

# Failboxes: Provably Safe Exception Handling<sup>★</sup>

Bart Jacobs<sup>\*\*</sup> and Frank Piessens

Department of Computer Science, Katholieke Universiteit Leuven, Belgium  
{bart.jacobs, frank.piessens}@cs.kuleuven.be

**Abstract.** The primary goal of exception mechanisms is to help ensure that when an operation fails, code that depends on the operation's successful completion is not executed (a property we call *dependency safety*). However, the exception mechanisms of current mainstream programming languages make it hard to achieve dependency safety, in particular when objects manipulated inside a try block outlive the try block.

Many programming languages, mechanisms and paradigms have been proposed that address this issue. However, they all depart significantly from current practice. In this paper, we propose a language mechanism called failboxes. When applied correctly, failboxes have no significant impact on the structure, the semantics, or the performance of the program, other than to eliminate the executions that violate dependency safety. Specifically, programmers may create failboxes dynamically and execute blocks of code in them. Once any such block fails, all subsequent attempts to execute code in the failbox will fail. To achieve dependency safety, programmers simply need to ensure that if an operation  $B$  depends on an operation  $A$ , then  $A$  and  $B$  are executed in the same failbox. Furthermore, failboxes help fix the unsafe interaction between locks and exceptions and they enable safe cancellation and robust resource cleanup. Finally, the Fail Fast mechanism prevents liveness issues when a thread is waiting on a failed thread.

We give a formal syntax and semantics of the new constructs, and prove dependency safety. Furthermore, to show that the new constructs are easy to reason about, we propose proof rules in separation logic. The theory has been machine-checked.

## 1 Introduction

If a program is seen as a state machine, a programmer's job may be seen as writing code to deal with each of the states that the program may reach. However, programmer time is limited and some states are less likely to occur during production than others. Therefore, in many projects it is useful to designate the most unlikely states as *failure states* and to deal with all failure states in a uniform way, while writing specific code only for non-failure (or *normal*) states.

An extreme form of this approach is to simply ignore failure states and not care what the program does when it reaches a failure state (i.e., when it *fails*).

---

<sup>\*</sup> We used the term *subsystems* in preliminary work.

<sup>\*\*</sup> Bart Jacobs is a Postdoctoral Fellow of the Research Foundation - Flanders (FWO).

This is often what happens when subroutines indicate failure conditions as special return values, and programmers have no time to write code at call sites to check for them.

A major problem with this approach is that it is *unsafe*: a failure may lead to the violation of any and all of the program's intended safety properties. Specifically, the approach violates *dependency safety*, the property which says that when an operation fails, code that depends on the operation's successful completion is not executed.

To fix this, modern programming languages offer constructs that make it easy for programmers to indicate that a state is a failure state, and deal with failure states by terminating the program by default. The underlying assumption is that termination is always safe. For example, in Java, a failure state is indicated by throwing an unchecked exception. We will focus on the Java language in this paper; the related work section discusses other languages.

Whereas by default, when a program throws an exception it terminates immediately, the programmer can override this default through the use of try-catch statements and try-finally statements. Furthermore, in a multithreaded program, when a thread's main method completes abruptly (i.e., an exception was thrown and not caught during its execution), only that thread, not the entire program, is terminated. Also, when a synchronized block's body completes abruptly, the lock is released before the exception is propagated further.

These deviations from strict termination behavior are useful and are used for two reasons. Firstly, not all exceptions indicate failure. Sometimes, programmers throw and catch exceptions to implement the program's functional behavior. Typically, in Java, checked exceptions are used for this. Secondly, programmers sometimes wish to increase the program's robustness by not considering the program to be a single unit of failure but rather by identifying multiple smaller units of failure. Common examples are extensible programs, where poorly written or malicious plugins (such as applets or servlets) should not affect the base system; and command-processing applications (such as request-response-based servers, GUI applications, or command-line shells) where a failure during the processing of a command should simply cause an error response to be returned, while continuing to process other commands normally.

However, by continuing to execute after a failure, the risk of safety violations reappears. In particular, safety violations are likely if the code that fails leaves a data structure in an inconsistent state and this data structure is then accessed during execution of a finally block or after the exception is caught, or by another thread. In other words, there is a safety risk if a try block manipulates an object that outlives the try block. More generally, if we define *dependency safety* as the property that if an operation fails, no code that depends on the operation's successful completion is executed, then dependency safety may be violated if pieces of code outside a try block depend on particular pieces of code inside the try block either not executing at all or executing to completion successfully. This is the problem we address in this paper.

To remedy this, we propose a language mechanism called *failboxes*. Programmers may create failboxes dynamically and execute blocks of code in them. Once any such block fails, all subsequent attempts to execute code in the failbox will fail. To achieve dependency safety, programmers simply need to ensure that if an operation  $B$  depends on an operation  $A$ , then  $A$  and  $B$  are executed in the same failbox. Furthermore, failboxes help fix the unsafe interaction between locks and exceptions and they enable safe cancellation and robust resource cleanup. Finally, the Fail Fast mechanism prevents liveness issues in the presence of failure in cooperating concurrent computations.

Failboxes are very lightweight: a failbox can be implemented as an object with a boolean field indicating if the failbox has failed, and a parent pointer. Executing a code block in a failbox essentially means that before and after executing the block, the thread-local variable that designates the current failbox is updated, and before a failbox is made current, it is checked that it has not failed.

We give a formal syntax and semantics of the new constructs, and prove dependency safety. Furthermore, to show that the new constructs are easy to reason about, we propose separation logic proof rules and prove their soundness.

The rest of the paper is structured as follows. In Section 2, we illustrate the problem with an example and discuss existing approaches. In Section 3, we introduce failboxes. We show additional aspects and benefits of the approach for multithreaded programs in Section 4. Section 5 briefly discusses how the approach enables safe cancellation and robust compensation. To show that it is easy to reason about the new constructs, we propose separation logic proof rules for the envisaged usage patterns in Section 6. We end the paper with sections on implementation issues (Section 7), related work (Section 8), and a conclusion (Section 9).

The theory of this paper has been machine-checked using the Coq proof assistant [12].

## 2 Problem Statement

Consider the example program in Figure 1. It shows a program that continuously receives commands and processes them. The code for processing commands is not shown, except that it involves calls of *compute* and calls of *addEntry* on a *Database* object *db* that is shared across all command executions. If the processing of a command fails, e.g. because it requires too much memory, the exception is caught, an error message is shown to the user, and the next command is received.

This program is unsafe. Specifically, some executions of this program violate the intended safety property that at the start of each loop iteration, object *db* is *consistent*, i.e., satisfies the property that *count* is not greater than the length of *entries*. In particular, consider an execution where method *addEntry* is called in a state where *entries* is full. This means *count* equals *entries.length*. As a result, after incrementing *count*, *addEntry* will attempt to allocate a new,

```

class Database {
    int count;
    int[] entries := new int[10];
    /* invariant: count ≤ entries.length */
    void addEntry(int entry) {
        count++;
        if (count = entries.length + 1) {
            int[] es := new int[count * 2]; // *** A ***
            System.arraycopy(entries, 0, es, 0, entries.length);
            entries := es;
        }
        entries[count - 1] := entry; // *** B ***
    } ... }
class Program {
    public static void main(String[] args) {
        Database db := new Database();
        while (true)
            /* invariant: db is consistent */
            {
                String cmd := readCommand();
                try {
                    ... compute(cmd); ...
                    ... db.addEntry(...); ...
                } catch (Throwable e) { showErrorMessage(e); }
            } ... }
}

```

**Fig. 1.** An unsafe program

larger array. Now assume there is not enough memory for this new array and an *OutOfMemoryError* occurs at location *A*. At this point, *count* is greater than the length of *entries* and the *Database* object is inconsistent. Next, the exception is caught in method *main* and the loop is continued, violating the safety property.

Note: In this case, the safety violation results in an *ArrayIndexOutOfBoundsException* at location *B* in each subsequent call of *addEntry*; however, in general, safety violations might remain undetected and lead to data corruption, incorrect results, or sending incorrect commands to hardware devices.

