

Developing Concurrent Web Applications with Effect Handlers in Links

Steven Chang



4th Year Project Report
Electronics and Computer Science
School of Informatics
University of Edinburgh

2023

Abstract

This project focuses on the study of effect handlers in Links, a functional programming language designed for developing web applications from the database to the client. Effect handlers are a programming paradigm that provide abstraction for implementing a variety of programming features, such as concurrency, within a programming language. The main objective of this project is to investigate the potential of modular development in concurrent programming with effect handlers in Links. Specifically, the project aims to develop an effect handler switching mechanism that enables users to define and switch between different effect handlers within the same application with minimal changes to the code.

To achieve this objective, I used a web application that visualizes user-level threads as lines on a canvas as a base application to demonstrate the behavior of the user-level threads scheduler implemented by effect handlers. Then, I modified the base application to enhance user experience and introduce new functionalities for concurrent programming. Furthermore, I modularized the codebase to improve maintainability and readability. Finally, I implemented schedulers with different scheduling mechanisms to enable seamless switching between schedulers.

The results show that the final application has a significantly reduced codebase, and the effect handlers switching mechanism works seamlessly among different applications with minimal code changes. However, the limitations of the project include the mismatch between effect handlers, event listeners for web component interaction, and actors in the actor model in concurrent computation. The project concludes that the objectives have been successfully achieved, and it provides valuable insights for future work in this field.

Research Ethics Approval

This project was planned in accordance with the Informatics Research Ethics policy. It did not involve any aspects that required approval from the Informatics Research Ethics committee.

Declaration

I declare that this thesis was composed by myself, that the work contained herein is my own except where explicitly stated otherwise in the text, and that this work has not been submitted for any other degree or professional qualification except as specified.

(Steven Chang)

Acknowledgements

I would like to express my sincere gratitude to my dissertation supervisor, Sam Lindley, for his constant support and invaluable feedback throughout my undergraduate dissertation project. His expertise, insights, and constructive feedback have been instrumental in shaping my work and helping me overcome technical challenges. Without his continuous support, this project would not have been possible.

Additionally, I would like to express my appreciation to Daniel Hillerström for his assistance and technical support in understanding Links and effect handlers, which greatly aided me in my dissertation work.

Table of Contents

1	Introduction	1
1.1	Motivations	1
1.2	Objectives	2
1.3	Contributions	3
1.4	Report Structure	3
2	Background	4
2.1	Modern Web Development	4
2.2	Links	5
2.2.1	XML and DOM	5
2.2.2	Event Listeners	5
2.2.3	Foreign Function Interface	6
2.3	Effect Handlers	6
2.3.1	Programming with Effect Handlers in Links	7
2.3.2	Code Example in Links	8
2.4	Concurrent Programming with Effect Handlers	10
2.5	Related work	11
2.5.1	React and React Fiber	11
2.5.2	racing_lines.links	11
3	Base Application	13
3.1	Overview	13
3.2	Three Priority Level	14
3.3	User Interface Implementation	15
3.4	Editing Priority of the Fiber	18
4	Modularization	19
4.1	JavaScript FFIs	19
4.2	Separation of Concerns	21
4.3	Extract Effect Interface	22
5	Schedulers and Case Studies	23
5.1	Time-based Scheduler	24
5.2	Step-based Scheduler	25
5.3	Probability-based Scheduler	27
5.4	Switch Schedulers	29

5.5	Case Studies Implementation	31
5.5.1	Case Study 1	31
5.5.2	Case Study 2	32
5.5.3	Case Study 3	33
6	Evaluation	34
6.1	Case Studies	34
6.1.1	Case Study 1	34
6.1.2	Case Study 2	35
6.1.3	Case Study 3	35
6.2	Triumph	36
6.3	Challenges: Mismatch Between Event Listeners and Effect Handlers .	36
7	Conclusion and Future Work	39
7.1	Future Work	39
	Bibliography	41
A	Complete code for racing lines demo	44
A.1	Case Study 1	44
A.2	Case Study 2	48
A.3	Case Study 3	52
A.4	fiberInterface.links	56
A.5	lineDrawer.links	57
A.6	runtime.js	58
A.7	Time-based Scheduler	59
A.8	Step-based Scheduler	64
A.9	Probability-based Scheduler	67

Chapter 1

Introduction

In this project, I explored the use of effect handlers in web development using the Links programming language. The primary objective is to enhance the modularity and accessibility of concurrent programming with effect handlers by allowing users to easily define and switch their own effect handlers within the client code. I achieved this by implementing various user-level thread schedulers using Links and effect handlers, with a focus on enabling users to switch between different effect handlers with minimal changes to the client code. This project successfully demonstrated a proof of concept that encapsulates effect handlers as independent modules, offering a more efficient and flexible approach for concurrent programming in the context of web development using effect handlers. The code related to this project is available at https://github.com/Steven-Chang1114/effect_handler_links.

1.1 Motivations

Links [22] is a functional programming language that aims to simplify web programming by providing a unified language for all tiers of a web application, which includes the client-side, server-side, and database. To achieve this, Links translates code into the appropriate languages for each tier, such as JavaScript for the browser, Java for the server, and SQL for the database. Additionally, Links supports features that facilitate web development, including manipulation of the Document Object Model (DOM) and event listeners [5]. Furthermore, Links allows for concurrency on both the client and server-side [5], which enables users to create web applications that can handle multiple tasks simultaneously.

In contrast to conventional web development technologies, Links offers effect handlers [12] as an abstraction for managing effects such as exceptions, input/output operations, and state changes in a modular and composable way. An effect is defined as an operation that can be performed, while a handler is a function that specifies how to interpret those operations [12]. By separating the definition of effects from their implementation, effect handlers provide a modular and composable way of handling operations in a program, resulting in more flexible and maintainable code. One of the practical applications of effect handlers is their ability to express concurrency in programming languages.

Previous research on the use of effect handlers for concurrent programming has been mostly focused on the system level, despite the growing popularity of effect handlers in recent years. For instance, Multicore OCaml has provided a concurrency library that is implemented with effect handlers [6]. Links, on the other hand, is the first programming language that can build client-side web applications incorporating concurrency through the use of effect handlers.

Meanwhile, the field of web development has experienced a rapid growth over the past three decades, evolving from plain HTML, CSS, and JavaScript in 1995, to more powerful libraries like jQuery [30] in 2006, and then to modern front-end frameworks such as React [28] in 2013, followed by Vue [32] a year later. Nowadays, these frameworks are exploring new possibilities, with Vue introducing its concurrency library [31] and React exploring functional programming paradigms. Specifically, React developers have been inspired by algebraic effects and have made progress integrating similar features into the React framework, such as React Hooks [27]. However, these features are more of a mimic of the effect handlers structure, as JavaScript does not yet fully support this feature.

Hence, compared with the other functional programming languages and web development technologies mentioned above, Links has its natural advantage as a functional programming language that is designed for building web applications and already supports the implementation of the effect handlers. As a result, this project aims to explore the possibility of effect handlers using Links in the context of web development. The primary focus of this project is to explore concurrent programming with effect handlers, which is becoming an important aspect of modern web development with all popular frameworks aiming to integrate concurrent mode into their production workflow.

1.2 Objectives

Daniel Hillerström and Nicole Meng demonstrated that effect handlers in Links can be used for concurrent programming by implementing user-level threads and their schedulers in Links, as presented in their previous work called `line.links` [11]. However, there is still room for improvement by implementing modular development for concurrent programming with effect handlers in Links. This project aimed to enhance this aspect by allowing users to easily and conveniently define and switch their own handlers within the client code.

To achieve this, the primary objective of this project was to develop various user-level thread schedulers using Links and effect handlers, each having different scheduling mechanisms. Afterwards, these effect handlers were integrated into the main application, allowing users to easily switch between them with minimal changes to the client code. Furthermore, several applications with distinct behaviors were implemented as case studies to test and demonstrate the versatility of this effect handler switching mechanism.

1.3 Contributions

In this project, I adapted the existing codebase of `racing_line.links` [11] and accomplished the following objectives:

- Enhanced the user experience of `racing_line.links`
- Expanded the scheduler to support three different priorities of the user-level threads, and the ability to modify priority levels
- Performed the separation of concerns to make the application code more readable and maintainable
- Extracted the effect interface into separate module
- Designed and implemented three different types of schedulers
- Modified the client to allow seamless adding, removing, and switching between the effect handlers
- Designed and implemented three different applications for visualizing user-level threads as case studies to evaluate the switching mechanism

With these contributions, I have successfully demonstrated a proof-of-concept demo that allows users to plug in different effect handlers into the client code with minimal effort. This highlights the feasibility of encapsulating effect handlers as independent modules and the ability to plug in different effect handlers for the same computational context without the need to rewrite the entire client code. The new structure offers a more efficient and flexible approach by keeping the client code generic when changing effect handlers for different behaviors.

1.4 Report Structure

The remainder of the report is structured as follows:

Chapter 2 provides an introduction to modern web development techniques and an overview of the Links programming language, including its syntax and relevant features. It also covers a discussion of effect handlers and concurrent programming using effect handlers, along with related work to provide context for the project.

Chapter 3 provides a detailed discussion of the modifications made to the base application to enhance user experience and introduce new functionalities.

Chapter 4 presents the steps taken to modularize the codebase to improve its overall quality, readability, maintainability, and modularity.

Chapter 5 presents the design and implementation of schedulers with different scheduling mechanisms and introduces different case studies to evaluate the schedulers.

Chapter 6 evaluates the final work and discusses the challenges encountered during the development process.

Chapter 7 concludes the work and suggests possible future improvements.

Chapter 2

Background

2.1 Modern Web Development

The emergence of modern web development technologies marked a significant change in the evolution of the World Wide Web. With the advent of HTML, CSS, and JavaScript, these three technologies formed the backbone of modern web development. Hypertext Markup Language (HTML) [15] provided the basic structure and content of web pages, while Cascading Style Sheets (CSS) [33] is used to add styling to web pages, such as font, colors, and layout. JavaScript [16], on the other hand, enables the creation of interactive and responsive web applications. The integration of these three technologies has paved the way for the creation of functional and engaging websites that could be accessed from any device, anywhere in the world.

One of the very important and fundamental concepts in web development is the Document Object Model (DOM) [7]. It is an essential feature that allows developers to interact with HTML or XML documents dynamically. The DOM is a tree-like structure where each node represents an element in the document. With the DOM, developers can dynamically change the style, content, and structure of a web page in response to user interactions or other events.

Another fundamental concept in web development that is used closely with DOM is event listeners [34] which allow developers to write code that can respond to user events on a web page. An event is simply an action that happens in the browser, such as clicking, scrolling, or hovering over an element. By accessing these elements through the DOM and adding event listeners, developers can detect these events and trigger specific behaviors in response. This functionality allows web applications to respond to user input in real-time, enhancing the overall user experience.

In contrast to HTML, Links uses XML to create web pages and has its own set of DOM operations for XML documents and event listeners using l-event attributes [5]. Further details about the key features of Links will be explained in the next section.

2.2 Links

In traditional web development, developers are required to be proficient in a variety of programming languages, including HTML, JavaScript, Java, Python, and SQL, to build full-stack applications. For beginners, this can be particularly challenging, and the process of linking these languages together can be overwhelming. Moreover, as developers attempt to integrate the client, server, and database, they often encounter the impedance mismatch problem [5], which arises when data must be converted to the corresponding acceptable data types between each tier. This can be a significant limitation for developers seeking to build modern, efficient web applications.

Links was developed to address these issues. As a strict, typed functional programming language, it aims to overcome the challenges of mastering multiple programming languages and the impedance mismatch problem in web development. It provides a single source code for all tiers of a web application, including client-side, server-side, and database. However, due to its research-oriented nature, there are limited resources available for learning Links, and most of the available resources are in the form of example codes and Links Wiki page [23] on GitHub. The following sections aim to introduce some of the essential features in Links to enhance the understanding of this project.

2.2.1 XML and DOM

In contrast to HTML, Links uses XML notations to construct web pages [5]. The syntax for the XML notations is similar to HTML, with the tag name enclosed in `<#>...</#>` syntax. The XML document is maintained by the client as a tree data structure, which is comparable to the DOM. Links offers two types of operations: DOM, which is mutable, and XML, which is for inspection only [5]. It also provides a set of conversion functions to convert custom data types to XML and insert them into the selected nodes. The full operations and functions reference can be accessed through the Links documentation [20].

2.2.2 Event Listeners

In order to enable interaction with the DOM, Links has provided syntax for defining event listeners in the XML code using attributes known as `l-event` attributes [20]. The format of these attributes is `l:name`, with `l:` serving as the prefix and `name` representing the event name. These `l-event` attributes must be assigned with a function that will be executed when the event is triggered [5]. For example,

```
<button l:onclick="onClick()">Click me</button>
```

This indicates the `onClick()` function is invoked when the `button` element is clicked. The complete list of event listeners available in Links can be found in the Links documentation [20].

2.2.3 Foreign Function Interface

In the context of web development, Links has the ability to call JavaScript functions directly within the Links program through the use of Foreign Function Interface (FFI) [21]. The FFI is made possible because modern web browsers run on JavaScript engines, and Links code on the client side is ultimately compiled into JavaScript code during runtime. By providing this feature, it offers extensive flexibility during the development process, enabling smooth collaboration between Links and JavaScript features, opening up new opportunities for developers.

A quick walkthrough on how to use JavaScript FFI in Links can begin by defining a function in a JavaScript file. The following code illustrates an example of how this can be done:

```
// js/log.js
function _logMessage(msg) {
    return console.log(msg);
}

var logMessage = LINKS.kify(_logMessage);
```

In the above example, `_logMessage(msg)` is a standard JavaScript function that logs the value of the `msg` variable in the console of the browser when it is called. This function is then wrapped inside a conversion function called `LINKS.kify` [21], which produces another function called `logMessage` that can be accessed in the Links program.

On the Links side, the JavaScript functions can be imported as a module (`Log`) for the entire program to access. To achieve this, the `alien` keyword is used in the module block that specifies the external language (`javascript`) and the file path of the JavaScript function. The following code block shows an example:

```
// log.links
module Log {
    alien javascript "js/log.js" {
        logMessage : (String) ~%~> ();
    }
}
```

Inside the `alien` block, the name and the types of the foreign function should be provided. Then the `logMessage` function can be assessed within the Links program by calling `Log.logMessage("Hello World")`.

2.3 Effect Handlers

Algebraic effects handlers are a feature in functional programming where the concept of algebraic effects was first introduced by Plotkin and Power [24], and later, algebraic effect handlers were introduced by Power and Pretnar [25]. The purpose of the algebraic effects handlers is to structure effectful programs in a modular and compositional way. This is accomplished by separating the effects into operations for expressions and

handlers for implementations (discussed in Section 2.3.1.1 and 2.3.1.2) using delimited continuations [9]. This approach enables the programs to pause, resume, and switch between different computation contexts, allowing them to express complex control-flow operations such as I/O, exception, state management, and concurrency. For consistency, the terms *effect* and *effect handlers* will be used instead of *algebraic effect* and *algebraic effect handlers*.

The concept of effect handlers was mostly focused on its theoretical aspects in the early stages of the development. However, in recent years, there has been a growing interest in the applicability of effect handlers. As a result, a variety of implementations of effect handlers have emerged, ranging from programming languages to libraries [3, 6, 13, 18].

In the effect handlers implementation, an effect is conceptualized as a signature of operations [19]. Operations are defined by the developers, and they can be thought of as abstract interfaces that describe the desired effect such as `Get` or `Set`. The concrete implementation of the operations, on the other hand, is defined by the effect handlers. Effect handlers provide a way to interpret these operations by defining a custom function that implements the behavior for the effectful operation. The following section will delve into the implementation details of effect handlers in Links.

