

Evaluating IoT Service Composition Mechanisms for the Scalability of IoT Systems

Damian Arellanes*, Kung-Kiu Lau

Department of Computer Science, The University of Manchester, Manchester M13 9PL, United Kingdom

ARTICLE INFO

Keywords:
Scalability
IoT service composition
Orchestration
Choreography
Dataflows
Workflows

ABSTRACT

The Internet of Things (IoT) is an emerging paradigm where practically every (physical and virtual) thing will be interconnected through innovative distributed services. Since the number of connected things is rapidly growing, IoT systems will require the composition of plenty of services into complex workflows. Thus, scalability in terms of the size of IoT systems becomes a significant concern. In this paper, we review and evaluate the fundamental semantics of existing IoT service composition mechanisms to determine how well they fulfil the scalability requirements of IoT systems. We identify scalability desiderata and, accordingly, our findings show that dataflows, orchestration and choreography do not fully satisfy such desiderata, unlike a novel composition mechanism called DX-MAN.

1. Introduction

The Internet of Things (IoT) is an emerging paradigm that promises the interconnection of (physical and virtual) things through innovative distributed services. Like traditional enterprise services, IoT services interact in many different ways via the Internet, in order to realise a global system workflow. However, unlike traditional enterprise systems, IoT systems will require the interaction of billions of services as the number of connected things (and therefore services) is rapidly growing [1, 2]. Thus, scalability becomes a crucial concern.

Scalability is typically considered as a system capability to handle increasing workloads [3, 4, 5, 6, 7, 8]. In particular, vertical scalability [9, 10, 11] refers to the addition or removal of computing resources in a single IoT node, while horizontal scalability [2, 12, 13] involves the addition or removal of IoT nodes. These kinds of scalability have been addressed by an extensive body of research [4, 10, 11, 14, 15, 16, 17, 18, 19, 20], unlike scalability in terms of the number of services composed in an IoT system, which we refer to as *functional scalability*.

Existing service composition mechanisms were primarily designed for the integration of static enterprise services, not for the physical world. For that reason, they may not address the functional scalability challenges that IoT systems pose. Early IoT systems were deployed in closed environments, using private Application Programming Interfaces (APIs) and private data. However, future IoT systems will be deployed in open environments (also known as software ecosystems [21]) with an overwhelming number of available services, as a result of the huge number of connected things [22]. For that reason, billions of IoT services will be composed into complex IoT systems [23, 24, 25, 26, 27, 28]. This raises the challenging question of *How to construct IoT systems composed of an ultra-large number of services?*

In this paper, we study the ability of service composition

mechanisms to handle any number of IoT services. Our aim is to determine *which service composition mechanism best fulfils the functional scalability requirements of IoT systems*. To answer this, we have reviewed the fundamental semantics (not specific implementation technologies) of existing composition mechanisms, and proposed an evaluation framework that considers six crucial desiderata: (i) explicit control flow [29, 30]; (ii) distributed workflows [23, 31, 32]; (iii) location transparency [20, 33]; (iv) decentralised data flows [23, 32, 34]; (v) separation of control, data and computation [35, 36, 37, 38]; and (vi) workflow variability [23, 39, 40, 41]. Our evaluation framework serves as a guideline for defining the semantics of future IoT service composition mechanisms.

The remainder of the paper is structured as follows. In the next section, we present an overview of IoT services, scalability and service composition. Section 3 outlines the related work on IoT service composition mechanisms. Section 4 introduces a motivating IoT scenario based on a smart parking system. The scenario is then considered in Section 5 to comprehensively describe the requirements for functional scalability and the rationale behind them. Section 6 analyses service composition mechanisms on the basis of the functional scalability requirements. Section 7 presents an evaluation that determines how well service composition mechanisms fulfil the scalability requirements. Our findings are then discussed in Section 8. Finally, Section 9 draws the conclusions and presents the directions for future research.

2. Background

This section presents a background for the rest of the paper, which includes an overview of IoT services, service composition and scalability.

2.1. IoT Services

Kevin Ashton coined the term *Internet of Things* (IoT) in a presentation made in 1999 at Procter and Gamble [42], referring to the interconnection of everything via the Internet for the creation of an ubiquitous computing environment [43]. As per the recommendation of ITU-T Y.4000, IoT has been

*Corresponding author.

✉ damian.arellanesmolina@manchester.ac.uk (D. Arellanes);
kung-kiu.lau@manchester.ac.uk (K. Lau)
ORCID(s): 0000-0002-0074-390X (D. Arellanes)

recently redefined as “a global infrastructure for the information society, enabling advanced services by interconnecting (physical and virtual) things based on existing and evolving interoperable information and communication technologies” [44].

A *thing* is practically a physical or virtual construct of the real-world, which is capable of being identified and integrated into communication networks through specific protocols. The difference between physical and virtual things lies in their tangibility [45]. A physical thing is a tangible object of the physical world, which is capable of being sensed, actuated and connected, e.g., home appliances, robots, buildings, plants and people. Contrastingly, a virtual thing is a non-tangible construct formed from a human idea which only exists in the information world, e.g., Clouds and enterprises.

The broad range of available things inevitably requires dealing with a high degree of heterogeneity in a distributed environment. Accordingly, a Service-Oriented Architecture (SOA) represents the best way of dealing with this issue [46]. It is a *logical way of designing a software system to provide services either to end-user applications or other services distributed in a network, via published and discoverable interfaces* [47]. Thus, it is expected that physical and virtual things will provide services to expose behaviour via interfaces [5, 48, 49, 50, 51, 52].

According to the Oxford dictionary, the word *service* can be a noun or a verb referring to the *action of helping or doing work for someone*. Considering a service as a noun allows the encapsulation of behaviour (i.e., actions) in the form of operations described as verbs. Thus, a *software service is a distributed software component that provides a set of operations through network-accessible endpoints* [53, 54]. In general, IoT services are virtual representations of the behaviour of things, which can be combined with other services into more complex behaviours to yield complex service-oriented IoT systems [52, 55, 56, 57]. Thus, things are integrated through the composition of the services they provide (see Section 2.2).

Resource-constrained things (e.g., pulse sensors) typically provide fine-grained services for basic functionality (e.g., fetching sensor data), whilst non-resource constrained things (e.g., a Cloud) may offer coarse-grained services in addition (e.g., services for geolocation or complex industrial processes). Enterprise services are typically coarse-grained as they are deployed on resource-rich infrastructures, whilst services of physical things are often fine-grained because they are usually deployed on resource-constrained things.

Figure 1 shows the relationship between things, services and operations. Figure 1(a) depicts a *washing machine* (i.e., a resource-constrained physical thing) that offers the *Washing* and *Drying* fine-grained services with two operations each (for starting and stopping the respective processes). Figure 1(b) shows a *City Council Cloud* (i.e., a non-resource constrained virtual thing) that offers the services *CouncilTax* and *Parking*. The *CouncilTax* service provides the operations *pay* (to pay a tax bill) and *check* (to query council tax information). The *Parking* service offers the operations *getNearest*

(for getting the closest parking space from a driver’s location) and *registerDisabled* (for registering an impaired driver).




Figure 1: Relationship between things, IoT services and operations.

For the rest of the paper, we use the notation $S.O$ to denote an operation O in service S , e.g., *Parking.getNearest* refers to the *getNearest* operation provided by the *Parking* service.

2.2. IoT Service Composition Revisited

An IoT service is a distributed unit of composition, which constitutes the virtual representation of a thing’s behaviour, and can be either atomic or composite. An *atomic service* is a well-defined and self-contained piece of behaviour that cannot be divided into other services [58, 59, 60]. A *composite service*, on the other hand, is a more complex unit that provides value-added functionality and is formed by the combination of (atomic or composite) services [52, 55, 57, 58, 59, 61, 62]. For example, a humidity sensing service can be combined with a temperature service into an air conditioning composite [63].

The ability of combining services is referred to as *compositionality* and is realised by a *composition mechanism* [59]. Thus, an IoT system requires a things infrastructure, the definition of what a service is and the selection of a composition mechanism [64]. In any scenario, composition is done regardless of both the technologies being used and the things infrastructure. Service technologies include REST [65, 66], WS-* [67, 68], OSGi [69, 70] and many others.¹

A service composition mechanism defines a meaningful interaction between services [59] by considering two functional dimensions: control flow and data flow [71, 72]. Control flow refers to the order in which interactions occur [73, 74], whilst data flow defines how data is moved among services [71]. In this paper, we focus on service composition mechanisms that define behaviour by workflows, namely (centralised and distributed) dataflows, (centralised and distributed) orchestration, choreography, and a novel composition mechanism called DX-MAN. Section 6 provides a detailed description of these mechanisms.

A workflow describes a series of discrete steps for the realisation of a computational activity, which can be control-driven, data-driven or hybrid [75]. In a control-driven workflow, steps (also known as tasks [76], actors [77], transitions [78], procedures [79], thorns [80], activities [81] and units [82]) are executed according to explicit control flow constructs that define sequencing, looping, branching or parallelising. A data-driven workflow invokes steps whenever

¹In RESTful services, operations are exposed as resources [65].

data becomes available without explicitly defining any control flow constructs [83]. In a hybrid workflow, some parts are control-driven, while others are data-driven [82]. Figure 2 illustrates a generic workflow that executes the operation op_1 , then decides to invoke either op_2 or op_3 and, finally, triggers the operations op_4 and op_5 in parallel. Operation invocations happen regardless of the workflow kind.




Figure 2: A generic workflow.

Workflows are increasingly important for IoT systems because they allow the integration of IoT services into complex tasks that automate a specific context [84, 85, 86, 87, 88]. For example, a smart home can be automated with a workflow that regulates the temperature of a room according to environmental changes. In the domain of smart agriculture, a workflow can be defined to analyse data coming from harvest sensors, predict diseases and react accordingly.

2.3. Scalability of IoT Systems

With the advent of hardware technologies, the number of IoT services is rapidly growing due to the excessive increase in the number of connected things. Currently, there are about 19 billion connected things, and it is predicted that this number will grow exponentially in the coming years [1, 2]. Thus, unlike traditional enterprise systems, scalability becomes a crucial concern for the full realisation of IoT systems.

Typically, scalability is the capability to handle increasing workloads in a IoT system [3, 4, 5, 6, 7, 8]. In this case, it is a metric that indicates how system performance improves over time. Workloads are typically measured in terms of either the number of requests dispatched [3] or the data streams generated [7]. The overall goal of scalable solutions is to enhance the Quality of Service (QoS) for guaranteeing a certain level of performance under the presence of high workloads, e.g., by minimising bandwidth, energy, latency and response time while maximising throughput. To quantitatively measure QoS, several network aspects of a service are considered such as jitter, throughput, packet loss and availability [8].

Currently, there are two kinds of scalability: vertical and horizontal.² Vertical scalability (or scaling up) [9, 10, 11] refers to the addition or removal of computing resources in a single thing, e.g., adding more memory to increase buffer size or adding more processor capacity to speed up processing. On the other hand, horizontal scalability (or scaling out) [2, 12, 13] involves the addition or removal of things in an IoT system. Its goal is to distribute the workload over multiple things to decrease individual loads, minimise response time and

²IoT cloud environments benefit from dynamically scaling vertically, horizontally or both.

enhance concurrency. Figure 3 depicts the contrast between vertical and horizontal scalability.

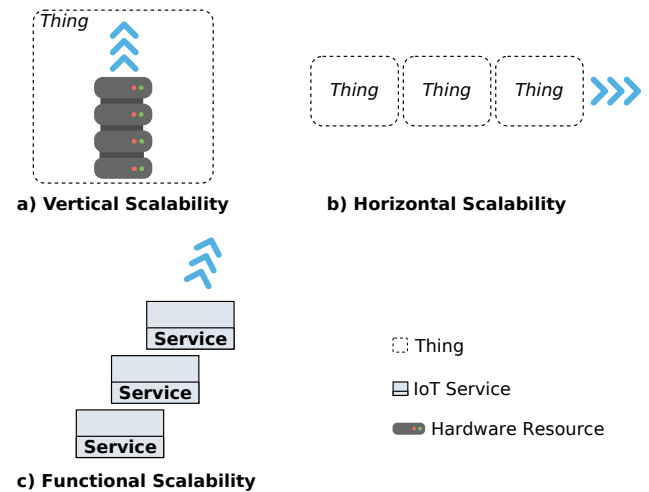


Figure 3: Scalability of IoT systems.