The following approaches exist to deal with this complication:

- **Never catch unchecked exceptions.** Never catching unchecked exceptions makes it easier to preserve safety properties, since the many implicit control flow paths created by catching unchecked exceptions are avoided. However, catching unchecked exceptions can be useful, as in the example. Note also that **try-finally** blocks are equivalent to **try-catch** blocks that catch unchecked exceptions; specifically, assuming  $S_1$  does not jump out of

the **try** block, a statement

$$\mathbf{try} \{ S_1 \} \mathbf{finally} \{ S_2 \}$$

is equivalent to

$$\mathbf{try} \{ S_1 \} \mathbf{catch} (\textit{Throwable } t) \{ S_2 \mathbf{ throw } t; \} S_2$$

and is subject to the same complication:  $S_2$  might depend on the successful completion of certain sub-computations within  $S_1$ . Never catching unchecked exceptions would imply never using **try-finally** blocks, or modifying their semantics so that they ignore unchecked exceptions. The semantics of synchronized blocks would need to be updated similarly.

- **Always maintain consistency.** It is often possible to ensure that objects used across try-catch blocks, like the *Database* object in the example, are in a consistent state at all times. Often it is sufficient to reorder assignments; e.g., in the example, moving the count increment after the assignment to *entries* preserves consistency. Another approach is to use a functional programming-like approach, where a new object state is built up separately and then installed into the object using a single assignment. In the example, method *addEntry* would return a new *Database* object rather than updating the existing one. Yet another approach is to use transaction-like technologies, such as software transactional memory [19, 5]. However, these approaches either require the programmer to perform non-trivial additional reasoning and/or programming work, or impose a potentially significant performance overhead.
- **Never fail during critical sections.** It might be possible in some cases to guarantee absence of failure at points where failure would violate safety. This requires careful programming to avoid operations that might encounter resource or implementation limitations, such as heap or stack memory allocations or operations on bounded integers, or to move these operations out of the critical section. Furthermore, this might require virtual machine support if the virtual machine may perform resource allocations implicitly. For example, the .NET Framework’s JIT compiler may allocate memory at any time to store a newly compiled piece of code. Therefore, starting with version 2, the .NET Framework offers constructs to “prepare” a piece of code that must execute without failure [21]. However, this approach imposes a significant burden on the programmer.
- **Ensure dependent code is not executed.** In this approach, steps are taken to ensure that if a computation fails with an unchecked exception, then no computations that depend on the failed computation’s successful completion ever get to run. There are at least two ways to achieve this:
  - **Use separate threads.** In this approach, threads are adopted as the units of failure. Within a thread, unchecked exceptions are never caught; that is, an exception in the thread causes the entire thread to die. All data structures are local to threads. Instead of running a block of code in a try-catch block, it is run in a separate thread. During this time, the

original thread waits for the termination of the child thread; additionally, the original thread may accept messages on a message queue. If the child thread needs to perform an operation whose failure should cause the parent thread to fail (such as an *addEntry* call on the *Database* object), the child thread may perform a remote procedure call into the parent thread via the parent thread's message queue. This is more or less the approach used in operating systems, in the Erlang language [1], and in the SCOOP multithreading approach for Eiffel [17].

- **Guard dependent code manually.** The programmer can manually arrange to ensure that dependent code is not executed. For example, the programmer could associate a boolean flag with each object used across try-catch blocks that tracks whether the object is in a consistent state, and check this flag before accessing the object [13]. If the flag is false, an exception is thrown.

In this paper we present a new approach in the fourth category, which, like the use of separate threads and manually guarding dependent code, supports catching exceptions and does not require that consistency be maintained always or that failures be avoided, but which has less programming and run-time overhead than the use of separate threads and which has less programming overhead than manually guarding dependent code.

### 3 Failboxes

In our approach, the language is extended with a notion of *failboxes*. Constructs are added for creating a new failbox and for running a piece of code in a designated failbox. As soon as one such piece of code fails (i.e., completes abruptly with an unchecked exception), any subsequent attempt to run code in the failbox fails. To ensure dependency safety, the programmer simply needs to ensure that if a computation *B* depends on a computation *A*, then *A* and *B* run in the same failbox.

To facilitate composition of program modules, failboxes are ordered hierarchically. When creating a new failbox, a parent may be specified. If an exception occurs in a failbox, both it and its transitive children are marked as failed.

#### 3.1 Syntax and semantics

The syntax of the new constructs is as follows:

$$\begin{aligned}
 s ::= & \dots \\
 & | x := \mathbf{currentfb}; \quad | x := \mathbf{newfb}; \quad | x' := \mathbf{newfb}(x); \\
 & | \mathbf{enter} (x) \{ \bar{s} \} \mathbf{catch} \{ \bar{s}' \}
 \end{aligned}$$

where *s* ranges over statements,  $\bar{s}$  ranges over sequences (i.e., sequential compositions) of statements, and *x* and *x'* range over local variable names.

Note: For simplicity, we ignore checked exceptions and exception objects in the formal developments.

A program state is a tuple of the form

$$(L, \Sigma, \Phi, T)$$

where  $L$ , the *lock map*, is a partial function that contains a pair  $(o, t)$  if thread  $t$  holds the lock of object  $o$ ;  $\Sigma$  is a partial function that maps each allocated failbox to its parent failbox (and a root failbox to itself);  $\Phi$  is the set of *failed failboxes*; and  $T$  is a partial function that maps each thread to its current state. (We omit the heap since our constructs do not interact with it.)

A thread state is a tuple of the form

$$(f, V, \bar{s}, \bar{b}, \bar{F})$$

where  $f$  is the thread's *current failbox*,  $V$  is a total function that maps each variable name to a value,  $\bar{s}$ , the *continuation*, is the sequence of statements to be executed by the thread,  $\bar{b}$  is the sequence of *enclosing blocks*, and  $\bar{F}$  is the sequence of enclosing activation records.

The syntax of an enclosing block is as follows:

$$b ::= \mathbf{enter} (f) \mathbf{catch} \{ \bar{s} \} \bar{s}' \mid \mathbf{synchronized} (o); \bar{s}$$

where an enclosing *enter* block records the failbox  $f$  that was current prior to the *enter* statement (*not* the failbox that was entered), the *catch* block body  $\bar{s}$ , and the statements  $\bar{s}'$  that are to be executed after completion of the *enter* statement; and an enclosing *synchronized* block records the object  $o$  whose lock was acquired, and the statements  $\bar{s}$  that are to be executed after completion of the *synchronized* statement.

In the initial program state of a program with main method body  $\bar{s}$ , the lock map is empty, there is a single failbox  $f$ , whose parent is itself, no failbox is marked as failed, and there is one thread  $t$  whose current failbox is  $f$ ; all of the thread's local variables are bound to **null**, and it has no enclosing blocks and no enclosing activation records:

$$\frac{\text{main } \bar{s}}{\text{initial } (\emptyset, \{(f, f)\}, \emptyset, \{(t, (f, (\lambda x. \mathbf{null}), \bar{s}, \epsilon, \epsilon)\})}$$

The statement  $x := \mathbf{currentfb}$ ; assigns the current failbox to variable  $x$ :

$$\frac{\text{CURRENTFB} \quad (t, (f, V, x := \mathbf{currentfb}; \bar{s}, \bar{b}, \bar{F})) \in T}{(L, \Sigma, \Phi, T) \rightarrow (L, \Sigma, \Phi, T(t := (f, V(x := f), \bar{s}, \bar{b}, \bar{F})))}$$

The statement  $x := \mathbf{newfb}$ ; creates a new root failbox and assigns it to  $x$ :

$$\frac{\text{NEWFB-ROOT} \quad (t, (f, V, x := \mathbf{newfb}; \bar{s}, \bar{b}, \bar{F})) \in T \quad f' \notin \text{dom}(\Sigma) \quad \Sigma' = \Sigma(f' := f')}{(L, \Sigma, \Phi, T) \rightarrow (L, \Sigma', \Phi, T(t := (f, V(x := f'), \bar{s}, \bar{b}, \bar{F})))}$$

