

High-level effect handlers in C

MARIO ALVAREZ-PICALLO, Huawei Research Centre, United Kingdom

TEODORO FREUND, Huawei Research Centre, United Kingdom

DAN R. GHICA, Huawei Research Centre, United Kingdom

SAM LINDLEY, The University of Edinburgh, United Kingdom

Effect handlers are an expressive language feature allowing the programmer to define custom computational effects in a compositional and principled way. We introduce **libseff**, a library that brings the power of effect handlers to C. In contrast to other existing effect handler libraries, **libseff** is designed to be used directly from C by C programmers writing idiomatic, direct-style code. Through a series of examples and benchmarks we demonstrate the expressiveness that effect handlers can bring without sacrificing readability or performance.

1 INTRODUCTION

Effect handlers [25] are an increasingly popular programming feature that empowers programmers to define and use custom computational effects, ranging from exceptions to generators to lightweight threads, in a structured way. With an effect handler oriented programming language or library the programmer can define custom effectful operations whose semantics is specified later by a suitable *effect handler*. The power of handlers lies in their ability to support fine-grained customisation (a given effectful computation can be handled by different handlers that give it different behaviours, such as implementing a different scheduling strategy), and their composability (handlers can be composed to allow using multiple different effects in the same program).

A central aspect of effect handlers is that when handling an operation they are provided with an explicit representation of the *continuation* of the code that performed the operation (that is the rest of the computation from that point up to the point at which the handler was installed). A continuation is a first-class object that can be resumed immediately, aborted entirely, or delayed for later execution. In this sense, effect handlers can be seen as providing a form of first-class resumable exceptions, and allow for the implementation of sophisticated forms of control flow, such as *async/await*, exceptions, generators and varied forms of lightweight concurrency, entirely as user-defined libraries.

Though effect handlers are often deployed in the context of high-level functional programming languages such as OCaml [28], we believe that lower-level languages also stand to gain much from such features. Indeed, if one enumerates all of the features that are enabled by the introduction of effect handlers, the only language in common use today to lack all of these is C. On the other hand, the C ecosystem is rife with ad-hoc implementations of complex control-flow operators that are intended to support exactly these features, often on a per-project basis.

There already exist two effect handler libraries for C, **libhandler** [18] and **libmpeff** [19]. However, they are both geared towards compiler writers, with the explicit goal of providing a compilation target for high-level languages with effects, rather than being used directly from C by C programmers. In contrast, in this paper we introduce and evaluate **libseff**, a small effect handler library designed to be used as part of a C codebase to write efficient code that looks and feels as much as possible like idiomatic C.

The **libseff** library differs from prior approaches in several respects:

- Unlike **libhandler** which relies on stack-copying (unsafe in C as there may be pointers into the stack) and **libmpeff** which relies on virtual memory (not feasible for embedded systems), **libseff** supports segmented stacks for resizing stacks. (Stack resizing is often important for applications such as web servers that spawn many lightweight threads, each of which needs its own stack.)

- Unlike traditional effect handler implementations **libseff** is oriented around mutable coroutine objects rather than immutable continuation objects. This offers a simple way of avoiding allocating a new continuation object every time an effectful operation (such as yielding to another thread) is performed. Moreover, it provides a more familiar interface for C programmers, who may treat **libseff** like a conventional coroutine library and integrate the effectful features as necessary.
- Unlike traditional effect handler implementations there is no special form for dispatching on effects. Instead performed effects are reified as request objects which are then typically dispatched on using a standard `switch` statement.

The main contributions of this paper are the following:

- The design of **libseff**, illustrated through a series of examples that introduce techniques for programming with effects and handlers in C using **libseff** (§2).
- The implementation of **libseff**, including a description which details the runtime representation, low-level primitives, and stack-management strategy (§3).
- An empirical evaluation of the performance profile of **libseff** through a series of benchmarks that demonstrate that the abstraction and expressiveness offered by effect handlers can be implemented in C in concert with competitive performance (§4).

§5 discusses related work and §6 concludes and outlines planned improvements for **libseff**.

The supplementary material includes an appendix with a formal calculus and abstract machine that specifies the semantics of the variant of effect handlers underlying the design of **libseff**.

2 DESIGN

We introduce **libseff** and motivate its design by way of a series of examples that illustrate the features and common idioms of the library.

2.1 Mutable state

To illustrate the core features of the library we begin with mutable state as a simple, albeit somewhat artificial (C has built-in support for mutable state), example. The following code declares two new *effects* for reading and writing an integer state value.

```
1 DEFINE_EFFECT(get, 0, int64_t, {});
2 DEFINE_EFFECT(put, 1, void, { int64_t new_value; });
```

In order to define an effect we use the macro `DEFINE_EFFECT(name, tag, ret_ty, { param_decls... })`, which takes an effect name (`name`), a tag (`tag`), a return type (`ret_ty`), and a possibly empty collection of parameter declarations (`param_decls`). The snippet above declares effect `get`, which returns a value of type `int64_t` and takes no parameters, and an effect `put`, which does not return a value and takes a single parameter `new_value` of type `int64_t`. At this stage these effects have type signatures, but no implementation. Together they can be thought of as providing an interface to integer state.

Tags. As C macros do not provide a mechanism for generating fresh numeric tags, we require the user to manually provide a tag for each defined effect. It is the responsibility of the user to ensure that no two effects are assigned the same tag. In fact, different effects with identical tags may be used safely, provided that no code performs one effect within the scope of a handler for another effect that is assigned the same tag. Due to **libseff**'s use of 64-bit wide bitsets to represent handled effects, only numbers 0-63 may be used as effect tags.

Terminology. More properly, `get` and `put` are *effect operations* and conceptually we might group them together to form an interface for a single integer state *effect*. However, as in OCaml 5 [28]

99 **libseff** does not explicitly group such operations, and we refer to each individual effect operation
 100 as an *effect*. Elsewhere effect operations are sometimes referred to as *commands* [3, 9].

101 The following code uses the `get` and `put` effects to implement a countdown loop.

```
102 5 void* counter(void* parameter) {
103 6     int64_t counter;
104 7     do {
105 8         counter = PERFORM(get);
106 9         printf("Counter is %ld\n", counter);
107 10        PERFORM(put, counter - 1);
108 11    } while (counter > 0);
109 12    return NULL;
110 13 }
```

111 As C lacks closures and parametric polymorphism, any handled code must be defined inside a
 112 top-level function (here `counter`) conforming to the prototype `void* fn(void*)`. In order to perform
 113 an effect, we use the `PERFORM(name, {arg...})` macro, which takes an effect name and a possibly
 114 empty collection of arguments. This macro provides a convenient wrapper over the lower-level,
 115 untyped `seff_perform` primitive which we describe in detail in §3.2. From the perspective of an
 116 end-user of **libseff**, an invocation of `PERFORM` looks much like a function call whose parameter
 117 and return types match those declared by the corresponding `DEFINE_EFFECT` macro. In particular, the
 118 parameter and return types are checked by the C compiler.

119 If we were to call `counter` directly as a normal function at the top level, then this would result in
 120 a runtime error when line 8 is reached as it performs the `get` effect outside the scope of a handler
 121 for `get` (analogous to raising an exception outside the scope of an exception handler). The following
 122 code illustrates how to handle the effects inside `counter` by instantiating `counter` as a coroutine and
 123 then repeatedly resuming the coroutine inside an event loop that handles the performed effects.

```
124 14 int main(void) {
125 15     effect_set handles_state = HANDLES(get) | HANDLES(put);
126 16     seff_coroutine_t *k = seff_coroutine_new(counter, NULL);
127 17     seff_request_t req = seff_handle(k, NULL, handles_state);
128 18     int64_t state = 100;
129 19     bool done = false;
130 20     while (!done) {
131 21         switch (req.effect) {
132 22             CASE_EFFECT(req, get, { req = seff_handle(k, (void *)state, handles_state); break; })
133 23             CASE_EFFECT(req, put, {
134 24                 state = payload.new_value; req = seff_handle(k, NULL, handles_state);
135 25                 break; })
136 26             CASE_RETURN(req, {
137 27                 printf("The handled code has finished executing\n"); done = true;
138 28                 break; })
139 29         }
140 30     }
141 31     seff_coroutine_delete(k);
142 32     return 0;
143 33 }
```

143 The `handles_state` effect set encapsulates the ability to handle the `get` and `put` effects. The call
 144 `seff_coroutine_new(counter, NULL)` allocates a new coroutine object pointed to by `k` which when
 145 resumed will run the `counter` function with the argument `NULL`. The call `seff_coroutine(k, NULL,`
 146 `handles_state)` resumes the coroutine pointed to by `k` and handles the `get` and `put` effects. In fact, it

147

148 only handles them to the extent that if performed, the coroutine will be suspended and they will be
 149 packaged up in the returned request object `req`. The actual handling code appears in the enclosing
 150 context, here an event loop which dispatches on `req.effect`. The mutable integer state is stored in
 151 the `state` variable. Inside the `switch` statement there is one clause (expressed using the `CASE_EFFECT`
 152 macro) for each of the possible effects that the coroutine may perform and a distinguished return
 153 clause (expressed using the `CASE_RETURN` macro) for the case where the coroutine returns normally
 154 without performing any effects. A `get` effect is handled by resuming the coroutine, passing in the
 155 current state (recall that the return type of `get` is `int64_t`). A `put` effect is handled by updating the
 156 current state and resuming the coroutine with a `NULL` argument (recall that the return type of `put` is
 157 `void`). The special `payload` variable contains the new state passed to the `put` effect. If the coroutine
 158 returns without performing an effect then a message is printed and the event loop is exited. Finally
 159 the coroutine object is deleted using `seff_coroutine_delete`.

160
 161 *Decoupling effect interception from handling code.* Formally, the handler is simply the code that
 162 intercepts effects in the given effect set, yielding a corresponding request object. However, it is
 163 natural to refer to the code in the surrounding context that dispatches on the request object as a
 164 handler and we frequently do so. Conventional effect handlers fuse these two phases together, much
 165 like exception handlers, but we opt for a decoupled approach in `libseff` in order to circumvent
 166 the awkwardness of encoding a bespoke dispatch mechanism in C.

