

for the synchronous, blocking echo server. In similar vein, we have implemented an effect-based asynchronous I/O library, `aeio` [1], that exposes a direct-style API to the clients. Experimental results (Fig. 1) of using the library for as a backend for `httpaf`, a full-featured OCaml web server demonstrates that effect handlers perform on par with highly optimised monadic concurrency libraries, while retaining the simplicity of direct-style code.

3 Asynchronous effects

Effect handlers thus provide a nice abstraction for expressing user-level cooperative fiber schedulers. However, if a fiber runs a long running pure computation, it leads to an undesirable situation where the fiber hogs the core and does not let other fibers run until the pure computation is complete. The standard solution is to preempt the fiber periodically by installing an interval timer. However, with a user-level scheduler, it is unclear how the signal handler for the timer interrupt (which itself is an asynchronous computation) could get hold of the continuation of the main computation. One could envision supplying the signal handler function with the current continuation, but this makes the API more awkward [6].

We make the observation that various asynchronous *up-calls* from the OS such as timer expiry, interrupts, signals, I/O completion notifications can be treated as asynchronous effects. Rather than handling the timer interrupt in a signal handler, the runtime raises an effect `TimerTick` in the current fiber, which can be handled like a synchronous `Yield` effect to implement preemptive scheduling:

```
| effect TimerTick k →
  (* add current fiber to scheduler *)
  enqueue (fun () → continue k ());
  run_next()
```

If the `TimerTick` is unhandled, it is a noop.

3.1 Signal handling

Signals are also modeled as asynchronous effects. On Unix-like systems, when the user of a command-line program presses `Ctrl-z`, the `SIGTSTP` signal is sent to the running program. By default, this suspends the program. However, we might want a different behaviour for `Ctrl-z` where we preempt the running fiber if we are in the scheduler, and suspend the program otherwise. We can achieve this by handling the asynchronous `Suspend` event in the scheduler to handle the `SIGTSTP` signal:

```
match f () with
| v → ...
| ...
| effect Suspend k → (* handle SIGTSTP *)
  enqueue (fun () → continue k ());
  run_next()
```

This implementation, however, has a subtle bug; if `Ctrl-z` was received while the control is in the body of clause for the `Suspend` effect, then the program is suspended. This is not the intended behaviour as we are in the middle of preempting a thread, and we do not expect to be suspended. In order to eliminate this possibility we require a means for temporarily disabling asynchronous exceptions. We follow the good advice of Marlow et al. in the design of GHC Haskell’s asynchronous exceptions [4], and prefer *scoped masking combinators*:

```
mask (fun () →
  match unmask f with
  | v → ...
  | ...
```

```
| effect Suspend k →
  enqueue (fun () → continue k ());
  run_next()
```

The changes to the masking state made by `mask` and `unmask` apply only to one scope, and are automatically updated when entering and leaving the scope. Scoped combinators are also exception safe: if `run_next` raises an exception which escapes the handler scope, then the masking state is automatically reset to that of the parent scope.

3.2 Asynchronous I/O notifications

Operating systems provide a number of different interfaces with which to perform I/O. The simplest is the direct-style *blocking I/O*, in which the program calls I/O functions provided by the operating system, which do not return until the operation completes. This allows a straightforward style of programming in which the sequence of I/O operations matches the flow of the code. We aim to preserve this style of programming, but implement it using alternative operating system interfaces that allow multiple I/O operations to be overlapped.

Earlier we saw one way of accomplishing this with effects, by using operating system multiplexing mechanisms like `select`, `poll`, etc., which block until one of several file descriptors is ready. An alternative interface is *asynchronous I/O*, in which multiple operations are submitted to the operating system, which overlaps their execution. However, applications written using asynchronous I/O tend to have complex control flow which does not clearly explain the logic being implemented, due to the complexity of handling the operating system’s asynchronous notifications of I/O completion.

We propose effects and handlers as a means of writing direct-style I/O code, but using the asynchronous operating system interfaces. We introduce two new effect operations: `Delayed`, which describes an operation that has begun and will complete later, and `Completed`, which describes its eventual completion. Both of these take an integer parameter, which is an ID number identifying the particular operation.

Potentially long-running operations like `read` perform the `Delayed` effect, indicating that the operation has been submitted to the operating system but has not yet completed. Later, upon receipt of an operating-system completion notification, the asynchronous effect `Completed` is performed.

Using this mechanism, support for asynchronous completions can be added to the scheduler by adding clauses for the `Delayed` and `Completed` effects, where `ongoing_io` is an initially empty hash table:

```
| effect (Delayed id) k →
  Hashtbl.add ongoing_io id k
| effect (Completed id) k →
  let k' = Hashtbl.find ongoing_io id in
  Hashtbl.remove ongoing_io id;
  enqueue (fun () → continue k ());
  continue k' ()
```

In this sample, the continuation `k` of the `Delayed` effect is the continuation of the code performing the I/O operation, which instead of being immediately invoked is stored in a hash table until it can be invoked without blocking.

The continuation `k` of the `Completed` effect is the continuation of whichever fiber was running when the I/O completed. This scheduler chooses to preempt that fiber in favour of the fiber that performed the I/O, by retrieving the continuation `k'` from the hashtable and continuing it. Equally, the scheduler’s policy could be to give priority to the already running fiber, by swapping `k` and `k'` in the last two lines.

References

- [1] Aeio: An asynchronous, effect-based I/O library, 2017. Accessed: 2017-05-05 09:21:00.
- [2] S. Dolan, L. White, and A. Madhavapeddy. Multicore OCaml. OCaml Workshop, 2014.
- [3] S. Dolan, L. White, K. Sivaramakrishnan, J. Yallop, and A. Madhavapeddy. Effective concurrency through algebraic effects. OCaml Workshop, 2015.
- [4] S. Marlow, S. P. Jones, A. Moran, and J. Reppy. Asynchronous exceptions in Haskell. In *Proceedings of the ACM SIGPLAN 2001 Conference on Programming Language Design and Implementation*, PLDI '01, pages 274–285, New York, NY, USA, 2001. ACM.
- [5] Y. Minsky, A. Madhavapeddy, and J. Hickey. *Real World OCaml - Functional Programming for the Masses*. O'Reilly, 2013.
- [6] Signal Handling in MLton Threads. <http://www.mlton.org/MLtonSignal>, 2016. accessed: 1-Jun-2017.
- [7] S. Peyton Jones. *Tackling the awkward squad: monadic input/output, concurrency, exceptions, and foreign-language calls in Haskell*, pages 47–96. IOS Press, January 2001.
- [8] J. Vouillon. Lwt: A cooperative thread library. In *Proceedings of the 2008 ACM SIGPLAN Workshop on ML*, ML '08, pages 3–12, New York, NY, USA, 2008. ACM.