If  $x$  is bound to a failbox  $f'$  and  $f'$  is not marked as failed, the statement  $x' := \mathbf{newfb}(x)$ ; creates a new child failbox of  $f'$  and assigns it to  $x'$ :

$$\begin{array}{c} \text{NEWFB-CHILD} \\ (t, (f, V, x' := \mathbf{newfb}(x); \bar{s}, \bar{b}, \bar{F})) \in T \\ \hline V(x) = f' \quad f' \notin \Phi \quad f'' \notin \text{dom}(\Sigma) \quad \Sigma' = \Sigma(f'' := f') \\ \hline (L, \Sigma, \Phi, T) \rightarrow (L, \Sigma', \Phi, T(t := (f, V(x' := f''), \bar{s}, \bar{b}, \bar{F}))) \end{array}$$

If  $x$  is bound to a failbox  $f'$ , and  $f'$  is not marked as failed, then statement  $\mathbf{enter}(x) \{ \bar{s}' \} \mathbf{catch} \{ \bar{s}'' \}$  records the current failbox, the catch block body  $\bar{s}'$ , and the current continuation in a new enclosing block, makes  $f'$  the current failbox, and starts executing the enter block body  $\bar{s}'$ :

$$\begin{array}{c} \text{ENTER} \\ (t, (f, V, \mathbf{enter}(x) \{ \bar{s}' \} \mathbf{catch} \{ \bar{s}'' \} \bar{s}, \bar{b}, \bar{F})) \in T \\ \hline V(x) = f' \quad f' \notin \Phi \quad \bar{b}' = (\mathbf{enter}(f) \mathbf{catch} \{ \bar{s}'' \} \bar{s}) \cdot \bar{b} \\ \hline (L, \Sigma, \Phi, T) \rightarrow (L, \Sigma, \Phi, T(t := (f', V, \bar{s}', \bar{b}', \bar{F}))) \end{array}$$

On normal completion of an enter block body, the former current failbox is restored and the catch block is skipped, provided that the former current failbox is not marked as failed:

$$\begin{array}{c} \text{ENTER-COMPLETE-NORMAL} \\ (t, (f, V, \epsilon, (\mathbf{enter}(f') \mathbf{catch} \{ \bar{s}' \} \bar{s}'') \cdot \bar{b}, \bar{F})) \in T \quad f' \notin \Phi \\ \hline (L, \Sigma, \Phi, T) \rightarrow (L, \Sigma, \Phi, T(t := (f', V, \bar{s}'', \bar{b}, \bar{F}))) \end{array}$$

where  $\epsilon$  denotes the empty sequence.

We model the occurrence of an exception as the replacement of the current continuation with a throw statement. An exception can occur at any time; this reflects the fact that in Java a *virtual machine error* can be thrown at any time [10, §11.3.2].

$$\begin{array}{c} \text{FAIL} \\ (t, (f, V, \bar{s}, \bar{b}, \bar{F})) \in T \quad \bar{s} \neq \mathbf{throw}; \\ \hline (L, \Sigma, \Phi, T) \rightarrow (L, \Sigma, \Phi, T(t := (f, V, \mathbf{throw};, \bar{b}, \bar{F}))) \end{array}$$

If variable  $x$  is not bound to a failbox, or it is bound to a failbox but the failbox is marked as failed, then both  $x' := \mathbf{newfb}(x)$ ; and  $\mathbf{enter}(x) \{ \bar{s} \} \mathbf{catch} \{ \bar{s}' \}$  throw an exception (of type *FailboxException*); this is covered by rule FAIL.

On abrupt completion of an enter block body with an exception, the current failbox and its descendants are marked as failed, the former current failbox is restored, and the catch block is executed, provided the former current failbox is not marked as failed:

$$\begin{array}{c} \text{ENTER-COMPLETE-ABRUPT} \\ (t, (f, V, \mathbf{throw};, (\mathbf{enter}(f') \mathbf{catch} \{ \bar{s}' \} \bar{s}'') \cdot \bar{b}, \bar{F})) \in T \\ \hline \Phi' = \Phi \cup (\Sigma^{-1})^*(f) \quad f' \notin \Phi' \\ \hline (L, \Sigma, \Phi, T) \rightarrow (L, \Sigma, \Phi', T(t := (f', V, \bar{s}' \bar{s}'', \bar{b}, \bar{F}))) \end{array}$$



where  $(\Sigma^{-1})^*(f)$  denotes the set of  $f$ 's descendants, including  $f$  itself.

On normal completion of an enter block body, if the former current failbox is marked as failed, it is restored but the catch block is skipped and a *FailboxException* exception is thrown:

$$\frac{\text{ENTER-COMplete-NORMAL-FAIL} \\ (t, (f, V, \epsilon, (\mathbf{enter} (f') \mathbf{catch} \{ \bar{s}' \} \bar{s}'') \cdot \bar{b}, \bar{F})) \in T \quad f' \in \Phi}{(L, \Sigma, \Phi, T) \rightarrow (L, \Sigma, \Phi, T(t := (f', V, \mathbf{throw};, \bar{b}, \bar{F})))}$$

On abrupt completion of an enter block body with an exception, if after marking the current failbox as failed, the former current failbox is marked as failed, the former current failbox is restored but the catch block is skipped and a *FailboxException* exception is thrown:

$$\frac{\text{ENTER-COMplete-ABRUPT-FAIL} \\ (t, (f, V, \mathbf{throw};, (\mathbf{enter} (f') \mathbf{catch} \{ \bar{s}' \} \bar{s}'') \cdot \bar{b}, \bar{F})) \in T \\ \Phi' = \Phi \cup (\Sigma^{-1})^*(f) \quad f' \in \Phi'}{(L, \Sigma, \Phi, T) \rightarrow (L, \Sigma, \Phi', T(t := (f', V, \mathbf{throw};, \bar{b}, \bar{F})))}$$

### 3.2 Syntactic sugar

We remove try-catch statements and try-finally statements from the language as separate statements. Instead, we define them as syntactic sugar over the new constructs. Specifically, the statement

$$\mathbf{try} \{ \bar{s} \} \mathbf{catch} \{ \bar{s}' \}$$

is defined as

$$x := \mathbf{currentfb}; x' := \mathbf{newfb}(x); \mathbf{enter} (x') \{ \bar{s} \} \mathbf{catch} \{ \bar{s}' \}$$

where  $x$  and  $x'$  are fresh. That is, a try-catch statement executes the try block in a new child failbox of the current failbox.

The statement

$$\mathbf{try} \{ \bar{s} \} \mathbf{finally} \{ \bar{s}' \}$$

is defined as

$$\mathbf{try} \{ \bar{s} \} \mathbf{catch} \{ \bar{s}' \mathbf{throw}; \} \bar{s}'$$

This means that a try-finally statement executes its try block in a new child failbox of the current failbox.

Furthermore, we define the following shorthands:

$$\mathbf{enter} (x) \{ \bar{s} \} \equiv \mathbf{enter} (x) \{ \bar{s} \} \mathbf{catch} \{ \mathbf{throw}; \} \\ \mathbf{reenter} (x) \{ \bar{s} \} \equiv \mathbf{enter} (x) \{ \bar{s} \} \mathbf{catch} \{ \}$$

In words, an enter statement propagates exceptions, and a reenter statement does not. Note: in real implementations, a reenter statement would not cause exception information to be lost, since the exception object would be associated with the failbox at the time the failbox is marked as failed, and an API would be provided to retrieve the stored exception object of a failed failbox.

### 3.3 Terminology

We use the following terminology: We say that an event in a thread  $t$  occurs *in a failbox*  $f$  or a statement is executed (or executes) in  $f$  if the event occurs or the statement execution starts at a time when  $f$  is the current failbox of  $t$ . We say that a *failure occurs* in  $t$  when an unchecked exception is thrown (i.e., the continuation of  $t$  is a throw statement). We say that a statement execution *fails* if it completes abruptly because of an unchecked exception. We say that a failbox  $f$  *fails* when a failure occurs in  $f$ . We say that an execution step *enters* a failbox  $f$  if  $f$  is the current failbox after the step and was not the current failbox before the step. Similarly, we say that an execution step *leaves* a failbox  $f$  if it is not the current failbox after the step and was the current failbox before the step.

### 3.4 Example

The approach is illustrated and motivated by the example in Figure 2. (Note: In the examples we use a more conventional syntax.) It shows how the unsafe program of Figure 1 can be made safe using failboxes. A failbox  $f$  is created and then both the main loop and calls of `addEntry` are executed in  $f$ . This ensures that if a call of `addEntry` fails, the main loop terminates.

