



**HAL**  
open science

## No More, No Less - A Formal Model for Serverless Computing

Maurizio Gabbrielli, Saverio Giallorenzo, Ivan Lanese, Fabrizio Montesi,  
Marco Peressotti, Stefano Pio Zingaro

► **To cite this version:**

Maurizio Gabbrielli, Saverio Giallorenzo, Ivan Lanese, Fabrizio Montesi, Marco Peressotti, et al.. No More, No Less - A Formal Model for Serverless Computing. COORDINATION 2019 - 21th International Conference on Coordination Languages and Models, Jun 2019, Kongens Lyngby, Denmark. pp.148-157, 10.1007/978-3-030-22397-7\_9. hal-02365509

**HAL Id: hal-02365509**

<https://inria.hal.science/hal-02365509v1>

Submitted on 15 Nov 2019

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

# No more, no less

## A formal model for serverless computing

Maurizio Gabbrielli<sup>1</sup>, Saverio Giallorenzo<sup>2</sup>, Ivan Lanese<sup>1</sup>, Fabrizio Montesi<sup>2</sup>,  
Marco Peressotti<sup>2</sup>, and Stefano Pio Zingaro<sup>1</sup>

<sup>1</sup> INRIA, France / Università di Bologna, Italy

{maurizio.gabbrielli,ivan.lanese,stefanopio.zingaro}@unibo.it

<sup>2</sup> University of Southern Denmark, Denmark

{saverio, fmontesi, peressotti}@imada.sdu.dk

**Abstract** Serverless computing, also known as Functions-as-a-Service, is a recent paradigm aimed at simplifying the programming of cloud applications. The idea is that developers design applications in terms of functions, which are then deployed on a cloud infrastructure. The infrastructure takes care of executing the functions whenever requested by remote clients, dealing automatically with distribution and scaling with respect to inbound traffic.

While vendors already support a variety of programming languages for serverless computing (e.g. Go, Java, Javascript, Python), as far as we know there is no reference model yet to formally reason on this paradigm. In this paper, we propose the first core formal programming model for serverless computing, which combines ideas from both the  $\lambda$ -calculus (for functions) and the  $\pi$ -calculus (for communication). To illustrate our proposal, we model a real-world serverless system. Thanks to our model, we capture limitations of current vendors and formalise possible amendments.

## 1 Introduction

Serverless computing [24], also known as Functions-as-a-Service, narrows the development of Cloud applications to the definition and composition of stateless functions, while the provider handles the deployment, scaling, and balancing of the host infrastructure. Hence, although a bit of a misnomer — as servers are of course involved — the “less” in serverless refers to the removal of some server-related concerns, namely, their *maintenance*, *scaling*, and expenses related to a sub-optimal management (e.g. idle servers). Essentially, serverless pushes to the extreme the per-usage model of Cloud Computing: in serverless, users pay only for the computing resources used at each function invocation. This is why recent reports [18, 24] address serverless computing as the actual realisation of the long-standing promise of the Cloud to deliver *computation as a commodity*. AWS Lambda [4], launched in 2014, is the first and most widely-used serverless implementation, however many players like Google, Microsoft, Apache, IBM, and also open-source communities recently joined the serverless market [3, 16,

19, 21, 22, 29]. Current serverless proposals support the definition of functions — written in mainstream languages such as Go, Java, Javascript or Python — activated by specific events in the system, like a user request to a web gateway, the delivery of content from a message broker or a notification from a database. The serverless infrastructure transparently handles the instantiation of functions, as well as monitoring, logging, and fault tolerance.

Serverless offerings have become more and more common, yet the technology is still in its infancy and presents limitations [6, 18, 24] which hinder its wide adoption. For example, current serverless implementations favour operational flexibility (asynchrony and scalability) over developer control (function composition). Concretely, they do not support the direct composition of functions, which must call some stateful service in the infrastructure (e.g. a message broker) which will take care of triggering an event bound to the callee. On the one hand, that limitation is beneficial, since programmers must develop their functions as highly fine-grained, re-usable components (reminiscent of service-oriented architectures and microservices [12]). On the other hand, such openness and fine granularity increases the complexity of the system: programmers cannot assume sequential consistency or serialisability among their functions, which complicates reasoning on the semantics of the transformations applied to the global state of their architecture. This holds true also when estimating resource usage/costs, due to the complexity of unfolding all possible concurrent computations.

