

# Type-Safe Transactional Monadic IO in Java 17

Leonid Meshcheriakov

Innopolis University, Innopolis, Tatarstan Republic, Russia

**Abstract.** Most popular Java approaches rely on imperative constructs, exception handling with try-catch blocks, and transaction frameworks. Many of them are thread-bound, which create challenges for asynchronous workflows and system composability. Inspired by functional languages like Haskell and Scala, we adapt the IO monad abstraction to pure Java using modern language features including sealed interfaces and records and add support for transaction management in this abstraction. Our implementation provides a declarative, type-safe mechanism for sequencing computations with side effects through monadic composition, structured exception flow control, and transaction management powered by lifecycle hooks. The library also supports inner transactions, the Saga pattern, and asynchronous computations. The final implementation shows that functional concepts such as the IO monad can be seamlessly integrated into Java 17 without custom runtime or syntax, making code more composable with preserved type safety.

**Keywords:** Java · Side Effects · Functional Programming · IO Monad · Transactions.

## 1 Introduction

In modern software engineering, reliable management of side effects and transactions is a fundamental requirement for building scalable and maintainable systems. Java, while widely used in enterprise systems, relies primarily on imperative constructs and framework-level mechanisms for handling side effects and transactions, such as exceptions and annotations. Moreover, widely used Java frameworks for transactions, such as Spring Transaction Management or Java Transaction API, are thread-bound, meaning that a transaction initiated in one thread cannot be safely continued, committed or rolled back in another. This limitation makes it hard to design asynchronous or parallel transactional workflows.

In contrast, functional programming languages such as Haskell and Scala have addressed this problem through the abstraction of the IO monad, which provides a declarative and type-safe mechanism for sequencing computations with possible side effects, and the *Handle* pattern, where resource and context management lies outside the computation. These approaches allow developers to model side effects explicitly, improving reasoning about program behavior, composability, and testability.

Despite the proven advantages of monadic abstractions in functional programming, there is no native or standardized mechanism in Java for encapsulating and composing side effects in a type-safe and purely functional manner. Although some researchers have attempted to introduce algebraic effects and handlers in Java [1,2], these solutions typically rely on non-standard language extensions, complex syntax, or advanced functional abstractions, which increase the barrier to entry and limit their usability in practical applications.

This research explores whether the idea of the IO monad, as used in functional languages, can be effectively adapted to Java 17 to create a type-safe, composable, and transactional model for managing side effects.

The goal of the research is to design and implement a library for Java 17 that enables the representation of computations as declarative IO actions with explicit and composable error handling and support for asynchronous execution and transaction management. Moreover, the library should be user-friendly for the average Java developer; in other words, it should have the lowest possible barrier to entry.

The library is implemented using modern features from Java 17, such as sealed interfaces and records, which help model data in a type-safe and concise way. The design follows the principles of IO monads from Haskell and Scala but extends them with built-in transactional support.

To evaluate the library, unit tests will be used to check compliance with monad laws and correctness of transactional workflows. Additionally, a test program will be implemented to validate the functionality, correctness and performance of the library.

## 2 Literature Review

### 2.1 The Limitations of Exception and try-catch Mechanisms

Exception handling using try-catch blocks has been a key component of error management in object-oriented programming languages such as Java. However, many studies show that, although this mechanism aims to improve reliability and maintainability, in practice it often leads to hidden complexity and problems with modular design. The work by Miller and Tripathi [3] was one of the first to analyse this problem, in which the authors state that traditional exception handling conflicts with object-oriented principles such as abstraction, inheritance, and encapsulation. Exceptions are propagated implicitly, and they can be handled far from the point where they were thrown. This makes control flow less transparent and introduces dependencies between modules that are not visible in the program structure. As a result, it becomes harder to reason about system behavior or to modify and reuse components safely. [3] concludes that exception handling, although designed to improve fault tolerance, can in fact reduce robustness if it is not designed properly.

