Serverless Cloud Computing Beyond FaaS: Programming Models and Abstractions

Philipp Haller

KTH Royal Institute of Technology Stockholm, Sweden



2nd Vienna Software Seminar (VSS) Vienna, Austria, Aug 29, 2019

Philipp Haller

Scala Background

- 2005-2014 Scala language team
 - 2012-2014 Typesafe, Inc.
- Co-author Scala language specification



- 2019: ACM SIGPLAN Programming Languages Software Award for Scala

Core contributors: **Martin Odersky**, Adriaan Moors, Aleksandar Prokopec, Heather Miller, Iulian Dragos, Nada Amin, Philipp Haller, Sebastien Doeraene, Tiark Rompf

Scala Actors and Akka

LEASES

JUN21

2006

JUN5

2007

OCT 13 2007

DEC 19

2007

00

Philipp Haller creates Scala Actors (the original standard library actors). His work becomes a major influence to Akka and the main reason for Jonas Bonér to choose Scala as the platform for Akka.

First commit: https://github.com/scala/scala/commit/0d8b14c6055e76c0bff3b65d0f428d711abe1f5a

CONFERENCE MPER Actors that Unify Threads and Events

Haller, Philipp: Odersky, Martin Editors: Vitek, Jan: Murphy, Anty L

Presented at: International Conference on Coordination Models and Languages, Paptos, Cyprus, 5-8 June 2007 Published itt: Proceedings of the Shi International Conference on Coordination Models and Languages (COORDINATION), p. 171-190 Series: Letture Notes in Computer Science (LNCS) 4467 : Somarr, 2007

There is an impedance mismatch between message-passing concurrency and virtual machines, such as the JVM. Whis usually map their threads to heavyweight OS processes. Without a lightweight process abstraction, users are often forced to write parts of concurrent applications in an evert-driven

Scala Actors used, e.g., in core message queue system of Twitter:

Philipp Haller publishes his influential paper on Scala's Actors; 'Actors that Unify Threads and Events'. http://infoscience



https://www.lightbend.com/akka-five-year-anniversary^{19/hotswap-cod}

ERLANG http://

Scala



Blog post:

http://jonasboner.com/2007/10/30/interview-with-joe-armstrong

Jonas starts tinkering with Scala Actors.



Current directions Ongoing work





- LaCasa: lightweight affine types and object capabilities in Scala [Haller & Loiko 2016]
 - Static reasoning about capabilities and resources
- Types for safe distribution
 - Closures [Miller et al. 2014], eventual consistency [Zhao & Haller 2019]
- Concurrent and distributed programming
 - Deterministic concurrency [Haller et al. 2016], function passing [Haller et al. 2018], asynchronous streams [Haller & Miller 2019]

Context Cloud computing

- Public cloud infrastructure integral part of numerous large-scale, commercial applications.
 - Amazon Web Services introduced > 12 years ago.
- Support for enterprise services: databases, queueing systems, object storage, etc.

So, cloud computing is now essentially a legacy enterprise service, right?

Cloud computing **Unused potential?**

So, cloud computing is now essentially a legacy enterprise service, right?

The cloud is...

"the biggest assemblage of data capacity and distributed computing power ever available to the general public, managed as a service." [1]

[1] Hellerstein et al. *Serverless Computing: One Step Forward, Two Steps Back.* CIDR 2019

Philipp Haller

NO!!!

Example event: "a commit was pushed to branch X of repository Y."

- Developers upload their code (functions) to the cloud.
 - Cloud platform executes these functions in response to events.
- No need for operating or provisioning servers.

Pay per use!



- Users only pay for compute resources used when their code is executed.
- Function execution is **autoscaling**: execution scales according to demand.

"Where is the catch?" Important restrictions

- Functions are **stateless**.
 - Must use external storage for any data/state that needs to survive multiple function executions.
- Function execution duration limited.
 - AWS Lambda: all function executions must complete within 300 seconds.