The above criticisms pushed us to investigate a core calculus for serverless computing, to reason on the paradigm, to model desirable features of future implementations, and to formalise guarantees over programs. In Section 2 we introduce the Serverless Kernel Calculus (SKC); as far as we know, the first core formal model for serverless computing. SKC combines ideas from both the  $\lambda$ -calculus (for functions) and the  $\pi$ -calculus (for communication). In Section 2, we also extend SKC to capture limitations of current serverless implementations. In Section 3 we use our extension to model a real-world serverless architecture [1], implemented on AWS Lambda. Finally, in Section 4 we discuss future developments of SKC.

## 2 A Serverless Kernel Calculus

Our kernel calculus defines a serverless architecture as a pair  $\langle S, \mathcal{D} \rangle$ , where  $S$  is the system of *running functions* and  $\mathcal{D}$  is a *definition repository*, containing function definitions. The repository  $\mathcal{D}$  is a partial function from function names  $f$  to function bodies  $M$ .  $M$  includes function application ( $M \ M'$ ), asynchronous execution of new functions (**async**  $M$ ), function names  $f$ , and values  $V$ . Values include variables  $x$ ,  $\lambda$ -abstractions  $\lambda x. M$ , named *futures* [5, 17, 32]  $c$ , and the unit value  $()$ . A system  $S$  contains *running functions*  $c \blacktriangleleft M$ , where  $c$  will contain the result of the computation of the function  $M$ . Systems can be composed in parallel  $|$  and include the empty system  $\emptyset$ . Futures can be restricted in systems via  $\nu c S$ .

$$S, S' ::= c \blacktriangleleft M \mid S \mid S' \mid \nu c S \mid \emptyset \quad (\text{Systems})$$

$$M, M' ::= M \ M' \mid \mathbf{async} \ M \mid f \mid V \quad (\text{Functions})$$

$$\begin{array}{c}
 \frac{\langle \mathcal{E}[(\lambda x. M) V], \mathcal{D} \rangle \longrightarrow \langle \mathcal{E}[M\{V/x\}], \mathcal{D} \rangle \quad [\beta]}{c \notin \text{fn}(M)} \quad \frac{\mathcal{D}(f) = M}{\langle \mathcal{E}[f], \mathcal{D} \rangle \longrightarrow \langle \mathcal{E}[M], \mathcal{D} \rangle} \quad [\text{RET}] \\
 \frac{\langle \mathcal{E}[\mathbf{async} M], \mathcal{D} \rangle \longrightarrow \langle \nu c(\mathcal{E}[c] \mid c \blacktriangleleft M), \mathcal{D} \rangle \quad [\text{ASYNC}]}{S_0 \equiv S'_0} \quad \frac{\langle \nu c(S \mid c \blacktriangleleft V), \mathcal{D} \rangle \longrightarrow \langle S\{V/c\}, \mathcal{D} \rangle \quad [\text{PUSH}]}{S'_1 \equiv S_1} \\
 \frac{\langle S_0, \mathcal{D} \rangle \longrightarrow \langle S_1, \mathcal{D}' \rangle \quad \langle S'_1, \mathcal{D}' \rangle \longrightarrow \langle S'_1, \mathcal{D}' \rangle}{\langle S_0, \mathcal{D} \rangle \longrightarrow \langle S_1, \mathcal{D}' \rangle} \quad [\text{STR}] \\
 \frac{\langle S, \mathcal{D} \rangle \longrightarrow \langle S', \mathcal{D}' \rangle}{\langle \nu c S, \mathcal{D} \rangle \longrightarrow \langle \nu c S', \mathcal{D}' \rangle} \quad [\text{RES}] \quad \frac{\langle S_1, \mathcal{D} \rangle \longrightarrow \langle S'_1, \mathcal{D}' \rangle}{\langle S_1 \mid S_2, \mathcal{D} \rangle \longrightarrow \langle S'_1 \mid S_2, \mathcal{D}' \rangle} \quad [\text{LPAR}]
 \end{array}$$