Empirical evidence supports these theoretical claims, for example, the authors of [4] found that in large software product lines exceptions flow frequently

goes beyond the scope of individual features and modules in ways that were not originally intended. Their study discovered that many exceptions are thrown by variable components, but caught in core parts of the system or sometimes not caught at all. This leads to unplanned exception propagation and unpredictable program behavior. Handlers are often generic, located far from the original source of the error, and therefore unable to recover properly. The authors conclude that try-catch mechanisms, when used in large modular systems, do not ensure the composability or predictability required for reliable software design. Moreover, in a large empirical study of Eclipse and Apache Tomcat projects, the authors analyzed 220 exception handling bugs and surveyed 154 professional developers [5]. They found that exception handling is often ignored in practice: only 27% of organizations reported having exception handling policies and standards, and 70% had no specific tests for exception handling code. Common problematic practices were overly generic error handlers or empty `catch` blocks, which made systems more fragile and harder to debug.

## 2.2 Monads and Side Effects

Studies [6,7] show that functional concepts, especially monads, can be effectively integrated into object-oriented and imperative languages. The authors of [6] focus on representing effectful computations with type-level monadic structures that make side effects explicit. Their library translates imperative constructs into typed bind chains, ensuring that each effect of `Option`, `State`, `Writer`, and `List` is safely composed. In [7] the Funk library was implemented in C#, introducing types such as `Maybe`, `OneOf`, and `Exc` to express optionality, branching, and exceptions as explicit monadic values. Both approaches prove that effect-safe functional abstractions can coexist with object-oriented design. This integration can improve code composability, clarity, and reliability and reduce boilerplate and failures by forcing developers to handle side effects properly.

Simon Peyton Jones [8] explains how Haskell deals with real-world side effects such as IO, concurrency, exception handling, and foreign language calls, while keeping the language pure and functional. The author shows that monads make it possible to describe side effects as typed values, so that programs stay predictable and composable. The IO monad is used to control execution order and isolate side effects inside a safe type structure. The paper also shows how monads make concurrent and exception-handling code easier to reason about.

In [9], the author presents the theoretical idea that exception handling can be seen as a special case of monadic binding, not as separate control-flow construct. He shows that the usual `catch` operation works in the same way as the monadic `bind` when it is used with an error monad. This means that exceptions can be handled as first-class values inside a typed computational context.

## 3 Design

To be more user-friendly for the average Java developer, the Java standard API and `Mono` from Project Reactor [10]—one of the widely used Java libraries among

those that adapt functional programming concepts—were considered while designing API.

### 3.1 Constraints of Java Type System

Java does not support higher-kinded types, but they can be introduced in type-safe way [12]. However, this approach makes the library more complex and harder to understand, which increases the barrier to entry. Moreover, this approach leads to boilerplate code and a loss of type information: if an operation is defined generically over all monads, its result will be only generic kind without type-specific methods, such as error recovery from the IO. In this work, the focus is limited to the implementation of the IO monad, because a separate library for a general Monad abstraction introduces more drawbacks than benefits.

### 3.2 Exceptions as IO Error

According to Java language standards [13], developers should throw `Exception` and its subclasses only to indicate that a problem has occurred. Therefore, we can consider an exception as an error state of our monad, that can be recovered from. If an exception is thrown inside a transformation function, it will be also treated as recoverable error state of the IO monad, ensuring that exception do not propagate outside.

### 3.3 Support of Asynchronous Computation

Modern applications often require asynchronous computations to handle I/O operations, external API calls, and concurrent tasks efficiently without blocking threads. To support such workflows, the library should rely on `CompletableFuture`—the standard Java API for asynchronous computations—so that the developer can seamlessly integrate existing libraries. This can be achieved by representing the result of each computation step as a `CompletableFuture`, to which subsequent operations are then composed. Additionally, to give developers control over thread management, it is important to provide the ability to explicitly specify which `Executor` should be used to run asynchronous computations when executing an IO.

### 3.4 Debugging Support

Since IO computations are lazy and asynchronous, standard Java stack traces become uninformative—they capture only the point where the computation was executed, such as a thread pool worker, rather than where the IO was actually constructed. To solve this problem, the library should capture stack traces at each IO step, so that the developer can trace the entire composition chain and identify which specific operation failed and how it was built.

### 3.5 Lack of Syntactic Sugar

Haskell has `do`-blocks, just like Scala `for`-comprehension, which make code more concise and less overloaded, when working with monads. Basically, these constructs are syntactic sugar over the `bind` operator. This kind of functionality is important for writing concise code. Since the library uses the standard Java without syntax changes, this problem can be solved by creating static methods.