Both vertical and horizontal scalability have been extensively addressed in the literature [4, 10, 11, 14, 15, 16, 17, 18, 19, 20], unlike *functional scalability* which we refer to as the capability to accommodate growth in terms of the number of services composed in an IoT system (see Figure 3(c)). In particular, it enables the composition of any number of services, without severely impacting global system properties such as performance, maintenance, evolution and monitoring. Hence, functional scalability is crucial for dealing with IoT systems composed of billions of services. Figure 4 shows that it is orthogonal to vertical and horizontal scalability.




Figure 4: Scalability dimensions.

Like the other kinds of scalability, functional scalability requires the definition of metrics to measure the degree of satisfaction for accommodating new services. In this paper, we propose six qualitative metrics which we discuss in Section 5. We do not claim that such metrics are complete since quantitative metrics for QoS, identified for vertical and horizontal scalability, can also be important. However, quantitative metrics are only applicable to specific implementations. As service composition is an abstraction rather than a concrete implementation, we strongly argue that qualitative metrics are the best ones to measure the degree of satisfaction of functional scalability in service composition mechanisms. For the rest of the paper, the terms scalability and functional scalability are used interchangeably.

3. Related Work

This section presents the current IoT service composition mechanisms that define behaviour by workflows, namely (centralised and distributed) dataflows, (centralised and distributed) orchestration, choreography, and a novel composition mechanism called DX-MAN. We particularly focus on the fundamental semantics of these mechanisms instead of addressing specific implementation technologies. This is because semantics constitutes general theory that defines how to compose services conceptually rather than a concrete implementation (that can only be evaluated in specific scenarios). Significantly, fundamental semantics underlies so-called composition algorithms [8, 89, 90, 91, 92, 93, 94, 95, 96, 97, 98, 99, 100, 101, 102], programming frameworks [84, 103, 104, 105, 106, 107, 108], languages [109, 110] and platforms [111, 112, 113, 114, 115, 116, 117], which have been somehow confusingly included in existing “IoT service composition” surveys [3, 5, 24, 118, 119].

It is also essential to mention that IoT service composition is just another name for traditional SOA composition and it is done regardless of so-called service “architectures” such as the ones defined for Microservices. Microservice Architecture [120, 121, 122] has gained considerable attention in the last few years and is becoming increasingly important and popular for the development of IoT systems [123, 124, 125, 126, 127]. Every Microservice Architecture is an SOA, but not the other way round [128]. Hence, the service composition mechanisms presented in this sub-section can be used interchangeably in both Microservices and traditional SOA services. In fact, there are no composition mechanisms specifically aimed for Microservices.

Dataflows, or *Flow-Based Programming* [129, 130, 131], is a composition mechanism that defines implicit workflows as directed graphs where vertices receive input data streams, carry out some computation and pass the result to other vertices via an edge for further processing [71, 132]. A centralised approach [83, 116] fully coordinates the exchange of data streams exogenously (i.e., from outside services), whereas a distributed one [37, 38, 133, 134] partitions a complex workflow dataflow over multiple coordinators. Currently, dataflows is the most popular IoT service composition mechanism, so there are plenty of available technologies for defining data-driven IoT workflows [83, 126, 135].

Orchestration defines explicit workflow control flow structures to coordinate the invocation of service operations exogenously and, like dataflows, it can be centralised or distributed. A centralised orchestration [72, 136, 137, 138] has full control over all the services composed, whereas a distributed approach (also known as “decentralised orchestration” [139, 140, 141, 142, 143, 144, 145, 146, 147, 148]) defines sub-workflows for a collaborative exchange of workflow control flow over the network. Although orchestration is not as popular as dataflows, there are some implementations available for this mechanism to support the definition of control-driven IoT workflows [149, 150, 151].

Choreography [136, 137, 138, 152, 153] is another composition mechanism that defines workflow control flows for

the invocation of operations in services. Unlike orchestration, it relies on a public protocol which specifies a global “service conversation” via decentralised interactions, and it is modelled using a choreography modelling language [122, 138, 153]. When a choreography is enacted, the composed services exchange control according to the protocol to realise decentralised workflows. It is because of this decentralised nature that there is an increasing trend of implementing technologies for choreographing IoT services [154, 155, 156]. Notably, some of these technologies (e.g., [155]) implement choreographies based on the data-driven paradigm with no notion of public protocols, which do not define any composite service but a bunch of interactions via events or decentralised message passing. So, they do not follow the standard definition as noted by [35].

DX-MAN [157, 158, 159] is a model that uses exogenous composition operators [160, 161] to algebraically compose IoT services in a hierarchical bottom-up manner. The result of composition is not a workflow, but an IoT composite service which is semantically equivalent to a potentially infinite family of (explicit) workflow control flows. As DX-MAN takes the best properties from choreography and orchestration, it enables decentralised data exchanges over the network while decoupling services via (exogenous) workflow control flows [162]. Currently, there is only one platform available to support the definition of algebraic IoT systems [163].

To continue the discussion, Section 6 describes the above composition mechanisms in detail and presents a concrete analysis in terms of functional scalability requirements.

4. A Large-Scale IoT Scenario: Smart Parking System (SPark)

To illustrate the context for scalability requirements, this section presents a large-scale IoT scenario in the smart parking domain. The scenario tackles a common problem of large cities and is described as follows.

With the increase of population, large cities have had to deal with daily traffic congestion caused by drivers actively searching for parking spaces, especially during rush hours. As a consequence, there are increased carbon emissions as well as waste of commuters’ time and money [164, 165]. This section presents a large-scale smart parking system, *SPark*, for self-driving vehicles which efficiently find (and reserve) the nearest parking space in a smart city. *SPark* thus helps to improve parking space utilisation, shorten parking search time, reduce environmental pollution, minimise parking costs and fuel consumption, and alleviate traffic congestion [164].

SPark operates in a smart city with plenty of parking spaces equipped with occupancy sensors whose data is managed by Infostations. Although it is an urban infrastructure device able to collect up-to-date occupancy status from all sensors in range, an infostation only pulls data from the sensors near a vehicle. Figure 4 illustrates *SPark* with four self-driving cars, where a vehicle gets its location and requests a parking space from the nearest Infostation. The InfoStation then pulls data from the nearby sensors to determine the near-




Figure 6: *SPark* services.

est parking space that is free. To avoid two different vehicles chasing the same parking space, the space is reserved and paid for in advance. Finally, the vehicle displays the desired route and drives towards the selected parking space. Figure 5 depicts the general workflow of *SPark*.

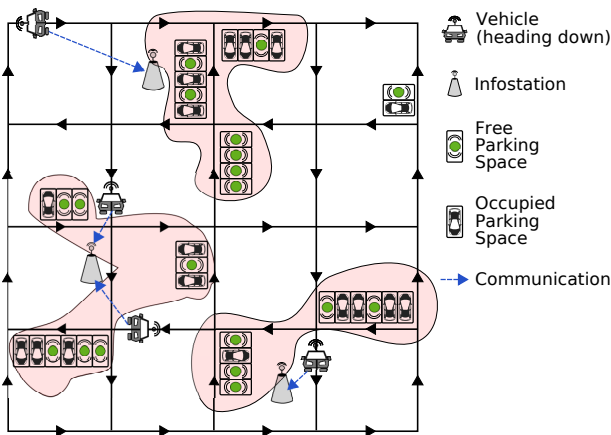


Figure 4: A smart city with *SPark*.




Figure 5: *SPark* workflow.

To achieve its goal, *SPark* composes a huge amount of IoT services distributed across a city. It is then an ultra large-scale system [23, 24, 25, 26] because sub-systems (i.e., services) may potentially integrate other sub-systems (i.e., services) and so on. Figure 6 depicts the relationship between services and sub-services in *SPark*. We distinguish between *atomic services* and *composite services*. An atomic service is the most primitive unit with no internal structure, whilst composite services integrate sub-services and can be integrated into even more complex composites.

To hide complexity and protect behaviour integrity, we assume that the composite services of our example are encapsulated. As IoT composite services potentially reside on differ-

ent business domains [166], *SPark* cannot be “decomposed” using a top-down approach, but it should be composed in a bottom-up fashion by reusing existing compositions [27, 31]. Figure 6 shows the resulting hierarchical service relationship.

Note that ellipses indicate that services may provide multiple operations or compose many other services. Also note that, in practice, atomic services can be defined as composite services. For example, *DrivingCtr* could be a complex service that integrates services for planning a path and controlling a vehicle. As another example, the *Display* service may internally use a web mapping service and a visualisation service for displaying a route. An alternative (larger) view of *SPark* can be found in Appendix A. For clarity, this section only shows a simplified version of *SPark* which is incrementally described below.

The *SensorNetwork* composite is a wireless network that composes a group of dedicated atomic *Sensor* services. It provides the *SensorNetwork.pullStates* operation as an interface for collecting parking space status in parallel. The collection is done by invoking the *Sensor.pullState* operations from the sensors near the vehicle.

The *CityManagement* composite service composes multiple *SensorNetwork* composites and one atomic *Booking* service. It provides the *CityManagement.getParking* operation for finding and reserving the best (i.e., the free and nearest) parking space. To do so, it collects sensor states with the *SensorNetwork.pullStates* operation, and then determines which parking spaces are free. Finally, it reserves the nearest free parking space using the *Booking.book* operation.

The *Payment* composite is an online electronic payment service that composes two payment providers: *Provider1* and *Provider2*. It chooses which payment method to use when the operation *Payment.pay* is invoked. On the one hand, *Provider1* is an atomic service with the operations *payVisa* and *payMaster* (for paying with Visa or Mastercard). On the other hand, *Provider2* is an atomic service with the operation *payWallet* (for paying with an eWallet).

The *Vehicle* composite encompasses two different behaviours, and composes the atomic services *GPS*, *Display* and *DrivingCtr*. The *Vehicle.getCurrentLocation* operation is an interface for the behaviour of *GPS.getLocation*, which gets geospatial positioning information. The *Vehicle.driveVehicle* operation is more complex, as it invokes *Display.showMap* so as to plan, compute and visualise the route on the vehicle’s display. Then, it executes *DrivingCtr.drive* to autonomously drive the vehicle towards the desired parking space.

SPark is the most complex composite service, since it composes the services *CityManagement*, *Payment* and *Vehicle*, and provides the *SPark.findBestParking* operation for the encapsulation of the whole system's behaviour. The behaviour is a workflow for the sequential invocation of *Vehicle.getCurrentLocation*, *CityManagement.getParking*, *Payment.pay* and *Vehicle.driveVehicle*. Appendix B describes the complete workflow of our scenario.

As the *SPark* workflow spans multiple administrative domains, our motivating example requires the distribution of services across different geographical locations. Table 1 shows a possible physical deployment configuration. In particular, *CityCouncilCloud* is a virtual thing where services *CityManagement*, *Payment* and *Booking* are deployed. *ParkingSpace* is a physical thing whose occupancy status is provided by the *Sensor* service. *Infostation* is a physical thing that collects data from sensors within a wireless range, via the *SensorNetwork* service. *ProviderServer1* and *ProviderServer2* are physical things of different payment provider companies, where services *Provider1* and *Provider2* are deployed. *Vehicle* is a physical thing that moves around the city, searching for a parking space, which provides the services *SPark*, *Vehicle*, *GPS*, *Display* and *DrivingCtr*. Note that deploying a composite service (e.g., *SPark*) does not necessarily mean that the composed services should be deployed on the same thing. This is because IoT services operate on different administrative domains.

Thing	Composite Services	Atomic Services
<i>CityCouncilCloud</i>	<i>CityManagement</i> , <i>Payment</i>	<i>Booking</i>
<i>ParkingSpace</i>		<i>Sensor</i>
<i>Infostation</i>	<i>SensorNetwork</i>	
<i>ProviderServer1</i>		<i>Provider1</i>
<i>ProviderServer2</i>		<i>Provider2</i>
<i>Vehicle</i>	<i>SPark</i> , <i>Vehicle</i>	<i>GPS</i> , <i>Display</i> , <i>DrivingCtr</i>

Table 1
SPark service distribution.

5. Functional Scalability Requirements

This section presents the functional scalability requirements of IoT systems in terms of *SPark*. These requirements were derived from an extensive review of the literature on large-scale IoT systems and serve as the foundation of our evaluation framework. The requirements are: (i) explicit control flow [29, 30]; (ii) distributed workflows [23, 31, 32]; (iii) location transparency [20, 33]; (iv) decentralised data flows [23, 32, 34]; (v) separation of control, data and computation [35, 36, 37, 38]; and (vi) workflow variability [23, 39, 40, 41].