**Figure 1.** SKC reduction semantics.

$$V, V' ::= x \mid \lambda x. M \mid c \mid () \quad (\text{Values})$$

We assume futures to appear only at runtime and not in initial systems. Moreover, we consider a standard structural congruence  $\equiv$  that supports changing the scope of restrictions to avoid name capture, and where parallel composition is associative, commutative, and has  $\emptyset$  as neutral element.

$$\begin{array}{l}
 \nu c \nu c' S \equiv \nu c' \nu c S \quad \nu c(S \mid S') \equiv \nu c S \mid S' \quad \text{if } c \notin \text{fn}(S') \\
 S \equiv S \mid \emptyset \quad S \mid S' \equiv S' \mid S \quad (S \mid S') \mid S'' \equiv S \mid (S' \mid S'')
 \end{array}$$

We define the semantics of our calculus using evaluation contexts  $\mathcal{E}$  and  $\mathcal{E}_\lambda$ , to evaluate, respectively, systems and functions.

$$\mathcal{E} ::= c \blacktriangleleft \mathcal{E}_\lambda \quad \mathcal{E}_\lambda ::= [-] \mid (\lambda x. M)\mathcal{E}_\lambda \mid \mathcal{E}_\lambda M$$

We report in Figure 1 the semantics of serverless architectures  $\langle S, \mathcal{D} \rangle$ , expressed as reduction rules. Rule  $[\beta]$  is the traditional function application of  $\lambda$ -calculus. Rule  $[\text{RET}]$  retrieves the body of function  $f$  from the definition repository  $\mathcal{D}$ . Rule  $[\text{ASYNC}]$  models the execution of new functions: it creates a fresh future  $c$  and, in parallel, it executes function  $M$  so that  $c$  will store the evaluation of  $M$ . When the evaluation of a function reduces to a value, rule  $[\text{PUSH}]$  returns the value to the associated future and removes both the terminated function and its restriction. Rules  $[\text{STR}]$ ,  $[\text{RES}]$ , and  $[\text{LPAR}]$  perform the closure under, respectively, structural congruence, restriction, and parallel composition. We include in SKC standard components (conditionals, etc.) and extend evaluation contexts ( $\mathcal{E}$ ) accordingly:

$$\begin{array}{l}
 M ::= \dots \mid \mathbf{if} M \mathbf{then} M' \mathbf{else} M'' \mid \mathbf{fst} M \mid \mathbf{snd} M \\
 V ::= \dots \mid \text{True} \mid \text{False} \mid (V, V')
 \end{array}$$

We define standard macros for **fixpoint**, **let** and **let rec** declarations, and pairs.

$$\begin{array}{l}
 \mathbf{fix} \triangleq \lambda f. (\lambda x. f(xx))(\lambda x. f(xx)) \quad \mathbf{let} x = M \mathbf{in} M' \triangleq (\lambda x. M') M \\
 \mathbf{let rec} x = M \mathbf{in} M' \triangleq \mathbf{let} x = \mathbf{fix} \lambda x. M \mathbf{in} M' \\
 \lambda(x, y). M \triangleq \lambda z. (\lambda x. \lambda y. M) (\mathbf{fst} z) (\mathbf{snd} z)
 \end{array}$$

## 2.1 SKC<sub>σ</sub> - A stateful extension of SKC

SKC considers static definition repositories, i.e. no rules mutate the state of  $\mathcal{D}$ . We now present SKC<sub>σ</sub>, an extension of SKC which includes two primitives to define transformations on definition repositories. As shown in Section 3, SKC<sub>σ</sub> is powerful enough to encode stateful services, like databases and message queues.

