Functional at Scale

**APPLYING
FUNCTIONAL
PROGRAMMING
PRINCIPLES TO
DISTRIBUTED
COMPUTING
PROJECTS**

MARIUS ERIKSEN, TWITTER

Modern server software is demanding to develop and operate: it must be available at all times and in all locations; it must reply within milliseconds to user requests; it must respond quickly to capacity demands; it must process a lot of data and even more traffic; it must adapt quickly to changing product needs; and in many cases it must accommodate a large engineering organization, its many engineers the proverbial cooks in a big, messy kitchen.

What's more, the best computers for these applications—whether you call them clouds, data centers, or warehouse computers—are really bad. They are complex and unreliable, and prone to partial failures. They have asynchronous interconnects, and deep memory hierarchies, and leave a lot of room for operator error.¹

Cloud computing thus forces us to confront the full complexity of distributed computing, where seemingly simple problems require complicated solutions. While much of this complexity is inherent—the very nature of the problem precludes simpler solutions—much of it is also incidental, a simple consequence of using tools unfit for the purpose.

At Twitter, we use ideas from functional programming to tackle many of the complexities of modern server software, primarily through the use of higher-order functions and effects. *Higher-order functions*, or those that return other functions, let us combine simpler functions to define more complex ones, building up application functionality in a piecemeal fashion. *Effects* are values that represent some side-effecting operation; they are used in conjunction with higher-order functions to build complex effects from simpler ones.⁷

Higher-order functions and effects help build scalable software in two ways: first, building complex software from simple parts makes it easy to understand, test, reuse, and replace individual components; second, effects make side-effecting operations tractable, promoting modularity and good separation of concerns.

This article explores three specific abstractions that follow this style of functional programming: *futures* are effects that represent asynchronous operations; *services* are functions that represent service boundaries; and *filters* are functions that encapsulate application-independent behavior. In turn, higher-order functions provide the glue used to combine these to create complex systems from simple parts.

Futures, services, and filters are thus combined to build server software in a piecemeal fashion. They let programmers build up complex software while preserving their ability to reason about the correctness of its constituent parts.

By consistently applying these principles, programmers

*“Concurrent
programs
wait faster.”
—Tony Hoare*

can construct systems that are at once simpler, more flexible, and performant.

CONCURRENT PROGRAMMING WITH FUTURES

Concurrency is a central topic in server-software development.¹² Two sources of concurrency prevail in this type of software. First, scale implies concurrency. For example, a search engine may split its index into many small pieces (shards) so that the entire corpus can fit in main memory. To satisfy queries efficiently, all shards must be queried concurrently. Second, communication between servers is asynchronous and must be handled concurrently for efficiency and safety.

Concurrent programming is traditionally approached by employing threads and locks.³ Threads furnish the programmer with concurrent threads of execution, while locks coordinate the sharing of (mutable) data across multiple threads.

In practice, threads and locks are notoriously difficult to get right.⁹ They are hard to reason about, and they are a stubborn cause of nasty bugs. What’s more, they are difficult to compose: you cannot safely and arbitrarily combine a set of threads and locks to construct new functionality. Their semantics of computation are wrapped up in the mechanics of managing concurrency.

At Twitter, we instead structure concurrent programs around futures. A future is an effect that represents the result of an asynchronous operation. It’s a type of reference cell that can be in one of three states: *incomplete*, when the future has not yet taken on a value; *completed with a value*, when the future holds the result of

a successful operation; and *completed with a failure*, when the operation has failed. Futures can undergo at most one state transition: from the incomplete state to either the success or failure state

In the following example, using Scala, the future `count` represents the result of an integer-valued operation. We *respond* to the future's completion directly: the block of code after `respond` is a callback that is invoked when the future has completed. (As you'll see shortly, we rarely respond directly to future completions in this way.)

```
val count: Future[Int] = getCount()
count.respond {
  case Return(value) =>
    println(s"The count was $value")
  case Throw(exc) =>
    println(s"getCount failed with $exc")
}
```

Futures represent just about every asynchronous operation in Twitter systems: RPC (remote procedure call), timeout, reading a file from disk, receiving the next event from an event stream.

With the help of a set of higher-order functions (called *combinators*), futures can be combined freely to express more complex operations. These combinations usually fall into one of two composition categories: sequential or concurrent.