2.3.1 Programming with Effect Handlers in Links

In Links, a handler function is defined with one parameter: the program's computation context. The handler function then defines the operation definition for each effectful operation as well as its concrete implementation. The operation definition accepts two kinds of arguments: the payload passed into the function, and the captured continuation. The upcoming sections will provide a more detailed explanation of operations and handlers in Links.

2.3.1.1 Operation

Operation is like a function that will produce an output after being performed [12]. The operation itself does not have any semantical meaning towards the program, and the actual implementation of the operation is defined entirely by the handlers. By convention, when defining an operation, the name should start with a capital letter and should not be left unhandled in the program or otherwise an error will be raised. In order to perform an operation, the syntax is `do Operation(arg1, arg2, ..., argn)`.

2.3.1.2 Handlers

The syntax for implementing handlers in Links is similar to the switch statement in imperative programming or pattern matching in functional programming:

```
handle (m) {
  case <Op1 => k> -> # Implementation for Op1
  case <Op2(p1, p2, ..., pn) => k> -> # Implementation for Op2
  case <Op3(p1, p2) => k> -> # Implementation for Op3
  ...
}
```

```

    case x -> x
}

```

The `handle` function takes in the computation context `m` as input. For every operation that is performed within the computation `m`, it will be mapped to the corresponding operation definition inside the `case` block for execution. The `case` block is composed of two elements: the operation definition, and its concrete implication. The operation definition can have any number of parameters, followed by a captured continuation `k` at the end, using the `=>` syntax. `k` is represented as a function that can take one or more arguments when invoking. By calling `k`, the control flow of the program is transferred back to the point in `m` where the operation was performed [12]. Finally, the `case x -> x` statement is an implicit return statement that will be executed after the computation `m` has completed its execution.

2.3.2 Code Example in Links

A straightforward way to explain the implementation of effect handlers in Links is through a series of examples. We begin with a program that prints two strings “Hello” and “World” in sequence. Without effect handlers, it is straightforward to implement in Links:

```

fun printMessage() {
    print("Hello\n");
    print("World\n");
}

```

Which yields

```

Hello
World
() : ()

```

When you call `printMessage()`

2.3.2.1 Printing with Effect Handlers

To integrate effect handlers into this example, the `print` method can be transformed into an operation named `Print`. Then, a handler called `forward` can be defined to handle the `printMessage` function as the computation context and interpret the `Print` operation inside the `printMessage` function.

```

fun printMessage() {
    do Print("Hello\n");
    do Print("World\n")
}

fun forward(m) {
    handle(m()) {
        case <Print(val) => k> -> print(val); k(())
    }
}

```

```

        case x -> ()
    }
}

```

To execute the function, we can call `forward(printMessage)` and it will yield the same result as the previous example:

```

Hello
World
() : ()

```

It first replaces `print` method into `Print` operation inside `printMessage` function, and then when executing the program, `printMessage` is passed to the `forward` handler as the computation `m`. Therefore, by invoking `printMessage` inside the handler, `Print` operation will be performed, then captured and handled by the `forward` handler.

The concrete implementation for each operation is defined in the case patterns inside `forward` handler. In this example, after performing `Print` operation, it will first be mapped to the `Print` case in the handler where it receives two arguments, `val` as the actual parameter being passed into `Print` operation, and `k` which is the continuation function. In the `Print` case, `print` method is invoked before the continuation function `k`, which means the execution will be in order. Therefore in the program "Hello" is printed before "World".

2.3.2.2 Continuation

It might seem unnecessary to convert two lines of code into more complicated structures with effect handlers. However, a powerful feature in effect handlers is that the developers can manipulate the functions' order of executions in a program by making use of the continuation `k`.

Now on top of the previous example, if we want to print the same strings but show them in the reverse order without changing the order of invocation in `printMessage`, the continuation in effect handlers can be handy when dealing with this issue:

```

fun reverse(m) {
    handle(m()) {
        case <Print(val) => k> -> k(); print(val)
        case x -> ()
    }
}

```

We defined a new handler called `reverse`, and then when interpreting `Print` operation inside the `reverse` handler, the continuation `k` is called first, which means `print(val)` will only be invoked after `k` is finished executing. By wrapping the program continuation into a function variable enables the developer to easily take control over the call stack of the program. Now if we call `reverse(printMessage)`, the result will be:

```

World
Hello

```

() : ()

As demonstrated, by simply switching different effect handlers to interpret the same operation, the order of printing can be controlled without modifying the `printMessage` function.

2.4 Concurrent Programming with Effect Handlers

While the section above only provides simple examples of effect handlers, it demonstrates some of their powerful features. Modularity, particularly, is highlighted as one of the strengths of effect handlers because they provide a way to separate the implementation from the expression of the effectful computation. Additionally, with the use of delimited continuation, which enables the pausing, resuming, and switching between computational contexts, effect handlers are ideal for implementing a user-level thread scheduler that can achieve concurrency.

Effect handlers also offer the advantage of writing concurrent programs in direct style, making them more readable and easier to debug. Compared with continuous-passing styles (CPS) which is callback-oriented, using direct-style code can usually reach better performance, as there are fewer function calls and less overhead associated with managing continuations [2]. In Multicore Ocaml, relevant work had already been deployed by implementing asynchronous effects and their handlers to construct an asynchronous I/O library using direct style programming [6].

In Links, concurrent programming is supported by using message passing, known as actor model [10]. In the actor model, the basic building blocks of the concurrent computation are “actors”, which can be thought of as lightweight agents or objects that can react to incoming messages and perform actions based on their internal state. Actors communicate with each other through asynchronous message-passing, which means that they send and receive messages without waiting for a response. This enables concurrency and prevents potential deadlocks and race conditions that can occur in traditional shared-memory concurrency models. Finally, Links programs are compiled into JavaScript using CPS translation [14] in order to preserve concurrency on the client side.

Given the fact that Links supports effect handlers, it presents an alternative to actor models for achieving concurrency in Links. Specifically, effect handlers in Links enable us to implement a user-level thread scheduler that can facilitate concurrent programming. The following section will first demonstrate a state-of-art implementation deployed in the field of web development that enables concurrent operations, followed by a web application developed in Links that illustrates the practicality of implementing concurrency using effect handlers in Links.

2.5 Related work

2.5.1 React and React Fiber

React is a popular JavaScript library for building user interfaces [28]. One of the key features of React is the creation of a virtual DOM [29], which allows it to efficiently update the browser DOM only where necessary, resulting in faster and more efficient rendering of components.

React Fiber [17] is a new implementation of React’s core algorithm, which aims to make React faster, more efficient, and more intelligent. It is a complete rewrite of the older versions of React’s reconciliation algorithm and is now the default reconciler for React 16 and newer versions. The name “Fiber” comes from the use of fibers to represent nodes in the DOM tree.

One of the primary features of React Fiber is its ability to pause, resume, and prioritize updates during the reconciliation process [8]. This allows React to interrupt an update in progress to handle a higher-priority update, and then resume the original update when the higher-priority update is complete. This makes React more responsive and efficient, especially in cases where there are a large number of updates or complex animations.

React Fiber’s ability to prioritize updates is also critical for applications with complex UIs. By giving priority to updates that require immediate attention, React Fiber can ensure that the user interface remains responsive even when there are multiple updates happening simultaneously. Overall, React Fiber represents a significant improvement to the React library, making it more efficient, more responsive, and more capable of handling complex and dynamic user interfaces.

2.5.2 racing_lines.links

As with the React Fiber example discussed earlier, there is a similar implementation in Links built by Daniel Hillerström and Nicole Meng in the official Links GitHub Repository called `racing_lines.links` [11] that mimics the interruption and resumption behaviors of user-level threads with different priorities. This program visualizes the user-level threads, also known as fibers, as lines that are rendered on the canvas. Throughout this report, the term “fiber” will be used to refer to these lightweight user-level threads for consistency.

The user interface of `racing_lines.links` contains a rectangular canvas, two dropdown menus, and one button. The rectangular canvas at the top is where the fibers are displayed. The dropdown menu offers 6 colors and 2 priority levels, namely, “Low” and “High”. Users can then select their desired line color and priority level for the fiber from the dropdown menu. Once they have made their selections, they can render the chosen fiber as a horizontal line on the canvas by clicking the “Draw Line” button.

In Figure 2.1 and 2.2, the process of creating four fibers is shown. First, the green fiber with low priority is created, followed by the red fiber with high priority, then the blue fiber with low priority, and finally the yellow fiber with high priority. It can be observed that, although the green fiber is the first to be created, it is interrupted



Figure 2.1: Rendering process of the High-priority fibers




Figure 2.2: Rendering process of the Low-priority fibers

during its rendering process by the red fiber because the red fiber has a higher priority. Similarly, the rendering process of the blue fiber is also blocked due to the fact that it has lower priority than the red fiber. On the other hand, since the yellow fiber has the same priority as the red fiber, both fibers will be rendered alternatively. The green and blue fibers can only resume their rendering process after the red and yellow fibers finish rendering and give control back to the lower-priority fibers. Then, they will resume their rendering process and render alternatively until they complete their rendering process.

The `racing_lines.links` program is implemented using effect handlers to create a fiber scheduler. The scheduler uses two operations, `Fork` and `Yield`, to create fiber instances and control their interruption and resumption. The fibers are maintained in a scheduler queue and will be popped out from the queue once they finish their execution.

Chapter 3

Base Application

This section builds upon the foundation laid out in Section 2.5.2, which serves as the base application for our project. Here, I present the specific details of the design and implementation of modifications that were made to the base application with the aim of improving user experience and introducing new functionalities.

3.1 Overview

The user interface design of the application has been revised to enhance the user experience that facilitates the testing and evaluation process. To simplify the process of creating fibers with different priorities, the color and priorities dropdown menu has been merged into one. Instead of allowing users to select from six different colors, each priority level can be assigned a specific color, such as red for Low priority and green for High priority. This change allows users to create fibers as lines on the canvas more efficiently with just one click instead of three. Additionally, each priority level has been assigned a specific color, making it easier for users to identify the state of the fibers. These improvements enable users to focus on the main task of the application, which is to create fibers and observe their behaviors when encountering fibers with different priorities that may preempt the control flow. A detailed description of the user interface implementation will be provided in Section 3.3.

To enhance the user experience of the application and show the strength of Links in concurrent programming, the application can be further improved by providing more priority levels for fibers. In addition to Low and High, a medium-priority Fiber can be introduced, represented by the color blue. This new priority level would allow the fiber to preempt the control of a low-priority fiber but would need to give up the control to high-priority fibers. This expansion of priority levels offers more possibilities for visualizing fiber behavior and allows users to explore a wider range of scenarios. More importantly, by introducing three priority levels, Links can demonstrate its superiority in concurrent programming compared to other modern front-end frameworks. For instance, React's Concurrent Mode [26] only supports two levels of priority due to the single-threaded nature of JavaScript. The specific implementation of the user interface with three priority levels will be discussed in Section 3.2.

In order to expand the capabilities of Links in concurrent programming, it is necessary to increase the level of control over the fiber scheduler. The goal is to have better management over the fibers, such as the ability to retrieve and modify the priority level of fibers. To achieve this objective, two new operations, namely `GetPrio` and `SetPrio`, can be defined and added to the effect handler. Currently, the effect handler only provides the functionality to create a new fiber and yield control from one fiber to another. These new operations provide developers with more options to manipulate the fibers. Further details regarding the implementation of `SetPrio` and `GetPrio` will be presented in Section 3.4, highlighting the potential benefits of Links in concurrent programming.

3.2 Three Priority Level

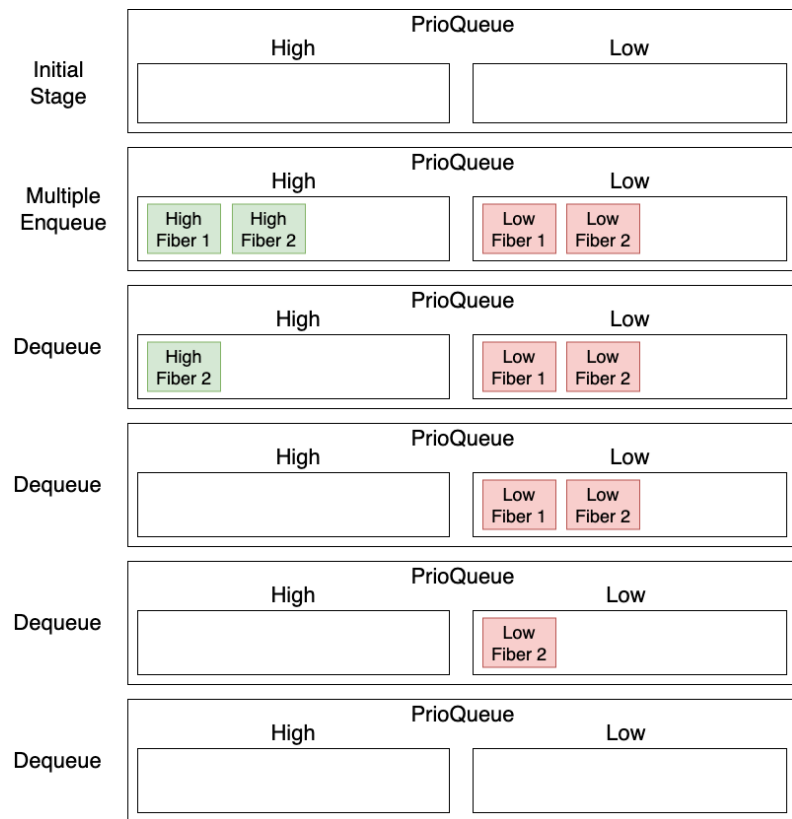


Figure 3.1: PrioQueue behavior for fibers with different priorities

The addition of one more priority level is being implemented first as it is the core functionality of the program. Before discussing the technical details of how to add the additional priority level, a brief overview of how fibers are created, stored, and scheduled by the schedulers is necessary. Each fiber has two states, `prio` and `f`, representing its priority level and the function that will be executed when the fiber is scheduled. In the two priority levels setting, fibers are stored in separate priority queues for Low and High priority. These two priority queues are combined into one data structure called `PrioQueue` for better management. During the enqueue stage, fibers are enqueued to their respective priority queues based on their `prio` value. During the dequeue stage,

high-priority fibers are dequeued first, followed by low-priority fibers only if the High Priority Queue is empty. This ensures that high-priority fibers can always run before low-priority fibers when performing the `Yield` operation. A diagram illustrating how the `PrioQueue` work is presented in Figure 3.1.

The fiber scheduler, which is responsible for managing all the fibers, is implemented using effect handlers and interprets two operations: `Fork` and `Yield`. While Section 5 provides a detailed explanation of the fiber scheduler, this section focuses on priority processing. In the `Yield` operation definition, its implementation processes the rendering operations on the canvas and yields control to other fibers depending on the scheduling mechanism. In the `Fork` operation definition, it determines whether the newly forked fiber should interrupt the current running fiber or be stored in the `PrioQueue` for later use using a switch statement. To add another priority level, it is necessary to modify the `PrioQueue` data structure to accommodate the new priority level, adjust the enqueue and dequeue logic accordingly, and change the implementation of the `Fork` and `Yield` operations.

The modifications made include defining a new queue for fibers with medium priority to the `PrioQueue` data structure, as well as modifying the `priorityEnqueue` and `priorityDequeue` functions to accommodate this change. The structure of the new `PrioQueue` has been illustrated in Figure 3.2.




Figure 3.2: new `PrioQueue` structure

Furthermore, in order to integrate the new `PrioQueue` structure to the fiber scheduler, the `Fork` operation definition was also modified to incorporate the relationship between low, medium, and high-priority fibers, allowing newly created fibers to preempt lower-priority fibers and enqueued into the `PrioQueue` if the current running fiber has a higher priority. The `Yield` operation was also updated to include the rendering process for fibers with medium priority. A more detailed explanation of the fiber schedulers can be found in Section 5.