Also, while designing the library, boilerplate code was observed when specifying the `IO` type with its parameters—context, error and result. Not only did the type itself become overly long and impractical to use, but it also needed to be repeated in multiple locations. Although this problem can be avoided by using the `var` keyword instead of a type, so it will be inferred, but it can be applied only to variables.

This problem can be solved in Haskell and Scala through type aliasing mechanism, which is not available in Java. However, it can be implemented in at least two ways. First, the developer defines a type alias via an interface, for example: `MyIO<R> extends IO<Nothing, Exception, R>`, and uses it in the code. Second, either the code is transformed at compile time by substituting `IO` instead of these user-defined interfaces, or the library generates classes implementing these interfaces at runtime. The first approach provides better runtime performance, but its implementation is significantly more complex than the second approach. This complexity arises from the requirement for advanced static analysis and code transformation. Moreover, it is not clear whether it will provide a significant performance improvement compared to the second solution. Hence, the second approach was chosen.

### 3.6 Transactions

Transaction handling should be designed so that transaction handlers, which configure environment, are defined in the same place as the code that requires it, rather than elsewhere. Therefore, the end user does not need to know, which implementation of service is used, to handle transactions correctly, which aligns with the object-oriented principle of inheritance. This approach enables easy combination of different modules with service implementations, reduces complexity and simplifies testing. Additionally, there should be support for isolated transaction for workflows that require it.

This behavior can be achieved by allowing the attachment of transactional hooks to an `IO` instance. The attached hooks should be executed before computation of that particular `IO` starts and at the end of its computation, providing information about whether it completed successfully. To properly isolate one transaction from another, their hooks must be handled separately from each other.

Unfortunately, some systems lack transaction mechanisms, for example, API calls, or cannot rely on them due to various reasons, such as the absence of atomicity, consistency, isolation or durability (ACID) guarantees. For such scenarios, there should be support for the Saga pattern [11]—a transaction pattern

that maintains data consistency through a sequence of local transactions. In this pattern, the developer can define a compensating action for each step that will be executed on failure to undo the effects of previously completed steps. This approach ensures consistency even when rollback mechanisms are not available.

### 3.7 Stack Overflow Problem

The most popular IO implementations in Scala—ZIO [14] and Cats Effect[15]—use a trampolining technique, where computations are executed step by step in a loop instead of recursive function calls, preventing stack overflow. Stack safety of asynchronous computations is ensured by storing the continuation and resuming it later via an event loop rather than through recursive calls.

In practice, Java programs are typically written in an object-oriented and imperative style, where stack overflow rarely occurs and usually happens only when recursion is used incorrectly. Moreover, we need to work with `CompletableFuture` that is not stack-safe, so that the developer can integrate libraries based on it. Due to these factors, this implementation was not designed to be stack-safe.

## 4 Implementation

We present results of our implementation<sup>1</sup> of the IO monad. This implementation demonstrates the feasibility of combining functional programming concepts with transaction management in a type-safe manner using an object-oriented language. While the core monadic operations and basic transactional management are in place, a number of advanced features are still under development.

### 4.1 Core Design

The `IO<C, E extends Exception, R>` interface defines a computation that is parameterized by a context type `C`, an exception type `E`, and a result type `R`.

This interface defines three core abstract methods that must be implemented: the `prepare` method simply returns an instance of `IORunnable` that represents the computation, the `getHooks` method returns the list of transactional hooks attached to this instance, and the `getInitializationTrace` method returns the stack trace of this instance initialization, needed for easing debugging. To run the computations, one of the execute methods, such as `tryExecute` or `tryExecuteAsync`, should be used.

```

1 public interface IO<C, E extends Exception, R> {
2     // ...
3     @NotNull
4     IORunnable<C, R> prepare(@NotNull C context);
5

```

---

<sup>1</sup> The implementation and documentation are available at <https://gitlab.com/worldm/functional/io>.

```

6  @NotNull
7  RecList<IOHook<? super C>> getHooks();
8
9  @NotNull
10 RecList<IOInitializationTrace> getInitializationTrace();
11
12 @NotNull
13 default Try<Exception, R> tryExecute(
14     @NotNull C Context
15 ) {
16     // ...
17 }
18
19 @NotNull
20 default CompletableFuture<Try<Exception, R>>
21     tryExecuteAsync(
22         @NotNull C Context
23 ) {
24     // ...
25 }
26 }

```

Besides these methods, the `IO` interface provides a comprehensive set of default methods for monadic composition, exception flow control, the type aliasing mechanism, transaction management, and utility operations.

