Implementing Joins using Extensible Pattern Matching

Philipp Haller[1], Tom Van Cutsem[2][*]

[1] LAMP-REPORT-2007-004
EPFL, 1015 Lausanne, Switzerland
`firstname.lastname@epfl.ch`
+41 21 693 6483, +41 21 693 6660
[2] Programming Technology Lab, Vrije Universiteit Brussel, Belgium

Abstract. Join patterns are an attractive declarative way to synchronize both threads and asynchronous distributed computations. We explore joins in the context of extensible pattern matching that recently appeared in languages such as F# and Scala. Our implementation supports Ada-style rendezvous, and constraints. Furthermore, we integrated joins into an existing actor-based concurrency framework. It enables join patterns to be used in the context of more advanced synchronization modes, such as future-type message sending and token-passing continuations.

Keywords: Concurrent Programming, Join Patterns, Actors

1 Introduction

Recently, the pattern matching facilities of languages such as Scala and F# have been generalized to allow representation independence for objects used in pattern matching [4,15]. Extensible patterns open up new possibilities of implementing abstractions in libraries which were previously only accessible as language features. More specifically, we claim that extensible pattern matching eases the construction of declarative approaches to synchronization in libraries rather than languages. To support this claim, we show how a concrete declarative synchronization construct, join patterns, can be implemented in Scala, a language with extensible pattern matching. Join patterns [6,7] offer a declarative way of synchronizing both threads and asynchronous distributed computations that is simple and powerful at the same time. They form part of functional languages such as JoCaml [5] and Funnel [10]. Join patterns have also been implemented as extensions to existing languages [2,17].

Recently, Russo [13] and Singh [14] have already shown that advanced programming language features, such as generics or software transactional memory, make it feasible to provide join patterns as libraries rather than language extensions. As we will argue in section 2, an implementation using extensible pattern matching improves upon these previous approaches by providing a better integration between library and language. More concretely, we make the following contributions:

– Our implementation technique overcomes several limitations of previous library-based designs and language extensions. In all library-based implementations that we know of, pattern variables are represented implicitly as parameters of join continuations. Mixing up parameters of the same type inside the join body may lead to obscure errors that are hard to detect. Our design avoids these errors by using the underlying pattern matcher to bind variables that are explicit in join patterns. The programmer may use a rich pattern syntax to express constraints using nested patterns and guards. However, efficiently supporting general guards in join patterns is currently an open problem, and we do not attempt to solve it.

– We present a complete implementation of our design as a Scala library[3] that supports Ada-style rendezvous and constraints. Moreover, we integrate our library into an existing event-based, non-blocking concurrency framework. This enables expressive join patterns to be used in the context of more advanced synchronization modes, such as future-type message sending and token-passing continuations. Our integration is unique in the sense that the new library provides a conservative syntax extension. That is, existing programs continue to run without change when compiled or linked against the extended framework.

The rest of this paper is structured as follows. In the following section we briefly highlight join patterns as a declarative synchronization abstraction, how they have been integrated in other languages before, and how combining them with pattern matching can improve this integration. Section 3.1 shows how to synchronize threads using join patterns written using our library. Section 3.2 shows how to use join patterns with actors. In section 4 we discuss the concrete Scala implementation of Joins for threads and actors. Section 5 discusses related work, and section 6 concludes.

2 Motivation

Background: Join Patterns A join pattern consists of a body guarded by a linear set of events. The body is executed only when *all* of the events in the set have been signaled to an object. Threads may signal synchronous or asynchronous events to objects. By signaling a synchronous event to an object, threads may implicitly suspend. The simplest illustrative example of a join pattern is that of an unbounded FIFO buffer. In $C\omega$, it is expressed as follows [2]:

```
public class Buffer {
  public async Put(int x);
  public int Get() & Put(int x) { return x; }
}
```

A detailed explanation of join patterns is outside the scope of this paper. For the purposes of this paper, it suffices to understand the operational effect of a join pattern. Threads may put values into a buffer b by invoking b.Put(v). They may also read values from the buffer by invoking b.Get(). The join pattern Get() & Put(int x) (called a *chord* in

---

[3] Available at `http://lamp.epfl.ch/~phaller/joins/`.

Cω) specifies that a call to Get may only proceed if a Put event has previously been signaled. Hence, if there are no pending Put events, a thread invoking Get is automatically suspended until such an event is signaled.