The example motivates why on entry to a try block, the failbox in which the try-catch statement executes is no longer considered the current failbox. This ensures that failures in method `compute` are properly caught by the try-catch statement, and do not cause the program to terminate.

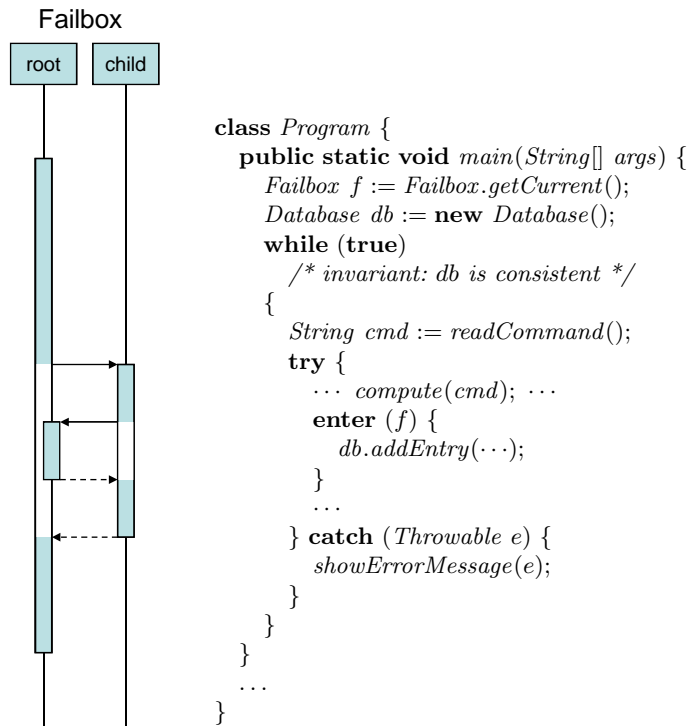
## 4 Multithreading

### 4.1 Synchronized statements: safety issues

One common way that the strict termination approach of dealing with failures is overridden, is through the use of **synchronized** blocks. A **synchronized** ( $o$ )  $S$  block in Java acquires the lock of object  $o$ , executes statement  $S$ , and then releases the lock of  $o$ , even if  $S$  failed. This helps prevent deadlocks, but it creates a safety risk. In particular, if a failure occurs while  $o$  is inconsistent, the commonly intended safety property that shared objects whose lock is not held are consistent, is violated.

The problem is illustrated by the example program in Figure 3. It is a multi-threaded version of the original example in Figure 1. Rather than processing each command before receiving the next command, the program receives a command, spawns a thread to process it, and immediately receives the next command. The `Database` object is shared by all command processing threads; accesses to the object are synchronized using a **synchronized** block.

This program is unsafe. In particular, in some executions, the intended safety property that whenever a shared object's lock is not held by any thread, the object is consistent, is violated. This property is relied on to guarantee that method



**Fig. 2.** The example of Figure 1, fixed using failboxes. When an *addEntry* call fails, failbox *f* is marked as failed. When control subsequently exits the try block, this is considered an attempt to enter *f*; therefore, a *FailboxException* is thrown. As a result, the catch block is skipped, the loop is exited, and the program terminates safely. The sequence diagram shows the failbox transitions.

*addEntry* is called only on objects that are consistent. Specifically, suppose a failure occurs in method *addEntry* while the *Database* object is inconsistent. This causes the lock to be released. Subsequent command processing threads that acquire the lock will then see the *Database* object in an inconsistent state.

#### 4.2 Failboxes approach for safe synchronized statements

Failboxes can be used to write safe lock-based multithreaded programs, by associating each shared object with a failbox and running the code that accesses a shared object within the associated failbox. This way, when a failure occurs, the failbox is marked as failed, so that when another thread subsequently attempts to enter the failbox in order to access the object, an exception is thrown and the thread is prevented from seeing inconsistent state. The modified safety property is that whenever no thread holds a shared object's lock, either the object is consistent or its associated failbox is marked as failed.

```

class Program {
  public static void main(String[] args) {
    final Database db := new Database();
    while (true) {
      final String cmd := readCommand();
      new Thread() {
        public void run() {
          try {
            ... compute(cmd); ...
            ... synchronized (db) { db.addEntry(...); } ...
          } catch (Throwable e) { showErrorMessage(e); }
        }
      }.start();
    }
  }
  ...
}

```

**Fig. 3.** An unsafe program. A failure in *compute* is handled correctly, but if a failure occurs in method *addEntry* while the *Database* object is inconsistent, the object’s lock is released, causing threads that subsequently acquire the lock to see the object in an unexpected state, violating safety.

The approach is illustrated in Figure 4. It is the example of Figure 3, made safe using failboxes. Specifically, the example uses an *enter* statement to execute the *addEntry* calls in the main thread’s root failbox. When an *addEntry* call fails, this failbox is marked as failed before the lock is released. When another thread subsequently acquires the lock and attempts to enter the failbox, an exception is thrown, so that the thread is prevented from unsafely calling *addEntry*.

### 4.3 Multithreaded failboxes

In a multithreaded program, it is possible for computations in multiple threads to be executing in the same failbox *f* concurrently. If this happens, we say *f* is multithreaded. The question then arises as to what happens when one of these computations fails. There are two distinct concerns involved in this matter: preserving the program’s intended safety properties, and ensuring useful progress.

### 4.4 Multithreaded failboxes: Safety

In a well-written program, a failure in one thread should not have safety implications for operations executing concurrently in other threads. Specifically, in a *data-race-free* program, where the program synchronizes accesses to shared memory using the language’s synchronization constructs, an operation can see

```

class Program {
  public static void main(String[] args) {
    final Failbox f := Failbox.getCurrent();
    final Database db := new Database();
    while (true) {
      final String cmd := readCommand();
      new Thread() {
        public void run() {
          try {
            ... compute(cmd); ...
            ... synchronized (db) { enter (f) { db.addEntry(...); } } ...
          } catch (Throwable e) { showErrorMessage(e); }
        }
      }.startInCurrentFailbox();
    }
  }
  ...
}

```

**Fig. 4.** The example of Figure 3, made safe using failboxes. If a call of *addEntry* fails, failbox *f* is marked as failed and subsequent attempts by other threads to enter the failbox will fail. Furthermore, by the Fail Fast feature, a *stop f* signal is sent to all threads running in the failed failbox *f* or a descendant of *f*. In the example, this means the program terminates.

the data that was being manipulated by a computation that failed only if the operation is not concurrent with the failure, i.e., the operation was synchronized with the failed computation. (Formally, the failure *happens-before* the operation.) Therefore, to ensure safety, it is sufficient that synchronization constructs perform the necessary failboxes bookkeeping to ensure that if a failure happens in a failbox, no operation that is ordered after this failure through synchronization runs in this failbox. To achieve this, we specify the semantics of synchronized statements with respect to failboxes as follows: after acquiring the lock, the statement checks that the current failbox has not failed; otherwise, it throws a *FailboxException*. Furthermore, before releasing the lock, if the body completed abruptly with an exception, the current failbox is marked as failed. The step rules are shown in Figure 5.

#### 4.5 Properties

We are now ready to state and sketch the proof of the main properties of the failboxes approach.