Sequential composition permits defining a future as a function of another, such that the two are executed sequentially. This is useful where data dependency exists

between two operations: the result of the first future is needed to compute the second future. For example, when a user sends a Tweet, we first need to see if that user is within the hourly rate limits before writing the Tweet to a database. In the following example, the future `done` represents this composite operation. (For historical reasons, the sequencing combinator is called `flatMap`.)

```
def isRateLimited(user: Long): Future[Boolean] = ...
def writeTweet(user: Long, tweet: String): Future[Unit] = ...

val user = 12L
val tweet: String = "just setting up my twitter"

val done: Future[Unit] =
  isRateLimited(user).flatMap {
    case true => Future.exception(new RateLimitingError)
    case false => writeTweet(user, tweet)
  }
```

This also shows how failures are expressed in futures: `Future.exception` returns a future that has already completed in a failure state. In the case of rate limiting, `done` becomes a failed future (with the exception `RateLimitingError`). Failure short-circuits any further composition: if, in the previous example, the future returned by `isRateLimited(user)` fails, then `done` is immediately failed; the closure passed to `flatMap` is not run.

Another set of combinators defines concurrent composition, allowing multiple futures to be combined

when no data dependencies exist among them. Concurrent combinators turn a list of futures into a future of a list of values. For example, you may have a list of futures representing RPCs to all the shards of a search index. The concurrent combinator `collect` turns this list of futures into a future of the list of results.

```
val results: List[Future[String]] = ...  
val all: Future[List[String]] =  
  Future.collect(results)
```

Independent futures are executed concurrently by default; execution is sequenced only where data dependencies exist.

Future combinators never modify the underlying future; instead, they return a *new* future that represents the composite operation. This is an important tool for reasoning: composite operations do not change the behavior of their constituent parts. Indeed, a single future can be used and reused in many different compositions.

Let's modify the earlier `getCount` example to see how composition allows building complex behavior piecemeal. In distributed systems, it is often preferable to degrade gracefully (e.g., where a default value or guess may exist) than to fail an entire operation.⁸ This can be implemented by using a timeout for the `getCount` operation, and, upon failure, returning a default value of 0. This behavior can be expressed as a compound operation among different futures. Specifically, you want the future that represents the winner of a timeout-with-default and the `getCount` operation.

```
val count: Future[Int] = getCount()  
// Future.sleep(x) returns a future which completes  
// after the given amount of time.  
val timeout: Future[Unit] = Future.sleep(5.seconds)  
  
// The map method on Future constructs a new Future  
// which, after successful completion, applies the given  
// function to the result. It returns a new Future  
// representing this composite operation. In this case,  
// we simply return a default value of 0 after the timeout.  
val default: Future[Int] = timeout.map(unit => 0)  
  
// Select composes two Futures, returning a new  
// Future which represents the first future to complete.  
val finalCount: Future[Int] = Future.select(count, default)
```

The future `finalCount` now represents the composite operation as described. This example has a number of notable features. First, we have created a composite operation using simple, underlying parts—futures and functions. Second, the constituent futures preserve their semantics under composition—their behavior does not change, and they may be used in multiple different compositions. Third, nothing has been said about the mechanics of execution; instead, the composite computation is expressed as the combination of a number of underlying parts. No threads were explicitly created,

nor was there any explicit communication between them—it is all implied by data dependencies.

This neatly illustrates how futures liberate the application's semantics (what is computed) from its mechanics (how it is computed). The programmer composes concurrent operations but needn't specify how they are scheduled or how values are communicated. This is a good separation of concerns: application logic is not entangled with the minutiae of runtime concerns.

Futures can sometimes free programmers from having to use locks. Where data dependencies are witnessed by composition of futures, the implementation is responsible for concurrency control. Put another way, futures can be composed into a dependency graph that is executed in the manner of data-flow programming.¹¹ While we need to resort to explicit concurrency control from time to time, a large set of common use cases are handled directly by the use of futures.

At Twitter, we implement the machinery required to make concurrent execution with futures work in our open-source Finagle^{4,5} and Util⁶ libraries. These take care of mapping execution onto OS threads through a pluggable scheduler mechanism. Some teams at Twitter have used this capacity to construct scheduling strategies that better match their problem domain and its attendant tradeoffs. We have also used this capability to add features such as maintaining runtime statistics and Dapper-style RPC tracing to our systems, without changing any existing APIs or modifying existing user code.

“We should have some ways of coupling programs like garden hose—screw in another segment when it becomes necessary to massage data in another way. This is the way of I/O also.”
—Doug McIlroy

PROGRAMMING WITH SERVICES AND FILTERS

Modern server software abounds with essential complexity—that which is inherent to the problem and cannot be avoided—as well as complexity of the more self-inflicted variety (did we really need that product feature)?

