

1 **Effectful Software Contracts**

2 CAMERON MOY, PLT @ Northeastern University, USA

3 CHRISTOS DIMOULAS, PLT @ Northwestern University, USA

4 MATTHIAS FELLEISEN, PLT @ Northeastern University, USA

5
6 Software contracts empower programmers to describe functional properties of components. When it comes to
7 constraining effects, though, the literature offers only one-off solutions for various effects. It lacks a universal
8 principle. This paper presents the design of an effectful contract system in the context a language with effect
9 handlers. A key metatheorem shows that contracts cannot unduly interfere with a program's execution. An
10 implementation of this design, along with an evaluation of its generality, demonstrates that the theory can
11 guide practice.

12 CCS Concepts: • **Software and its engineering** → **Semantics**.

13 Additional Key Words and Phrases: effect handlers, software contracts

14 **ACM Reference Format:**

15 Cameron Moy, Christos Dimoulas, and Matthias Felleisen. 2023. Effectful Software Contracts. *Proc. ACM Pro-*
16 *gram. Lang.* 8, POPL (November 2023), 39 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

17 **1 CONTRACTS AND EFFECTS: UBIQUITOUS YET IGNORED**

18
19 For many years, functional programming languages have included constructs for expressing and
20 checking higher-order behavioral contracts [Findler and Felleisen 2002; Keil and Thiemann 2015b;
21 Xu 2014; Xu et al. 2009]. With such contracts, programmers state function specifications and the
22 language's runtime checks them.¹ Concretely, a contract describes the promises a library makes
23 about exported values, and the expectations it imposes on uses [Meyer 1988, 1992]. Put differently,
24 contracts represent agreements between modules about values that flow from one to the other.

25
26 Although these contract systems can deal with a wide range of functional properties, none
27 can systematically express and check properties concerning effects. For example, a library that
28 parallelizes map computations [Dean and Ghemawat 2008] should enforce—but often does not—
29 that the function argument to map is pure. Similarly, when a module exports a function that mutates
30 a hash table, its interface should promise client modules—but often cannot—that it modifies only
31 the given table.

32 The literature is teeming with *ad hoc* solutions: affine contracts [Tov and Pucella 2010] to in-
33 teroperate with substructural type systems; framing contracts [Shinnar 2011] to limit mutation;
34 temporal contracts [Disney et al. 2011] to monitor protocols; authorization contracts [Moore et al.
35 2016] to enforce access control; size-change contracts [Nguyễn et al. 2019] to guarantee termina-
36 tion; trace contracts [Moy and Felleisen 2023] to check properties across multiple calls. All of these

37
38 ¹The presentation here focuses on run-time checks, but some tools [Nguyễn et al. 2018; Xu 2012] can partially verify
39 higher-order contracts at compile time and generate residual run-time checks for unverified properties.

40
41 Authors' addresses: Cameron Moy, PLT @ Northeastern University, Boston, MA, USA; Christos Dimoulas, PLT @ North-
42 western University, Evanston, IL, USA; Matthias Felleisen, PLT @ Northeastern University, Boston, MA, USA.

43 Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee
44 provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and
45 the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses,
46 contact the owner/author(s).

47 © 2023 Copyright held by the owner/author(s).

48 2475-1421/2023/11-ART

49 <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

50 systems use effects in *contracts* to constrain effects in *code*. No existing work supplies a unified
51 approach for doing so, however.

52 This paper presents *effect-handler contracts*, a universal mechanism for expressing and monitor-
53 ing properties of effectful code (Section 2). Its central contribution is a formal semantics of effectful
54 software contracts (Section 3). The model consists of a language where effectful operations are ex-
55 pressed in terms of effect requests and effect handlers [Plotkin and Pretnar 2009], not as primitive
56 operations; in the context of such a language, effect-handler contracts suffice to check a broad
57 class of constraints. An extension to the model (Section 4) formalizes dependent variants of these
58 contract forms. The model is carefully constructed to satisfy an *erasure* property (Section 5), mean-
59 ing that contracts cannot interfere with a program’s computation, other than signaling an error
60 and stopping the world. It also satisfies *blame correctness*, meaning contracts correctly identify
61 components serving values that break the contract assertion.

62 A secondary contribution is an implementation based on this design. The implementation is
63 a standalone language within the Racket ecosystem [Felleisen et al. 2018] that has both effect
64 handlers and effect-handler contracts (Section 6). A thorough literature survey confirms that effect-
65 handler contracts subsume many existing constructs from prior work (Section 7).

66 2 EFFECT-HANDLER CONTRACTS, INFORMALLY

67 This section presents an effect-handler language that extends a functional core with constructs
68 for requesting effects, interpreting effects, and contracts governing effects. While adding ordinary
69 higher-order functional contracts to such a language is straightforward, extending it with contracts
70 on effects requires careful language design.

71 The first subsection presents the syntax of the model, while the second subsection illustrates
72 the semantics informally, using a series of code snippets that add up to a complete example.

73 2.1 Syntax and Informal Semantics

74 The model’s syntax will be presented in three stages: the untyped by-value λ -calculus [Plotkin
75 1975]; an extension with functional contracts [Dimoulas and Felleisen 2011]; and an extension
76 with effect handlers [Plotkin and Pretnar 2009] that also includes syntax for contracts governing
77 effects.

78 **CORE**

79 $e \in \text{Expr} = x \mid b \mid f \mid \langle e, e \rangle \mid \text{if } e e e \mid e e$
80 $b \in \text{Bool} = \text{true} \mid \text{false}$
81 $f \in \text{Fun} = o \mid \lambda x. e$
82 $o \in \text{Op} = \text{fst} \mid \text{snd}$
83 $x, y, z \in \text{Var}$

84 The functional CORE language comes with three built-in data types: Booleans, functions, and
85 pairs. They are eliminated by conditionals, application, and projections, respectively.

86 **CONTRACTS** extends CORE

87 $e \in \text{Expr} = \dots \mid \kappa \mid \text{mon}_j^{k,l} e e$
88 $\kappa \in \text{Con} = b \mid f \mid \langle e, e \rangle \mid e \longrightarrow e$
89 $j, k, l \in \text{Lab}$

90 The CONTRACTS extension reinterprets the base data types as contracts. As a contract, true and
91 false accept and reject all values, respectively. Functions, when used as a contract, are predicates
92 that describe *flat* first-order constraints. A contract pair $\langle e_1, e_2 \rangle$ checks the first component of a
93

value pair with e_1 and the second component with e_2 . A function contract $e_1 \longrightarrow e_2$ protects functions by checking that arguments satisfy e_1 and results satisfy e_2 . A monitor $\text{mon}_j^{k,l} e_1 e_2$ attaches the contract e_1 to the value of e_2 , called the *carrier*. Labels j , k , and l name the contract-defining, value-providing, and value-consuming parties, respectively. These labels are used in contract error messages to *blame* the party responsible for a violation [Findler and Felleisen 2002].

EFFECTS extends **CONTRACTS**

$e \in \text{Expr} = \dots \mid \text{handle}^m e \text{ with } e \mid \text{do } e$

$\kappa \in \text{Con} = \dots \mid e \triangleright e \mid \diamond e$

$m \in \text{Mode} = \triangleright \mid \diamond$

The **EFFECTS** extension introduces two new pieces of syntax related to effect handlers: `handle` and `do`. Evaluating `do v` requests the effect described by v . The evaluation of a request proceeds by searching for the matching handler in the enclosing evaluation context and supplying it with v . Handlers come in two flavors:

`handle▷ e with e_h` is a *main-effect handler*. It interprets only effects performed by ordinary code in the body expression e using the handler e_h .

`handle◇ e with e_h` is a *contract-effect handler*. It interprets only effects performed by contract-checking code in the body expression e using the handler e_h .

Note. Two handler forms are needed to eliminate effect interference. If `handle▷` were to interpret effects at the contract level, a contract could use this channel of communication to change the outcome of a program. By interpreting effects at different levels with different handlers, contract code cannot affect the result of a program. Thus, if a flat contract requests an effect via `do v`, it is not interpreted by a `handle▷` form, even if it is the nearest enclosing handler. ■

Symmetrically, *effect-handler contracts* also demand two constructs, one per level. Both of these monitor a function f that may request effects:

$(e_1 \triangleright e_2)$ is a *main-effect contract*. It ensures that effects performed during the application of f satisfy e_1 and values received from the handler satisfy e_2 .

$(\diamond e)$ is a *contract-handler contract*. It handles, using e , effect requests during the application of f that occur during the dynamic extent of a contract check.

2.2 Examples, Informally

The model suffices to establish essential metaproperties, but illustrating the ideas with such a spartan syntax is too cumbersome. Hence, this section uses ML-like syntactic sugar to present simple examples that illustrate the informal semantics of contracts, effect handlers, and effect-handler contracts. For interesting examples, rather than synthetic ones, see Section 7.1.

Higher-Order Contracts. The RSA algorithm is widely used for sending encrypted messages [Rivest et al. 1978]. Crucially, RSA relies on the difficulty of factoring prime numbers.

Here is the sketch of an RSA-key-generating function, using first-class contracts for higher-order functions to describe the primality constraint:

```
let p_gen_c = is_unit → ⟨is_prime, is_prime⟩
```

```
let k_gen_c = is_unit → ⟨is_key, is_key⟩
```

```
let rsa_c = p_gen_c → k_gen_c
```

```
let rsa : rsa_c = — elided —
```

The contract on `rsa`, attached with a colon, tells the reader that `rsa` is a function that accepts a pair-of-primes-generating thunk and returns a key-pair-generating thunk. Contracts are first-class values and can be defined using `let`. The `· → ·` combinator protects functions by composing an

148 argument contract and a result contract. Furthermore, unlike a type for such a function, contracts
 149 can employ user-defined predicates, e.g., `is_prime`, to check the validity of arguments and results.

150 If the runtime discovers a contract violation—possibly in a distant client module—an error is
 151 signaled identifying the violated contract and blaming the responsible party. Given an invalid
 152 `p_gen` function—say, one that does not generate primes—the contract system identifies the source
 153 of the violation like this:

```
154 > let bad_p_gen () = ⟨3, 4⟩
155 > rsa bad_p_gen ()
156 rsa: contract violation
157   expected: is_prime
158   given: 4
159   blaming: bad_p_gen
160   (assuming the contract is correct)
```

161 **Note.** Effect-handler contracts on their own are about safety properties; they do not suffice to
 162 establish security properties. But, abstractions that enforce security can be built on top of effect-
 163 handler contracts (Section 6.3). ■

164 *Main-Effect Handlers.* A key requirement of RSA is that the generated prime numbers are ran-
 165 dom. To generate random primes, there must be some way to generate ordinary random numbers.

166 A pseudorandom number generator (PRNG) is a deterministic algorithm for generating num-
 167 bers with properties similar to truly random numbers. The interface to most PRNGs is effectful:
 168 generating a random number causes the PRNG’s internal state to change.

169 In a language with effect handlers, a PRNG function collaborates with an effect handler via
 170 effect-request messages to realize state changes. The implementation, of which, is entirely stand-
 171 ard [Pretnar 2015]:

```
172
173 data gen = Gen
174 let rand () = do Gen
175
176 let prng_h req kont =
177   match req with
178   | Gen → λs.kont (prng_extract s) (prng_next s)
179   | _   → λs.kont (do req) s
180
181 let run_with_prng thk seed =
182   (handle▷ (let r = thk () in λ_.r) with prng_h) (prng_init seed)
```

183 The entry point is the `run_with_prng` function. Given a thunk and a random seed, it runs
 184 the thunk in a context that makes random-number generation available. To provide this service,
 185 `run_with_prng` applies the thunk inside `handle▷` with `prng_h`, an effect handler that interprets re-
 186 quests for generating a random number. This handler function takes two arguments: the requested
 187 effect and a continuation that reifies the computation between the origin of the effect request and
 188 the handler (inclusive).²

189 When `thk` needs a random number, it applies `rand`, which in turn, issues a `do Gen` effect re-
 190 quest. The handler of a request packages up the request (`Gen`) and the delimited continuation to
 191 give the handler function. Once the handler function receives these values, it constructs a λ that,

192 ²Effect handlers come in two flavors: *deep* [Cartwright and Felleisen 1994; Plotkin and Pretnar 2009] and *shallow* [Hiller-
 193 ström and Lindley 2018]. In the deep setting, the delimited continuation includes the handler itself; in the shallow one, it
 194 does not. The `handle▷` form uses the deep flavor.

when given the PRNG state, invokes the continuation (kont) with the next random number. The context then applies this function to the PRNG state. If some other effect is requested, the handler propagates the request to an outer handler. Propagation works by applying the continuation to a renewed request. Since `do req` is not a value, the call-by-value semantics ensures that the request is handled before the continuation is resumed.

As a reminder, effect composition is the key benefit of using an effect-handler-based language instead of a language with primitive effects. Since an effect-handler language expresses effects *uniformly*, it is straightforward to reinterpret them, too. Concretely, a programmer can replace or supplement the default PRNG provided by `run_with_prng`, without changing the computation (thk) at all. For example, here is a handler that biases random numbers toward extreme values by squaring them:

```

209 let bias_h req kont =
210   let bias x = if req = Gen then x * x else x in
211   kont (bias (do req))
212
213 let run_with_bias thk = handle▷ thk () with bias_h

```

Assuming the original generator produces reals in $[0, 1]$, this new handler can be composed with the original PRNG to yield a biased generator:

```

217 run_with_prng (λ_.run_with_bias (λ_.rand ())) 0

```

Main-Effect Contracts. In the presence of I/O effects, the contract for `rsa` does not suffice. A program may accidentally (or intentionally) use a prime-generating function that reveals more information than desired:

```

223 let bad_p_gen_v2 () =
224   let ⟨p,q⟩ = — elided — in
225   do (Write "secret.txt" p); ⟨p,q⟩
226

```

Concretely, this prime-generating function writes the secret prime p to a file and thus compromises the RSA key. A contract for `rsa` should prohibit the use of effectful arguments such as `bad_p_gen_v2`.

With main-effect contracts, expressing this constraint is straightforward:

```

232 let p_gen_c_v2 = p_gen_c ⊓ (is_gen ▷ is_real)
233

```

This revised contract is a conjunction; the \sqcap combinator applies each of the two conjuncts, one after another. Consequently, the prime-generating function must satisfy both. While the first conjunct is the original `p_gen_c` contract, the second one describes a constraint on effects. In this example, `is_gen ▷ is_real` ensures that effect requests satisfy `is_gen`, and that the handler passes only values to the continuation if they satisfy `is_real`. Since `is_gen` returns true only for `Gen`, but not `Write`, a use of `bad_p_gen_v2` signals the desired contract violation.³

Main-effect contracts are active only during the dynamic extent of the protected function, and not at any other point. Consider the following handler:

³This example assumes that `data` generates a tag-checking predicate for each variant.

```

246 let printer_h req kont =
247   match req with
248   | Gen -> let res = do Gen in
249             do (Write "secret.txt" res);
250               kont res
251   | _   -> kont (do req)
252
253 let run_with_printer thk = handle▷ thk () with printer_h

```

254 This handler intercepts all random number requests and writes them to the filesystem. In the same
 255 way as `bad_p_gen_v2`, this handler can be used to expose information:
 256

```

257 run_with_prng (λ_.run_with_printer (λ_.rsa p_gen)) 0
258

```

259 However, this program does not result in a contract violation even when the contract of the prime-
 260 generating function is `p_gen_c_v2`. When the prime-generating function requests a random num-
 261 ber, evaluation moves to the body of the handler `printer_h`, which is outside the prime gener-
 262 ator’s dynamic extent. Therefore `is_gen ▷ is_real` is no longer active when `printer_h` writes to
 263 the filesystem.

264 The above behavior is by design; it is critical for assigning correct blame. Recall that a contract
 265 establishes an agreement between a client and server module. According to `p_gen_c_v2`, the `p_gen`
 266 function is responsible only for ensuring that its code does not perform forbidden effects directly,
 267 or indirectly by calling other functions. Client code, including the code that calls `p_gen` and the code
 268 that handles the legitimate effects `p_gen` performs, is not restricted by this part of the contract. In
 269 other words, blaming `p_gen_c_v2` would be wrong even though `printer_h` writes to the filesystem;
 270 it would violate the Blame Correctness Theorem (see Section 5.3).

