import`.

Fig. 1: MIR overview. MIR analyzes individual libraries, many of which may be subvertible, to generate permissions. It then enforces these permissions at runtime in order to (quantifiably) lower the privielge of these libraries over the rest of the application and its surrounding environment.

libraries—is possible due to several compounding factors. Libraries offer a great deal of functionality, but only a small fraction of this functionality may be used by any one particular client [30]. Default-allow semantics give any library unrestricted access to all of the features in a programming language—*e.g.*, accessing global variables, introspecting and rewriting core functionality, and even importing other libraries [11]. A library's intended use is a small space well-understood by its developers, but unexpected or pathological use covers a much larger space, typically understood by the library's clients [13]. For example, only the client of a deserialization library knows whether it will be fed non-sanitized input coming directly from the network; the library's developers cannot make such assumptions about its use.

To address dynamic compromise, MIR augments a module system with a model for specifying, enforcing, inferring, and quantifying the privilege available to libraries in a backward-compatible fashion (Fig. 1). MIR's key insight is that libraries cannot be subverted at runtime to exploit functionality to which they do not already have access. Coupling default-deny semantics with explicit and statically-inferrable whitelisting, MIR minimizes the effects of dynamic compromise—regardless of the behavior exercised by a library and its clients.

Specification   MIR's permission model allows specifying fine-grained *read-write-execute permissions* (RWX) that guard access to individual object fields available within the scope of

a library. Such objects and fields are available through both explicit library imports and built-in language features—*e.g.*, globals, process arguments, **import** capabilities. In the aforementioned de-serialization library, two example permissions would be (1) R only on `process.env.PWD`, only allowing read-access to the `PWD` environment variable, and (2) X only on `serialize`, disallowing execute-access to the—unused but exploitable—de-serialization function.

Inference Unfortunately, manually specifying permissions is challenging, even for security-paranoid developers. Typical challenges include (1) naming issues, such as variable aliasing, poor variable names, and non-local use, (2) library-internal code, possibly not intended for humans, and (3) continuous codebase evolution, which requires updating the specification for every code change. To address these problems, MIR analyzes name uses within a library to automatically infer permissions. The majority of the analysis is static, augmented with a short phase of import-time analysis addressing common runtime meta-programming concerns.

Quantification MIR improves security by restricting the permissions granted to libraries by default; at the same time, inferrring 100% of a library's permissions is intractable for dynamic languages such as JavaScript. To address this conundrum, we propose a quantitative privilege reduction metric to evaluate the permission prevention that MIR exercises. This quantification is achieved by comparing the permissions granted by MIR to those a library would have by default—*i.e.*, by statically counting all the names available in a library's lexical scope.

Enforcement To enforce RWX permissions, MIR transforms libraries to add security monitors. MIR's transformations are enabled by the fact that JavaScript features a library-import mechanism that loads code at runtime as a string. Such lightweight load-time code transformations operate on the string representation of the module, as well as the context to which it is about to be bound, to insert permission-checking code into the library before it is loaded.

Evaluation MIR's evaluation shows that only 7 out of 1044 (0.67%) unique accesses performed in the tests of 50 libraries were not automatically inferred. MIR defends against 16 real attacks, only requiring a manual modification of the infered permissions for one attack. In terms of runtime performance, MIR's static analysis averages 2.5 seconds per library and its enforcement overhead remains around 1.93% of the unmodified runtime.

Sections 4–8 present our key contributions:

- *Permission model and language*: a simple yet effective permission model and associated expression language for controlling the functionality available to modules (§4),
- *Automated permission inference*: a permission inference component that aids developers by analyzing libraries to generate their default permissions (§5),
- *Privilege reduction quantification*: a conceptual framework

```
1 function srv(req, res) {      1 let lg = import("log");
2 let srl, obj;                 2 lg.LVL = lg.levels.WARN;
3 srl = import("serial");       3 module.export = {
4 obj = srl.dec(req.body);      4   dec: (str) => {
5 dispatch(obj, res);           5     let obj;
6 }                             6     lg.info("srl:dec");
                               7     obj = eval(str);
                               8     return obj
                               9   },
                              10   enc: (obj) => {...}
                              11 }
         main                          serial
```

Fig. 2: Use of third-party modules. The `main` module (left) **import**s off-the-shelf serialization implemented by the `serial` *third-party* module (right), vulnerable to remote code execution.

for quantifying privilege reduction, which allows reasoning about the privilege achieved by applying MIR (§6),

- *Backward-compatible runtime enforcement*: a series of program transformations applied upon library load to insert permission checks, interposing on interactions with other libraries and core language structures (§7),
- *Implementation and evaluation*: an open-source implementation (§8) as an easily pluggable, backward-compatible library, along with an extensive compatibility (§8.1), security (§8.2–8.3), and performance (§8.4) evaluation.

Other sections include an example illustrating dynamic compromise and how MIR addresses it (§2), a discussion of the threat model (§3), MIR's relation to prior work (§9). We conclude (§10) that MIR's automation and performance characteristics make it an important addition to a developer's modern toolkit, in many circumstances working in tandem with defenses that focus on other threats.

2 Background and Overview

This section uses a server-side JavaScript application to present (1) security issues related to third-party code (§2.1), and (2) an overview of how MIR addresses these issues (§2.2).

2.1 Example: A De-serialization Library

Fig. 1 presents a simplified schematic of a multi-library program. Dotted boxes correspond to the context of different third-party modules of varying trust, one of which is used for (de-)serialization. This module is fed client-generated strings, often without proper sanitization [58], which may enable remote code execution (RCE) attacks. RCE problems due to serialization have been identified in widely used libraries [2,3,7] as well as high-impact websites, such as PayPal [60] and LinkedIn [61]. Injection and insecure de-serialization are respectively ranked number one and eight in OWASP's ten most critical web application security risks [49].

For clarity of exposition, Fig. 2 zooms into only two fragments of the program. The **main** module (left) **import**s off-

the-shelf serialization functionality through the **serial** module, whose **dec** method de-serializes strings using **eval**, a native language primitive. The **serial** module **import**s **log** and assigns it to the `lg` variable.[2]

Although **serial** is not actively malicious, it is subvertible by attackers at runtime, who can use the input `str` for several attacks: (1) overwrite `info`, affecting all (already loaded) uses of **log**.`info` across the entire program; (2) inspect or rewrite top-level **Object** handlers, built-in libraries, such as **crypto**, and the **cache** of loaded modules; (3) access global or pseudo-global variables such as **process** to reach into environment variables; and (4) load other modules, such as **fs** and **net**, to exfiltrate files over the network.

2.2 Overview: Applying MIR on Serial

Our work address these security problems by intercepting applications at the boundaries between libraries. MIR's permission model annotates functionality that is not part of the current library with RWX permissions. A part of this functionality comes from imported libraries; for example, among other permissions, **serial** needs to be able to execute `info` from module **log**—*i.e.*, **log**.`info` needs X. Another part of this functionality comes directly from the programming language; for example, **serial** clearly needs X for **import** and **eval**. It is not **serial** that provides this functionality, but rather the language and its runtime environment.

These three permissions are a part of the total nine required for **serial**'s normal operation. To aid developers in identifying the remaining permissions, MIR comes with a static inference component that analyzes how libraries use the available names. Fig. 3 shows a small example of this analysis: **levels** and **WARN** are read, thus are annotated as R; **LEVEL** is written, thus W; and `info` is executed, thus X. The analysis also infers permissions for **import**, **eval**, **module**, and **exports**. Names that do not show up—even if they are built-in language objects—get no permissions.

After extracting all necessary permissions, the developer can start the program using MIR's runtime enforcement component. MIR shadows all variable names that cross a boundary with variables that point to modified values. When accessing a modified value, MIR checks the permissions before forwarding access to the original value. If a module does not have permission to access a value, the approach throws a special exception that helps the developer in diagnosing its cause.

The attacks described at the end of the previous section (§2.1) are now impossible: (1) overwriting `info` from **serial** will throw a W violation, (2) inspecting **Object** handlers, built-in libraries, such as **crypto**, and the **cache** of loaded mod-

---

[2]Naming is important in this paper: we differentiate between the module **log** and the variable `lg`. MIR tracks permissions at the level of modules, irrespective of the variables they are assigned to. To aid the reader, modules, and more broadly, contexts, are typeset in **purple sans serif**, fields in `olive teletype`, and plain variables in uncolored teletype fonts.

```
var lg = import("log");
lg.LEVEL = lg.levels.WARN;
lg.info("srl:dec");
```
◻Read(R) ◻Write(W) ◻Execute(X)

Fig. 3: MIR Inference. MIR uses static analysis to infer module permissions. In this code snippet, the different fields of the **log** module are given distinct permissions based on how they are used.

ules will all throw R violations, and (3) accessing global or pseudo-global variables, such as **process**, to reach into the environment will also throw R violations. Shielding against (4) module loading depends on a refinement on the base RWX model (§4); this refinement does not make automated static inference intractable, as MIR just notes that **import** is only executable with argument log.

3 Threat Model

MIR focuses on the dynamic compromise of possibly buggy or vulnerable libraries, *i.e.*, ones that are not actively malicious or hiding their intentions. Such libraries are subvertible by attackers providing payloads through web interfaces, programmatic APIs, or shared interfaces. Of particular interest are libraries that offer some form of object de-serialization or runtime code evaluation, where attackers can inject or execute arbitrary code by passing carefully constructed payloads. This is because these libraries implement their features using runtime interpretation, and thus are subvertible when receiving potentially attacker-controlled inputs (see §8).