167 *Function signatures.* Type signatures for the three primitive functions seen so far are as follows.

```
168 seff_coroutine_t *seff_coroutine_new(void *(*fn)(void*), void *arg);
169 void seff_coroutine_delete(seff_coroutine_t* k);
170 seff_request_t seff_handle(seff_coroutine_t* k, void* arg, effect_set handled);
```

171
 172 The API does not differentiate between starting and resuming a coroutine. However, when called
 173 on a coroutine for the first time `arg` is ignored (the underlying function has already been applied to
 174 an argument supplied to `seff_coroutine_new`), whereas on subsequent calls the continuation of the
 175 coroutine is applied to `arg`, which corresponds to the value returned by the effect.

176
 177 *Coroutines as mutable continuations.* Traditional accounts of effect handlers do not take coroutines
 178 as primitive, but rather *continuations*. A continuation (also sometimes called a resumption) is an
 179 immutable object that represents the rest of a computation. In effect a continuation is like an
 180 immutable `seff_coroutine_t`, but in `libseff` we always manipulate coroutines as pointers to a
 181 mutable `seff_coroutine_t` object which is updated in place whenever an effect is handled.

182
 183 *Handlers in libseff are sheep handlers.* Traditional effect handlers are classified as deep or
 184 shallow [13]. A deep handler implicitly wraps itself around the continuation of a suspended effect,
 185 ensuring that all effects in a computation must be handled uniformly; a shallow handler does not.
 186 Following WasmFX [24], handlers in `libseff` are a hybrid sometimes called *sheep* handlers. Sheep
 187 handlers are: like shallow handlers in that the original handler need not be installed each time a
 188 continuation is resumed; and like deep handlers in that some handler (though not necessarily the
 189 original one) must be installed every time a continuation is resumed. In `libseff` this behaviour
 190 manifests as the need to supply an effect set every time we call `seff_handle` on a coroutine.

191 2.2 Lightweight concurrency

192 A much more compelling application of effect handlers, and the central motivation behind the
 193 initial development of `libseff`, is *lightweight concurrency*. We begin by defining two effects.

```
194 1 DEFINE_EFFECT(fork, 0, void, { void *(*fn)(void *); void *arg; });
195 2 DEFINE_EFFECT(yield, 1, void, {});
```

196

197 The `fork` effect takes a function pointer (`fn`) and an argument to apply it to (`arg`); it spawns a new
 198 thread that invokes `fn(arg)`. (In a language with closures we would typically implement `fork` as a
 199 one argument effect.) The `yield` effect suspends the current thread.

200 We write a small example application that initialises a root thread which is responsible for
 201 spawning 10 worker threads. These threads then each print 10 messages to the screen.

```
202 1 void *root(void *param) {
203 2     for (int64_t i = 1; i <= 10; i++) PERFORM(fork, worker, (void *)i);
204 3     return NULL;
205 4 }
206 5 void *worker(void *param) {
207 6     int64_t id = (int64_t)param;
208 7     for (int64_t iteration = 0; iteration < 10; iteration++) {
209 8         printf("Worker %ld, iteration %ld\n", id, iteration);
210 9         PERFORM(yield);
211 10    }
212 11    return NULL;
213 12 }
```

214 To run this code, we need to define a handler for the `yield` and `fork` effects which amounts to
 215 implementing a custom scheduler. The ability of effect handlers to describe APIs to communicate
 216 with a scheduler is at the heart of effect handlers' applications to concurrency [6, 7, 9, 24, 28, 29].

```
217 1 void with_scheduler(seff_coroutine_t *initial_coroutine) {
218 2     effect_set handles_scheduler = HANDLES(yield) | HANDLES(fork);
219 3     tl_queue_t queue;
220 4     tl_queue_init(&queue, 5);
221 5     tl_queue_push(&queue, initial_coroutine);
222 6     while (!tl_queue_empty(&queue)) {
223 7         seff_coroutine_t *next = (seff_coroutine_t *)tl_queue_steal(&queue);
224 8         seff_request_t req = seff_handle(next, NULL, handles_scheduler);
225 9         switch (req.effect) {
226 10            CASE_EFFECT(req, yield, {
227 11                tl_queue_push(&queue, (struct task_t *)next); break; })
228 12            CASE_EFFECT(req, fork, {
229 13                tl_queue_push(&queue, (struct task_t *)next);
230 14                seff_coroutine_t *new = seff_coroutine_new(payload.fn, payload.arg);
231 15                tl_queue_push(&queue, (struct task_t *)new);
232 16                break; })
233 17            CASE_RETURN(req, {
234 18                seff_coroutine_delete(next);
235 19                break; })
236 20        }
237 21    }
238 22 }
239 23 int main(void) { with_scheduler(seff_coroutine_new(root, (void*)0)); return 0; }
```

239 As in §2.1, the body of the handler is a `switch` statement nested inside a loop. The main difference
 240 with the state example is that now a variable number of coroutines are managed simultaneously by
 241 the scheduler, and these are stored in the task queue `queue`. On each iteration, the scheduler pops
 242 a coroutine off the head of the queue and proceeds to resume it with `seff_handle`. A `fork` or `yield`
 243 request is handled by pushing the suspended coroutine to the back of the queue. The `CASE_RETURN`
 244 clause is responsible for releasing the coroutine structures as they finish execution.

245

246 *One-shot continuations.* In performance-oriented implementations of effect handlers [24, 28] it is
 247 common to restrict continuations to be invoked at most once. This restriction simplifies the runtime
 248 system by precluding the duplication of continuations (which would involve creating a copy of
 249 the stack frame captured by the continuation). A similar limitation applies in **libseff**, which
 250 provides no facilities to copy stack frames. Doing so in C is inherently unsafe, as programmers often
 251 manipulate pointers into the stack which would be invalidated if the stack was copied elsewhere.
 252 However, in **libseff** there is no way to resume a continuation twice, as continuations per se are
 253 not exposed by the API – each time we handle a coroutine its continuation changes. On the other
 254 hand, a new kind of bug can occur if a coroutine pointer is copied accidentally (recall that we
 255 always refer to coroutines via a `seff_coroutine_t` pointer). For example, in the scheduler code above,
 256 if the programmer duplicated line 11 by accident, the coroutine `next` would be enqueued twice. This
 257 would not cause an immediate crash, but would lead to surprising behaviour: every time a thread
 258 were to yield it would subsequently be scheduled to run twice as often. However, once finished its
 259 coroutine object would be deleted and further attempts to dereference the other copy of the pointer
 260 in the queue would fail. It is important with **libseff** for the programmer to take care to manually
 261 manage the lifetime of coroutines, but this is quite standard for heap-allocated objects in C.

262 2.3 Resources

264 One technique supported by handlers, which we have thus far not seen, is the ability to “delay”
 265 a computation to be performed *after* an effect has been handled. This can be done by having the
 266 handler explicitly maintain a stack keeping track of all the effects that have been handled so far
 267 which is then “unwound” after a coroutine finishes execution. A more elegant approach is to write
 268 our handler as a recursive function, rather than a direct imperative loop, and writing additional
 269 code after the recursive call.

270 As a motivating example, we implement scoped resource handling using a single `defer` effect,
 271 whose purpose is to schedule a clean-up function `defer_fn` to be called with argument `defer_arg`
 272 when the enclosing coroutine ends its execution. We will also define our own variants of resource-
 273 allocating primitives (for this example, `malloc` and `fopen`), which immediately perform the `defer`
 274 effect to ensure that the corresponding clean-up function is called in a timely fashion.

```
275 1 DEFINE_EFFECT(defer, 0, void, { void (*defer_fn)(void*); void *defer_arg; });
276 2 void *malloc_scoped(size_t size) {
277 3     void *ptr = malloc(size); PERFORM(defer, free, ptr); return ptr;
278 4 }
279 5 FILE *fopen_scoped(const char *path, const char *mode) {
280 6     FILE *f = fopen(path, mode); PERFORM(defer, fclose, f); return f;
281 7 }
```

282 These functions may be used as drop-in replacements for `malloc` and `fopen`, the only caveat being
 283 that any code that uses them must be run inside a coroutine that handles the `defer` effect.

```
284 1 void *uses_resources(void *arg) {
285 2     ... void *ptr1 = malloc_scoped(256); ... void *ptr2 = malloc_scoped(512);
286 3     ... FILE *f = fopen_scoped("example", "r"); ...
287 4 }
```

289 Calling any of these scoped resource acquisition functions will result in the `defer` effect being
 290 performed, communicating the need for resource clean-up to any installed handler. One possible
 291 implementation for such a handler is given by the recursive function `handle_defer` below.

```
292 1 void *handle_defer(seff_coroutine_t *k) {
293 2     seff_request_t req = seff_handle(k, NULL, HANDLES(defer));
```

294

```

295 3     switch (req.effect) {
296 4         CASE_EFFECT(req, defer, {
297 5             void *result = handle_defer(k);
298 6             // Run the clean-up function
299 7             payload.defer_fn(payload.defer_arg);
300 8             return result; })
301 9         CASE_RETURN(req, { return payload.result; })
302 10    }
303 11 }

```

Observe that the structure is similar to a recursive version of the event loop of §2.2, with the crucial difference that the recursive call does not take place in tail position; instead, it is followed by a call to the deferred function. At runtime, the call stack of `handle_defer` will match the order in which the different invocations of `defer` were performed, and the corresponding clean-up functions will be called starting from the last.

We abstract away the creation and management of the coroutine object inside a helper function which takes as an argument the function pointer to be run within the scope of the `defer` handler. We can now run `uses_resources` like so:

```

312 1 void *run_with_handle_defer(void *(*fn)(void*), void *arg) {
313 2     seff_coroutine_t *k = seff_coroutine_new(fn, arg); handle_defer(k); seff_coroutine_delete(k);
314 3 }
315 4 int main(void) { run_with_handle_defer(uses_resources, NULL); }
316

```

2.4 Composition

An important property of effect handlers is their *composability* [13][11, Chapter 2]. This allows different libraries to define different effects which programmers can then mix within the same function. To illustrate effect handler composition, we use the `defer` effect from the previous section together with a new effect for defining generators. Throughout the rest of this subsection we assume that all the definitions from §2.3 are still in scope.