271 *Contract-Effect Handlers.* The `p_gen_c_v2` contract guarantees that the thunk always receives a
 272 real from the PRNG handler in response to its requests, but it gives no assurance that these reals are
 273 (somewhat) random. A PRNG function that always returns $\frac{1}{2}$ does not cause an error, but yields
 274 a useless prime generator. Statistical tests exist to detect faulty PRNGs [Bassham et al. 2010]; a
 275 contract can employ such tests to detect obviously bad PRNG implementations.
 276

277 Consider the simple-minded test that ensures two consecutive random numbers are different:

```

278 data check = Check is_real
279
280 let rec diff_h prev req =
281   match req with
282   | Check cur → ⟨prev = cur, diff_h cur⟩
283   | _       → ⟨do req, diff_h prev⟩
284
285 let run_with_diff_check thk = handle◊ thk () with (diff_h 0)
286

```

287 When executed via `run_with_diff_check`, contracts can use this test to determine whether the
 288 generated random number differs from the previously generated one. The `handle◊` form is a re-
 289 stricted handler that interprets effects requested in the dynamic extent of a contract check. It does
 290 not get to invoke the delimited continuation of the effect request; instead, the handler function
 291 is expected to return a pair of values: the effect result and a new handler to replace the current
 292 one. Critically, `handle◊` affects only contract-checking code because it can transfer values only to
 293 contract code.
 294

Note. This restriction is similar to that of a runner [Ahman and Bauer 2020] where, informally, a handler *may* invoke the continuation at most once in tail position. Here, the handler *must* invoke the continuation exactly once in tail position.

Direct access to the delimited continuation would permit tampering with the program result and would thus allow interference between program code and contract code. For example, a handler could ignore the continuation completely and return an arbitrary value. ■

Despite the restriction, `handle◇` is quite useful. Adapting `p_gen_c_v2` yields a contract with the desired test:

```
let diff_real x = is_real x && do (Check x)
let p_gen_c_v3 = p_gen_c  $\sqcap$  (is_gen  $\triangleright$  diff_real)
```

Here, `diff_real` requests an effect whose purpose is to check whether the latest argument to a function differs from the most recent one. With this contract, and its corresponding effect handler installed, a PRNG that always returns $\frac{1}{2}$ signals a contract error.

The `p_gen_c_v3` example illustrates why contracts themselves may need to perform effects. Moreover, these effects cannot be locally encapsulated within the contract. In this example, state should persist across multiple calls to the prime-generating function. If the state was locally contained to `p_gen_c_v3`, then subsequent invocations of the prime-generating function would reset the state. Such a limitation would not suffice to express many of the systems in Section 7.

Contract-Handler Contracts. Suppose, for unit testing, the author of `rsa` wants to use a predetermined pool of random numbers for random generation, instead of a PRNG. As such, it is important that the number of times a program requests a random number does not exceed the length of the pool. Thus, the contract needs to keep track of this information:

```
data remaining = Remaining

let rec rem_h k req =
  match req with
  | Remaining  $\rightarrow$   $\langle k, \text{rem}_h (k - 1) \rangle$ 
  | _  $\rightarrow$   $\langle \text{do req}, \text{rem}_h k \rangle$ 

let has_rem req = not (is_gen req) || (do Remaining) > 0
let pool_c k = (is_unit  $\rightarrow$  is_any)  $\sqcap$  has_rem  $\triangleright$  is_real  $\sqcap$   $\diamond(\text{rem}_h k)$ 
let run_with_pool (xs : is_list) (thk : pool_c (length xs)) = elided
```

Like `diff_h` in the previous example, `rem_h` is a contract-effect-handler function. It stores the number of values remaining in the random number pool. Instead of being installed directly using `handle◇`, it is installed by `pool_c`. Specifically, the contract-handler contract $\diamond(\text{rem}_h k)$ installs the interpreting function `rem_h k` using `handle◇`. As such, `run_with_pool` executes `thk` in a context where the `Remaining` effect is interpreted by `rem_h` initialized with the size of the pool.

On its own, a contract-handler contract cannot signal a violation. Rather, it supports other contracts that can. Here, that check happens in the \triangleright conjunct of `pool_c`. When the `thk` requests a random number, `has_rem` checks if there are still random numbers left in the pool. If so, the request is forwarded. Otherwise, an error is raised.

Note. The order of the conjuncts in `pool_c` is relevant. Since `has_rem` requires that the `rem_h` handler is installed, it must come earlier in the list of conjuncts than $\diamond(\text{rem}_h k)$. Since \sqcap applies contracts left-to-right, the right-most conjunct creates the outermost wrapper. ■

3 A FORMAL MODEL OF EFFECT HANDLER CONTRACTS

Defining a semantics amounts to defining an evaluation function that maps programs to answers. Specifying such a function with a reduction relation provides an easy way to prove metatheorems. Following tradition, the specification starts with an extension of the model's syntax to an evaluation syntax (Section 3.1). Next, the presentation of the reduction relation consists of three subsections, corresponding to the three layers of syntax (Sections 3.2 to 3.5). The reduction rules for contracts differ a bit from conventional definitions—namely, flat contracts have *cascading* behavior where, instead of just a Boolean, they can return any arbitrary contract that is then applied. This difference and its purpose are examined in detail (Section 3.6).

3.1 Evaluation Syntax

To start with, the evaluation syntax extends the grammar of expressions and defines the set of values:

EFFECTS (EVAL) extends EFFECTS

$$e \in \text{Expr} = \dots \mid \text{mark}_j^{k,l} v e \mid \text{err}_j^k$$

$$v \in \text{Val} = b \mid f \mid \langle v, v \rangle \mid v \longrightarrow v \mid v \triangleright v \mid \diamond v$$

Expressions include marks and errors, which can arise during evaluation but cannot be expressed in written programs. The expression $\text{mark}_j^{k,l} v e$ states that effects requested by e , and their fulfillment, must satisfy v_κ . In other words, effect requests “passing through” the mark must satisfy the contract. These marks are installed by main-effect contracts.

Next comes the grammar of evaluation contexts. The reduction relation requires three different kinds of evaluation context, each with a different role:

E^\triangleright is the set of *main-executing contexts*, which contain ordinary code and are handled with $\text{handle}^\triangleright$.
 E^\diamond is the set of *contract-executing contexts*, which describe the dynamic extent of contract code and are handled with handle^\diamond .

E is the set of *unrestricted contexts*, which is the union of the (disjoint) sets E^\triangleright and E^\diamond .

Here are the elements of the grammar that are shared between each kind of evaluation context:

EFFECTS (EVAL) extends EFFECTS

$$E \in \text{Ctx} = \langle E, e \rangle \mid \langle v, E \rangle \mid \text{if } E e e \mid E e \mid v E \mid \text{handle}^m E \text{ with } v \mid \text{do } E \mid E \longrightarrow e$$

$$\mid v \longrightarrow E \mid E \triangleright e \mid v \triangleright E \mid \diamond E \mid \text{mon}_j^{k,l} v E \mid \text{mark}_j^{k,l} v E$$

$$E^\triangleright \in \text{Ctx}^\triangleright = \text{--- the above mutatis mutandis ---}$$

$$E^\diamond \in \text{Ctx}^\diamond = \text{--- the above mutatis mutandis ---}$$

And here are the elements that differ between the evaluation contexts:

EFFECTS (EVAL) extends EFFECTS

$$E \in \text{Ctx} = \dots \mid \square \mid \text{mon}_j^{k,l} E e \mid \text{handle}^m e \text{ with } E$$

$$E^\triangleright \in \text{Ctx}^\triangleright = \dots \mid \square \mid \text{handle}^\triangleright e \text{ with } E^\triangleright$$

$$E^\diamond \in \text{Ctx}^\diamond = \dots \mid \square \mid \text{mon}_j^{k,l} E e \mid \text{handle}^\diamond e \text{ with } E \mid \text{handle}^\triangleright e \text{ with } E^\diamond$$

Contract code executes in two syntactic positions: e_κ in $\text{mon}_j^{k,l} e_\kappa e$, and e_h in $\text{handle}^\diamond e$ with e_h . While the former is obvious, the latter might be a surprise. Recall the purpose of handle^\diamond : it interprets effect requests that originate in contract code. By implication, e_h may receive and execute higher-order values originating from contract code. Therefore, it *must* be considered contract code.

The definition of evaluation contexts reflects this reasoning. In particular, E^\triangleright omits productions of the shape $\text{mon}_j^{k,l} E^\triangleright e$ and $\text{handle}^\diamond e$ with E^\triangleright , while E^\diamond omits those for \square . This restriction on E^\diamond ensures that fully formed E^\diamond contexts contain either $\text{mon}_j^{k,l} E e$ or $\text{handle}^\diamond e$ with E .

Finally, formulating the reduction rules and an evaluator requires two more definitions:

$\boxed{\text{EFFECTS (EVAL)}}$ extends EFFECTS

$$\begin{aligned} U \in \text{unhandled} &= \{E_1 \mid E_1 \neq E_2[\text{handle}^m E_3^m \text{ with } v_h]\} \\ s \in \text{stuck} &= \{E[v_f v] \mid v_f \notin \text{Fun}\} \\ &\cup \{E[o v] \mid \delta(o, v) \text{ is undefined}\} \\ &\cup \{E[\text{do } v] \mid E \in \text{unhandled}\} \\ &\cup \{E[\text{handle}^\diamond e \text{ with } v_h] \mid v_h \neq \langle v_a, v_b \rangle, v_h \neq f\} \end{aligned}$$

An unhandled evaluation context lacks a handler for any effect requests that may originate from an expression plugged into the hole. The set of stuck expressions describes those to which no reduction rule applies, i.e., they are not in the domain of the reduction relation. Examples are the application of non-functions to values or an effect request in the hole of an unhandled context. Instead of dealing with stuck expressions in the reduction relation, the evaluator is defined to produce a sensible error when the reduction relation (transitively) reduces to a stuck expression.

The next three subsections present the one-step reduction relation for complete programs using the evaluation contexts [Felleisen et al. 2009; Felleisen and Hieb 1992]. The evaluator is defined by the reflexive-transitive closure of the union of these relations.

3.2 Core Reduction Rules

$$\begin{array}{ll} \text{IF-TRUE} & E[\text{if } v e_1 e_2] \mapsto E[e_1] \text{ if } v \neq \text{false} \\ \text{IF-FALSE} & E[\text{if } \text{false } e_1 e_2] \mapsto E[e_2] \\ \text{APP-LAMBDA} & E[(\lambda x. e) v] \mapsto E[e[v/x]] \\ \text{APP-OP} & E[o v] \mapsto E[\delta(o, v)] \end{array}$$

$$\delta(o, v) = \begin{cases} v_1 & \text{if } o = \text{fst}, v = \langle v_1, v_2 \rangle \\ v_2 & \text{if } o = \text{snd}, v = \langle v_1, v_2 \rangle \end{cases}$$

Fig. 1. Core Reduction Rules

Figure 1 displays the core reduction rules, which are entirely standard [Plotkin 1975]. Conditionals and applications of λ are reduced in the expected manner. A δ metafunction interprets primitive operations [Barendregt 1981]; this choice renders the reduction relation easily extensible.

3.3 Contract Reduction Rules

Figure 2 presents the rules governing contract monitors. Following Dimoulas and Felleisen [2011], $\text{mon}_j^{k,l} e_\kappa e$ monitors the value of e with the contract expression e_κ .

The reduction of many contract expressions is defined by two related rules, prefixed with MON and GRD, respectively. A MON rule checks a first-order property of the carrier. In this model, checking first-order properties amounts to checking whether the top-level *shape* of the carrier value is

442	MON-TRUE	$E[\text{mon}_j^{k,l} \text{ true } v] \mapsto E[v]$
443	MON-FALSE	$E[\text{mon}_j^{k,l} \text{ false } v] \mapsto E[\text{err}_j^k]$
444	MON-FLAT	$E[\text{mon}_j^{k,l} f v] \mapsto E[\text{mon}_j^{k,l} (f v) v]$
445	MON-PAIR	$E[\text{mon}_j^{k,l} \langle v_1, v_2 \rangle v] \mapsto E[\text{err}_j^k]$ if $v \neq \langle v_3, v_4 \rangle$
446	GRD-PAIR	$E[\text{mon}_j^{k,l} \langle v_1, v_2 \rangle \langle v_3, v_4 \rangle] \mapsto E[\langle \text{mon}_j^{k,l} v_1 v_3, \text{mon}_j^{k,l} v_2 v_4 \rangle]$
447	MON-FUN	$E[\text{mon}_j^{k,l} (v_1 \longrightarrow v_2) v] \mapsto E[\text{err}_j^k]$ if $v \notin \text{Fun}$
448	GRD-FUN	$E[\text{mon}_j^{k,l} (v_1 \longrightarrow v_2) f] \mapsto E[\lambda x. \text{mon}_j^{k,l} v_2 (f (\text{mon}_j^{l,k} v_1 x))]$
449		
450		
451		
452		
453		
454		
455		
456		
457		

Fig. 2. Contract Reduction Rules

the expected one. For example, MON-FUN checks whether the carrier is a function. A GRD rule assumes the shape check is satisfied and constructs a wrapper to check either higher-order (as in GRD-FUN) or nested properties (as in GRD-PAIR).

The MON-TRUE and MON-FALSE rules immediately succeed and fail, respectively. When the contract is a predicate f , MON-FLAT applies the predicate to the carrier and uses the result as a contract. Since true and false double as contracts, this cascading of checks works as expected. It is possible to return non-Boolean contracts as well; Section 3.6 explains why this matters.

The GRD rules cover values that need “deep” checking. A pair of contracts distributes over a pair of values. A function contract yields a wrapper value that checks the argument and result contract when the wrapper is applied. The blame labels on the argument monitor are swapped since the argument position of a function contract is contravariant [Finder and Felleisen 2002].

3.4 Effect-Handler Reduction Rules

Take a look at the reduction rules for effect handlers given in Figure 3. When the body of any handler is a value, the effect computation has run its course and the handler is eliminated. Otherwise, one of the Do rules may apply. The unhandled side condition in all of these rules ensures that only the innermost handler is matched with an effect request. Note how Do^\triangleright and Do^\diamond use the special evaluation contexts from Section 3.1 to ensure that the requested effect ($\text{do } v$) originates from either main-program code or contract code.

The Do^\triangleright rule specifies main-program handlers as deep. Concretely, the handler is applied to the effect request and a delimited continuation that includes the handler itself. The evaluation context E^\triangleright may contain marks deposited by main-effect contracts. Two metafunctions, \uparrow and \downarrow , collect the contracts for main-effect requests and fulfillment, respectively. Plugging the raw effect request v into the context created by \uparrow produces an expression that performs all of the necessary contract checks. The same goes for x and \downarrow .

The blame labels flip for \downarrow . The return value, given to the continuation, comes from a handler, which exists in the context of an effect request. As such, swapping the labels is necessary so that blame assignment points to the party that violated the contract [Dimoulas et al. 2011].

By contrast, the Do-PAIR^\diamond and Do-FUN^\diamond rules specify handlers that have no control over the continuation. Furthermore, two rules are needed to distinguish the two contract cases, analogous to the rules for Boolean contracts and predicate contracts. Specifically, the handler e_h in $\text{handle}^\diamond e$ with e_h can reduce to either a function or a pair:

491	HANDLE	$E[\text{handle}^m v \text{ with } v_h] \mapsto E[v]$
492	DO [▷]	$E[\text{handle}^▷ E^▷[\text{do } v] \text{ with } v_h] \mapsto E[v_h e_v (\lambda x. \text{handle}^▷ E^▷[e_x] \text{ with } v_h)]$
493		if $E^▷ \in \text{unhandled}$
494		where $e_v = (\uparrow E^▷)[v]$, $e_x = (\downarrow E^▷)[x]$
495		if $E^▷ \in \text{unhandled}$
496	DO-PAIR [◊]	$E[\text{handle}^◊ E^◊[\text{do } v] \text{ with } \langle v_1, v_2 \rangle] \mapsto E[\text{handle}^◊ E^◊[v_1] \text{ with } v_2]$
497		if $E^◊ \in \text{unhandled}$
498	DO-FUN [◊]	$E[\text{handle}^◊ E^◊[\text{do } v] \text{ with } f] \mapsto E[\text{handle}^◊ E^◊[\text{do } v] \text{ with } (f v)]$
499		if $E^◊ \in \text{unhandled}$
500		
501		
502		
503	$\boxed{\uparrow : \text{Ctx} \rightarrow \text{Ctx}}$	$\boxed{\downarrow : \text{Ctx} \rightarrow \text{Ctx}}$
504	$\uparrow \square = \square$	$\downarrow \square = \square$
505	$\uparrow \langle E, e \rangle = \uparrow E$	$\downarrow \langle E, e \rangle = \downarrow E$
506	$\uparrow \langle v, E \rangle = \uparrow E$	$\downarrow \langle v, E \rangle = \downarrow E$
507
508	$\uparrow(\text{mark}_j^{k,l}(v_1 \triangleright v_2) E) = \text{mon}_j^{k,l} v_1 (\uparrow E)$	$\downarrow(\text{mark}_j^{k,l}(v_1 \triangleright v_2) E) = (\downarrow E) [\text{mon}_j^{l,k} v_2 \square]$
509		
510		
511		

Fig. 3. Effect-Handler Reduction Rules

- (1) In DO-PAIR[◊], the first component is plugged into the evaluation context, which is the continuation of the effect request, and the second component becomes the next handler.
- (2) In DO-FUN[◊], the function is applied to the effect request with the expectation that this new contract expression eventually reduces to a pair. Like MON-FLAT, this rule ensures that contract code is always executed in one syntactic position.