The advantage of join patterns is that they allow a *declarative* specification of the synchronization between different threads. Often, the join patterns correspond closely to a finite state machine that specifies the valid states of the object [2]. Section 3.1 provides a more illustrative example of the declarativeness of join patterns.

Limitations of library-based designs In Cω, join patterns are supported as a language extension through a dedicated compiler. This ensures that join patterns are perfectly integrated in the language. With the introduction of generics in C#, Russo has made join patterns available as a regular C# library called Joins [13]. In that library, the Buffer example is encoded as follows:

```
public class Buffer {
  // Declare (a)synchronous channels
  public readonly Asynchronous.Channel<int> Put;
  public readonly Synchronous<int>.Channel Get;
  public Buffer() {
    Join join = Join.Create();
    join.Initialize(out Put); join.Initialize(out Get); // initialize channels
    join.When(Get).And(Put).Do(delegate(int x) { return x; });
  }
}
```

In C# Joins, join patterns consist of linear combinations of channels and a delegate (a function object) which encapsulates the body. Join patterns are triggered by invoking channels, which are special delegates.

Even though the synchronization between Get and Put is still readily apparent in the above example, the Joins library design has some drawbacks. First and foremost, the way in which arguments are passed between the channels and the body is very implicit: the delegate is implicitly invoked with the value passed via the Put channel. Contrast this with the Cω example in which the variable x is explicitly tied to the Put message. Furthermore, because Joins are defined by means of an ad hoc combination mechanism, it is impossible to declaratively specify additional pattern matches or even guards. For example, it is not possible to add a join pattern triggering only on calls to Put where x is 0 or x > 100. Instead, one would have to add an if-test to the body of the join which partially defeats the declarative nature of the synchronization. In section 3, we show how these drawbacks can be eliminated by integrating join patterns with a host language's standard support for pattern matching.

Joins for Actors While join patterns have been successfully used to synchronize threads, to the best of our knowledge, they have not yet been applied in the context of an actor-based concurrency model. In Scala, actor-based concurrency is supported by means of a library [8]. Because we provide join patterns as a library extension as well, we have created the opportunity to combine join patterns with the event-driven concurrency model

offered by actors. We give a detailed explanation of this combination in section 3.2. However, in order to understand this integration, we first briefly highlight how to write concurrent programs using Scala's actor library.

Scala's actor library is largely inspired by Erlang's model of concurrent processes communicating by message-passing [1]. New actors are defined as classes inheriting the `Actor` class. The actor's life cycle is described by its `act` method. The following code shows how to implement the unbounded buffer as an actor:

```scala
class Buffer extends Actor {
  override def act() { loop(Nil) }
  def loop(buf: List[Int]) {
    receive {
      case Put(x) => loop(buf ::: List(x)) // append x to buf
      case Get() if !buf.isEmpty => reply(buf.head); loop(buf.tail) }
  }
}
```

The `receive` method allows an actor to selectively wait for certain messages to arrive in its mailbox. Interacting with a buffer actor occurs as follows:

```scala
val buffer = new Buffer; buffer.start()
buffer ! Put(42) // asynchronous send, returns nothing
println(buffer !? Get()) // synchronous send, waits for reply
```

Synchronous message sends make the sending process wait for the actor to reply to the message (by means of the `reply(value)` method). Scala actors also offer more advanced synchronization patterns such as futures [9,19]. `actor !! msg` denotes an asynchronous send that immediately returns a future object. In Scala, a future is a nullary function that, when applied, returns the future's computed result value. If the future is applied before the value is computed, the caller is blocked.

In the above example, the required synchronization between `Put` and `Get` is achieved by means of a *guard*. The guard in the `Get` case disallows the processing of any `Get` message while the `buf` queue is empty. In the implementation, all cases are sequentially checked against the incoming message. If no case matches, or all of the guards for matching cases evaluate to false, the actor keeps the message stored in its mailbox and awaits other messages.

Even though the above example remains simple enough to implement, the synchronization between `Put` and `Get` remains very implicit. The actual *intention* of the programmer, i.e. the fact that an item can only be produced when the actor received both a `Get` *and* a `Put` message, remains implicit in the code. Hence, even actors can benefit from the added declarative synchronization of join patterns, as we will illustrate in section 3.2.