We first define some terms. An *execution* is a finite or countably infinite sequence of program states. An *execution point* is a nonnegative integer that

$$\begin{array}{c}
\text{SYNCHRONIZED} \\
\frac{(t, (f, V, \mathbf{synchronized}(x) \{ \bar{s}' \} \bar{s}, \bar{b}, \bar{F})) \in T \quad V(x) = o \quad o \notin \text{dom}(L) \quad f \notin \Phi \quad \bar{b}' = (\mathbf{synchronized}(o); \bar{s}) \cdot \bar{b}}{(L, \Sigma, \Phi, T) \xrightarrow{t:\text{acq}(o)} (L(o := t), \Sigma, \Phi, T(t := (f, V, \bar{s}', \bar{b}', \bar{F})))} \\
\\
\text{SYNCHRONIZED-REENTRANT} \\
\frac{(t, (f, V, \mathbf{synchronized}(x) \{ \bar{s}' \} \bar{s}, \bar{b}, \bar{F})) \in T \quad V(x) = o \quad (o, t) \in L}{(L, \Sigma, \Phi, T) \rightarrow (L, \Sigma, \Phi, T(t := (f, V, \bar{s}' \bar{s}, \bar{b}, \bar{F})))} \\
\\
\text{SYNCHRONIZED-COMPLETE-NORMAL} \\
\frac{(t, (f, V, \epsilon, (\mathbf{synchronized}(o); \bar{s}) \cdot \bar{b}, \bar{F})) \in T}{(L, \Sigma, \Phi, T) \xrightarrow{t:\text{rel}(o)} (L \setminus \{(o, t)\}, \Sigma, \Phi, T(t := (f, V, \bar{s}, \bar{b}, \bar{F})))} \\
\\
\text{SYNCHRONIZED-COMPLETE-ABRUPT} \\
\frac{(t, (f, V, \mathbf{throw};, (\mathbf{synchronized}(o); \bar{s}) \cdot \bar{b}, \bar{F})) \in T \quad \Phi' = \Phi \cup (\Sigma^{-1})^*(f) \quad T' = T(t := (f, V, \mathbf{throw};, \bar{b}, \bar{F}))}{(L, \Sigma, \Phi, T) \xrightarrow{t:\text{rel}(o)} (L \setminus \{(o, t)\}, \Sigma, \Phi', T')}
\end{array}$$

**Fig. 5.** Step rules for synchronized statements

serves as an index into an execution. A *thread execution point*  $(k, t)$  is a pair of an execution point  $k$  and a thread identifier  $t$ .

**Definition 1 (Happens-Before).** *The happens-before relation  $\xrightarrow{\text{hb}_E}$  on thread execution points of an execution  $E = C_0, C_1, \dots$  is the smallest transitive relation that satisfies the following properties:*

- Any thread execution point of a thread  $t$  happens-before any subsequent thread execution point of  $t$

$$k_1 < k_2 \Rightarrow (k_1, t) \xrightarrow{\text{hb}_E} (k_2, t)$$

- If execution step  $k_1$  is a release of some lock  $o$  by some thread  $t_1$ , and subsequent execution step  $k_2$  is an acquire of  $o$  by some thread  $t_2$ , then  $(k_1, t_1)$  happens-before  $(k_2 + 1, t_2)$

$$C_{k_1} \xrightarrow{t_1:\text{rel}(o)} C_{k_1+1} \Rightarrow C_{k_2} \xrightarrow{t_2:\text{acq}(o)} C_{k_2+1} \Rightarrow k_1 < k_2 \Rightarrow (k_1, t_1) \xrightarrow{\text{hb}_E} (k_2 + 1, t_2)$$

- If at execution step  $k$  thread  $t$  starts a new thread  $t'$  (see Figure 6), then  $(k, t)$  happens-before  $(k + 1, t')$

$$C_k \xrightarrow{t:\text{fork}(t')} C_{k+1} \Rightarrow (k, t) \xrightarrow{\text{hb}_E} (k + 1, t')$$

The Main Lemma states that once an exception occurs in a failbox, no code executes in that failbox “afterwards”.

**Lemma 1 (Main Lemma).** *Consider an execution  $E$  of a program  $\pi$  of the extended language, and consider two thread execution points  $(k_1, t_1)$  and  $(k_2, t_2)$  in  $E$ , such that  $(k_1, t_1)$  happens-before  $(k_2, t_2)$ . If  $t_1$  is executing in some failbox  $f_1$  in state  $k_1$ , and  $t_2$  is executing in some descendant  $f_2$  of  $f_1$  in state  $k_2$ , then if  $t_1$  is failing in state  $k_1$ , then  $t_2$  is failing in state  $k_2$ .*

$$\begin{aligned} \text{exec}(\pi, E) \Rightarrow (k_1, t_1) \xrightarrow{\text{hb}_E} (k_2, t_2) \Rightarrow \\ C_{k_1} = (L_1, \Sigma_1, \Phi_1, T_1) \Rightarrow T_1(t_1) = (f_1, V_1, \mathbf{throw};, \bar{b}_1, \bar{F}_1) \Rightarrow \\ C_{k_2} = (L_2, \Sigma_2, \Phi_2, T_2) \Rightarrow T_2(t_2) = (f_2, V_2, \bar{s}_2, \bar{b}_2, \bar{F}_2) \Rightarrow \\ f_2 \in (\Sigma_2^{-1})^*(f_1) \Rightarrow \bar{s}_2 = \mathbf{throw}; \end{aligned}$$

*Proof.* It suffices to prove for every prefix of some path from  $(k_1, t_1)$  to  $(k_2, t_2)$  in the happens-before graph, that at the thread execution point  $(k_3, t_3)$  at the end of the prefix, one or more of the following hold:

- the thread is failing and the current failbox is  $f_1$
- failbox  $f_1$  and its descendants have been marked as failed and one or more of the following hold:
  - the current failbox is not  $f_1$  or a descendant of  $f_1$ , or
  - the thread is failing.

This can be proved easily by induction on the length of the prefix and case analysis on the last edge.

Now consider an execution  $E$  of a program  $\pi$  and a *dependency relation*  $D$  on the thread execution points of  $E$ . We say  $E$  *uses failboxes correctly* with respect to  $D$ , if whenever thread execution point  $p_2$  depends on thread execution point  $p_1$ , the current failbox at  $p_2$  is a descendant of the current failbox at  $p_1$ . We say  $E$  is *dependency-safe* with respect to  $D$  if whenever  $p_2$  depends on  $p_1$ , and  $p_1$  happens-before  $p_2$ , and  $p_1$  is failing, then  $p_2$  is failing. We then have the Soundness Theorem: if  $E$  uses failboxes correctly with respect to  $D$ , then  $E$  is dependency-safe with respect to  $D$ . This follows directly from the Main Lemma.

A machine-checked proof of these properties is available online [12].

#### 4.6 Multithreaded failboxes: Ensuring useful progress

Even if a computation is safe, it might not be contributing to the useful work of the application. Specifically, if multiple computations are running in the same failbox, then this is taken to mean that they depend on each other for useful progress. As a result, if one of them fails, there is no point for the others to continue, so they should be stopped to free up CPU cycles, memory, and other resources these computations may be using. Therefore, in our approach, at the time a failbox  $f$  is marked as failed, a *stop  $f$*  signal is sent to all threads currently running in  $f$  or a descendant of  $f$ . When the signal arrives, this results in a *FailboxException* being thrown in the target thread, provided it is still running in  $f$  or a descendant. To allow efficient implementations, we do not impose timing constraints on the delivery of the signal, except that it must arrive eventually.

We call this mechanism the Fail Fast mechanism (after the Fail Fast principle [20]).

The usefulness of the Fail Fast mechanism is illustrated by the example in Figure 4. Once failbox  $f$  has failed, all subsequent attempts to access the database fail. Assuming most commands access the database, this means the program's functionality is severely degraded. Therefore, it seems appropriate to escalate the failure and terminate the program. This typically signals a system administrator or service management daemon to restart the program in a clean state, hopefully restoring full service. In the example, this behavior is achieved by running not just the `addEntry` calls, but the main loop as well, in failbox  $f$ . When an `addEntry` call fails, an asynchronous exception is thrown in the main thread, which causes the loop to terminate.

In fact, since the existing command processing threads are unlikely to be able to run to completion successfully, it makes sense to terminate these as well. This is achieved in the example by running the command processing threads in failbox  $f$  as well, by using method `startInCurrentFailbox` (added by our language extension) instead of `start` to start these threads. (To ensure backward compatibility, method `start` starts the new thread in a newly created root failbox, so that failure of the new thread does not cause a stop signal to be sent to the original thread.)

In the example, the failbox hierarchy is as follows. Failbox  $f$ , a root failbox, has one child for each try block execution. This ensures, as before, that exceptions in method `compute` do not cause the program to terminate.