3.5 Effect-Handler Contract Reduction Rules

522	MON-HANDLE [▷]	$E[\text{mon}_j^{k,l}(v_1 \triangleright v_2) v] \mapsto E[\text{err}_j^k] \text{ if } v \notin \text{Fun}$
523	GRD-HANDLE [▷]	$E[\text{mon}_j^{k,l}(v_1 \triangleright v_2) f] \mapsto E[\lambda x. \text{mark}_j^{k,l}(v_1 \triangleright v_2)(f x)]$
524	MARK	$E[\text{mark}_j^{k,l} v_\kappa v] \mapsto E[v]$
525	MON-HANDLE [◊]	$E[\text{mon}_j^{k,l}(\diamond v_h) v] \mapsto E[\text{err}_j^k] \text{ if } v \notin \text{Fun}$
526	GRD-HANDLE [◊]	$E[\text{mon}_j^{k,l}(\diamond v_h) f] \mapsto E[\lambda x. \text{handle}^◊(f x) \text{ with } v_h]$
527		
528		
529		
530		
531		
532		
533		
534		
535		

Fig. 4. Effect-Handler Contract Reduction Rules

Finally, Figure 4 presents the reduction rules governing both kinds of effect-handler contract. The MON rules ensure that the carriers are functions; if not, they signal a violation. If the carriers are functions, the contracts act in a higher-order manner via GRD-HANDLE[▷] and GRD-HANDLE[◊].

The $\text{GRD-HANDLE}^\triangleright$ rule simply installs a mark that constrains effects performed in f . Actually checking these contracts is delegated to Do^\triangleright . Once the dynamic extent of a mark expression is over, the mark itself can be eliminated via the MARK rule.

The $\text{GRD-HANDLE}^\diamond$ rule wraps the carrier in a contract-effect handler, where v_h becomes the handler function. As such, v_h also becomes contract-checking code.

3.6 On the Importance of Cascading Contracts

Flat contracts in this model generalize the ones from the literature to allow cascading. In particular, a flat contract can return any contract, not just a Boolean.

Generalizing flat contracts in this manner is highly useful. Take affine contracts [Tov and Pucella 2010]. An affine contract guarantees that a function is called at most once by keeping track of how many times the function has previously been called. It does so with mutable state. Here is the code for a contract that allows a function to be called at most n times:

```

556 let  $x \xrightarrow{n} y =$ 
557    $\lambda\_.$ let  $r = \text{do } (\text{Ref } n)$  in
558      $((\text{unused}/c \ r) \sqcap x) \longrightarrow y$ 
559
560 let  $\text{unused}/c \ r =$ 
561    $\lambda\_.$ if  $\text{is\_zero } (\text{do } (\text{Get } r))$  then
562     false
563   else
564     do  $(\text{Update } r \ (\lambda n.n - 1))$ ; true
565
566 let  $\text{run\_with\_mut\_refs thk} = \text{handle}^\diamond \text{ thk } ()$  with — elided —

```

The run_with_mut_refs function grants contract code the ability to create, read from, and write to mutable references. Accordingly, the $x \xrightarrow{n} y$ contract specifies a function $x \longrightarrow y$ that can be called at most n times. This property is maintained by allocating a reference containing the remaining number of calls permitted and decrementing that count each time the function is applied.

The $x \xrightarrow{n} y$ contract is a flat contract, not a function contract. When applied to a function, $x \xrightarrow{n} y$ ignores its argument (the function) allocates a cell initialized with n , and *then* returns a function contract. Due to the cascading behavior, this allocation happens exactly once for each function carrier whose monitor enforces the “call at most n times” constraint. Without cascading, this kind of contract is not expressible [Felleisen 1991] in terms of existing contract forms.

4 DEPENDENT CONTRACTS

The contract forms considered thus far cannot deal with dependencies. For example, the result part of a function contract might have to depend on the actual argument. Such dependencies arise frequently in practice.

This section extends the model with dependency: both traditional dependent function contracts, written as $e_1 \implies e_2$, and new dependent main-effect contracts, written as $e_1 \blacktriangleright e_2$. Formally, the syntax is extended as follows:

In this example, `is_in_range` matches on the effect request to determine the upper bound of the random number and uses this to construct a predicate that ensures the generated number is within bounds.

Formalizing dependent main-effect contracts requires a few adjustments to the original semantics. First, two additional rules are needed to reduce monitors containing dependent main-effect contracts. These are analogous to the ones for ordinary main-effect contracts:

$$\begin{array}{ll} \text{MON-HANDLE}^\triangleright & E[\text{mon}_j^{k,l}(v_1 \blacktriangleright v_2) v] \mapsto E[\text{err}_j^k] \text{ if } v \notin \text{Fun} \\ \text{GRD-HANDLE}^\triangleright & E[\text{mon}_j^{k,l}(v_1 \blacktriangleright v_2) f] \mapsto E[\lambda x. \text{mark}_j^{k,l}(v_1 \blacktriangleright v_2)(f x)] \end{array}$$

Next, the \downarrow metafunction has to be extended permit dependencies:

$$\boxed{\downarrow : \text{Val} \times \text{Ctx} \rightarrow \text{Ctx}}$$

$$\begin{array}{l} v \downarrow \square = \square \\ v \downarrow \langle E, e \rangle = v \downarrow E \\ v \downarrow \langle v_1, E \rangle = v \downarrow E \\ \dots \\ v \downarrow (\text{mark}_j^{k,l}(v_1 \triangleright v_2) E) = (v \downarrow E)[\text{mon}_j^{l,k} v_2 \square] \\ v \downarrow (\text{mark}_j^{k,l}(v_1 \blacktriangleright v_2) E) = (v \downarrow E)[\text{mon}_j^{l,k}(v_2 e) \square] \\ \text{where } e = \text{mon}_j^{k,j} v_1 ((\uparrow E)[v]) \end{array}$$

With this revision, \downarrow has access to the raw effect request v . When a mark contains a dependent contract, it must generate the wrapper needed for the effect response. To do so, it applies v_2 to e , where e is the protected effect request. In a lax semantics, $e = (\uparrow E)[v]$. For `indy`, e must also protect v with v_1 where the client label is the contract-defining party.

Finally, given this adapted metafunction, the `Do▷` rule must be adjusted accordingly:

$$\begin{array}{l} \text{Do}^\triangleright \quad E[\text{handle}^\triangleright E^\triangleright [\text{do } v] \text{ with } v_h] \mapsto E[v_h e_v (\lambda x. \text{handle}^\triangleright E^\triangleright [e_x] \text{ with } v_h)] \\ \text{if } E^\triangleright \in \text{unhandled} \\ \text{where } e_v = (\uparrow E^\triangleright)[v], e_x = (v \downarrow E^\triangleright)[x] \end{array}$$

Here, e_x uses the updated metafunction (highlighted) with the raw effect request v supplied.

5 SEMANTIC PROPERTIES

At this point, the definition of a partial evaluation function, also known as an *evaluator*, is straightforward:

$$\boxed{\text{DEPENDENT (PROOF)}} \text{ extends } \text{DEPENDENT (EVAL)}$$

$$\boxed{\text{eval} : \text{Prog} \rightarrow \text{Ans}}$$

$$p \in \text{Prog} = \{e \mid e \text{ is closed}\}$$

$$a \in \text{Ans} = b \mid \text{opaque} \mid \text{err}_j^k \mid \text{err}^\circ$$

$$\text{eval}(e) = \begin{cases} b & \text{if } e \mapsto^* b \\ \text{opaque} & \text{if } e \mapsto^* v, v \notin \text{Bool} \\ \text{err}_j^k & \text{if } e \mapsto^* E[\text{err}_j^k] \\ \text{err}^\circ & \text{if } e \mapsto^* s \end{cases}$$

687 Programs, i.e. closed expressions, are the input to the evaluator. Answers are the output of the
 688 evaluator. Concretely, if a program reduces to a Boolean, the answer is the same Boolean. All other
 689 values yield the opaque token.⁴ This behavior matches that of most REPLs where function values
 690 are printed as an opaque symbol. Two kinds of error can occur during execution: contract errors,
 691 which produce err_j^k , and language errors,⁵ which produce err° .

692 5.1 Well-Definedness

693 Following convention, the first theorem states two properties that ensure the sanity of the reduction
 694 relation. Specifically, eval is a partial function because the reduction relation relates each
 695 program to at most one answer, and because there are programs with unbounded reduction sequences.
 696

697 **Theorem 5.1 (Functional Evaluation).** Two facts about the evaluator hold:

- 700 (1) The eval relation is a partial function.
 701 (2) If e is a program, then either (i) $\text{eval}(e)$ is defined or (ii) the reduction sequence starting
 702 with e is unbounded.

703 **PROOF.** See Appendix B. □

704 5.2 Erasure

705 The key property of interest for the model is *contract erasure*. Contracts serve one purpose, namely,
 706 to detect violations of specifications. Therefore, the output of a correct program should not depend
 707 on the presence or absence of contracts. In short, contracts must not interfere with program
 708 execution—other than possibly signaling an error. Non-interference in the presence of effects is
 709 critical for modular reasoning [Oliveira et al. 2012].

710 Stating the erasure theorem requires the definition of the (natural) erasure function \mathcal{E} for contract
 711 monitoring:

$$712 \boxed{\mathcal{E} : \text{Expr} \rightarrow \text{Expr}}$$

$$713 \mathcal{E}(b) = b$$

$$714 \mathcal{E}(x) = x$$

$$715 \mathcal{E}(\lambda x. e) = \lambda x. \mathcal{E}(e)$$

716 ...

$$717 \mathcal{E}(\text{mon}_j^{k,l} e_\kappa e) = \mathcal{E}(e)$$

$$718 \mathcal{E}(\text{handle}^\diamond e \text{ with } e_h) = \text{handle}^\diamond \mathcal{E}(e) \text{ with } \mathcal{E}(e_h)$$

$$719 \boxed{\mathcal{E}^+ : \text{Expr} \rightarrow \text{Expr}}$$

$$720 \mathcal{E}^+(b) = b$$

$$721 \mathcal{E}^+(x) = x$$

$$722 \mathcal{E}^+(\lambda x. e) = \lambda x. \mathcal{E}^+(e)$$

723 ...

$$724 \mathcal{E}^+(\text{mon}_j^{k,l} e_\kappa e) = \mathcal{E}^+(e)$$

$$725 \mathcal{E}^+(\text{handle}^\diamond e \text{ with } e_h) = \mathcal{E}^+(e)$$

726 **Theorem 5.2 (Erasure).** If $\text{eval}(e) = b$ then $\text{eval}(\mathcal{E}(e)) = b$.

727 **PROOF.** The proof of erasure proceeds by a simulation argument with the following simulation:

$$728 \lambda x. f x \sim \tilde{f}$$

$$729 \text{handle}^\diamond e \text{ with } e_h \sim \text{handle}^\diamond \tilde{e} \text{ with } \tilde{e}_h$$

$$730 \text{handle}^\diamond e \text{ with } e_h \sim \tilde{e}$$

731 ⁴For simplicity, the function turns pairs into opaque, too.

732 ⁵In essence, such errors are violations of the runtime system's contracts.

$$\begin{aligned} \text{mon}_j^{k,l} e_\kappa e &\sim \tilde{e} \\ \text{mark}_j^{k,l} v e &\sim \tilde{e} \\ &\dots \end{aligned}$$

By convention, a metavariable with a tilde such as \tilde{e} is in simulation with its plain counterpart e . See Appendix C for details. \square

Technically, non-termination is the one contract effect that can affect a program's behavior. So long as contracts contain code in a Turing-complete language, this effect is unavoidable. As stated, Theorem 5.2 holds because the antecedent rules out non-terminating contracts.

While the syntax design already clarifies that there are two separate, disjoint levels of effect handling, the proof for Theorem 5.2 confirms this claim. Additionally, main code cannot be serviced by contract-effect handlers. A small adjustment to the erasure function, defined above as \mathcal{E}^+ , makes it possible to state the claim formally.

Corollary 5.3 (No Effect Interference). If $\text{eval}(e) = b$ then $\text{eval}(\mathcal{E}^+(e)) = b$.

PROOF. Follows directly from the proof of Theorem 5.2. \square

Establishing the erasure theorem is straightforward in a pure setting, yet difficult to achieve in a language with effects. Ensuring erasure means contract code must not interfere with the main program directly or indirectly via effects. A language with effect-handler contracts poses the additional problem of having to grant contract code the right to interact with effects, while also imposing constraints on such interactions.

A physicist would describe the model as being in an “unstable equilibrium;” an author of a types textbook would use the word “brittle” and compare the design to the Hindley-Milner algorithm for type inference. Directly put, designing a language semantics that satisfies contract erasure demands balancing expressive power with preventing interference. The model presented here achieves this delicate balance, as the theorem and Section 7 show. Limiting the expressive power any further makes programming inconvenient and would neglect many existing use cases. However, experiments adding more power to the model show that many natural extensions violate erasure.

For example, consider a naive design where the reduction relation for handlers merges the two levels of effect handling:

$$E[\text{handle } E_k[\text{do } v] \text{ with } v_h] \mapsto E[v_h v (\lambda x. \text{handle } E_k[x] \text{ with } v_h)] \text{ if } E_k \in \text{unhandled}$$

Instead of restricting the evaluation context in the body of the handler, this rule uses the unrestricted context E_k . Such a rule violates contract erasure as the following program demonstrates:

$$\text{handle } (\text{mon}_j^{k,l} (\text{do false}) \text{ true}) \text{ with } \lambda x. \lambda y_k. x$$

While the original program evaluates to false, erasing the contract yields a variant whose value is true. Similarly, modifying Do^\triangleright to use E , or modifying Do-FUN^\diamond to give direct access to the continuation, both result in counterexamples to erasure because they produce a rule equivalent to the one above.

As remarked in Section 2.1, introducing main-handler contracts like this

$$E[\text{mon}_j^{k,l} (\triangleright v_h) f] \mapsto E[\lambda x. \text{handle}^\triangleright (f x) \text{ with } v_h]$$

also violates erasure. Here is a counterexample:

$$\text{handle}^\triangleright ((\text{mon}_j^{k,l} (\triangleright (\lambda y. \lambda z_k. \text{false}))) (\lambda x. \text{do } x)) \text{ true}) \text{ with } \lambda y. \lambda z_k. y$$

785 Again, the original program evaluates to false, while its erased variant yields true. In short, GRD-
 786 HANDLE[◇] cannot be generalized.

787

788 5.3 Blame Correctness

789 The final property to consider is blame correctness, that is, whether a failing monitor assigns blame
 790 to the component that serves a faulty value. In the context of the model, the Do[▷] reduction deserves
 791 particular attention. Like the rule for monitoring first-class functions, the reduction for main-effect
 792 handling switches the order of blame labels as it pushes the relevant contracts down the handler's
 793 continuation ($v \downarrow E^\triangleright$). The question is—as it was for the original work on higher-order (dependent)
 794 function contracts [Findler and Felleisen 2002]—whether this switch is correct. As Dimoulas et al.
 795 [2011] show, the answer is a blame correctness theorem.

796 By now, the strategy for proving blame correctness is reasonably standard. The first step is to
 797 introduce ownership labels on expressions, values, and evaluation contexts. Intuitively, an expres-
 798 sion $|e|^l$ denotes that the owner of e is component l . See Appendix D for details.

799 The second step is to adjust the reduction relation so that ownership changes when a value
 800 crosses from one component to another. Crossing may either add or drop a label from a value.
 801 The reduction drops a label when the crossing involves a contract check, meaning the value is
 802 vetted and “absorbed” by a new host component. A blame label is added when the crossing does
 803 not involve a check, meaning the value becomes co-owned by several distinct components. It is
 804 critical that the ownership labels do not affect the semantics proper.

805 The third and final step is to show that when a monitor is about to check a value, the latest
 806 ownership label of the value is the same one that the monitor uses to assign blame.

807

808 **Theorem 5.4 (Blame Correctness).** For all e if $l_o; \emptyset \vdash e$, if $e \mapsto_o^* E[\text{mon}_j^{k,l} v_\kappa v]$, then $v = |v'|^k$.

809

810 **PROOF.** The proof uses the standard subject-reduction technique [Curry and Feys 1958; Wright
 811 and Felleisen 1994] and a consistency judgment for the ownership annotations. The judgment
 812 $l; \Gamma \vdash e$ says that e is well-formed if its owner is l , given an environment Γ that maps variables to
 813 their owners. Importantly, if a program is well-formed under the default owner l_o , then for any
 814 monitors it contains, the owner of the carrier matches the server label of the monitor. Subject
 815 reduction shows that this consistency is preserved across reduction sequences, and hence, if a
 816 monitor check fails, blame is assigned to the correct component. See Appendix D for the proof. \square

817 The labeled reduction semantics is indeed equivalent to the unlabeled one after erasing owner-
 818 ship labels ($\mathcal{O}(\cdot)$) from the first one.