3.3 User Interface Implementation

Now that the implementation of the three priority levels is complete, the focus can now shift towards improving the user interface of the application. As shown in Figure 3.3, the current process of creating a new fiber requires users to interact with two separate dropdown menus to select the color and priority of the fiber, followed by clicking the "Draw Line" button, which involves three clicks. This process can be further simplified by reducing the number of clicks required to create a fiber to one, providing users with greater convenience and more straightforward control over the creation of fibers.

In the current implementation of the application, the `buttonPressed` function is linked to the `button` tag in the code using `l-event` attributes. This function will be executed when a click event on the button is detected. It retrieves the values for priority and color using the `getValueFromSelection` function, and then it uses a `switch` statement to determine the priority level of the fiber, with the appropriate action taken based on the priority level. The code for the `buttonPressed` function is shown below:

```
fun buttonPressed() {
  var prio = getValueFromSelection("prio");
  var f = setUpLineDrawing(getValueFromSelection("line-color"));
  switch(prio) {
    case "High" -> sysEnqueue(makeFiber(High, f))
    case "Medium" -> sysEnqueue(makeFiber(Medium, f))
    case "Low" -> sysEnqueue(makeFiber(Low, f))
    case _ -> ()
  }
}
```

Based on the design choice discussed in Section 3.1, the two dropdown menus and the "Draw Line" button can be simplified into one single button for users. This new button should enable users to create a fiber with a fixed color and priority level. In order to achieve this, three buttons are needed, with each button responsible for creating fibers of a specific color and priority level. For example, green represents high-priority fibers, blue represents medium-priority fibers, and red represents low-priority fibers. The implementation can be accomplished by replacing the HTML code for the dropdown menus and button with the new code for the three buttons. Additionally, since the dropdown menus are no longer necessary, the `buttonPressed` function can be updated to hardcode the values for priority and color. As a result, the existing code can be replaced from:

```
<div class="selection margin-10 center">
  <select id="line-color">
    <option value="red">Red</option>
    <option value="green">Green</option>
    <option value="blue">Blue</option>
    <option value="yellow">Yellow</option>
    <option value="#801638">Berry </option>
    <option value="#027878">Teal</option>
  </select>
  <select id="prio">
    <option value="Low">Low</option>
    <option value="Medium">Medium</option>
    <option value="High">High</option>
  </select>
</div>
<button class="block button center" l:onclick="{buttonPressed()}">
  Draw Line
</button>
```

to the new code as shown below:

```
<div class="selection margin-10 center">
  <button
    class="block button center"
    l:onclick="{buttonPressed("Low", "red")}">
    Draw Low Line
  </button>
  <button
    class="block button center"
    l:onclick="{buttonPressed("Medium", "blue")}">
    Draw Medium Line
  </button>
  <button
    class="block button center"
    l:onclick="{buttonPressed("High", "green")}">
    Draw High Line
  </button>
</div>
```

The proposed modifications to the user interface simplify the process of creating fibers and improve user experience by making the interface more intuitive and user-friendly. Additionally, these changes result in cleaner and more maintainable code. It is important to note that the function signature of `buttonPressed` is modified to receive hardcoded color and priority values as arguments, which leads to a revised implementation for `buttonPressed`. The new function implementation is shown below:




Figure 3.3: Before modification




Figure 3.4: After modification

```
fun buttonPressed(prio, color) {
  var f = setUpLineDrawing(color);
  switch(prio) {
```

```

        case "High" -> sysEnqueue(makeFiber(High, f))
        case "Medium" -> sysEnqueue(makeFiber(Medium, f))
        case "Low" -> sysEnqueue(makeFiber(Low, f))
        case _ -> ()
    }
}

```

Finally, the user interface before and after the modifications are displayed in Figure 3.3 and Figure 3.4, respectively.

3.4 Editing Priority of the Fiber

As I described in Section 2.3.1.2, defining an operation inside the effect handlers requires providing both of its definition and implementation. Defining the `GetPrio` operation is straightforward since the state of the fiber can be accessed by having an additional parameter in the `handle` function. Therefore, every time this operation is captured by the handlers, its current priority level can be returned upon resumption.

The `SetPrio` operation is more challenging to implement than `GetPrio` since it requires updating both the priority level in the fiber's state and its position in the correct Priority Queue inside the `PrioQueue` data structure shown in Figure 3.2. To implement `SetPrio`, a parameter `p` for priority level is defined in the operation definition, and a helper function called `newPrioToFiber` is defined. This function takes `p` as an argument, creates a new fiber with the same information as the original fiber but with the priority level changed to the selected level `p`, and enqueues this new fiber into the appropriate Priority Queue in `PrioQueue`. The scheduler will then yield control to the next fiber according to its scheduling mechanism.

The implementation code for `GetPrio` and `SetPrio` operations is presented below. The variable `state` holds the current state of the fiber's scheduler, including its priority level and the `PrioQueue` data structure. This state is passed to the handler function, so that the operations can be executed within the context of the current fiber's state. As the fiber interacts with the operations defined in the program, the `state` parameter is updated throughout the execution of the fiber.

```

handle(fiber.f()) (state <- (prio=fiber.prio, runQ=runQ)){
  case x -> runNext(poll(state.runQ))
  case <GetPrio => resume> -> resume(state.prio, state)
  case <SetPrio(p) => resume> ->
    var q = fiberEnqueue(
      newPrioToFiber(resume, p),
      state.runQ
    );
    runNext(poll(q))
  case <Fork(f) => resume> -> ...
  case <Yield => resume> -> ...
}

```


Chapter 4

Modularization

In the current implementation, all functionalities of the application, including the HTML, CSS code for the web interface, the JavaScript FFI for handling the drawing function on the canvas, Links code for maintaining the scheduler queue, and the fiber scheduler implemented using effect handlers, are all contained in one file. It would be more desirable to separate these functionalities into different modules, which aligns with standard software engineering coding practices [1]. This would not only make the code more organized and easier to maintain, but also improve its reusability and extensibility in the long run. To achieve this, the application can be first broken down into separate functionalities as follows:

- HTML and CSS code for the web user interface
- Links code for the fibers and fiber schedulers
- Links code for the scheduler queues, including its enqueue and dequeue functionalities
- Links code for the user interaction
- Links code for handling drawing functionalities
- JavaScript FFIs for accessing the actual canvas rendering functionalities and maintaining the fibers

The process of separation has been divided into two sections. Section 4.2 focuses on the modularization process, which separates each individual functionality into separate modules. On the other hand, Section 4.3 focuses on the separation process specifically in the context of effect handlers.

4.1 JavaScript FFIs

Before proceeding with the separation process, it is necessary to discuss the role played by JavaScript FFIs in the application. The main application involves two JavaScript FFIs: one to render the lines on the canvas using Canvas API (`canvas.js`), and the other to manage fibers in the JavaScript runtime by providing a similar scheduler queue

structure in the `runtime.js`. While the purpose of `canvas.js` is straightforward, which is to provide `Links` access to the drawing methods in the `Canvas` API, the purpose of `runtime.js` may not be immediately obvious. It might seem redundant to define another scheduler queue in the JavaScript runtime when a `PrioQueue` structure is already defined within the `Links` program.

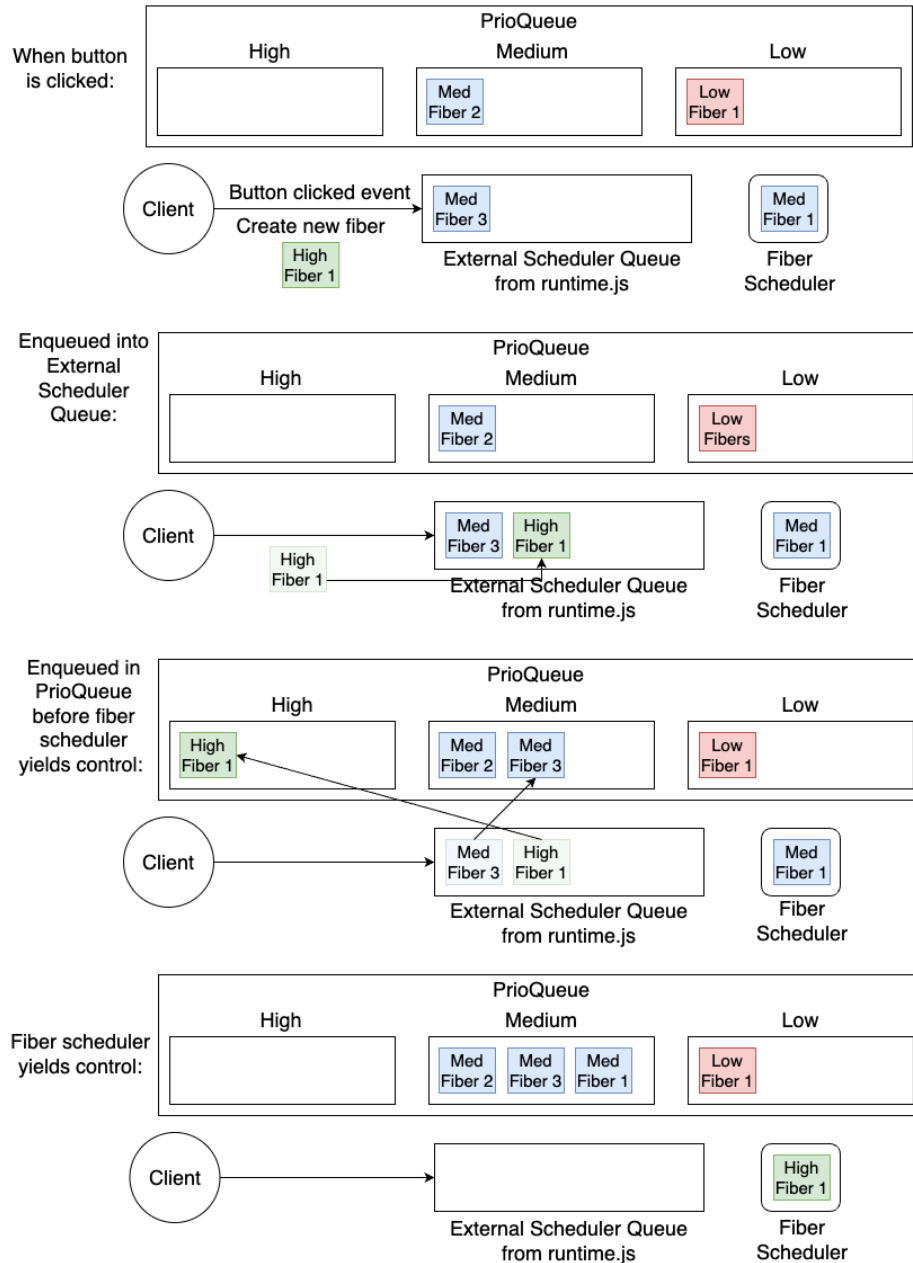


Figure 4.1: How external scheduler queue interacts with `PrioQueue`


The main purpose of having this additional scheduler queue also known as the external scheduler queue in `runtime.js` is to temporarily store the newly created fibers triggered by button-click events. The reason for requiring two separate queues to maintain fibers is explained in detail in Section 6.3. This external scheduler queue includes its own enqueue and dequeue functions. When a button is clicked, the application creates a

new fiber using the `makeFiber` function. Instead of performing a `Fork` operation, the `enqueue` function of the external scheduler queue is called, and the new fiber is added to this queue. Then, when the scheduler decides to yield control to other fibers, it first calls the `dequeue` function of the external scheduler queue to load the fibers from the external queue into the actual `PrioQueue`. The scheduler then selects the next appropriate fiber in the `PrioQueue` to run according to its scheduling mechanism. The entire process of fiber creation and yielding has been presented in Figure 4.1.

4.2 Separation of Concerns

Prior to modularizing the application, it is important to analyze the relationships and dependencies between its functionalities. According to the functionalities listed earlier in Section 4, each one can be separated into its own module. For example, the main application should only contain the user interaction code, while other operations, such as drawing and styling, should be separated into independent modules. Additionally, the scheduler can be divided into separate modules for the scheduling mechanism and the scheduler queue. This approach can help to manage the codebase more effectively and improve overall code quality.

After the separation process, the file `racingLines.links` will have dependencies on the files `racingLines.css` and `lineDrawer.links`, while `scheduler.links` will depend on `queue.links`. Furthermore, the two JavaScript files, namely `canvas.js` and `runtime.js`, are being used for the JavaScript FFI, which will be utilized in both `scheduler.links` and `racingLines.links`. The file structure of the application before and after the separation of concerns is shown in the left and right panels, respectively.



The separation of the code into individual modules not only allows developers to work on specific functionalities without conflicts, but also makes it easier to identify and resolve bugs, since issues can be isolated to specific modules, thereby improving the stability and reliability of the application.

4.3 Extract Effect Interface

To achieve the primary objective of providing a more convenient approach for the application to use different effect handlers on the same effectful operations, it is essential to extract the effect interface, which represents the signatures of the effectful operations, into its own module. This separation ensures the application executes on the same operation sets, regardless of the effect handlers (schedulers) that integrate. This allows to integrate various implementations of scheduling mechanisms on the same operations.

The application includes four operations, namely, `Fork` for creating a new fiber, `Yield` for yielding control to another fiber, and `GetPrio` and `SetPrio` for getting and setting the priority level of the fiber. These four operations are wrapped by corresponding effect interfaces and functions, which are extracted into a new file called `fiberInterface.links`. Additionally, certain utility functions such as `makeFiber` are also included in this file, as they are related to the functionality of fiber creation.

After this modification, the `racingLines.links` and `scheduler.links` modules will incorporate the functionalities and types offered by the `fiberInterface.links` module to create and manipulate the fibers. The updated file structure is as follows:

```
/
├── /js
│   ├── runtime.js
│   └── canvas.js
├── /css
│   └── racingLines.css
├── scheduler.links
├── LineDrawer.links
├── queue.links
├── fiberInterface.links
└── racingLines.links
```

Chapter 5

Schedulers and Case Studies

The heart of this project is the scheduler, which works in conjunction with the `PrioQueue`. To achieve our objective stated in Section 1.2, it is necessary to implement multiple schedulers with different fiber scheduling mechanisms. This serves as the foundation for the later implementation in Section 5.4 which focuses on enabling the easy switching of schedulers within the application. The objective is to make the application run smoothly while also making it easy for developers to switch between different effect handlers with minimal code changes. Ultimately, the aim is to simplify the process enough that only a few lines of code are necessary to switch between different effect handlers. Additionally, to demonstrate the versatility and applicability of the effect handlers switching approach, three case study applications with distinct rendering behaviors are designed and implemented.