$$M, M' ::= \dots \mid \mathbf{set} \ f \ M \mid \mathbf{take} \ f$$

The first primitive included in SKC<sub>σ</sub> is  $\mathbf{set} \ f \ M$ , which updates the definition repository  $\mathcal{D}$  to map  $f$  to  $M$ : users can use the  $\mathbf{set}$  primitive to deploy new function definitions or update/override existing ones. The second primitive is  $\mathbf{take} \ f$ , which removes the definition of  $f$  from  $\mathcal{D}$ , returning it to the caller. We report below the semantics of the new primitives.

$$\frac{\text{futures}(M) = \emptyset}{\langle \mathcal{E}[\mathbf{set} \ f \ M], \mathcal{D} \rangle \longrightarrow \langle \mathcal{E}[f], \mathcal{D}[f \mapsto M] \rangle} \text{[SET]}$$

$$\frac{\mathcal{D}(f) = M}{\langle \mathcal{E}[\mathbf{take} \ f], \mathcal{D} \rangle \longrightarrow \langle \mathcal{E}[\mathbf{let} \ \mathbf{rec} \ f=M \ \mathbf{in} \ M], \mathbf{undef}(\mathcal{D}, f) \rangle} \text{[TAKE]}$$

The only restriction on the application of rule [SET] is that the body  $M$  of the newly deployed function  $f$  does not contain futures ( $\text{futures}(M)$  is the set of futures occurring in  $M$ ). This preserves the semantics of restriction of futures in function evaluations (cf. rules [ASYNC] and [PUSH]). In the reductum, the rule returns the name of the deployed function, useful to invoke it in the continuation. Rule [TAKE] removes the definition  $M$  of a deployed function  $f$ . For simplicity, we define [TAKE] applicable only if  $f$  is defined. In the reductum, the caller of the  $\mathbf{take}$  obtains the **recursive let** declaration of the function (useful for internal application) while the association for  $f$  is removed from  $\mathcal{D}$  by function  $\mathbf{undef}$ .

## 2.2 SKC<sub>e</sub> - Event-based function composition in SKC

We present an idiom of SKC, called SKC<sub>e</sub>, which models event-based function composition. SKC<sub>e</sub> captures one of the main limitations of current serverless vendors: the lack of support for direct function invocation, replaced by an event-handling/event-triggering invocation model. Indeed, current serverless implementations, such as AWS Lambda, work as follows: they include infrastructural stateful services, such as API gateways, that we can model using our stateful extension SKC<sub>σ</sub>, and these services throw events. User-defined functions are invoked as handlers of these events. User-defined functions can then invoke the infrastructural services above. Notably, a user-defined function cannot directly invoke another user-defined function. We will see an instance of the event-based pattern in Section 3, while we describe below event handling mechanisms.

We model events ( $e$  and variations thereof) inside SKC as function names associated with peculiar function bodies in the repository  $\mathcal{D}$  that asynchronously evaluate the corresponding event handler and discard the handler result. For convenience,  $i$ ) we package the asynchronous call of an event handler in the helper

function `callHandler` below (hereafter, we assume that  $\mathcal{D}$  contains `callHandler`) and *ii*) we write `_` for unused variable symbols in binding constructs.

$$\text{callHandler} \mapsto \lambda h. \lambda x. \mathbf{let} \_ = \mathbf{async} (h () x) \mathbf{in} ()$$

Event  $e$  is defined in  $\mathcal{D}$  as  $e \mapsto \text{callHandler } \lambda \_ . h_e$  and its event handler as  $h_e \mapsto M_e$ ; we wrap the name  $h_e$  in a lambda abstraction to avoid expansion (via  $\text{[RET]}$ ) since function names are not values. Raising an event  $e$  with some parameter  $v$  results in asynchronously executing the corresponding handler, as shown by the derivation below (we abbreviate  $\langle S, \mathcal{D} \rangle \rightarrow \langle S', \mathcal{D} \rangle$  as  $S \rightarrow_{\mathcal{D}} S'$  and label reductions with the names of the most relevant applied rules).