Also, the library provides several standard implementations such as `IO.success(R result)` and `IO.error(E error)` for pure values, `IO.of(Supplier<R> supplier)` for simple effectful operations, `IO.fromCompletionStage(...)`, and so on.

```

1 IO<Nothing, Exception, String> io1 = IO.success("abc");
2 // Java allows type inference via "var" keyword.
3 // Type of io2 is IO<Object, RuntimeException, Long>
4 var io2 = IO.of(
5     () -> System.currentTimeMillis()
6 );

```

## 4.2 Monadic Composition

The implementation satisfies monad laws through the `map` and `flatMap` methods. The `map` method simply transforms results, while the `flatMap` method creates a sequence of `IO` computations. Method signatures ensure compatibility with other context types and allow flexible result type transformations. Moreover, it is impossible to compose `IO` instances with different context types directly, it is required to explicitly change the context through the `mapContext` method. If any of these methods produce an exception inside an user-defined transformation function, it will be caught, and the `IO` computation will be in a recoverable error state.

```

1 IO<Nothing, Exception, String> io1 = IO.success("abc");
2 IO<Nothing, Exception, Integer> io2 = io1.map(
3     s -> s.length()
4 );
5 // Type inference works even on transformations.
6 // Type of io3 is IO<Nothing, Exception, Integer>
7 var io3 = io2.flatMap(
8     i -> IO.success(i + 1)
9 );
10 IO<Connection, Exception, Integer> io4 = io3.mapContext(
11     c -> Nothing.INSTANCE
12 );

```

### 4.3 Exception Flow Control

Structured exception flow control is provided through three methods: `recover` converts exceptions into successful results, `flatRecover` allows exception handlers to return an `IO` computation for complex recovery strategies, and `mapError` transforms exception types for adapting to different error contexts. The type system tracks exception types through composition chains, providing compile-time safety.

```

1 IO<Nothing, SQLException, String> io1 = IO.error(
2     new SQLException("abc")
3 );
4 IO<Nothing, SQLException, String> io2 = io1.recover(
5     e -> e.getMessage()
6 );
7 IO<Nothing, SQLException, String> io3 = io1.flatRecover(
8     e -> new Random().nextBoolean()
9         ? IO.success(e.getMessage())
10        : IO.error(e)
11 );
12 IO<Nothing, Exception, String> io4 = io1.mapError(
13     e -> (Exception) e
14 );

```

### 4.4 Asynchronous Computations

The implementation is built on top of `CompletableFuture` to support asynchronous computations, allowing the wrapping of existing asynchronous code based on the standard Java API as `IO` computations, thus gaining the benefits of this library, such as transaction management. It is worth noting that there is a simple optimization: if an `IO` is executed asynchronously, the entire computation will be performed in the same thread that called the execution until it reaches the first asynchronous computation. Therefore, if an `IO` has no asynchronous operations, the entire computation runs synchronously in the same thread.

```

1 IO<ApiClient, Exception, String> io = IO.fromCompletionStage
  (context -> {
2   String password = context.getApiPassword();
3
4   CompletableFuture<String> future = apiService.getToken(
5     password);
6
7   return future;
8 });

```

Additionally, the developer can specify the `Executor` that should be used to run asynchronous computations by providing an execution context that implements the `ExecutionContext` interface.

```

1 class MyContext implements ExecutionContext {
2   private final Executor executor;
3
4   public MyContext(@NotNull Executor executor) {
5     this.executor = executor
6   }
7
8   @Override
9   @Nullable
10  public Executor getExecutor() {
11    return executor;
12  }
13
14  @NotNull
15  public ExecutionContext withExecutor(
16    @Nullable Executor executor
17  ) {
18    return new MyContext(executor);
19  }
20 }

```

```

1 IO<Object, Exception, String> io = IO.success("abc")
2   .flatMap(s -> IO.fromCompletionStage(...))
3   .map(o -> o.toString());
4
5 Executor executor = Executors.newSingleThreadExecutor();
6
7 // Asynchronous computation will use the executor,
8 // provided to the context
9 Try<Exception, String> result = io.tryExecute(
10   new MyContext(executor)
11 );