3   A Scala Joins Library

We now discuss a Scala library (henceforth called Scala Joins) providing join patterns implemented via extensible pattern matching. First, we explain how Scala Joins enables the declarative synchronization of threads, postponing joins for actors until section 3.2.

3.1 Joining Threads

Scala Joins draws on Scala's extensible pattern matching facility [4]. This has several advantages: first of all, the programmer may use Scala's rich pattern syntax to express constraints using nested patterns and guards. Moreover, reusing the existing variable binding mechanism avoids typical problems of other library-based approaches where the order in which arguments are passed to the function implementing the join body are merely conventional, as explained in section 2. Similar to C# Joins's channels, joins in Scala Joins are composed of synchronous and asynchronous *events*. Events are strongly typed and can be invoked using standard method invocation syntax. The FIFO buffer example is written in Scala Joins as follows:

```scala
class Buffer extends Joins {
  val Put = new AsyncEvent[Int]
  val Get = new SyncEvent[Int]
  join { case Get() & Put(x) => Get reply x }
}
```

To enable join patterns, a class inherits the `Joins` class. Events are declared as regular fields. They are distinguished based on their (a)synchrony and the number of arguments they take. For example, `Put` is an asynchronous event that takes a single argument of type `Int`. Since it is asynchronous, no return type is specified (it immediately returns `unit` when invoked). In the case of a synchronous event such as `Get`, the first type parameter specifies the return type. Therefore, `Get` is a synchronous event that takes no arguments and returns values of type `Int`.

Joins are declared using the `join { ... }` construct. This construct enables pattern matching via a list of `case` declarations that each consist of a left-hand side and a right-hand side, separated by =>. The left-hand side defines a join pattern through the juxtaposition of a linear combination of asynchronous and synchronous events. As is common in the joins literature, we use & as the juxtaposition operator. Arguments of events are usually specified as variable patterns. For example, the variable pattern x in the `Put` event can bind to any value (of type `Int`). This means that on the right-hand side, x is bound to the argument of the `Put` event when the join pattern matches. Standard pattern matching can be used to constrain the match even further (an example of this is given below).

The right-hand side of a join pattern defines the join body (an ordinary block of code) that is executed when the join pattern matches. Like JoCaml, but unlike Cω and C# Joins, Scala Joins allows any number of synchronous events to appear in a join pattern. Because of this, it is impossible to use the return value of the body to implicitly reply to the single synchronous event in the join pattern. Instead, the body of a join pattern explicitly replies to all of the synchronous events that are part of the join pattern on the left-hand side. Synchronous events are replied to by invoking their `reply` method. This wakes up the thread that originally signalled that event.

To demonstrate how join patterns can be combined with ordinary pattern matching, consider the traditional problem of synchronizing multiple concurrent readers with one or more writers who need exclusive access to a resource. A multiple reader/one writer lock can be implemented in our library as follows[4]:

---

[4] This implementation is based on that of Cω [2] and Russo's Joins library for C# [13].

```scala
class ReaderWriterLock extends Joins {
  private val Idle = new NullaryAsyncEvent
  private val Sharing = new AsyncEvent[Int]
  val Exclusive, ReleaseExclusive = new NullarySyncEvent
  val Shared, ReleaseShared = new NullarySyncEvent
  join {
    case Exclusive() & Idle() => Exclusive reply
    case ReleaseExclusive() => { Idle(); ReleaseExclusive reply }
    case Shared() & Idle() => { Sharing(1); Shared reply }
    case Shared() & Sharing(n) => { Sharing(n+1); Shared reply }
    case ReleaseShared() & Sharing(1) => { Idle(); ReleaseShared reply }
    case ReleaseShared() & Sharing(n) => { Sharing(n-1); ReleaseShared reply }
  }
  Idle() }
```

In the above example, events are used to encode the state of the reader-writer lock. The last statement ensures that the lock starts off in an `Idle` state (no active threads). A writer can signal a synchronous `Exclusive` event to acquire the lock. Concurrent readers are represented by means of a `Sharing(n)` event which encodes the number of currently active readers.

In the join pattern `ReleaseShared() & Sharing(1)`, regular pattern matching is used to constrain the pattern only to `Sharing` events whose argument equals 1, thus ensuring that this pattern only triggers when the *last* reader releases the lock. If join patterns would not be integrated with pattern matching, code like this would require additional tests in the body of more general join patterns.