Focus: MIR focuses on the confidentiality (*e.g.*, read global state, load other libraries, exfiltrate data) and integrity (*e.g.*, write global state or tamper with the library cache) of data and code. These concerns extend to the broader environment within which a program is executing—including environment variables, the file system, or the network. Such ambient over-privilege stems from: (1) common features in programming languages (*e.g.*, call stack inspection, reflection capabilities, monkey patching); (2) unusual language features or deficiencies (*e.g.*, in JavaScript: default-is-global, mutability attacks); (3) implementation concerns (*e.g.*, library cache, import capabilities); and (4) authority confusions, where all parts of a program have equal access rights (*e.g.*, read `process.env` or `process.args`, write to the file-system or network).

MIR aims to mitigate the aforementioned attack vectors by allowing users to control ambient over-privilege at a fine granularity—that of individual names, functions, and fields in the scope of a module. MIR does not limit itself to the interface exposed explicitly by a module, but rather captures the full observable behavior around it: any name that resolves to a value defined outside the module is protected—including language features, built-in names, and environment constructs accessible programmatically.

Assumptions: MIR's static analysis is assumed to be per-

$$
\begin{array}{rcll}
s, m & \in & \text{String} \\
r & := & \text{R} \mid \varepsilon & \text{ReadPerm} \\
w & := & \text{W} \mid \varepsilon & \text{WritePerm} \\
x & := & \text{X} \mid \varepsilon & \text{ExecPerm} \\
i & := & \text{I} \mid \varepsilon & \text{ImpoPerm} \\
\mu & := & [\, r\, w\, x\, i\,] & \text{Mode} \\
f & := & f.s \mid \star.f \mid f.\star \mid s & \text{ObjPath} \\
p & := & f : \mu \mid f : \mu, p & \text{ModPermSet} \\
\omega & := & m : \{p\} \mid m : \{p\}, \omega & \text{FullPermSet} \\
\end{array}
$$

Fig. 4: MIR's permission DSL. The DSL captures the set of annotations used for specifying permissions across libraries (*Cf.*§4).

formed prior to execution, otherwise a malicious library can rewrite the code of a benign library upon load. For the same reason, MIR's runtime enforcement component is assumed to be loaded prior to any other library. MIR places trust in the language runtime and built-in modules, such as `fs`: a minimum of trusted functionality is needed from the module system to locate and load permissions.

MIR does not consider native libraries written in lower-level languages, such as C/C++, or available as binaries. These libraries are out of scope for two reasons. First, they cannot be analyzed by MIR's static analysis, which operates on source code. Second, they can bypass MIR's language-based runtime protection, which depends on memory safety. Any operation violating memory safety could access arbitrary memory regions. MIR also does not consider availability, denial-of-service, and side-channel attacks.

4 Permission Model and Language

MIR's goal is to reduce the privilege that libraries possess. At the core of our approach is the ability to explicitly specify a subset of the rights granted to modules by default upon import. This specification is expressible per-library using a domain-specific language (DSL, Fig. 4) that focuses on read (R), write (W), execute (X), and import (I) permissions.

Core Permission Model The core of MIR's permission model and associated DSL is a per-library permission set: ModPermSet maps names accessible within the library context to a Mode, *i.e.*, a set of access rights encoded as RWX permissions. Ignoring ⋆-constructs for now, names represent access paths within the object graph reachable from within the scope of the library—*e.g.*, **String.toUpperCase**. Access paths start from a few different points that can be grouped into two broad classes. The first class contains a closed set of known root points that are provided by the language, summarized in the first four rows of Tab. 1. These names are available by default through (and shared with) the library's outer context, *i.e.*, resolving to a scope outside that of a library and pervasively accessible from any point in the code. Examples

include top-level objects and functions, such as **process.args** and **eval**, functions to perform I/O, such as **console.log**, and the library-**import** ability *itself*.

The second class contains access paths that start from explicitly importing a new library into the current scope. Such an import results in multiple names available through the imported library's (equivalent of) **export** statement. Examples of such paths from Fig. 2 include **log.info** and **srl.dec** (§2.1).

MIR's model can thus be thought as an object-path protection service: access rights are expressed as permissions associated with a path from the program's context roots to the field currently accessed. Values created within the scope of a library or a function are *not* part of this model: MIR does not allow specifying or enforcing access restrictions on, say, arbitrary objects or function return values.

Semantics: The semantics behind the core set of permissions can be summarized as follows:

- A read permission (R) grants clients the ability to read a value, including assigning it to variables and passing it around to other modules.
- A write permission (W) grants clients the ability to modify a value, and includes the ability to delete it. The modification will be visible by all modules that have read permissions over the original value.
- An execute permission (X) grants clients the ability to execute a value, provided that it points to an executable language construct, such as a function or a method. It includes the ability to invoke the value as a constructor (typically prefixed by `new`).

RWX permissions are loosely based on the Unix permission model, with a few key differences. Reading a field of a composite value `x.f` requires R permissions on the value `x` *and* the field `f`—that is, an R permission allows only a single dereference. Reading or copying a function only requires an R-permission, but performing introspection requires X permissions over its subfields due to introspection facilities being provided by auxiliary methods (*e.g.*, `toString` method). A W permission on the base pointer allows discarding the entire object. While a base write may look like it bypasses all permissions, modules holding pointers to fields of the original value will not see any changes.

Example: To illustrate the base permission model on the de-serialization example (§2), consider `main`'s permissions:

```
1  main:
2    import: RX
3    import("serial").dec: RX
```

The set of permissions for `serial` is more interesting:

```
1  serial:
2    eval: RX
3    import: RX
4    module: R
5    module.exports: W
```

4

Tab. 1: Object path roots. Objects resolving to a scope outside that of a module can be reached through a few starting points: (1) core built-in objects, (2) the standard library, (3) implementation-specific objects, (4) module-locals, and (5) global variables.

Root Context	Example Names
es	`Math`, `Number`, `String`, `JSON`, `Reflect`, ...
node	`Buffer`, `process`, `console`, `setImmediate`, ...
lib-local	`exports`, `module.exports`, `__dirname`, ...
globs	`GLOBAL`, `global`, `Window`
import	`import(`*`lib`*`)`,

```
6   import("log").levels: R
7   import("log").levels.WARN: R
8   import("log").info: RX
9   import("log").LVL: W
```

Importing   A simple `X` permission to the built-in `import` function gives libraries too much power. Thus, MIR needs to allow specifying which imports are permitted from a library.

This is achieved through an additional `I` permission. This permission is provided to an `ObjPath` that explicitly specifies the absolute file-system path of a library.[3] Using the absolute file-system path is a conscious decision: the same library name imported from different locations of a program may resolve to different libraries residing in different parts of the file system and possibly corresponding to different versions.

Using a separate permission `I` provides additional flexibility by distinguishing from `R`. Libraries are often imported by a part of the program only for their side-effects (*i.e.*, not for their interface). In these cases, their fields should not necessarily be accessible by client code. Typical examples include singleton configuration objects and stateful servers.

Wildcards: It is often helpful or necessary to provide a single permission mode to *all* possible matches of a segment within a path. To achieve this, MIR offers *wildcards*: `*`.`f` assigns a mode to all fields named `f` reachable from any object, and `o`.`*` assigns a mode to all fields of an object (or path) `o`. These forms may also be combined, as in `o`.`*`.`f`.

Wildcards have many practical uses. The primary use case is when fields or objects are altered through runtime meta-programming. In such cases, the fields are not necessarily accessible from a single static name and might depend on dynamic information. Often, these fields (not just the paths) are constructed at runtime, which means that they are not available for introspection by MIR at library-load time.

5   Permission Inference

To aid users in expressing permissions, MIR provides an analysis that automatically infers permissions for a library and its dependencies.

---

[3]For portability, MIR prefixes records with a `__CWD__` variable that can be instantiated to different values across environments.

Static Analysis   The core of the analysis is an intra-procedural, flow-sensitive forward data flow static analysis. The analysis is conservative in the sense that if it infers a permission, then there exists a path in the analyzed module that uses the permission. In contrast, the analysis may miss required permissions, and we evaluate how often this happens in practice.

The analysis visits each statement of a module by traversing a control flow graph of each function. During these visits, it updates two data structures. First, it updates the set C of (API, permission) pairs that eventually will be reported as the inferred permission set. The set C grows monotonically during the entire analysis, *i.e.*, the analysis adds permissions until all uses of third-party code have been analyzed. Second, the analysis updates a map *DefToAPI*, which maps definitions of variables and properties to the fully qualified API that the variable or property points to after the definition. For example, when visiting a definition `x = import("foo").bar`, the analysis updates *DefToAPI* by mapping the definition of `x` to "foo.bar". The *DefToAPI* map is a helper data structure discarded when the analysis completes analyzing a function.

Table 2 summarizes how the analysis updates C and *DefToAPI* when visiting specific kinds of statements. The updates to C reflect the way that the analyzed module uses library-external names. Specifically, whenever a module reads, writes, or executes an API a, then the analysis adds to C a permission (a, R), (a, W), or (a, X), respectively. The updates to *DefToAPI* propagate the information about which APIs a variable or property points to. For example, suppose that the analysis knows that variable `a` points to a module "foo" just before a statement `b = a.bar`; then it will update *DefToAPI* with the fact that the definition of `b` now points to "foo.bar".

While traversing the control flow graph, the analysis performs the updates in Table 2 for every statement. On control flow branches, it propagates the current state along both branches. When the control flow merges again, then the analysis computes the union of the C sets and the union of the *DefToAPI* maps of both incoming branches. MIR handles loops by unrolling each loop once, which is sufficient in practice for analyzing uses of third-party code, because loops typically do not re-assign references to third-party APIs.

Qualified Interfaces   The update functions in Table 2 rely on a helper function *getAPIs*. Given a reference, *e.g.*, a variable or property access, this function returns the set of fully qualified APIs that the reference may point to. For example, after the statement `obj.x = import("foo").bar`, *getAPIs*($obj.x$) will return the set $\{\text{"}foo.bar\text{"}\}$. When queried with a variable that does not point to any API, *getAPIs* simply returns the empty set. Algorithm 1 presents the *getAPIs* function in more detail. We distinguish four cases, based on the kind of reference given to the function. Given a direct import of a module, *getAPIs* simply returns the name of the module. Given a variable, the function queries pre-computed reaching-definitions information (see below) to

Tab. 2: MIR's analysis updates. Updates performed by the static analysis when visiting specific kinds of statements.