```

#### 4.5 Type Aliasing

The implementation contains type aliasing mechanism that allows the use of domain-specific type definitions via interfaces. The `withTypeAlias` method adapts

IO to the user-defined interface type without losing type-safety, allowing the developer to write concise code. Additionally, the developer can override methods to change the type of the return value to the corresponding type alias. All type constraints work at compile time; hence, the reliability of the type aliasing mechanism is significantly increased.

```

1 interface MyIO<R> extends IO<Nothing, Exception, R> {
2     @Override
3     default <NR> MyIO<NR> map(
4         CheckedFunction<
5             ? extends SQLException,
6             ? super R,
7             ? extends NR
8         > function
9     ) {
10         return IO.super.map(function)
11             .withTypeAlias(new TypeRef<>() {
12                 });
13     }
14 }

1 IO<Nothing, Exception, String> io1 = IO.success("abc");
2 MyIO<String> io2 = io1.withTypeAlias(new TypeRef<>() {});
3 MyIO<Integer> io3 = io2.map(s -> s.length());
4 IO<Nothing, Exception, Integer> io4 = io3;

```

## 4.6 Transaction Management

Transaction lifecycle hooks (onStart, onEnd, onException) attach to IO instances by using the addHook method and are preserved through monadic composition. It is guaranteed that a hook will be considered only once if it correctly defines a key—a combination of objects that uniquely identifies this hook. Moreover, the library ensures correct handling of hooks after IO instance transformations. The current implementation supports single-phase transactions with proper resource management and commit-rollback logic. Additionally, the library supports nested transactions by introducing the isolate method that runs a new transaction independently of the current one.

```

1 class SqlHook implements IOHook<Connection> {
2     @Override
3     public void onStart(@NotNull Connection connection) {
4         connection.setAutoCommit(false);
5     }
6
7     @Override
8     public void onEnd(@NotNull Connection connection) {
9         connection.commit();
10        connection.close();
11    }

```

```

12
13     @Override
14     public void onException(@NotNull Connection connection) {
15         connection.rollback();
16         connection.close();
17     }
18
19     @NotNull
20     public IOHookKey getKey() {
21         // Hook does not have any arguments or dependencies,
22         // so the key is just the class
23         return new IOHookKey(getClass());
24     }
25 }

1 IO<Connection, SQLException, Unit> io1 = IO.of(connection ->
2 {
3     Statement statement = connection.prepareStatement(
4         "INSERT INTO users(name, age) VALUES (?, ?)"
5     );
6     statement.setString(1, "Ivan");
7     statement.setInt(2, 10);
8     statement.executeUpdate();
9     return Unit.INSTANCE;
10 }).addHook(new SqlHook());
11
12 IO<Nothing, Exception, Unit> io2 = IO.success(Unit.INSTANCE)
13     .flatMap(unit ->
14         io1.isolate(() ->
15             DriverManager.getConnection("..."))
16     );
17
18 Try<Exception, Unit> result = io2.tryExecute(
19     Nothing.INSTANCE
20 );

```

## 4.7 Saga Pattern

The library provides the Saga pattern for systems, where resources lack proper transactional support. The developer can use the `compensate` method to associate compensating actions with an `IO` instance, so that when a failure occurs, they execute automatically in reverse order, ensuring consistency across the system.

```

1 Map<Integer, String> storage = new ConcurrentHashMap<>();
2 AtomicInteger lastId = new AtomicInteger();
3
4 IO<Nothing, Exception, String> io = IO.defer(() -> {

```

```

5     int id = lastId.incrementAndGet();
6     return IO.of(() -> {
7         String data = "data";
8         storage.put(id, data);
9         return data;
10    }).compensate(IO.of(() -> {
11        storage.remove(id);
12    }));
13 });

```

#### 4.8 Debugging Support

In the library, each `IO` instance captures the stack trace of its initialization, maintaining a recursive list of the full transformation history. When an `IO` computation raises an exception, these initialization stack traces are automatically filtered and attached as suppressed exceptions, allowing developers to trace the failure. Moreover, developers can control what should be filtered out by providing an execution context that implements the `InitializationTraceContext` interface.

```

1 class MyContext implements InitializationTraceContext {
2     @Override
3     @NotNull
4     public IOTraceIgnoreTrie getStackIgnoreTrie() {
5         return new IOTraceIgnoreTrie()
6             .addStartsWith("java.util.concurrent.");
7     }
8
9     @Override
10    public int getInitializationTraceDepth() {
11        return 20;
12    }
13 }

1 IO<Object, Exception, String> io = IO.success("abc")
2     .map(s -> s.repeat(2))
3     .flatMap(s -> IO.error(new Exception));
4
5 // Contains an suppressed exception with information
6 // about the initialization trace of the IO
7 Try<Exception, String> result = io.tryExecute(
8     new MyContext()
9 );

```

#### 4.9 Composition Utilities

The library provides several utilities to simplify composition and transformation of `IO` instances, improving code clarity and maintaining type safety.