A generator is a function that yields a stream of multiple values, suspending its execution each time a value is produced and resuming from the same place next time it is invoked. In languages without native support for generators, they can be simulated by a global transformation. With effect handlers we can implement them directly using a single effect.

```

328 1 DEFINE_EFFECT(yield_str, 1, void, { char *elt; });

```

In this case, the `yield_str` effect yields a string. As we wish to compose it with `defer` (whose id is 0) we have taken care to give it the id 1.

Any function can now be turned into a generator by having it perform the `yield_str` effect. For example, we now define a generator that yields squares up to a certain number, formatted as heap-allocated strings. We use the previously-defined `malloc_scoped` function to reserve memory.

```

335 1 void *squares(void *arg) {
336 2     int64_t limit = (size_t)arg;
337 3     for (int64_t i = 0; i < limit; i++) {
338 4         char *str = malloc_scoped(32);
339 5         sprintf(str, "%5lu", i * i);
340 6         PERFORM(yield_str, str);
341 7     }
342 8     return NULL;
343 9 }

```

344 In order to access the elements of this generator, we must define a handler for it. A more sophisticated
 345 generator library could provide iteration combinators for consuming the elements of a generator.
 346 Here we simply define a `print_all` function that prints every element produced by this generator in
 347 sequence.

```
348 1 void *print_all(void *arg) {
349 2     seff_coroutine_t *k = seff_coroutine_new(squares, arg);
350 3     while (true) {
351 4         seff_request_t req = seff_handle(k, NULL, HANDLES(yield_str));
352 5         switch (req.effect) {
353 6             CASE_EFFECT(req, yield_str, { puts(payload.elt); break; })
354 7             CASE_RETURN(req, { seff_coroutine_delete(k); return NULL; })
355 8         }
356 9     }
357 10 }
```

358 If we run `print_all` directly, then it crashes on the first call to `malloc_scoped`, as there is no handler
 359 for `defer` in scope. Instead, we use the `run_with_handle_defer` combinator from §2.3.

```
360 1 int main(void) { run_with_handle_defer(print_all, (void*)50); }
```

361 This code prints the squares of all integers from 0 to 50, while also ensuring that all of the memory
 362 allocated by the underlying generator is freed. Notice that the handlers for the `yield_str` and `defer`
 363 are completely independent — they can be defined in separate modules and combined freely by the
 364 programmer.
 365

366 2.5 Overriding and default handlers

367 An effect `eff` is always handled by the innermost handler whose effect set includes `eff`. In contrast
 368 to function calls, where the callee is determined statically at compile-time, this allows us to redefine
 369 the handling of effects at runtime, providing a form of dynamic binding.

370 Consider a `print` effect for printing strings, along with a function `print_point` that formats a point
 371 given by two coordinates and prints it, and an example function that prints two points.
 372

```
373 1 DEFINE_EFFECT(print, 0, void, { char *msg; });
374 2 void print_point(int64_t x, int64_t y) {
375 3     char buffer[256];
376 4     sprintf(buffer, "{ x: %ld, y: %ld }");
377 5     PERFORM(print, buffer);
378 6 }
379 7 void *example(void *arg) { print_point(0, 0); print_point(1, 2); }
```

380 If `print` was simply a function then the behaviour would be fixed, but because it is an effect we can
 381 substitute in different implementation as runtime.

382 As an example, we define one handler that simply prints to standard output, and another one that
 383 redirects all output to a designated buffer. However, it would be cumbersome and inefficient if we had
 384 to install a handler every time we print. In this case, there is a reasonable default way to implement
 385 `print`: simply send the payload to `stdout`. The `libseff` library supports *default handlers* [7] which
 386 are functions of type `void (*)(void *)` that handle a given effect if no other handler is in scope.
 387 Default handlers, however, do not interrupt normal control flow of execution; instead, they are
 388 executed exactly as a normal function would, with control returning to the caller code immediately
 389 after executing the body of the handler. We now define a function `default_print` that is used as the
 390 default handler, as well as a more sophisticated handler that stores all output in a buffer instead.

```
391 1 void *default_print(void *print_payload) {
```

392

High-level effect handlers in C

```
393 2     EFF_PAYLOAD_T(print) payload = *(EFF_PAYLOAD_T(print) *)(print_payload);
394 3     fputs(payload.msg, stdout);
395 4     return NULL;
396 5 }
397 6 void *with_output_to_buffer(char *buffer, void *(*fn)(void*), void *arg) {
398 7     seff_coroutine_t *k = seff_coroutine_new(fn, arg);
399 8     while (true) {
400 9         seff_request_t req = seff_handle(k, NULL, HANDLES(print));
401 10        switch (req.effect) {
402 11            CASE_EFFECT(req, print, {
403 12                strcpy(buffer, payload.msg); buffer += strlen(payload.msg); break; })
404 13            CASE_RETURN(req, { seff_coroutine_delete(k); return payload.result; })
405 14        }
406 15    }
407 16 }
```

Note that the API for establishing default handlers is not type-safe: the payload of the handled effect is passed as a `void` pointer that must be manually cast to the correct type through the `EFF_PAYLOAD_T` macro, which desugars to the payload type of the given effect tag.

We can install `default_print` as a default handler by calling `seff_set_default_handler` and providing the id of the effect to be handled. For convenience, we provide the `EFF_ID` macro which expands to the id of the given effect.

```
413 1 int main(void) {
414 2     seff_set_default_handler(EFF_ID(print), default_print);
415 3     example(NULL);
416 4     char buffer[256];
417 5     with_output_to_buffer(buffer, example, NULL);
418 6 }
```

After installing the default handler, the direct call to `example` prints its output to the screen, whereas the call inside `with_output_to_buffer` is instead output to `buffer`.

3 IMPLEMENTATION

This section provides an overview of the implementation strategy for `libseff`, and some of the tradeoffs involved. Unlike other implementations [9, 17, 28] `libseff` does not keep a separate stack of handlers, but instead handlers coincide with coroutines: the context that resumed a coroutine becomes the handler for any effects that may be performed within the coroutine. As a coroutine executes, it keeps a pointer to its parent coroutine, creating a runtime configuration where the currently active coroutine acts as the top of a linked list of coroutines. This list plays a role analogous to the handler stack in other implementations, obviating the bookkeeping and additional allocations involved in keeping track of both continuations and handlers.

3.1 Runtime representation

During the execution of the program, any effectful computation is instantiated as an object of type `seff_coroutine_t`, which keeps track of the execution state of the coroutine and its environment as well as the set of effects that can be handled from it. More in detail, each coroutine object contains:

- The *state* of the coroutine, which can be one of `RUNNING`, `SUSPENDED` or `FINISHED`. These names are somewhat misleading and the values should be understood as preconditions to the `libseff` API: a value of `SUSPENDED` indicates that a coroutine can be resumed via `seff_handle` and a value of `RUNNING` indicates only that a coroutine can be suspended; multiple coroutines

can simultaneously be in the `RUNNING` state even in single-threaded applications despite only one of them being executed at a given point in time (This can happen when a coroutine spawns and resumes another coroutine: at this point, both parent and child are in the `RUNNING` state). Similarly, the child of a coroutine in the `SUSPENDED` state can itself be in the `RUNNING` state, indicating that it could suspend.

- The set of *handled effects*, which should not be understood as the effects that can be handled by this coroutine, but instead as the effects that can be handled by suspending this coroutine.
- A pointer to a *parent coroutine* used when performing an effect, to locate the corresponding handler.
- A *resumption state* containing the execution state when the coroutine was last resumed or suspended. More precisely, when the coroutine is in the `RUNNING` state, this field contains the execution state of the context that last resumed it, and is used for suspending the coroutine. When the coroutine is in the `SUSPENDED` state, this field instead holds the execution state of the coroutine at the moment of suspending, and is used for resuming it. The specific contents of the execution context are architecture-dependent but for x86-64 Linux, the only architecture currently supported, it consists of the instruction, stack and frame pointers as well as all callee-saved registers according to the standard System V calling convention.
- A pointer to a region in the heap containing the allocated *stack space* for the coroutine. As we shall explain in more detail in §3.3, `libseff` can use multiple approaches to stack management, depending on which this may be a pointer to a fixed-size, heap-allocated stack, or to a linked list of heap-allocated “stacklets”.

A pointer to the coroutine being currently executed (if any), as well as a pointer to the location of the system stack are also stored in global (thread-local, more precisely) variables. As we shall explain in §3.3.2, this information is used for avoiding allocating larger stack frames when calling library code from a coroutine. For a concrete example, consider the following code.

```

442 1 DEFINE_EFFECT(eff1, 0, void, {});
443 2 DEFINE_EFFECT(eff2, 1, void, {});
444 3 void *g(void *arg) { PERFORM(eff1); PERFORM(eff2); }
445 4 void *f(void *arg) {
446 5     seff_coroutine_t *k2 = seff_coroutine_new(g, NULL);
447 6     seff_request_t req1 = seff_handle(k2, NULL, HANDLES(eff2)); seff_request_t req2 = seff_handle(
448 7     k2, NULL, HANDLES(eff2));
449 7 }
450 8 void main() {
451 9     seff_coroutine_t *k1 = seff_coroutine_new(f, NULL);
452 10    seff_request_t req1 = seff_handle(k1, NULL, HANDLES(eff1) | HANDLES(eff2));
453 11    seff_request_t req2 = seff_handle(k1, NULL, HANDLES(eff1) | HANDLES(eff2));
454 12 }

```

It sets up two nested coroutines and performs effects `eff1`, `eff2` from the innermost one. After both coroutines have been created and started, and immediately before the call to `PERFORM(eff1)`, the state of the system is as depicted in Figure 1. Both coroutines have been instantiated and are in the `RUNNING` state, with the `current_coroutine` variable pointing to `k2`. Since both `k1` and `k2` are currently running, the resumptions stored in them contain the program state immediately before they were started (the resumption for `k1` points to the state right before the `seff_handle` call in line 11, and the resumption for `k2` to that in line 6).