Anyhow, since we can’t seem to keep our applications simple, we must instead cope with their complexities.

Thus, the modern software engineering practice is centered on how to contain and manage complexity—to package it up in ways that allow us to reason about our application’s behavior. In this endeavor the goal is to balance simplicity and clarity with reusability and modularity.

What’s more, server software must account for the realities of distributed systems. For example, a search engine might simply omit results from a failing index shard so that it can return a partial result rather than failing the query in its entirety (as seen in the `getCount` example). In this case, the user would not usually be able to tell the difference—the search engine is still useful without a complete set of results. Applying application-level knowledge in this way is often essential to creating a truly resilient application.

Distributed applications are therefore organized around *services* and *filters*. Services represent RPC service endpoints. They are a function of type `A => Future[R]` (i.e., a function that takes a single argument of type `A` and returns an `R`-typed future). These functions are asynchronous: they return immediately; deferred actions are captured by the returned future. Services are symmetric: RPC clients call services to dispatch RPC calls;

servers implement services to handle them.

The following example defines a simple dictionary service that looks up a key in a map, which is stored in memory. The service takes a `String` argument and returns a `String` value, or else null if the key is not found.

```
// This is the dictionary data, encoded in a
// string-to-string map.
val map: Map[String, String] = Map(...)

val dictionary: Service[String, String] =
  // Construct a new service from an
  // anonymous function, invoked
  // for every request issued.
  Service(key: String => {
    // Return the lookup results. (Map.get
    // returns an optional string.) Future.value
    // creates a Future that is already
    // satisfied with the given value.
    if (map.contains(key))
      Future.value(map(key))
    else
      Future.value(null)
  })
```

Note that the service, `dictionary`, is again a value like any other. Once defined, it can be made available for dispatch over the network through a favorite RPC protocol. On the server, services are exported as

```
Rpc.serve(dictionary, ":8080")
```

while clients can bind and call remote instances of the service:

```
val service: Service[String, String] =  
  Rpc.bind("server:8080")  
  
val result: Future[String] =  
  service("key")
```

Observe the symmetry between client and server: both are conversing in terms of services, which are location-transparent; it is an implementation detail that one is implemented locally while the other is bound remotely.

Services are easily composed using ordinary functional composition. For example, the following service performs a scatter-gather lookup across multiple dictionary services. The dictionary can therefore be distributed over multiple shards.

```
def multipleLookup(services: Seq[Service[String, String]])  
  : Service[String, String] =  
  Service(key => {  
    val results: Seq[Future[String]] =  
      services.map(_.apply(key))  
    val all: Future[Seq[String]] =  
      Future.collect(results)  
    all.map { results: Seq[String] =>  
      results.indexWhere(_ != null) match {  
        case -1 => null  
        case i => results(i)  
      }  
    }  
  })
```

This example has a lot going on. First, each underlying service is called with the desired key, obtaining a sequence of results—all futures. `Future.collect` composes futures concurrently, turning a sequence of futures into a single one containing a sequence of the values of the (successful) completion of the constituent futures. The code then looks for the first non-null result and returns it.

Filters are also asynchronous functions that are combined with services to modify their behavior. Their type is $(A, \text{Service}[A, R]) \Rightarrow \text{Future}[R]$ (i.e., a function with two arguments: a value of type `A` and a service `Service[A, R]`); the filter then returns a future of that service's return type, `R`. The filter's type indicates that it is responsible for satisfying a request, given a service. It can thus modify both how the request is dispatched to the service (e.g., it can modify it) and how it is returned (e.g., add a timeout to the returned future).

Filters are used to implement functionality such as timeouts and retries, failure-handling strategies, and authentication. Filters may be combined to form compound filters—for example, a filter that implements retry logic can be combined with a filter implementing timeouts. Finally, filters are combined with a service to create a new service with modified behavior.

Building on the previous example, the following is a filter that downgrades failures from lookup services to null. This kind of behavior is often useful when constructing

resilient services—it's sometimes better to return partial results than to fail altogether.

```
val downgradeToNull: Filter[String, String] =  
  Filter((key, service) => {  
    service.apply(key).transform {  
      case Return(value) => value  
      case Throw(exception) => null  
    }  
  })
```

Another filter short-circuits requests for useless keys (these are known as *stopwords* in search engines), so that they don't inflict needless load on the underlying service:

```
val stopwords: Set[String] ...  
val stopwordFilter: Filter[String, String] =  
  Filter((key, service) => {  
    if (stopwords.contains(key))  
      Future.value(null)  
    else  
      service.apply(key)  
  })
```

Finally, the two filters are combined and then applied to all of the services.

```
val resiliencyFilter: Filter[String, String] =
  stopwordFilter.andThen(downgradeToNull)

def resilientMultipleLookup(services: Service[String, String]) = {
  val resilientServices =
    services.map { service => resiliencyFilter.andThen(service) }
  multipleLookup(resilientServices)
}

val resilientService: Service[String, String] =
  resilientMultipleLookup(
    Rpc.bind("server1:8080"),
    Rpc.bind("server2:8080"),
    ...
  )
```

The dictionary lookup service `resilientService` implements scatter-gather lookup in a resilient fashion. The functionality was built from individual well-defined components that are composed together to create a service that behaves in a desirable way.