$$\begin{aligned} e \ v \xrightarrow{\text{[RET]}}_{\mathcal{D}} \text{callHandler } \lambda \_ . h_e \ v \xrightarrow{\text{[RET]}}_{\mathcal{D}} (\lambda h. \lambda x. \mathbf{let} \_ = \mathbf{async} (h () x) \mathbf{in} ()) \ \lambda \_ . h_e \ v \\ \xrightarrow{[\beta], [\beta]}_{\mathcal{D}} \mathbf{let} \_ = \mathbf{async} (\lambda \_ . h_e () v) \mathbf{in} () \xrightarrow{\text{[ASYNC]}}_{\mathcal{D}} \nu c(\mathbf{let} \_ = c \mathbf{in} () \mid c \blacktriangleleft \lambda \_ . h_e () v) \\ \xrightarrow{[\beta], [\beta]}_{\mathcal{D}} \nu c(c \mid c \blacktriangleleft h_e v) \xrightarrow{\text{[RET]}}_{\mathcal{D}} () \mid \nu c(c \blacktriangleleft M_e v) \end{aligned}$$

### 3 An Illustrative Example

Let  $\text{SKC}_{\sigma_e}$  be the compound of  $\text{SKC}_{\sigma}$  and  $\text{SKC}_e$  presented in Section 2. Here, we illustrate how  $\text{SKC}_{\sigma_e}$  can capture real-world serverless systems by encoding a relevant portion (depicted in Figure 2) of Tailor [1], an architecture for user registration, developed by Autodesk over AWS Lambda. Tailor mixes serverless functions with vendor-specific services: *API Gateways*, key-value databases (*DynamoDB*), and queue-based notification services (*SNS*). In the architecture, each function defines a fragment of the logic of a user-registration procedure, like the initiation of registration requests (`talr-receptionist`), request validation (`talr-validator`), etc. To model Figure 2 in  $\text{SKC}_{\sigma_e}$ , first, we install in  $\mathcal{D}$  the event handlers for the *API Gateway*, the *DynamoDB*, and *SNS* services<sup>3</sup>:

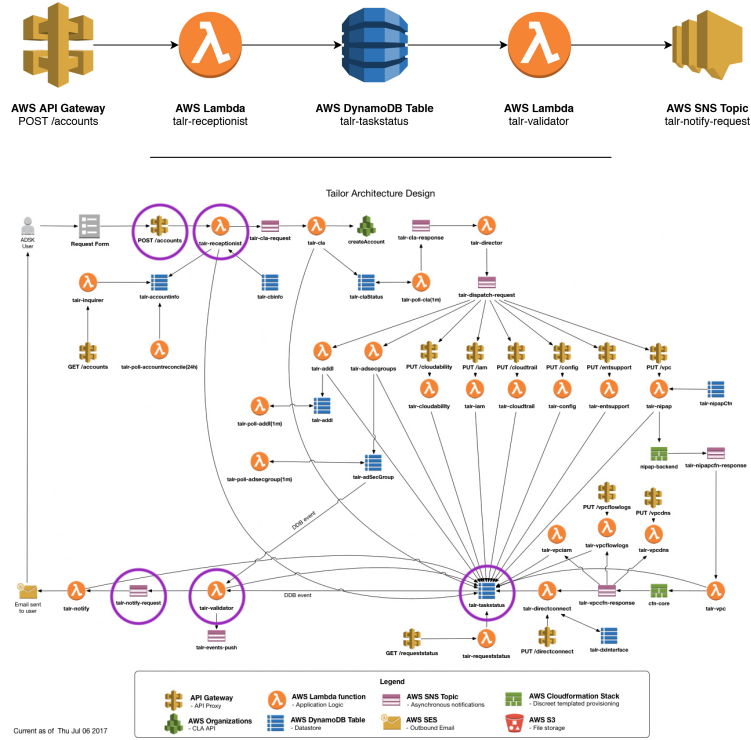
$$\begin{aligned} e_{\text{API}} \mapsto \text{callHandler}(\text{talr-receptionist}) \quad e_{\text{SNS}} \mapsto \text{callHandler}([\dots]) \\ e_{\text{DB}} \mapsto \text{callHandler}(\text{talr-validator}) \end{aligned}$$