When performing `eff1`, the linked list of coroutines is traversed upwards, starting at `k2`, to locate a suitable handler. In this case, `eff1` is featured in the handled effect set of `k1`, so `PERFORM(eff1)` immediately suspends `k1` and relinquishes control to its environment, which is then responsible

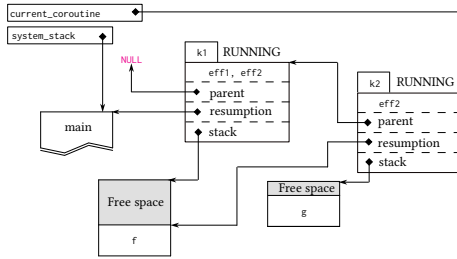


Fig. 1. Before `PERFORM(eff1)`

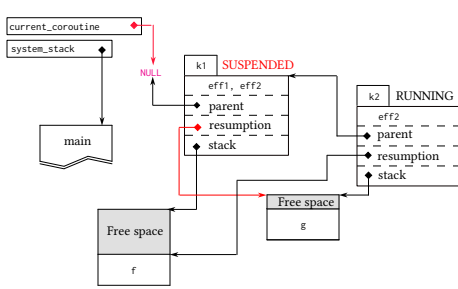


Fig. 2. After `PERFORM(eff1)`

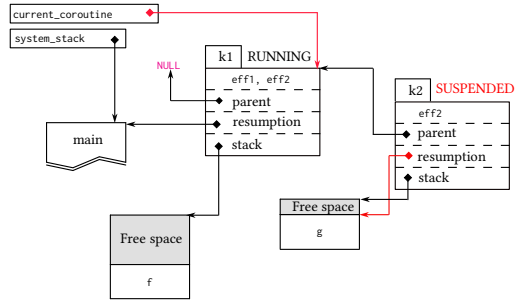


Fig. 3. After `PERFORM(eff2)`

for handling the effect. The system state at this point is depicted in Figure 2: coroutine `k1` is now suspended, and its resumption stores the program state immediately preceding the call to `PERFORM(eff1)`, inside the stack frame of the call to `g`. Note that `k2` remains unchanged and is still considered to be in a running state.

After the effect `eff1` has been (trivially) handled, execution of the suspended `k1` is resumed in line 12 and continues until the call to `PERFORM(eff2)` is reached. At that point, the stack of active coroutines is traversed again until a handler for `eff1` is reached; in this case, this effect can be handled directly by `k2`, and so `k2` is suspended and control is transferred back to line 6 in `f`. This corresponds to the diagram in Figure 3. Note how the resumption for `k1` is updated to hold the program state of the now-paused coroutine.

In order to make it efficient to perform an effect, `libseff` takes particular care when passing the payload of a command from the coroutine that performs the effect to the context that handles it. Effect payloads are marshalled, together with an effect tag, into a `seff_request_t` struct which effectively functions as an untyped discriminated union. In many high-level languages, creating such a data structure would involve allocating memory on the heap and thus incur a significant performance overhead. To avoid this, `libseff` uses two low-level tricks: first, the effect payload is allocated directly on the stack of the coroutine performing it, and the handler receives a pointer into this stack-allocated payload, which also saves the overhead of copying. Second, the `seff_request_t` struct consists of only two 64-bit fields, namely the tag and a pointer to the aforementioned payload, hence it can be returned from `seff_handle` directly via processor registers.

3.2 Primitives

Throughout the examples in the previous section, we have shown only the higher-level interface provided by `libseff`, which is intended for general use and provides convenience and some degree

540 of compiler checking of input and output types. Internally, operations such as the `PERFORM` macro
 541 are implemented in terms of a simpler set of primitives, that we describe here.

542 At the lowest level, we have three primitives for performing context-switching. These have the
 543 single task of resuming or suspending an active coroutine and are written directly in assembly.

```
544 1 seff_request_t seff_handle(seff_coroutine_t *k, void *arg, effect_set handled);
545 2 void *seff_yield(seff_coroutine_t *self, effect_id effect, void *payload);
546 3 void seff_exit(seff_coroutine_t *self, effect_id effect, void *payload);
```

547 We have already discussed `seff_handle`. `seff_yield` is responsible for suspending the given corou-
 548 tine `self` and returning control to the point where this coroutine was last resumed. The coroutine
 549 to be suspended is provided as an explicit argument to this function, and the caller is responsible
 550 for ensuring that, at the moment of invoking `seff_yield`, the coroutine to be suspended is either
 551 the current coroutine or an ancestor of it, otherwise the call to `seff_yield` will result in undefined
 552 behaviour¹. `seff_exit` behaves similarly to `seff_yield`, with the difference that a coroutine that is
 553 suspended via `seff_exit` is considered terminated and can no longer be resumed. This means that
 554 the execution context does not need to be saved, and so `seff_exit` is more efficient in general.

555 Note that both of the functions `seff_yield`, `seff_exit` take an `effect_id` argument, which is used
 556 to construct the `seff_request_t` object that will result from suspending the coroutine. However, this
 557 argument is not used to locate an appropriate handler. Instead, control is always relinquished to
 558 the last resumer of the given coroutine, whether or not it is able to handle the given effect.

559 A separate set of primitives is provided for looking up the appropriate handler in scope (if any)
 560 when selecting which coroutine to suspend.

```
561  

  562 1 seff_coroutine_t *seff_locate_handler(effect_id effect);
  563 2 void *seff_perform(effect_id effect, void *payload);
  564 3 void seff_throw(effect_id effect, void *payload);
```

565 Dynamic dispatch is taken care of by `seff_locate_handler`, which walks through the stack of
 566 currently active coroutines until it finds the first whose `handled_effects` bitset covers the effect
 567 effect. As explained before, this indicates the context which resumed that coroutine last is able to
 568 handle the corresponding effect.

569 `seff_perform` and `seff_throw` are analogous to `seff_yield` and `seff_exit`, except that they use the
 570 given `effect_id` to select which coroutine to suspend. Effectively, `seff_perform(e, p)` is equivalent
 571 to `seff_yield(seff_locate_handler(e), e, p)`, with the only difference that, if no appropriate handler
 572 can be found in scope, `seff_perform` will invoke a default handler, whereas the `seff_yield` version
 573 will dereference a null pointer.

574 The `PERFORM` macro (illustrated in §2) is the preferred method of performing an effect. It is defined
 575 as a thin wrapper over `seff_perform`. A call to `PERFORM(eff, args...)` simply constructs a payload
 576 object of type `EFF_PAYLOAD_T(eff)` on the current stack frame and initialises it with the provided
 577 arguments, then calls `seff_perform` with the effect and a pointer to the stack-allocated payload.
 578 Unlike in other systems with resizable stacks [23, 31], `libseff` guarantees that the stack area for a
 579 given coroutine always remains at the same location, hence pointers into the stack of a coroutine
 580 will remain valid while the coroutine is suspended.

581 3.3 Stack management

582
 583 One of the most important technical decisions when implementing stackful coroutines is how stack
 584 frames are allocated and, most importantly, resized. When designing `libseff`, we considered four

585 ¹It is possible to check this condition at runtime and fail gracefully if it is not satisfied, by traversing the list of coroutines
 586 and ensuring that the coroutine that is being suspended is reachable from the currently active one, but this would impose a
 587 prohibitive overhead.

589 different approaches, which we detail below. Currently, the first two (fixed and segmented stacks)
 590 are implemented and can be switched between via a build flag (but should not be mixed together in
 591 the same project). The third approach (overcommitting) is planned but currently unimplemented
 592 and we have deemed the last (stack copying) to be unsuitable for our target setting.
 593

594 **3.3.1 Fixed-size stacks.** The simplest approach to stack management is to reserve a fixed-size block
 595 of memory to hold the coroutine stack. This has the dual advantages of being simple to implement
 596 and introducing no any additional runtime overhead. However, it can result in a significant waste
 597 of memory. Given that it is hard to determine in advance how much stack space a given program
 598 will eventually need, the programmer must preemptively allocate larger stacks than necessary in
 599 order to mitigate against the risk of stack overflow.
 600

601 **3.3.2 Segmented stacks.** Segmented stacks, also called split stacks, replace the traditional con-
 602 tiguous fixed-size stack area by a linked list of stack segments or “stacklets”. The compiler then
 603 instruments every function with a small prelude that checks whether the current stack is large
 604 enough to accommodate the stack frame of the current function. If not, a new stacklet is allocated
 605 to hold the new frame.
 606

607 Conveniently, support for segmented stacks is provided by both GCC and Clang via the `-fsplit-`
 608 `stack` flag, which will add stack overflow checks to every function preamble. As shown in Figure 4,
 609 the compiler-generated prelude checks for a potential stack overflow and, if required, calls a routine
 610 `__morestack` which is responsible for allocating a new segment, copying any parameters that were
 611 passed through the stack, and setting the return address to point to an epilogue that frees the
 612 newly-allocated stacklet. A simple implementation of this routine is provided by Clang, but `libseff`
 613 defines its own instead in order to give us finer-grained control over memory allocation.

614 Though they enable the programmer to
 615 write code without concerning themselves
 616 with stack frame sizes, segmented stacks are
 617 not without disadvantages. If no memory needs
 618 to be allocated, the overhead of the function
 619 prelude is mostly negligible; however, it is possible
 620 for a function call inside a tight loop to
 621 require the repeated allocation and deallocation
 622 of a large segment, resulting in a significant
 623 slowdown. This is sometimes known as the
 624 “hot split” problem and caused Go to move
 625 away from segmented stacks [23]. `libseff` miti-
 626 gates this issue by holding its stacklets in a
 627 doubly-linked list; when a stacklet is no longer
 628 necessary, instead of being released immedi-
 629 ately it is simply kept at the end of this list.
 630 Then, if a later function call requires the allocation
 631 of additional stack memory, this stacklet can
 632 be recycled, avoiding the need for an additional
 633 allocation. As we shall see in §4.1.2, with this
 634 optimisation, the cost of calling a function inside
 635 a hot split loop is 9x the cost of a normal function
 636 call.
 637

```

1  lea -0x108(%rsp),%r11
2  cmp %fs:0x70,%r11
3  ja 8
4  mov $0x108,%r10d
5  mov $0x0,%r11d
6  call <__morestack>
7  ret
8  push %rbp
9  mov %rsp,%rbp
10 ...
11 pop %rbp
12 ret
    
```

} Stack overflow check

} Segment and argument size

```

1 void split_stack() {
2     char buffer[256];
3     ...
4 }
    
```

Fig. 4. Segmented stack prelude in Clang-12

633 We argue that, in practice, this is not a significant problem: if the cost of the hot split overhead
 634 dominates the execution time of the called function, then it is likely that this is a small function
 635 that should get inlined by the compiler. Even then, there is a lot of space for further improvement:
 636 micro-optimisations such as lowering the segment reuse code path to assembly, analysis-based
 637

638 optimisations like preemptively inlining functions that are likely to cause a hot split, or even more
 639 sophisticated runtime detection of these cases [21].