The step rules for thread creation are shown in Figure 6. In the formal language, statement **fork** corresponds with method `startInCurrentFailbox`, and **fork\*** corresponds with method `start`.

$$\begin{array}{c}
\text{FORK} \\
\frac{(t, (f, V, \mathbf{fork} \{ \bar{s}' \} \bar{s}, \bar{b}, \bar{F})) \in T \quad t' \notin \text{dom}(T) \quad T' = T(t := (f, V, \bar{s}, \bar{b}, \bar{F}), t' := (f, V, \bar{s}', \epsilon, \epsilon))}{(L, \Sigma, \Phi, T) \xrightarrow{t:\mathbf{fork}(t')} (L, \Sigma, \Phi', T')} \\
\\
\text{THREAD-COMPLETE-ABRUPT} \\
\frac{(t, (f, V, \mathbf{throw}; \epsilon, \epsilon)) \in T \quad f \notin \Phi \quad \Phi' = \Phi \cup (\Sigma^{-1})^*(f)}{(L, \Sigma, \Phi, T) \rightarrow (L, \Sigma, \Phi', T)} \\
\\
\mathbf{fork}^* \{ \bar{s} \} \equiv \mathbf{fork} \{ x := \mathbf{newfb}; \mathbf{reenter}(x) \{ \bar{s} \} \} \quad \text{where } x \text{ is fresh}
\end{array}$$

**Fig. 6.** Step rules for thread creation



## 4.7 Wait dependency safety

A sub-concern of the concern of ensuring useful progress is the concern of ensuring progress. Specifically, one of the correctness properties that are difficult to achieve in the presence of unchecked exceptions is *wait dependency safety*, the property that if, in a given program execution, a wait operation  $W$  depends on a computation  $A$ , then, assuming that  $W$  terminates if  $A$  does not fail,  $W$  terminates. Analogously to the dependency relation used in the definition of dependency safety, the wait dependency relation used here is an application-specific relation; the intention is that if a wait operation  $W$  depends on a computation  $A$ , this means that, abstractly speaking,  $W$  waits for a signal to be sent by  $A$ . In Java, a typical example of this is when  $W$  is an *Object.wait* call on some object  $o$  and  $A$  at some point performs an *Object.notifyAll* call on  $o$ .

Failboxes can be used to achieve wait dependency safety. We say that a program *uses failboxes correctly* for the purpose of wait dependency safety if whenever in a given program execution, a wait operation  $W$  depends on a computation  $A$ , then  $A$  runs in a failbox  $f$  and  $W$  runs in a descendant of  $f$ . We then have the property that if a program uses failboxes correctly for the purpose of wait dependency safety, then the program is wait-dependency-safe. Indeed, if  $A$  fails, a stop signal is sent to the thread that is running  $W$ . As a result, when the signal arrives, either  $W$  has already terminated, or  $W$  is terminated by the *FailboxException* thrown by the Fail Fast mechanism. We call this property the soundness of the Fail Fast mechanism.

A machine-checked proof of this property is available online [12].

## 5 Cancellation and Compensation

We propose the use of failboxes in programs to make them safe for failures. However, it turns out that if failboxes are applied correctly in a program, then this also enables safe *cancellation* of computations, with no extra effort, and without the need for polling, through the Fail Fast mechanism. In order to enable cancellation of a computation, the program runs it in a dedicated failbox; to cancel the computation, it calls the *Failbox* object's *cancel* method, which simulates the occurrence of a failure in the failbox and triggers the Fail Fast mechanism. This achieves the convenience of the deprecated *Thread.stop* approach, without the safety risk.

Consider for example the program of Figure 4. The main loop repeatedly receives a command and starts a command thread to process it. The processing is done inside a try-catch statement, and therefore in a per-command child failbox of the root failbox. This program could be extended to enable cancellation of commands as follows. In order to cancel a command, the program calls the command failbox's *cancel* method. If the command thread is executing in the command failbox, it is stopped; however, if it is executing inside the database, it is allowed to continue to execute until it leaves the root failbox and re-enters the command failbox, at which point an exception is thrown. Contrast this with

calling *stop* on the command thread, which would stop the thread even if it was running in the database, causing the entire program to fail.

The failboxes mechanism also enables *safe compensation*. By compensation, we refer to the scenario where a client computation invokes a service offered by a provider computation, which changes the provider's state. This imposes the obligation on the client to invoke a compensating service to restore the provider's state, after the client is done using the service. The conventional approach to compensation is through try-finally statements. However, an unchecked exception can cause the compensation action to be skipped, if the exception occurs after the action that is to be compensated, but before the try block is entered, or if it occurs after the finally block is entered, but before the compensation action completes.

This may be addressed using the failboxes mechanism by performing the following transformation:

$$\begin{array}{l}
 \text{init}(); \\
 \mathbf{try} \{ \\
 \quad // \text{Use the service} \\
 \} \mathbf{finally} \{ \\
 \quad \text{compensate}(); \\
 \}
 \end{array}
 \Rightarrow
 \begin{array}{l}
 \mathbf{enter} (\text{provider}) \{ \\
 \quad \text{init}(); \\
 \quad \mathbf{reenter} (\text{client}) \{ \\
 \quad \quad // \text{Use the service} \\
 \quad \} \\
 \quad \text{compensate}(); \\
 \}
 \end{array}$$

Before invoking the service, the thread running the client computation enters the provider's failbox. After the service is invoked, it re-enters the client failbox using a nested reenter statement where the client uses the service. When the client is done using the service, it leaves the nested enter statement, causing the thread to re-enter the provider failbox, perform the compensating action, and finally leave the outer enter statement, re-entering the client failbox. This approach guarantees that either the compensation occurs or the provider failbox is marked as failed. If an exception occurs while the client uses the service, the client failbox is marked as failed, but the exception is not propagated by the reenter statement. This ensures that compensation is not skipped. When the thread leaves the outer enter statement, it enters the client failbox, which was marked as failed, and therefore the exception is propagated from that point, as in the case of the try-finally statement.

## 6 Proof rules

To show that it is easy to reason about programs that use failboxes, in this section we propose separation logic proof rules for the main envisaged usage patterns.

Recall the semantics of separation logic assertions: **emp** describes the empty heap, and the separate conjunction  $P * Q$  describes a heap that can be split into one that satisfies  $P$  and one that satisfies  $Q$ :

$$s, h \models \mathbf{emp} \Leftrightarrow h = \emptyset \quad s, h \models P * Q \Leftrightarrow \exists h_1, h_2 \bullet h = h_1 \uplus h_2 \wedge s, h_1 \models P \wedge s, h_2 \models Q$$

We extend the syntax of correctness judgments (but not the syntax of assertions) to be failboxes-aware. Specifically,

$$\Sigma; f \vdash \{P\} \bar{s} \{Q\}$$

denotes the correctness of statement list  $\bar{s}$  under commitment list  $\Sigma$ , current failbox  $f$ , precondition  $P$ , and postcondition  $Q$ . The syntax of commitment lists is as follows:

$$\Sigma ::= \epsilon \mid \Sigma, f : P$$

We say that assertion  $P$  is committed to failbox  $f$ . Informally, this means that to access the resources of  $P$ ,  $f$  must first be entered. Failboxes are denoted using logical variables.

The above correctness judgment implies the following validity statement:

$$\llbracket \Sigma \rrbracket * P \Rightarrow \text{valid}(\bar{s}, \llbracket \Sigma \rrbracket * Q, \llbracket \Sigma \rrbracket * \mathbf{true})$$

(under the assumption that  $\bar{s}$  does not assign to any variables that  $\Sigma$  depends on) where  $\Sigma$  is here interpreted as a separation logic assertion as follows:

$$\llbracket \epsilon \rrbracket \equiv \mathbf{emp} \quad \llbracket \Sigma, f : P \rrbracket \equiv \llbracket \Sigma \rrbracket * (f \in \Phi \vee P)$$

i.e., for each commitment  $f : P$ , either  $P$  holds (and is owned by the current thread) or  $f$  has failed.  $\text{valid}(\bar{s}, Q, R)$  is true under a given heap, failed set, and variable environment, if after executing  $\bar{s}$  in this state, upon normal completion  $Q$  holds and upon abrupt completion  $R$  holds.