Kind of statement	Updates	Example
Assignment *lhs* = *rhs* at location *l*:		
For each $a \in getAPIs(lhs)$	Add $(a, \texttt{W})$ to contract C	`someModule.foo = 5`
For each $a \in getAPIs(rhs)$	Add $(a, \texttt{R})$ to contract C	`x = import("someModule")`
	Add *lhs at l* $\mapsto$ *a* to *DefToAPI*	
Call of function *f*:		
For each $a \in getAPIs(f)$	Add $(a, \texttt{X})$ to contract C	`someModule.foo()`
Any other statement that contains a reference *r*:		
For each $a \in getAPIs(r)$	Add $(a, \texttt{R})$ to contract C	`foo(someModule.bar)`

Data: Reference r
Result: Set of APIs that r may point to
if *r* is an import of module "m" **then**
   | **return** { "m" }
end
if *r* is a variable **then**
   | $A \leftarrow \emptyset$
   | $defs \leftarrow$ get reaching definitions of r
   | **for** *each d in defs* **do**
      | $A \leftarrow A \cup DefToAPI(d)$
   | **end**
   | **return** A
end
if *r* is a property access *base.prop* **then**
   | $A_{base} \leftarrow getAPIs(base)$
   | **return** $\{a + "." + prop \mid a \in A_{base}\}$
end
return $\emptyset$

Algorithm 1: Helper function *getAPIs*.

obtain possible definitions of the variable, and then looks up the APIs these variables point to in *DefToAPI*. Given a property access, *e.g.*, `x.y`, the function recursively calls itself with the reference to the base object, *e.g.*, `x`, and then concatenates the returned APIs with the property name, *e.g.*, "y". Finally, for any other kind of reference, *getAPIs* returns an empty set. This case includes cases that we cannot handle with an intra-procedural analysis, *e.g.*, return values of function calls. In practice, these cases are negligible, because real-world code rarely passes around references to third-party APIs via function calls. We therefore have chosen an intra-procedural analysis, which ensures that the static contract inference scales well to large code-bases.

To find the APIs a variable may point to, Algorithm 1 gets the reaching definitions of the variable. This part of the analysis builds upon a standard intra-procedural may-reach definitions analysis, which MIR pre-computes for all functions in the module. To handle nested scopes, *e.g.*, due to nested function definitions, MIR builds a stack of definition-use maps, where each scope has an associated set of definition-use pairs. To find the reaching definitions of a variable, the analysis first queries the inner-most scope, and then queries the surrounding scopes until the reaching definitions are found. To handle

built-in APIs of JavaScript, *e.g.*, `console.log`, MIR creates an artificial outer-most scope that contains the built-in APIs available in the global scope.

Returning to the running example in Figure 2. For **main**, the static analysis results in the following permission set:

$$\{("serial'', \texttt{R}), ("serial.dec'', \texttt{R}), ("serial.dec'', \texttt{X})\}$$

As illustrated by the example, the inferred contract allows the intended behavior of the module, but prevents any other, unintended uses of third-party APIs. Our evaluation shows that the static analysis is effective also for larger, real-world modules (§8.2).

Limitations In line with MIR's design goal of being conservative in granting permissions, the analysis infers a permission only if there exists a path that uses the permission. In contrast, the analysis may miss permissions that a module requires. For example, missed permissions may results from code that passed a reference to a module across functions:

```
1 x = require("foo");
2 bar(x);
```

In this example, the analysis misses any permissions on "foo" that are required by `bar`. Tracking object references across function boundaries would require an inter-procedural analysis, which is difficult to scale to a module and its potentially large number of transitive dependencies [68]. Another example of potentially missed permissions is code that dynamically computes property names:

```
1 x = require("foo");
2 p = "ba" + "r";
3 x[p] = ...;
```

In this example, the analysis misses the `W` permission for "foo.bar". Tracking such dynamically computed property names is known to be a hard problem in static JavaScript analysis [57].

Import-time Permission Analysis To reduce the number of required permissions not inferred by its static analysis, MIR adds a short phase of dynamic import-time analysis. This analysis is performed by simply importing a library—but without invoking the library's interface—and only adds to the number of permissions inferred by the analysis described

earlier. The underlying insight is that many libraries simply wrap or re-export existing interfaces using dynamic meta-programming. The following code snippet demonstrates a simple but common pattern:

```
1 for (let k in fs) {
2   module.exports[k] = fs[k];
3 }
```

Inferring such meta-programming permissions statically poses a challenge due to the limitations discussed earlier, and thus simply loading the library should be enough to enable a more complete view into the library's behavior. Import-time analysis does not depend on the existence of library tests or any consuming code, as it does not call any library interfaces.

6 Quantifying Privilege Reduction

Any policy—whether automated or manual—on existing programs aims at striking a balance between compatibility and security: an ideal policy would allow only the necessary accesses but no more. Unfortunately, statically inferring such an ideal policy in the context of a dynamic language—even with the addition of import-time permissions—is impossible, due to the language's very nature. However, some analyses are better than others, *i.e.*, they infer policies with fewer accesses, even if they do not infer the ideal policy. To be able to quantitatively evaluate the security benefits offered by such analyses, we propose a novel privilege reduction metric.

Privilege Reduction   Informally, the single-library privilege reduction is calculated as the ratio of disallowed permissions over the full set of permissions available by default within the lexical scope of the library. The default permission set is calculated by statically expanding all names available in scope; the disallowed set of permissions is calculated by subtracting the allowed permissions from the default permission set. Single-library privilege reductions across the full dependency tree are then combined into a single reduction metric for the program's current decomposition. The following paragraphs explain the details.

The quantitative definition of privilege reduction depends on a few key notions. The first notion is a set of *target critical resources M_t* that are used to restrict a library's (or program's) access. The second notion is the set of *subjects M_s* that can access these resources through specific means (the permissions). MIR's quantitative definition of privilege reduction is general in that it can be applied to any scenario with a finite number of subjects and critical resources.

Informal Development   Before formalizing privilege reduction, we use the de-serialization example to build an intuition. Let's assume that from the two modules presented in Fig. 2, we are only interested in quantifying **main**'s prvilege; thus, $M_s = \{$**main**$\}$. As implied earlier (§4), the set of critical resources contains many paths available to **main**. For simplicity, we now assume it only contains **globals**, **fs**, and **import**; thus,

$M_t = \{$**globals**, **fs**, **import**$\}$. Module **main** needs an X permission on **import** to be able to load **serial**, and an X permission on **serial.dec** to be able to call the dec function. With this simple configuration, MIR disallows all accesses except for $P(M_s, M_t) = \{\langle$**import**$, X\rangle, \langle$**serial.dec**$, X\rangle\}$.

MIR's goal is to quantify this privilege with respect to the default permissions. If **main** was executed without additional protection, its privilege would be $P_{base}(M_s, M_t) = \{\langle$**globals.∗**$, RWX\rangle, \langle$**fs.read**$, RWX\rangle, ...\}$.

Formal Development   More formally, by default at runtime any module has complete privilege on all exports of any other module. Thus, for any modules m_1, m_2 the baseline privilege that m_1 has on m_2 is:

$$P_{base}(m_1, m_2) = \{\langle a, \rho\rangle | a \in API_{m_2}, \rho \in P\}$$

where $\rho \in \mathcal{P}$ is a set of orthogonal permissions on a resource, which for MIR is $\mathcal{P} = \{R, W, X\}$. Name a can be any field that lies arbitrarily deeply within the values exported by another module.

MIR reduces privilege by disallowing all but the white-listed permissions at module boundaries:

$$P_S(m_1, m_2) = \{\langle a, \rho\rangle | a \in API_{m_2}, S \text{ gives } m_1 \ \rho \text{ on } a\}$$

To calculate the privilege reduction across a program that contains several different modules, we lift the privilege definition to a set of subject and target modules:

$$P(M_s, M_t) = \bigcup_{\substack{m_1 \in M_s \\ m_2 \in M_t}} P(m_1, m_2)$$

Based on this, we can define privilege reduction, a metric of the permissions restricted by a privilege-reducing system S, such as MIR:

$$PR(M_s, M_t) = \frac{|P_{base}(M_s, M_t)|}{|P_S(M_s, M_t)|}$$

A higher reduction factor implies a smaller attack surface since the subjects are given privilege to a smaller portion of the available resources. P_{base} is an under-approximation of base privileges, as a source module can in principle import and use any other malicious module that is installed in the execution environment. Consequently, the measured privilege reduction is actually a lower bound of the privilege reduction that MIR achieves in practice.

Permission importance   An important observation is that different permissions $\rho \in R$ may not be equally important. In principle, X may be considered more dangerous than W, which in turn may be considered more dangerous than R. Similarly, in terms of "base" APIs, **console.log** may be considered less important than **fs.write**. For these reasons, MIR's model abstracts these values as coefficients; the same coefficients should be used for meaningful comparisons across different

```
1 let old_srl = srl;
2 srl = {}
3 srl.dec = (...args) => {
4   if (σ(srl.dec, perms.X)) {
5     return old_srl.dec(args);
6   }
7 }
```

(a) Object-wrapping fragment

```
1 var CONTEXT = {
2   eval: mir.wrap(eval, {X}),
3   import: mir.wrap(import, ["log"]),
4   Number: mir.wrap(Number),   //denied
5   Array: mir.BT(Array),       //denied
6   toString: mir.BT(toString),//denied
7   // [another 150 entries denied]
8 }
```

(b) Custom context creation

```
1 function (cxt) {              10   dec: (str) => {
2   var eval = cxt.eval;        11     log.info("[start]");
3   var require = ctx.require;  12     let obj = eval(str);
4   var Number = ctx.Number;    13     return obj
5   var Array = ctx.Array;      14   },
6   var toString = ctx.toString;15   enc: (obj) => {...}
7 let lg = require("log");      16 }
8 lg.LVL = lg.levels.WARN;      17 }
9 exports = {
```

(c) Context rebinding

Fig. 5: MIR's runtime enforcement transformations. MIR's basic `wrapping` traverses objects and wraps fields with inline monitors (a, line 4). A new modified context is created by `wrapping` all values available in a module's original context (b). The modified context is bound to the module by enclosing the module source (half-visible code fragment) in a function that redefines all non-local variable names as function-local ones, pointing to values from the modified context (*Cf.*§7).

configurations. In this paper we set all coefficients to the value "1", considering all permissions to be equally important.