640 One more concern about code using segmented stacks is interoperability with library code. The
 641 use of segmented stacks relies on instrumenting every function with an overflow check, but any
 642 functions that are compiled separately (including the standard library, unless the user builds it from
 643 scratch with support for segmented stacks enabled) will lack this prelude and any stack overflow
 644 will cause a crash or, in the worst case scenario, silent memory corruption. To avoid this issue,
 645 Clang’s implementation of segmented stacks conservatively requests a much larger amount of
 646 stack space if a function calls any other functions that have been compiled without segmented
 647 stack support. This is a sensible compromise, but it can lead to much higher memory consumption
 648 than necessary.

649 When using `libseff`, this overhead can usually be avoided: a function that was not compiled
 650 with segmented stacks enabled cannot make use of the context-switching features of `libseff`,
 651 therefore it can be run directly on the system stack instead of the stack of whichever coroutine
 652 happens to be executing. This obviates the need to preemptively allocate a larger segment. For this
 653 purpose, `libseff` defines the `MAKE_SYSCALL_WRAPPER` macro, which wraps a given function in code
 654 that handles switching to and from the system stack.

```
655 1 MAKE_SYSCALL_WRAPPER(int, puts, const char *s);
656 2 // Expands to:
657 3 int __attribute__((no_split_stack)) puts_syscall_wrapper(const char *c);
658 4 __asm__("puts_syscall_wrapper:"
659 5         "movq %rsp, %fs:_seff_paused_coroutine_stack@TPOFF;"
660 6         "movq %fs:_seff_system_stack@TPOFF, %rsp;"
661 7         "movq %fs:0x70, %rax;"
662 8         "movq %rax, %fs:_seff_paused_coroutine_stack_top@TPOFF;"
663 9         "movq $0, %fs:0x70;"
664 10        "callq puts;"
665 11        "movq %fs:_seff_paused_coroutine_stack@TPOFF, %rsp;"
666 12        "movq %fs:_seff_paused_coroutine_stack_top@TPOFF, %rcx;"
667 13        "movq %rcx, %fs:0x70;"
668 14        "retq;");
```

669 In the example above, a new function `puts_syscall_wrapper` is defined which has the same interface
 670 as the standard library function `puts` but will switch to the system stack instead of allocating a stack
 671 segment. We warn the user that this macro is only correct when the wrapped function takes all
 672 parameters and returns its result through processor registers. In addition, the wrapped function
 673 must only be called from within a coroutine (outside a coroutine, the original function should be
 674 called instead).

675 **3.3.3 Overcommitting.** Another approach to avoid stack overflow without the need for physically
 676 resizing a coroutine’s stack is to use overcommitting and reserve a large amount of (virtual) memory
 677 for each coroutine, leaving it to the operating system to allocate physical memory as necessary.
 678 This approach is used by `libmprompt`, striking an excellent balance between performance and
 679 convenience in systems that support it. However, we intend `libseff` to be also deployable in
 680 embedded systems, which do not necessarily provide virtual memory or a large address space.
 681 Thus, while we are planning to eventually provide virtual memory-based stack management for
 682 `libseff`, it is not among our top priorities.

684 **3.3.4 Stack copying.** Finally, a popular approach in managed languages is *stack copying*: coroutines
 685 are initialised with a small, fixed-size stack and dynamic checks for stack overflow are inserted
 686

687 (much like in the case of segmented stacks). However, whenever a coroutine requires more stack
 688 space than is available, instead of initialising a new segment, an entire contiguous region is allocated
 689 to serve as the new stack and the contents of the old stack are copied onto it. This approach avoids
 690 the hot split problem, although it incurs the extra cost of stack copying when a resize is needed.
 691 However, it is unsuitable for a low-level language like C since the process of copying the stack
 692 necessarily invalidates any pointers into it.

693 Other solutions exist. Go automatically rewrites any pointers into the stack as it is being copied,
 694 but this relies on an amount of runtime information which is simply not available to a C program.

695

696 4 EVALUATION

697 We evaluate **libseff** on a range of benchmarks comparing it to other effect handler implementa-
 698 tions as well as other concurrency mechanisms. All benchmarks were run on an Intel® Xeon® Gold
 699 6154 x86-64 running Ubuntu 20.04, with the clang 12.0.0 compiler. Except when stated otherwise
 700 we used **libseff** with segmented stacks.

701

702 4.1 Microbenchmarks

703 All benchmarks in this section are single-threaded.

704

705 *4.1.1 State.* Our first microbenchmark is based on the mutable state example of §2.1.

706

```
707 1 void *stateful(void *depth) {
708 2   if (depth == 0){ for (int i = PERFORM(get); i > 0; i = PERFORM(get)) PERFORM(put, i - 1);
709 3   } else {
710 4     seff_coroutine_t *k = seff_coroutine_new(stateful, (void *) (uintptr_t)(depth - 1));
711 5     seff_handle(k, NULL, HANDLES(error)); seff_coroutine_delete(k);
712 6   }
713 7   return NULL;
714 8 }
```

714 The *stateful* function recursively builds a stack of nested handlers for the error effect up to a
 715 specified depth. In the base case a counter, implemented using *get* and *put*, is decremented in a loop.

716 In any effect handler framework, performing an effect involves two steps: (a) locating the
 717 appropriate handler; and (b) transferring control to the handler. This benchmark measures the cost
 718 of both steps and how they scale depending on the number of times an effect is performed and the
 719 depth of the target handler.

720 In order to separate out the cost of locating the handler from that of transferring control to the
 721 handler, we implement two versions of the benchmark. The first is the one above, where every
 722 execution of an effect triggers a search for its handler. The second is an optimised version that
 723 arises from observing that the handlers for the *get* and *put* effects never change during execution of
 724 the loop, which allows us to locate the handlers once and then yield directly to the coroutine that
 725 handles them. This is shown in the code below, where **YIELD** wraps *seff_yield*, in the same way that
 726 **PERFORM** wraps *seff_perform*. If **libseff** were used as a backend for a higher-level language with
 727 effects, a compiler could apply this optimisation.

728

```
729 3   seff_coroutine_t *put_handler = seff_locate_handler(EFF_ID(put));
730 4   seff_coroutine_t *get_handler = seff_locate_handler(EFF_ID(get));
731 5   for (int i = YIELD(get_handler, get); i > 0; i = YIELD(get_handler, get))
732 6     YIELD(put_handler, put, i - 1);
```

733 We compare against several libraries. For each library we implement a general case and an
 734 optimal case to compare against both of our implementations: **native** is plain C without effect

735





Fig. 5. State benchmark results.

Multiplications	0	5	10	15	20
native	1.30	1.00	1.00	1.00	1.00
libseff fixed	1.00	1.00	1.00	1.00	1.00
libseff baseline	1.10	1.00	1.00	1.00	1.00
libseff hot split	9.00	1.83	1.06	1.00	1.00
libseff dealloc	32.22	6.47	3.68	2.50	2.08

(a) Relative execution time of the hot split benchmark.



(b)

Fig. 6. Hot split results

handlers or any kind of dynamic dispatching of operations; **cpp-effects** [9] is a C++ library for effect handlers; **libhandler** [17, 18] and **libmpeff** [19] are other C libraries for effect handlers.

The **cpp-effects** optimal avoids handler lookups in a similar fashion to **libseff**, but it also eliminates context switching by requiring that handler for `get` and `put` resume immediately and do not need to capture a continuation. The **libmpeff** and **libhandler** optimal cases also similarly avoid context switches, but they do not allow for caching handler lookups.

Figure 5a shows the general case. All effect handler implementations degrade significantly as the number of installed handlers increases, with **libseff** consistently the fastest. Figure 5b shows the optimal case. The elimination of traversing the stack of handlers gives **libseff** and **cpp-effects** a distinct advantage. For both **libhandler** and **libmpeff** the optimal case is still affected by the level of recursion. Whereas **libhandler** speeds up by avoiding copying the stack in this case, **libmpeff** shows little improvement over the general implementation when nested handlers are introduced. Figure 5c shows the optimal case with depth smaller than 10. Whereas, **libmpeff** and **libhandler** are initially faster than both **cpp-effects** and **libseff**, the cost of searching for handlers quickly becomes a bottleneck, becoming slower than **libseff** after depth 3.

4.1.2 Hot Split. The next benchmark is designed to quantify the cost of the hot split problem, as discussed in §3.3.2. It forces a function call to require more stack space than available in the current segment, and therefore request a larger one every time it is called. This function is then repeatedly called from a tight loop executing 10^8 times.

We compare four different configurations for **libseff** against the optimal case in plain C without segmented stacks, where a function call translates to exactly one assembly call operation. We vary the called function slightly to include a number of floating point multiplications, ranging from 0 and 20. Figure 7 shows the resulting compiled code for the **native** case, where the `movsd` instructions are only present if there is at least one multiplication and in between there is a fixed number of

785 `mulsd` operations. Lines 1 and 9 reserve and release a large array on the stack. When using `libseff`
 786 with segmented stacks, this function includes runtime checks similar to those of Figure 4.

```

787
788
789 1 sub    $0x788,%rsp
790 2 movsd  0x191(%rip),%xmm0
791 3 movsd  0x2099(%rip),%xmm1
792 4 mulsd  %xmm0,%xmm1
793 5 ...
794 6 mulsd  %xmm0,%xmm1
795 7 movsd  %xmm1,0x2041(%rip)
796 8 lea   -0x80(%rsp),%rax
797 9 add    $0x788,%rsp
798 10 retq
  
```

798 Fig. 7. Hot-split function body in assembly. m is
 799 the number of multiplications

800
 801 increase the number of operations executed by, the function the relative overhead incurred by the
 802 segment switching rapidly diminishes. The cost of a mere 10 multiplications dominates this over-
 803 head, and the difference in cost of the function call becomes negligible, as shown in Figure 6b.
 804 The results for *dealloc* illustrate the significant performance difference that recycling segments
 805 provides.