3.2   Joining Actors

We now describe an integration of our Joins library with Scala's actor library [8]. The following example shows how to re-implement the unbounded buffer example using Joins:

```scala
val Put = new Join1[Int]
val Get = new Join
class Buffer extends JoinActor {
  def act() {
    receive { case Get() & Put(x) => Get reply x }
  } }
```

It differs from the thread-based bounded buffer using joins in the following ways:

– The `Buffer` class inherits the `JoinActor` class to declare itself to be an actor capable of processing join patterns.
– Rather than defining `Put` and `Get` as synchronous or asynchronous *events*, they are all defined as *join messages* which may support both kinds of synchrony (this is explained in more detail below).
– The `Buffer` actor overrides `act` and awaits incoming messages by means of `receive`. Note that it is still possible for the actor to serve regular messages within the `receive` block. In fact, regular actor messages can be regarded as unary join patterns.

We illustrate below how the buffer actor can be used as a coordinator between a consumer and a producer actor. The producer sends an asynchronous `Put` message while the consumer awaits the reply to a `Get` message by invoking it synchronously (using `!?`)[5].

```scala
val buffer = new Buffer; buffer.start()
val prod = actor { buffer ! Put(42) }
val cons = actor { (buffer !? Get()) match { case x:Int => /* process x */ } }
```

By applying joins to actors, the synchronization dependencies between `Get` and `Put` can be specified declaratively by the buffer actor. The actor will receive `Get` and `Put` messages by queuing them in its mailbox. Only when all of the messages specified in the join pattern have been received is the body executed by the actor. Before processing the body, the actor atomically removes all of the participating messages from its mailbox. Replies may be sent to any or all of the messages participating in the join pattern. This is similar to the way replies are sent to events in the thread-based Joins library described previously.

Contrary to the way events are defined in the thread-based joins library, an actor does not explicitly define a join message to be synchronous or asynchronous. We say that join messages are "synchronization-agnostic" because they can be used in different synchronization modes between the sender and receiver actors. However, when they are used in a particular join-pattern, the sender and receiver actors have to agree upon a valid synchronization mode. In the previous example, the `Put` join message was sent asynchronously, while the `Get` join message was sent synchronously. In the body of a join pattern, the receiver actor replied to `Get`, but not to `Put`.

The advantage of making join messages synchronization agnostic is that they can be used in arbitrary synchronization modes, including more advanced synchronization modes such as ABCL's future-type message sending [19] or Salsa's token-passing continuations [16]. Every join message instance has an associated *reply destination*, which is an output channel on which processes may listen for possible replies to the message. How the reply to a message is processed is determined by the way the message was sent. For example, if the message was sent purely asynchronously, the reply is discarded; if it was sent synchronously, the reply awakes the sender. If it was sent using a future-type message send, the reply resolves the future.

4   Integrating Joins and Extensible Pattern Matching

Our implementation technique for joins is unique in the way events interact with an extensible pattern matching mechanism. We explain the technique using a concrete implementation in Scala. However, we expect that implementations based on, e.g., the active patterns of F# [15] would not be much different. In the following we first talk about pattern matching in Scala. After that we dive into the implementation of events which crucially depends on properties of Scala's extensible pattern matching. Finally, we highlight how joins have been integrated into Scala's actor framework.

---

[5] Note that the `Get` message has return type `Any`. The type of the argument values is recovered by pattern matching on the result, as shown in the example.

Partial Functions In the previous section we used the `join { ... }` construct to declare a set of join patterns. It has the following form:

```
join {
   case pat₁ ⇒ body₁
   ...
   case patₙ ⇒ bodyₙ
}
```

The patterns pat_i consist of a linear combination of events evt_1 & ... & evt_m. Threads synchronize over a join pattern by invoking one or several of the events listed in a pattern pat_i. When all events occurring in pat_i have been invoked, the join pattern matches, and its corresponding join $body_i$ is executed.

In Scala, the pattern matching expression inside braces is treated as a first-class value that is passed as an argument to the `join` function. The argument's type is an instance of `PartialFunction`, which is a subclass of `Function1`, the class of unary functions. The two classes are defined as follows.

```
abstract class Function1[A, B] {
   def apply(x: A): B }
abstract class PartialFunction[A, B] extends Function1[A, B] {
   def isDefinedAt(x: A): Boolean }
```

Functions are objects which have an `apply` method. Partial functions are objects which have in addition a method `isDefinedAt` which tests whether a function is defined for a given argument. Both classes are parameterized; the first type parameter `A` indicates the function's argument type and the second type parameter `B` indicates its result type.