5.1. Explicit Control Flow

Control flow defines the execution order of composed services [71], according to sequencing, branching or parallelising constructs [75]. It is explicit when there is a visible construct defining control flow in a service composition, whilst it is implicit when it needs to be inferred from the collaborative

interaction of the composed services [71, 74, 138].³ Figure 7 describes the difference between such concepts. The left side depicts visible constructs for the sequential execution of $A.op_{A1}$ and $B.op_{B1}$. Contrastingly, the right side does not show any control flow constructs, as the workflow logic is hardcoded in the computation of the composed services.




Figure 7: Explicit control flow vs implicit control flow.

The number of composed services may become overwhelmingly large because of both the huge number of things involved and the complexity of workflow control flow. Consequently, execution failures become unavoidable, more challenging to manage and may potentially unleash catastrophic consequences (for individuals or societies) [167]. For that reason, explicit control flow becomes crucial to tackle functional scalability, since it facilitates monitoring, tracking, verification, maintenance and evolution of large-scale IoT composite services [29, 30, 122, 168, 169]. For instance, we can leverage explicit control flow to detect abnormalities in a *SPark* execution [170, 171, 172], or we could easily obfuscate workflow control flow in order to avoid malicious reverse-engineering [173, 174].

Imagine that *SPark* suddenly stops working because of a bottleneck in some service, so developers want to analyse the system execution flow to find out where the problem is. For this, they can leverage explicit control flow to display a visual representation of the execution paths that *SPark* has taken [29, 74, 175, 176]. By looking at the blueprint, developers can identify the services that perform poorly and react accordingly.

Implicit control flow has historically been a barrier for functional scalability since it hinders control flow visualisation, especially when the number of services increases [29, 30, 121, 122, 177, 178]. Although there are attempts to visualise control flow [178, 179, 180, 181, 182], transforming implicit control flow into an explicit one is still a challenging task, especially when the system workflow is complex. This issue has made software development companies stop composing services with implicit control flow. Netflix [183] is the most prominent case, which has particularly expressed its concern for monitoring composite service executions when control flow is implicit. To tackle this, it has recently developed Conductor [29] to move from compositions with implicit control to compositions with explicit control.

³For example, control flow is explicit in imperative languages and implicit in declarative ones.

5.2. Location Transparency

Large-scale IoT systems are inherently dynamic and uncertain due to the presence of churn in the operating environment [20, 25, 27, 33, 184, 185, 186]. Churn means that things (and their services) dynamically connect and disconnect from the network as a result of auto-scaling, software upgrades, failures, poor connection and mobility. It is particularly evident when an IoT system uses resource-constrained things with a poor connection [187], or when there are mobile things involved [185, 188]. Hence, churn results in physical service locations (i.e., IP addresses) changing over time frequently. For example, the Spark's *Sensor* service is a resource-constrained thing with limited connectivity, which is likely to experience network disconnections. Similarly, the service *Vehicle* may change its IP address because of its high mobility. The same happens to *Provider1*, *Provider2* and *Booking* as they are deployed on the Cloud.

To deal with churn, IoT systems require location transparency, in order to ensure that atomic services are unaware of the physical location of other services [74]. Figure 8 illustrates this concept. In particular, the lower section shows that the (non-static) IP address of $B.op_{B1}$ is hardcoded in $A.op_{A1}$. The upper section shows a scenario where there is support for location transparency.




Figure 8: Location Transparency.

Location transparency is typically achieved with a service discovery mechanism that dynamically queries a central service registry [52]. In client-side discovery, service providers register at startup in a registry which is later queried by service consumers. In server-side discovery, a router (e.g., a service bus like a gateway [112]) queries the registry and forwards requests to an available service provider on behalf of service consumers.

A service composition with no location transparency requires service consumers to know the location of service providers in advance. For example, without any support for location transparency, *SPark* would have to be updated every time a *Vehicle* changes location. Nonetheless, this is worrying because there is a massive number of vehicles constantly moving around a city and, therefore, changing location. Addressing this issue with a service discovery mechanism entails the addition of “intrusive” elements to the composition (e.g., a service registry or a naming service) which are not part of the semantics of a composition mechanism. In fact, an atomic service computation would be tightly coupled with

the “intrusive” element, since the former needs to be updated whenever the location of the latter changes. We shall keep this in mind in Section 6.

5.3. Distributed Workflows

IoT service composition can define a *centralised* or a *distributed* workflow. It is centralised when a single entity governs the entire workflow control flow. Contrastingly, it is distributed when multiple entities collaboratively define control flow [74]. Figure 9 illustrates the contrast between a centralised and a distributed workflow. The upper part shows a central composite service that defines a workflow for the sequential invocation of $A.op_{A1}$, $A.op_{B1}$, $A.op_{C1}$, $A.op_{D1}$ and $A.op_{E1}$. The lower part shows a possible distribution of the same workflow control flow over three composite services deployed on different things.

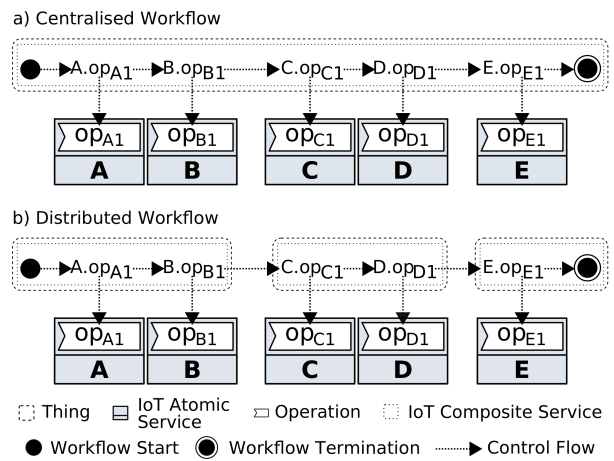


Figure 9: Centralised workflow vs distributed workflow.

Although a centralised workflow facilitates the design and maintenance of an IoT system, implementing a large-scale IoT system using such an approach is inefficient and infeasible [133]. This is because a central entity constitutes a single point of failure and attack, and leads to a performance bottleneck since all control and data goes through it [27, 34, 140, 142, 144, 148, 189, 190, 191]. Moreover, a central entity negatively impacts the availability of an IoT system.

The distributed nature of IoT requires control and data to pass through geographically dispersed entities (potentially deployed on different business domains) [32, 37, 133, 166]. For that reason, no single domain should govern the entire composition workflow, since multiple domains (perhaps in the order of millions) may want control over their own workflow part [74, 190, 192]. For instance, the *CityManagement* composite may potentially be controlled by the City Council, and a payment provider company could operate the *Payment* composite. This issue clearly implies that every administrative domain should be able to manage their own workflow definition. Thus, a *distributed workflow* should be part of any IoT service composition mechanism, since it allows interoperability between different domains and improves load balancing, throughput and availability of an IoT system [74, 87].

5.4. Decentralised Data Flows

Data flow defines the way data elements are moved from one service to another in a service composition [71]. This process can be done in three different ways: *centralised*, *distributed* or *decentralised*.

In the centralised approach [112, 136, 193, 194], a single composite mediates the exchange of data between services (see Figure 10(a)). Although it facilitates data management, the mediator becomes a potential bottleneck and introduces additional network hops for data exchange, since all data passes through it. This negatively impacts QoS by increasing the response time, and leads to network congestion [34, 141, 142, 190, 195], especially when there are plenty of IoT services exchanging huge amounts of data continuously [24, 25, 187, 190, 196].

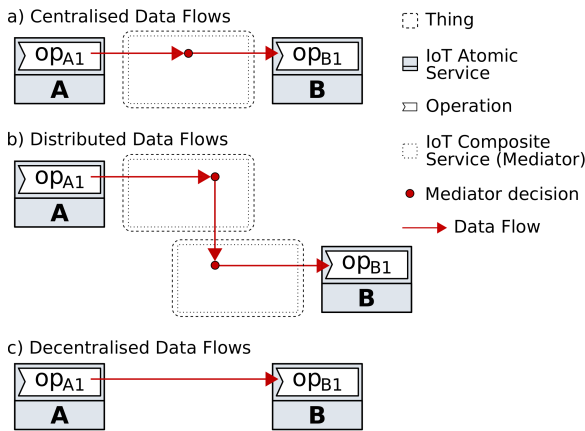


Figure 10: Data flow approaches.

To avoid a single bottleneck, the distributed approach [133, 139, 140, 141, 142, 143, 144, 146] removes the central composite and distributes the load of data over multiple composites (see Figure 10(b)). Although this improves load balancing, it introduces unnecessary network overhead as data passes through many mediators, even if data is unimportant for them, i.e., the more mediators, the more network overhead.

IoT services must exchange data as efficiently as possible (by minimising network hops), in order to avoid performance bottlenecks, achieve better response time and improve throughput [147]. A decentralised approach provides the most suitable data exchange for service composition since it requires only one network hop to pass data directly from a service producer to a service consumer (see Figure 10(c)) [34, 117, 152, 162, 195, 197, 198, 199]. For example, the data generated by *Booking.book* should be passed to *Display.showMap*, without passing through any other entity that does not require the produced data (e.g., the *CityManagement* composite or the *Vehicle* composite). In general, the amount of data transmitted in *SPark* may potentially be huge due to the large number of services involved of which sensors generate data continuously. Therefore, *SPark* requires a composition mechanism with support for decentralised data flows to provide an optimal QoS.

5.5. Separation of Control, Data and Computation

Large-scale IoT systems may potentially span multiple administrative domains and exhibit a high degree of heterogeneity in many different forms [25, 74, 200]. For instance, there may be different service providers (e.g., Amazon AWS IoT and IBM Watson), a wide variety of programming languages (e.g., Swift and embedded C), multiple operating systems (e.g., Contiki and TinyOS) and different network communication protocols (e.g., CoAP and MQTT). For that reason, service composition mechanisms must provide the means to deal with such heterogeneity, in order to compose large-scale cross-domain IoT systems. This is, in fact, one of the major challenges for building IoT systems [200].

To enable flexible service composition in such heterogeneous environments, control flow, data flow and computation should be orthogonal [35, 36, 38, 76, 144, 201, 202]. This would enable separate reasoning of concerns so system developers can focus on IoT service composition, while service developers focus on efficient service functionality [37, 74, 76]. Consequently, services are workflow agnostic because workflow control flow is never embedded in the computation of many services. Moreover, the separation of control and computation facilitates workflow monitoring [29].

Figure 11 illustrates a separation of control, data and computation. There are two computations defined in *A.opA1* and *B.opB1*. The control flow part defines a workflow for the sequential invocation of *A.opA1* and *B.opB1*, without considering data flow and computation. The data flow part defines intermediate processing of the output from *A.opA1*, before sending it to the input of *B.opB1*. When control is mixed with data, such processing represents an extra control flow step which is intrusive to the original workflow.

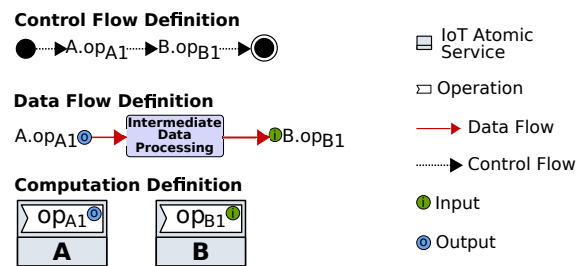


Figure 11: Separation of control, data and computation.

The separation of data flow will additionally enable separate data management for improving performance without considering control flow [35, 162]. Similarly, control flow can be used in isolation for defining efficient deployment strategies that do not consider data flow. Moreover, different technologies can be used for implementing control flow, data flow and computation separately [35]. For instance, developers of *SPark* could use CoAP for passing control between services, the Blockchain for managing data flows and embedded C for implementing service computation. Providing independent reasoning of concerns also enables a separate validation and verification (V & V) of services.

When control is mixed with computation, a service

provider and a service consumer are tightly coupled because invocations are originated in the consumer's computation. Thus, changing workflow control flow in one service requires further changes in the respective computation. This rigidity evidently limits the scalability of IoT systems because it is difficult to add, remove, change and replace services (and system workflow control flow). For that reason, reusing services is not possible at scale when control is mixed with computation [29, 37, 71, 74, 203].

Suppose there is a branching control flow structure in the computation of *Booking.book*, for invoking either *Provider1.payVisa* or *Provider2.payWallet*. As a result, *Booking.book* is not reusable as it is since it must be adapted to the requirements of other systems, e.g., it must be updated when two new payment providers come into play (e.g., *Provider3* and *Provider4*). As another example, SPark could need the parallel execution of *Payment.Pay* and *Vehicle.driveVehicle*, in order to improve concurrency. In this case, the sequence of invocations hardcoded in *SPark.findBestParking* must be replaced with a parallel control flow structure.