Transitive Permissions Fig. 2's **main** is not allowed to directly call **eval**; however, it can call **eval** indirectly by executing **serial.dec**. This is a limitation of attempting to calculate the privilege reduction dynamically within an environment that specifies opaque abstraction boundaries, such as libraries. Accurately quantifying such transitive privilege requires tracking transitive calls across such boundaries, which requires heavyweight information flow analysis. MIR's privilege reduction quantification does not attempt such an analysis to keep runtime overheads low. As a result, MIR's estimate is necessarily conservative—*i.e.*, MIR reports a lower number than the one achieved in practice.

7 Runtime Permission Enforcement

During program execution, MIR's runtime component enforces the chosen permissions—automatically inferred, developer-provided, or a combination thereof. MIR's enforcement is enabled by the fact that JavaScript features a module-import mechanism that loads code at runtime as a string. Lightweight load-time code transformations operate on the string representation of each module, as well as the context to which it is about to be bound, to insert enforcement-specific wrappers into the module before it is loaded.

MIR's transformations can be grouped into four phases. The first phase simply modifies **import** so that calls yield into MIR rather than the built-in locate-and-load mechanism. For each module, the second phase creates a fresh copy of the runtime context—*i.e.*, all the name-value mappings that are available to the module by default. The third phase binds the modified context with the module, using a source-to-source transformation that re-defines names in the context as library-local ones and assigns to them the modified values. After interpreting the module, the fourth phase further transforms the module's interface so that its client can only access the names—*e.g.*, methods, fields—it is allowed to access.

Base Transform These transformations have a common structure that traverses objects recursively—a base transformation `wrap`, which we review first. At a high level, `wrap` takes an object O and a permission set p and returns a new object O'. Every field f of O is wrapped with a method f'

defined to enclose the permissions for f. Every f' implements a security monitor—a level of indirection that oversees accesses to the field and ensures that they conform to the permissions corresponding to that field. At runtime, f' checks f's permission for the current access type: if the access is allowed, it forwards the call to f; otherwise, it throws a special exception, `AccessControlException`, that contains contextual information for diagnosing root cause—including the type of violation (*e.g.*, R), names of the modules involved, names of accessed functions and objects, and a stack trace.

The result of applying the `wrap` transformation to the object (returned by) `serial` is shown in Fig. 5a.

Context Creation To prepare a new context to be bound to a library being loaded, MIR first creates an auxiliary hash table (Fig. 5b), mapping names to (newly transformed) values: names correspond to implicit modules—globals, language built-ins, module-locals, *etc.* (Tab. 1); transformed values are created by traversing individual values in the context using the `wrap` method to insert permission checks.

User-defined global variables are stored in a well-known location (*i.e.*, a map accessible through a global variable named `global`). However, traversing the global scope for built-in objects is generally not possible. To solve this problem, MIR collects such values by resolving well-known names hardcoded in a list. Using this list, MIR creates a list of pointers to unmodified values upon startup.

Care must be taken with module-local names—*e.g.*, the module's absolute filename, its `export`ed values, and whether the module is invoked as the application's `main` module: each module refers to its own copy of these variables. Attempting to access them directly from within MIR's scope will fail subtly, as they will end up resolving to module-local values of MIR *itself*—and specifically, the module within MIR applying the transformation. MIR solves this issue deferring these transformations for the context-binding phase (discussed next).

Fig. 5b shows the creation of `serial`'s modified context.

Context Binding To bind the code whose context is being transformed with the freshly created context, MIR applies a source-to-source transformation that wraps the module with a function closure. By enclosing and evaluating a closure, MIR leverages JavaScript's lexical scoping to inject a non-bypassable step in the variable name resolution mechanism.

The closure starts by redefining default-available non-local

Tab. 3: Invalid accesses and reasons for incompatibility. For libraries that showed some incompatibility, the table shows the number of inferred permissions (col. 2, 3), the number of invalid permissions (col. 4, 5), and the primary reasons for each incompatibility (col. 6). Permissions include two configurations—base static analysis (col. 2, 4) and augmented with import-time inference (col. 3, 5).

| | Inferred | | Missing | | |
	base	+iti	base	+iti	Reason for incompatibility
through2	18	24	2	0	Gets (inherits) properties through a testing library
set-value	14	14	3	3	Dynamic access to `process.env.npm_package_keywords_15`
periods	6	20	14	0	Function call in `export` index
node-slug	9	10	1	0	Dynamic property access to `module.exports.*`
is-generator	7	7	2	2	`fn.constructor` does not see `fn` as a parameter
he	12	17	6	2	Access to `process.env.npm_package_keywords_4`, `process.env.npm_package_keywords_5`
fs-promise	14	180	135	0	Dynamically computed across all methods of `require("fs")`
file-size	9	10	1	0	Assignment at return: `return exports = plugin()`
total	89	282	175	7	95.5% Improvement due to import-time inference

names as module-local ones, pointing to transformed values that exist in the newly-created context. It accepts as an argument the customized context and assigns its entries to their respective variable names in a preamble consisting of assignments that execute before the rest of the module. Module-local variables (a challenge outlined earlier) are assigned the return value of a call to `wrap`, which will be applied only when the module is evaluated and the module-local value becomes available. MIR evaluates the resulting closure, invokes it with the custom context as an argument, and applies further `wrap` transformations to its return value.

The result of such a source-to-source linking of `serial`'s context is shown in Fig. 5c.

8 Implementation & Evaluation

We built a prototype of MIR for JavaScript, targeting the Node.js ecosystem, and available via `npm i -g @blindorg/mir`. This command makes a few components available. MIR's static analysis component (§5) becomes available as `mir-sa`, implemented as a compiler pass in the Google Closure Compiler [21] in about 2.1 KLoC. MIR's runtime enforcement component (§7) is available as `mir-da`, implemented in about 2.8 KLoC of JavaScript. It supports the configuration of several parameters—for example, `depth` is the number of levels MIR augments, starting from the root of the top objects (Tab. 1). Permissions (§4) exchanged between `mir-sa` and `mir-da` are specified as auxiliary JSON files. To evaluate MIR, we investigate the following questions:

- **Q1** How compatible is MIR with existing code—*i.e.*, what is the danger of breaking legacy programs? Only 7 out of 1044 (0.67%) unique accesses performed in the tests of 50 libraries were not inferred by MIR (§8.1).
- **Q2** How effective is MIR in theory, in terms of privilege reduction? MIR provides 15.6–706× (avg: 224.5×) reduction in terms of permissions over the default execution (§8.2).
- **Q3** How effective is MIR in practice, against *in vivo* and

in vitro attacks? Applied on (1) 16 real attacks in real packages, and (2) several hundred crafted attacks against a highly-vulnerable module, MIR mitigates all attacks that fall under its threat model. In all cases except one, MIR applied default permissions inferred automatically (§8.3).
- **Q4** How efficient and scalable are MIR's inference and enforcement components? MIR's static analysis averages 2.5 seconds per library, comparably to a widely-used linter, and its runtime overhead averages about 1.93% (§8.4).

The remainder of the evaluation section presents the supporting evidence and details underlying these findings.

Libraries and workloads To address Q1, Q2, and Q4 we focus on a set of 50 packages from a curated list [47] that (i) provide comprehensive tests suites executable via `npm test` and (ii) are small enough to enable us to inspect their behavior, and in particular, their use of other packages. Criterion (ii) is crucial for not only reporting quantitative but also qualitative results about the compatibility provided by MIR. As the 50 libraries are quite popular (averaging 4.8M weekly downloads (total: 227M) and 656 dependees (total: 30K)), development investment in their testing infrastructure has been significant, having a variety of different tests that stress different parts of the library—in some cases tests were 10× the size of the corresponding library. To answer Q3, we use a different set of benchmarks that contain vulnerable code. A set of 16 vulnerabilities in real packages that were found *in vivo* [56]; and a series of attacks against a single module crafted *in vitro* to represent the most vulnerable threat within MIR's threat model. To further evaluate the scalability of the static analysis in Q4, we also apply the analysis to 986 npm packages, which comprise 5,826,357 LoC.

Setup Experiments were conducted on a modest server with 4GB of memory and 2 Intel Core2 Duo E8600 CPUs clocked at 3.33GHz, running a Linux kernel version 4.4.0-134. The JavaScript setup uses Node.js v12.19, bundled with V8 v7.8.279.23, LibUV v1.39.0, and `npm` version v6.14.8. MIR's static inference run on Java SE 1.8.0_251, linked with Google Closure v20200927. Libraries were tested on their
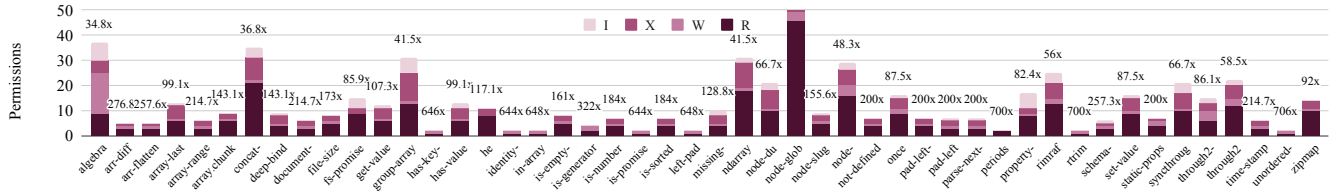
Fig. 6: Inferred Permissions. Number of inferred permissions. Stacked bars show the number of R, W, X, and I permissions; the number above the bar shows the privilege reduction over their default permissions (*Cf.*§8.2).

latest versions at the time of this writing.