What is it good for? Which use cases are well-supported?

Depending on the patterns of function invocation [1]:

- Fully independent function invocations.
 - Scale up or down on demand:
 "invocations never wait for each other"
- Event-driven workflows connected via queueing systems or object stores.
 - High latency due to task handling and state management.

"Embarrassingly

parallel"

Key limitations

- Communication through slow storage: Functions not directly network-addressable, all communication via external services
- Functions are short-lived
 - Cannot service repeated requests via internal caches.
 - Cannot implement general distributed systems.
- I/O bottlenecks

invoking a **ication latency** no-op Lambda function

on a 1KB argument

Latency of "communicating" 1KB:

Lambda I/O Lambda I/O EC2 I/O EC2 NW Func. Invoc. EC2 I/O (DynamoDB) (1KB)(S3)(DynamoDB) (S3)(0MQ)303ms 108ms $11 \mathrm{ms}$ 290µs Latency 106ms $11 \mathrm{ms}$ Compared to best $1,045 \times$ 372× 37.9× 37.9× $365 \times$ $1 \times$

[1] Hellerstein et al. *Serverless Computing: One Step Forward, Two Steps Back.* CIDR 2019

write+read from "long-running" function

1KB network

message roundtrip

Back to the roots **Re-thinking distributed systems building**

- Re-think fundamental building blocks.

- Improve distributed systems stack.
- Devise and study programming models, languages, and systems.

Informed by SE and systems!

Challenge Programming model

- From data-shipping to **function-shipping**
 - Enable entirely different classes of applications: big data, ML model training.
- Principled fault-tolerance based on lineages.
 - Guarantee properties related to fault tolerance.
 - *Example:* program execution should never "get stuck" if at most N-1 out of 2N replicas fail.
 - Requires foundations for fault-tolerant programming.

Distributed programming with functional lineages a.k.a. function passing

 New data-centric programming model for functional processing of distributed data.

- Key ideas:

- Utilize lineages for fault injection *and* recovery
- Provide lineages by programming abstractions
- Keep data stationary (if possible), send functions

Introducing The function passing model

Consists of 3 parts:

- Silos: stationary, typed, immutable data containers
- SiloRefs: references to local or remote Silos.
- **Spores:** safe, serializable functions.

Some visual intuition of **The function passing model**





Two parts.

- SiloRef. Handle to a Silo.
- Silo. Typed, stationary data container.

User interacts with SiloRef.

SiloRefs come with 4 primitive operations. Philipp Haller



Primitive: apply

def apply[S](fun: T => SiloRef[S]): SiloRef[S]

- Takes a function that is to be applied to the data in the silo associated with the SiloRef.
- Creates new silo to contain the data that the userdefined function returns; evaluation is *deferred*



Enables interesting computation DAGs Philipp Haller



Primitive: send
def send(): Future[T]

- Forces the built-up computation DAG to be sent to the associated node and applied.
- Future is completed with the result of the computation.





Silo factories:

```
object SiloRef {
  def populate[T](host: Host, value: T): SiloRef[T]
  def fromTextFile(host: Host, file: File): SiloRef[List[String]]
  ...
}
```

Creates silo on given host populated with given value/text file/...



Basic idea: apply/send



Let's make an interesting DAG!