Generally speaking, large-scale IoT systems require separate reasoning of control, data and computation for independent maintenance, validation, verification, reuse and evolution of those concerns. This separation of concerns could also result in reduced time to market and reduced software production and maintenance costs [74].

5.6. Workflow Variability

Large-scale IoT systems operate in highly dynamic environments subjected to variability caused by external or internal factors [39, 41, 204, 205]. External factors are beyond the scope of the system and include changes in requirements and increasing workloads. Internal factors are associated with the system operation and include system failures and sub-optimal behaviours.

A service composition mechanism must support the definition of alternative behaviours, in order to adapt a composite to changes in both the external and the internal environment. As manually choosing alternative behaviours is a costly and inefficient management process, when there are many services in the composite, behaviours must be selected autonomously (i.e., with no human intervention) [206]. Thus, variability of behaviour is a crucial desideratum for the realisation of large-scale autonomous IoT systems [23, 39, 40, 41, 87, 207].

Workflow variability [157, 208] allows the definition of alternative control flow constructs (i.e., behaviours) in a service composition (see Figure 12), and it is particularly useful for autonomously changing workflows at runtime. For instance, the *SensorNetwork* composite may define variable parallel behaviours for pulling data, depending on a vehicle's location. For one vehicle, it may pull data from *Sensor1* and *Sensor2*, whilst for another one it could pull data from *Sensor10*, *Sensor29* and *Sensor34*. As another example, consider multiple payment service providers that offer different QoS (which is always fluctuating according to different workloads). For this, a variable branching construct can be used to

dynamically choose the *Provider* service with the best QoS.

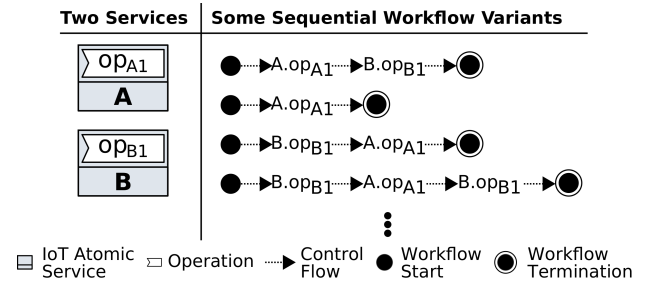


Figure 12: Workflow variability.

When there is a family of related compositions, a *Greatest Common Denominator* [208] is a base variability model that can be reused for defining alternative compositions (also known as configurations). For example, suppose a SPark vehicle can be either *manual* or *self-driving*. A manual vehicle requires the services *GPS* and *Display*, whilst a self-driving one uses the service *DrivingCtr* in addition. For this, the *Vehicle* composite of Figure 6 would be the *Greatest Common Denominator*, which can be reused for defining two alternative behaviours for two different vehicles.

Workflow variability is not only useful to accommodate the requirements of different vehicles, but also meaningful for different cities and users. Consider the case of Northern Ireland where many cities offer free on-street parking, and a few others (e.g., Belfast, Lisburn and Newry) impose a tariff [209]. Here, some *SPark* workflow variants may require a *Payment* composite, while others not. At the city scale, it would become very difficult to define workflows from scratch since many services would need to be changed and customised. Users can also require different workflows according to their needs, e.g., one user may require pre-payment, whereas another one may require post-payment. However, manually changing behaviour at runtime to accommodate different user requirements is infeasible. Thus, workflow variability represents a suitable solution for tackling the above scenarios.

In order to define workflow variability, a service composition mechanism must enable *total compositionality* by which all behaviours (e.g., operations) from the composed services are semantically available in a composite service [157, 159]. This concept contrasts with *partial compositionality* where only named and selected behaviours are semantically preserved, leading to a combinatorial explosion of behaviours as the number of services increases [159].⁴ Figure 13 illustrates the dichotomy between total and partial compositionality. The lower section shows the preservation of operations *A.opA1* and *B.opB2* in the composite service, for the creation of a single workflow involving those operations. Remarkably, the upper section depicts a composite service that preserves all the operations from the sub-services, which means that alternative workflows can be created. For instance, we could define a sequential workflow for the invocation of *B.opB1*

⁴For example, there are $2^k - 1$ possible scenarios for pulling sensor data, where k is the number of sensors.

and $A.op_{A2}$, or an alternative sequence for the execution of $A.op_{A1}$, $B.op_{B2}$ and $A.op_{A1}$. These behaviours are impossible to achieve with partial compositionality which allows the definition of only one fixed workflow at a time.

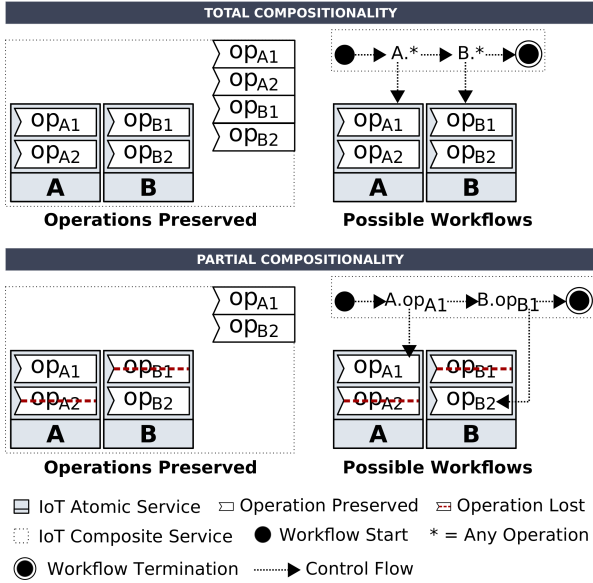


Figure 13: Total compositionality vs partial compositionality.

6. Analysis of IoT Service Composition Mechanisms

This section reviews and analyses the fundamental semantics of current IoT service composition mechanisms, namely (centralised and distributed) dataflows, (centralised and distributed) orchestration, choreography and DX-MAN. The goal is to determine how well they fulfil the functional scalability requirements of IoT systems. For this analysis, we investigate the following research questions per mechanism:

- **RQ1:** Is there any architectural entity explicitly defining control flow?
- **RQ2:** Are atomic services location-agnostic of other atomic services?
- **RQ3:** Is it possible to distribute a workflow over multiple entities?
- **RQ4:** Do atomic services exchange data directly?
- **RQ5:** Do atomic services perform computation without passing any data or control to other atomic services?
- **RQ6:** Is there any notion of workflow variants?

For $R1-R4$ and $R6$, we use a tick mark to indicate that a specific mechanism fulfils the requirement being analysed, or a cross mark to indicate the opposite. $RQ5$ is the only requirement that uses a textual representation for showing

which concerns (i.e., control, data and computation) are independent. To answer $RQ6$, we need to determine the compositionality of a mechanism which can be either total or partial (see Section 5.6). To do so, we investigate the following research questions:

- **RQ7:** What is the resulting type from composition?
- **RQ8:** How many workflows does the composition mechanism define?

6.1. Dataflows

Dataflows, or *Flow-Based Programming* [129, 130, 131], is a composition mechanism that defines a workflow using data transformations (e.g., filter, split, union and sort) as well as exogenous data exchange between services [71, 132]. A dataflow description is a directed graph where vertices are asynchronous data processing units (invoking service operations), and edges are connections for passing data streams between vertices via the network (by message passing or events). A vertex explicitly defines input ports and output ports. When it receives data from all inputs, it performs some computation and writes results in output ports. The resulting data is then moved to other vertices via an edge. This process is illustrated in Figure 14.

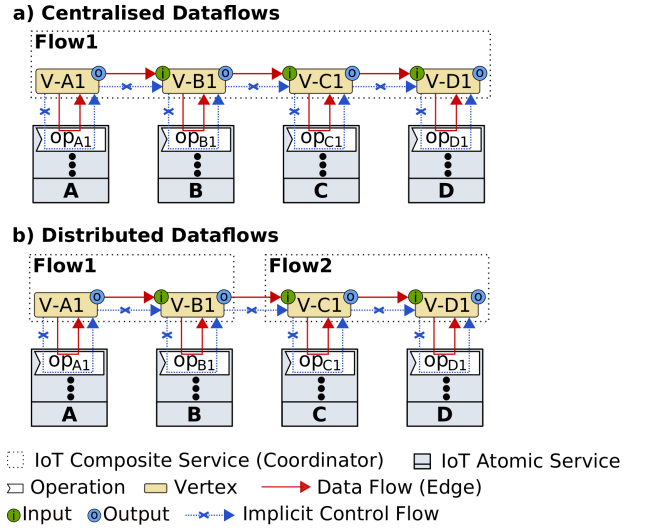


Figure 14: Composition by dataflows.

Dataflows can be centralised or distributed. A centralised dataflow [83, 116] defines a single coordinator for managing an entire graph and exogenously invokes service operations. A distributed dataflow [37, 38, 133, 134] partitions and distributes a complex graph over multiple coordinators that interact directly by exchanging data between vertices. Figure 14(a) shows a centralised dataflow for a pipeline of services A , B , C and D . When the coordinator $Flow1$ is triggered, vertex $V-A1$ invokes $A.op_{A1}$ and passes the result to the vertex $V-B1$ which, in turn, executes $B.op_{B1}$. Next, the result of $V-B1$ is passed to the vertex $V-C1$ which executes $C.op_{C1}$. Finally, the data of $V-C1$ is moved to the vertex $V-D1$ and then processed by $D.op_{D1}$. A distributed version of the same pipeline

Composition Mechanism	Composition Unit	Resulting Type from Composition	Compositionality	Number of Workflows
Centralised Dataflows	IoT Service	Workflow	Partial	1
Distributed Dataflows	IoT Service	Workflow	Partial	1

Table 2
Compositionality of dataflows.

	Centralised Dataflows	Distributed Dataflows
Explicit Control Flow	✗	✗
Service Location Transparency	✓	✓
Distributed Workflows	✗	✓
Decentralised Data Flows	✗	✗
Separation of Control/Data/Computation	Data/Computation	Data/Computation
Workflow Variability	✗	✗

Table 3
Analysis of composition by dataflows w.r.t. the scalability desiderata.

is shown in Figure 14(b), where there is an edge between *V-B1* and *V-C1* for moving data from the *Flow1* composite to the *Flow2* composite.

Dataflows are increasingly popular for composing IoT systems. In particular, they are widely used for the Internet of Data (IoD) [210] which involves data collection from multiple sources (e.g., sensors), data analysis and control of the physical world. This paradigm has been referred to as Sense-Compute-Control (SCC) [211].⁵ Currently, there are many platforms for composing IoT services using dataflows. Examples include Node-RED [83], COMPOSE [214], Glue.Things [116], LabVIEW [135], Paraimpu [215], Virtual Sensors [111], SpaceBrew [216], FogFlow [18], ASU VIPLE [217], ThingNet [126], Calvin [218], IoT Services Orchestration Layer [219], NoFlo [220] and many others [221, 222, 223, 224].

As the Web 2.0 became more data-centric and user-friendly [65, 71], dataflows have gained popularity for IoD through mashups. Mashups [225] are realised by dataflows [71, 226], and they allow the composition and visualisation of data streams on a graphical user interface displayed on the Web [65, 71, 227]. Examples of mashup tools include WoTKit [228], IoTaaS [196] and Clickscript [56].

Dataflows have been accepted as coordination languages since a graph is defined in a coordinator that exogenously invokes services according to a dataflow description [129, 130, 133]. A dataflow graph is typically created with a graphical editor and executed by an engine (i.e., the coordinator), and it is triggered by either timing constraints or events.

A dataflow coordinator is the only entity aware of service locations, and provides separation between data and computation. This separation allows service developers to focus on data stream computation while system developers wire up

⁵Do not confuse data analysis tools [212, 213] with dataflows. A data analysis tool is a software that allows the collection, storing, indexing, processing, monitoring and visualisation of data. On the other hand, dataflows specify how data is passed between services according to a dataflow graph specification.

vertices exogenously [133]. However, despite the aforementioned advantages, decentralised data flows are not supported as data streams always pass through data flow coordinators.

Although passing data between vertices is explicitly defined in a dataflow graph, control flow is implicit in the collaborative data stream exchange. This is because control flow statements are not visible in the graph specification.

In general, a graph specification is a single flat workflow of named and selected service operations. It might seem that a distributed dataflow provides multiple workflows. However, this is not true because a distributed dataflow graph is just a single nested workflow, distributed over different coordinators. Thus, dataflows only provide partial compositionality and do not support workflow variability (see Table 2).

Table 3 summarises the results of our analysis of data flows w.r.t. the scalability desiderata. Both centralised dataflows and distributed dataflows have similar characteristics. The only difference is that the latter provides support for distributed workflows by partitioning a dataflow graph over multiple coordinators.