A throw statement always satisfies partial correctness.

$$\begin{array}{c} \text{C-THROW} \\ \Sigma; f \vdash \{P\} \mathbf{throw}; \{Q\} \end{array}$$

For verifying a try-catch statement, the heap is split into two parts: part  $P_f$  is accessed by the try block only inside  $\mathbf{enter}(f)$  statements, and part  $P$  is accessed freely. The second premise of the rule ensures soundness for normal completion of the try block. The third is for the case where the try block fails.

$$\begin{array}{c} \text{C-TRYCATCH} \\ \frac{\forall f' \bullet \Sigma, f : P_f; f' \vdash \{P\} \bar{s} \{Q\} \quad P_f * Q \Rightarrow Q' \quad \Sigma; f \vdash \{P_f\} \bar{s}' \{Q'\}}{\Sigma; f \vdash \{P_f * P\} \mathbf{try} \{ \bar{s} \} \mathbf{catch} \{ \bar{s}' \} \{Q'\}} \end{array}$$

(under the assumption that  $P_f$  does not depend on any variables that  $\bar{s}$  assigns to).

An enter block can access the piece of heap associated with the failbox being entered.

$$\begin{array}{c} \text{C-ENTER} \\ \frac{\Sigma; f \vdash \{P * P_f\} \bar{s} \{Q * P_f\}}{\Sigma, f : P_f; f' \vdash \{P \wedge x = f\} \mathbf{enter}(x) \bar{s} \{Q\}} \end{array}$$

The compensation pattern can be verified as follows.

$$\begin{array}{c}
\text{C-COMPENSATION} \\
\frac{\Sigma; f \vdash \{P * P_f\} \bar{s}_1 \{Q_1 * P'_f \wedge y = f'\} \quad \Sigma, f : P'_f; f' \vdash \{Q_1\} \bar{s}_2 \{Q\} \quad \Sigma; f \vdash \{P'_f\} \bar{s}_3 \{P_f\}}{\Sigma, f : P_f; f' \vdash \{P \wedge x = f\} \mathbf{enter} (x) \{ \bar{s}_1 \mathbf{reenter} (y) \{ \bar{s}_2 \} \bar{s}_3 \} \{Q\}}
\end{array}$$

(under the assumption that  $P'_f$  does not care about any variables that  $\bar{s}_2$  assigns to). The compensation pattern allows the commitment  $f : P_f$  to be replaced temporarily with the commitment  $f : P'_f$ .

A machine-checked soundness proof of these proof rules is available online [12].

We developed a prototype verifier based on these ideas [12].

## 7 Implementation Issues

We created a prototype implementation of the approach on the .NET Framework as a C# 3.0 library. C# 3.0's lambda expression syntax can be used to write reasonably concise enter statements.

A major complication for achieving a fully correct implementation of the approach in the form of a library, is the fact that the .NET Framework Common Language Runtime may throw an exception at any program point, due to an internal resource limit being reached or an internal error being discovered within the execution engine [21]. (The same holds for the Java Virtual Machine. See the Java Virtual Machine Specification, Second Edition [14], Section 2.16.2.) Specifically, if an enter block completes abruptly with an exception, no internal exception must intervene between catching the exception and marking the failbox as failed; otherwise, the enter statement completes without marking the failbox as failed, breaking dependency safety.

Version 2.0 of the .NET Framework introduced constructs specifically for writing code that must execute reliably in the presence of internal exceptions [21]. We used these constructs in our prototype implementation to ensure that on abrupt completion of the body of an enter statement, the failbox and its descendants are marked as failed and stop signals are sent to other threads executing in the failbox or its descendants. Specifically, we used the following API:

*ExecuteCodeWithGuaranteedCleanup*( $t, c, u$ )

where  $t$  and  $c$  are delegates (similar to function pointers in C) and  $u$  is arbitrary user data that is passed to  $t$  and  $c$ . The API first executes  $t$ . When  $t$  completes, either normally or abruptly, the cleanup delegate  $c$  is executed. The API guarantees that no internal exceptions occur during the execution of  $c$ , provided that  $c$  satisfies certain constraints, such as: no heap memory allocation, and no unbounded call stack memory allocation. Unfortunately, these constraints have not been spelled out very precisely anywhere; we had to make some assumptions as to what can reasonably be executed without the risk of internal exceptions.

We have performed a few microbenchmark performance tests. These indicate the following approximate timings for the following statements:

Statement	Timing	Timing*
<code>try {} catch {}</code>	13 $\mu$ s	1.9 $\mu$ s
<code>try { enter (f) {} } catch {}</code>	23 $\mu$ s	3.4 $\mu$ s

To measure the impact of the *ExecuteCodeWithGuaranteedCleanup* construct, we replaced it with a dummy that uses a simple try-finally statement. The resulting timings are shown in the third column. It turns out that the overhead of this construct dominates the run time.

Even though the current performance is probably acceptable for most real-world applications, we believe it can still be improved significantly, in particular if the constructs are implemented directly in the virtual machine rather than as a library. Performing such an implementation is future work.

We have also prepared a prototype implementation of failboxes as a library on the Java virtual machine. However, due to the absence of constructs to prevent internal or asynchronous exceptions on this platform, the implementation is not safe in the presence of such exceptions.

The prototype implementations are available on line [12].

## 8 Related Work

To the best of our knowledge, failboxes are the first approach for programmers to achieve dependency safety of their Java-like programs that combines low programming overhead, low performance overhead, and low reasoning overhead, and is compositional (i.e. failboxes can be nested arbitrarily).

*Languages as operating systems* Many extensions of Java have been proposed that support running multiple programs or *tasks* in the same virtual machine. These can typically be used to enforce dependency safety. However, in contrast to failboxes, all of these have goals beyond dependency safety, typically including protection against malicious code, and accounting of memory and other resources. As a result, they impose greater programming and performance overhead on communication between tasks than the overhead of switching between failboxes.

Perhaps the most closely related such system is Luna [11]. To support memory accounting and immediate guaranteed memory reclamation when a task is killed, the heap is logically partitioned among the tasks; the only way for one task to access an object belonging to another task is through a *remote pointer*, which is distinguished from local pointers through its type. When a task is killed, remote pointers pointing into it are *revoked*, so that if the task was holding a lock, other tasks do not see inconsistent state. Failboxes offer no memory accounting or guaranteed memory reclamation, but in turn impose a lower programming and performance overhead. Specifically, passing data across tasks requires either copying or the use of remote pointers, both of which incur a programming and

performance overhead; failboxes, in contrast, allow data to be passed around freely.

DrScheme [8, 7] is a Scheme environment designed for programs that serve as platforms for other programs. In DrScheme, it is possible for two child programs to share a mutable data structure and yet be killed independently. The solution is to host the data structure in a separate thread, and to access it only via message passing with this thread. DrScheme's contribution is that it enables two untrusted child programs to set up such a shared data structure without circumventing resource policies and without the need for the shared structure to be trusted by the kernel. However, from a dependency safety point of view, the situation is as in Java: DrScheme requires either the use of message passing between separate threads or manually guarding dependent code.

Erlang [1] is a language focused on reliability. Inconsistent data structures within a process are ruled out because the language has no destructive update. Processes communicate through asynchronous message passing. Fail-fast is achieved by linking processes: when a process dies, an exit signal is sent to linked processes, causing those to die as well by default.

*Non-compositional approaches* Marlow et al. [16] propose an extension of concurrent Haskell with constructs that make it possible to write safe programs where one thread throws an asynchronous exception in another thread. The `block e` construct disables asynchronous exceptions during execution of `e`; `e` can use `unblock e'` to re-enable them during execution of a sub-expression `e'`. Unlike failboxes, the `block` construct is not compositional; for example, in the program of Figure 4, the `addEntry` call could be protected against cancellation of jobs using `block`; however, imagine the command processing program is part of a larger system. Then one may want to cancel the program as a whole, including any `addEntry` calls. This is possible with failboxes (by cancelling failbox `f`, which cancels its descendants as well), but not with the `block` construct. Also, the construct does not help in dealing with failures; for example, a failure during the `addEntry` call would not prevent further accesses to the database. However, the `block` construct, or something similar, is useful and even necessary to be able to robustly implement failboxes as a library in a given language.

Starting with version 2, the .NET Framework includes reliability features that make it possible to write cleanup routines that are guaranteed to execute even in the presence of failure or cancellation [21]. However, like the `block` construct, the approach is not compositional: these cleanup routines cannot be cancelled; furthermore, they must be carefully coded to rule out failures within the cleanup routines themselves since those are not dealt with safely. The mechanism is intended only for manipulation of execution environment resources; it is not for general application use.