8.1 Compatibility Analysis (Q1)

To a large extent, backward-compatibility drives practical adoption of tools such as MIR—if a tool requires significant effort to address compatibility, chances are developers will avoid it despite of any security benefits it provides. This is especially true for MIR, which follows a static analysis approach to meet its threat model of dynamic compromise, rather than one based on dynamic analysis.

Big Picture MIR applies an average of 346 transformations per library, wrapping a total of 25,609 fields. The distribution of accesses is bimodal: (1) not all fields are accessed at runtime—on average only 20.88 (6.06%) of fields are accessed per library; (2) the ones that do get accessed, are accessed multiple times: 794.9 times per field, on average.

Fig. 7 shows a histogram of the number of packages and the number of unique invalid accesses, *i.e.*, the number of benign accesses that happen in the tests but were not found by MIR's inference component. From the 50 libraries on which we applied MIR, only 3 showed some divergence between statically inferred permissions and runtime behavior. All of these 3 libraries diverged on 3 accesses or less.
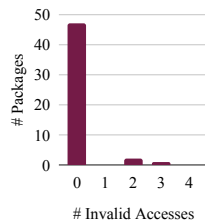

Fig. 7: Invalid accesses. Only 3 out of 50 libraries encountered an invalid access, totalling 7 out of 1044 unique accesses (0.67%) in less than 0.1% of tests (*Cf.*§8.1).

These libraries were tested under multiple workloads each triggering different accesses. Counting with respect to tests, rather than libraries, about 99.9% of all tests performed no invalid accesses. From 1044 total unique accesses, only 7 (0.67%) were invalid; if repetitions are taken into account, 12 out of 16,598 (0.07%) were invalid.

Incompatibility Tab. 3 zooms into modules that show some degree of incompatibility between inferred and required permissions. It shows the number of inferred permissions (col. 2, 3), the number of invalid permissions (col. 4, 5), and the primary reasons for each incompatibility, under two inference configurations: base static analysis (col. 2, 4) and augmented with import-time inference (col. 3, 5). The key take-away is

that MIR's import-time inference is enough to improve results by 95.5%: it brings the total number of non-inferred accesses to seven (from 175, due to static-only inference). We next focus on a few interesting cases.

Modules `zipmap` and `concat-stream` create a local **toString** that they get via the **Object** prototype. MIR's runtime enforcement is configured to not wrap the prototype chain by default, missing the R. Moreover, the runtime reports **toString** as global (capturing a R on `global` and W on **toString**), because it is defined via `var`. (Changing `var` to `let` leads to base incompatibility of 1.) Modules `static-props`, `set-value`, `fs-promise is-generator`, and `he` compute properties they access dynamically. Module `fs-promise` is a good example of this behavior: it traverses the built-in `fs` interface to generate promise-fied wrappers.

These cases highlight a connection between the size of the interface manipulated and the type of incompatibility showing up. When developers use many or all of the properties of an object, they are likely to access them programmatically via runtime reflection; conversely, if they use reflection, they are likely to be accessing many (or all of the) properties. In these cases, static analysis is unlikely to extract these behaviors, which is where import-time inference aids.

The appendix discusses the sources of permission incompatibility between static analysis and dynamic enforcement (§A) in more detail.

Tab. 4: *In vivo* **attacks defended by MIR.** The first column contains vulnerable modules × number of exploits applied to each module. The last column shows the number of permissions inferred by MIR (*Cf.*§8.3).

Module	Version	Attack	CVE	CWE	Perms
Ser/pt	< 1.0	Arbitrary Code Execution	2017-5941	502	13
Nod/ze	0.0.4	Arbitrary Code Execution	2017-5954	502	7
Ser/js	< 3.1	Arbitrary Code Injection	2020-7660	94	58
Saf/al ×4	< 0.4	Sandbox Escaping	(multiple)	265	10
Stat/al ×2	< 2.0	Arbitrary Code Execution	2017-16226	94	9
Fast-redact	< 1.5	Arbitrary Code Execution	—	94	32
Mathjs ×5	< 1.5	Arbitrary Code Execution	(multiple)	94	1
Morgan	< 1.9	Arbitrary Code Injection	2018-3784	502	27

8.2 Permissions and Privilege Reduction (Q2)

In this section we measure the privilege reduction achieved by MIR's static inference.

Number of Permissions Fig. 6 shows MIR's number of inferred R, W, X, and I permissions. Their number ranges between 3–54 (avg: 11.9), spread unevenly between R (6.85), W (1.54), X (3.42), and I (1.52). MIR infers a small number of permissions because developers use only a small subset of the APIs provided by libraries and the broader language environment.

Privilege Reduction Fig. 6 also shows the privilege reduction achieved by MIR—*i.e.*, the ratio of allowed permissions over a library's default set of permissions (§6). Privilege is reduced by up to three orders of magnitude, ranging between 15.6×–706× (avg: 224.5×). MIR's privilege reduction is high because developers use only a small fraction of the available APIs.

8.3 Practical Security Impact (Q3)

This section complements the analysis of MIR's privilege reduction with an investigation of its effectiveness on vulnerable code by applying it to *in vivo* and *in vitro* vulnerabilities. The former set focuses on packages found in the wild [56], for which Table 4 includes the vulnerable versions and MITRE's common enumeration identifiers [42, 43]. The latter set focuses on attacks against a carefully crafted scenario that allows completely unrestricted runtime code evaluation, capturing the most dangerous vulnerability within MIR's threat model. MIR defends against all these vulnerabilities; apart from *Mathjs*, no permissions required manual intervention.

In Vivo For these vulnerabilities, we use the proof-of-concept exploit (PoC) attached to the original vulnerability report; many of them include multiple CVEs or PoCs, all of which we attempt. *Ser/pt*, *Nod/ze*, and *Ser/js* all perform some form of unsafe serialization; their PoCs either (1) import **child_process** to call ls or id, or (2) invoke **console.log**. As none of these is part of the library's permission set, MIR disallows access to these APIs.

Saf/al and *Stat/al* sanitize input prior to calling **eval**. The PoC payloads are similar to the ones described earlier, against which MIR protects. The case of *Stat/al* is interesting because it accepts ASTs rather than strings; the PoC passes **process.env** through the Esprima parser making it vulnerable—which MIR solves by disallowing access to **process.env**.

The *Fast-redact* PoC uses a **constructor** to reach into the **Function** prototype and invoke **open** via **process** bindings. MIR allows the first two steps, crucial for compatibility, but blocks the call to both **bindings** and **open**.

The popular *Mathjs* module includes a math evaluator, which the PoC exploits to **console.log** a message. As *Mathjs* is compiled into a universal-module-definition by webpack (*i.e.*, features no imports and thus is a pure function), we pro-

vided its permissions manually. Despite the fact that *Mathjs* does not import any modules, the PoC calls into Node's APIs via **cos.constructor** (à la *fast-redact*). While MIR does not block the constructor call, it does block **process.env** and **fs.read** from within *Mathjs*.

For most of these benchmarks, MIR blocks the PoC at multiple levels—*e.g.*, even if *Nod/ze*'s import of **child_process** was allowed, MIR would still block **exec**.

In Vitro The goal of this experiment is to evaluate MIR against the worst-case, most-vulnerable single module that falls under its threat model. As MIR aims at protecting against dynamic compromise, the ultimate dynamic-compromise vulnerability would be to expose fully unrestricted access to runtime code evaluation. This is the most general threat that falls within MIR's threat model in the context of a

Attack
write x, global.x
read require.cache
execute process.argv
read process.env
fs.read /etc/passwd
chi/ocess.spawn
call Math.log
read os.EOL

Fig. 9: *In vitro* **attacks prevented by** MIR. MIR defends against handcrafted remote code execution attacks, a subset of which is shown here.

single module, as smaller-scope vulnerabilities for reading and writing fields are subsumed by runtime code execution—and code execution via **eval** comes with the fewest protections: any string passed to **eval** is evaluated at runtime—effectively exposing a runtime interpreter to the attacker.

To implement this scenario, we construct a library e that only exposes an **eval** function. We craft several expressions that attempt to access sensitive interfaces—over 800 attacks amounting to >11K possible accesses targeting built-in JavaScript, Node, and Common.js interfaces. A small subset of these is summarized in Fig. 9: attempts to access the file-system, exfiltrate data through the network, overwrite the module cache, access environment variables, and read the process arguments.

MIR infers two permissions: {e: {eval: RX, exports: W}}. Having MIR enforce these permissions allows only primitive operations within e, such as arithmetic on primitive values and basic string manipulation. It disallows all other operations with side-effects outside e, mitigating these attacks.

8.4 Efficiency and Scalability (Q4)

Static Analysis To evaluate the efficiency and scalability of the static analysis, we run it on 986 npm packages gathered from the most depended-upon packages and the 50 benchmark libraries described above. In total, the analyzed packages comprise 5,826,357 LoC. The static analysis successfully analyses all packages without errors, except for malformed files on which the Google Closure Compiler without our analysis also fails, *e.g.*, files containing syntax errors.