Then, we define the functions called by the handlers installed above, using the same names of the *AWS Lambda* functions in Figure 2. Handler  $e_{\text{API}}$  calls function `talr-receptionist`, which validates the request and inserts the information of the user in the key/value database. For brevity, we omit the behaviour of `talr-receptionist` in case of invalid requests and the definition of auxiliary functions `validate_request`, `get_key`, `get_value` in  $\mathcal{D}$ :

$$\begin{aligned} \text{talr-receptionist} \mapsto \lambda x. \mathbf{if} \ \text{validate\_request} \ x \ \mathbf{then} \\ \quad \text{write\_db} (\text{get\_key} \ x, \text{get\_value} \ x) \ \mathbf{else} \ [\dots] \end{aligned}$$

Handler  $e_{\text{DB}}$  invokes function `talr-validator`, which retrieves from the database the `status` of task  $x$ , checks if it is complete, and sends a notification on *SNS*. We omit the definitions of functions `check` and `compose_msg` and of the `else` branch.

<sup>3</sup> We omit the name of the function called by  $e_{\text{SNS}}$ , excluded in the excerpt of Figure 2.



**Figure 2.** Scheme of the Autodesk Tailor system. Top, excerpt considered in the example. Bottom, full architecture (circled elements belong to the excerpt).

```
tailr-validator  $\mapsto$   $\lambda x$ .let status = read_db x in
    if check status then push (compose_msg x) else [...]
```

We conclude illustrating the definitions of functions `write_db`, `read_db`, and `push` in  $\mathcal{D}$ , which exemplify how  $SKC_{\sigma_e}$  can encode stateful, event-triggering services. Keys are represented as function names and values are stored in  $\mathcal{D}$ ; thus keys are passed around wrapped in lambda abstractions ( $\lambda_.k$ ) as done for events.

```
write_db  $\mapsto$   $\lambda(x,v)$ .eDDB (set (x ()) v)    read_db  $\mapsto$   $\lambda x$ .x ()
push  $\mapsto$   $\lambda(x,v)$ .eSNS (set (x ()) v)
```

Function `write_db` takes a key (wrapped as  $x = \lambda_.k$ ) and a value  $v$  as parameters, writes on the database by **setting** to  $v$  the body of a function called  $k$ , and notifies the write, invoking  $e_{DDB}$ <sup>4</sup>. Function `read_db` simply unwraps the key thus enabling

<sup>4</sup> More involved variants of the database are possible. E.g. to avoid clashes among services using the same key for different elements, we can either use scoping or prefix key names with service names — e.g. Tailor uses service-specific tables in *DynamoDB*.

retrieval from  $\mathcal{D}$ . Similarly to `write_db`, function `push` publishes (`set`) a message  $v$  on an *SNS* topic (represented as a function name) and triggers  $e_{\text{SNS}}$ .

*Remark 3.1.* The example illustrates how SKC can capture (but not be restricted by) one of the most prominent limitations of current serverless platform [18], i.e. that *i*) user-defined functions can be only invoked by raising an event that executes a new function (as done by `callHandler`, using the `async` primitive) and *ii*) functions can invoke other functions only by interacting with some event-triggering infrastructural service (e.g. a database, represented by function `write_db`, or a notification queue, represented by function `push`).

## 4 Discussion and Conclusion

We propose SKC, the first core formal model to reason on serverless computing. While the design of SKC strives for minimality, it captures the main ingredients [18, 24] of serverless architectures: *i*) the deployment and instantiation of event-triggered, stateless functions and *ii*) the desiderata of direct function-to-function invocation based on futures — in Section 3 we show how this mechanism is powerful enough to cover also the current setting of serverless vendors, where function invocation must rely on third-party services that handle event triggering.