819

820 **Proposition 5.5 (Ownership Erasure).** For all labeled e , $e \mapsto_o^* e'$ iff $\mathcal{O}(e) \mapsto^* \mathcal{O}(e')$.

821

822 **Note.** The stronger *complete monitoring* property states that all channels of communication
 823 between modules can be monitored using contracts [Dimoulas et al. 2012]. The presented model
 824 does not satisfy complete monitoring. As Section 7 explains, the intent of effect-handler contracts
 825 is to be a low-level mechanism for implementing other constructs. Complete monitoring is more
 826 relevant to prove for these higher-level contract systems, not the low-level target. \blacksquare

827

828 6 EFFECT RACKET

829 Rapidly moving from a model to a full-fledged programming language calls for (1) a programmable
 830 production-level language with (2) linguistic constructs for realizing effect handlers easily and (3) a
 831 well-developed higher-order contract system. Racket is such a language [Felleisen et al. 2018; Find-
 832 ler and Felleisen 2002; Flatt and PLT 2010; Flatt et al. 2007]. This section presents effect/racket,

833

```

834 #lang effect/racket
835
836 (effect ref (v))
837 (effect ref-get (r))
838 (effect ref-set (r v))
839
840 ;; Store → Service
841 (define (ref-service store)
842   (handler
843     ;; Creates a reference
844     [(ref v)
845      (define-values (r new-store) (store-allocate store init))
846      (with ((ref-service new-store))
847            (continue* r))]
848
849     ;; Returns a reference's value
850     [(ref-get r)
851      (continue (store-get store r))]
852
853     ;; Sets a reference's value
854     [(ref-set r v)
855      (with ((ref-service (store-set store r v)))
856            (continue* (void)))]))
857
858 > (with ((ref-service empty-store))
859     (define r (ref 0))
860     (ref-set r (add1 (ref-get r)))
861     (ref-get r))
862 1

```

(a) Program

(b) REPL

Fig. 5. Mutable References with effect/racket

a language with effect handlers and a full contract system (Section 6.1). Following the precedent of typed/racket, the language is implemented as a library [Tobin-Hochstadt et al. 2011] (Section 6.2). The language implementation validates that the model can be realized and therefore may help guide implementers of other effect-handler languages.

6.1 The Language, By Example

The section is organized like Section 2.2, but uses different examples to keep things interesting.

Main-Effect Handlers. As an introductory example, consider implementing ML's first-class mutable references, using effect handlers. References come with a `ref` constructor and two elimination forms: `ref-get` and `ref-set`. In `effect/racket`, each form demands the declaration of a corresponding effect: one for allocating a reference cell, one for getting its value, and yet another for assigning to a cell. Declaring an effect makes the effect name available both for requesting the effect and, within a handler, for interpreting the effect.

883 Figure 5 displays the code for both the effect declarations and the effect handler. The handler
 884 function, dubbed a *service* for references, comes with three clauses, one per declared effect; all
 885 other effects are propagated automatically. Furthermore, the handler form binds two identifiers to
 886 delimited continuations: `continue`, for resuming handling in a deep manner; and `continue*`, for
 887 handling effects in a shallow manner. Otherwise, the handler function uses standard techniques
 888 for implementing a store in this setting [Cartwright and Felleisen 1994; Pretnar 2015].

889 Any language in the Racket ecosystem, including `effect/racket`, is easily equipped with a read-
 890 eval-print loop (REPL). By running the `effect/racket` program, the definition of effects and serv-
 891 ices becomes available for interactive experimentation. Figure 5b shows how to install the handler
 892 function using the `with` form. In the context of this `with` expression, it is now possible to allocate
 893 a numeric reference cell, to increase its value by 1, and then to retrieve this value.

894 *Main-Effect Contracts.* Suppose a programmer wishes to write a library function that guaran-
 895 tees a frame condition. To make this concrete, the function guarantees that it manipulates only a
 896 specific, given reference cell during the dynamic extent of any call. A good name for this contract
 897 would be `mutates-only/c`, and here is how the library’s interface would state that guarantee:

```
899 (provide
900   (contract-out
901     [ref-restore
902       ;; Runs (f r), restores the content of r, and
903       ;; returns the value of r that f stores there.
904       (->i ([r ref?]
905            [f (and/c (mutates-only/c r)
906                    (-> ref? any/c))])
907            [result any/c])]))
```

908 The function contract is a standard indy contract [Dimoulas et al. 2011] that governs two arguments—
 909 `r` and `f`—and promises nothing about its result. The new part is the contract for `f`, which says that
 910 (1) `f` is a function from a reference cell to any value and (2) it may mutate only `r`.

911 The frame contract is a rather straightforward instance of a main-effect contract:

```
912 (define (mutates-only/c r-ok)
913   (define (effect-ok? e)
914     (match e
915       [(ref-set r _) (equal? r r-ok)]
916       [_ true]))
917   (->e effect-ok? any/c))
```

918 The `mutates-only/c` function takes a reference cell as an argument and returns a main-effect
 919 contract that permits only writing to the given cell and no other one. The two-part contract tells
 920 a reader that requested effects must satisfy the `effect-ok?` predicate and that values returned by
 921 the handler can be anything. According to `effect-ok?`, any write effect must be to a reference cell
 922 that is equal to `r-ok`. All other effects are permitted.

923 *Contract-Effect Handlers.* Equipped with reference cells, it is now possible to transliterate the
 924 affine-function contract from Section 3.6 into running code. Figure 6 shows the implementation of
 925 \multimap as a contract exported from a library.

926 Since the contract relies on the use of reference cells at the contract level, it is mandatory to lift
 927 the service from the main level to the contract level; see Figure 6 (lines 27–35). The contract-handler
 928 form does not make the delimited continuations available; instead each arm must return a pair of

931

```

932 1 #lang effect/racket
933 2
934 3 (provide
935 4   ;; Store  $\rightarrow$  Service
936 5   ;; Service to be installed for uses of  $\multimap$ .
937 6   ref-contract-service
938 7
939 8   ;; Natural Contract Contract  $\rightarrow$  Contract
940 9   ;; Returns a contract for a function that is called at most n times.
941 10   $\multimap$ )
942 11
943 12 (define ( $\multimap$  n dom cod)
944 13   (self/c
945 14     ( $\lambda$  _
946 15       (define r (ref n))
947 16       ( $\rightarrow$ i ([x dom])
948 17         #:pre (unused? r)
949 18         [result cod])))
950 19
951 20 (define (unused? r)
952 21   ( $\lambda$  _
953 22     (define m (ref-get r))
954 23     (cond
955 24       [(zero? m) false]
956 25       [else (ref-set r (sub1 m)) true])))
957 26
958 27 (define (ref-contract-service store)
959 28   (contract-handler
960 29     [(ref v)
961 30       (define-values (r new-store) (store-allocate store v))
962 31       (values r (ref-contract-service new-store))]
963 32     [(ref-get r)
964 33       (values (store-ref store r) (ref-contract-service store))]
965 34     [(ref-set r v)
966 35       (values (void) (ref-contract-service (store-set store r v)))]))

```

Fig. 6. Affine-Function Contracts with effect/racket

values: the value to be supplied to the delimited continuation, and a new handler to be installed around the continuation.

Using `ref-contract-service`, both \multimap and `unused?` can be defined using Racket's existing contract library, with reference effects performed as needed; see Figure 6 (lines 12–25). As in Section 3.6, \multimap must use cascading to allocate a reference for affine functions at the right time; the presented code realizes this constraint using the `self/c` combinator, which when protecting a value `v`, applies a function to `v` and uses the result to protect `v`—just like flat contracts in the model. Here, the function given to `self/c` returns the expected function contract. For \multimap , the value `v` is not needed and is discarded.

Contract-Handler Contracts. A function is *reentrant* if it can call itself recursively, directly or indirectly. A contract-handler contract can check for non-reentrancy by prohibiting recursive calls during a function’s dynamic extent. Implementing such a constraint requires both a contract-handler to mark the dynamic extent of a function call and a contract-handler contract.

Here is the contract handler and the `nre?` effect declaration:

```

986 (effect nre? ())
987
988 (define nre-service
989   (contract-handler
990     [(nre?) (values false nre-service)]))

```

The contract for a non-reentrant function `f` installs this handler, like thus:

```

993 1 (provide
994 2   (contract-out
995 3     [f (and/c
996 4       (with/c nre-service)
997 5         (->i ([l vector?])
998 6           #:pre (nre? #:fail true)
999 7             [result vector?]))]))

```

When a client applies `f`, the precondition requests an `nre?` effect (line 6). If this returns `false`, the function might already be running; otherwise the `#:fail` option, which provides a default value if no matching handler is installed, returns `true`. Once the precondition check passes, the second wrapper sets up a contract-handler contract (line 4). Thus, if `f` were to call itself, the contract prohibits it because the installed `nre-service` supplies `false`.

6.2 The Implementation, An Overview

The implementation of `effect/racket` consists of about 1,100 lines of code. Most of these lines compose elements from existing libraries. For example, effect handlers themselves are implemented as thin wrappers around Racket’s existing library of delimited control operators [Flatt et al. 2007]. Other pieces of the implementation ensure that Racket’s effectful primitive operations are inaccessible to programs in `effect/racket`. After all, a main-effect contract would be meaningless if certain primitive effects cannot be reinterpreted.

One critical aspect of the implementation concerns the key assumption of the model in Section 3, which demands that handlers can detect whether an effect request originates from within main code or contract code. Formally, this idea is encoded via special evaluation contexts; see Section 3.1. As it turns out, Racket’s contract system already provides a mechanism for determining whether code is executing inside a contract [Andersen et al. 2018]. Specifically, contracts set up continuation marks [Clements et al. 2001] that delineate contract-specific code from user code. Thus, the implementation of the effect handler forms inspect the delimited continuation and look for this mark to determine whether the effect should be handled. As a result, `effect/racket` does not necessitate any modifications to Racket’s contract system.

As a language in Racket’s ecosystem, `effect/racket` inherits the module system, too, which raises the interoperability issue. Indeed, the preceding examples already rely on the module system, showing that `effect/racket` modules can export functions with effectful software contracts. In addition to full interoperability with other `effect/racket` modules, the language has a shallow form of interoperability with plain Racket modules. Following the terminology of Matthews and Findler [2007], the interoperability uses a first-order natural, higher-order lump-embedding. By implication, first-order values can freely flow from an `effect/racket` module to a foreign module

1030 and back; in contrast, higher-order values are wrapped in an opaque structure so they become
 1031 unusable.

1032 In summary, the implementation effort reveals that the addition of effectful software contracts
 1033 to an effect-handler language is rather straightforward, with the exception of effect stratification.
 1034 Assuming the erasure property is desirable, an implementer must add a mechanism that demar-
 1035 cates the dynamic extent of contract-checking code.

1036 6.3 Restricting Handlers

1037 As presented, handlers have unlimited access to interpose and reinterpret all effects. This means
 1038 library authors have no guarantee about how their effects are interpreted. Others have recognized
 1039 this lack of abstraction safety and have proposed solutions, especially in typed settings [Biernacki
 1040 et al. 2017; Brachthäuser et al. 2022; Leijen 2013; Xie et al. 2020; Zhang and Myers 2019].

1041 An alternative design can rectify this problem easily. Racket gives programmers the ability to
 1042 attach metadata to continuations via continuation marks [Clements and Felleisen 2004; Flatt and
 1043 Dybvig 2020]. To prevent other parties from arbitrarily tampering with this information, the lan-
 1044 guage only permits access to continuation marks via *keys*. Racket also uses this mechanism to
 1045 limit how much of the continuation a program can capture and abort [Felleisen 1988; Flatt et al.
 1046 2007; Sitaram and Felleisen 1990]. These keys are first-class unforgeable values. If a module does
 1047 not export its key, then no other party can view or update the information associated with that
 1048 key. To prevent interception, a module can just keep an effect key internal. Effectively, a key is a
 1049 capability [Dennis and Van Horn 1966].

1050 Instead of an arbitrary match pattern, a handler could be restricted to explicitly provide a set of
 1051 “effect keys” that it can interpret. Similarly, main-effect contracts would have to include a set of
 1052 effect keys instead of an arbitrary predicate over all effect requests.

1053 There is a downside to this approach; it eliminates a useful class of contracts like those for
 1054 purity. A contract that guarantees purity must, by definition, be able to interpose on *all* effects.
 1055 This includes ones that are kept hidden. Unsurprisingly, there is a trade-off between security and
 1056 expressiveness.

1057 If desired, though, this kind of restriction can be built on top of effect/racket; the language is
 1058 flexible so that it can serve as foundation upon which other abstractions can be constructed. Thus,
 1059 language implementors can choose the design that fits their situation.

1060 7 EVALUATION AND RELATED WORK

1061 The introduction of this paper claims that effect-handler contracts are a universal mechanism. An
 1062 evaluation of such a claim must show that the model and its full-scale implementation cover all
 1063 existing work.⁶ Additionally, such related research must be analyzed and systematically compared.
 1064 As such, this section consists of two pieces: (1) an evaluation of effect-handler contracts with
 1065 respect to a survey of existing literature; and (2) a summary of each piece of related research and
 1066 how it compares to this paper.

1067 7.1 Analysis

1068 Table 1 presents an overview of the existing literature. It explicates the many overlapping problems
 1069 that the various papers address. Rows correspond to existing pieces of literature, and columns
 1070 correspond to concrete properties that at least one system can express.

1071 ⁶Effect-handler contracts subsume only the low-level contract aspects of existing work—nothing more. All of these papers
 1072 build sophisticated systems on top of these low-level constructs. These contributions are orthogonal to, and not subsumed
 1073 by, effect-handler contracts.

Table 1. Detailed Comparison Matrix

	ALLOW CALL	EXCEPTIONS	FRAMING	GHOST STATE	MUST CALL	NON-REENTRANT	PURE	RESTRICTED EFFECT	TERMINATION	UNION CONTRACTS
Chalin et al. [2006]		•	•	•			•		•	
Tov and Pucella [2010]				•						
Shinnar [2011]		•	•	•			•			
Disney et al. [2011]	•		•	•	•	•				
Keil and Thiemann [2015a]										•
Scholliers et al. [2015]	•			•	•					
Moore et al. [2016]	•			•	•	•				
Dimoulas et al. [2016]	•			•	•					
Bañados Schwerter [2016]		•					•	•		
Williams et al. [2018]										•
Nguyễn et al. [2019]									•	
Moy and Felleisen [2023]	•		•	•	•	•				
Effect-Handler Contracts	•	•	•	•	•	•	•	•	•	•

A concise description of these properties follows:

ALLOW CALL A function may be called only during the dynamic extent of another function.

EXCEPTIONS Only specified exceptions may be raised during a function call. This property is the dynamic analogue to Java’s checked exceptions.

FRAMING Mutations are restricted to specified memory locations.

GHOST STATE Values are associated with a mutable reference that is used to check conformance with a protocol.

MUST CALL A function *must* be called during the dynamic extent of a call to another function.

NON-REENTRANT A function must not call itself recursively.

PURE No effects—other than non-termination and error signals—are permitted.

RESTRICTED EFFECT Effects are restricted at a fine-grained level.

TERMINATION A function call must terminate. Specifically, a call graph keeps track of changes to the size of arguments.

UNION CONTRACTS Given a set of contracts, guarantee that a value always satisfies at least one those contracts. Checking the union of flat contracts is easy, but checking the union of higher-order contracts relies on state to keep track of violations and assign blame.

Most papers illustrate these properties with a plethora of examples, all of which can be implemented in *effect/racket*.⁷

The cells of the table have the following rough meaning. A • indicates that the presented contract framework supports this property and comes with an illustrative example. Note that the number of • columns do not indicate anything about the “power” of the presented system. It merely means that a paper with fewer • entries may focus on a narrower set of properties. Also, these papers differ in other, significant ways that are not communicated by the • markings.

⁷See the artifact for all of the code.

1128 The effect/racket language can faithfully express all but one existing contract. This exception
1129 is Scholliers et al. [2015]’s contract that prohibits a function from being called during the dynamic
1130 extent of a call to another function. Effect-handler contracts can achieve this behavior as long as
1131 the excluded function comes with a contract that enables the system to monitor it; if not, this
1132 contract is impossible to realize without invasive monitoring techniques. For details, see the next
1133 subsection.