6.2. Orchestration

Orchestration can be centralised or distributed. Centralised orchestration [72, 136, 137, 138] describes interactions between services from the perspective of a central coordinator (also known as orchestrator) which has control over all parties involved. It explicitly defines workflow control flow to coordinate the invocation of service operations, in order to realise some complex function that cannot be achieved by any individual service [69, 70, 229]. In a distributed orchestration, also known as “decentralised orchestration” [139, 140, 141, 142, 143, 144, 145, 146, 147, 148], multiple coordinators collaboratively define workflow control flow.

An orchestration is typically defined using a workflow language such as BPEL [70, 81, 230, 231, 232, 233, 234, 235] or BPMN [86, 236]. The resulting workflow has tasks for passing control among services according to explicit control flow constructs (for sequencing, parallelising, branching and

Composition Mechanism	Composition Unit	Resulting Type from Composition	Compositionality	Number of Workflows
Centralised Orchestration	IoT Service	Workflow	Partial	1
Distributed Orchestration	IoT Service	Workflow	Partial	1

Table 4
Compositionality of orchestration.

	Centralised Orchestration	Distributed Orchestration
Explicit Control Flow	✓	✓
Service Location Transparency	✓	✓
Distributed Workflows	✗	✓
Decentralised Data Flows	✗	✗
Separation of Control/Data/Computation	Control/Computation	Control/Computation
Workflow Variability	✗	✗

Table 5
Analysis of composition by orchestration w.r.t. the scalability desiderata.

looping) [71]. In distributed orchestration, the interaction between coordinators can be done in three different ways: (i) *with an extra task* [140], (ii) *with two extra tasks* [141, 142, 148] or (iii) *without any extra task* [139, 144, 145]. In (i), an orchestration invokes the interface of another one using an external task to the system’s workflow control flow. In (ii), there are two different tasks for receiving and passing control (and data) between two orchestrations. Finally, in (iii), two orchestrations interact by moving control (and data) directly between the tasks of the system’s workflow control flow.

An orchestration engine is responsible for executing a workflow process by invoking service operations in a given order. Although traditional engines can be used (e.g., Camunda BPM workflow engine [86, 237], Activiti [238, 239] and AWS Step Functions [151]), recently we have seen the emergence of orchestration engines particularly designed for IoT systems (e.g., PROtEUS [149] and [150]).⁶

Figure 15(a) illustrates a centralised orchestration for the services *A*, *B*, *C* and *D*, where the coordinator *Orch1* defines a “composite service” for the sequential invocation of *A.op_{A1}*, *B.op_{B1}*, *C.op_{C1}* and *D.op_{D1}*. Three distributed versions are depicted in Figure 15(b). In the former, *Orch1* defines a “composite service” for the sequential execution of *A.op_{A1}* and *B.op_{B1}*, and then uses an extra task to pass control (and data) to *Orch2*. *Orch2* defines another “composite service” to sequentially invoke *C.op_{C1}* and *D.op_{D1}*. The second distributed version uses two extra tasks for passing and receiving control (and data) between *Orch1* and *Orch2*. Finally, in the last distributed version, control and data are passed directly from *B.op_{B1}* to *C.op_{C1}*.

A glance at Figure 15 reveals that services are workflow agnostic because an orchestrator is the only entity aware of the location of other services. Having workflow agnostic services implies that an orchestrator provides separation between control flow and computation. This separation allows service

⁶A workflow engine can be deployed on either a specialised server [237] or a service bus like a Gateway [112, 193, 194].

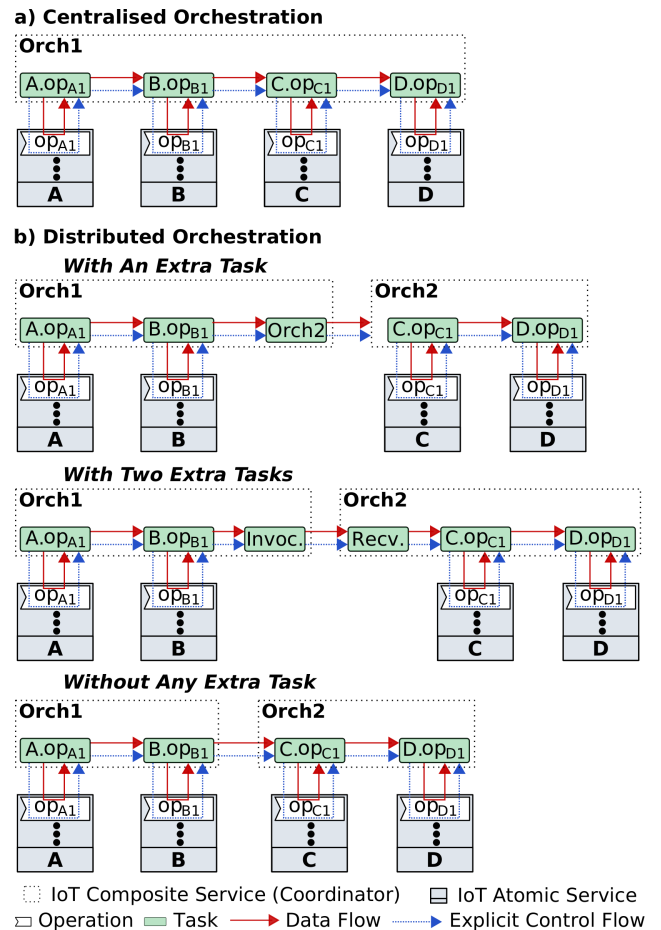


Figure 15: Composition by orchestration.

developers to focus on efficient service functionality (i.e., computation) while system developers focus on workflow control flow in orchestrator(s).

Although a central coordinator facilitates the management of control flow logic, it easily becomes a performance

Composition Mechanism	Composition Unit	Resulting Type from Composition	Compositionality	Number of Workflows
Choreography	IoT Service	Workflow	Partial	1

Table 6
Compositionality of choreography.

bottleneck because all data passes through it [140, 142]. This is because data follows control [71, 162]. Furthermore, the resulting composite service is a single flat workflow, for the invocation of named and selected operations, which must be transformed into a service [159, 240]. For that reason, orchestration only provides partial compositionality and does not support workflow variability (see Table 4).

Table 5 summarises the results of our analysis of orchestration w.r.t. the scalability desiderata, where we can see that centralised orchestration and distributed orchestration have similar characteristics. The only difference is that the latter supports distributed workflows via the partition of workflow control flow over multiple coordinators.

6.3. Choreography

A choreography describes service interactions from a global perspective using a public contract (also known as protocol) [136, 137, 138, 152, 153]. The contract specifies a “conversation” among participants via decentralised message exchanges, which can be modelled by a global observer using a choreography modelling language [122, 138, 153].⁷ An interaction-based model allows the definition of event-driven or request-response messages by connecting required and provided interfaces. Examples include WS-CDL [241] and Let’s Dance [242]. An interconnected interface model, on the other hand, allows the specification of control flow per participant. Examples include BPEL4Chor [240], Web Service Choreography Interface (WSCI) [243] and BPMN [244].⁸

A protocol defines roles for the collaborative realisation of a global workflow with no control over the internal details of the participants involved [229]. A role explicitly describes a participant workflow control flow in terms of expected and produced messages. When a concrete service instance plays a role, it must behave accordingly by exchanging messages with other instances, using either direct message passing (e.g., invoking REST APIs) or events [63, 74, 117, 122].⁹ This process is known as choreography enactment.

IoT is moving towards a more decentralised environment to reduce the bottleneck caused by centralised environments. As choreographies represent a natural way of dealing with such decentralisation, there are currently some platforms for composing IoT services by choreographies, e.g.,

⁷A Microservice architecture prefers choreography over orchestration to support decentralised workflows [122].

⁸The choice of the contract depends on the type of participants involved which can be either atomic services or orchestrations.

⁹Service participants are tightly coupled in terms of dependencies. In choreographies based on direct message-passing, services hardcode invocation calls in service computation. In event-driven choreographies, services are tightly coupled because senders and receivers agree a topic queue in advance [245].

CHOReVOLUTION [154], ChorSystem [246], Actorsphere [156], *BeC*³ [117] and TraDE [35].

Figure 16 illustrates a sequential choreography for the services *A*, *B*, *C* and *D*, where a protocol (defined with standard BPMN 2.0 notation [236]) specifies that *B.op_{B1}* expects a message from *A.op_{A1}*, *C.op_{C1}* a message from *B.op_{B1}* and *D.op_{D1}* a message from *C.op_{C1}*. When the choreography is enacted, there is a chain reaction that starts with the invocation of *A.op_{A1}* and finishes with the execution of *D.op_{D1}*.

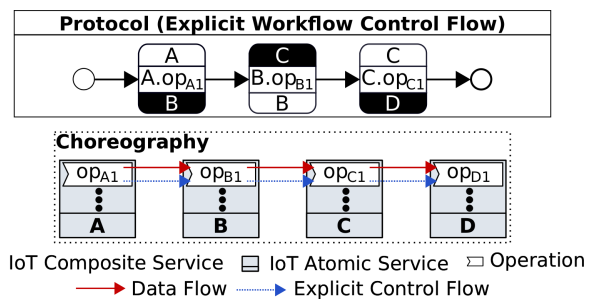


Figure 16: Composition by choreography.

Figure 16 shows that workflow control flow is explicitly defined in the protocol. During enactment, control is passed alongside data in every invocation, so services need to be aware of the location of other services. Of course, a service registry can be used, but this entails adding an “intrusive” element external to the composition (see Section 5.2).

Like orchestration, the resulting composition is a flat workflow for the invocation of selected and named operations (specified in the protocol). If the resulting workflow needs to be further composed, a choreography needs to be transformed into a service [159, 240]. For that reason, a choreography is partially compositional and workflow variability is not supported (see Table 6). Table 7 summarises the results of our analysis of choreography w.r.t. the scalability desiderata.

	Choreography
Explicit Control Flow	✓
Service Location Transparency	✗
Distributed Workflows	✓
Decentralised Data Flows	✓
Separation of Control/Data/Computation	None
Workflow Variability	✗

Table 7
Analysis of composition by choreography w.r.t. the scalability desiderata.

6.4. DX-MAN

DX-MAN [157, 158, 159] is an algebraic model where IoT services and service composition operators are first-class semantic entities. A service is a stateless distributed unit of composition which can be either atomic or composite, and its semantics is a workflow space (i.e., a family of workflow variants). A *composition operator* defines variable control flows between families of workflow variants (see Figure 17(b)).

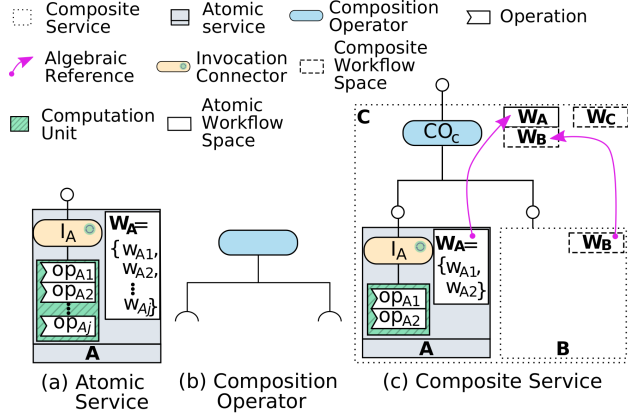


Figure 17: DX-MAN model.

An *atomic service* is formed by connecting an *invocation connector* with a *computation unit* (see Figure 17(a)). It is a finite workflow space whose elements invoke a different operation implemented in the computation unit via an invocation connector. As a computation unit is semantically identical to a traditional SOA service (because it is a set of operations), DX-MAN atomic services and SOA atomic services are semantically different (cf. Section 2.1).

A *composite service* connects a composition operator with multiple (atomic or composite) sub-services (see Figure 17(c)), which is equivalent to connecting a composition operator with multiple sub-workflow spaces. The result is a composite workflow space whose workflow variants invoke elements of atomic sub-workflow spaces and/or entire composite sub-workflow spaces, according to the control flow definition of the composition operator being used. There are composition operators for sequencing (i.e., *sequencer*), branching (i.e., *inclusive selector* and *exclusive selector*) and parallelism (i.e., *paralleliser*). Figure 18 shows that the sequencer and paralleliser operators define infinite workflow variants, whilst the branching operators define $2^n - 1$ variants s.t. n is the total number of atomic sub-service operations plus the number of composite sub-workflow spaces. For further details on this matter, see [157].