**For-Comprehension** The `IODo` class provides Scala-like for-comprehension syntax to simplify the composition of multiple `IO` instances. It allows developers to express sequential operations in a declarative style, reducing syntactic overhead while maintaining type safety and monadic semantics.

```

1 IO<Nothing, Exception, String> io = IO.success("a");
2
3 IO<Nothing, Exception, List<String>> result = IODo.of(
4   io,
5   s -> IO.success(1),
6   (s, i) -> IO.success(Collections.nCopies(i, s))
7 ).yield((s, i, l) -> l);
```

**Result Aggregation** The `IOZip` class provides the ability to execute multiple `IO` computations in sequence and zip their results together into a single value in a type-safe manner. It provides a more concise alternative to `IODo` when the goal is simply to collect and combine multiple independent results.

```

1 IO<Nothing, Exception, String> io1 = IO.success("a"),
2 IO<Nothing, Exception, Integer> io2 = IO.success(1),
3 IO<Nothing, Exception, Long> io3 = IO.success(2L),
4
5 IO<Nothing, Exception, List<String>> result = IOZip.of(
6   io1,
7   io2,
8   io3
9 ).yield((s, i, l) -> new MyEntity(s, i, l));
```

**Monadic Transformations** The `IOOps` class provides static methods for several monadic operations. It includes the `Optional` monad transformer for concise handling of nullable values, as well as the `sequence` method that converts a collection of `IO` instances into a single `IO` instance that contains a collection of the results.

```

1 List<IO<Nothing, Exception, String>> list = List.of(
2   // ...
3 );
4
5 IO<Nothing, Exception, List<String>> io1 = IOOps.sequence(
6   list
7 );
8
9 IO<Nothing, Exception, Optional<String>> io2 = IO.of(
10  Optional.of("a")
11 );
12
13 IO<Nothing, Exception, Optional<String>> io3 =
14  IOOps.flatMapOptional(
```

```

15           io2,
16           string -> IO.success(string + "b")
17     );

```

## 4.10 Extension Libraries

To demonstrate the practical applicability and extensibility of the IO library, two extension libraries were created: a Java Database Connectivity (JDBC) integration layer—JDBC is the standard Java API for working with databases—and a remote procedure call framework supporting distributed transactions.

**io-jdbc** This library<sup>2</sup> allows developers to easily wrap actions with database based on JDBC to the IO instances with transaction support:

```

1 record User(int id, @NotNull String name, int age) {
2
3 }
4
5 String name = "foo";
6
7 SqlIO<SqlContext, List<User>> io1 = JdbcQuery.select(
8         "SELECT id, age FROM users WHERE name = ?"
9         )
10        .argument(name)
11        .build(new ListOutcome<>(
12            resultSet ->
13            new User(
14                resultSet.getInt("id"),
15                name,
16                resultSet.getInt("age")
17            )
18        )
19    );
20
21 IO<SqlContext, SQLException, List<User>> io2 = io1;

```

**io-rpc** This library<sup>3</sup> allows the developer to delegate execution of an IO computation to a remote server in a single distributed transaction via any backend, that the user can implement. The library contains an implementation for HTTP and AMQP backends. Distributed transactional behavior is achieved through the following steps:

---

<sup>2</sup> The implementation and documentation are available at <https://gitlab.com/worldm/functional/io-jdbc>.

<sup>3</sup> The implementation and documentation are available at <https://gitlab.com/worldm/functional/io-rpc>.

1. The client sends the first RPC request for this server in transaction A.
2. The server stores the already existing context if the first RPC request in this transaction was sent by this server; otherwise, it creates and stores a new one.
3. The server calls the method with arguments, provided by the client, runs the received IO and sends the result to the client.
4. The server asynchronously waits for another RPC request or a request that indicates whether the transaction should be rolled back or committed.
5. If waiting exceeds the timeout, the transaction will be rolled back.