806 When no multiplications are inserted, we observed that changing the position of the functions in
 807 the compiled code, by adding a few `nop` instructions, could affect performance by up to 40%, which
 808 explains why both *fixed* and *baseline* are faster than *native*. A similar behaviour was noted in [28].
 809 when evaluating the cost of low level operations.

810 It is worth noting that the hot split problem is only observable at all if functions are not inlined
 811 or completely optimised away. Modern compilers invariably do inline functions that are simple
 812 enough that the call is the dominant cost. To force the compiler to generate the code from Figure
 813 7 without writing the assembly instructions by hand, we had to deliberately disable inlining and
 814 introduce empty inline assembly blocks so it would still compile both the function and the actual
 815 call.

816 4.2 Macrobenchmarks

817 In this section we benchmark `libseff` against other systems running whole applications.

818
 819 4.2.1 *HTTP Server*. Our first macrobenchmark is a simple HTTP server as used to benchmark
 820 OCaml 5 (formerly Multicore OCaml) [28, 30]. The server receives `GET` requests and respond to
 821 them asynchronously with a `text/plain` message. It serves each request with a single coroutine that
 822 is released when the connection is closed.

823 We compare against three alternative implementations:

- 824 • *nethttp_go* is built using Go’s `net/http` package.
- 825 • *rust_hyper* is a server built on top of `Hyper`, a highly performant HTTP library for Rust for
 826 the Tokio runtime, a state of the art runtime for Rust `async/await` concurrency.
- 827 • *cohttp_eio* is a server implemented for OCaml 5 over an effect based I/O library [29] and an
 828 HTTP library built on top of it [22].

829
 830 The three variants are rather diverse in that they include an extremely simple to implement
 831 server (*nethttp_go*), a low-level highly performant server (*rust_hyper*), and a server built on top of
 832 effect handlers like ours (*cohttp_eio*).

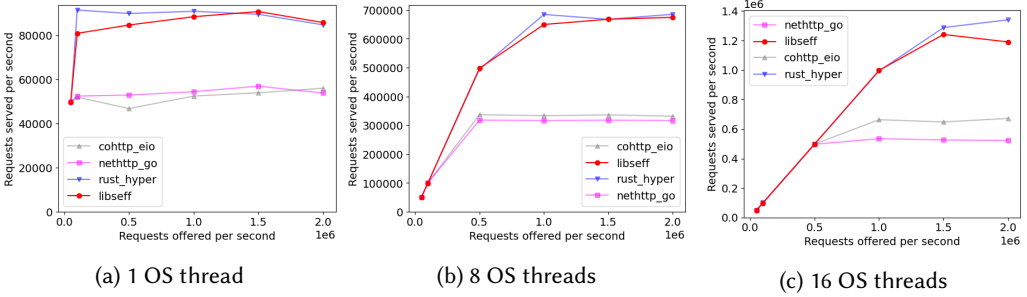


Fig. 8. Requests per second served per offered, with 1000 live connections.




Fig. 9. Maximum memory consumed, with 800000 requests offered per second.

Figure 8 shows the speeds running on 1, 8, and 16 OS threads. Both **libseff** and **rust_hyper** perform consistently better than **nethttp_go** and **cohttp_eio**, regardless of the number of OS threads. Figure 9 shows maximum memory consumption. We observe the maximum memory used by each implementation by varying the number of live connections, which coincide with the maximum number of coroutines spawned by each implementation.

The **libseff** implementation is built on top of a multi-threaded work-stealing scheduler based on effect handlers that takes advantage of the fact that **libseff**'s coroutines can be safely moved between OS threads. We used async I/O functions built on top of that scheduler and a small but complete HTTP request parser [26].

4.2.2 Prefetching. Our next benchmark, inspired by Jonathan et al. [12], uses C++'s coroutines to improve performance of memory heavy applications by alternating multiple concurrent runs and prefetching memory locations to cache before executing reads. The application executes multiple binary searches of different values over the same array. The array is big enough to not fit entirely in cache, and accesses to memory are not linear, making cache misses a significant part of the cost.

The naïve version executes searches sequentially; both the C++ coroutine implementation and the **libseff** implementation interleave multiple searches. Each search hints to the CPU to prefetch some address from memory and then, in a round-robin manner, moves on to the next search. Before execution returns to the coroutine that requested the prefetch, the values will already be stored in the cache and ready to be read, minimising cache




Fig. 10. Prefetch benchmark results. Large shapes mark the fastest execution for each framework.

883 misses. By varying the number of concurrent searches we can minimise the waiting time between
 884 execution returning to the search and the read being completed.

885 Whereas the C++'s coroutine implementation is explicit about prefetching, then yielding execu-
 886 tion, and finally reading the memory upon return, ours considers dereferencing a memory location
 887 to be an effect; it is the responsibility of the handler to prefetch the memory location, suspend the
 888 coroutine for some time and eventually provide it with the contents of the memory.

```
889
890 1 bool seff_binary_search(int const *first, size_t len, int val) {
891 2   while (len > 0) {
892 3     size_t half = len / 2; int x = PERFORM(deref, first + half);
893 4     if (x < val) { first += half + 1; len = len - half - 1; }
894 5     else { len = half; }
895 6     if (x == val) return true;
896 7   }
897 8   return false;
898 9 }
```

899 This code is a simplified version of binary search from the **libseff** benchmark. The main difference
 900 from regular binary search is the effectful computation of the dereference operation `PERFORM(deref,`
 901 `first + half)`. Figure 10 shows the results. They show that **libseff** incurs an overhead over the
 902 naïve sequential implementation whenever too few or too many streams are used but significantly
 903 improves upon it at its best point, taking around 2/3 of the time. The version with C++ coroutines
 904 is noticeably faster and allows for more concurrent searches to be executed simultaneously; this
 905 is unsurprising as C++ stackless coroutines have full compiler support and leverage a smaller
 906 memory footprint and stack allocation for better cache locality. Nonetheless, these results show
 907 that **libseff** effects are lightweight and efficient enough to materialise performance gains from
 908 cache prefetching, without these being obscured by context-switching overhead.

909 5 RELATED WORK

910 *Effect handlers for C and C++.* The **libhandler** [17, 18] and **libmpeff** [19] are existing effect
 911 handlers libraries for C. Unlike **libseff** they are designed as targets for compiler writers rather
 912 than for writing code directly in C. Each of these C libraries uses a different stack-management
 913 strategy: **libhandler** copies stacks into a temporary structure before restoring them on resumption,
 914 **libmpeff** uses virtual memory to allow stacks to grow without moving in memory, and **libseff**
 915 uses segmented stacks. The **cpp-effects** library [9] is a C++ effect handlers library which heavily
 916 relies on C++ features both in its implementation and its API.

917 *Coroutines in C/C++.* There exist many different coroutine libraries for C and C++, including
 918 Boost coroutines [2], **libco** [31], **libmill** [1], and C++20 stackless coroutines.

919 *Varieties of coroutine.* de Moura and Ierusalimsky [5] give a comprehensive classifications of
 920 the different notions of coroutine. The kind of coroutines provided by **libseff** are *asymmetric*
 921 *first-class stackful* coroutines. Moreover, **libseff** provides stacks that are guaranteed not to be
 922 moved and coroutines that can migrate between system threads.

923 *Effect handlers as coroutines.* A distinctive aspect of effect handlers in **libseff** is their foundation
 924 on mutable asymmetric coroutines rather than immutable continuations. Nonetheless, the close
 925 connections between asymmetric coroutines and effect handlers have been exploited, in a somewhat
 926 different way, elsewhere. Kawahara and Kameyama [15] show how to translate one-shot effect
 927 handlers into asymmetric coroutines. Phipps-Costin et al. [24] exploit essentially the same encoding

928

to implement effect handlers on top of the Wasmtime Fiber API [4], which implements coroutines for the Wasmtime engine for WebAssembly.

Optimising effect handlers. Much of the research on effect handlers has focused on programming and reasoning about them. Nonetheless, there have also been various attempts to compile effect handlers to efficient code. Kammar et al. [13] take advantage of Haskell’s aggressive inlining of type classes to speed up an implementation based on a continuation monad. Wu and Schrijvers [32] explain the essence of this optimisation as an instance of fusion. Kiselyov and Ishii [16] introduce so-called “Freer monads” as another means to speed up implementations of effect handlers in Haskell. Karachalias et al. [14] optimise effect handler code by aggressively inlining as many handlers as possible. Schuster et al. [27] achieve a similar end by way of staged computation. Xie and Leijen [33] and Ghica et al. [9] apply an instance of the optimisation we describe in Section 4 to avoid searching the handler stack, or indeed context-switching at all, when a handler is known to be “tail-resumptive” meaning that it immediately invokes the continuation in tail-position. Another optimisation performed by both of the latter two systems is for the case in which the continuation is never invoked (as in exception handlers).

6 CONCLUSION AND FUTURE WORK

We have described the design and implementation of **libseff**, a library for effect handlers in C. While other effect handlers for C exist, these are primarily designed as targets for compilers, whereas **libseff** is the first library to provide an idiomatic interface for programming with effect handlers in C. The key challenge we had to overcome is C’s lack of modern features, especially closures and generics. This led us to a design based on sheep handlers, coroutines, and explicit request objects, which enables writing handlers as simple, direct-style loops that should be familiar to any C developer.

Our benchmarks demonstrate that effect handler programs, even without special compiler support, can be compiled to efficient code that is competitive with other state-of-the-art approaches, notably Rust’s stackless coroutines. The **libseff** library outperforms most other libraries in this space due to simpler handler dispatch logic and hand-written context-switching code. It is also, to our knowledge, the first such C library to offer a choice of stack management strategies, currently supporting both segmented and fixed-size stacks, with planned support for a third approach based on overcommitting of virtual memory.