A pattern matching expression { **case** p_1 => e_1; ...; **case** p_n => e_n } is then a partial function whose methods are defined as follows.

- The `isDefinedAt` method returns `true` if one of the patterns p_i matches the argument, `false` otherwise.
- The `apply` method returns the value e_i for the first pattern p_i that matches its argument. If none of the patterns match, a `MatchError` exception is thrown.

Join patterns as partial functions. Whenever a thread invokes an event e, each join pattern in which e occurs has to be checked for a potential match. Therefore, events have to be associated with the set of join patterns in which they participate. As shown before, this set of join patterns is represented as a partial function. Invoking `join(pats)` associates each event occurring in the set `pats` of join patterns with `pats`.

When a thread invokes an event, the `isDefinedAt` method of `pats` is used to check whether any of the associated join patterns match. If yes, the corresponding join body is executed by invoking the `apply` method of `pats`. A question remains: what argument is passed to `isDefinedAt` and `apply`, respectively? To answer this question, consider the simple buffer example from the previous section. It declares the following join pattern:

```
join { case Get() & Put(x) ⇒ Get reply x }
```

Assume that no events have been invoked before, and a thread t invokes the `Get` event to remove an element from the buffer. Clearly, the join pattern does not match, which

causes t to block since Get is a synchronous event (more on synchronous events later). Assume that after thread t has gone to sleep, another thread s adds an element to the buffer by invoking the Put event. Now, we want the join pattern to match since both events have been invoked. However, the result of the matching does not only depend on the event that was last invoked but also on the fact that *other events* have been invoked previously. Therefore, it is *not* sufficient to simply pass a Put message to the isDefinedAt method of the partial function the represents the join patterns. Instead, when the Put event is invoked, the Get event has to somehow "pretend" to also match, even though it has nothing to do with the current event. While previous invocations can simply be buffered inside the events, it is non-trivial to make the pattern matcher actually consult this information during the matching, and "customize" the matching results based on this information. To achieve this customization we use extensible pattern matching.

Extensible Pattern Matching Emir et al. [4] recently introduced *extractors* for Scala that provide representation independence for objects used in patterns. Extractors play a role similar to *views* in functional programming languages [18,12] in that they allow conversions from one data type to another to be applied implicitly during pattern matching. As a simple example, consider the following object that can be used to match even numbers:

```scala
object Twice {
  def apply(x: Int) = x*2
  def unapply(z: Int) = if (z%2 == 0) Some(z/2) else None }
```

Objects with apply methods are uniformly treated as functions in Scala. When the function invocation syntax Twice(x) is used, Scala implicitly calls Twice.apply(x). The unapply method in Twice reverses the construction in a pattern match. It tests its integer argument z. If z is even, it returns Some(z/2). If it is odd, it returns None. The Twice object can be used in a pattern match as follows:

```scala
val x = Twice(21)
x match {
  case Twice(y) => println(x+" is two times "+y)
  case _ => println("x is odd") }
```

To see where the unapply method comes into play, consider the match against Twice(y). First, the value to be matched (x in the above example) is passed as argument to the unapply method of Twice. This results in an optional value which is matched subsequently[6]. The preceding example is expanded as follows:

```scala
val x = Twice.apply(21)
Twice.unapply(x) match {
  case Some(y) => println(x+" is two times "+y)
  case None => println("x is odd") }
```

---

[6] The optional value is of parameterized type Option[T] that has the two subclasses Some[T](x: T) and None.

Extractor patterns with more than one argument correspond to unapply methods returning an optional tuple. Nullary extractor patterns correspond to unapply methods returning a boolean.

In the following we show how extractors can be used to implement the matching semantics of join patterns. In essence, we define appropriate unapply methods for events which get implicitly called during the matching.