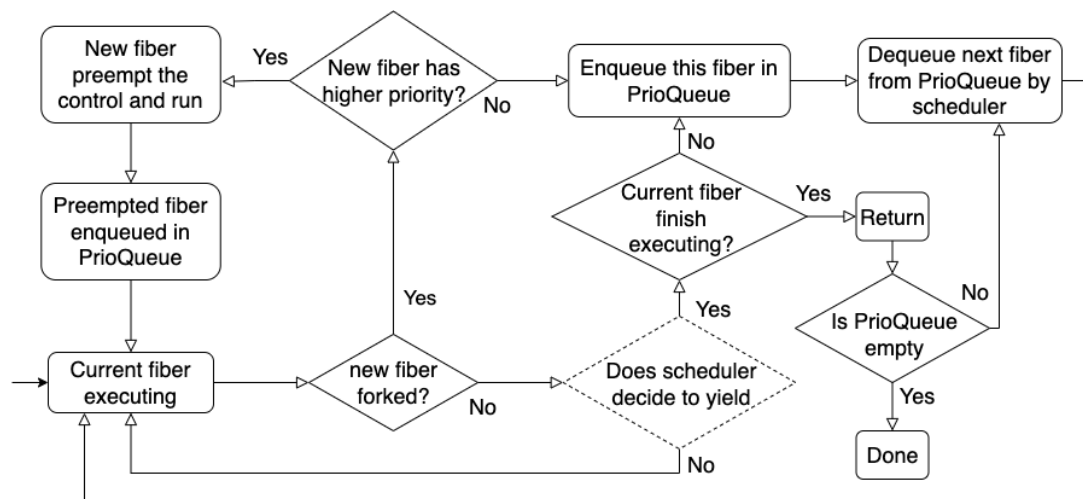


Figure 5.1: Flowchart of how the fiber scheduler operates

Before delving into the specific design and implementation of each different fiber scheduler, it is important to provide an overview of how a fiber scheduler operates in general. When a fiber is created, the scheduler compares its priority level with the currently running fiber to determine whether it should be executed. Then the unselected fiber is enqueued into the `PrioQueue`, while the selected fiber is executed by rendering

a slice of lines on the canvas until the scheduler decides to yield control. If the executing fiber is not finished when control is yielded, it is re-enqueued to the `PrioQueue`. On the other hand, if it finishes its execution, it simply returns. The scheduler then selects the next fiber from the `PrioQueue` by calling its `dequeue` function. The entire process repeats until the queue is empty. The flowchart of the scheduler operation is illustrated in Figure 5.1. In this figure, the conditional block enclosed with the dashed line will be determined by the scheduling mechanisms from each scheduler. With this basic understanding, we can now proceed to explore the specific design and implementation of each different fiber scheduler.

The following sections present the design for the three schedulers that are developed: time-based, steps-based, and probability-based schedulers. Each section provides an overview of how the scheduler operates, along with a diagram to aid in understanding. Moreover, I also present the actual implementation of the scheduler and explain the modifications made to the `PrioQueue` and effect handlers to achieve it.

5.1 Time-based Scheduler

The time-based scheduler is designed to allocate different execution times to fibers based on their priority levels. When yielded by the scheduler, high-priority fibers are allocated 400ms, medium-priority fibers are allocated 200ms, and low-priority fibers are allocated 100ms for execution. This approach enables higher-priority fibers to execute their operations for a longer period of time before yielding control, enabling them to complete their execution more quickly than lower-priority fibers.

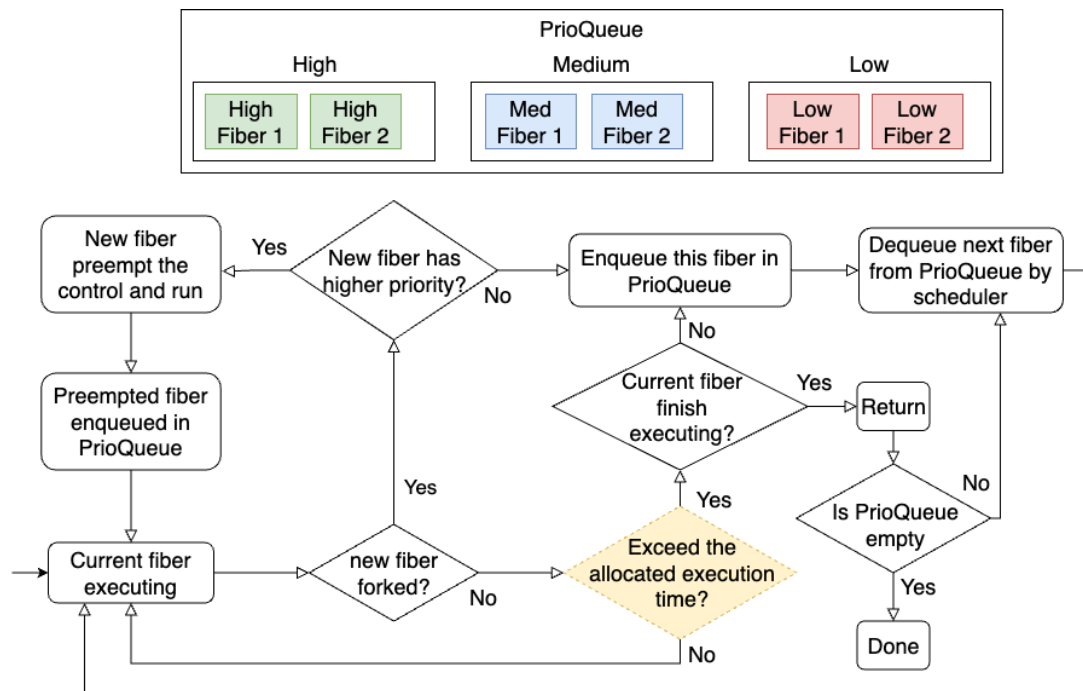


Figure 5.2: Time-based Scheduler

The structure of the `PrioQueue` and the `enqueue` and `dequeue` functions remain the

same as described in Section 3.2, where the scheduler prioritizes selecting fibers with the highest possible priority levels over the lower-priority fibers. The precise process of this scheduling mechanism and the `PrioQueue` structure is illustrated in Figure 5.2. The modified part inside the scheduler is highlighted in yellow.

As can be observed in Figure 5.2, the modifications required for implementing the time-based scheduler are narrowed down to the `Yield` operation and the `state` structure attached to the handler, and the `PrioQueue` remains unchanged. Specifically, a new state field named `startTime`, which represents the start time of the fiber's execution in milliseconds, needs to be added. Upon yielding control, the `startTime` of the next running fiber is updated. Then, execution times are assigned to fibers according to their priority levels, and the scheduler checks if the fiber has used up its allocated execution time. If not, the fiber's execution is resumed; otherwise, control is yielded to the next fiber. The main code implementing these modifications is provided below.

```
handle(fiber.f()) ( state <- (
    prio=fiber.prio,
    runQ=runQ,
    startTime=clientTimeMilliseconds()
  ){
  ...
  case <Yield => resume> ->
    var currentTime = clientTimeMilliseconds();
    var buffer = switch(state.prio){
      case High -> 400
      case Medium -> 200
      case Low -> 100
    };
    if (currentTime - state.startTime <= buffer) {
      resume(), state
    } else{
      var q = fiberEnqueue(
        resumptionToFiber(resume, state.prio),
        state.runQ
      );
      runNext(poll(q))
    }
  }
}
```

5.2 Step-based Scheduler

Step-based schedulers allocate different numbers of steps for fibers based on their priority levels, with high-priority fibers receiving 40 steps, medium-priority fibers receiving 20 steps, and low-priority fibers receiving 10 steps. Each step represents a segment of a line. Allocating more steps to a fiber allows it to render a longer line during its execution time.

This scheduler introduces a significant modification to the enqueue and dequeue behavior in the `PrioQueue`. Instead of prioritizing the dequeuing of higher-priority fibers, it implements a first-in, first-out (FIFO) policy like a classic queue. As a result, each fiber gets to run alternately without interruption from higher-priority fibers. The structure of the `PrioQueue` and the scheduling mechanisms is illustrated in Figure 5.3.

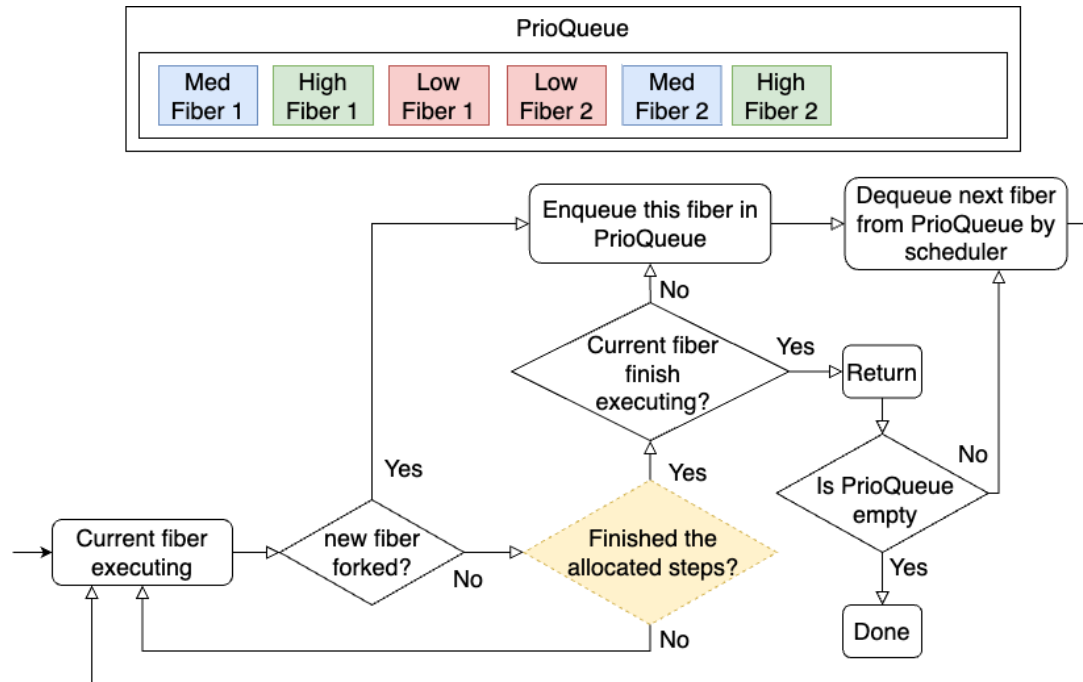


Figure 5.3: Step-based Scheduler

As depicted in Figure 5.3, this scheduler operates in a simpler way as there is no preemption when a new fiber is created, and the `PrioQueue` behaves like a standard queue. Regarding the scheduler aspect, I made modifications to both the `Fork` and `Yield` operations. The implementation for the `Fork` operation is much simpler, where the new fibers are directly enqueued into `PrioQueue` regardless of their priority levels. For `Yield` operation, a new state called `step` is added to keep track of the number of steps taken by the currently executing fiber. The scheduler yields control only when the fiber has completed all its allocated steps. The main code implementing these modifications is provided below.

```
handle(fiber.f())(state <- (prio=fiber.prio, runQ=runQ, step = 0)){
  ...
  case <Fork(f) => resume> ->
    var q = poll(state.runQ);
    var qq = fiberEnqueue(fiber0ToFiber(f), q);
    resume(), (state with runQ = qq)
  case <Yield => resume> ->
    var steps = switch(state.prio){
      case High -> 40
      case Medium -> 20
      case Low -> 10
```



```

        case None -> 0
    };
    var c = state.step + 1;
    if (c < steps) {
        resume(), (state with step = c)
    } else{
        var q = fiberEnqueue(
            resumptionToFiber(resume, state.prio),
            state.runQ
        );
        runNext(poll(q))
    }
}

```

5.3 Probability-based Scheduler

Like rolling a dice, the probability-based scheduler selects the next fiber to run by randomly choosing the priority level and dequeuing the fiber from the corresponding priority queue. The High Priority Queue has a 50% chance of being chosen, the Medium Priority Queue has a 30% chance, and the Low Priority Queue has a 20% chance. The scheduling mechanism is similar to the time-based scheduler discussed in Section 5.1, with the difference being that each fiber now has the same amount of allocated time (200ms) during its execution period. This setup ensures that even though each fiber has the same allocated execution time, higher-priority fibers have a higher chance of being chosen and finishing their rendering sooner.

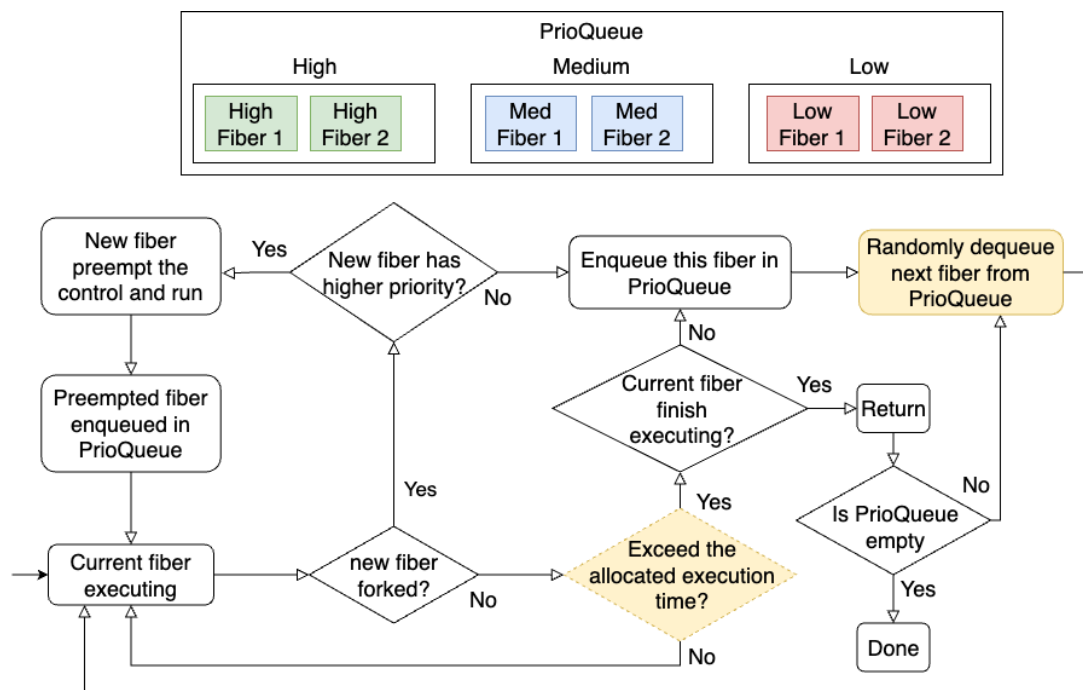


Figure 5.4: Probability-based Scheduler

The `PrioQueue` structure remains unchanged in the probability-based scheduler, and the only modification is made to the dequeue mechanism. Instead of dequeuing the fibers from the highest nonempty priority queue, this scheduler dequeues fibers from a particular priority queue determined by random selection based on the specified probabilities for each priority level. If the selected priority queue is empty, the scheduler attempts to dequeue fibers from the highest priority queues to the lowest until the entire `PrioQueue` is emptied. For example, if the selected priority level is `Medium` but the corresponding priority queue is empty, the scheduler will attempt to dequeue fibers from the `High Priority Queue`. If that queue is also empty, then the scheduler will attempt to select the fiber from the `Low Priority Queue`. The structure of the `PrioQueue` and the scheduling mechanism is demonstrated in Figure 5.4.

Figure 5.4 illustrates that the handler function's state is modified in a manner similar to the time-based scheduler discussed in Section 5.1. The implementation of the `Yield` operation is also similar to the time-based scheduler with a minor change in which fibers from all priority levels are allotted 200ms for execution. A crucial addition is the helper function called `getNextPrio`, which is included in the implementations of every operation that involves the invocation of `runNext`. `getNextPrio` randomly selects the priority level, which is then passed to `runNext` as an additional argument. Inside `runNext`, the dequeue function of the `PrioQueue` also uses this specified priority level as an additional argument to dequeue the fiber from the corresponding priority queue. The implementation of `getNextPrio`, `runNext` and the `Yield` operation is shown below:

```
fun getNextPrio() {
    var prob = floatToInt(random() *. 100.0);
    if (prob <= 20) {
        Low
    } else if (prob <= 50) {
        Medium
    } else {
        High
    }
}

fun runNext(q, prio){
    switch(fiberDequeue(q, prio)){
        case (Nothing, _) -> ()
        case (Just((fiber, prio)), q) ->
            fiber.f(makeSchedulerState(prio, q))
    }
}

case <Yield => resume> ->
    var currentTime = clientTimeMilliseconds();
    if (currentTime - state.startTime <= 200) {
        resume((), state)
    }
}
```

```

} else{
    var nextPrio = getNextPrio();
    var q = fiberEnqueue(
        resumptionToFiber(resume, state.prio),
        state.runQ
    );
    runNext(poll(q), nextPrio)
}

```

In summary, after evaluating the three schedulers, it has been observed that the time-based scheduler may not be completely fair. This is because this scheduler prioritizes higher-priority fibers over lower-priority fibers, which may lead to the lower-priority fibers waiting longer for their execution. Consequently, the execution of lower-priority fibers may be delayed or even starved, which means that the time-based scheduler does not provide equal opportunities for all fibers to execute. Therefore, the time-based scheduler cannot be considered fair based on the fairness definition of concurrency.