We are currently actively working on porting **libseff** to other architectures including ARM and 32-bit Intel processors. We also plan to support more approaches to stack management: both virtual memory and some form of arena allocation of stack segments are under consideration.

Perhaps unsurprisingly for a C library, the interface provided by **libseff** is prone to certain kinds of errors: using coroutines non-linearly or performing unhandled effects, can crash the program at runtime. The C type system is not rich enough to encode the necessary constraints to avoid these errors, but we would like to be developing a set of Rust bindings on top of **libseff** that will leverage Rust’s rich type system and borrow checker to ensure safety at compile time.

Finally, we plan to explore further low-level improvements to the **libseff** implementation. A common optimisation in other effect handler libraries is to avoid creating new continuations or stack frames for certain effects where the continuation is either never invoked (such as exceptions) or invoked immediately at the end of the handler (such as mutable state or dynamic binding). This promises significant additional performance for such use-cases allowing us to efficiently take fuller advantage of the expressive power of effect handlers.

REFERENCES

- [1] [n. d.]. libmill. <https://github.com/sustrik/libmill>.
- [2] Boost. [n. d.]. Boost.coroutine library. https://www.boost.org/doc/libs/1_83_0/libs/coroutine/doc/html/index.html.
- [3] Lukas Convent, Sam Lindley, Conor McBride, and Craig McLaughlin. 2020. Doo bee doo bee doo. *J. Funct. Program.* 30 (2020), e9.
- [4] Alex Crichton. 2021. Wasmtime Fiber API. https://docs.wasmtime.dev/api/wasmtime_fiber/index.html. Accessed 2023-11-15.
- [5] Ana Lúcia de Moura and Roberto Ierusalimsky. 2009. Revisiting coroutines. *ACM Trans. Program. Lang. Syst.* 31, 2 (2009), 6:1–6:31.
- [6] KC Sivaramakrishnan Deepali Ande, Sudha Parimala. 2023. Effectively Composing Concurrency Libraries. Draft. https://kcsrk.info/papers/composable_concurrency.pdf.
- [7] Stephen Dolan, Spiros Eliopoulos, Daniel Hillerström, Anil Madhavapeddy, KC Sivaramakrishnan, and Leo White. 2017. Concurrent System Programming with Effect Handlers. In *TFP (Lecture Notes in Computer Science, Vol. 10788)*. Springer, 98–117.
- [8] Matthias Felleisen and Daniel P. Friedman. 1986. Control Operators, the SECD-machine, and the λ -Calculus. In *Formal Description of Programming Concepts III*. Elsevier, 193–217.
- [9] Dan R. Ghica, Sam Lindley, Marcos Maroñas Bravo, and Maciej Piróg. 2022. High-level effect handlers in C++. *Proc. ACM Program. Lang.* 6, OOPSLA2 (2022), 1639–1667.
- [10] Daniel Hillerström, Sam Lindley, and Robert Atkey. 2020. Effect handlers via generalised continuations. *J. Funct. Program.* 30 (2020), e5.
- [11] Daniel Hillerström. 2021. *Foundations for Programming and Implementing Effect Handlers*. Ph. D. Dissertation. School of Informatics, The University of Edinburgh, Scotland, UK.
- [12] Christopher Jonathan, Umar Farooq Minhas, James Hunter, Justin J. Levandoski, and Gor V. Nishanov. 2018. Exploiting Coroutines to Attack the "Killer Nanoseconds". *Proc. VLDB Endow.* 11, 11 (2018), 1702–1714. <https://doi.org/10.14778/3236187.3236216>
- [13] Ohad Kammar, Sam Lindley, and Nicolas Oury. 2013. Handlers in action. In *ICFP*. ACM, 145–158.
- [14] Georgios Karachalias, Filip Koprivec, Matija Pretnar, and Tom Schrijvers. 2021. Efficient compilation of algebraic effect handlers. *Proc. ACM Program. Lang.* 5, OOPSLA (2021), 1–28.
- [15] Satoru Kawahara and Yukiyoshi Kameyama. 2020. One-Shot Algebraic Effects as Coroutines. In *TFP (Lecture Notes in Computer Science, Vol. 12222)*. Springer, 159–179.
- [16] Oleg Kiselyov and Hiroshi Ishii. 2015. Freer monads, more extensible effects. In *Haskell*. ACM, 94–105.
- [17] Daan Leijen. 2017. Implementing Algebraic Effects in C – "Monads for Free in C". In *APLAS (Lecture Notes in Computer Science, Vol. 10695)*. Springer, 339–363.
- [18] Daan Leijen. 2019. libhandler. <https://github.com/koka-lang/libhandler>.
- [19] Daan Leijen and KC Sivaramakrishnan. 2023. libmprompt and libmpeff. <https://github.com/koka-lang/libmprompt>.
- [20] Paul Blain Levy, John Power, and Hayo Thieleck. 2003. Modelling environments in call-by-value programming languages. *Inf. Comput.* 185, 2 (2003), 182–210.
- [21] Zhiyao Ma and Lin Zhong. 2023. Bringing Segmented Stacks to Embedded Systems. In *Proceedings of the 24th International Workshop on Mobile Computing Systems and Applications, HotMobile 2023, Newport Beach, California, February 22–23, 2023*. ACM, 117–123. <https://doi.org/10.1145/3572864.3580344>
- [22] Mirage. [n. d.]. Eio. <https://github.com/mirage/ocaml-cohttp>.
- [23] Daniel Morsing. [n. d.]. How Stacks are Handled in Go.
- [24] Luna Phipps-Costin, Andreas Rossberg, Arjun Guha, Daan Leijen, Daniel Hillerström, KC Sivaramakrishnan, Matija Pretnar, and Sam Lindley. 2023. Continuing WebAssembly with Effect Handlers. *Proc. ACM Program. Lang.* 7, OOPSLA2 (2023), 460–485.
- [25] Gordon D. Plotkin and Matija Pretnar. 2013. Handling Algebraic Effects. *Logical Methods in Computer Science* 9, 4 (2013).
- [26] The H20 Project. [n. d.]. PicoHTTPParser. <https://github.com/h2o/picohttpparser>.
- [27] Philipp Schuster, Jonathan Immanuel Brachthäuser, and Klaus Ostermann. 2020. Compiling effect handlers in capability-passing style. *Proc. ACM Program. Lang.* 4, ICFP (2020), 93:1–93:28.
- [28] KC Sivaramakrishnan, Stephen Dolan, Leo White, Tom Kelly, Sadiq Jaffer, and Anil Madhavapeddy. 2021. Retrofitting effect handlers onto OCaml. In *PLDI*. ACM, 206–221.
- [29] OCaml Multicore team. [n. d.]. Eio. <https://github.com/ocaml-multicore/eio>.
- [30] OCaml Multicore team. 2021. Multicore OCaml HTTP benchmarks. <https://github.com/ocaml-multicore/retro-httpaf-bench>.
- [31] Tencent. [n. d.]. libco. <https://github.com/Tencent/libco>.

1030 [32] Nicolas Wu and Tom Schrijvers. 2015. Fusion for Free - Efficient Algebraic Effect Handlers. In *MPC (Lecture Notes in*
1031 *Computer Science, Vol. 9129)*. Springer, 302–322.

1032 [33] Ningning Xie and Daan Leijen. 2021. Generalized evidence passing for effect handlers: efficient compilation of effect
1033 handlers to C. *Proc. ACM Program. Lang.* 5, ICFP (2021), 1–30.

1034

1035

1036

1037

1038

1039

1040

1041

1042

1043

1044

1045

1046

1047

1048

1049

1050

1051

1052

1053

1054

1055

1056

1057

1058

1059

1060

1061

1062

1063

1064

1065

1066

1067

1068

1069

1070

1071

1072

1073

1074

1075

1076

1077

1078

1079 A SEMANTICS

1080 In this appendix we give an abstract characterisation of the variant of effect handlers that **libseff**
 1081 is based on. Following the approach of Hillerström et al. [10] we do so by way of a CEK [8] abstract
 1082 machine for a fine-grain call-by-value [20] lambda calculus. Our calculus is untyped whereas theirs
 1083 is simply-typed. Other than that, the only substantive difference between our account and that of
 1084 Hillerström et al. [10] is the treatment of effects and handlers. We return to these differences after
 1085 presenting our calculus and abstract machine. Again following Hillerström et al. [10], we diverge
 1086 somewhat from **libseff** by basing the effect handlers in this section on continuations rather than
 1087 coroutines. We make no attempt here to prevent continuations from being invoked more than once
 1088 in the abstract machine, but it would be entirely straightforward to do so.

1089 The syntax of our calculus is given by the following grammar.

1090 Values	$V, W ::= x \mid k \mid c \mid \lambda x. M \mid \mathbf{rec} f x. M \mid \langle \rangle \mid \langle V, W \rangle \mid \mathbf{inj}_\ell V$
1091 Computations	$M, N ::= V W \mid \mathbf{let} \langle x, y \rangle = V \mathbf{in} N \mid \mathbf{case} V \{ \mathbf{inj}_\ell x \mapsto M; y \mapsto N \}$ 1092 $\mid \mathbf{return} V \mid \mathbf{let} x \leftarrow M \mathbf{in} N$ 1093 $\mid \mathbf{newcont} V \mid \mathbf{resume} \mathcal{L} V W \mid \mathbf{perform} \ell V$

1094 We let V range over value terms, M range of computation terms, x range over value term variables,
 1095 k range of literals, c range of primitive operations (e.g. addition), ℓ range over individual effects,
 1096 and \mathcal{L} range over sets of effects. Being fine-grained, there are different productions for value and
 1097 computation terms. Apart from **newcont**, **resume**, and **perform** computation term constructors,
 1098 everything else is standard.

1099 The term **newcont** V converts a function value V into a continuation value. It is an idealised
 1100 analogue of `seff_coroutine_new(f, NULL)`, where V represents f . The term **resume** $\mathcal{L} V W$ resumes
 1101 continuation V with argument W handling effects \mathcal{L} . It is an idealised analogue of `seff_handle(k,`
 1102 `arg, effs)` where \mathcal{L} represents `effs`, V represents `k`, and W represents `arg`. The term **perform** ℓV
 1103 performs effect ℓ with argument V . It is an idealised analogue of `seff_perform(eff, arg)` where ℓ
 1104 represents `eff` and V represents `arg`.