1134

1135

7.2 Related Work

1136

1137

1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

1149

1150

1151

1152

1153

1154

1155

1156

1157

1158

1159

1160

1161

1162

1163

1164

1165

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

The Java Modeling Language (JML) [Chalin et al. 2006] is a specification language for stating and verifying properties of objects in Java. It encompasses a broad range of features including assertions, class invariant statements, frame conditions, purity constraints, termination constraints, and ghost state declarations—just to name a few. Property checking takes place in one of two modes: static deductive verification (DV) or dynamic runtime-assertion checking (RAC). Some properties, such as termination, can be checked only using DV. JML differs from higher-order contract systems in three major ways. First, properties are described using a restrictive set of “well-defined” terms, a limitation compared to contracts written with ordinary constructs. Second, JML supports only first-order properties. Finally, JML lacks a blame assignment component, meaning developers are on their own when a contract check fails.

Tov and Pucella [2010]’s research on interoperability between a language with a substructural type system and one with a plain structural type system relies on an affinity check for function arguments. Specifically, the boundary employs a run-time check to ensure that a function argument is *affine*, meaning it can be applied at most once. This check uses a mutable Boolean field associated with each function value, i.e. ghost state, which indicates whether a function has been applied. Dimoulas et al. [2016] also use ghost state. They define a general-purpose DSL that uses ghost state to check protocol conformance. As described in Section 3.6, contract-effect handlers can easily introduce and manipulate ghost state.

For interoperability, a language with a sound gradual type-and-effect system relies on a run-time enforcement mechanism to restrict the effects performed by untyped code. The contracts for such a language are formulated in terms of two primitive operations [Bañados Schwerter et al. 2014]: *has* (for checking the privileges granted by the current context) and *restrict* (for restricting privileges of an expression). In the effect-handler language, these primitives are just main-effect contracts.

Shinnar [2011] takes some of the constructs from JML, in particular framing contracts, and adapts them to Haskell. The implementation uses delimited checkpointing to keep track of state. A delimited checkpoint is a snapshot of memory captured using software transactional memory (STM). Framing contracts can detect and restrict writes to transactional references by comparing memory snapshots. Shinnar proves erasure for a limited model of Haskell with delimited checkpoints. This work is similar to those pieces of research [Findler and Felleisen 2001; Strickland et al. 2012] that consider erasure for only a few restricted effects.

Disney et al. [2011]’s higher-order temporal (HOT) contracts and Moy and Felleisen [2023]’s trace contracts check properties of sequences of argument and returns values for functions and methods. While the two differ in many respects, from the perspective of effectful software contracts they fall into the same class of extended higher-order contracts. Describing constraints over sequences amounts to a writing a predicate that “folds over” the sequence incrementally, storing intermediate state in a mutable reference. As such, contract-effect handlers can supply the needed mutable references to such contracts. Indeed, Disney et al. [2011] present some examples that are *more directly* expressed using effect-handler contracts than HOT contracts. For example, their HOT contract for non-reentrancy does not suffice in the presence of control effects, whereas an effect-handler contract implementation of the same property is robust.

1177 Scholliers et al. [2015]’s computational contracts instantiate aspect-oriented programming for
1178 the contract world. Critically, such contracts can prohibit or enforce that a function f is called
1179 in a particular dynamic extent. Due to the intrusiveness of aspect-oriented programming, compu-
1180 tational contracts do not require that f is aware of the contract. Indeed, without aspect-oriented
1181 programming or a similarly invasive mechanism, there is no way to interpose on function applica-
1182 tions in a dynamic extent, which is why the effect-handler language cannot fully realize this form
1183 of checking.

1184 Moore et al. [2016]’s authorization contracts enforce access control with contracts about granted
1185 privileges. Specifically, authorization contracts can capture, check, and restore access privileges
1186 via an authority environment that records access privileges. Moore et al. [2016]’s model is essen-
1187 tially a variant of contract-handler contracts topped off with a DSL for authorization management.
1188 Effectful contracts alone do not implement any of the security aspects of the system. However, au-
1189 thorization contracts could be built on top of contract-handler contracts given the secure design
1190 described in Section 6.3.

1191 Nguyễn et al. [2019] provide a run-time check for termination by monitoring the SCP of func-
1192 tions dynamically. Any diverging function must exhibit an SCP violation, causing a contract vi-
1193 olation. They turn this run-time check into a static one, using existing contract verification tech-
1194 niques [Nguyễn et al. 2018]. To guarantee termination, they use continuation marks to store size-
1195 change information on the stack. Since parameters are a thin layer around continuation marks [Flatt
1196 and Dybvig 2020], porting this contract to parameter contracts is straightforward.

1197 While the literature on higher-order contracts tends to mention intersection and union con-
1198 tracts, implementing those in general is a serious challenge. Indeed, Racket rejects or/c contracts
1199 if the disjuncts are not “first-order distinguishable.” Several researchers [Freund et al. 2021; Keil
1200 and Thiemann 2015a; Williams et al. 2018] have studied this problem, and all come to the con-
1201 clusion that effects are needed. For example, Williams et al. [2018] use a mutable blame state to
1202 keep track of contract violations. A contract-effect handler can be used to implement this blame
1203 state. Moreover, erasure guarantees that such an implementation does not have adverse effects
1204 on a program’s result. This property is critically important because even benign-looking contract
1205 effects can have unintended consequences. Such has been observed in practice [Lazarek et al. 2020,
1206 section 6.1].

1207 8 IGNORED NO LONGER

1209 In the real world, developers use contracts with effects; in papers, researchers study how to employ
1210 effects in contracts. What has been lacking is a general framework for combining contracts and
1211 effects. As a result, existing extensions solve specific problems and do not generalize.

1212 This paper offers the first general model of effectful software contracts. As such, it synthesizes
1213 a model of effect handlers with a model of contracts. In this combination, contracts can check
1214 effects, contracts can request effects, and contracts can handle contract-requested effects. Yet, since
1215 contracts should not affect the main program—other than signal violations—the model is designed
1216 to avoid interference between contract-level effects and main-level code. Hence, in addition to
1217 well-definedness and blame correctness, the theory satisfies an erasure theorem.

1218 Beyond theoretical explorations, the formalism also provides guidance for implementation ef-
1219 forts. A fully faithful implementation, `effect/racket`, exists as a standalone language within the
1220 Racket ecosystem. This language demonstrates that the design can be realized. It is an open ques-
1221 tion how to modify an existing contract system to support all of the model’s expressive power in
1222 a backwards compatible manner. Still, the theory can serve as a roadmap for others who wish to
1223 combine effects and contracts in a principled way. And, as effect handlers go mainstream [Chan-
1224 drasekaran et al. 2018], the theory may find many more practical uses.

1225

ACKNOWLEDGMENTS

This work was supported by NSF grant SHF 2116372. The authors thank Robby Findler, participants of the NII Shonan Meeting 203 on effect handlers, and the anonymous POPL reviewers.

REFERENCES

- Danel Ahman and Andrej Bauer. 2020. Runners in Action. In *European Symposium on Programming (ESOP)*. https://doi.org/10.1007/978-3-030-44914-8_2
- Leif Andersen, Vincent St-Amour, Jan Vitek, and Matthias Felleisen. 2018. Feature-Specific Profiling. *Transactions on Programming Languages and Systems (TOPLAS)*. <https://doi.org/10.1145/3275519>
- Felipe Bañados Schwerter. 2016. Side Effects Take the Blame. In *Software Language Engineering (SLE)*. <https://doi.org/10.1145/2997364.2997381>
- Felipe Bañados Schwerter, Ronald Garcia, and Éric Tanter. 2014. A Theory of Gradual Effect Systems. In *International Conference on Functional Programming (ICFP)*. <https://doi.org/10.1145/2692915.2628149>
- Hendrik Pieter Barendregt. 1981. *The Lambda Calculus*. North-Holland Publishing Co.
- Lawrence E. Bassham, Andrew L. Rukhin, Juan Soto, James R. Nechvatal, Miles E. Smid, Elaine B. Barker, Stefan D. Leigh, Mark Levenson, Mark Vangel, David L. Banks, Nathanael Alan Heckert, James F. Dray, and San Vo. 2010. *A Statistical Test Suite for Random and Pseudorandom Number Generators for Cryptographic Applications*. Technical Report. National Institute of Standards and Technology. <https://csrc.nist.gov/publications/detail/sp/800-22/rev-1a/final>.
- Dariusz Biernacki, Maciej Piróg, Piotr Polesiuk, and Filip Sieczkowski. 2017. Handle with Care: Relational Interpretation of Algebraic Effects and Handlers. In *Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/3158096>
- Matthias Blume and David McAllester. 2006. Sound and Complete Models of Contracts. *Journal of Functional Programming (JFP)*. <https://doi.org/10.1017/S0956796806005971>
- Jonathan Immanuel Brachthäuser, Philipp Schuster, Edward Lee, and Aleksander Boruch-Gruszecki. 2022. Effects, Capabilities, and Boxes: From Scope-Based Reasoning to Type-Based Reasoning and Back. In *Object-Oriented Programming, Systems, Languages and Applications (OOPSLA)*. <https://doi.org/10.1145/3527320>
- Robert Cartwright and Matthias Felleisen. 1994. Extensible Denotational Language Specifications. In *Theoretical Aspects of Computer Software (TACS)*. https://doi.org/10.1007/3-540-57887-0_99
- Patrice Chalin, Joseph R. Kiniry, Gary T. Leavens, and Erik Poll. 2006. Beyond Assertions: Advanced Specification and Verification with JML and ESC/Java2. In *Formal Methods for Components and Objects*. https://doi.org/10.1007/11804192_16
- Sivaramakrishnan Krishnamoorthy Chandrasekaran, Daan Leijen, Matija Pretnar, and Tom Schrijvers. 2018. Algebraic Effect Handlers go Mainstream (Dagstuhl Seminar 18172). *Dagstuhl Reports*. <https://doi.org/10.4230/DagRep.8.4.104>
- Stephen Chang, Eli Barzilay, John Clements, and Matthias Felleisen. 2011. From Stack Traces to Lazy Rewriting Sequences. In *Proc. Implementation and Application of Functional Languages*. https://doi.org/10.1007/978-3-642-34407-7_7
- John Clements and Matthias Felleisen. 2004. A Tail-Recursive Machine with Stack Inspection. In *Transactions on Programming Languages and Systems (TOPLAS)*. <https://doi.org/10.1145/1034774.1034778>
- John Clements, Matthew Flatt, and Matthias Felleisen. 2001. Modeling an Algebraic Stepper. In *European Symposium on Programming (ESOP)*. https://doi.org/10.1007/3-540-45309-1_21
- H.B. Curry and R. Feys. 1958. *Combinatory Logic, Volume I*. North-Holland, Amsterdam.
- Jeffrey Dean and Sanjay Ghemawat. 2008. MapReduce: Simplified Data Processing on Large Clusters. *Communications of the ACM (CACM)*. <https://doi.org/10.1145/1327452.1327492>
- Jack B. Dennis and Earl C. Van Horn. 1966. Programming Semantics for Multiprogrammed Computations. *Communications of the ACM (CACM)*. <https://doi.org/10.1145/365230.365252>
- Christos Dimoulas and Matthias Felleisen. 2011. On Contract Satisfaction in a Higher-Order World. *Transactions on Programming Languages and Systems (TOPLAS)*. <https://doi.org/10.1145/2039346.2039348>
- Christos Dimoulas, Robert Bruce Findler, Cormac Flanagan, and Matthias Felleisen. 2011. Correct Blame for Contracts: No More Scapegoating. In *Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/1926385.1926410>
- Christos Dimoulas, Max S. New, Robert Bruce Findler, and Matthias Felleisen. 2016. Oh Lord, Please Don't Let Contracts Be Misunderstood (Functional Pearl). In *International Conference on Functional Programming (ICFP)*. <https://doi.org/10.1145/2951913.2951930>
- Christos Dimoulas, Sam Tobin-Hochstadt, and Matthias Felleisen. 2012. Complete Monitors for Behavioral Contracts. In *European Symposium on Programming (ESOP)*. https://doi.org/10.1007/978-3-642-28869-2_11
- Tim Disney, Cormac Flanagan, and Jay McCarthy. 2011. Temporal Higher-Order Contracts. In *International Conference on Functional Programming (ICFP)*. <https://doi.org/10.1145/2034773.2034800>
- Matthias Felleisen. 1988. The Theory and Practice of First-Class Prompts. In *Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/73560.73576>

- 1275 Matthias Felleisen. 1991. On the Expressive Power of Programming Languages. *Science of Computer Programming*. [https://doi.org/10.1016/0167-6423\(91\)90036-W](https://doi.org/10.1016/0167-6423(91)90036-W)
- 1276
- 1277 Matthias Felleisen, Robert Bruce Findler, and Matthew Flatt. 2009. *Semantics Engineering with PLT Redex*. MIT Press.
- 1278 Matthias Felleisen, Robert Bruce Findler, Matthew Flatt, Shriram Krishnamurthi, Eli Barzilay, Jay McCarthy, and Sam Tobin-Hochstadt. 2018. A Programmable Programming Language. *Communications of the ACM (CACM)*. <https://doi.org/10.1145/3127323>
- 1279
- 1280 Matthias Felleisen and Robert Hieb. 1992. The Revised Report on the Syntactic Theories of Sequential Control and State. In *Theoretical Computer Science*. [https://doi.org/10.1016/0304-3975\(92\)90014-7](https://doi.org/10.1016/0304-3975(92)90014-7)
- 1281
- 1282 Robert Bruce Findler and Matthias Blume. 2006. Contracts as Pairs of Projections. In *Functional and Logic Programming (FLP)*. https://doi.org/10.1007/11737414_16
- 1283
- 1284 Robert Bruce Findler and Matthias Felleisen. 2001. Contract Soundness for Object-Oriented Languages. In *Object-Oriented Programming, Systems, Languages and Applications (OOPSLA)*. <https://doi.org/10.1145/504311.504283>
- 1285
- 1286 Robert Bruce Findler and Matthias Felleisen. 2002. Contracts for Higher-Order Functions. In *International Conference on Functional Programming (ICFP)*. <https://doi.org/10.1145/581478.581484>
- 1287
- 1288 Matthew Flatt and R. Kent Dybvig. 2020. Compiler and Runtime Support for Continuation Marks. In *Programming Language Design and Implementation (PLDI)*. <https://doi.org/10.1145/3385412.3385981>
- 1289
- 1290 Matthew Flatt and PLT. 2010. *Reference: Racket*. Technical Report PLT-TR-2010-1. PLT Design Inc. <https://racket-lang.org/tr1/>.
- 1291
- 1292 Matthew Flatt, Gang Yu, Robert Bruce Findler, and Matthias Felleisen. 2007. Adding Delimited and Composable Control to a Production Programming Environment. In *International Conference on Functional Programming (ICFP)*. <https://doi.org/10.1145/1291151.1291178>
- 1293
- 1294 Teodoro Freund, Yann Hamdaoui, and Arnaud Spiwack. 2021. Union and Intersection Contracts Are Hard, Actually. In *Dynamic Languages Symposium (DLS)*. <https://doi.org/10.1145/3486602.3486767>
- 1295
- 1296 Martin Gashichler and Michael Sperber. 2005. Integrating User-Level Threads with Processes in Scsh. *Higher-Order and Symbolic Computation*. <https://doi.org/10.1007/s10990-005-4879-2>
- 1297
- 1298 Michael Greenberg, Benjamin C. Pierce, and Stephanie Weirich. 2010. Contracts Made Manifest. In *Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/1706299.1706341>
- 1299
- 1300 Daniel Hillerström and Sam Lindley. 2018. Shallow Effect Handlers. In *Asian Symposium on Programming Languages and Systems (APLAS)*. https://doi.org/10.1007/978-3-030-02768-1_22
- 1301
- 1302 Matthias Keil and Peter Thiemann. 2015a. Blame Assignment for Higher-Order Contracts with Intersection and Union. In *International Conference on Functional Programming (ICFP)*. <https://doi.org/10.1145/2784731.2784737>
- 1303
- 1304 Matthias Keil and Peter Thiemann. 2015b. TreatJS: Higher-Order Contracts for JavaScripts. In *European Conference on Object-Oriented Programming (ECOOP)*. <https://doi.org/10.4230/LIPICs.ECOOP.2015.28>
- 1305
- 1306 Lukas Lazarek, Alexis King, Samanvitha Sundar, Robert Bruce Findler, and Christos Dimoulas. 2020. Does Blame Shifting Work?. In *Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/3371133>
- 1307
- 1308 Chin Soon Lee, Neil D. Jones, and Amir M. Ben-Amram. 2001. The Size-Change Principle for Program Termination. In *Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/360204.360210>
- 1309
- 1310 Daan Leijen. 2013. *Koka: Programming with Row-Polymorphic Effect Types*. Technical Report MSR-TR-2013-79. Microsoft Research. <https://www.microsoft.com/en-us/research/publication/koka-programming-with-row-polymorphic-effect-types/>.
- 1311
- 1312 Jacob Matthews and Robert Bruce Findler. 2007. Operational Semantics for Multi-Language Programs. In *Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/1498926.1498930>
- 1313
- 1314 Bertrand Meyer. 1988. *Object-Oriented Software Construction*. Prentice Hall.
- 1315
- 1316 Bertrand Meyer. 1992. Applying “Design by Contract”. *Computer*. <https://doi.org/10.1109/2.161279>
- 1317
- 1318 Scott Moore, Christos Dimoulas, Robert Bruce Findler, Matthew Flatt, and Stephen Chong. 2016. Extensible Access Control with Authorization Contracts. In *Object-Oriented Programming, Systems, Languages and Applications (OOPSLA)*. <https://doi.org/10.1145/2983990.2984021>
- 1319
- 1320 Cameron Moy and Matthias Felleisen. 2023. Trace Contracts. <https://khoury.northeastern.edu/~camoy/pub/trace-contract.pdf>.
- 1321
- 1322 Phúc C. Nguyễn, Thomas Gilray, Sam Tobin-Hochstadt, and David Van Horn. 2018. Soft Contract Verification for Higher-Order Stateful Programs. In *Principles of Programming Languages (POPL)*. <https://doi.org/10.1145/3158139>
- 1323
- 1324 Phúc C. Nguyễn, Thomas Gilray, Sam Tobin-Hochstadt, and David Van Horn. 2019. Size-Change Termination as a Contract. In *Programming Language Design and Implementation (PLDI)*. <https://doi.org/10.1145/3325984>
- 1325
- 1326 Bruno C. D. S. Oliveira, Tom Schrijvers, and William R. Cook. 2012. MRI: Modular Reasoning About Interference in Incremental Programming. *Journal of Functional Programming (JFP)*. <https://doi.org/10.1017/S0956796812000354>
- 1327
- 1328 Gordon Plotkin. 1975. Call-by-name, call-by-value and the λ -calculus. *Theoretical Computer Science*. [https://doi.org/10.1016/0304-3975\(75\)90017-1](https://doi.org/10.1016/0304-3975(75)90017-1)