The analysis takes only 42 minutes in total, *i.e.*, an average of 2.5 seconds per npm package. To put these results into
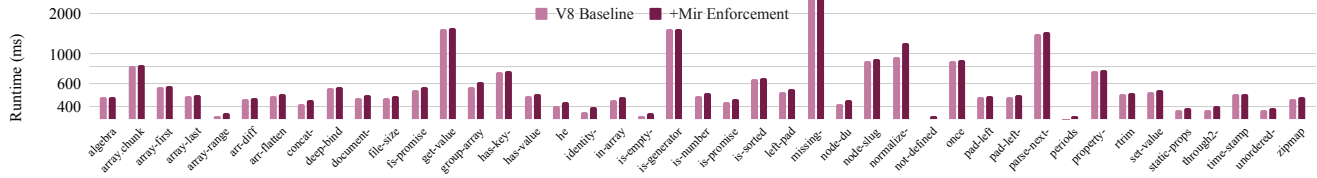
Fig. 8: Performance of MIR's runtime enforcement component. Comparison between vanilla runtime performance with (1) full `RWX` with 3-level depth and across all contexts, and (2) 3-level depth across global-only variables. On average, +3.3ms per library (1.93% slowdown) and +3.2ms (1.62% slowdown) (*Cf.*§8.4).

Tab. 5: Number of inferred `R`, `W`, `X`, and `I` permissions. $1^{st}Q$, $2^{nd}Q$ and $3^{3r}Q$ represent the first quartile, the median, and the third quartile.

Type	Min	$1^{st}Q$	$2^{nd}Q$	$3^{3r}Q$	Max
R	0	13	39	156	27736
W	0	1	4	19	9018
X	0	6	20	78	16638
I	0	2	6	29	16639
Total	0	25	73	295	66559

context, we also measure the performance of a popular linter—a lightweight static analysis pass flagging common human errors. We use `eslint` (v5.0.1) [67] and each library's default linting configuration, falling back to Google's style rules when no such configuration exists. `eslint` ranges between 0.94–6.017*s* across the set of our benchmark libraries, averaging 1.34*s* per library. These results show that the static inference scales well to real-world modules and that its efficiency is comparable to other tools that developers use regularly.

On average, the analysis infers one permission per 10 lines of library code In Table 5, we show statistics about the different types of permissions inferred by MIR. Most of them are `R` and `X` permissions (50.33% and 22.97%, respectively), showing that the client packages rarely write to references outside of their boundaries, *e.g.*, to global variables or the API of the packages they use. While this is not entirely surprising, it shows that the probability of certain types of errors, *e.g.*, concurrent writes to the same global variable name, as addressed by ConflictJS [48], is low for npm packages. The fact that a typical package in our experiments only imports six other entities is somehow consistent with the previous findings of Zimmermann et al. [68] that show that the direct number of dependencies is relatively low for npm packages, but their transitive number is high due to the fragmentation of the ecosystem.

Dynamic Enforcement Fig. 8 shows the performance of MIR's runtime enforcement component for the 50 library benchmarks, as compared with with the performance of the unmodified runtime (V8 baseline).

MIR's full `RWX` enforcement adds between 0.13–4.14*ms* slowdown to executions that range between 324*ms* and 2.77*s*. Slowdowns average about 3.3*ms* per library causing 1.93% slowdown. Based on these results, we do not anticipate a need

for users to trade in runtime security to gain performance.

To better understand the sources of these overheads, we perform a series of micro-benchmarks with tight loops calling several ES-internal libraries without any I/O. The primary source of overhead is the use of JavaScript's `with` construct (§7), The `with` construct dominates overheads, as it (1) interposes on too many accesses, only a fraction of which are relevant, and (2) remains significantly unoptimized, since its use is strongly discouraged by the JavaScript standards [44].

9 Related Work

MIR's techniques touch upon a great deal of previous work in several distinct domains.

Privilege Reduction A number of works have addressed privilege reduction [4, 8, 11, 12, 17, 22, 25, 35, 50, 52, 53, 66], often offering significant automation. This automation often comes at the cost of *lightweight annotations* on program objects—*e.g.*, configurations in Privman [25], `priv` directives in Privtrans [12], tags in Wedge [11], and compartmentalization hypotheses in SOAAP [22]. TRON [10] introduced a permission model similar to MIR, but at the level of processes rather than libraries.

Wedge and SOAAP stand out as offering some automation via dynamic and static analysis, respectively. However, Wedge still requires altering programs manually to use its API, and SOAAP mostly checks rather than suggests policies. In comparison to these works, MIR (1) leverages existing boundaries, and (2) offers significantly more automation.

To ameliorate manual annotations on individual objects, more recent library-level compartmentalization [28, 36, 64] exploits runtime information about module boundaries to guide compartment boundaries. These systems automate the creation and management of compartments, but do not automate the specification of policies through some form of inference. MIR (1) focuses on benign-but-buggy libraries, rather than actively malicious ones, and (2) offers a simplified `RWX` permission model rather than more expressive (often Turing-complete) policies—both in exchange for significant automation in terms of the permissions.

Pyxis [15] and PM [32] reduce the problem of boundary inference to an integer programming problem by defining several performance and security metrics. These systems are

complementary to MIR, as they focus on separating the application code into a sensitive and insensitive compartment to minimize these metrics, while MIR tries to automatically infer and restrict the permissions between different libraries.

The static permission inference of MIR relates to work on statically inferring permissions that an application requires in the Java permission model [27]. Jamrozik et al. [24] describe a dynamic analysis to infer pairs of Android permissions and UI events that trigger the need for a permission. We rely on static inference instead, to avoid the problem of automatically exercising the analyzed code. An important difference to both above approaches is that MIR focuses on `RWX` permissions for specific access paths, instead of the more coarse-grained permissions supported by Java and Android.

Language-based Isolation MIR draws inspiration from language-based isolation techniques [6, 65]: rather than depending on the operating system for protection, MIR enforces access control from with the language runtime—even for libraries that access the file-system and the network. Software fault isolation [65] modifies object code of modules written in *unsafe* languages to prevent them from writing or jumping to addresses outside their domains. Singularity's software-isolated processes [6] ensure isolation through verification. Leveraging memory safety, MIR can be applied in environments with *runtime code evaluation*, for which verification and static transformation might not be an option.

JavaScript Isolation There is significant prior work on JavaScript protection [5, 40, 41, 59, 62], motivated by the pervasive use of multi-party mashups on the client. MIR is unusual in its model, inference, and enforcement: (1) it only allows simple `RWX` permissions rather than Turing-complete functions; (2) its inference is static rather than dynamic; (3) it leverages the library-import mechanism to apply both source- and context-transformations, thus wrapping a library's full observable behavior rather than just its interface.

NodeSentry [63] proposes powerful policies mediating between libraries. Different from MIR, these policies are Turing-complete, are applied manually, and focus on inter-library APIs rather than a library's full observable behavior.

Capabilities Capability systems [31, 55] and object-capability systems [19, 41] place access restrictions by restricting the ability to *name* a resource, essentially treating the object reference graph as an access graph. To make capabilities benefits more widespread, efforts such as Joe-E for Java [39] and Caja for JavaScript [41] restrict popular languages to object-capability-safe subsets. Similar to capability systems, MIR interferes with the program's ability to name a resource—but rather than disabling naming, it augments it with a permission check. MIR does not focus on a language subset, and its static-analysis offers significant automation.

Formalization Prior work has developed formal frameworks for stating and proving strong isolation properties in the context of new languages or subsets of existing languages [1, 18, 20, 37, 38]. MIR takes a different approach, working with an existing language and offering a quantification of privilege reduction rather than an all-or-nothing property. MIR's `RWX` model can be seen as overlaying a type system over the base language semantics, albeit one that is simple and geared towards existing codebases.

Software De-bloating Functionality elimination [51] and, more recently, software de-bloating [9, 23, 26] lower potential vulnerabilities by eliminating unused code in a program. These techniques focus on making subvertible, dormant code inactive—thus are similar in terms of goal to MIR, but approach the problem differently. Rather than eliminating code, MIR makes the code inaccessible at runtime.

Ecosystem-focused Non-academic responses to the challenges of third-party libraries [46, 54, 56] focus on tools that check the program's dependency chain for known vulnerabilities. Dependency locking [45] does not rule out security problems; on the contrary, users forego valuable bug and vulnerability fixes, while experiencing a more convoluted dependency management. Deno, a new JavaScript/TypeScript runtime, offers a coarse-grained allow-deny permission model focusing on the file-system and the network [16]. While this addition confirms that practitioners are still lacking a solution, it lacks MIR's advanced automation and fine granularity.

10 Conclusion

Dynamic library compromise due to bugs or other problems—even when libraries themselves are not malicious—poses a serious security threat in modern software development. MIR addresses it by augmenting the module system with a fine-grained read-write-execute (`RWX`) permission model that allows specifying privilege at the boundaries of libraries. MIR performs an analysis to infer these permissions to not overburden developers and then enforces them using a lightweight load-time code transformation. Finally this paper proposes privilege reduction, a metric that can be used to quantitatively evaluate MIR's impact on prevented permissions.

This paper shows that MIR manages to prevent real attacks without breaking library functionality while (i) requiring minimal user effort, and (ii) imposing minimal overhead on execution and compilation time. MIR is therefore meant to be a lightweight addition to any contemporary developer's toolkit, complementing—rather than replacing—other defense mechanisms. MIR is available for system-wide installation via `npm` [blinded], and its source code is available on GitHub:

```
github.com/blindorg/mir
```

References

[1] Carmine Abate, Arthur Azevedo de Amorim, Roberto Blanco, Ana Nora Evans, Guglielmo Fachini, Catalin

Hritcu, Théo Laurent, Benjamin C. Pierce, Marco Stronati, and Andrew Tolmach. When good components go bad: Formally secure compilation despite dynamic compromise. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, CCS '18, pages 1351–1368, New York, NY, USA, 2018. ACM.

[2] Ajin Abraham. Snyk: Arbitrary code execution in node-serialize. https://snyk.io/vuln/npm:node-serialize:20170208, 2017. Accessed: 2020-03-19.