5.4 Switch Schedulers

The next step after defining and implementing the three schedulers is to integrate them into the application and allow users to switch between them. To achieve this, the schedulers need to be parameterized as variables that can be passed as arguments to the application. My solution is to define each scheduler as a subroute inside the main function and pass it as an argument to the page generation function. To facilitate this, three clickable links are added to the user interface, allowing users to switch between schedulers by clicking the appropriate link. Upon clicking a link, the browser refreshes and navigates to the designated subroute for the selected scheduler. The scheduler is then passed to the page generation function and used to initialize the scheduler. The modified code to enable this functionality has been provided.

```

- fun main_page(_) {
+ fun main_page(schedule) {
    - var pId = spawnClient{
        Scheduler.schedule(makeFiber(High, start))
    };
    + var pId = spawnClient{
        schedule(FiberInterface.makeFiber(High, start))
    };

page
<html>
...
<body>
    + <a href="/scheduler1.links">Scheduler 1</a>
    + <a href="/scheduler2.links">Scheduler 2</a>
    + <a href="/scheduler3.links">Scheduler 3</a>
</body>

```

```
</html>
}

fun main() {
    ...
    - addRoute("/", main_page);
    + addRoute("/", fun(_) {
        main_page(Scheduler.schedule)
    });
    + addRoute("/scheduler1.links", fun(_) {
        main_page(Scheduler1.schedule)
    });
    + addRoute("/scheduler2.links", fun(_) {
        main_page(Scheduler2.schedule)
    });
    + addRoute("/scheduler3.links", fun(_) {
        main_page(Scheduler3.schedule)
    });
    servePages()
}
```




Figure 5.5: Improved user interface with new scheduler switch functionality

This modification greatly simplifies the process of adding or removing schedulers from the application. The developer only needs to add a subroute in the main function and a hyperlink on the UI that links to this subroute, which only requires two lines of code. Users can then switch between schedulers during runtime by clicking the hyperlink. This approach achieves our objective of making it easy for developers to switch and

integrate different effect handlers with minimal changes to the code. The updated user interface is presented in Figure 5.5.

5.5 Case Studies Implementation

The objective now is to validate the general applicability of our scheduler switching approach by building three different applications for case studies that exhibit distinct rendering behaviors. These case studies are necessary because they ensure that our scheduler switching approach is not limited to a specific type of application and can work effectively across various contexts. Once the applications are developed, they will be evaluated and tested in detail in Chapter 6 to demonstrate the ease and robustness of incorporating different schedulers and ensuring the correct functioning of the applications.

In general, the application works as follows: when the user clicks a button, the corresponding event listener triggers the `buttonPressed` function. This function creates the fiber based on the priority level, and the fiber is enqueued into the external scheduler queue in `runtime.js` through JavaScript FFI. The fiber is later dequeued and stored in the actual `PrioQueue` when the fiber scheduler yields control. The fiber has two properties: its priority level and a drawing function responsible for rendering lines on the canvas. During the fiber's execution, the rendering function is invoked, which draws the line slice by slice by calling the `drawCustomUnit` function in `canvas.js` through JavaScript FFI. The scheduler tries to yield control every time a slice of line is rendered. When the rendering process is complete, the function's execution finishes, and the fiber is returned.

This section provides a detailed description and implementation of each case study, along with figures that visualize the rendering behavior of the application. These figures use various symbols to help readers understand the behavior of the application, such as removing the border of the button to indicate a button-pressed action, using arrows to indicate the rendering direction of lines, using pause signs to indicate the fibers that are being interrupted and paused, and lines without any signs indicate that a line has reached the end. Throughout this and the following chapters, the terms fiber and line are used interchangeably, as each fiber is represented by a line on the canvas. Additionally, each case study uses the time-based scheduler to maintain consistency in demonstrating and visualizing rendering behaviors.

5.5.1 Case Study 1

The first case study has a general and intuitive rendering behavior, where each fiber starts from the left and proceeds toward the right of the canvas. Once the line reaches the right edge of the canvas, it is considered finished and the fiber returns from the effect handlers. A visual representation of this behavior is provided in Figure 5.6. As illustrated, the low-priority fiber is created first and starts rendering. However, its execution is interrupted when the higher-priority fiber is created and takes control of the execution. The higher-priority fiber then renders until completion before the lower-priority fiber resumes rendering until it is also finished.




Figure 5.6: Rendering behavior of Case Study 1

The rendering function begins by retrieving the start and end coordinates of the line to be drawn. When the fiber is executing, the function `drawLineInColor` is called, which renders the line with the specified color and then calls the `drawHorizontalLine` function. This function first selects the canvas using DOM operation and invokes the inner function called `drawSlice`, which is responsible for rendering the line on the canvas slice by slice using the `drawCustomUnit` function through JavaScript FFI. The slices of the line are drawn progressively from left to right, with control being yielded to other fibers after each slice is rendered. Then, after the resumption of execution, `drawSlice` delays the execution of the next slice of the line, allowing the line to be drawn progressively. This process repeats recursively with the coordinates being updated at each iteration until the entire line is drawn.

5.5.2 Case Study 2



Figure 5.7: Rendering behavior of Case Study 2

The second case study is an extension of the first case study, but with a modification in the rendering behavior. In this case, after the line is rendered from the left to the right edge of the canvas, instead of returning the corresponding fiber, the line reverses its direction and moves back and forth repeatedly. As a result, the fiber never returns. Figure 5.7 provides a visual representation of this rendering behavior to aid

in understanding. Since the lines are rendered back and forth, the low-priority fiber is always paused because the high-priority fiber is continuously rendering and never yields control.

Similarly, the implementation of this rendering behavior is also an extension of the implementation of Case Study 1, with the primary modification being made to the `drawLineInColor` function. Once the initial line is rendered from left to right using `drawHorizontalLine`, another additional invocation of `drawHorizontalLine` is made, but this time the line is rendered from right to left. After the line is drawn in both directions, `drawLineInColor` is recursively called to render the line back and forth continuously.

5.5.3 Case Study 3

In this case study, I investigate the effects of changing the priority level of fibers during runtime. Similar to Case Study 2, the line is rendered back and forth repeatedly after reaching the end. However, the key difference is that when the line is rendered from right to left, the priority level of the fiber is set to Low. When the line starts rendering from left to right again, the priority level is reset back to its original level. To illustrate this behavior clearly, Figure 5.8 has been included. This modification allows us to examine how changes in priority levels can affect the rendering behavior of the application.



Figure 5.8: Rendering behavior of Case Study 3

In the given figure, the high-priority fiber interrupts the execution of the low-priority fiber when it renders the line from left to right. When the high-priority fiber renders the line from right to left, both fibers render alternatively as the priority level of the high-priority fiber has temporarily set to low. To achieve this behavior, in addition to the implementation of Case Study 2, the `var curPrio = FiberInterface.getPrio()` and `FiberInterface.setPrio(curPrio)` functions are used before and after the second `drawHorizontalLine` function respectively.

Chapter 6

Evaluation

6.1 Case Studies

This section evaluates the three schedulers described in Chapter 5 and the scheduler switch approach presented in Section 5.4 using the case study applications developed in Section 5.5. In this evaluation, I will create fibers in the following order for each scheduler: Low, Medium, High, Medium, High, Low. This ordering is chosen to demonstrate the interruption and resumption of the lower-priority fibers and validate the effectiveness of the scheduler switching approach across various applications. Additionally, videos of each scheduler's behavior on each case study have been recorded and posted on YouTube as a playlist [4].

6.1.1 Case Study 1

In the time-based scheduler, the two high-priority lines are rendered first, causing interruption to all other fibers, including the lower-priority fibers that started rendering before them. Once the high-priority fibers finish execution, the medium-priority fibers start executing, followed by the low-priority fibers after the medium-priority fibers finish execution.

In the step-based scheduler, fibers are executed alternately in the same order as they were created. Typically, the high-priority fibers finish their execution faster due to the fact that they have more steps in their execution period. The medium-priority fibers finish in second place, followed by the low-priority fibers. However, it is possible for lower-priority fibers to finish rendering faster than higher-priority fibers by creating the higher-priority fibers much later than the lower-priority fibers since this scheduler has no interruption.

The probability-based scheduler is less predictable, with higher-priority fibers having a higher chance of finishing first, but with the possibility of lower-priority fibers finishing first in some cases. Under certain extreme test cases, such as creating one low-priority fiber followed by fourteen high-priority fibers, the low-priority fiber has a high probability of finishing first. This is because, although the High Priority Queue has an 80% chance of being selected (due to the empty Medium Priority Queue so the

scheduler will also dequeue fibers from the High Priority Queue when the medium priority level is randomly selected), the fourteen high-priority fibers render alternately, which can slow down their individual rendering process. Meanwhile, the low-priority fiber is the only fiber in the Low Priority Queue and is thus always dequeued to render more efficiently when the low-priority level is randomly selected by the scheduler.

6.1.2 Case Study 2

In the time-based scheduler, due to the repetitive rendering of lines back and forth, only the high-priority lines are able to continue rendering on the canvas while the lower-priority lines remain paused indefinitely. This is because the higher-priority fibers never get returned and popped out from the `PrioQueue`, resulting in the lower-priority lines being unable to resume rendering.

The step-based scheduler renders fibers alternately and repeatedly back and forth. However, due to the fact that higher-priority fibers have more steps allocated for their execution period, they can render faster than lower-priority fibers.

In the probability-based scheduler, all fibers are rendered repeatedly back and forth, and the probability of a fiber being selected for rendering depends on its priority level. Therefore, higher-priority fibers have a higher chance of getting selected by the scheduler, resulting in faster rendering times. However, in certain cases, such as having one low-priority fiber and fourteen high-priority fibers, the low-priority fiber can render faster than the high-priority fibers since the probability of selection is distributed among all the high-priority fibers.

6.1.3 Case Study 3

In the time-based scheduler, the execution of higher-priority fibers interrupts the execution of lower-priority fibers. As a result, two high-priority fibers reach the end before the lower-priority fibers. When two high-priority fibers reach the end, they do not return and are not dequeued from the `PrioQueue`. Instead, they wait at the end until the two medium-priority fibers are rendered to the end. Then, all the fibers start executing because the higher-priority fibers' priority level is changed to low as they reach the end. The rendering process for all fibers continues until one of the higher-priority fibers reaches back to the starting point, where its priority level is set back to its original level. This will again pause the rendering process of lower-priority fibers including the higher-priority fibers whose priority level has been temporarily changed to low.

In the step-based scheduler, the behavior is similar to the step-based scheduler in Case Study 2, but with a key difference. When higher-priority fibers reach the end of the canvas, their priority level is temporarily changed to low. This results in slower rendering when the fibers start rendering back.

In the probability-based scheduler, the behavior is similar to the one in Case Study 2, with the only difference being that the higher-priority fibers render slower when rendering back due to their priority levels being temporarily changed to low, making it less likely for them to be selected by the scheduler. This feature also affects the

extreme test case where there is one low-priority fiber and fourteen high-priority fibers, making it less likely for the low-priority fiber to win the race. This is because the high-priority fibers set their priority levels to low when they reach the end, which makes the low-priority fiber less likely to be selected by the scheduler and thus rendering much slower.

6.2 Triumph

This project has successfully implemented the switching of effect handlers, allowing developers to interpret the same effectful operations with different implementations under the same application context. Additionally, the project improved the maintainability and readability of the application code by modularizing different functionalities into separate modules. Finally, by implementing the change of priority functionalities in the fiber scheduler, this work highlights the potential for developing the pause and resume functionalities of fibers during runtime in the future.

Our application, compared with the original `racing.links` example [11], offers a more user-friendly experience by providing a simple button for users to create a fiber, along with additional features such as the ability to create medium-priority fibers and use different schedulers for different behaviors. From an implementation perspective, our application has significantly reduced the amount of code needed to implement these features, with only 129 lines of code, including comments, which is only 29% of the original code length of 444 lines. This modularized codebase enables faster debugging process by precisely pinpointing errors to specific modules.

6.3 Challenges: Mismatch Between Event Listeners and Effect Handlers

As this project mainly focuses on Links, minimizing or eliminating the dependency on JavaScript FFI has been attempted. Despite this effort, the use of `canvas.js` could not be avoided due to the lack of support for the Canvas API in Links, making it dependent on JavaScript. On the other hand, a significant amount of time and effort was invested in attempting to remove `runtime.js` from the application. This file contains functionalities such as managing fibers in the JavaScript runtime before enqueueing them into `PrioQueue`, which can, in theory, be replaced by the Links code.

Based on the introduction of `runtime.js` in Section 4.1 and the interaction between the external scheduler queue and the `PrioQueue` illustrated in Figure 4.1, it appears that the current implementation has resulted in the `Fork` operation being bypassed altogether. This goes against the principle of using effect handlers, as every action affecting the fiber should be carried out using operations. The use of `runtime.js` as JavaScript FFI for fiber management is a workaround that negates the use of effectful operations. Therefore, there is a requirement to address the interpretation of the `Fork` operation and eliminate the use of the external scheduler queue in order to comply with the principles of using effect handlers.

One potential solution to address the bypassing of the `Fork` operation is to replace all `enqueue` functions with `Fork` operations. However, this approach may not work as expected since the event listener of the button is not handled as part of the computational context managed by the scheduler. As a result, this approach leads to compile-time errors. Thus, a more nuanced solution is needed, which involves handling the event listener within the context of the scheduler to properly interpret the `Fork` operation.

An alternative approach to solving the problem of bypassing the `Fork` operation is to use the actor model, which was introduced in Section 2.4 and is built into `Links`. The actor model replaces the event listener with the message-passing mechanism where each button clicks event dispatches a message to a receiver that executes the corresponding function based on the message. A comparison between the same implementation of the actor model-based receiver and the event listener has been presented below.

```
// Actor model-based
fun receiver(){
  Scheduler.yield();
  receive {
    case CreateLowLine -> buttonPressed("Low", "red")
    case CreateMediumLine -> buttonPressed("Medium", "blue")
    case CreateHighLine -> buttonPressed("High", "green")
  };
  receiver()
}

fun main_page(_){

  var pId = spawnClient{
    Scheduler.schedule(Scheduler.makeFiber(High, receiver))
  };
  ...
  <button l:onclick="{pId ! CreateHighLine}">
    Draw High Line
  </button>
  ...
}

// Event listener based
fun main_page(){

  var pId = spawnClient{
    schedule(FiberInterface.makeFiber(High, start))
  };
  ...
  <button l:onclick="{buttonPressed("High", "green")}">
    Draw High Line
  </button>
}
```

```

    ...
}

```

By merging the three event listeners into one `receiver` function, it facilitates the handling of this function by the scheduler, which enables the possibility to perform the Fork operation. This approach eliminates the need to use `runtime.js`.