val persons: SiloRef[List[Person]] = ... Silo[List[Person]] persons: SiloRef[List[Person]]





```
val persons: SiloRef[List[Person]] = ...
val adults =
    persons.apply(spore { ps =>
        val res = ps.filter(p => p.age >= 18)
        SiloRef.populate(currentHost, res)
    })
val vehicles: SiloRef[List[Vehicle]] = ...
// adults that own a vehicle
val owners = adults.apply(...)
```



```
val persons: SiloRef[List[Person]] = ...
                                                                                                     Silo[List[Person]]
                                                              persons:
val adults =
                                                       SiloRef[List[Person]]
  persons.apply(spore { ps =>
   val res = ps.filter(p => p.age >= 18)
                                                                                               .....
    SiloRef.populate(currentHost, res)
 })
val vehicles: SiloRef[List[Vehicle]] = ...
// adults that own a vehicle
                                                        adults
                                                                                 vehicles
val owners = adults.apply(...)
val sorted =
  adults.apply(spore { ps =>
                                                   sorted
    SiloRef.populate(currentHost,
      ps.sortWith(p => p.age))
                                                                    owners
  })
val labels =
  sorted.apply(spore { ps =>
    SiloRef.populate(currentHost,
                                                   labels
      ps.map(p \Rightarrow "Hi, " + p.name))
  })
```

```
val persons: SiloRef[List[Person]] = ...
                                                                                                 Silo[List[Person]]
                                                            persons:
val adults =
                                                     SiloRef[List[Person]]
 persons.apply(spore { ps =>
   val res = ps.filter(p => p.age >= 18)
   SiloRef.populate(currentHost, res)
 })
val vehicles: SiloRef[List[Vehicle]] = ...
// adults that own a vehicle
                                                      adults
                                                                             vehicles
val owners = adults.apply(...)
val sorted =
  adults.apply(spore { ps =>
                                                 sorted
    SiloRef.populate(currentHost,
      ps.sortWith(p => p.age))
                                                                 owners
  })
val labels =
  sorted.apply(spore { ps =>
    SiloRef.populate(currentHost,
                                                 labels
     ps.map(p => "Hi, " + p.name))
  })
    so far we just staged
    computation, we haven't yet
    "kicked it off".
```



Putting lineages to work A functional design for fault-tolerance

- A SiloRef is a lineage, a persistent (in the sense of functional programming) data structure.
- The lineage is the DAG of operations used to derive the data of a silo.
- Since the lineage is composed of spores [2], it is serializable. This means it can be **persisted** or **transferred** to other machines.

[2] Miller, Haller, and Odersky. *Spores: a type-based foundation for closures in the age of concurrency and distribution.* ECOOP '14

Putting lineages to work A functional design for fault-tolerance

Next: we formalize **lineages**, a concept from the database + systems communities, in the context of PL. Natural fit in context of functional programming!

Formalization: typed, distributed core language with spores, silos, and futures.

Abstract syntax

```
t ::=
                                           terms:
                                           variable
       х
                                           value
       v
                                           application
       t t
       t \oplus t
                                           integer operator
       spore { \overline{x:T=t} ; (x:T) \Rightarrow t } spore
       populate(t,t)
                                          populate silo
       apply(t,t)
                                           apply
       send(t)
                                           send
       await(t)
                                           await future
       respond(h,t,t)
                                           respond
```

Local reduction and lineages

$\frac{R\text{-IntOp}}{E[v \oplus v'] \mid \sigma \to E[v''] \mid \sigma}$	R-AppAbs $E[((x:T) \Rightarrow t) v] \mid \sigma \rightarrow^{h} E[[x \mapsto v]t] \mid \sigma$				
$\begin{array}{l} \text{R-AppSpore} \\ E[(\texttt{spore} \{ \overline{x: T = v}; (x: T) \Rightarrow t \}) v] \mid \sigma \to^{h} E[\overline{[x \mapsto v]}[x \mapsto v]t] \mid \sigma \end{array}$					
$\frac{\text{R-Await}}{E[\text{await}(\iota)] \mid \sigma \rightarrow^{h} E[\nu] \mid \sigma}$	$\frac{\text{R-APPLY}}{r = \text{Ref}(l, h') l' = \text{Applied}((h, i), l, p) i \text{ fresh}}{E[\text{apply}(r, p)] \mid \sigma \rightarrow^{h} E[\text{Ref}(l', h')] \mid \sigma}$				

$$p ::= \text{spore} \{ \overline{x:T = v}; (x:T) \Rightarrow t \}$$

$$l ::= \qquad lineages: \\ Mat(\iota) \qquad materialized \\ | \text{Applied}(\iota, l, p) \qquad lineage with apply \\ r ::= \text{Ref}(l, h) \text{ where } h \in \mathscr{H} \text{ silo reference} \\ \iota ::= (h, i) \text{ where } h \in \mathscr{H} \text{ and } i \in \mathbb{N} \text{ decentralized identifier} \end{cases}$$