Matching Join Patterns As shown previously, a set of join patterns is represented as a partial function. Its isDefinedAt method is used to find out whether one of the join patterns matches. In the following we are going to explain the code that the Scala compiler produces for the body of this method. Let us revisit the join pattern that we have seen in the previous section:

```
Get() & Put(x)
```

In our library, the & operator is an extractor (see previous section) that defines an unapply method; therefore, the Scala compiler produces the following matching code:

```
&.unapply(m) match {
  case Some((u, v)) =>
    u match {
      case Get() => v match {
        case Put(x) => true
        case _ => false }
      case _ => false }
  case None => false }
```

We defer a discussion of the argument m that is passed to the & operator. For now, it is important to understand the general scheme of the matching process. Basically, calling the unapply method of the & operator produces a pair of intermediate results wrapped in Some. Standard pattern matching decomposes this pair into the variables u and v. These variables, in turn, are matched against the events Get and Put. Only if both of them match, the overall pattern matches.

Since the & operator is left-associative, matching more than two events proceeds by first calling the unapply methods of all the & operators from right to left, and then matching the intermediate results with the corresponding events from left to right.

Since events are objects that have an unapply method, we can expand the code further:

```
&.unapply(m) match {
  case Some((u, v)) =>
    Get.unapply(u) match {
      case true => Put.unapply(v) match {
        case Some(x) => true
        case None => false }
      case false => false }
  case None => false }
```

As we can see, the intermediate results produced by the unapply method of the & operator are passed as arguments to the unapply methods of the corresponding events. Since the Get event is parameterless, its unapply method returns a Boolean, telling whether it matches or not. The Put event, on the other hand, takes a parameter; when the pattern matches, this parameter gets bound to a concrete value that is produced by the unapply method.

The unapply method of a parameterless event such as Get essentially checks whether it has been invoked previously. The unapply method of an event that takes parameters such as Put returns the argument of a previous invocation (wrapped in Some), or signals failure if there is no previous invocation. In both cases, previous invocations have to be buffered inside the event.

Firing join patterns. As mentioned before, executing the right-hand side of a pattern that is part of a partial function amounts to invoking the apply method of that partial function. Basically, this repeats the matching process, thereby binding any pattern variables to concrete values in the pattern body. When firing a join pattern, the events' unapply methods have to dequeue the corresponding invocations from their buffers. In contrast, invoking isDefinedAt does not have any effect on the state of the invocation buffers. To signal to the events in which context their unapply methods are invoked, we therefore need some way to propagate out-of-band information through the matching. For this, we use the argument that is passed to the isDefinedAt and apply methods of the partial function. The & operator propagates this information verbatim to its two children. Eventually, this information is passed to the events' unapply methods.

4.1 Implementation Details

Events. Events are represented as classes that contain queues to buffer invocations. The Event class is the super class of all synchronous and asynchronous events[7]:

```scala
abstract class Event[R, Arg](owner: Joins) {
  val argQ = new Queue[Arg]
  def test(): R
  def apply(arg: Arg): R = synchronized { argQ += arg; test() }
  def unapply(isDryRun: Boolean): Option[Arg] =
    if (isDryRun && !argQ.isEmpty)
      Some(argQ.front)
    else if (!isDryRun && owner.takes(tag))
      Some(argQ.dequeue())
    else None
}
```

Events have a unique owner which is an instance of the Joins class. This class provides the join method that we used in the buffer example to declare a set of join patterns. An event can appear in several join patterns declared by its owner. The test method is

---

[7] In our actual implementation the fact whether an event is parameterless is factored out for efficiency. Due to lack of space, we show a simplified class hierarchy.

used to run synchronization-specific code when the event is invoked. The Event class has two type parameters R and Arg that indicate the result type and parameter type of event invocations, respectively. Whenever the event is invoked via its apply method, we append the provided argument to the argQ. In the unapply method we test whether matching occurs during a dry run. If it does not, we ask the owner whether the event belongs to a matching join pattern in which case an event invocation is dequeued.

Synchronous events are implemented as follows:

```
abstract class SyncEvent[R, Arg] extends Event[R, Arg] {
  val waitQ = new Queue[SyncVar[R]]
  def test(): R = { val res = new SyncVar[R]
    waitQ += res; owner.test(); res.get }
  def reply(res: R) { waitQ.dequeue().set(res) }
}
```

Synchronous events contain a logical queue of waiting threads, waitQ, which is implemented using the implicit wait set of synchronous variables[8]. The test method is run whenever the event is invoked. It creates a new SyncVar and appends it to the waitQ. Then, the owner's test method is invoked to check whether the event invocation triggers a complete join pattern. After that, the current thread waits for the SyncVar to become initialized by accessing it. If the owner detects (during owner.test()) that a join pattern triggers, it will apply the join, thereby re-executing the pattern match (binding variables etc.) and running the join body. Inside the body, synchronous events are replied to by invoking their reply method. Replying means dequeuing a SyncVar and setting its value to the supplied argument.