Futures [5, 17], which are the main communication mechanism in SKC, are becoming one of the de-facto standards in asynchronous systems [13, 15, 35, 37]. We considered using named channels (as in CCS/ $\pi$ -calculus [30, 34]) instead of futures, but we found them too general for the needs of the serverless model (they are bi-directional and re-usable). Besides, futures can encode channels [32].

The work closest to ours is [23], appeared during the submission of this work, in the form of a technical report. It presents a detailed operational semantics that captures the low-level details of current serverless implementations (e.g. cold/warm components, storage, and transactions are primitive features of their model) whereas SKC identifies a kernel model of serverless computing. Another work close to SKC is [32], where the authors introduce a  $\lambda$ -calculus with futures. Since the aim of [32] is to formalise and reason on a concurrent extension of Standard ML, their calculus is more involved than SKC, as it contains primitive operators (handlers and cells) to encode safe non-deterministic concurrent operations, which can be encoded as macros in SKC. An interesting research direction is to investigate which results from [23, 32] can be adapted to SKC.

Being the first core framework to reason on serverless architectures, SKC opens multiple avenues of future research. For example, current serverless technologies offer little guarantee on sequential execution across functions, which compels the investigation of new tools to enforce sequential consistency [28] or serialisability [33] of the transformations of the global state [18]. That challenge can be tackled developing static analysis techniques and type disciplines [2, 20] for SKC. Another direction concerns programming models, which should give to programmers an overview over the overall logic of the distributed functions and capture the loosely-consistent execution model of serverless [18]. Choreographic



Programming [10, 31] is a promising candidate for that task, as choreographies are designed to capture the global interactions in distributed systems [26], and recent results [9, 11, 14] confirmed their applicability to microservices [12], a neighbouring domain to that of serverless architectures. Other possible research directions, that we do not discuss for space constraints, include monitoring, various kinds of security analysis including “self-DDoS attacks” [25, 27, 36] and performance analysis. This last one is particularly relevant in the per-usage model of serverless architectures, yet requires to extend SKC with an explicit notion of time in order to support quantitative behavioural reasoning for timed systems [7, 8].

## Bibliography

- [1] Alan Williams: Tailor - the AWS Account Provisioning Service. <https://github.com/alanwill/aws-tailor>, online; acc. 02/2019
- [2] Ancona, D., et al.: Behavioral types in programming languages. *Foundations and Trends in Programming Languages* **3**(2-3), 95–230 (2016)
- [3] Apache: OpenWhisk. <https://github.com/apache/incubator-openwhisk>, online; acc. 02/2019
- [4] AWS: Lambda. <https://aws.amazon.com/lambda/>, online; acc. 02/2019
- [5] Baker Jr, H.C., Hewitt, C.: The incremental garbage collection of processes. In: *ACM Sigplan Notices*. vol. 12(8), pp. 55–59. ACM (1977)
- [6] Baldini, I., et al.: Serverless computing: Current trends and open problems. In: *Research Advances in Cloud Computing*, pp. 1–20. Springer (2017)
- [7] Brengos, T., Peressotti, M.: A uniform framework for timed automata. In: *CONCUR. LIPIcs*, vol. 59, pp. 26:1–26:15. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2016)
- [8] Brengos, T., Peressotti, M.: Behavioural equivalences for timed systems. *Logical Methods in Computer Science* **15**(1) (2019)
- [9] Carbone, M., Montesi, F.: Deadlock-freedom-by-design: multiparty asynchronous global programming. In: *POPL*. pp. 263–274. ACM (2013)
- [10] Cruz-Filipe, L., Montesi, F.: A core model for choreographic programming. In: *FACS*. pp. 17–35. *Lecture Notes in Computer Science* (2016)
- [11] Dalla Preda, M., et al.: Dynamic choreographies: Theory and implementation. *Logical Methods in Computer Science* **13**(2) (2017)
- [12] Dragoni, N., et al.: Microservices: Yesterday, today, and tomorrow. In: *Present and Ulterior Software Engineering.*, pp. 195–216. Springer (2017)
- [13] EcmaScript 2018 language specification. <http://ecma-international.org/ecma-262/9.0/index.html>, online; acc. 02/2019
- [14] Giallorenzo, S., Montesi, F., Gabbrielli, M.: Applied choreographies. In: *FORTE*. pp. 21–40. Springer (2018)
- [15] Goetz, B., et al.: *Java concurrency in practice*. Pearson Education (2006)
- [16] Google: Cloud Functions. <https://cloud.google.com/functions>, online; acc. 02/2019
- [17] Halstead Jr, R.H.: Multilisp: A language for concurrent symbolic computation. *ACM Transactions on Programming Languages and Systems (TOPLAS)* **7**(4), 501–538 (1985)