- 1324 Gordon Plotkin and Matija Pretnar. 2009. Handlers of Algebraic Effects. In *European Symposium on Programming (ESOP)*.
1325 https://doi.org/10.1007/978-3-642-00590-9_7
- 1326 Matija Pretnar. 2015. An Introduction to Algebraic Effects and Handlers. In *Mathematical Foundations of Programming*
1327 *Semantics (MFPS)*. <https://doi.org/10.1016/j.entcs.2015.12.003>
- 1328 R. L. Rivest, A. Shamir, and A. Adleman. 1978. A Method for Obtaining Digital Signatures and Public-Key Cryptosystems.
1329 In *Communications of the ACM (CACM)*. <https://doi.org/10.1145/359340.359342>
- 1330 Christophe Scholliers, Éric Tanter, and Wolfgang De Meuter. 2015. Computational Contracts. *Science of Computer Program-*
1331 *ming*. <https://doi.org/10.1016/j.scico.2013.09.005>
- 1332 Avraham Ever Shinnar. 2011. *Safe and Effective Contracts*. Ph. D. Dissertation. Harvard University.
- 1333 Dorai Sitaram and Matthias Felleisen. 1990. Control Delimiters and Their Hierarchies. *Lisp and Symbolic Computation*.
1334 <https://doi.org/10.1007/BF01806126>
- 1335 T. Stephen Strickland, Sam Tobin-Hochstadt, Robert Bruce Findler, and Matthew Flatt. 2012. Chaperones and Impersonators:
1336 Run-Time Support for Reasonable Interposition. In *Object-Oriented Programming, Systems, Languages and Applications*
1337 *(OOPSLA)*. <https://doi.org/10.1145/2384616.2384685>
- 1338 Sam Tobin-Hochstadt, Vincent St-Amour, Ryan Culpepper, Matthew Flatt, and Matthias Felleisen. 2011. Languages as
1339 Libraries. In *Programming Language Design and Implementation (PLDI)*. <https://doi.org/10.1145/1993316.1993514>
- 1340 Jesse A. Tov and Riccardo Pucella. 2010. Stateful Contracts for Affine Types. In *European Symposium on Programming*
1341 *(ESOP)*. https://doi.org/10.1007/978-3-642-11957-6_29
- 1342 Jack Williams, J. Garrett Morris, and Philip Wadler. 2018. The Root Cause of Blame: Contracts for Intersection and Union
1343 Types. In *Object-Oriented Programming, Systems, Languages and Applications (OOPSLA)*. <https://doi.org/10.1145/3276504>
- 1344 Andrew K. Wright and Matthias Felleisen. 1994. A Syntactic Approach to Type Soundness. *Information and Computation*.
1345 <https://doi.org/10.1006/inco.1994.1093>
- 1346 Ningning Xie, Jonathan Immanuel Brachthäuser, Daniel Hillerström, Philipp Schuster, and Daan Leijen. 2020. Effect Hand-
1347 lers, Evidently. In *International Conference on Functional Programming (ICFP)*. <https://doi.org/10.1145/3408981>
- 1348 Dana N. Xu. 2012. Hybrid Contract Checking via Symbolic Simplification. In *Partial Evaluation and Program Manipulation*
1349 *(PEPM)*. <https://doi.org/10.1145/2103746.2103767>
- 1350 Dana N. Xu. 2014. Dynamic Contract Checking for OCaml. <http://gallium.inria.fr/~naxu/research/camlcontract.pdf>.
- 1351 Dana N. Xu, Simon Peyton Jones, and Koen Claessen. 2009. Static Contract Checking for Haskell. In *Principles of Program-*
1352 *ming Languages (POPL)*. <https://doi.org/10.1145/1480881.1480889>
- 1353 Yizhou Zhang and Andrew C. Myers. 2019. Abstraction-Safe Effect Handlers via Tunneling. In *Principles of Programming*
1354 *Languages (POPL)*. <https://doi.org/10.1145/3290318>

1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372

1373 A PROOF SYNTAX

1374 $\boxed{\text{DEPENDENT (PROOF)}}$ extends DEPENDENT (EVAL)

1375 $t \in \text{Ter} = v \mid \text{err}_j^k$

1376 $r \in \text{Redex} = \text{if } v e e \mid v v \mid \text{mon}_j^{k,l} v v \mid \text{handle}^m v \text{ with } v \mid \text{handle}^m E^m[\text{do } v] \text{ with } v$

1379 B FUNCTIONAL EVALUATION PROOF

1380 **Theorem 5.1 (Functional Evaluation).** Two facts about the evaluator hold:

- 1382 (1) The eval relation is a partial function.
 1383 (2) If e is a program, then either (i) $\text{eval}(e)$ is defined or (ii) the reduction sequence starting
 1384 with e is unbounded.

1385 **PROOF.** These facts are established by a few lemmas.

- 1386 (1) By Lemma B.1.
 1387 (2) By interleaved application of Lemma B.3 and Lemma B.4.

1389 \square

1390 **Lemma B.1 (Deterministic Reduction).** If $e \mapsto e'$ and $e \mapsto e''$ then $e = e''$.

1391 **PROOF.** By Lemma B.2, every reducible expression can be decomposed into a unique evaluation
 1392 context and a unique redex. Inspecting the reduction relation, each pair of rules is disjoint. In
 1393 particular, DO and ERR-DO are kept disjoint by the definition of unhandled. Additionally, there
 1394 is only one way for DO to apply to a given expression; it applies to only the innermost handler
 1395 because of the side condition that $E^\triangleright \in \text{unhandled}$. \square

1396 **Lemma B.2 (Unique Decomposition).** For $e \in \text{Expr}$, either $e \in \text{Ter}$ or there exists unique evalu-
 1397 ation context E and unique redex r such that $e = E[r]$.

1400 **PROOF.** By structural induction on e .

1401 **Case $e = b$.**

1402 Booleans are terminal, so $e \in \text{Ter}$.

1403 **Case $e = \text{if } e_g e_t e_f$.**

1404 The inductive hypothesis applied to e_g yields three cases. If $e_g \in \text{Val}$ then $E = \square$ and
 1405 $r = e$. If $e_g = \text{err}_j^k$ then $E = \text{if } \square e_t e_f$ and $r = e_g$. Otherwise, $e_g = E_g[r]$, in which case
 1406 $E = \text{if } E_g e_t e_f$.

1407 **Otherwise.**

1408 The remaining cases are similar to one of the above. \square

1409 **Lemma B.3 (Progress).** If e is closed then either $e \in \text{Ter}$ or $e \mapsto e'$.

1410 **PROOF.** By Lemma B.2, either $e \in \text{Ter}$ or e can be decomposed into a unique evaluation context
 1411 and a unique redex. Suppose $e = E[r]$. Conclusion follows by cases on r . \square

1412 **Lemma B.4 (Preservation).** If e is closed and $e \mapsto e'$ then e' is closed.

1413 **PROOF.** By cases on $e \mapsto e'$. \square

1414 C ERASURE PROOF

1415 **Theorem 5.2 (Erasure).** If $\text{eval}(e) = b$ then $\text{eval}(\mathcal{E}(e)) = b$.

$\sim \subseteq \text{Ctx}^\triangleright \times \text{Ctx}^\triangleright$	
handle $^\diamond$ E^\triangleright with $e_h \sim$ handle $^\diamond$ $\widetilde{E}^\triangleright$ with e	if $E^\triangleright \sim \widetilde{E}^\triangleright$
handle $^\diamond$ E^\triangleright with $e_h \sim \widetilde{E}^\triangleright$	if $E^\triangleright \sim \widetilde{E}^\triangleright$
mon $^{k,l}_j$ $E^\triangleright e_c \sim \widetilde{E}^\triangleright$	if $E^\triangleright \sim \widetilde{E}^\triangleright$
mark $^{k,l}_j$ $v E^\triangleright \sim \widetilde{E}^\triangleright$	if $E^\triangleright \sim \widetilde{E}^\triangleright$
...	

Fig. 7. Simulation on Contexts

PROOF. By Lemma C.1, let $\widetilde{e} = \mathcal{Z}(e)$ where $e \sim \widetilde{e}$. It suffices to show that $\widetilde{e} \mapsto^* b$. By induction on $e \mapsto^* b$.

Case $e = b$.

Booleans are preserved by the simulation so $\widetilde{e} = b$, hence $\widetilde{e} \mapsto^* b$.

Case $e \mapsto^+ b$.

By Lemma C.2, $e \mapsto^+ e' \mapsto^* b$ and $\widetilde{e} \mapsto^* e_i \simeq \widetilde{e}'$ for $e' \sim \widetilde{e}'$. The inductive hypothesis yields that $\widetilde{e}' \mapsto^* b$. Thus, $\widetilde{e} \mapsto^* e_i \simeq \widetilde{e}' \mapsto^* b$ implies $\widetilde{e} \mapsto^* b$. \square

Lemma C.1 (Erasure Inclusion). For all $e \in \text{Expr}$, $e \sim \mathcal{Z}(e)$.

PROOF. By induction on e . \square

Lemma C.2 (Simulation). If $e \mapsto^+ v$ then for all \widetilde{e} there exists e', \widetilde{e}' such that $e \mapsto^+ e' \mapsto^* v$ and $\widetilde{e} \mapsto^* \widetilde{e}'$.

PROOF. Since e evaluates to a value, not an error, that means $e = E[e_r] \mapsto E[e_c]$ for some evaluation context E and expressions e_r, e_c .

Suppose $E \notin \text{Ctx}^\triangleright$, or equivalently $E = E^\diamond$. Assume too that $E^\diamond[e_r] \sim \widetilde{e}_j$. By Lemma C.3, $E^\diamond[e_c] \sim \widetilde{e}_j$ as needed. Otherwise, take $E = E^\triangleright$. By cases on $E^\triangleright[e_r] \mapsto E^\triangleright[e_c]$.

Case $E^\triangleright[\text{if } v e_t e_f] \mapsto E^\triangleright[e_t], v \neq \text{false}$.

By Lemma C.7, $E^\triangleright[e_r] \sim \widetilde{E}^\triangleright[\text{if } \widetilde{v} \widetilde{e}_t \widetilde{e}_f]$ for $e_t \sim \widetilde{e}_t$. Since \sim preserves non-false values, $\widetilde{E}^\triangleright[\text{if } \widetilde{v} \widetilde{e}_t \widetilde{e}_f] \mapsto \widetilde{E}^\triangleright[\widetilde{e}_t]$. By Lemma C.8, $E^\triangleright[e_t] \sim \widetilde{E}^\triangleright[\widetilde{e}_t]$. All of the remaining cases use Lemma C.7 and Lemma C.8 in a similar way.

Case $E^\triangleright[(\lambda x. e) v] \mapsto E^\triangleright[e[v/x]]$.

By Lemma C.9.

Case $E^\triangleright[\text{handle}^\triangleright v \text{ with } v_h] \mapsto E^\triangleright[v]$.

Let $\widetilde{e} = \widetilde{E}^\triangleright[\text{handle}^\triangleright \widetilde{v} \text{ with } \widetilde{v}_h]$. Then $\widetilde{e} \mapsto \widetilde{E}^\triangleright[\widetilde{v}]$ as needed.

Case $E^\triangleright[\text{handle}^\diamond v \text{ with } v_h] \mapsto E^\triangleright[v]$.

If $\widetilde{e} = E^\triangleright[\text{handle}^\diamond \widetilde{v} \text{ with } e_h]$ then $\widetilde{e} \mapsto E^\triangleright[\widetilde{v}]$. If $\widetilde{e} = E^\triangleright[\widetilde{v}]$ already then no step is needed.

Case $E^\triangleright[\text{handle}^\triangleright E_k^\triangleright[\text{do } v] \text{ with } v_h] \mapsto E^\triangleright[v_h v_d (\lambda x. \text{handle}^\triangleright E_k^\triangleright[e_c] \text{ with } v_h)]$.

By Lemma C.4 and Lemma C.5.

Case $E^\triangleright[\text{handle}^\diamond E_k^\diamond[\text{do } v] \text{ with } \langle v_d, v_h \rangle] \mapsto E^\triangleright[\text{handle}^\diamond E_k^\diamond[v_c] \text{ with } v_h]$.

By Lemma C.3.

Case $E^\triangleright[\text{handle}^\diamond E_k^\diamond[\text{do } v] \text{ with } f] \mapsto E^\triangleright[\text{handle}^\diamond E_k^\diamond[\text{do } v] \text{ with } (f v_d)]$.

Follows directly from the simulation.

Case $E^\triangleright[\text{mon}_j^{k,l} \text{ true } v] \mapsto E^\triangleright[v]$.

Let $\widetilde{e} = \widetilde{E}^\triangleright[\widetilde{v}]$ for $v \sim \widetilde{v}$. Since $E^\triangleright[v] \sim \widetilde{E}^\triangleright[\widetilde{v}]$ that implies $E^\triangleright[v] \sim \widetilde{e}$.

1471 **Case** $E^\triangleright[\text{mon}_j^{k,l} \text{false } v] \mapsto E^\triangleright[\text{err}_j^k]$.

1472 Contradiction since err_j^k does not evaluate to a value.

1473 **Case** $E^\triangleright[\text{mon}_j^{k,l} f v] \mapsto E^\triangleright[\text{mon}_j^{k,l} (f v) v]$.

1474 Let $\tilde{e} = \widetilde{E^\triangleright}[\tilde{v}]$ for $v \sim \tilde{v}$. Since $E^\triangleright[\text{mon}_j^{k,l} (f v) v] \sim \widetilde{E^\triangleright}[\tilde{v}]$ that implies $E^\triangleright[\text{mon}_j^{k,l} (f v) v] \sim \tilde{e}$.

1476 **Case** $E^\triangleright[\text{mon}_j^{k,l} (v_d \implies v_c) f] \mapsto E^\triangleright[\lambda x. \text{---}]$.

1477 This step produces an expression that is still in simulation with \tilde{e} :

$$\begin{aligned}
 1478 \quad E^\triangleright[\lambda x. \text{---}] &= E^\triangleright[\lambda x. \text{let } x_j = \text{mon}_j^{l,j} v_d \text{ x in} \\
 1479 \quad &\quad \text{let } x_k = \text{mon}_j^{l,k} v_d \text{ x in} \\
 1480 \quad &\quad \text{mon}_j^{k,l} (v_c x_j) (f x_k)] \\
 1481 \quad &\sim \widetilde{E^\triangleright}[\lambda x. \text{let } x_j = x \text{ in} \\
 1482 \quad &\quad \text{let } x_k = x \text{ in} \\
 1483 \quad &\quad \tilde{f} x_k] \\
 1484 \quad &\simeq \widetilde{E^\triangleright}[\lambda x. \tilde{f} x] \\
 1485 \quad &\simeq \widetilde{E^\triangleright}[\tilde{f}] \\
 1486 \quad &= \tilde{e}
 \end{aligned}$$

1491 **Case** $E^\triangleright[\text{mon}_j^{k,l} (v_d \blacktriangleright v_c) f] \mapsto E^\triangleright[\lambda x. \text{mark}_j^{k,l} (v_d \blacktriangleright v_c) (f x)]$.