Composition Operator	Number of Workflows
<i>Paralleliser</i> (PAR)	∞
<i>Sequencer</i> (SEQ)	∞
<i>Inclusive Selector</i> (SEL)	$2^n - 1$
<i>Exclusive Selector</i> (XSEL)	$2^n - 1$

Figure 18: DX-MAN composition operators.

In addition to composition operators, DX-MAN provides special transformation operators (called adapters) for sequencing, branching, parallelising, looping and guarding over individual workflow spaces. Hence, DX-MAN is Turing complete [247]. Currently, there is only one platform that implements the DX-MAN model [163], and it is available at <https://github.com/damianarellanes/dxman>.

In DX-MAN, a composition is done incrementally in a bottom-up fashion. So, a hierarchical connection structure of operators sits on top of atomic services. Figure 19(a) shows a DX-MAN composition that involves four atomic services and three composite services. The first step is to model the atomic services A, B, C and D with the invocation connectors I_A, I_B, I_C and I_D . This process results in the atomic workflow spaces W_A, W_B, W_C and W_D , respectively. In the next hierarchy level, we create the composite services E and F . To do so, we use the composition operator SEQ_E to define a composite workflow space W_E from the atomic sub-workflow spaces W_A and W_B . Likewise, we define the composite workflow space W_F using the operator SEQ_F which operates on the atomic sub-workflow spaces W_C and W_D . Finally, we create the top-level composite service G . For this, we use the composition operator SEQ_G which defines the composite workflow space W_G from W_E and W_F .

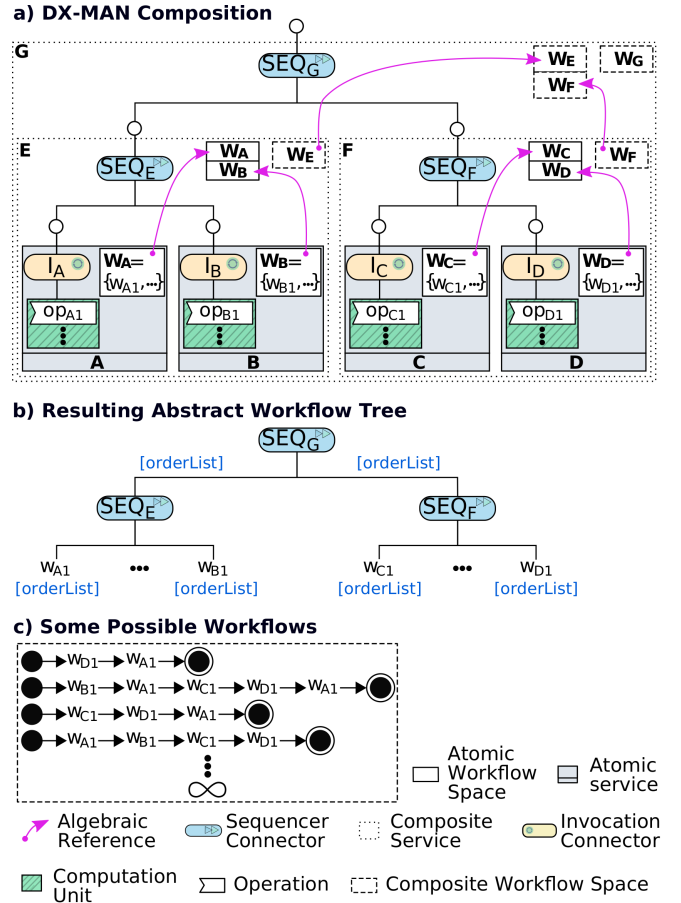


Figure 19: Composition by DX-MAN.

Note that workflow spaces from sub-services are available

Composition Mechanism	Composition Unit	Resulting Type from Composition	Compositionality	Number of Workflows
DX-MAN	IoT Service	IoT Service	Total	$[1, \infty]$

Table 8
Compositionality of DX-MAN.

in the next hierarchy level thanks to the algebraic semantics of the model. Consequently, the result of algebraic composition is not a single workflow, but a (potentially infinite) workflow space (i.e., a service) of workflow variants. Thus, DX-MAN supports total compositionality (see Table 8).

When a composition is done, an abstract workflow tree is automatically derived, which represents the hierarchical control flow structure of the composite (see Figure 19(b)). To select a particular variant, a concrete workflow tree must be created, which is just a projection function over a composite workflow space. See [157, 158] for further details.

Figure 19(b) shows the abstract workflow tree of our example from which we can choose infinite sequential workflow variants (see Figure 19(c)). One case is depicted in Figure 20, where a concrete workflow tree is created for the sequential execution of $A.op_{A1}$, $B.op_{B1}$, $C.op_{C1}$ and $D.op_{D1}$. Figure 21 shows another variant where a different concrete workflow tree is defined to invoke $A.op_{A1}$ and $D.op_{D1}$ sequentially.

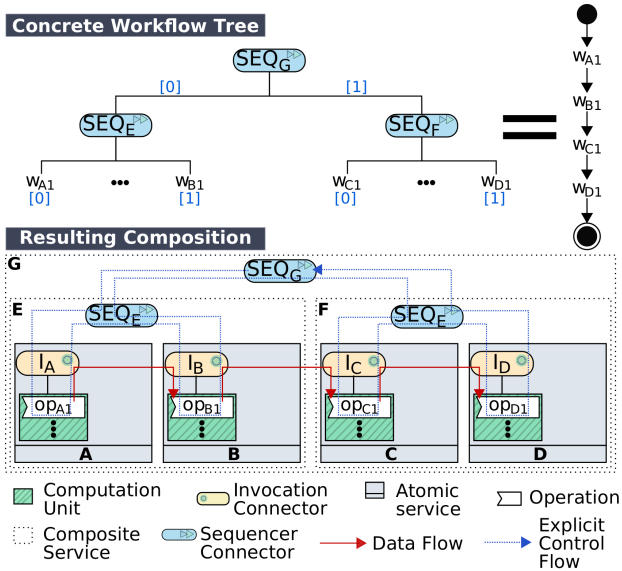


Figure 20: A workflow variant derived from Figure 19 for sequentially executing $A.op_{A1}$, $B.op_{B1}$, $C.op_{C1}$ and $D.op_{D1}$.

Note that a DX-MAN composition has composition operators for the coordination of workflow control flow only. This is because data flow and control flow are modelled separately for each workflow variant. The paper [162] describes further details on how to define data flows per variant.

A glance at Figures 20 and 21 reveals that data, control and computation are independent concerns. Thus, data is exchanged in a P2P fashion between atomic services, while composition operators coordinate workflow execution by passing control only. The coordination can be done in a distributed

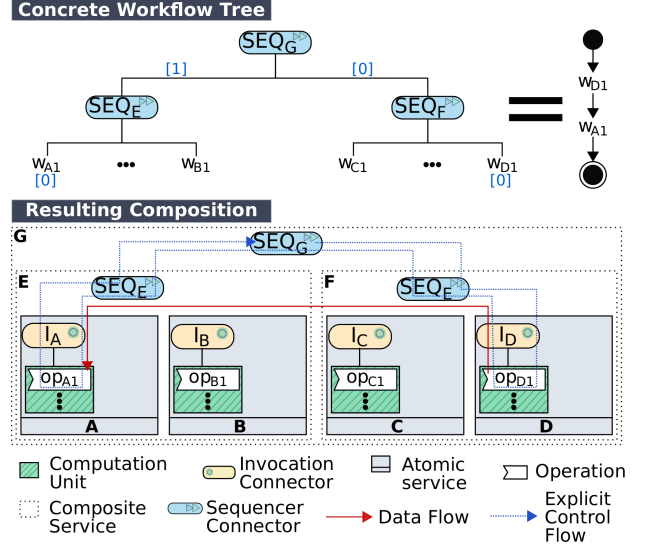


Figure 21: A workflow variant derived from Figure 19 for sequentially executing $D.op_{D1}$ and $A.op_{A1}$.

way since composition operators can be deployed on different things [159]. As coordination happens from outside services, computation units do not interact with one another, which results in independent distributed computations.

When an invocation connector receives control, it reads data from a decentralised data space (i.e., the Blockchain in the current implementation), invokes a service operation and writes results in the space. For that reason, only invocation connectors know the location of the connected computation units (i.e., service implementations). Furthermore, as invocation connectors perform operations on the data space, composition operators never exchange data during workflow execution. This (transparent) decentralised data exchange is achieved by the separation of control and data [162].

Table 9 summarises the results of our analysis of DX-MAN w.r.t. the scalability desiderata.

	DX-MAN
Explicit Control Flow	✓
Service Location Transparency	✓
Distributed Workflows	✓
Decentralised Data Flows	✓
Separation of Control/Data/Computation	Control/Data/Computation
Workflow Variability	✓

Table 9
Analysis of composition by DX-MAN w.r.t. the scalability desiderata.

Composition Mechanism	Composition Unit	Resulting Type from Composition	Compositionality	Number of Workflows
Centralised Dataflows	IoT Service	Workflow	Partial	1
Distributed Dataflows	IoT Service	Workflow	Partial	1
Centralised Orchestration	IoT Service	Workflow	Partial	1
Distributed Orchestration	IoT Service	Workflow	Partial	1
Choreography	IoT Service	Workflow	Partial	1
DX-MAN	IoT Service	IoT Service	Total	$[1, \infty]$

Table 10
Compositionality of IoT composition mechanisms.

	Centralised Dataflows	Distributed Dataflows	Centralised Orchestration	Distributed Orchestration	Choreography	DX-MAN
Explicit Control Flow	✗	✗	✓	✓	✓	✓
Service Location Transparency	✓	✓	✓	✓	✗	✓
Distributed Workflows	✗	✓	✗	✓	✓	✓
Decentralised Data Flows	✗	✗	✗	✗	✓	✓
Separation of Control/Data/Computation	Data/Computation	Data/Computation	Control/Computation	Control/Computation	None	Control/Data/Computation
Workflow Variability	✗	✗	✗	✗	✗	✓

Table 11
Analysis of scalability desiderata of IoT composition mechanisms.

7. Evaluation

This section presents an evaluation of service composition mechanisms w.r.t. compositionality and the functional scalability requirements of IoT systems.

Table 10 summarises the analysis on compositionality presented in Section 6 to answer the research questions *RQ7* and *RQ8*. It shows that DX-MAN is the only mechanism that enables total compositionality since algebraic composition yields a service with a potentially infinite number of workflow variants. The other mechanisms define only one workflow at a time as the composition result is not a service, but a single flat workflow that invokes selected and named operations.

Research questions *RQ1-RQ6* have also been studied in Section 6. They enable us to analyse how well service composition mechanisms fulfil the scalability requirements: (i) explicit control flow; (ii) location transparency; (iii) distributed workflows; (iv) decentralised data flows; (v) separation of control, data and computation; and (vi) workflow variability.

Requirements (i), (ii), (iii), (iv) and (vi) are binary because they can be either supported (i.e., a tick mark) or not supported (i.e., a cross mark). Accordingly, we use Equation 1 to determine the satisfaction degree of binary requirements for a specific composition mechanism,

$$r_b(b) = \frac{b}{5} \quad (1)$$

where r_b is the satisfaction degree of binary requirements and b is the number of supported binary requirements by the mechanism s.t. $b \in (\mathbb{N} \cap [0, 5])$ and $r_b \in (\mathbb{Q} \cap [0, 1])$.

The requirement (v) is quinary since it admits $2^3 - 3$ possible results: *None*, *Control/Data*, *Control/Computation*,

Data/Computation and *Control/Data/Computation*. This is because (v) considers three different concerns (i.e., control, data and computation) and discards options involving only one concern. When (v) is *None*, the requirement support becomes zero. Accordingly, Equation 2 determines the degree of separation of concerns for a specific composition mechanism,

$$r_c(c) = \frac{c}{3} \quad (2)$$

where r_c is the degree of separation of concerns and c is the number of independent concerns supported by the mechanism s.t. $c \in (\mathbb{N} \cap \{0, 2, 3\})$ and $r_c \in (\mathbb{Q} \cap [0, 1])$.

To determine the satisfaction degree of a composition mechanism w.r.t. scalability desiderata, Equations 1 and 2 are used in Equation 3. The result is a percentage that takes into account five binary requirements and one quinary requirement. A higher percentage means a higher satisfaction.

$$s(r_b, r_c) = (r_b \times \frac{5}{6} + r_c \times \frac{1}{6}) \times 100 \quad (3)$$

where s is the overall satisfaction degree w.r.t. all scalability requirements s.t. $s \in [0, 100]$.

Table 11 summarises our analysis of scalability desiderata and Table 12 shows the respective satisfaction degrees. Our interpretation of the results is presented below.