Moreover, multiple parallel distributed transactions on one server instance are supported.

Unfortunately, there is a case that can lead the system to an inconsistent state. If the timeout is exceeded at one of the servers and it then receives a commit request, the request will fail. However, it will not affect commit requests to other servers, because in the current implementation the IO library ignores exceptions from transactional hooks. To handle such situations correctly, a two-phase commit protocol must be implemented in the future.

```

1 // API module, shared between the server and its clients
2 public interface HashService {
3     @NotNull
4     IO<Nothing, Exception, String> hash(
5         @NotNull String s
6     );
7 }
8
9 // Server code
10 // -- Configuration
11 ServerRpcIO<Nothing> serverRpcIO = new ServerRpcIO<>(
12     new ServerId("my-service"), // Server ID
13     () -> Nothing.INSTANCE, // Context factory
14     Duration.ofSeconds(5) // Transaction timeout
15 );
16
17 serverRpcIO.addCallHandler(
18     HashService.class,
19     new Sha256HashService() // Implementation
20 );
21
22 // -- Usage
23 String rpcJsonFromClient = "...";
24
25 RpcClientMessage<?> request = serverRpcIO
26     .get.ObjectMapper()
27     .readValue(
28         rpcJsonFromClient,
29         RpcClientMessage.class
30     );

```

```
23 CompletableFuture<RpcResponseMessage<RpcServerMessage>>
24         response = serverRpcIO.handle(request);
25
26
27 // Wait for response, serialize it and send
28 // to the client
29
30 // Client side
31 // -- Configuration
32 RpcBackend rpcBackend = ...;
33
34
35 ClientRpcIO<Nothing> clientRpcIO = new ClientRpcIO(
36         rpcBackend,
37         Nothing.class // Context class
38 );
39
40
41 HashService remoteHashServer = clientRpcIO.implement(
42         HashService.class
43 );
44
45
46 // -- Usage
47 String hashedData = IO.success("data")
48         .flatMap(p -> remoteHashServer.hash(p))
49         .execute(Nothing.INSTANCE);
```

## 4.11 Constraints

The current implementation has several constraints that must be considered when using the library.

As discussed in Section 3.7, the implementation is not stack-safe. This design focuses on practical needs: stack overflow rarely happens in typical Java code, and compatibility with the existing standard Java API for asynchronous computations—which itself is not stack-safe—is required, since there is a huge code base built on top of it.

All methods enforce non-null semantics: null values are forbidden in all arguments and results in transformation functions even for the `map` method. This is achieved Developers must use `Optional` types when representing absent values.

Exception handling currently captures only `Exception` and its subclasses, meaning other `Throwable` types such as `Error` are not caught and will propagate unchecked. While this aligns with standard Java practices, where `Error` indicates a state where it is impossible to recover, some third-party libraries inappropriately use `Error` subclasses for recoverable failures. To properly support badly designed libraries, we plan to refactor the library to handle all `Throwable` types while maintaining appropriate semantics for unrecoverable errors.

The transaction hook system has a critical limitation in error handling: exceptions thrown by `onEnd` and `onException` hooks are only logged and not handled properly. This creates a problem with consistency when several hooks

manage different resources. For example, if a file system and an SQL connection hooks have been initialized, and one fails during commit, the other proceeds independently, potentially creating an inconsistent state. To address this issue, there must be support for two-phase commits to ensure the atomicity principle, which is our primary focus.

To change or define the behavior of the effects, the code should rely on the execution context provided by the developer. However, due to the inheritance principle, developers should declare the most general context possible, without exposing implementation details. For example, if we have a `UserRepository` service with in-memory and database implementations, the context type of the interface cannot be `SqlContext`. Otherwise, the code becomes non-extensible: for instance, it is impossible to create an implementation for MongoDB without using type-unsafe context casts. Therefore, the developer should use parameterized types for the context in all abstractions and implementations to ensure fully type-safe operations. It is worth noting that in Java it is possible to obtain a union type only by using parameterized types. This approach is demonstrated in the code example below.

```

1 // Service and context definitions
2 interface UserRepository<C> {
3     @NotNull
4     IO<C, Exception, User> getUser(int id);
5 }
6
7 interface SqlContext {
8     // ...
9 }
10
11 class PostgresUserRepo implements UserRepository<SqlContext> {
12     // ...
13 }
14
15 class MemoryUserRepo implements UserRepository<Object> {
16     // ...
17 }
18
19 interface PermissionService<C> {
20     @NotNull
21     IO<C, Exception, Boolean> isAdmin(int id);
22 }
23
24 interface PermissionContext {
25     // ...
26 }
27
28 class UserPermissionService<
29     // Union type
30     C extends URC & PermissionContext,
31     // UserRepository context type