[3] Ajin Abraham. Snyk: Arbitrary code execution in serialize-to-js. https://snyk.io/vuln/npm:serialize-to-js:20170208, 2017. Accessed: 2020-03-19.

[4] Mike Accetta, Robert Baron, William Bolosky, David Golub, Richard Rashid, Avadis Tevanian, and Michael Young. Mach: A New Kernel Foundation for UNIX Development. In *USENIX Technical Conference*, 1986.

[5] Pieter Agten, Steven Van Acker, Yoran Brondsema, Phu H. Phung, Lieven Desmet, and Frank Piessens. Jsand: Complete client-side sandboxing of third-party javascript without browser modifications. In *Proceedings of the 28th Annual Computer Security Applications Conference*, ACSAC '12, pages 1–10, New York, NY, USA, 2012. ACM.

[6] Mark Aiken, Manuel Fähndrich, Chris Hawblitzel, Galen Hunt, and James Larus. Deconstructing process isolation. In *Proceedings of the 2006 Workshop on Memory System Performance and Correctness*, MSPC '06, pages 1–10, New York, NY, USA, 2006. ACM.

[7] Unknown Author. Snyk: Arbitrary code injection in serialize-javascript. https://snyk.io/vuln/SNYK-JS-SERIALIZEJAVASCRIPT-570062, 2020. Accessed: 2020-03-19.

[8] Niels Avonds, Raoul Strackx, Pieter Agten, and Frank Piessens. Salus: Non-hierarchical memory access rights to enforce the principle of least privilege. In *International Conference on Security and Privacy in Communication Systems*, pages 252–269. Springer, 2013.

[9] Babak Amin Azad, Pierre Laperdrix, and Nick Nikiforakis. Less is more: quantifying the security benefits of debloating web applications. In *28th {USENIX} Security Symposium ({USENIX} Security 19)*, pages 1697–1714, 2019.

[10] Andrew Berman, Virgil Bourassa, and Erik Selberg. Tron: Process-specific file protection for the unix operating system. In *Proceedings of the USENIX 1995 Technical Conference Proceedings*, TCON'95, pages 14–14, Berkeley, CA, USA, 1995. USENIX Association.

[11] Andrea Bittau, Petr Marchenko, Mark Handley, and Brad Karp. Wedge: Splitting applications into reduced-privilege compartments. In *Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation*, NSDI'08, pages 309–322, Berkeley, CA, USA, 2008. USENIX Association.

[12] David Brumley and Dawn Song. Privtrans: Automatically partitioning programs for privilege separation. In *Proceedings of the 13th Conference on USENIX Security Symposium - Volume 13*, SSYM'04, pages 5–5, Berkeley, CA, USA, 2004. USENIX Association.

[13] John N Buxton and Brian Randell. *Software Engineering Techniques: Report on a Conference Sponsored by the NATO Science Committee*. NATO Science Committee; available from Scientific Affairs Division, NATO, 1970.

[14] Mircea Cadariu, Eric Bouwers, Joost Visser, and Arie van Deursen. Tracking known security vulnerabilities in proprietary software systems. In *Software Analysis, Evolution and Reengineering (SANER), 2015 IEEE 22nd International Conference on*, pages 516–519. IEEE, 2015.

[15] Alvin Cheung, Owen Arden, Samuel Madden, and Andrew C Myers. Automatic partitioning of database applications. *arXiv preprint arXiv:1208.0271*, 2012.

[16] Ryan Dahl and the Deno Contributors. Deno. https://deno.land/manual/getting_started/permissions, 2019. Accessed: 2020-06-11.

[17] Willem De Groef, Fabio Massacci, and Frank Piessens. Nodesentry: Least-privilege library integration for server-side javascript. In *Proceedings of the 30th Annual Computer Security Applications Conference*, ACSAC '14, pages 446–455, New York, NY, USA, 2014. ACM.

[18] Christos Dimoulas, Scott Moore, Aslan Askarov, and Stephen Chong. Declarative policies for capability control. In *2014 IEEE 27th Computer Security Foundations Symposium*, pages 3–17. IEEE, 2014.

[19] Sophia Drossopoulou and James Noble. The need for capability policies. In *Proceedings of the 15th Workshop on Formal Techniques for Java-like Programs*, FTfJP '13, pages 6:1–6:7, New York, NY, USA, 2013. ACM.

[20] Sophia Drossopoulou, James Noble, Mark S. Miller, and Toby Murray. Permission and authority revisited, towards a formalisation. In *Proceedings of the 18th Workshop on Formal Techniques for Java-like Programs*, FTfJP'16, New York, NY, USA, 2016. Association for Computing Machinery.

[21] Inc Google. Closure. https://developers.google.com/closure/, 2009. Accessed: 2019-06-11.

[22] Khilan Gudka, Robert NM Watson, Jonathan Anderson, David Chisnall, Brooks Davis, Ben Laurie, Ilias Marinos, Peter G Neumann, and Alex Richardson. Clean application compartmentalization with soaap. In *Proceedings of the 22nd ACM SIGSAC Conference on Computer and Communications Security*, pages 1016–1031. ACM, 2015.

[23] Kihong Heo, Woosuk Lee, Pardis Pashakhanloo, and Mayur Naik. Effective program debloating via reinforcement learning. In *Proceedings of the 2018 ACM SIGSAC Conference on Computer and Communications Security*, pages 380–394, 2018.

[24] Konrad Jamrozik, Philipp von Styp-Rekowsky, and Andreas Zeller. Mining sandboxes. In Laura K. Dillon, Willem Visser, and Laurie A. Williams, editors, *Proceedings of the 38th International Conference on Software Engineering, ICSE 2016, Austin, TX, USA, May 14-22, 2016*, pages 37–48. ACM, 2016.

[25] Douglas Kilpatrick. Privman: A Library for Partitioning Applications. In *USENIX Annual Technical Conference, FREENIX Track*, pages 273–284, 2003.

[26] Hyungjoon Koo, Seyedhamed Ghavamnia, and Michalis Polychronakis. Configuration-driven software debloating. In *Proceedings of the 12th European Workshop on Systems Security*, pages 1–6, 2019.

[27] Larry Koved, Marco Pistoia, and Aaron Kershenbaum. Access rights analysis for java. In Mamdouh Ibrahim and Satoshi Matsuoka, editors, *Proceedings of the 2002 ACM SIGPLAN Conference on Object-Oriented Programming Systems, Languages and Applications, OOPSLA 2002, Seattle, Washington, USA, November 4-8, 2002*, pages 359–372. ACM, 2002.

[28] Benjamin Lamowski, Carsten Weinhold, Adam Lackorzynski, and Hermann Härtig. Sandcrust: Automatic sandboxing of unsafe components in rust. In *Proceedings of the 9th Workshop on Programming Languages and Operating Systems*, PLOS'17, pages 51–57, New York, NY, USA, 2017. ACM.

[29] Tobias Lauinger, Abdelberi Chaabane, Sajjad Arshad, William Robertson, Christo Wilson, and Engin Kirda. Thou shalt not depend on me: Analysing the use of outdated javascript libraries on the web. 2017.

[30] Sebastian Lekies, Ben Stock, Martin Wentzel, and Martin Johns. The unexpected dangers of dynamic javascript. In *Proceedings of the 24th USENIX Conference on Security Symposium*, SEC'15, pages 723–735, Berkeley, CA, USA, 2015. USENIX Association.

[31] H. M. Levy. *Capability Based Computer Systems*. Digital Press, 1984.

[32] Shen Liu, Dongrui Zeng, Yongzhe Huang, Frank Capobianco, Stephen McCamant, Trent Jaeger, and Gang Tan. Program-mandering: Quantitative privilege separation. 2019.

[33] SS Jeremy Long. Owasp dependency check. 2015.

[34] Michael Maass. *A Theory and Tools for Applying Sandboxes Effectively*. PhD thesis, Carnegie Mellon University, 2016.

[35] Marcela S Melara, Michael J Freedman, and Mic Bowman. Enclavedom: Privilege separation for large-tcb applications in trusted execution environments. *arXiv preprint arXiv:1907.13245*, 2019.

[36] Marcela S Melara, David H Liu, and Michael J Freedman. Pyronia: Redesigning least privilege and isolation for the age of iot. *arXiv preprint arXiv:1903.01950*, 2019.

[37] Darya Melicher. *Controlling Module Authority Using Programming Language Design*. PhD thesis, Carnegie Mellon University.

[38] Darya Melicher, Yangqingwei Shi, Valerie Zhao, Alex Potanin, and Jonathan Aldrich. Using object capabilities and effects to build an authority-safe module system: Poster. In *Proceedings of the 5th Annual Symposium and Bootcamp on Hot Topics in the Science of Security*, HoTSoS '18, pages 29:1–29:1, New York, NY, USA, 2018. ACM.

[39] Adrian Mettler, David Wagner, and Tyler Close. Joe-e: A security-oriented subset of java. In *Networked and Distributed Systems Security*, volume 10 of *NDSS'10*, pages 357–374, 2010.

[40] James Mickens. Pivot: Fast, synchronous mashup isolation using generator chains. In *2014 IEEE Symposium on Security and Privacy*, pages 261–275. IEEE, 2014.

[41] Mark S Miller, Mike Samuel, Ben Laurie, Ihab Awad, and Mike Stay. Caja: Safe active content in sanitized javascript, 2008. *Google white paper*, 2009.

[42] MITRE. Common vulnerability enumeration. 2006. Accessed: 2018-11-18.

[43] CWE MITRE. Common weakness enumeration. 2006. Accessed: 2018-11-18.

[44] Mozilla Developer Network. with statement. https://developer.mozilla.org/en-US/docs/Web/JavaScript/Reference/Statements/with, 2020. Accessed: 2020-03-19.