Centralised dataflows is the worst mechanism because it supports only one binary requirement (i.e., location transparency) and separates two concerns (i.e., data and computation). This means that the satisfaction degree of binary requirements is 0.20, while the satisfaction degree of separation of concerns is 0.66. Thus, the overall satisfaction degree

	b	c	r_b	r_c	s
Centralised Dataflows	1	2	0.20	$0.\overline{66}$	27.78%
Distributed Dataflows	2	2	0.40	$0.\overline{66}$	44.44%
Centralised Orchestration	2	2	0.40	$0.\overline{66}$	44.44%
Distributed Orchestration	3	2	0.60	$0.\overline{66}$	61.11%
Choreography	3	0	0.60	0.00	50.00%
DX-MAN	5	3	1.00	1.00	100.00%

Table 12

Satisfaction degree of IoT composition mechanisms w.r.t. scalability desiderata.

of centralised dataflows is 27.78%. Distributed dataflows provide distributed workflows in addition. Consequently, it possesses a (higher) satisfaction degree of binary requirements equal to 0.40 and, therefore, a (higher) overall satisfaction degree of 44.44%.

Although centralised orchestration has the same satisfaction degree as distributed dataflows, it supports different binary requirements (i.e., explicit control flow and location transparency) and separates different concerns (i.e., control and computation). Distributed orchestration is similar, yet different. It offers distributed workflows in addition for a (higher) satisfaction of binary requirements of 0.60 and, therefore, a (higher) overall satisfaction degree of 61.11%.

Choreography covers three binary requirements (i.e., explicit control flow, distributed workflows and decentralised data flows) and does not provide any separation of concerns since control and data are mixed in service computation. This means that the satisfaction degree of binary requirements is 0.60, with a null separation of concerns. Overall, choreography fulfils scalability requirements to a degree of 50%.

DX-MAN is the only mechanism that fulfils all binary requirements and provides the separation of control, data and computation. It is also the only one that supports workflow variability because of total compositionality (see Table 10). Accordingly, the satisfaction degrees of binary requirements and the separation of concerns are both 1. Thus, DX-MAN best fulfils the desiderata with a satisfaction degree of 100%.

8. Discussion

This section discusses the results presented in Section 7 and covers additional issues concerning compositionality, scalability requirements and the relationship between them.

Our results show that explicit control flow is supported by the majority of the mechanisms. In particular, orchestration defines control flow in orchestrators, choreography in a protocol and DX-MAN in composition operators. Dataflows is the only one that does not support such a requirement.

Workflow distribution is also met by the majority of the mechanisms, except centralised dataflows and centralised orchestration where a coordinator fully governs a workflow.

Almost all composition mechanisms use a coordinator to exogenously define workflow(s) and, therefore, ensuring the

separation of at least two concerns. In particular, orchestration and DX-MAN separate control and computation, while dataflows orthogonalises data and computation. Such separation enables location transparency as only coordinators are aware of atomic service locations. Choreography is the only mechanism without any support for location transparency, since control and data are mixed with computation.

There is a special choreography implementation based on the data-driven paradigm in which some computation is triggered once input data becomes available [199]. As a protocol must explicitly define workflow control flow, the analysis presented in Section 6 is also applicable.

Generally speaking, avoiding coordinators allows choreography styles to support decentralised data flows. So, there is a trade-off between coordination and decentralised data flows. DX-MAN obviates this problem by separating data in addition to control and computation. This separation enables the realisation of decentralised data flows, without considering neither control nor computation. So, data never passes through composition operators (i.e., coordinators) [162].

Some orchestration approaches [34, 146, 195, 198] partially separate control from data algorithmically. To achieve this, coordinators pass data references alongside control, rather than exchanging data values. However, analysing data dependencies to extract references is a challenging task because data and control are still semantically entangled.

For orchestration and dataflows, a coordinator is commonly referred to as a composite service. However, it is just a composition of specific operations (i.e., a workflow) that must be transformed into a service, rather than a composition of entire services (with all their provided operations). Only DX-MAN achieves an actual composition of services (not operations) as it defines a composite service without any transformation step while preserving all service operations from which multiple workflow variants can be derived. Thus, algebraic composition equates to total compositionality which, in turn, implies workflow variability. A DX-MAN composite service is in fact semantically equivalent to a (potentially infinite) set of orchestrations. This is because each element in a composite workflow space is a different composition workflow as regarded by existing mechanisms. For any existing composition mechanism to be equivalent to a composite workflow space, it would require to include all possible combinations of execution paths in the workflow definition, leading to a combinatorial explosion problem. While it is true a configurable workflow [208] can deal with this issue, it does not define multiple workflows at a time but just a single workflow that can be manually configured.

In this paper, we analyse the fundamental semantics of current service composition mechanisms that allow the definition of IoT workflows. However, other mechanisms must be considered even if they do not allow the explicit definition of behaviour, e.g., port connections [248] and ensemble-based composition [249]. Moreover, we did not analyse service interactions where there is no definition of composite services such as direct message passing [250], broker-based interactions [112] and event-driven interactions [251]. For this, we

refer the reader to another paper [74].

Different composition mechanisms can coexist in the same system. For example, it is possible to build a system where some services are composed using orchestration while others are composed by choreography [155, 252]. Our intention is not to analyse the combination of different composition mechanisms, but to analyse them individually. Clearly, our analysis results apply to such a combination.

9. Conclusions

Functional scalability becomes a crucial concern for IoT systems as the number of available services increases. For that reason, we need to look back at the foundations of service composition in order to tackle this challenging problem. In this paper, we analysed the semantics of current IoT service composition mechanisms, using an evaluation framework that considers six functional scalability requirements: (i) explicit control flow; (ii) distributed workflows; (iii) location transparency; (iv) decentralised data flows; (v) separation of control, data and computation; and (vi) workflow variability. Our results suggest that DX-MAN is the composition mechanism that best fulfils the functional scalability requirements of IoT systems. This is not surprising since such a model was designed with scalability desiderata in mind.

It is important to note that, except DX-MAN, there are no new composition mechanisms developed in the last decade. This crisis is a worrying situation that must be prioritised and remedied because of the inherent scale that IoT systems pose. Only DX-MAN was designed with this in mind, but we expect further developments on this in the coming years.

We do not claim that the scalability requirements we analyse here are complete, as other characteristics must also be considered. Nevertheless, we believe our evaluation framework has included the most critical ones and provides a useful starting point for further extensions that consider other functional scalability desiderata such as dynamic reconfiguration, the number of messages exchanged and the support for implicit/explicit data flows. The framework can thus be refined by other researchers when conducting further studies on IoT service composition mechanisms.

Appendix A

A larger view of *SPark* is shown in Figure A.1, where there are new composite services representing *Private Parkings* of a smart city. Also, the *Display* service has been converted into a composite that uses the atomic services *WebMap* and *Visual*. There is also a new composite service called *PathPlanning* that composes multiple *TrafficMonitor* composites and multiple atomic *People* services. In particular, a *TrafficMonitor* composite has multiple atomic services that represent congestion sensors and cameras (distributed across a city). Furthermore, the *Vehicle* composite has been endowed with multiple atomic *ChargSt* services (for charging the battery of a self-driving vehicle). Finally, the *DrivingCtr* service has been converted into a composite service with the atomic services *Brake* and *Steering*.

It is important to note that the larger view shown in Figure A.1 can be expanded into an even larger composition which may become infeasible to depict in a single document.

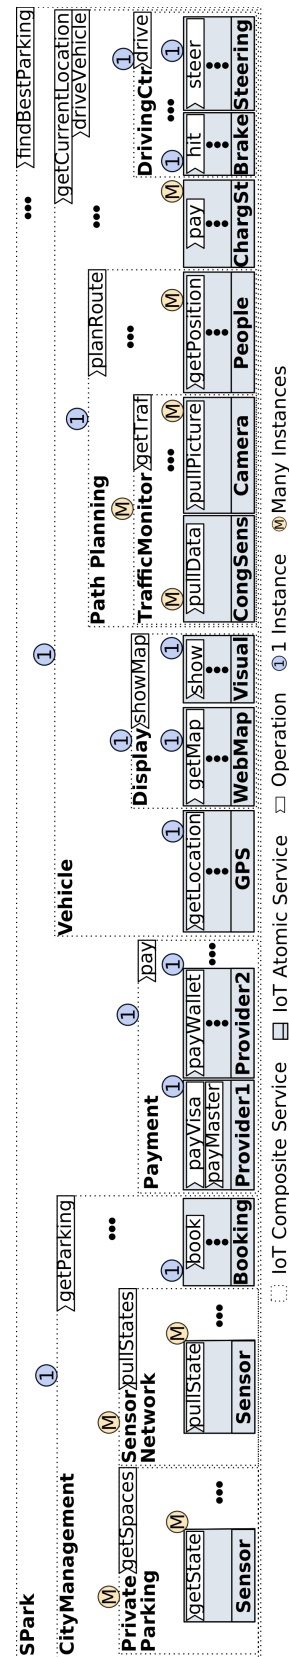


Figure A.1: A larger view of *SPark* services.

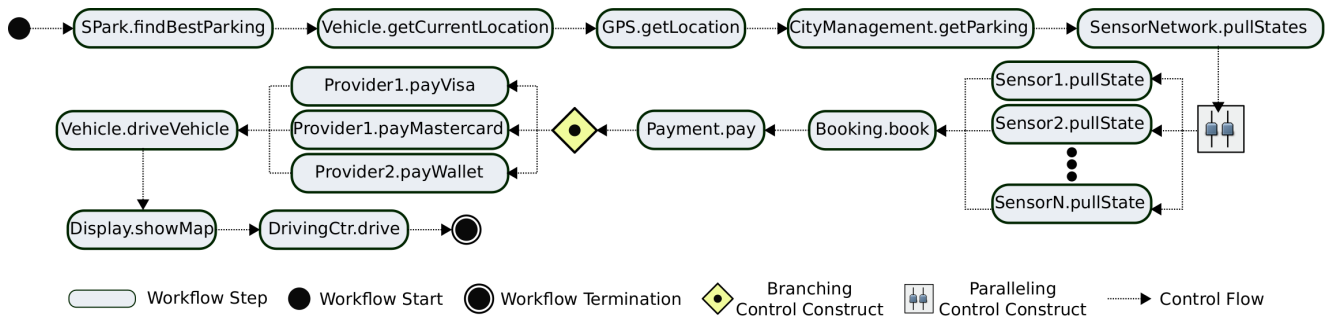


Figure B.1: *SPark* workflow control flow.

Appendix B

Figure B.1 illustrates the complete workflow of *SPark* which starts with the on-demand execution of *SPark.findBestParking*. Next, the *SPark* composite gets the current vehicle's location by invoking the operation *Vehicle.getCurrentLocation* which, in turn, invokes *GPS.getLocation*. Then, it invokes the *CityManagement.getParking* operation which internally executes the *SensorNetwork.pullStates* operation from the nearest InfoStation. In particular, *SensorNetwork.pullStates* pulls the states (in parallel) from the sensors close to the vehicle's location, using the respective *Sensor.pullState* operations. Next, the *CityManagement.getParking* operation uses the sensor states to determine the best (i.e., the free and nearest) parking space and then reserves that parking space using *Booking.book*. Then, control is passed to the *Payment.pay* operation which decides to invoke *Provider1.payVisa*, *Provider1.payMastercard* or *Provider2.payWallet*. Finally, the *Vehicle.driveVehicle* is invoked for implicitly calling *Display.showMap* and *DrivingCtr.drive*, in that order.