1105 Before giving the transition relation for the machine we spell out the grammar for abstract
 1106 machine syntax.

1107 Configurations	$C ::= \langle M \mid \gamma \mid \kappa \rangle$
1108 Environments	$\gamma ::= \emptyset \mid \gamma[x \mapsto v]$
1109 Machine values	$v, w ::= x \mid k \mid c \mid (\gamma, \lambda x. M) \mid (\gamma, \mathbf{rec} f x. M) \mid \langle \rangle \mid \langle v, w \rangle \mid \mathbf{inj}_\ell v \mid (\kappa, \sigma)$
1110 Continuations	$\kappa ::= [] \mid (\sigma, \mathcal{L}) :: \kappa$
1111 Pure continuations	$\sigma ::= [] \mid (\gamma, x, N) :: \sigma$

1112 The configurations (C) of a CEK machine are triples: C (here ranged over by M, N) stands for
 1113 control (the program, that is, current computation term), E (here ranged over by γ) for environment
 1114 (a mapping from variables to machine values), K (here ranged over by κ) for kontinuation (what to
 1115 do next).

1116 The machine values are mostly quite standard, including corresponding forms for each basic
 1117 term value form. Indeed, we define an interpretation $\llbracket V \rrbracket \gamma$ for value term V as a machine value,
 1118 where free variables are given by the environment γ .

1119 $\llbracket x \rrbracket \gamma = \gamma(x)$	1120 $\llbracket \lambda x. M \rrbracket \gamma = (\gamma, \lambda x. M)$	1121 $\llbracket \langle \rangle \rrbracket \gamma = \langle \rangle$
1122 $\llbracket k \rrbracket \gamma = k$	1123 $\llbracket \mathbf{rec} f x. M \rrbracket \gamma = (\gamma, \mathbf{rec} f x. M)$	1124 $\llbracket \langle V, W \rangle \rrbracket \gamma = \langle \llbracket V \rrbracket \gamma, \llbracket W \rrbracket \gamma \rangle$
1125 $\llbracket c \rrbracket \gamma = c$	1126 $\llbracket \mathbf{inj}_\ell V \rrbracket \gamma = \mathbf{inj}_\ell (\llbracket V \rrbracket \gamma)$	

1127 In particular, anonymous function terms and named recursive function terms are interpreted using
 1128 closures. The final machine value form (κ, σ) is used to represent a continuation value (as returned

by **newcont** or when performing an effect). Let us defer explaining why continuation values are represented this way to the point at which we consider the transition rules.

Following Hillerström et al. [10] a continuation (κ) is a list (stack) of pairs of pure continuations σ and handlers \mathcal{L} (here actually effects sets which denote the effects handled by a handler). A pure continuation (σ) is a list (stack) of let-binding closures. (A traditional CEK machine for coarse-grained call-by-value would need many more. The advantage of fine-grain call-by-value – or ANF or SSA or CPS – is that because the result of every intermediate step must be explicitly named we know that pure computation can only proceed through another let-binding.)

Now we present the transition relation for the abstract machine.

1128	M-LAM		$\langle V \ W \mid \gamma \mid \kappa \rangle \longrightarrow \langle M \mid \gamma' [x \mapsto \llbracket W \rrbracket \gamma] \mid \kappa \rangle,$
1132			if $\llbracket V \rrbracket \gamma = (\gamma', \lambda x. M)$
1133	M-REC		$\langle V \ W \mid \gamma \mid \kappa \rangle \longrightarrow \langle M \mid \gamma' [f \mapsto (\gamma', \mathbf{rec} \ f \ x. M),$
1134			$x \mapsto \llbracket W \rrbracket \gamma] \mid \kappa \rangle,$
1135			if $\llbracket V \rrbracket \gamma = (\gamma', \mathbf{rec} \ f \ x. M)$
1136	M-CONST		$\langle V \ W \mid \gamma \mid \kappa \rangle \longrightarrow \langle \mathbf{return} \ (\ulcorner c \urcorner (\llbracket W \rrbracket \gamma)) \mid \gamma \mid \kappa \rangle,$
1137			if $\llbracket V \rrbracket \gamma = c$
1138	M-SPLIT		$\langle \mathbf{let} \ \langle x, y \rangle = V \ \mathbf{in} \ N \mid \gamma \mid \kappa \rangle \longrightarrow \langle N \mid \gamma [x \mapsto v, y \mapsto w] \mid \kappa \rangle,$
1139			if $\llbracket V \rrbracket \gamma = \langle v, w \rangle$
1140	M-CASEMATCH		$\mathbf{case} \ V \ \{\mathbf{inj}_\ell \ x \mapsto M; y \mapsto N\} \mid \gamma \mid \kappa \longrightarrow \langle M \mid \gamma [x \mapsto v] \mid \kappa \rangle,$
1141			if $\llbracket V \rrbracket \gamma = \mathbf{inj}_\ell \ v$
1142	M-CASEDEF		$\mathbf{case} \ V \ \{\mathbf{inj}_\ell \ x \mapsto M; y \mapsto N\} \mid \gamma \mid \kappa \longrightarrow \langle N \mid \gamma [y \mapsto \mathbf{inj}_{\ell'} \ v] \mid \kappa \rangle,$
1143			if $\llbracket V \rrbracket \gamma = \mathbf{inj}_{\ell'} \ v$ and $\ell \neq \ell'$
1144	M-LET		$\langle \mathbf{let} \ x \leftarrow M \ \mathbf{in} \ N \mid \gamma \mid (\sigma, \mathcal{L}) :: \kappa \rangle \longrightarrow \langle M \mid \gamma \mid ((\gamma, x, N) :: \sigma, \mathcal{L}) :: \kappa \rangle$
1145	M-RETCONT		$\langle \mathbf{return} \ V \mid \gamma \mid ((\gamma', x, N) :: \sigma, \mathcal{L}) :: \kappa \rangle \longrightarrow \langle N \mid \gamma' [x \mapsto \llbracket V \rrbracket \gamma] \mid (\sigma, \mathcal{L}) :: \kappa \rangle$
1146	M-NEWCONT		$\langle \mathbf{newcont} \ V \mid \gamma \mid \kappa \rangle \longrightarrow \langle \mathbf{return} \ x \mid \gamma [x \mapsto ([], [(\gamma, y, V \ y)])] \mid \kappa \rangle$
1147	M-RESUME		$\langle \mathbf{resume} \ \mathcal{L} \ V \ W \mid \gamma \mid \kappa \rangle \longrightarrow \langle \mathbf{return} \ W \mid \gamma \mid \kappa' \# [(\sigma', \mathcal{L}) \# \kappa],$
1148			if $\llbracket V \rrbracket \gamma = (\kappa', \sigma')$
1149	M-PERFORM		$\langle \mathbf{perform} \ \ell \ V \mid \gamma \mid \kappa \rangle \longrightarrow \langle \mathbf{return} \ (\mathbf{inj}_\ell \ \langle V, x \rangle) \mid \gamma [x \mapsto (\kappa', \sigma')] \mid \kappa'' \rangle$
1150			if κ handles ℓ at $((\kappa', \sigma'), \kappa'')$
1151	M-RETHANDLER		$\langle \mathbf{return} \ V \mid \gamma \mid ([], \mathcal{L}) :: \kappa \rangle \longrightarrow \langle \mathbf{return} \ (\mathbf{inj}_{\mathbf{ret}} \ V) \mid \gamma \mid \kappa \rangle$

The first six rules are routine. We write $\ulcorner c \urcorner$ for the function that implements c on machine values.

The M-LET rule reifies a let-binding at the head of the current pure continuation. The M-RETCONT rule binds a returned value in the body of the reified let-binding at the head of the current pure continuation.

The M-NEWCONT rule allocates a new continuation value, binding it in the environment. This continuation value simply applies the function V to its argument. The M-RESUME rule resumes a continuation value by concatenating it onto the front of the continuation component of the configuration. It is now that we see why a continuation value comprises a pair of a continuation and a pure continuation. Really (κ', σ') represents a continuation $\kappa' \# [(\sigma', X)]$ with a hole X in it that is here replaced by the effect set \mathcal{L} . The M-PERFORM rule performs an effect by reifying it as a labelled variant value containing a pair of the payload and the continuation. The auxiliary relation κ handles ℓ at $((\kappa', \sigma'), \kappa'')$ splits the current continuation κ into two parts where (κ', σ') is the continuation object up to the handler for ℓ and κ'' is the remainder of the continuation.

$$\frac{\ell \in \mathcal{L}}{(\sigma, \mathcal{L}) :: \kappa \text{ handles } \ell \text{ at } (([], \sigma), \kappa)} \qquad \frac{\ell \notin \mathcal{L} \quad \kappa \text{ handles } \ell \text{ at } ((\kappa', \sigma'), \kappa'')}{(\sigma, \mathcal{L}) :: \kappa \text{ handles } \ell \text{ at } ((\sigma, \mathcal{L}) :: \kappa', \sigma'), \kappa''}$$

1177 The M-RETHANDLER rule reifies a top-level return as a labelled variant value with a special ret
1178 label which denotes that the computation returned normally.

1179 *Comparison with standard effect handler calculi and abstract machines.* Whereas the calculus
1180 of Hillerström et al. [10] includes both deep and shallow handlers ours provides hybrid sheep
1181 handlers [24]. A deep handler automatically wraps the original handler around the body of each
1182 suspended continuation. A shallow handler does not. A sheep handler does not automatically
1183 wrap the original handler around the body of each continuation, but does require a handler to be
1184 explicitly installed whenever the continuation is resumed. Sheep handlers guarantee that some
1185 handler must be installed whenever a continuation is resumed, but not necessarily the original one.

1186 The other substantive difference between our calculus and more classical ones like that of
1187 Hillerström et al. [10] is that although **resume** specifies the effect set for a handler, there is no
1188 special construct for specifying a handler by dispatching on the effect. Instead the result of **resume**
1189 (either a normal return or a performed effect) is wrapped up in a variant value and the dispatch is
1190 implemented using **case**.
1191

1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225