As with futures, combining filters and services does not change the underlying, constituent components. It creates a new service with new behavior but does not change the meaning of either underlying filter or service. This again enhances reasoning since the constituent components stand on their own; we need not reason about their interactions after they are combined.

Futures, services, and filters form the foundation upon which server software is built at Twitter. Together

*“Don’t tie the hands
of the implementer.”*

—Martin Rinard

they support both modularity and reuse: we define applications and behaviors independently, composing them together as required. While their application is pervasive, two examples nicely illustrate their power. First, we implemented an RPC tracing system à la Dapper¹⁰ as a set of filters, requiring no changes to application code. Second, we implemented backup requests² as a small, self-contained filter.

CONCLUSION

Functional programming promotes thinking about building complex behavior out of simple parts, using higher-order functions and effects to glue them together. At Twitter we have applied this line of thinking to distributed computing, structuring systems around a set of core abstractions that express asynchrony through effects and that are composable. This allows building complex systems from components with simple semantics that, preserved under composition, makes it easier to reason about the system as a whole.

This approach leads to simpler and more modular systems. These systems promote a good separation of concerns and enhance flexibility, while at the same time permitting efficient implementations.

Functional programming has thus furnished essential tools for managing the complexity that is inherent in modern software—untying the hands of the implementer.

Acknowledgments

Jake Donham, Erik Meijer, Mark Compton, Terry Coatta, Steve Jenson, Kevin Oliver, Ruben Oanta, and Oliver Gould

provided excellent feedback and guidance on earlier drafts of this article. Early versions of the abstractions discussed here were designed together with Nick Kallen; numerous people at Twitter and in the open-source community have since worked to improve, expand, and solidify them.

References

1. Barroso, L. A., Clidaras, J., Hölzle, U. 2013. The datacenter as a computer: an introduction to the design of warehouse-scale machines. *Synthesis Lectures on Computer Architecture* 8(3): 1-154.
2. Dean, J., Barroso, L. A. 2013. The tail at scale. *Communications of the ACM* 56(2): 74-80.
3. Dijkstra, E. W. 1965. Solution of a problem in concurrent programming control. *Communications of the ACM* 8(9): 569.
4. Eriksen, M. 2013. Your server as a function. In *Proceedings of the Seventh Workshop on Programming Languages and Operating Systems*, ACM (November): 5.
5. Eriksen, M., Kallen, N. 2010. Finagle; <http://twitter.github.com/finagle>.
6. Eriksen, M., Kallen, N. 2010. Util; <http://twitter.github.com/util>.
7. Hughes, J. 1989. Why functional programming matters. *The Computer Journal* 32(2): 98-107.
8. Netflix. Hystrix; <https://github.com/Netflix/Hystrix>.
9. Ousterhout, J. 1996. Why threads are a bad idea (for most purposes). In presentation given at the Usenix Annual Technical Conference, vol. 5.
10. Sigelman, B.H., Barroso, L.A., Burrows, M., Stephenson, P., Plakal, M., Beaver, D., Jaspán, S., Shanbhag, C., 2010.

Dapper, a large-scale distributed systems tracing infrastructure. Technical report, Google: 36.

11. Smolka, G. 1995. The Oz programming model. In *Computer Science Today*, ed. Jan van Leeuwen, Lecture Notes in Computer Science, Volume 1000: 324-343. Berlin: Springer-Verlag.
12. Sutter, H. 2005. The free lunch is over: a fundamental turn toward concurrency in software. *Dr. Dobbs's Journal* 30[3]: 202-210.

Marius Eriksen is a principal engineer in Twitter's systems infrastructure group. He works on all aspects of distributed systems and server software and is currently working on data management and integration systems; he also chairs Twitter's architecture group. You can reach him at marius@twitter.com or [@marius](https://twitter.com/marius) on Twitter.

Copyright © 2016 held by owner/author. Publication rights licensed to ACM.