```

```

32         URC
33 > implements PermissionService<C> {
34
35     private final UserRepo<URC> repo;
36
37     public UserPermissionService(
38         @NotNull UserRepo<URC> repo
39     ) {
40         this.repo = repo;
41     }
42
43     // ...
44 }

1 // Configuration
2 interface CombinedContext extends
3     SqlContext,
4     PermissionContext {
5 }
6
7 UserRepo<SqlContext> repository =
8     new PostgresUserRepo();
9
10 PermissionService<CombinedContext> service =
11     new UserPermissionService<>(repository);

```

Unfortunately, there is one problem with contexts: the user can use a final class as a context type definition, which limits the ability to handle effects, since these classes cannot be overridden or extended with interface implementations. Addressing this limitation is a priority for future work.

## 4.12 Conclusion

In this work, the library for the functional IO monad with transaction support was implemented in Java 17, while preserving all the benefits of this approach, including type-safety, side effects handling, and robust workflow. The current implementation provides a solid foundation, but further work remains to be done to achieve more robust workflow.

## References

1. Brachthäuser, J., Schuster, P. & Ostermann, K. Effect handlers for the masses. *Proc. ACM Program. Lang.* **2** (2018,10), <https://doi.org/10.1145/3276481>
2. T. Mahler. Jiffy, <https://github.com/thma/jiffy>, Accessed: 2026-01-11
3. Miller, R. & Tripathi, A. Issues with exception handling in object-oriented systems. *ECOOP'97—Object-Oriented Programming*. pp. 85-103 (1997), <https://doi.org/10.1186/2195-1721-1-3>

4. Melo, H., Coelho, R., Kulesza, U. & Sena, D. In-depth characterization of exception flows in software product lines: an empirical study. *Journal Of Software Engineering Research And Development*. **1**, 3 (2013,10), <https://doi.org/10.1186/2195-1721-1-3>
5. Ebert, F., Castor, F. & Serebrenik, A. An exploratory study on exception handling bugs in Java programs. *Journal Of Systems And Software*. **106** pp. 82-101 (2015), <https://www.sciencedirect.com/science/article/pii/S0164121215000862>
6. Anderlind, J. & Åsberg, M. Monadic Programming in Imperative Languages. (2023), <https://hdl.handle.net/20.500.12380/306361>
7. Čerim, H. Extending C# with a Library of Functional Programming Concepts. (Univerzita Karlova, Matematicko-fyzikální fakulta,2020), <https://hdl.handle.net/20.500.11956/121344>
8. Peyton Jones, S. Tackling the Awkward Squad: monadic input/output, concurrency, exceptions, and foreign-language calls in Haskell. (2002), <https://www.microsoft.com/en-us/research/wp-content/uploads/2016/07/mark.pdf>
9. Malakhovski, J. Exceptionally Monadic Error Handling. (2019), <https://arxiv.org/abs/1810.13430>
10. VMWare. Project Reactor, <https://projectreactor.io/>, Accessed: 2026-01-11
11. Garcia-Molina, H. & Salem, K. Sagas. *ACM Sigmod Record*. **16**, 249-259 (1987), <https://doi.org/10.1145/38714.38742>
12. Higher-Kinded-J Maintainers. Higher-Kinded-J, <https://higher-kinded-j.github.io>, Accessed: 2026-01-11
13. Oracle Corporation. Throwing Exceptions, <https://dev.java/learn/exceptions/throwing/>, Accessed: 2026-01-11
14. ZIO Maintainers. ZIO Documentation, <https://zio.dev/reference>, Accessed: 2026-01-11
15. Typelevel. Cats Effect Documentation, <https://typelevel.org/cats-effect/docs/getting-started>, Accessed: 2026-01-11