The approach described above, although promising, resulted in unexpected issues due to a false assumption made in the initial design. It was assumed that the `receiver` function listens for the message asynchronously, but in reality, the synchronous listening process of the `receiver` function in Links blocks the rendering process, which in turn, blocks the entire scheduler when it is waiting for the message. As a result, if a button is clicked, the newly created fiber is only able to draw a portion of the line before yielding to the next fiber. Then, the entire execution is blocked at the `receive` block after the resumption until the user clicks the “Draw Line” button again. This behavior makes the implementation unsuitable for the intended purpose.

```

fun receiver() {
  Scheduler.yield();
  // Blocked here until the receiver receives new messages
  receive {
    case CreateLowLine -> buttonPressed("Low", "red")
    case CreateMediumLine -> buttonPressed("Medium", "blue")
    case CreateHighLine -> buttonPressed("High", "green")
  };
  receiver()
}

```

Although frustrating, this issue exposes the current limitations of wrapping a single effect handler around all event listeners in Links. In addition, there is a mismatch between effect handlers, event listeners, and actors, making it difficult to integrate the `receiver` function into the cooperative concurrency context of effect handlers. Therefore, currently, the only option is to use shared state and implement communication between event listeners and effect handlers through `runtime.js` using JavaScript FFI.

Chapter 7

Conclusion and Future Work

This project successfully implemented an approach to switch effect handlers with minimal effort to interpret the same effectful operations under the same application context. To achieve this, I improved the user experience of the application, expanded the scheduler to support three different priorities and the changing of priority levels for the fibers. I then performed the separation of concerns to make the application code more readable and maintainable, and extracted the effect interface into a separate module. Next, I implemented three different types of schedulers to serve as the foundation for the seamless addition, removal, and switching of effect handlers. Additionally, I designed and implemented three case studies to visualize user-level threads, which were then used to evaluate the effectiveness of the effect handlers switching mechanism.

While this project successfully achieved the main objective stated in Section 1.2, I was not fully able to eliminate the use of `runtime.js` as JavaScript FFI to maintain the newly created fibers due to the mismatch between effect handlers, event listeners, and actors in Links. Nevertheless, this work provided valuable contributions to improving the user experience of defining, plugging in, and switching between different effect handlers within the same application, allowing the same effectful operations to have different implementations.

7.1 Future Work

Scrollable Canvas At present, the canvas has a limit of displaying only up to 15 lines. Although it is still possible for users to create more fibers, the newly created fibers will not be visible on the canvas due to this limitation. To address this issue and enhance user experience, it is recommended to implement a scrollable canvas feature that allows users to view all the created fibers.

Pause and Resume of Fibers Given the successful implementation and presentation of the `GetPrio` and `SetPrio` effectful operations in Section 3.4, it may be possible to extend the capabilities of our application by implementing the pause and resume of fibers using these operations. Such an implementation would further demonstrate the potential of effect handlers in the development of concurrent web applications.

More Case Study Applications While the existing case study applications effectively demonstrate the versatility of our effect handler switching mechanism, the evaluation can be further improved by introducing case studies with completely different behaviors that deviate from the rendering of lines on the canvas. This will enhance the diversity of our case studies and provide more concrete and convincing evaluation results.

Benchmark Application The current evaluation of our application only compares it with the original `racing_lines.links` example. To provide a more comprehensive evaluation, it would be beneficial to develop a similar application in React or other popular frameworks used by developers and industries. A comparison between our Link-based application and the ones developed using these frameworks could be conducted to demonstrate the effectiveness and uniqueness of our approach.

Bibliography

- [1] Carliss Baldwin and Kim Clark. “*Design Rules Volume I: The Power of Modularity*”, volume 1. 01 2000. doi:10.2307/259400.
- [2] Luiz Barroso, Mike Marty, David Patterson, and Parthasarathy Ranganathan. “Attack of the Killer Microseconds”. *Commun. ACM*, 60(4):48–54, mar 2017. doi:10.1145/3015146.
- [3] Andrej Bauer and Matija Pretnar. “Programming with Algebraic Effects and Handlers”. *CoRR*, abs/1203.1539, 2012. URL: <http://arxiv.org/abs/1203.1539>, arXiv:1203.1539.
- [4] Steven Chang. “Developing Concurrent Web Applications with Effect Handlers in Links: Case Studies”, Apr. 2023. URL: <https://www.youtube.com/playlist?list=PLjXsWCuXyRhFALyJW0SNvDJBFtVavDcx7>.
- [5] Ezra Cooper, Sam Lindley, Philip Wadler, and Jeremy Yallop. “Links: web programming without tiers”. *Formal Methods for Components and Objects, 5th International Symposium, FMCO 2006, Amsterdam, The Netherlands, November 7-10, 2006, Revised Lectures.*, pages 266–296, 2006. doi:10.1007/978-3-540-74792-5_12.
- [6] Stephen Dolan, Spiros Eliopoulos, Daniel Hillerström, Anil Madhavapeddy, K.C. Sivaramakrishnan, and Leo White. “*Concurrent System Programming with Effect Handlers*”, pages 98–117. Apr. 2018. doi:10.1007/978-3-319-89719-6_6.
- [7] dom. “DOM Specification”. URL: <https://dom.spec.whatwg.org/>.
- [8] flexiple. “Introduction to React Fiber”. URL: <https://flexiple.com/react/react-fiber/>.
- [9] Yannick Forster, Ohad Kammar, Sam Lindley, and Matija Pretnar. “On the Expressive Power of User-Defined Effects: Effect Handlers, Monadic Reflection, Delimited Control”. *Proceedings of the ACM on Programming Languages*, 1, 10 2016. doi:10.1145/3110257.
- [10] Carl Hewitt. “Actor Model of Computation: Scalable Robust Information Systems”, 2015. arXiv:1008.1459.
- [11] Daniel Hillerström. “*racing_lines.links*”. URL: <https://github.com/links-lang/links/blob/master/examples/handlers/racing-lines.links>.

- [12] Daniel Hillerström. “Handlers for Algebraic Effects in Links”. Master’s thesis, Edinburgh, United Kingdom, 2015.
- [13] Daniel Hillerström and Sam Lindley. “Liberating Effects with Rows and Handlers”. In *Proceedings of the 1st International Workshop on Type-Driven Development, TyDe 2016*, page 15–27, New York, NY, USA, 2016. Association for Computing Machinery. doi:10.1145/2976022.2976033.
- [14] Daniel Hillerström, Sam Lindley, Robert Atkey, and K. C. Sivaramakrishnan. “Continuation Passing Style for Effect Handlers”. 84:18:1–18:19, 2017. URL: <http://drops.dagstuhl.de/opus/volltexte/2017/7739>, doi:10.4230/LIPIcs.FSCD.2017.18.
- [15] HTML. “HTML Specification”. URL: <https://html.spec.whatwg.org/s>.
- [16] JavaScript. “JavaScript Official Website”. URL: <https://www.javascript.com/>.
- [17] Karthik Kalyanaraman. “A deep dive into React Fiber”, 14 Mar, 2022. URL: <https://blog.logrocket.com/deep-dive-react-fiber/>.
- [18] Ohad Kammar, Sam Lindley, and Nicolas Oury. “Handlers in Action”. *International Conference on Functional Programming*, pages 145–158, 2013. doi:10.1145/2500365.2500590.
- [19] Sam Lindley. “Algebraic effects and effect handlers for idioms and arrows”. *WGP 2014 - Proceedings of the 2014 ACM SIGPLAN Workshop on Generic Programming*, 08 2014. doi:10.1145/2633628.2633636.
- [20] Links. “Links Documentation”. URL: <https://links-lang.org/quick-help.html>.
- [21] Links-lang. “JavaScript FFI”. URL: <https://github.com/links-lang/links/wiki/JavaScript-FFI>.
- [22] Links-lang. “Links Official Website”. URL: <https://links-lang.org/>.
- [23] Links-lang. “Links Wiki”. URL: <https://github.com/links-lang/links/wiki>.
- [24] Gordon Plotkin and John Power. “Adequacy for Algebraic Effects”. volume 2030, Jan. 2001. doi:10.1007/3-540-45315-6_1.
- [25] Gordon Plotkin and Matija Pretnar. “Handling Algebraic Effects”. *Logical Methods in Computer Science*, 9(4), 2013. URL: <https://doi.org/10.48550/arXiv.1312.1399>, doi:10.2168/lmcs-9(4:23)2013.
- [26] React. “Introducing Concurrent Mode”. URL: <https://17.reactjs.org/docs/concurrent-mode-intro.html>.
- [27] React. “Introducing Hooks”. URL: <https://reactjs.org/docs/hooks-intro.html>.
- [28] React. “React Official Website”. URL: <https://reactjs.org/>.

- [29] React. “Virtual DOM and Internals”. URL: <https://reactjs.org/docs/faq-internals.html>.
- [30] John Resig. “jQuery Official Website”. URL: <https://jquery.com/>.
- [31] Vue. “vue-concurrency Official Website”. URL: <https://vue-concurrency.netlify.app/>.
- [32] Vue. “Vue Official Website”. URL: <https://vuejs.org/>.
- [33] W3C. “Cascading Style Sheets home page”. URL: <https://www.w3.org/Style/CSS/Overview.en.html>.
- [34] W3Schools. “JavaScript HTML DOM EventListener”. URL: https://www.w3schools.com/js/js_html_dom_eventlistener.asp.

Appendix A

Complete code for racing lines demo

A.1 Case Study 1

```
import Scheduler;
import Scheduler2;
import Scheduler3;
import LineDrawer;
import FiberInterface;

typename Canvas = ();
typename Point = (x:Int, y:Int);

alien javascript "js/runtime.js" {
  delayExecution: (Int) ~> ();
  sysEnqueue: (a) ~> ();
}

alien javascript "js/canvas.js" {
  drawCustomUnit: (Int, Int, Int, Int, String, Canvas) ~> ();
}

var colors = ["green", "red", "blue", "yellow"];
var allColors = ["red", "green", "blue", "yellow", "#801638", "#027878"];

var priorities = [];

# Page Functionality

sig drawHorizontalLine: (Point, Point, String) {Yield:() |e}~> ()
fun drawHorizontalLine(start, end, color){
  var div = getNodeById("box");
  var context = jsGetContext2D(div);
  var diff = LineDrawer.difference(end, start);
```

```

fun drawSlice(n, start){
    drawCustomUnit(start.x, start.y, 1, 25, color, context);
    FiberInterface.yield();
    delayExecution(5);
    if (n > 0) {
        drawSlice(
            n - 1,
            (
                x = start.x + (diff.x / LineDrawer.absolute(diff.x)),
                y = start.y
            )
        )
    }
}
drawSlice(LineDrawer.absolute(diff.x), start)
}

sig drawLineInColor: (Point, Point, String) {Yield:() |e}~> ()
fun drawLineInColor(start, end, color){
    println("start drawing " ^^ color ^^ ", yield.");
    drawHorizontalLine((x=start.x, y=start.y),
                      (x=end.x, y=end.y), color);
    println("finish drawing " ^^ color)
}

sig setUpLineDrawing: (String) ~> () {Yield:() |e}~> ()
fun setUpLineDrawing(color){
    var points = LineDrawer.getPoints();
    switch(points){
        case Nothing -> fun(){
        case Just((start, end)) -> {
            LineDrawer.updatePoints(start.x,
                                    start.y + 40,
                                    end.x ,
                                    end.y + 40);

            fun(){
                drawLineInColor(start, end, color)
            }
        }
    }
}

fun start(){
    FiberInterface.fork(
        FiberInterface.makeFiber(Low, FiberInterface.forever)
    );
}

```

```

fun aux(colors, prios){
  switch((colors, prios)){
    case ([],[]) -> ()
    case (x::xs, y::ys) -> {
      FiberInterface.fork(
        FiberInterface.makeFiber(y, setUpLineDrawing(x))
      );
      aux(xs, ys)
    }
    case (_, _) -> ()
  }
}
aux(colors, priorities)
}

fun buttonPressed(prio, color){
  var f = setUpLineDrawing(color);
  switch(prio){
    case "High" -> sysEnqueue(FiberInterface.makeFiber(High, f))
    case "Medium" -> sysEnqueue(FiberInterface.makeFiber(Medium, f))
    case "Low" -> sysEnqueue(FiberInterface.makeFiber(Low, f))
    case _ -> ()
  }
}

# Web interface

fun main_page(schedule){

  var pId = spawnClient{
    schedule(FiberInterface.makeFiber(High, start))
  };

  page
  <html>
  <head>
    <link
      rel="stylesheet"
      href="css/racingLines.css"
      type="text/css"
    />
  </head>
  <body>
    <canvas
      id="box"
      width="400"

```

```

        height="600"
        class="box center block">
</canvas>
<div class="selection margin-10 center">
    <button
        class="block button center"
        l:onclick="{buttonPressed("Low", "red")}">
        Low
    </button>
    <button
        class="block button center"
        l:onclick="{buttonPressed("Medium", "blue")}">
        Med
    </button>
    <button
        class="block button center"
        l:onclick="{buttonPressed("High", "green")}">
        High
    </button>
</div>
<div class="links">
    <a href="/scheduler1.links">Time-based Scheduler</a>
    <a href="/scheduler2.links">Step-based Scheduler</a>
    <a href="/scheduler3.links">Probability-based Scheduler</a>
</div>
<div id="data" display="none">
    <div id="start-x" value="10"></div>
    <div id="start-y" value="10"></div>
    <div id="end-x" value="390"></div>
    <div id="end-y" value="10"></div>
</div>
</body>
</html>
}

sig main: () ~> ()
fun main() {
    addRoute("/", fun(_) {main_page(Scheduler.schedule)});
    addStaticRoute("/css", "css", [("css", "text/css")]);
    addRoute("/scheduler2.links", fun(_) {
        main_page(Scheduler2.schedule)
    });
    addRoute("/scheduler3.links", fun(_) {
        main_page(Scheduler3.schedule)
    });
    addStaticRoute("/js", "js", [("js", "text/javascript")]);
}

```

```

    servePages ()
}

main()

```

A.2 Case Study 2

```

import Scheduler;
import Scheduler2;
import Scheduler3;
import LineDrawer;
import FiberInterface;

typename Canvas = ();
typename Point = (x:Int, y:Int);

alien javascript "js/runtime.js" {
    delayExecution: (Int) ~> ();
    sysEnqueue: (a) ~> ();
}

alien javascript "js/canvas.js" {
    drawCustomUnit: (Int, Int, Int, Int, String, Canvas) ~> ();
}

var colors = ["green", "red", "blue", "yellow"];
var allColors = ["red", "green", "blue", "yellow", "#801638", "#027878"];

var priorities = [];

# Page Functionality

sig drawHorizontalLine: (Point, Point, String) {Yield:() |e}~> ()
fun drawHorizontalLine(start, end, color){
    var div = getNodeById("box");
    var context = jsGetContext2D(div);
    var diff = LineDrawer.difference(end, start);
    fun drawSlice(n, start){
        drawCustomUnit(start.x, start.y, 1, 25, color, context);
        FiberInterface.yield();
        delayExecution(5);
        if (n > 0) {
            drawSlice(
                n - 1,
                (

```

```

        x = start.x + (diff.x / LineDrawer.absolute(diff.x)),
        y = start.y
    )
    )
}
drawSlice(LineDrawer.absolute(diff.x), start)
}

sig drawLineInColor: (Point, Point, String) {Yield:() |e}~> ()
fun drawLineInColor(start, end, color){
    println("start drawing " ^^ color ^^ ", yield.");
    drawHorizontalLine((x=start.x, y=start.y), (x=end.x, y=end.y), color);
    print("reached top, draw back");
    drawHorizontalLine(
        (x=end.x, y=end.y),
        (x=start.x, y=start.y),
        "white"
    );
    println("finish drawing " ^^ color);
    drawLineInColor(start, end, color)
}

sig setUpLineDrawing: (String) ~> () {Yield:() |e}~> ()
fun setUpLineDrawing(color){
    var points = LineDrawer.getPoints();
    switch(points){
        case Nothing -> fun(){
        case Just((start, end)) -> {
            LineDrawer.updatePoints(start.x,
                                    start.y + 40,
                                    end.x ,
                                    end.y + 40);

            fun(){
                drawLineInColor(start, end, color)
            }
        }
    }
}

fun start(){
    FiberInterface.fork(
        FiberInterface.makeFiber(Low, FiberInterface.forever)
    );
    fun aux(colors, prios){
        switch((colors, prios)){