Three further reliability-related features in .NET Framework version 2 are the following. Firstly, cancellation is disabled during finally blocks. This enables safe cleanup in the presence of cancellation (but not failure). Secondly, an unhandled exception in one thread kills all other threads, without executing catch or finally blocks. However, in the thread that throws the unhandled ex-

ception, finally blocks are executed normally and locks are released, leaving a time window between the release of the lock and the time the exception reaches the toplevel (possibly after executing other finally blocks) where other threads can see inconsistent state. Thirdly, a method *Environment.FailFast* was added, which terminates the program immediately.

Rudys et al. [18] propose weaving code into an untrusted plugin (such as an applet) that polls a cancellation request flag to enable forcibly cancelling the plugin. The flag is also checked whenever the host system calls into the plugin. In our approach, a thread running in one failbox may protect itself from cancellation of its failbox by entering an ancestor failbox to which it has a reference; however, separate techniques (e.g., perhaps by associating permissions with failboxes) could be used to prevent this in case the thread is running untrusted code.

The SCOOP multithreading approach for Eiffel [17] has a notion of *subsystems*. A subsystem in SCOOP is a thread and a set of objects handled by that thread. Brooke and Paige [3] suggest marking an object as “dead” when the processing of an asynchronous incoming call fails, causing subsequent calls to fail immediately. SCOOP subsystems cannot be nested.

*Other related work* Garcia et al. [9] provide a survey of exception mechanisms. However, the authors do not discuss the dependency safety issue. In fact, most modern imperative and/or object-oriented languages have inherited the exception mechanism of CLU [15] and therefore suffer from the problems addressed by our approach.

*Class-handlers*, as proposed by Dony [4] and others, are exception handlers associated with classes rather than blocks of statements; they apply to all methods of the class. They would facilitate manually guarding dependent code. For example, a class-handler on the *Database* class could set a *failed* field to *true* when an unchecked exception is caught and then re-throw the exception. The field would still need to be checked manually on entry to each method.

Weimer and Necula [22] propose *compensation stacks* to make it easier to write effective cleanup code. However, they do not address the safety issues identified in Section 5.

Fetzer et al. [5] assume the viewpoint that “exception handling is only effective if the premature termination of a method due to an exception does not leave an object in an inconsistent state”. The paper proposes techniques to detect and “mask” *non-atomic exception handling*, i.e. violations against *failure atomicity*. The paper assumes that after catching an exception, the entire application should be in a consistent state, whereas we allow *failed failboxes* to remain in an inconsistent state, while preventing control from entering a failed failbox. The authors find a large number of Java methods that are not failure atomic. This would strengthen the case for failboxes, because it indicates that exceptions do indeed commonly leave objects in an inconsistent state.

An alternative way to deal with failures is to roll the state of the objects involved back to a consistent state, through the use of transactions (e.g. Shavit and Touitou [19], Welc et al. [23], Fetzer et al. [5]). However, this has a greater performance overhead; also, it presents problems when the computation that failed

performed I/O. Our failboxes approach is more conservative from a semantic and performance point of view.

This work was inspired by our research in program verification for Java-like languages that is sound in the presence of failures. To the best of our knowledge, no existing program verifiers for Java-like languages (including ESC/Java [6] and Spec# [2]) have this property. In Jacobs et al. [13], we propose a verification approach for Java programs where the programmer manually guards dependent code using flag variables that track an object's consistency. The present work addresses the programming overhead of that approach.

## 9 Conclusion

We propose a language extension, called *failboxes*, that facilitates writing sequential or multithreaded programs that provably preserve intended safety properties and that do not leak resources, even in the presence of failure, and that perform safe cancellation of computations. To the best of our knowledge, it is the first such extension of a Java-like language that combines low programming, performance, and reasoning overhead, and that is compositional.

Future work includes gaining experience with our prototype implementation, mainly to assess the applicability and the usability of the approach. We anticipate the possible need to facilitate the placement of enter blocks, perhaps through annotations on methods, classes, or packages, or through some inference scheme. Other work includes applying the failboxes idea to the problem of exception handling in asynchronous and callback patterns.

## Acknowledgements

The authors would like to thank Jan Smans, Marko van Dooren, Scott Owens, and the Cambridge programming languages group for their helpful comments. This research is partially funded by the Interuniversity Attraction Poles Programme Belgian State, Belgian Science Policy.

## Bibliography

- [1] Joe Armstrong. *Making reliable distributed systems in the presence of software errors*. PhD thesis, Royal Institute of Technology, Stockholm, Sweden, 2003.
- [2] Mike Barnett, Bor-Yuh Evan Chang, Robert DeLine, Bart Jacobs, and K. Rustan M. Leino. Boogie: A modular reusable verifier for object-oriented programs. In *Proc. FMCO 2005*, volume 4111 of *LNCS*. Springer, 2006.
- [3] Phillip J. Brooke and Richard F. Paige. Exceptions in Concurrent Eiffel. *Journal of Object Technology*, 6(10):111–126, November 2007.
- [4] Christophe Dony. Exception handling and object-oriented programming: a synthesis. In *Proc. OOPSLA*, 1990.



- [5] Christof Fetzer, Karin Högstedt, and Pascal Felber. Automatic detection and masking of non-atomic exception handling. In *Proc. Intl. Conf. Dependable Systems and Networks (DSN)*, 2003.
- [6] Cormac Flanagan, K. Rustan M. Leino, Mark Lillibridge, Greg Nelson, James B. Saxe, and Raymie Stata. Extended static checking for Java. In *Proc. PLDI*, pages 234–245, 2002.
- [7] Matthew Flatt and Robert Bruce Findler. Kill-safe synchronization abstractions. In *Proc. PLDI*, 2004.
- [8] Matthew Flatt, Robert Bruce Findler, Shriram Krishnamurthi, and Matthias Felleisen. Programming languages as operating systems (or Revenge of the son of the Lisp machine). In *Proc. Intl. Conf. on Functional Programming (ICFP)*, 1999.
- [9] Alessandro F. Garcia, Cecília M. F. Rubira, Alexander B. Romanovsky, and Jie Xu. A comparative study of exception handling mechanisms for building dependable object-oriented software. *Journal of Systems and Software*, 59(2):197–222, 2001.
- [10] James Gosling, Bill Joy, Guy Steele, and Gilad Bracha. *The Java Language Specification, Third Edition*. Prentice Hall PTR, 2005.
- [11] Chris Hawblitzel and Thorsten von Eicken. Luna: a flexible Java protection system. In *Proc. OSDI*, 2002.
- [12] Bart Jacobs and Frank Piessens. Failboxes: Prototype implementations, prototype verifier, machine-checked metatheory. <http://www.cs.kuleuven.be/~bartj/failboxes>, July 2008.
- [13] Bart Jacobs, Peter Müller, and Frank Piessens. Sound reasoning about unchecked exceptions. In *Proc. ICFEM*, 2007.
- [14] Tim Lindholm and Frank Yellin. *The Java Virtual Machine Specification, Second Edition*. Addison-Wesley, 1999. Online at <http://java.sun.com/docs/books/jvms/>.
- [15] Barbara Liskov and Alan Snyder. Exception handling in CLU. *IEEE Trans. Software Eng.*, 5(6):546–558, 1979.
- [16] Simon Marlow, Simon Peyton Jones, Andrew Moran, and John Reppy. Asynchronous exceptions in Haskell. In *Proc. PLDI*, 2001.
- [17] Bertrand Meyer. *Eiffel: The Language*. Prentice Hall, 1992.
- [18] Algis Rudys, John Clements, and Dan S. Wallach. Termination in language-based systems. In *Network and Distributed System Security Symposium (NDSS)*, February 2001.
- [19] N. Shavit and D. Touitou. Software transactional memory. In *Proc. PODC*, pages 204–213, 1995.
- [20] Jim Shore. Fail fast. *IEEE Software*, September 2004.
- [21] Stephen Toub. Keep your code running with the reliability features of the .NET Framework. *MSDN Magazine*, October 2005.
- [22] Westley Weimer and George C. Necula. Finding and preventing run-time error handling mistakes. In *Proc. OOPSLA*, pages 419–431, October 2004.
- [23] Adam Welc, Suresh Jagannathan, and Antony L. Hosking. Transactional monitors for concurrent objects. In *Proc. ECOOP*, 2004.