1492 Thus, $\tilde{e} = \widetilde{E^\triangleright}[\tilde{f}]$, and $e' = \widetilde{E^\triangleright}[\tilde{f}] \sim \widetilde{E^\triangleright}[\lambda x. \tilde{f} x] \sim \widetilde{E^\triangleright}[\tilde{f}]$.

1493 **Case** $E^\triangleright[\text{mark}_j^{k,l} (v_d \blacktriangleright v_c) v] \mapsto E^\triangleright[v]$.

1494 Thus, $\tilde{e} = \widetilde{E^\triangleright}[\tilde{v}]$, and $e' = E^\triangleright[v] \sim \widetilde{E^\triangleright}[\tilde{v}]$.

1495 **Case** $E^\triangleright[\text{mon}_j^{k,l} (\diamond v_h) f] \mapsto E^\triangleright[\lambda x. \text{handle}^\diamond (f x) \text{ with } v_h]$.

1496 Thus, $\tilde{e} = \widetilde{E^\triangleright}[\tilde{f}]$, and $e' = E^\triangleright[\lambda x. \text{handle}^\diamond (f x) \text{ with } v_h] \sim E^\triangleright[\lambda x. \tilde{f} x] \sim E^\triangleright[\tilde{f}]$.

1497 **Otherwise.**

1498 The remaining cases are similar to one of the above. □

1500 **Lemma C.3 (Diamond Irrelevance).** If $E^\diamond[e_s] \sim \tilde{e}$ then $E^\diamond[e_t] \sim \tilde{e}$.

1501 **PROOF.** There are only two situations that can occur during evaluation:

1502 **Case** $E^\diamond = E^\triangleright[\text{mon}_j^{k,l} E e_c]$.

1503 Therefore, $E^\diamond[e_s] = E^\triangleright[\text{mon}_j^{k,l} E[e_s] e_c] \sim \widetilde{E^\triangleright}[\tilde{e}_c] = \tilde{e}$. For the same reason, $E^\diamond[e_t] \sim \tilde{e}$.

1504 **Case** $E^\diamond = E^\triangleright[\text{handle}^\diamond e_b \text{ with } E]$.

1505 Similar to the above. □

1506 **Lemma C.4 (Push Empty).** If $E \sim \tilde{E}$ then $v \downarrow \tilde{E} = \square$.

1507 **PROOF.** By induction on $E \sim \tilde{E}$. □

1508 **Lemma C.5 (Pull Empty).** If $E \sim \tilde{E}$ then $\uparrow \tilde{E} = \square$.

1509 **PROOF.** By induction on $E \sim \tilde{E}$. □

1510 **Lemma C.6 (Unhandled Preservation).** If $E^\triangleright \in \text{unhandled}$ then $\widetilde{E^\triangleright} \in \text{unhandled}$.

1511 **PROOF.** By induction on $E^\triangleright \sim \widetilde{E^\triangleright}$. □

1512

1520 **Lemma C.7 (Simulation Decomposition).** If $e \sim \tilde{e}$ and $e = E^\triangleright[e_s]$ for $e_s \notin \text{Val}$, then exists $\widetilde{E^\triangleright}$
 1521 and \tilde{e}_s such that $\tilde{e} = \widetilde{E^\triangleright}[\tilde{e}_s]$ where $E^\triangleright \sim \widetilde{E^\triangleright}$ and $e_s \sim \tilde{e}_s$.

1522 PROOF. By induction on $e \sim \tilde{e}$. □

1524 **Lemma C.8 (Simulation Composition).** If $E^\triangleright \sim \widetilde{E^\triangleright}$ and $e \sim \tilde{e}$, then $E^\triangleright[e] \sim \widetilde{E^\triangleright}[\tilde{e}]$.

1526 PROOF. By induction on $E^\triangleright \sim \widetilde{E^\triangleright}$. □

1528 **Lemma C.9 (Substitution).** If $e \sim \tilde{e}$ and $v \sim \tilde{v}$ then $e[v/x] \sim \tilde{e}[\tilde{v}/x]$.

1529 PROOF. By induction on $e \sim \tilde{e}$. □

1531 D BLAME CORRECTNESS

1532 D.1 Syntax with Ownership

1534 DEPENDENT (ANNOTATED) extends DEPENDENT (EVAL)

1536 $e \in \text{Expr} = \dots \mid |e|^l$

1537 $f \in \text{Fun} = \dots \mid |f|^l$

1538 $v \in \text{Val} = \dots \mid |v|^l$

1539 $E \in \text{Ctx} = \dots \mid |E|^l$

1540 $E^\triangleright \in \text{Ctx}^\triangleright = \dots \mid |E^\triangleright|^l$

1541 $E^\diamond \in \text{Ctx}^\diamond = \dots \mid |E^\diamond|^l$

1543

1544 D.2 Ownership Metafunctions

1545

1546 $\uparrow_o : \text{Ctx} \rightarrow \text{Ctx}$

1547

1548 $\uparrow_o \square = \square$

1549 $\uparrow_o \langle E, e \rangle = \uparrow_o E$

1550 $\uparrow_o \langle v, E \rangle = \uparrow_o E$

1551 \dots

1552 $\uparrow_o |E|^l = |(\uparrow_o E)|^l$

1553 $\uparrow_o (\text{mark}_j^{k,l} (v_d \triangleright v_c) E) = \text{mon}_j^{k,l} v_d (\uparrow_o E)$

1554 $\uparrow_o (\text{mark}_j^{k,l} (v_d \blacktriangleright v_c) E) = \text{mon}_j^{k,l} v_d (\uparrow_o E)$

1555

1556

1557

1558

1559

1560 $\uparrow_o^- : \text{Ctx} \rightarrow \text{Ctx}$

1561

1562 $\uparrow_o^- \square = \square$

1563 $\uparrow_o^- \langle E, e \rangle = \uparrow_o^- E$

1564 $\uparrow_o^- \langle v, E \rangle = \uparrow_o^- E$

1565 \dots

1566 $\uparrow_o^- |E|^l = |(\uparrow_o^- E)|^l$

1567

1568

1546 $\downarrow_o : \text{Val} \times \text{Ctx} \rightarrow \text{Ctx}$

1547

1548 $v \downarrow_o \square = \square$

1549 $v \downarrow_o \langle E, e \rangle = v \downarrow_o E$

1550 $v \downarrow_o \langle v', E \rangle = v \downarrow_o E$

1551 \dots

1552 $v \downarrow_o |E|^l = (v \downarrow_o E)[|\square|^l]$

1553 $v \downarrow_o (\text{mark}_j^{k,l} (v_d \triangleright v_c) E) = (v \downarrow_o E)[\text{mon}_j^{l,k} v_c \square]$

1554 $v \downarrow_o (\text{mark}_j^{k,l} (v_d \blacktriangleright v_c) E) = (v \downarrow_o E)[\text{mon}_j^{l,k} v_c e \square]$

1555 where $e = \text{mon}_j^{k,l} v_d ((\uparrow_o E)[v])$

1556

1557

1558

1559 $\downarrow_o^- : \text{Ctx} \rightarrow \text{Ctx}$

1560

1561 $\downarrow_o^- \square = \square$

1562 $\downarrow_o^- \langle E, e \rangle = \downarrow_o^- E$

1563 $\downarrow_o^- \langle v, E \rangle = \downarrow_o^- E$

1564 \dots

1565 $\downarrow_o^- |E|^l = (\downarrow_o^- E)[|\square|^l]$

$$\boxed{\text{Outer}_o : \text{Ctx} \rightarrow \text{Lab}}$$

$$\text{Outer}_o \square = l_o$$

$$\text{Outer}_o \langle E, e \rangle = \text{Outer}_o E$$

$$\text{Outer}_o \langle v, E \rangle = \text{Outer}_o E$$

$$\dots$$

$$\text{Outer}_o |E|^l = l_o$$

$$\text{Outer}_o (\text{mon}_j^{k,l} E e) = l$$

$$\text{Outer}_o (\text{mon}_j^{k,l} v E) = l$$

$$\text{Outer}_o (\text{mark}_j^{k,l} (v_d \triangleright v_c) E) = l$$

$$\boxed{\text{Inner}_o : \text{Ctx} \times \text{Lab} \rightarrow \text{Lab}}$$

$$\text{Inner}_o (\square, l) = l$$

$$\text{Inner}_o (\langle E, e \rangle, l) = \text{Inner}_o (E, l)$$

$$\text{Inner}_o (\langle v, E \rangle, l) = \text{Inner}_o (E, l)$$

$$\dots$$

$$\text{Inner}_o (|E|^l, l') = l$$

$$\text{Inner}_o (\text{mon}_j^{k,l} E e, l') = j$$

$$\text{Inner}_o (\text{mon}_j^{k,l} v E, l') = k$$

$$\text{Inner}_o (\text{mark}_j^{k,l} (v_d \triangleright v_c) E, l') = k$$

$$\boxed{\mathcal{O} : \text{Expr} \rightarrow \text{Expr}}$$

$$\mathcal{O} x = x$$

$$\mathcal{O} b = b$$

$$\mathcal{O} \langle e_1, e_2 \rangle = \langle \mathcal{O} e_1, \mathcal{O} e_2 \rangle$$

$$\dots$$

$$\mathcal{O} (|e|^l) = \mathcal{O} e$$

D.3 Syntactic Sugar

$$\bar{l} = l_1 \dots l_n$$

$$\overleftarrow{l_1 \dots l_n} = l_n \dots l_1$$

$$|e|^{l_1 \dots l_n} = |\dots |e|^{l_1} \dots|^{l_n}$$

$$||e||^{l_1 \dots l_n} = |e|^{l_1 \dots l_n} \text{ where } e \neq |e'|^l$$

$$l_1 > E > l_2 \Leftrightarrow \text{Outer}_o E = l_1 \wedge \text{Inner}_o (E, l_o) = l_2$$

D.4 Annotated Reduction Rules

$$\text{IF-TRUE} \quad E[\text{if } v \ e_1 \ e_2] \mapsto_o E[e_1] \text{ if } v \neq \text{false}$$

$$\text{IF-FALSE} \quad E[\text{if false } e_1 \ e_2] \mapsto_o E[e_2]$$

$$\text{APP-LAMBDA} \quad E[(\lambda x. e) |^{l_1 \dots l_n} v] \mapsto_o E[|e|^{l_1 \dots l_n} / x]^{l_1 \dots l_n} \quad \text{where } l_o > E > l$$

$$\text{APP-OP} \quad E[|o|^{l_1} |v|^{l_2}] \mapsto_o E[|\delta(o, v)|^{l_1}]$$

$$\text{HANDLE} \quad E[\text{handle}^m v \text{ with } v_h] \mapsto_o E[v]$$

$$\text{DO}^\triangleright \quad E[\text{handle}^\triangleright E^\triangleright [\text{do } v] \text{ with } v_h] \mapsto_o E[v_h v_d (\lambda x. \text{handle}^\triangleright E^\triangleright [e_c] \text{ with } v_h)]$$

if $E^\triangleright \in \text{unhandled}$

where $v_d = (\uparrow_o E^\triangleright)[v]$, $e_c = (v \downarrow_o E^\triangleright)[x]$

$$\text{DO-PAIR}^\diamond \quad E[\text{handle}^\diamond E^\diamond [\text{do } v] \text{ with } \langle v_d, v_h \rangle] \mapsto E[\text{handle}^\diamond E^\diamond [v_c] \text{ with } v_h]$$

if $E^\diamond \in \text{unhandled}$

where $v_c = (\downarrow_o E^\diamond)[v_d]$

1618	DO-FUN $^\diamond$	$E[\text{handle}^\diamond E^\diamond[\text{do } v] \text{ with } f] \mapsto E[\text{handle}^\diamond E^\diamond[\text{do } v] \text{ with } (f v_d)]$
1619		if $E^\diamond \in \text{unhandled}$
1620		where $v_d = (\uparrow_o E^\triangleright)[v]$
1621		
1622	MON-TRUE	$E[\text{mon}_j^{k,l} \ \text{true}\ \bar{l} v] \mapsto_o E[v]$
1623		
1624	MON-FALSE	$E[\text{mon}_j^{k,l} \ \text{false}\ \bar{l} v] \mapsto_o E[\text{err}_j^k]$
1625		
1626	MON-FLAT	$E[\text{mon}_j^{k,l} f v] \mapsto_o E[\text{mon}_j^{k,l} (f v) v]$
1627		
1628	MON-PAIR	$E[\text{mon}_j^{k,l} \langle v_a, v_b \rangle v] \mapsto_o E[\text{err}_j^k]$ if $v \neq \langle v_c, v_d \rangle$
1629	GRD-PAIR	$E[\text{mon}_j^{k,l} \langle v_a, v_b \rangle \ \langle v_c, v_d \rangle\ \bar{l} v] \mapsto_o E[\langle \ \text{mon}_j^{k,l} v_a v_c\ \bar{l}, \ \text{mon}_j^{k,l} v_b v_d\ \bar{l} \rangle]$
1630		
1631	MON-FUN	$E[\text{mon}_j^{k,l} (v_d \longrightarrow v_c) v] \mapsto_o E[\text{err}_j^k]$ if $v \notin \text{Fun}$
1632	GRD-FUN	$E[\text{mon}_j^{k,l} (v_d \longrightarrow v_c) f] \mapsto_o E[\lambda x. \text{mon}_j^{k,l} v_c (f (\text{mon}_j^{l,k} v_d x))]$
1633		
1634	MON-DEP-FUN	$E[\text{mon}_j^{k,l} (v_d \implies v_c) v] \mapsto_o E[\text{err}_j^k]$ if $v \notin \text{Fun}$
1635	GRD-DEP-FUN	$E[\text{mon}_j^{k,l} (v_d \implies v_c) f] \mapsto_o E[\lambda x. \text{mon}_j^{k,l} v_c (\text{mon}_j^{l,j} v_d x) (f (\text{mon}_j^{l,k} v_d x))]$
1636		
1637	MON-HANDLE $^\triangleright$	$E[\text{mon}_j^{k,l} (v_d \triangleright v_c) v] \mapsto_o E[\text{err}_j^k]$ if $v \notin \text{Fun}$
1638	GRD-HANDLE $^\triangleright$	$E[\text{mon}_j^{k,l} (v_d \triangleright v_c) f] \mapsto_o E[\lambda x. \text{mark}_j^{k,l} (v_d \triangleright v_c) (f x)]$
1639		
1640	MON-HANDLE $^\blacktriangleright$	$E[\text{mon}_j^{k,l} (v_d \blacktriangleright v_c) v] \mapsto_o E[\text{err}_j^k]$ if $v \notin \text{Fun}$
1641	GRD-HANDLE $^\blacktriangleright$	$E[\text{mon}_j^{k,l} (v_d \blacktriangleright v_c) f] \mapsto_o E[\lambda x. \text{mark}_j^{k,l} (v_d \blacktriangleright v_c) (f x)]$
1642		
1643	MARK	$E[\text{mark}_j^{k,l} v_\kappa v] \mapsto_o E[v]$
1644	MON-HANDLE $^\diamond$	$E[\text{mon}_j^{k,l} (\diamond v_h) v] \mapsto_o E[\text{err}_j^k]$ if $v \notin \text{Fun}$
1645		
1646	GRD-HANDLE $^\diamond$	$E[\text{mon}_j^{k,l} (\diamond v_h) f] \mapsto_o E[\lambda x. \text{handle}^\diamond (f x) \text{ with } v_h ^j]$
1647		
1648		
1649		

D.5 Ownership Well-formedness

1650							
1651							
1652							
1653	$\frac{}{l; \Gamma \vdash b}$	$\frac{}{l; \Gamma \uplus x : l \vdash x}$	$\frac{l; \Gamma \vdash e_1 \quad l; \Gamma \vdash e_2}{l; \Gamma \vdash \langle e_1, e_2 \rangle}$	$\frac{l; \Gamma \uplus x : l \vdash e}{l; \Gamma \vdash \lambda x. e}$	$\frac{}{l; \Gamma \vdash o}$		
1654							
1655	$\frac{l; \Gamma \vdash e_1 \quad l; \Gamma \vdash e_2 \quad l; \Gamma \vdash e_3}{l; \Gamma \vdash \text{if } e_1 e_2 e_3}$	$\frac{l; \Gamma \vdash e_1 \quad l; \Gamma \vdash e_2}{l; \Gamma \vdash e_1 e_2}$	$\frac{l; \Gamma \vdash e_1 \quad l; \Gamma \vdash e_2}{l; \Gamma \vdash \text{handle}^m e_1 \text{ with } e_2}$	$\frac{l; \Gamma \vdash e}{l; \Gamma \vdash \text{do } e}$			
1656							
1657							
1658	$\frac{l; \Gamma \vdash e_1 \quad l; \Gamma \vdash e_2}{l; \Gamma \vdash e_1 \longrightarrow e_2}$	$\frac{l; \Gamma \vdash e_1 \quad l; \Gamma \vdash e_2}{l; \Gamma \vdash e_1 \implies e_2}$	$\frac{l; \Gamma \vdash e_1 \quad l; \Gamma \vdash e_2}{l; \Gamma \vdash e_1 \triangleright e_2}$	$\frac{l; \Gamma \vdash e_1 \quad l; \Gamma \vdash e_2}{l; \Gamma \vdash e_1 \blacktriangleright e_2}$	$\frac{l; \Gamma \vdash e}{l; \Gamma \vdash \diamond e}$		
1659							
1660							
1661	$\frac{}{l; \Gamma \vdash \text{err}_j^k}$	$\frac{j; \Gamma \vdash e_1 \quad k; \Gamma \vdash e_2}{l; \Gamma \vdash \text{mon}_j^{k,l} e_1 e_2 ^k}$	$\frac{j; \Gamma \vdash e_1 \quad k; \Gamma \vdash x}{l; \Gamma \vdash \text{mon}_j^{k,l} e_1 x}$	$\frac{j; \Gamma \vdash e_1 \quad k; \Gamma \vdash e_2 \quad k; \Gamma \vdash e_3}{l; \Gamma \vdash \text{mon}_j^{k,l} e_1 e_2 ^k e_3}$			
1662							
1663							
1664	$\frac{j; \Gamma \vdash e_1 \quad k; \Gamma \vdash e_2}{l; \Gamma \vdash \text{mark}_j^{k,l} e_1 e_2 ^k}$	$\frac{j; \Gamma \vdash e_1 \quad k; \Gamma \vdash e_2 \quad k; \Gamma \vdash e_3}{l; \Gamma \vdash \text{mark}_j^{k,l} e_1 e_2 ^k e_3}$	$\frac{k; \Gamma \vdash e}{l; \Gamma \vdash e ^k}$				
1665							
1666							

D.6 Theorems, Key Lemmas and their Proofs

Theorem 5.4 (Blame Correctness). For all e if $l_0; \emptyset \vdash e$, if $e \mapsto_o^* E[\text{mon}_j^{k,l} v_\kappa v]$, then $v = |v'|^k$.