[45] npm, Inc. npm-shrinkwrap: Lock down dependency versions. https://docs.npmjs.com/cli/shrinkwrap, 2012.

[46] Erlend Oftedal et al. Retirejs, 2016.

[47] Andrea Parodi. Awesome micro npm packages (latest commit: Oct 5, 2020; a302e14). https://git.io/JUpA4, 2009. Accessed: 2020-10-07.

[48] Jibesh Patra, Pooja N Dixit, and Michael Pradel. Conflictjs: finding and understanding conflicts between javascript libraries. In *Proceedings of the 40th International Conference on Software Engineering*, pages 741–751. ACM, 2018.

[49] Open Web Application Security Project. Owasp top ten project'17. https://www.owasp.org/index.php/Top_10-2017_Top_10, 2018. Accessed: 2018-09-27.

[50] Niels Provos, Markus Friedl, and Peter Honeyman. Preventing privilege escalation. In *Proceedings of the 12th Conference on USENIX Security Symposium - Volume 12*, SSYM'03, pages 16–16, Berkeley, CA, USA, 2003. USENIX Association.

[51] Martin Rinard. Manipulating program functionality to eliminate security vulnerabilities. In *Moving target defense*, pages 109–115. Springer, 2011.

[52] J. M. Rushby. Design and verification of secure systems. In *Proceedings of the Eighth ACM Symposium on Operating Systems Principles*, SOSP '81, pages 12–21, New York, NY, USA, 1981. ACM.

[53] Jerome H Saltzer. Protection and the control of information sharing in multics. *Communications of the ACM*, 17(7):388–402, 1974.

[54] Node Security. Continuous security monitoring for your node apps, 2016.

[55] Jonathan S Shapiro, Jonathan M Smith, and David J Farber. *EROS: a fast capability system*, volume 33. ACM, 1999.

[56] Snyk. Find, fix and monitor for known vulnerabilities in node.js and ruby packages, 2016.

[57] Manu Sridharan, Julian Dolby, Satish Chandra, Max Schäfer, and Frank Tip. Correlation tracking for points-to analysis of javascript. In *ECOOP 2012 - Object-Oriented Programming - 26th European Conference, Beijing, China, June 11-16, 2012. Proceedings*, pages 435–458, 2012.

[58] Cristian-Alexandru Staicu, Michael Pradel, and Benjamin Livshits. Synode: Understanding and automatically preventing injection attacks on node. js. In *Networked and Distributed Systems Security*, NDSS'18, 2018.

[59] Deian Stefan, Edward Z. Yang, Petr Marchenko, Alejandro Russo, Dave Herman, Brad Karp, and David Mazières. Protecting users by confining javascript with cowl. In *11th USENIX Symposium on Operating Systems Design and Implementation (OSDI 14)*, pages 131–146, Broomfield, CO, 2014. USENIX Association.

[60] Michael Stepankin. [demo.paypal.com] node.js code injection (rce). http://artsploit.blogspot.com/2016/08/pprce2.html, 2016. Accessed: 2018-10-05.

[61] Michael Stepankin. Snyk: Code injection in dustjs-linkedin. https://snyk.io/vuln/npm:dustjs-linkedin:20160819, 2016. Accessed: 2019-03-19.

[62] Jeff Terrace, Stephen R Beard, and Naga Praveen Kumar Katta. Javascript in javascript (js. js): sandboxing third-party scripts. In *Presented as part of the 3rd USENIX Conference on Web Application Development (WebApps 12)*, pages 95–100, 2012.

[63] Neline van Ginkel, Willem De Groef, Fabio Massacci, and Frank Piessens. A server-side javascript security architecture for secure integration of third-party libraries. *Security and Communication Networks*, 2019, 2019.

[64] Nikos Vasilakis, Ben Karel, Nick Roessler, Nathan Dautenhahn, André DeHon, and Jonathan M. Smith. Breakapp: Automated, flexible application compartmentalization. In *Networked and Distributed Systems Security*, NDSS'18, 2018.

[65] Robert Wahbe, Steven Lucco, Thomas E. Anderson, and Susan L. Graham. Efficient software-based fault isolation. In *Proceedings of the Fourteenth ACM Symposium on Operating Systems Principles*, SOSP '93, pages 203–216, New York, NY, USA, 1993. ACM.

[66] Yongzheng Wu, Sai Sathyanarayan, Roland HC Yap, and Zhenkai Liang. Codejail: Application-transparent isolation of libraries with tight program interactions. In *European Symposium on Research in Computer Security*, pages 859–876. Springer, 2012.

[67] Nicholas C. Zakas and ESLint contributors. Eslint—pluggable javascript linter. https://eslint.org/, 2013. Accessed: 2018-07-12.

[68] Markus Zimmermann, Cristian-Alexandru Staicu, Cam Tenny, and Michael Pradel. Small world with high risks: A study of security threats in the npm ecosystem. In
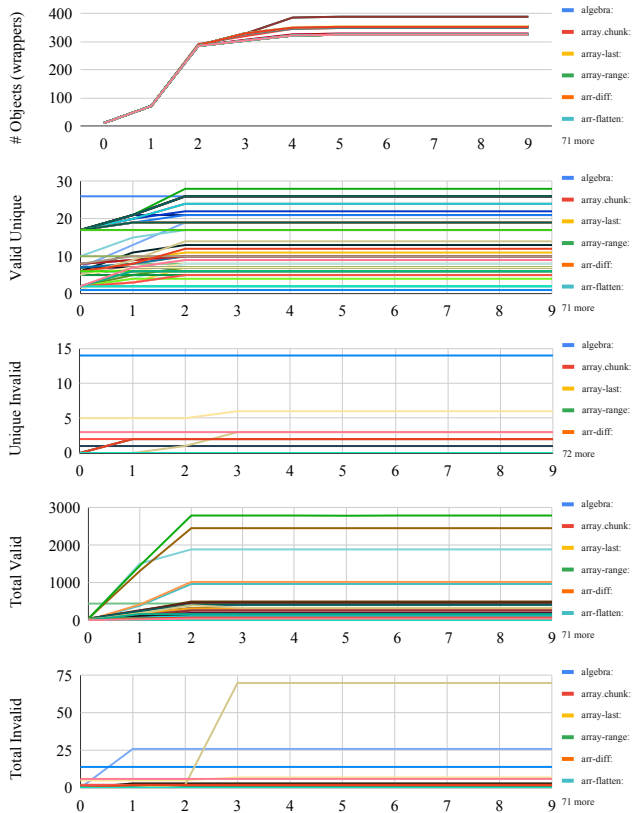
Fig. 10: Depth vs. compatibility. From the top: (a) number of MIR wrappers, (b) unique valid accesses, (c) unique invalid accesses, (d) total valid accesses, and (e) total invalid accesses—all as a function of depth (*Cf.*§8.1).

A   Zooming into Incompatibilities

We now zoom into the sources of incompatibility between the inferred and required permissions, assuming no import-time inference—*i.e.*, between static analysis and dynamic semantics.

Context   Context refers to the broad source of names that are available in the current scope (§4). This includes names defined by the EcmaScript standard (**es**), through an explicit import (**exports**), by the Node.js runtime (**node**), or via global variables (**globs**). A few names seem globally available but are in fact module-locals (**require**). User-defined global variables are not prefixed with global thus requiring special interposition.

Fig. 11 shows several statistics across all 50 libraries as a function of context. In terms of the number of objects wrapped, the vast majority comes from **node** which accounts for 91.1% of all wrappers. In terms of unique number of accesses, for both valid and invalid ones the majority comes from **require**. Taking the number of accesses into account, the majority of accesses comes from **es** and **node** names, whereas the majority of invalid ones comes from **exports**.

These results highlight that MIR gathers most of its correct results around **node** and **es** names. In retrospect, this makes sense: there is no dynamic resolution required to see, for example, that **Math** refers to the standard math library—and this is the most common to refer to it. The majority of invalid accesses come from module exports. This is expected, as a few meta-programming libraries use introspection and reflection capabilities to rewrite entire interfaces and re-export them dynamically.

Depth   Depth refers to how far MIR traverses references starting from the names of objects in the contexts. For example, the access **global.obj.x** is two levels deep and **fs.readFile** is one level deep.

Fig. 10 shows the same statistics as before but relative to depth. There are a few highlights worth noting. First, while the number of objects which MIR's runtime component keeps track of starts growing exponentially, it quickly reaches an average upper bound of about 400. Second, accesses grow quickly for the first couple of levels—objects for some of these libraries are hit several thousand times—but then remain stable. This is usually because many interfaces follow a mostly-flat format where all methods are defined in the top level. Lastly, the majority of invalid accesses remain constant across several levels of depth. This is mostly because these libraries exhibited dynamic behavior at exactly one point— *e.g.*, the export object or a dynamically-computed property (see Tab. 3).
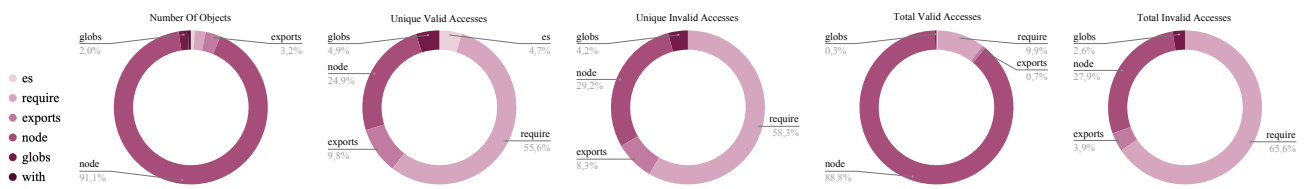
Fig. 11: Context vs. compatibility. From the left: (a) number of MIR wrappers, (b) unique valid accesses, (c) unique invalid accesses, (d) total valid accesses, and (e) total invalid accesses—all as a function of context (*Cf*.§8.1)