```

```

        case ([], []) -> ()
        case (x::xs, y::ys) -> {
            FiberInterface.fork(
                FiberInterface.makeFiber(y, setUpLineDrawing(x))
            );
            aux(xs, ys)
        }
        case (_, _) -> ()
    }
}
aux(colors, priorities)
}

fun buttonPressed(prio, color){
    var f = setUpLineDrawing(color);
    switch(prio){
        case "High" -> sysEnqueue(FiberInterface.makeFiber(High, f))
        case "Medium" -> sysEnqueue(FiberInterface.makeFiber(Medium, f))
        case "Low" -> sysEnqueue(FiberInterface.makeFiber(Low, f))
        case _ -> ()
    }
}

# Web interface

fun main_page(schedule){

    var pId = spawnClient{
        schedule(FiberInterface.makeFiber(High, start))
    };

    page
    <html>
    <head>
        <link
            rel="stylesheet"
            href="css/racingLines.css"
            type="text/css"
        />
    </head>
    <body>
        <canvas
            id="box"
            width="400"
            height="600"
            class="box center block">

```



```

</canvas>
<div class="selection margin-10 center">
  <button
    class="block button center"
    l:onclick="{buttonPressed("Low", "red")}">
    Low
  </button>
  <button
    class="block button center"
    l:onclick="{buttonPressed("Medium", "blue")}">
    Med
  </button>
  <button
    class="block button center"
    l:onclick="{buttonPressed("High", "green")}">
    High
  </button>
</div>
<div class="links">
  <a href="/scheduler1.links">Time-based Scheduler</a>
  <a href="/scheduler2.links">Step-based Scheduler</a>
  <a href="/scheduler3.links">Probability-based Scheduler</a>
</div>
<div id="data" display="none">
  <div id="start-x" value="10"></div>
  <div id="start-y" value="10"></div>
  <div id="end-x" value="390"></div>
  <div id="end-y" value="10"></div>
</div>
</body>
</html>
}

sig main: () ~> ()
fun main() {
  addRoute("/", fun(_) {main_page(Scheduler.schedule)});
  addStaticRoute("/css", "css", [("css", "text/css")]);
  addRoute("/scheduler2.links", fun(_) {
    main_page(Scheduler2.schedule)
  });
  addRoute("/scheduler3.links", fun(_) {
    main_page(Scheduler3.schedule)
  });
  addStaticRoute("/js", "js", [("js", "text/javascript")]);

  servePages()
}

```

```
}
main()
```

A.3 Case Study 3

```
import Scheduler;
import Scheduler2;
import Scheduler3;
import LineDrawer;
import FiberInterface;

typename Canvas = ();
typename Point = (x:Int, y:Int);

alien javascript "js/runtime.js" {
  delayExecution: (Int) ~> ();
  sysEnqueue: (a) ~> ();
}

alien javascript "js/canvas.js" {
  drawCustomUnit: (Int, Int, Int, Int, String, Canvas) ~> ();
}

var colors = ["green", "red", "blue", "yellow"];
var allColors = ["red", "green", "blue", "yellow", "#801638", "#027878"];

var priorities = [];

# Page Functionality

sig drawHorizontalLine: (Point, Point, String) {Yield:() |e}~> ()
fun drawHorizontalLine(start, end, color){
  var div = getNodeById("box");
  var context = jsGetContext2D(div);
  var diff = LineDrawer.difference(end, start);
  fun drawSlice(n, start){
    drawCustomUnit(start.x, start.y, 1, 25, color, context);
    FiberInterface.yield();
    delayExecution(5);
    if (n > 0) {
      drawSlice(
        n - 1,
        (
          x = start.x + (diff.x / LineDrawer.absolute(diff.x)),
          y = start.y
```

```

        )
    )
}
}
drawSlice(LineDrawer.absolute(diff.x), start)
}

sig drawLineInColor: (Point, Point, String) {Yield:() |%}~> ()
fun drawLineInColor(start, end, color){
    println("start drawing " ^^ color ^^ ", yield.");
    drawHorizontalLine((x=start.x, y=start.y), (x=end.x, y=end.y), color);
    print("reached top, draw back");
    var curPrio = FiberInterface.getPrio();
    FiberInterface.setPrio(Low);
    drawHorizontalLine(
        (x=end.x, y=end.y),
        (x=start.x, y=start.y),
        "white"
    );
    FiberInterface.setPrio(curPrio);
    drawLineInColor(start, end, color)
}

sig setUpLineDrawing: (String) ~> () {Yield:() |%}~> ()
fun setUpLineDrawing(color){
    var points = LineDrawer.getPoints();
    switch(points){
        case Nothing -> fun(){
        case Just((start, end)) -> {
            LineDrawer.updatePoints(
                start.x,
                start.y + 40,
                end.x,
                end.y + 40
            );
            fun(){
                drawLineInColor(start, end, color)
            }
        }
    }
}

fun start(){
    FiberInterface.fork(
        FiberInterface.makeFiber(Low, FiberInterface.forever)
    );
}

```

```

fun aux(colors, prios){
  switch((colors, prios)){
    case ([],[]) -> ()
    case (x::xs, y::ys) -> {
      FiberInterface.fork(
        FiberInterface.makeFiber(y, setUpLineDrawing(x))
      );
      aux(xs, ys)
    }
    case (_, _) -> ()
  }
}
aux(colors, priorities)
}

fun buttonPressed(prio, color){
  var f = setUpLineDrawing(color);
  switch(prio){
    case "High" -> sysEnqueue(FiberInterface.makeFiber(High, f))
    case "Medium" -> sysEnqueue(FiberInterface.makeFiber(Medium, f))
    case "Low" -> sysEnqueue(FiberInterface.makeFiber(Low, f))
    case _ -> ()
  }
}

# Web interface

fun main_page(schedule){

  var pId = spawnClient{
    schedule(FiberInterface.makeFiber(High, start))
  };

  page
  <html>
  <head>
    <link
      rel="stylesheet"
      href="css/racingLines.css"
      type="text/css"
    />
  </head>
  <body>
    <canvas
      id="box"
      width="400"

```

```

        height="600"
        class="box center block">
</canvas>
<div class="selection margin-10 center">
    <button
        class="block button center"
        l:onclick="{buttonPressed("Low", "red")}">
        Low
    </button>
    <button
        class="block button center"
        l:onclick="{buttonPressed("Medium", "blue")}">
        Med
    </button>
    <button
        class="block button center"
        l:onclick="{buttonPressed("High", "green")}">
        High
    </button>
</div>
<div class="links">
    <a href="/scheduler1.links">Time-based Scheduler</a>
    <a href="/scheduler2.links">Step-based Scheduler</a>
    <a href="/scheduler3.links">Probability-based Scheduler</a>
</div>
<div id="data" display="none">
    <div id="start-x" value="10"></div>
    <div id="start-y" value="10"></div>
    <div id="end-x" value="390"></div>
    <div id="end-y" value="10"></div>
</div>
</body>
</html>
}

sig main: () ~> ()
fun main() {
    addRoute("/", fun(_) {main_page(Scheduler.schedule)});
    addStaticRoute("/css", "css", [("css", "text/css")]);
    addRoute("/scheduler2.links", fun(_) {
        main_page(Scheduler2.schedule)
    });
    addRoute("/scheduler3.links", fun(_) {
        main_page(Scheduler3.schedule)
    });
    addStaticRoute("/js", "js", [("js", "text/javascript")]);
}

```

```

    servePages ()
  }

main()

```

A.4 fiberInterface.links

```

typename Priority = [|High|Medium|Low|None|];
typename Fiber0(e::Eff) = (prio: Priority,
                          f: () {Fork: (Fiber0({ |e})) => () |e}~> ());

sig fork : (Fiber0({ |e})) {Fork: (Fiber0({ |e})) => () |e}~> ()
fun fork(f) {
  do Fork(f)
}

sig yield: () {Yield:() |e}~> ()
fun yield(){
  do Yield
}

sig getPrio: () {GetPrio: () => Priority |e}~> Priority
fun getPrio(){
  do GetPrio
}

sig setPrio: (Priority) {SetPrio: (Priority) => () |e}~> ()
fun setPrio(p){
  do SetPrio(p)
}

sig makeFiber: (
  Priority,
  () {Fork: (Fiber0({ |e})) => () |e}~> ()
) -> Fiber0({ |e})
fun makeFiber(prio, f){
  (prio= prio, f=f)
}

sig forever: () {Yield:() |e}~> ()
fun forever(){ yield(); forever()}

```

A.5 lineDrawer.links

```

typename Point = (x:Int, y:Int);

sig difference: (Point, Point) -> Point
fun difference(a, b){
  (x = a.x - b.x, y= a.y - b.y)
}

sig absolute: (Int) -> Int
fun absolute(n){
  if(n < 0){
    n * -1
  } else{
    n
  }
}

sig parseInt: (String) ~> Maybe(Int)
fun parseInt(n){
  if (n =~ /^[0-9]+$/ ) {
    Just(stringToInt(n))
  } else{
    Nothing
  }
}

fun updatePoints(startx, starty, endx, endy){
  setValue(startx, "start-x", intToString);
  setValue(starty, "start-y", intToString);
  setValue(endx, "end-x", intToString);
  setValue(endy, "end-y", intToString)
}

sig getPoints: () ~> Maybe((Point, Point))
fun getPoints(){
  var startx = parseInt(getValue("start-x"));
  var starty = parseInt(getValue("start-y"));
  var endx = parseInt(getValue("end-x"));
  var endy = parseInt(getValue("end-y"));
  switch((startx, starty, endx, endy)){
    case (Just(sx),
          Just(sy),
          Just(ex),
          Just(ey)) -> Just((x = sx, y = sy), (x = ex, y = ey))
    case (_, _, _, _) -> Nothing
  }
}

```

```

    }
}

```

A.6 runtime.js

```

function delayExecution(delay, kappa){
  window.setTimeout(
    function(){
      return _$K.yield(kappa, _$Constants.UNIT);
    },
    delay
  );
  return _$Constants.UNIT;
}

const SystemQueue = (function(){

  let queue = _$List.nil;

  function enqueue(fiber){
    queue = _$List.cons(fiber, queue);
    return _$Constants.UNIT;
  }

  function dequeue(){
    let temp = queue;
    queue = _$List.nil;
    return temp;
  }

  function length(){
    return _$List.length(queue);
  }

  return { "enqueue": enqueue
          , "dequeue": dequeue
          , "length" : length }

})();

const sysEnqueue = _$Links.kify(SystemQueue.enqueue);
const sysDequeue = _$Links.kify(SystemQueue.dequeue);
const sysQueueLength = _$Links.kify(SystemQueue.length);

```


A.7 Time-based Scheduler

```

open import FiberInterface;
open import Queue;

alien javascript "js/runtime.js" {
  sysDequeue: () ~> [a];
}

typename PrioQueue(a) = (
  high: Queue(a),
  medium: Queue(a),
  low: Queue(a),
  none: Queue(a)
);

typename SchedulerState(a) = (
  runQ: PrioQueue(a),
  prio: Priority,
  startTime: Int
);

typename Fiber(e::Eff) = (prio:Priority,
                          f: (SchedulerState(Fiber({ |e}))) ~e~> ());
typename FiberQueue(e::Eff) = PrioQueue(Fiber({ |e}));

sig priorityEnqueue: (a, Priority, PrioQueue(a)) ~> PrioQueue(a)
fun priorityEnqueue(x, prio, pq){
  switch(prio){
    case High -> (pq with high = enqueue(x, pq.high))
    case Medium -> (pq with medium = enqueue(x, pq.medium))
    case Low -> (pq with low = enqueue(x, pq.low))
    case None -> (pq with none = enqueue(x, pq.none))
  }
}

sig priorityDequeue: (PrioQueue(a)) ~> (Maybe((a, Priority)), PrioQueue(a))
fun priorityDequeue(pq){
  switch(dequeue(pq.high)){
    case (Just(x), q) -> (Just((x, High)), (pq with high= q))
    case (Nothing,_) ->
      switch(dequeue(pq.medium)){
        case (Just(x), q) -> (
          Just((x, Medium)),
          (pq with medium= q)
        )
      }
  }
}

```

```

        case (Nothing, _) ->
            switch(dequeue(pq.low)) {
                case (Just(x), q) -> (
                    Just((x, Low)),
                    (pq with low=q)
                )
                case (Nothing, _) ->
                    switch(dequeue(pq.none)) {
                        case (Just(x), q) -> (
                            Just((x, None)),
                            (pq with none=q)
                        )
                        case (Nothing, _) -> (Nothing, pq)
                    }
            }
    }
}

sig fiberEnqueue: (Fiber({ |e}), FiberQueue({ |e})) ~> FiberQueue({ |e})
fun fiberEnqueue(f, q) {
    priorityEnqueue(f, f.prio, q)
}

sig fiberBulkEnqueue: ([Fiber({ |e})],
    FiberQueue({ |e})) ~> FiberQueue({ |e})
fun fiberBulkEnqueue(fibers, q) {
    switch(fibers) {
        case [] -> q
        case x::xs ->
            var pq = fiberEnqueue(x, q);
            fiberBulkEnqueue(xs, pq)
    }
}

sig fiberDequeue: (FiberQueue({ |e})) ~> (
    Maybe((Fiber({ |e}), Priority)),
    FiberQueue({ |e})
)
fun fiberDequeue(q) {
    priorityDequeue(q)
}

fun emptyPrioQueue() {
    (high= emptyQueue(),
    medium=emptyQueue(),

```

```

    low=emptyQueue(),
    none=emptyQueue())
}

sig fiberQueueLength: (FiberQueue({ |e})) -> Int
fun fiberQueueLength(q) {
  length(q.low.rear) +
  length(q.high.front) +
  length(q.high.rear) +
  length(q.medium.front) +
  length(q.medium.rear) +
  length(q.low.front)
}

#sig schedule: (Fiber0({ |e})) {Yield-, Fork- | e}~> ()
fun schedule(fib) {

  sig makeSchedulerState: (Priority, PrioQueue(a)) ~> SchedulerState(a)
  fun makeSchedulerState(prio, q) {
    (runQ = q, prio = prio, startTime = clientTimeMilliseconds())
  }

  fun runNext(q) {
    switch(fiberDequeue(q)) {
      case (Nothing, _) -> ()
      case (Just((fiber, prio)), q) -> {
        fiber.f(makeSchedulerState(prio, q))
      }
    }
  }

  fun resumptionToFiber(resume, prio) {
    (prio = prio, f = fun(state) { resume(), state })
  }

  fun newPrioToFiber(resume, prio) {
    (prio = prio,
     f = fun(state) {
       resume(), (state with prio = prio)
     })
  }

  mutual {
    sig fiber0ToFiber: (Fiber0({ |%e})) ~%~> Fiber({ |%e})
    fun fiber0ToFiber(fiber) {
      (prio = fiber.prio,