4.2   Implementation of Actor-based Joins

Actor-based joins integrate with Scala's pattern matching in essentially the same way as the thread-based joins, making both implementations very similar. We highlight how joins are integrated into the actor library, and how reply destinations are supported.

In the Scala actors library, receive is a method that takes a PartialFunction as a sole argument, similar to the join method defined previously. To make receive aware of join patterns, the abstract JoinActor class overrides these methods by wrapping the partial function into a specialized partial function that understands join messages. JoinActor also overrides send to set the reply destination of a join message. Message sends such as a!msg are interpreted as calls to a's send method.

```
abstract class JoinActor extends Actor {
  override def receive[R](f: PartialFunction[Any, R]): R =
    super.receive(new JoinPatterns(f))
  override def send(msg: Any, replyTo: OutputChannel[Any]) {
    setReplyDest(msg, replyTo)
    super.send(msg, replyTo) }
  def setReplyDest(msg: Any, replyTo: OutputChannel[Any]) {...} }
```

---

[8] A SyncVar is an atomically updatable reference cell; it blocks threads trying to access an uninitialized cell.

`JoinPatterns` is a special partial function that detects whether its argument message is a join message. If it is, then the argument message is transformed to include out-of-band information that will be passed to the pattern matcher, as is the case for events in the thread-based joins library. The boolean argument passed to the `checkJoinMessage` method indicates to the pattern matcher whether or not join message arguments should be dequeued upon successful pattern matching. If the `msg` argument is not a join message, `checkJoinMessage` passes the original message to the pattern matcher unchanged, enabling regular actor messages to be processed as normal.

```
class JoinPatterns[R](f: PartialFunction[Any, R])
  extends PartialFunction[Any, R] {
  def asJoinMessage(msg: Any, isDryRun: Boolean): Any =
    ...
  override def isDefinedAt(msg: Any) =
    f.isDefinedAt(checkJoinMessage(msg, true))
  override def apply(msg: Any) =
    f(checkJoinMessage(msg, false))
}
```

Recall from the implementation of synchronous events that thread-based joins used constructs such as `SyncVars` to synchronize the sender of an event with the receiver. Actor-based joins do not use such constructs. In order to synchronize sender and receiver, every join message has a reply destination (which is an `OutputChannel`, set when the message is sent in the actor's send method) on which a sender may listen for replies. The `reply` method of a `JoinMessage` simply forwards its argument value to this encapsulated reply destination. This wakes up an actor that performed a synchronous send (`a!?msg`) or that was waiting on a future (`a!!msg`).

5 Discussion and Related Work

Benton et al. [2] note that supporting general guards in join patterns is difficult to implement efficiently as it requires testing all possible combinations of queued messages to find a match. Side effects pose another problem. Benton et al. suggest a restricted language for guards to overcome these issues. However, to the best of our knowledge, there is currently no joins framework that supports a sufficiently restrictive yet expressive guard language to implement efficient guarded joins. Our current implementation does not handle general guards, although they are permitted in Scala's pattern syntax [11]. We find that guards often help at the interface to a component that uses private messages to represent values satisfying a guard, as in the following example:

```
join { case Put(x) if (x > 0) => this.PositivePut(x)
       case PositivePut(x) & Get() => Get reply x }
```

$C\omega$ [2] is a language extension of C# supporting *chords*, linear combinations of methods. In contrast to Scala Joins, $C\omega$ allows at most one synchronous method in a chord. The thread invoking this method is the thread that eventually executes the chord's body. The benefits of $C\omega$ as a language extension over Scala Joins are that chords can be enforced to be well-formed and that their matching code can be optimized ahead of time.