Distributed reduction

R-Local					
$t \mid \sigma \rightarrow^{h} t' \mid \sigma'$					
$\overline{\{(t,\sigma)^h\} \cup H \mid M \twoheadrightarrow \{(t',\sigma')^h\} \cup H \mid M}$					
R-Send					
$r = \operatorname{Ref}(l, h')$ $m = \operatorname{Req}(h, r, id(l))$ $M' = M \uplus \{h' \leftarrow m\}$					
$\overline{\{(E[send(r)], \sigma)^h\} \cup H \mid M \twoheadrightarrow \{(E[id(l)], \sigma)^h\} \cup H \mid M'$					
R-Populate					
$\iota = (h, i)$ <i>i</i> fresh $l = \operatorname{Mat}(\iota)$ $M' = M \uplus \{h' \leftarrow \operatorname{Res}(\iota, v)\}$					
$\overline{\{(E[\texttt{populate}(h',v)],\sigma)^h\} \cup H \mid M \twoheadrightarrow \{(E[\texttt{Ref}(l,h')],\sigma)^h\} \cup H \mid M'$					
R-Respond					
$M' = M \uplus \{h' \leftarrow \operatorname{Res}(\iota, v)\}$					
$\{(E[\texttt{respond}(h', \iota, \nu)], \sigma)^h\} \cup H \mid M \twoheadrightarrow \{(E[\texttt{unit}], \sigma)^h\} \cup H \mid M'$					
R-Process					
$\iota \notin dom(\sigma) M = M' \uplus \{h \leftarrow m\} process(h, m, \sigma) = (t, M'', \sigma')$					
$\overline{\{(E[await(\iota)], \sigma)^h\} \cup H \mid M \twoheadrightarrow \{(E[t; await(\iota)], \sigma')^h\} \cup H \mid M' \uplus M''}$					
R-Process-Val					
$M = M' \uplus \{h \leftarrow m\} process(h, m, \sigma) = (t, M'', \sigma')$					
$\{(v,\sigma)^h\} \cup H \mid M \twoheadrightarrow \{(t,\sigma')^h\} \cup H \mid M' \uplus M''$					