- [18] Hellerstein, J.M., et al.: Serverless computing: One step forward, two steps back. In: CIDR. [www.cidrdb.org](http://www.cidrdb.org) (2019)
- [19] Hendrickson, S., et al.: Serverless Computation with OpenLambda. In: USENIX. USENIX Association (2016)
- [20] Hüttel, H., et al.: Foundations of session types and behavioural contracts. *ACM Comput. Surv.* **49**(1), 3:1–3:36 (2016)
- [21] IBM: Cloud Functions. <https://www.ibm.com/cloud/functions>, online; acc. 02/2019
- [22] Iron.io: IronFunctions. <https://open.iron.io>, online; acc. 02/2019
- [23] Jangda, A., et al.: Formal foundations of serverless computing. CoRR **abs/1902.05870** (2019), <http://arxiv.org/abs/1902.05870>
- [24] Jonas, E., et al.: Cloud programming simplified: A berkeley view on serverless computing. Tech. rep., EECS Department, University of California, Berkeley (Feb 2019)
- [25] K-Optional Software: Serverless out of Control. <https://koptional.com/2019/01/22/serverless-out-of-control/>, online; acc. 02/2019
- [26] Kavantzias, N., Burdett, D., Ritzinger, G., Lafon, Y.: Web services choreography description language version 1.0, W3C candidate recommendation. Tech. rep., W3C (2005), <http://www.w3.org/TR/ws-cdl-10>
- [27] Kevin Vandendorpe: Serverless: A lesson learned. The hard way. <https://sourcebox.be/blog/2017/08/07/serverless-a-lesson-learned-the-hard-way/>, online; acc. 02/2019
- [28] Lamport, L.: How to make a multiprocessor computer that correctly executes multiprocess programs. *IEEE Trans. Comput.* **28**(9), 690–691 (Sep 1979)
- [29] Microsoft: Azure Functions. <https://azure.microsoft.com/services/functions>, online; acc. 02/2019
- [30] Milner, R.: A Calculus of Communicating Systems, Lecture Notes in Computer Science, vol. 92. Springer (1980)
- [31] Montesi, F.: Kickstarting choreographic programming. In: WS-FM/BEAT. pp. 3–10. Springer (2015)
- [32] Niehren, J., Schwinghammer, J., Smolka, G.: A concurrent lambda calculus with futures. *Theor. Comput. Sci.* **364**(3), 338–356 (2006)
- [33] Papadimitriou, C.H.: The serializability of concurrent database updates. *J. ACM* **26**(4), 631–653 (1979)
- [34] Sangiorgi, D., Walker, D.: The Pi-Calculus - a theory of mobile processes. Cambridge University Press (2001)
- [35] Summerfield, M.: Python in practice: create better programs using concurrency, libraries, and patterns. Addison-Wesley (2013)
- [36] Tom Wright: Beware “RunOnStartup” in Azure Functions — a serverless horror story. <http://blog.tdwright.co.uk/2018/09/06/beware-runonstartup-in-azure-functions-a-serverless-horror-story/>, online; acc. 02/2019
- [37] Williams, A.: C++ concurrency in action. Manning (2017)