In Scala Joins, the joins are only analyzed at pattern-matching time. The benefit of Scala Joins as a library extension is that it provides more flexibility, such as experimenting with multiple synchronous events. Russo's `Joins` library [13] exploits the expressiveness of C# 2.0's generics to implement Cω's synchronization constructs. Piggy-backing on an existing variable binding mechanism allows us to avoid problems with `Joins`' delegates where the order in which arguments are passed is merely conventional. Scala's implicit (constructor) arguments also help to alleviate some of the initialization boilerplate. CCR [3] is a C# library for asynchronous concurrency that supports join patterns without synchronous components. Join bodies are scheduled for execution in a thread pool. Our library integrates with JVM threads using synchronous variables, and supports event-based programming through its integration with Scala Actors. Singh [14] shows how a small set of higher-order combinators based on Haskell's software transactional memory (STM) can encode expressive join patterns. Salsa [16] is a language extension of Java supporting actors. In Salsa, actors may synchronize in an event-driven way upon the arrival of multiple messages by means of a *join continuation*. However, join continuations only allow an actor to synchronize on gathering replies to previously sent messages. Using joins, Scala actors may synchronize on any incoming message.

6 Conclusion

We presented a novel implementation of join patterns based on extensible pattern matching constructs of languages such as Scala and F#. The embedding into general pattern matching provides expressive features such as nested patterns and guards for free. The resulting programs are often as concise as if written in more specialized language extensions. We implemented our approach as a Scala library that supports Ada-style rendezvous and constraints and furthermore integrated it with the Scala Actors event-based concurrency framework without changing the syntax and semantics of existing programs.

References

1. Joe Armstrong, Robert Virding, Claes Wikström, and Mike Williams. *Concurrent Programming in Erlang, Second Edition*. Prentice-Hall, 1996.
2. Nick Benton, Luca Cardelli, and Cédric Fournet. Modern concurrency abstractions for C#. *ACM Trans. Program. Lang. Syst*, 26(5):769–804, 2004.
3. Georgio Chrysanthakopoulos and Satnam Singh. An asynchronous messaging library for C#. In *Proc. SCOOL Workshop, OOPSLA*, 2005.
4. Burak Emir, Martin Odersky, and John Williams. Matching objects with patterns. In Erik Ernst, editor, *Proc. ECOOP*, volume 4609 of *LNCS*, pages 273–298. Springer, 2007.
5. Cédric Fournet, Fabrice Le Fessant, Luc Maranget, and Alan Schmitt. JoCaml: A language for concurrent distributed and mobile programming. In *Advanced Functional Programming*, volume 2638 of *LNCS*, pages 129–158. Springer, 2002.
6. Cédric Fournet and Georges Gonthier. The reflexive chemical abstract machine and the join-calculus. In *Proc. POPL*, pages 372–385. ACM, January 1996.
7. Cédric Fournet, Georges Gonthier, Jean-Jacques Lévy, Luc Maranget, and Didier Rémy. A Calculus of Mobile Agents. In *CONCUR*, pages 406–421. Springer-Verlag, August 1996.

8. Philipp Haller and Martin Odersky. Actors that Unify Threads and Events. In *International Conference on Coordination Models and Languages*, LNCS, 2007.

9. Robert H. Halstead, Jr. Multilisp: a language for concurrent symbolic computation. *ACM Trans. Program. Lang. Syst.*, 7(4):501–538, 1985.

10. Martin Odersky. Functional Nets. In *European Symposium on Programming 2000*, Lecture Notes in Computer Science. Springer Verlag, 2000. Invited paper.

11. Martin Odersky and al. An Overview of the Scala Programming Language. Technical Report IC/2004/64, EPFL, Lausanne, Switzerland, 2004.

12. C. Okasaki. Views for Standard ML, 1998.

13. Claudio V. Russo. The Joins concurrency library. In *PADL*, pages 260–274, 2007.

14. Satnam Singh. Higher-order combinators for join patterns using STM. In *Proc. TRANSACT Workshop, OOPSLA*, 2006.

15. Don Syme, Gregory Neverov, and James Margetson. Extensible Pattern Matching via a Lightweight Language Extension. In *Proc. ICFP*, 2007.

16. Carlos Varela and Gul Agha. Programming dynamically reconfigurable open systems with SALSA. *ACM SIGPLAN Notices*, 36(12):20–34, 2001.

17. G.S. von Itzstein and David Kearney. Join Java: An alternative concurrency semantic for Java. Technical Report ACRC-01-001, University of South Australia, 2001.

18. Philip Wadler. Views: A way for pattern matching to cohabit with data abstraction. In *POPL*, pages 307–313, 1987.

19. Akinori Yonezawa, Jean-Pierre Briot, and Etsuya Shibayama. Object-oriented concurrent programming in ABCL/1. In *Proc. OOPSLA*, pages 258–268, 1986.