```

```

    f = fun(state) { runFiber(fiber, state.runQ ) }
  }

#sig poll: (FiberQueue({ |e})) ~> FiberQueue({ |e})
fun poll(q){
  var enqueue = map(fiber0ToFiber, sysDequeue());
  fiberBulkEnqueue(enqueue, q)
}

fun runFiber(fiber, runQ){
  println("runQ: " ^^ intToString(fiberQueueLength(runQ)));
  #dump(self());
  handle(fiber.f()) ( state <- (
    prio=fiber.prio,
    runQ=runQ,
    startTime=clientTimeMilliseconds()
  ){
    case x ->
      print("RETURN");
      runNext(poll(state.runQ))
    case <GetPrio => resume> -> resume(state.prio, state)
    case <SetPrio(p) => resume> ->
      var q = fiberEnqueue(newPrioToFiber(resume, p),
        state.runQ);
      runNext(poll(q))
    case <Fork(f) => resume> ->
      print("fork");
      var q = poll(state.runQ);
      # f is new fiber, state.prio is the current fiber
      switch ((f.prio, state.prio)) {
      case (High, _) ->
        # Enqueue the current fiber, run the new fiber
        print("High fiber enqueued");
        var qq = fiberEnqueue(
          resumptionToFiber(resume, state.prio),
          q
        );
        runFiber(f, qq)
      case (Medium, High) ->
        print("Medium high fiber enqueued");
        # Enqueue the new fiber, run the current fiber
        var qq = fiberEnqueue(fiber0ToFiber(f), q);
        resume(), (state with runQ = qq)
      case (Medium, _) ->
        print("Medium fiber enqueued");
        # Enqueue the new fiber, run the current fiber

```

```

        var qq = fiberEnqueue(
            resumptionToFiber(resume, state.prio),
            q
        );
        runFiber(f, qq)
    case (_, High) ->
        # Enqueue the new fiber, run the current fiber
        print("Low High fiber enqueued");
        var qq = fiberEnqueue(fiber0ToFiber(f), q);
        resume(), (state with runQ = qq)
    case (_, Medium) ->
        # Enqueue the new fiber, run the current fiber
        print("Low Medium fiber enqueued");
        var qq = fiberEnqueue(fiber0ToFiber(f), q);
        resume(), (state with runQ = qq)
    case (_, _) ->
        # Enqueue the current fiber, run the new fiber
        print("Low fiber enqueued");
        var qq = fiberEnqueue(
            resumptionToFiber(resume, state.prio),
            q
        );
        runFiber(f, qq)
    }
case <Yield => resume> ->
    var currentTime = clientTimeMilliseconds();
    var buffer = switch(state.prio){
        case High -> 400
        case Medium -> 200
        case Low -> 100
        case None -> 0
    };
    if (currentTime - state.startTime <= buffer) {
        resume(), state
    } else{
        var q = fiberEnqueue(
            resumptionToFiber(resume, state.prio),
            state.runQ
        );
        runNext(poll(q))
    }
}
}
}

runFiber(fib, emptyPrioQueue())

```

```
}

```

A.8 Step-based Scheduler

```
open import FiberInterface;
open import Queue;

alien javascript "js/runtime.js" {
  sysDequeue: () ~> [a];
}

typename PrioQueue(a) = (queue: Queue(a));
typename SchedulerState(a) = (
  runQ: PrioQueue(a),
  prio: Priority,
  step: Int
);

typename Fiber(e::Eff) = (
  prio:Priority,
  f: (SchedulerState(Fiber({ |e}))) ~e~> ()
);

typename FiberQueue(e::Eff) = PrioQueue(Fiber({ |e}));

sig fiberEnqueue: (Fiber({ |e}), FiberQueue({ |e})) ~> FiberQueue({ |e})
fun fiberEnqueue(f, pq){
  (pq with queue = enqueue(f, pq.queue))
}

sig fiberBulkEnqueue: ([Fiber({ |e})],
  FiberQueue({ |e})) ~> FiberQueue({ |e})
fun fiberBulkEnqueue(fibers, q){
  switch(fibers){
    case [] -> q
    case x::xs ->
      var pq = fiberEnqueue(x, q);
      fiberBulkEnqueue(xs, pq)
  }
}

sig fiberDequeue: (FiberQueue({ |e})) ~> (
  Maybe((Fiber({ |e}), Priority), FiberQueue({ |e})
)
fun fiberDequeue(pq){

```

```

    switch(dequeue(pq.queue)) {
      case (Nothing, _) -> (Nothing, pq)
      case (Just(x), q) -> (Just((x, x.prio)), (pq with queue=q))
    }
  }

fun emptyPrioQueue() {
  (queue=emptyQueue())
}

sig fiberQueueLength: (FiberQueue({ |e})) -> Int
fun fiberQueueLength(q) {
  length(q.queue.rear) + length(q.queue.front)
}

#sig schedule: (Fiber0({ |e})) {Yield-, Fork- | e}~> ()
fun schedule(fib) {

  sig makeSchedulerState: (Priority, PrioQueue(a)) ~> SchedulerState(a)
  fun makeSchedulerState(prio, q) {
    (runQ = q, prio = prio, step = 0)
  }

  fun runNext(q) {
    switch(fiberDequeue(q)) {
      case (Nothing, _) -> ()
      case (Just((fiber, prio)), q) -> fiber.f(
        makeSchedulerState(prio, q)
      )
    }
  }

  fun resumptionToFiber(resume, prio) {
    (prio = prio, f = fun(state) { resume(), state } )
  }

  fun newPrioToFiber(resume, prio) {
    (prio = prio,
     f = fun(state) { resume(), (state with prio = prio) } )
  }

  mutual {
    sig fiber0ToFiber: (Fiber0({ |%e})) ~%~> Fiber({ |%e})
    fun fiber0ToFiber(fiber) {
      (prio = fiber.prio,
       f = fun(state) { runFiber(fiber, state.runQ ) } )
    }
  }
}

```

```

}

fun poll(q){
  var enqueue = map(fiber0ToFiber, sysDequeue());
  fiberBulkEnqueue(enqueue, q)
}

fun runFiber(fiber, runQ){
  println("runQ: " ^^ intToString(fiberQueueLength(runQ)));
  #dump(self());
  handle(fiber.f()) ( state <- (
    prio=fiber.prio,
    runQ=runQ,
    step = 0)
  ){
    case x ->
      runNext(poll(state.runQ))
    case <GetPrio => resume> -> resume(state.prio, state)
    case <SetPrio(p) => resume> ->
      var q = fiberEnqueue(
        newPrioToFiber(resume, p),
        state.runQ
      );
      runNext(poll(q))
    case <Fork(f) => resume> ->
      var q = poll(state.runQ);
      # f is new fiber, state.prio is the current fiber
      var qq = fiberEnqueue(fiber0ToFiber(f), q);
      resume(), (state with runQ = qq))
    case <Yield => resume> ->
      var steps = switch(state.prio){
        case High -> 40
        case Medium -> 20
        case Low -> 10
        case None -> 0
      };
      var c = state.step + 1;

      if (c < steps) {
        resume(), (state with step = c)
      } else{
        var q = fiberEnqueue(
          resumptionToFiber(resume, state.prio),
          state.runQ
        );
        runNext(poll(q))
      }
  }
}

```



```

        }
    }
}

runFiber(fib, emptyPrioQueue())
}

```

A.9 Probability-based Scheduler

```

open import FiberInterface;
open import Queue;

alien javascript "js/runtime.js" {
  sysDequeue: () ~> [a];
}

typename PrioQueue(a) = (
  high: Queue(a),
  medium: Queue(a),
  low: Queue(a),
  none: Queue(a)
);

typename SchedulerState(a) = (
  runQ: PrioQueue(a),
  prio: Priority,
  startTime: Int
);

typename Fiber(e::Eff) = (
  prio: Priority,
  f: (SchedulerState(Fiber({ |e}))) ~e~> ()
)
;

typename FiberQueue(e::Eff) = PrioQueue(Fiber({ |e}));

var bufferTime = 200;

sig priorityEnqueue: (a, Priority, PrioQueue(a)) ~> PrioQueue(a)
fun priorityEnqueue(x, prio, pq){
  switch(prio){
    case High -> (pq with high = enqueue(x, pq.high))
    case Medium -> (pq with medium = enqueue(x, pq.medium))
    case Low -> (pq with low = enqueue(x, pq.low))
  }
}

```

```

        case None -> (pq with none = enqueue(x, pq.none))
    }
}

sig priorityDequeue: (PrioQueue(a), Priority) ~> (
    Maybe((a, Priority)),
    PrioQueue(a)
)

fun priorityDequeue(pq, prio){
    switch(prio) {
        case High -> switch(dequeue(pq.high)) {
            case (Just(x), q) -> (
                Just((x, High)),
                (pq with high= q)
            )
            case (Nothing, _) ->
                switch(dequeue(pq.medium)){
                    case (Just(x), q) -> (
                        Just((x, Medium)),
                        (pq with medium= q)
                    )
                    case (Nothing, _) ->
                        switch(dequeue(pq.low)){
                            case (Just(x), q) -> (
                                Just((x, Low)),
                                (pq with low=q)
                            )
                            case (Nothing, _) -> (
                                Nothing,
                                pq
                            )
                        }
                    }
                }
        }
        case Medium -> switch(dequeue(pq.medium)) {
            case (Just(x), q) -> (
                Just((x, Medium)),
                (pq with medium= q)
            )
            case (Nothing, _) ->
                switch(dequeue(pq.high)){
                    case (Just(x), q) -> (
                        Just((x, High)),
                        (pq with high= q)
                    )
                    case (Nothing, _) ->
                }
        }
    }
}

```

```

        switch(dequeue(pq.low)) {
            case (Just(x), q) -> (
                Just((x, Low)),
                (pq with low=q
            )
            case (Nothing, _) -> (
                Nothing,
                pq
            )
        }
    }
}
case Low -> switch(dequeue(pq.low)) {
    case (Just(x), q) -> (
        Just((x, Low)),
        (pq with low= q)
    )
    case (Nothing, _) ->
        switch(dequeue(pq.high)) {
            case (Just(x), q) -> (
                Just((x, High)),
                (pq with high= q)
            )
            case (Nothing, _) ->
                switch(dequeue(pq.medium)) {
                    case (Just(x), q) -> (
                        Just((x, Medium)),
                        (pq with medium=q)
                    )
                    case (Nothing, _) -> (
                        Nothing,
                        pq
                    )
                }
        }
}
case None -> switch(dequeue(pq.none)) {
    case (Just(x), q) -> (
        Just((x, None)),
        (pq with none=q)
    )
    case (Nothing, _) -> (Nothing, pq)
}
}
}

```

```

sig fiberEnqueue: (Fiber({ |e}),
                  FiberQueue({ |e})) ~> FiberQueue({ |e})
fun fiberEnqueue(f, q){
  priorityEnqueue(f, f.prio, q)
}

sig fiberBulkEnqueue: ([Fiber({ |e})],
                      FiberQueue({ |e})) ~> FiberQueue({ |e})
fun fiberBulkEnqueue(fibers, q){
  switch(fibers){
    case [] -> q
    case x::xs ->
      var pq = fiberEnqueue(x, q);
      fiberBulkEnqueue(xs, pq)
  }
}

sig fiberDequeue: (FiberQueue({ |e}), Priority) ~> (
  Maybe((Fiber({ |e}), Priority)),
  FiberQueue({ |e})
)
fun fiberDequeue(q, prio){
  priorityDequeue(q, prio)
}

fun emptyPrioQueue(){
  (high= emptyQueue(),
   medium=emptyQueue(),
   low=emptyQueue(),
   none=emptyQueue())
}

sig fiberQueueLength: (FiberQueue({ |e})) -> Int
fun fiberQueueLength(q){
  length(q.low.rear) +
  length(q.high.front) +
  length(q.high.rear) +
  length(q.medium.front) +
  length(q.medium.rear) +
  length(q.low.front)
}

#sig schedule: (Fiber0({ |e})) {Yield-, Fork- | e}~> ()
fun schedule(fib){

  sig makeSchedulerState: (Priority, PrioQueue(a)) ~> SchedulerState(a)

```

```

fun makeSchedulerState(prio, q){
    (runQ = q, prio = prio, startTime = clientTimeMilliseconds())
}

fun runNext(q, prio){
    switch(fiberDequeue(q, prio)){
        case (Nothing, _) -> ()
        case (Just((fiber, prio)), q) -> fiber.f(
            makeSchedulerState(prio, q)
        )
    }
}

fun resumptionToFiber(resume, prio){
    (prio = prio, f = fun(state) { resume(), state } )
}

fun newPrioToFiber(resume, prio){
    (
        prio = prio,
        f = fun(state) { resume(), (state with prio = prio)}
    )
}

fun getNextPrio() {
    var prob = floatToInt(random() *. 100.0);

    if (prob <= 20) {
        Low
    } else if (prob <= 50) {
        Medium
    } else {
        High
    }
}

mutual {
    sig fiber0ToFiber: (Fiber0({ |%e})) ~%~> Fiber({ |%e})
    fun fiber0ToFiber(fiber){
        (prio = fiber.prio,
         f = fun(state) { runFiber(fiber, state.runQ ) })
    }

    fun poll(q){
        var enqueue = map(fiber0ToFiber, sysDequeue());
        fiberBulkEnqueue(enqueue, q)
    }
}

```

```

}

fun runFiber(fiber, runQ){
  println("runQ: " ^^ intToString(fiberQueueLength(runQ)));
  #dump(self());
  handle(fiber.f()) ( state <- (
    prio=fiber.prio,
    runQ=runQ,
    startTime=clientTimeMilliseconds()
  ){
    case x ->
      print("RETURN");
      var nextPrio = getNextPrio();
      runNext(poll(state.runQ), nextPrio)
    case <GetPrio => resume> -> resume(state.prio, state)
    case <SetPrio(p) => resume> ->
      var q = fiberEnqueue(
        newPrioToFiber(resume, p),
        state.runQ
      );
      var nextPrio = getNextPrio();
      runNext(poll(q), nextPrio)
    case <Fork(f) => resume> ->
      print("fork");
      var q = poll(state.runQ);
      # f is new fiber, state.prio is the current fiber
      switch ((f.prio, state.prio)) {
        case (High, _) ->
          # Enqueue the current fiber, run the new fiber
          print("High fiber enqueued");
          var qq = fiberEnqueue(
            resumptionToFiber(resume, state.prio),
            q
          );
          runFiber(f, qq)
        case (Medium, High) ->
          print("Medium high fiber enqueued");
          # Enqueue the new fiber, run the current fiber
          var qq = fiberEnqueue(
            fiber0ToFiber(f),
            q
          );
          resume((), (state with runQ = qq))
        case (Medium, _) ->
          print("Medium fiber enqueued");
          # Enqueue the new fiber, run the current fiber

```

```

        var qq = fiberEnqueue(
            resumptionToFiber(resume, state.prio),
            q
        );
        runFiber(f, qq)
    case (_, High) ->
        # Enqueue the new fiber, run the current fiber
        print("Low High fiber enqueued");
        var qq = fiberEnqueue(
            fiber0ToFiber(f),
            q
        );
        resume((), (state with runQ = qq))
    case (_, Medium) ->
        # Enqueue the new fiber, run the current fiber
        print("Low Medium fiber enqueued");
        var qq = fiberEnqueue(fiber0ToFiber(f), q);
        resume((), (state with runQ = qq))
    case (_,_) ->
        # Enqueue the current fiber, run the new fiber
        print("Low fiber enqueued");
        var qq = fiberEnqueue(
            resumptionToFiber(resume, state.prio),
            q
        );
        runFiber(f, qq)
    }
    case <Yield => resume> ->
        var currentTime = clientTimeMilliseconds();
        if (currentTime - state.startTime <= bufferTime) {
            resume((), state)
        } else{
            var nextPrio = getNextPrio();
            var q = fiberEnqueue(
                resumptionToFiber(resume, state.prio),
                state.runQ
            );
            runNext(poll(q), nextPrio)
        }
    }
}

runFiber(fib, emptyPrioQueue())
}

```