PROOF. Direct consequence of Lemma D.1 and Lemma D.3. \square

Lemma D.1 (Well-formedness Preservation). For all e such that $l_0; \emptyset \vdash e$, if $e \mapsto_o^* e'$ then $l_0; \emptyset \vdash e'$.

PROOF. By induction on the length of $e \mapsto_o^* e'$ using Lemma D.2 for the inductive step. \square

Lemma D.2 (One-Step Well-formedness Preservation). For all e such that $l_0; \emptyset \vdash e$, if $e \mapsto_o e'$ then $l_0; \emptyset \vdash e'$.

PROOF. By case analysis on the reduction rules for \mapsto_o . For most rules:

- (1) By assumption, e , the left-hand side of the arrow, is well-formed and Lemma D.3 yields that the redex is well-formed.
- (2) By inversion of the corresponding last inference rule in the derivation of well-formedness for the redex, we derive that the various sub-expressions of the redex are also well-formed.
- (3) By (1), (2), and the relevant inference rules for well-formedness, the result of the reduction of the redex is well-formed.
- (4) By Lemma D.4, placing that last expression in the evaluation context gives a well-formed expression e' .

In addition to the above steps:

- For rule APP-LAMBDA, at step 3, use Lemma D.5 to establish the well-formedness of the body of the function after substitution.
- For rule DO[>], after step 2, $l; \emptyset \vdash E^>[\text{do } v]$ where $l_0 > E > l$. Without loss of generality, let $l > E^> > l'$. By Lemma D.3, $l'; \emptyset \vdash \text{do } v$, and hence, $l'; \emptyset \vdash v$. Lemma D.6 and $l'; \emptyset \vdash v$ yields $l; \emptyset \vdash v_d$. From Lemma D.7 and $l; \{x : l\} \vdash x$, it follows that $l'; \{x : l\} \vdash e_c$, which from Lemma D.4 entails $l; \{x : l\} \vdash E^>[e_c]$. The rest of the proof for this case proceeds with the general step 4 from above.
- For DO-PAIR[◇], after step 2, $l; \emptyset \vdash E^◇[\text{do } v]$ where $l_0 > E > l$. Without loss of generality, let $l > E^◇ > l'$. From Lemma D.3, it follows that $l'; \emptyset \vdash \text{do } v$, and hence, $l'; \emptyset \vdash v$. From Lemma D.9 and $l'; \emptyset \vdash v$, $l; \emptyset \vdash v_c$, which from Lemma D.4 entails $l; \emptyset \vdash E^◇[v_c]$. The rest of the proof for this case proceeds with the general step 4 from above.
- For DO-FUN[◇], after step 2, $l; \emptyset \vdash E^◇[\text{do } v]$ where $l_0 > E > l$. Without loss of generality, let $l > E^◇ > l'$. By Lemma D.3, $l'; \emptyset \vdash \text{do } v$, and hence, $l'; \emptyset \vdash v$. Lemma D.8 and $l'; \emptyset \vdash v$ yields $l; \emptyset \vdash v_d$. The rest of the proof for this case proceeds with the general step 4 from above.

Lemma D.3 (Well-Formed Expressions Decompose to Well-formed Expressions). For all e_1 such that $l_1; \Gamma \vdash e_1$, if $e_1 = E[e_2]$, $e_2 \neq |e_3|^l$ and $l_1 > E > l_2$ then $l_2; \Gamma \vdash e_2$.

PROOF. By induction on the structure of E. \square

Lemma D.4 (Replacement in Context Preserves Well-formedness). For all e_1 and e_2 such that $l_1; \Gamma \vdash e_1$ and $l_2; \Gamma \vdash e_2$, if $e_1 = E[e_3]$, $e_3 \neq |e_4|^l$ and $l_1 > E > l_2$ then $l_1; \Gamma \vdash E[e_2]$

PROOF. By induction on the structure of E and Lemma D.3. \square

1716 **Lemma D.5 (Substitution Preserves Well-formedness).** For all e and v such that $l; \Gamma \uplus x : l \vdash e$
 1717 and $l; \emptyset \vdash v, l; \Gamma \vdash e[v/x]$.

1718 PROOF. By induction on the structure of e . □
 1719

1720 **Lemma D.6 (Pull Preserves Well-formedness).** For all e and v such $l; \emptyset \vdash e$ and $l'; \emptyset \vdash e$, if
 1721 $e = E[e_1]$, $e_1 \neq |e_2|^{l_1}$ and $l > E > l'$ then $l; \emptyset \vdash (\uparrow_o E)[v]$
 1722

1723 PROOF. By induction on the structure of E . □
 1724

1725 **Lemma D.7 (Push Preserves Well-formedness).** For all e and e' such $l; \emptyset \vdash e$ and $l; \Gamma \vdash e'$, if
 1726 $e = E[e_1]$, $e_1 \neq |e_2|^{l_1}$, $l > E > l'$ and $l'; \emptyset \vdash v$ then $l'; \Gamma \vdash (v \downarrow_o E)[e']$
 1727

1728 PROOF. By induction on the structure of E using Lemma D.6 for the case where E is of the form
 1729 $\text{mark}_j^{k,l}(v_d \blacktriangleright v_c) E'$. □

1730 **Lemma D.8 (Pull-Minus-Marks Preserves Well-formedness).** For all e and v such $l; \emptyset \vdash e$ and
 1731 $l'; \emptyset \vdash e$, if $e = E[e_1]$, $e_1 \neq |e_2|^{l_1}$ and $l > E > l'$ then $l; \emptyset \vdash (\uparrow_o^- E)[v]$
 1732

1733 PROOF. By induction on the structure of E . □
 1734

1735 **Lemma D.9 (Push-Minus-Marks Preserves Well-formedness).** For all e and e' such $l; \emptyset \vdash e$
 1736 and $l; \Gamma \vdash e'$, if $e = E[e_1]$, $e_1 \neq |e_2|^{l_1}$ and $l > E > l'$ then $l'; \Gamma \vdash (\downarrow_o^- E)[e']$
 1737

1738 PROOF. By induction on the structure of E . □
 1739

1740 **Proposition 5.5 (Ownership Erasure).** For all labeled e , $e \mapsto_o^* e'$ iff $\mathcal{O}(e) \mapsto^* \mathcal{O}(e')$.

1741 PROOF. By induction on the length of $e \mapsto_o^* e'$ for one direction, and the length of $\mathcal{O}(e) \mapsto^*$
 1742 $\mathcal{O}(e')$ for the other. □

1743 E PARAMETER CONTRACTS IN RACKET

1744 Integrating effectful software contracts with a natively imperative language poses a steeper chal-
 1745 lenge than adding them to an effect-handler language. It remains an open question whether a
 1746 complete integration is possible without major changes to the existing language. In this spirit,
 1747 this section presents a backwards-compatible extension to Racket's existing contract system that
 1748 covers contract-handler contracts only.⁸
 1749

1750 E.1 Parameter Contracts, By Example

1751 As prior examples have demonstrated, contract effects essentially ensure that contracts can set up,
 1752 and refer to, markings of dynamic extent in a declarative manner. Racket programmers deal with
 1753 dynamic extent via parameters [Gasbichler and Sperber 2005]. A parameter is a value container
 1754 that can store a different value for the duration of the dynamic extent of an expression's evaluation;
 1755 no matter how this evaluation proceeds, the original value is placed back into the parameter when
 1756 the dynamic extent ends (even via an exception or continuation jump). The Racket implementation
 1757 of parameters uses the already-mentioned continuation marks [Clements et al. 2001].

1758 Given this context, an extension to the contract system should enable contracts to set and refer to
 1759 parameters, turning ad hoc contract effects into (almost) declarative specifications. This is precisely
 1760 the purpose of *parameter contracts*. The remainder of this section illustrates this point with two
 1761 examples: contracts for generator yielding and contracts for function termination.
 1762

1763 ⁸This extension is available in the released version of Racket (as of 8.8).
 1764

1765 *Generator and Yield.* A *generator* is a procedure that may call the one-argument `yield` procedure
 1766 in its dynamic extent. When it does so, the evaluation of the generator is suspended and the
 1767 value handed to `yield` becomes the result of the generator. Once the generator is called again, its
 1768 evaluation resumes the (hidden) suspension until the next call to `yield`.

1769 Here is a simplistic example of a generator that produces all even natural numbers:

```
1770 (define evens
1771   (generator/f
1772     (λ ()
1773       (for ([k (in-naturals)])
1774         (yield (* 2 k))))))
```

1775 A generator is created by passing a thunk to the `generator/f` function.

1776 The key constraint is that `yield` should be invoked only in the dynamic extent of the thunk.
 1777 Without effectful contracts, such a constraint is documented informally and checked using interspersed
 1778 defensive checks.

1779 A parameter contract can express this constraint on the generator interface:

```
1780 (provide
1781   (contract-out
1782     [generator/f
1783      (->i ([thunk (->i () #:param in-gen? true any/c)]) [result generator?])]
1784     [yield
1785      (->i ([v any/c]) #:pre (in-gen?) [result any/c])]))
```

1786 Parameter contracts tend to come in pairs: a `#:param` clause in an `->i` contract sets up the context
 1787 and a precondition clause checks the context for the relevant information. In this example, the `->i`
 1788 contract on the thunk handed to `generator/f` ensures that the `in-gen?` parameter is set to `true`
 1789 when the thunk is run. Symmetrically, `yield` checks the value of this parameter in its precondition
 1790 to ensure that it is called in the dynamic extent of a call to the thunk.

1791 *Termination.* The literature on static analysis occasionally relies on termination checking, and
 1792 such checks often encode the size-change property (SCP) [Lee et al. 2001]. Nguyễn et al. [2019]
 1793 present an ad hoc contract for checking this property. Roughly, the contract keeps track of a call
 1794 graph that includes information about non-descending paths.

1795 A parameter contract can express this termination contract directly. It uses parameters to update
 1796 the call graphs:

```
1797 (define-syntax (total-> stx)
1800   (syntax-parse stx
1801     [(_ arg-ctc ... res-ctc)
1802      #:with (param ...) (generate-temporaries #'(arg-ctc ...))
1803      #'(self/c
1804         (λ _
1805           (define CG (make-parameter empty-call-graph))
1806             (->i ([param arg-ctc] ...)
1807                #:pre (graph-update CG (list param ...))
1808                #:param G (graph-update CG (list param ...))
1809                [result res-ctc]))))])
```

1810 The `total->` macro produces a termination contract. Its pieces are the argument and result
 1811 contracts; its result is an `->i` contract.

1812

1814 As with affine contracts, `self/c` is used to create the CG parameter at an appropriate time. This
 1815 parameter initially contains the empty call graph. When called, the `total->` contract's precondition
 1816 first checks whether updating the call graph with the new arguments would violate the SCP.
 1817 If so, `graph-update` returns `false` and the program signals a contract violation. Otherwise, the
 1818 `#:param` option extends the call graph with the new information about the arguments.

1819 Here is how this contract may be used in practice:

```
1820 (define ack
1821   (invariant-assertion
1822     (total-> integer? integer? integer?)
1823     (λ (m n)
1824       (cond
1825         [(= 0 m) (+ 1 n)]
1826         [(= 0 n) (ack (- m 1) 1)]
1827         [else (ack (- m 1) (ack m (- n 1)))]))))))
```

1829 The `invariant-assertion` construct attaches a contract to a function that is checked for all calls,
 1830 including recursive calls. In this case, the assertion promises, and checks, that the recursive `ack`
 1831 function terminates on whatever arguments it is given.

1832 E.2 Parameter Contracts Implementation, An Overview

1833 The modification of Racket's contract system consists of about 230 lines of code in the implemen-
 1834 tation of the `->i` combinator. Like `effect/racket`, the modification takes advantage of Racket's
 1835 existing libraries; but because the contract system is one of the more foundational pieces of the
 1836 Racket implementation, the extension cannot exploit abstractions from high-level layers.

1837 Key is Racket's expressive support for proxy values [Strickland et al. 2012]. Proxy values are able
 1838 to maintain expected invariants, such as equality between the original value and the proxy. Im-
 1839 portantly, procedure proxies already support manipulating continuation marks upon application.
 1840 The patch to `->i` takes advantage of a special internal value that serves as the continuation-mark
 1841 key for all parameterizations [Flatt and Dybvig 2020]. The value of this key is a mapping between
 1842 parameters and their assignments. Another internal function updates this mapping. When the
 1843 `#:param` option is set, the modified `->i` contract generates code that installs an updated value for
 1844 the parameter continuation-mark key. When the `#:param` option is missing, the contract does not
 1845 generate this code. In short, parameter contracts are a pay-as-you-go construct.

1846 Like `effect/racket`, the modification of Racket's contract system can guide the effort of others
 1847 to add parameter contracts to an existing contract system. If the underlying language comes with
 1848 a mechanism like continuation marks, the effort is straightforward. Otherwise, the implementer
 1849 may wish to consider adding a continuation-mark mechanism, because it has proven useful in
 1850 different ways [Chang et al. 2011; Clements and Felleisen 2004; Clements et al. 2001; Flatt and
 1851 Dybvig 2020].

1852 E.3 Comparing Effect-Handler and Parameter Contracts

1853 Table 2 summarizes how the two implementations deal with the properties present in the literature.
 1854 Like Table 1, rows correspond to existing pieces of literature; columns, though, correspond to
 1855 the two implementations. Keep in mind that the effect-handler language is *not* Racket but a new
 1856 language that implements the model faithfully; parameter contracts extend the existing Racket
 1857 contract library with a partial realization of the model.

1858 The cells of Table 2 are marked with \checkmark , \times , or \sim . A \checkmark in the table indicates that the implementa-
 1859 tion is able to faithfully express the contracts presented in the respective paper. A \sim mark means
 1860

1862

Table 2. Summary Evaluation (Full ✓, Partial ~, None ×)

	Effect-Handler Contracts	Parameter Contracts
1863		
1864		
1865		
1866	Chalin et al. [2006]	✓
1867	Tov and Pucella [2010]	✓
1868	Shinnar [2011]	✓
1869	Disney et al. [2011]	✓
1870	Keil and Thiemann [2015a]	✓
1871	Scholliers et al. [2015]	~
1872	Moore et al. [2016]	✓
1873	Dimoulas et al. [2016]	✓
1874	Bañados Schwerter [2016]	✓
1875	Williams et al. [2018]	✓
1876	Nguyễn et al. [2019]	✓
1877	Moy and Felleisen [2023]	✓
1878		

that the implementation can check the property with caveats. Finally, an × admits a failure; the implementation is not able to express the contracts presented in the corresponding paper.

As mentioned in Section 7.1, effect/racket can completely express all but one of the contracts present in the literature. By contrast, the column for parameter contracts shows a lack of expressive power in the revised contract library. While the library can still express—to some degree—half of the existing constructs, it certainly cannot cover the whole terrain. In other words, this column raises the research question of how an existing contract library could implement the model faithfully and thus become a universal framework.