Type assignment

T-VAR	T-Int	T-IntOp				
$x:T\in\Gamma$	n integer literal	$\Gamma;\Sigma \vdash t$:Int	$\Gamma;\Sigma \vdash t':$ Int			
$\Gamma;\Sigma \vdash x:T$	$\Gamma;\Sigma \vdash n:$ Int	$\Gamma;\Sigma \vdash t$	$\oplus t'$:Int			
$\frac{\text{T-ABS}}{\Gamma; \Sigma \vdash t: T'}$ $\frac{\Gamma; \Sigma \vdash ((x:T) \Rightarrow t): T \Rightarrow T}{\Gamma; \Sigma \vdash ((x:T) \Rightarrow t): T \Rightarrow T}$			T-UNIT Γ;Σ⊢unit	:Unit		
T-Spore $\nabla \nabla = \overline{T} \cdot \overline{T}$	WE Smonel	WF-Store2		WF-STOP		WF-Lin1
$\Gamma; \Sigma \vdash \overline{t} : \overline{T}$	WF-Store1	$\emptyset; \Sigma \vdash v : T$	$\Sigma \vdash \sigma$	$\Sigma \vdash \sigma$	$\Sigma'\supseteq\Sigma$	$\iota \in dom(\Sigma)$
$\Gamma;\Sigma \vdash (spore\ \{$	∅⊢∅	$[\iota \mapsto T]\Sigma \vdash [$	$\iota \mapsto \nu]\sigma$	$\Sigma' \vdash$	σ	$\Sigma \vdash Mat(\iota)$
$\frac{\text{T-AppSpore}}{\Gamma; \Sigma \vdash t : T \Rightarrow T' \{ \text{type } G \\ \Gamma; \Sigma \vdash t : T \Rightarrow T' \}$		$\Sigma \vdash \text{Applie}$	$d(\iota, l, p)$			$\frac{ \begin{array}{c} \text{WF-R}_{\text{EF}} \\ \underline{\Sigma \vdash l h \in \mathscr{H}} \\ \hline \underline{\Sigma \vdash \text{Ref}(l,h)} \end{array} $
T-APPLY $\Gamma: \Sigma \vdash t: silepsf[T] = \Gamma: \Sigma \vdash t':$	WF-RES	WF-REO			WF-	HostConfig
$\frac{\mathbf{I}, \mathbf{\Sigma} + \mathbf{I} \cdot \text{SHORET}[\mathbf{I}] - \mathbf{I}, \mathbf{\Sigma} + \mathbf{I}}{\mathbf{\Sigma} \cdot \mathbf{\Sigma} + \mathbf{I}}$	$\Sigma(\iota) = T \emptyset; \Sigma \vdash v : T$	$r = \operatorname{Ref}(l,$	$h') \Sigma(id(l)) =$	$=\Sigma(\iota)$ Σ	$\vdash r \qquad \Sigma \vdash c$	$\sigma \exists \Gamma. \ \Gamma; \Sigma \vdash t : T$
$\frac{\begin{array}{c} \text{T-Apply} \\ \Gamma; \Sigma \vdash t : \texttt{siloRef}[T] \Gamma; \Sigma \vdash t' : \\ \hline \Gamma; \Sigma \vdash \texttt{apply}(t) \\ \text{T-Await} & \text{T-Interpretent} \end{array}$	$\Sigma \vdash \operatorname{Res}(\iota, v)$		$\Sigma \vdash \operatorname{Reg}(h, r, r)$	() ()		$\Sigma \vdash (t, \sigma)^h$
T-Await T-Id				.)		21 (1,0)
$\frac{\Gamma; \Sigma \vdash t: Future[T]}{\Gamma; \Sigma \vdash wwait(t): T} \qquad \overline{\Gamma; \Sigma}$ T-Pop	$\begin{array}{c} WF-Host1\\ \Sigma \vdash \emptyset \end{array} \qquad \begin{array}{c} WF-1\\ \Sigma \vdash \emptyset \end{array}$	$\frac{\text{Host}^2}{(t,\sigma)^h} \Sigma \vdash H}{\{(t,\sigma)^h\} \cup H}$		sages-Emf	$\sum_{\Sigma \vdash m} WF-M$	
$\Gamma;\Sigma$			WF-Config			
	$\Sigma \vdash H \Sigma \vdash M$					
	1		$\Sigma \vdash H \mid M$			
	Philipp Halle	٦r				35

Formalization **Properties of function passing model**

- Subject reduction theorem guarantees preservation of types under reduction, as well as preservation of lineage mobility.
- Progress theorem guarantees the finite materialization of remote, lineage-based data.

First correctness results for a programming model for lineage-based distributed computation.

Details, proofs, etc. **Paper**

Haller, Miller, and Müller. *A Programming Model and Foundation for Lineage-Based Distributed Computation.* Journal of Functional Programming 28 (2018) <u>https://infoscience.epfl.ch/record/230304</u>



Onward Ongoing and future work

Security & Privacy

Chaos Engineering

- Consistency, availability, partition tolerance
- Determinism

- Privacy-aware distribution
- Informationflow security
- Testing hypotheses about resilience in production systems

Conclusion

• Serverless computing

- Promising direction, intriguing properties
- Important limitations
- Foundations for function-shipping
 - Lineage-based distributed computation
 - First correctness results for a programming model based on lineages
- Goal: principles and foundations for a new distributed systems stack