References

- [1] IoT Analytics. State of the IoT 2018: Number of IoT devices now at 7b – Market accelerating, 2018. URL <https://iot-analytics.com/state-of-the-iot-update-q1-q2-2018-number-of-iot-devices-now-7b/>.
- [2] Chayan Sarkar, Akshay Uttama Nambi S. N, R. Venkatesha Prasad, Abdur Rahim, Ricardo Neisse, and Gianmarco Baldini. DIAT: A Scalable Distributed Architecture for IoT. *IEEE Internet of Things Journal*, 2(3):230–239, 2015.
- [3] Marzieh Hamzei and Nima Jafari Navimipour. Toward Efficient Service Composition Techniques in the Internet of Things. *IEEE Internet of Things Journal*, 5(5):3774–3787, 2018.
- [4] Asrin Vakili and Nima Jafari Navimipour. Comprehensive and systematic review of the service composition mechanisms in the cloud environments. *Journal of Network and Computer Applications*, 81: 24–36, 2017.
- [5] Parvaneh Asghari, Amir Masoud Rahmani, and Hamid Haj Seyyed Javadi. Service composition approaches in IoT: A systematic review. *Journal of Network and Computer Applications*, 120:61–77, 2018.
- [6] Steen Maarten Van and Andrew S. Tanenbaum. *Distributed Systems*. CreateSpace, London, UK, 3rd edition, 2017.
- [7] Xiang Sun and Nirwan Ansari. EdgeIoT: Mobile Edge Computing for the Internet of Things. *IEEE Communications Magazine*, 54(12): 22–29, 2016.
- [8] Ling Li, Shancang Li, and Shanshan Zhao. QoS-Aware Scheduling of Services-Oriented Internet of Things. *IEEE Transactions on Industrial Informatics*, 10(2):1497–1505, 2014.
- [9] Stelios Sotiriadis, Nik Bessis, Cristiana Amza, and Rajkumar Buyya. Elastic Load Balancing for Dynamic Virtual Machine Reconfiguration Based on Vertical and Horizontal Scaling. *IEEE Transactions on Services Computing*, 12(2):319–334, 2019.
- [10] Jesús Alejandro Cárdenes Cabré, Doina Precup, and Ricardo Sanz. Horizontal and Vertical Self-Adaptive Cloud Controller with Reward Optimization for Resource Allocation. In *International Conference on Cloud and Autonomic Computing (ICAC)*, pages 184–185. IEEE, 2017.
- [11] T. Ramalingeswara Rao, Pabitra Mitra, Ravindara Bhatt, and A. Goswami. The big data system, components, tools, and technologies: a survey. *Knowledge and Information Systems*, pages 1–81, 2018. Advance online publication.
- [12] Chii Chang, Satish Narayana Srirama, and Rajkumar Buyya. Internet of Things (IoT) and New Computing Paradigms. In Rajkumar Buyya and Satish Narayana Srirama, editors, *Fog and Edge Computing*, pages 3–23. Wiley Publishing, Hoboken, NJ, USA, 2019.
- [13] Jian Wu, Qianhui Liang, and Elisa Bertino. Improving Scalability of Software Cloud for Composite Web Services. In *International Conference on Cloud Computing (CLOUD)*, pages 143–146. IEEE, 2009.
- [14] Radu Calinescu, Lars Grunske, Marta Kwiatkowska, Raffaella Mirandola, and Giordano Tamburrelli. Dynamic QoS Management and Optimization in Service-Based Systems. *IEEE Transactions on Software Engineering*, 37(3):387–409, 2011.
- [15] Alti Ilari Maarala, Xiang Su, and Jukka Riekkki. Semantic Reasoning for Context-Aware Internet of Things Applications. *IEEE Internet of Things Journal*, 4(2):461–473, 2017.
- [16] Yi Xu and Abdelsalam Helal. Scalable Cloud–Sensor Architecture for the Internet of Things. *IEEE Internet of Things Journal*, 3(3): 285–298, 2016.
- [17] Roberto Girau, Salvatore Martis, and Luigi Atzori. Lysis: A Platform for IoT Distributed Applications Over Socially Connected Objects. *IEEE Internet of Things Journal*, 4(1):40–51, 2017.
- [18] Bin Cheng, Gürkan Solmaz, Flavio Cirillo, Ernő Kovacs, Kazuyuki Terasawa, and Atsushi Kitazawa. FogFlow: Easy Programming of IoT Services Over Cloud and Edges for Smart Cities. *IEEE Internet of Things Journal*, 5(2):696–707, 2018.
- [19] Emanuel Ferreira Coutinho, Flávio Rubens de Carvalho Sousa, Paulo Antonio Leal Rego, Danielo Gonçalves Gomes, and José Neuman de Souza. Elasticity in cloud computing: a survey. *Annals of Telecommunications*, 70(7):289–309, 2015.
- [20] Rajkumar Buyya and Amir Dastjerdi. *Internet of Things: principles and paradigms*. Morgan Kaufmann, Cambridge, MA, USA, 2016.
- [21] David G. Messerschmitt and Clemens Szyperski. *Software Ecosystem: Understanding an Indispensable Technology and Industry*. MIT Press, Cambridge, MA, USA, 1st edition, 2003.
- [22] Roy Want, Bill N. Schilit, and Scott Jenson. Enabling the Internet of Things. *Computer*, 48(1):28–35, 2015.
- [23] Peter Feiler, Richard P. Gabriel, John Goodenough, Rick Linger,

- Tom Longstaff, Rick Kazman, Mark Klein, Linda Northrop, Douglas Schmidt, Kevin Sullivan, and Kurt Wallnau. Ultra-Large-Scale Systems The Software Challenge of the Future. Technical report, Software Engineering Institute, Carnegie Mellon University, 2006.
- [24] Luigi Atzori, Antonio Iera, and Giacomo Morabito. The Internet of Things: A survey. *Computer Networks*, 54(15):2787–2805, 2010.
- [25] Flávia C. Delicato, Paulo F. Pires, and Thais Batista. The Resource Management Challenge in IoT. In Flávia C. Delicato, Paulo F. Pires, and Thais Batista, editors, *Resource Management for Internet of Things*, pages 7–18. Springer, Cham, Switzerland, 2017.
- [26] Damian Roca, Daniel Nemirovsky, Mario Nemirovsky, Rodolfo Milito, and Mateo Valero. Emergent Behaviors in the Internet of Things: The Ultimate Ultra-Large-Scale System. *IEEE Micro*, 36(6):36–44, 2016.
- [27] Hermann Kopetz. *Real-Time Systems: Design Principles for Distributed Embedded Applications*. Springer, New York, NY, USA, 2nd edition, 2011.
- [28] Martin Törngren and Paul T. Grogan. How to Deal with the Complexity of Future Cyber-Physical Systems? *Designs*, 2(4):1–16, 2018.
- [29] Netflix. Conductor, 2019. URL <https://netflix.github.io/conductor/>.
- [30] Martin Fowler. What do you mean by "Event-Driven"?, 2017. URL <https://martinfowler.com/articles/201701-event-driven.html>.
- [31] Ian Sommerville, Dave Cliff, Radu Calinescu, Justin Keen, Tim Kelly, Marta Kwiatkowska, John Mcdermid, and Richard Paige. Large-scale Complex IT Systems. *Communications of the ACM*, 55(7):71–77, 2012.
- [32] Reza Rezaei, Thiam Kian Chiew, and Sai Peck Lee. An interoperability model for ultra large scale systems. *Advances in Engineering Software*, 67:22–46, 2014.
- [33] Valérie Issarny, Nikolaos Georgantas, Sara Hachem, Apostolos Zarras, Panos Vassiliadis, Marco Autili, Marco Aurélio Gerosa, and Amira Ben Hamida. Service-oriented middleware for the Future Internet: state of the art and research directions. *Journal of Internet Services and Applications*, 2(1):23–45, 2011.
- [34] Adam Barker, Jon B. Weissman, and Jano I. Van Hemert. Reducing Data Transfer in Service-Oriented Architectures: The Circulate Approach. *IEEE Transactions on Services Computing*, 5(3):437–449, 2012.
- [35] Michael Hahn, Uwe Breitenbücher, Oliver Kopp, and Frank Leymann. Modeling and execution of data-aware choreographies: an overview. *Computer Science - Research and Development*, 33(3):329–340, 2018.
- [36] Gio Wiederhold, Peter Wegner, and Stefano Ceri. Towards Megaprogramming: A Paradigm for Component-Based Programming. *Communications of the ACM*, 35(11):89–99, 1992.
- [37] Nam Ky Giang, Rodger Lea, and Victor C. M. Leung. Exogenous Coordination for Building Fog-Based Cyber Physical Social Computing and Networking Systems. *IEEE Access*, 6:31740–31749, 2018.
- [38] Nam Ky Giang, Rodger Lea, and Victor C. M. Leung. Developing applications in large scale, dynamic fog computing: A case study. *Software: Practice and Experience*, pages 1–14, 2019. Advance online publication.
- [39] Matthias Galster, Uwe Zdun, Danny Weyns, Rick Rabiser, Bo Zhang, Michael Goedicke, and Gilles Perrouin. Variability and Complexity in Software Design: Towards a Research Agenda. *SIGSOFT Software Engineering Notes*, 41(6):27–30, 2017.
- [40] Danny Weyns, Gowri Sankar Ramachandran, and Ritesh Kumar Singh. Self-managing Internet of Things. In A Min Tjoa, Ladjel Bellatreche, Stefan Biffl, Jan van Leeuwen, and Jiří Wiedermann, editors, *SOFSEM 2018: Theory and Practice of Computer Science*, volume 10706 of *Lecture Notes in Computer Science*, pages 67–84, Cham, Switzerland, 2018. Springer.
- [41] Gerald Holl, Paul Grünbacher, and Rick Rabiser. A systematic review and an expert survey on capabilities supporting multi product lines. *Information and Software Technology*, 54(8):828–852, 2012.
- [42] Kevin Ashton. That 'Internet of Things' Thing. *RFID Journal*, page 1, 2009.
- [43] Mark Weiser. The Computer for the 21st Century. *Scientific American*, 265(3):94–105, 1991.
- [44] ITU-T. Overview of the Internet of things. Technical Report ITU-T Y.4000/Y.2060, International Telecommunication Union, 2012.
- [45] Karl Popper. Three Worlds. The Tanner Lecture on Human Values, 1978.
- [46] Ing-Ray Chen, Jia Guo, and Fenyee Bao. Trust Management for SOA-Based IoT and Its Application to Service Composition. *IEEE Transactions on Services Computing*, 9(3):482–495, 2016.
- [47] Michael P. Papazoglou, Paolo Traverso, Schahram Dustdar, and Frank Leymann. Service-Oriented Computing: State of the Art and Research Challenges. *Computer*, 40(11):38–45, 2007.
- [48] Sherif Abdelwahab, Bechir Hamdaoui, Mohsen Guizani, and Taieb Znati. Cloud of Things for Sensing-as-a-Service: Architecture, Algorithms, and Use Case. *IEEE Internet of Things Journal*, 3(6):1099–1112, 2016.
- [49] Robert Brzoza-Woch and Piotr Nawrocki. FPGA-Based Web Services – Infinite Potential or a Road to Nowhere? *IEEE Internet Computing*, 20(1):44–51, 2016.
- [50] Michele Nitti, Virginia Pilloni, Giuseppe Colistra, and Luigi Atzori. The Virtual Object as a Major Element of the Internet of Things: A Survey. *IEEE Communications Surveys & Tutorials*, 18(2):1228–1240, 2016.
- [51] Li Da Xu, Wu He, and Shancang Li. Internet of Things in Industries: A Survey. *IEEE Transactions on Industrial Informatics*, 10(4):2233–2243, 2014.
- [52] Dominique Guinard, Vlad Trifa, Stamatis Karnouskos, Patrik Spiess, and Domnic Savio. Interacting with the SOA-Based Internet of Things: Discovery, Query, Selection, and On-Demand Provisioning of Web Services. *IEEE Transactions on Services Computing*, 3(3):223–235, 2010.
- [53] Karl D. Gottschalk, Stephen Graham, Heather Kreger, and James Snell. Introduction to Web services architecture. *IBM Systems Journal*, 41(2):170–177, 2002.
- [54] Francisco Curbera, Matthew Duftler, Rania Khalaf, William Nagy, Nirmal Mukhi, and Sanjiva Weerawarana. Unraveling the Web Services Web: An Introduction to SOAP, WSDL, and UDDI. *IEEE Internet Computing*, 6(2):86–93, 2002.
- [55] Nissanka B. Priyantha, Aman Kansal, Michel Goraczko, and Feng Zhao. Tiny Web Services: Design and Implementation of Interoperable and Evolvable Sensor Networks. In *ACM Conference on Embedded Network Sensor Systems (SenSys)*, pages 253–266. ACM, 2008.
- [56] Dominique Guinard, Vlad Trifa, Friedemann Mattern, and Erik Wilde. From the Internet of Things to the Web of Things: Resource-oriented Architecture and Best Practices. In Dieter Uckelmann, Mark Harrison, and Florian Michahelles, editors, *Architecting the Internet of Things*, pages 97–129. Springer, Berlin, Germany, 2011.
- [57] Sylvain Cherrier and Yacine M. Ghamri-Doudane. The "Object-as-a-Service" paradigm. In *Global Information Infrastructure and Networking Symposium (GIIS)*, pages 1–7. IEEE, 2014.
- [58] Douglas K. Barry and David Dick. *Web Services, Service-Oriented Architectures, and Cloud Computing: The Savy Manager's Guide*. Morgan Kaufmann, Burlington, MA, USA, 2nd edition, 2013.
- [59] Kung-Kiu Lau and Simone Di Cola. *An Introduction to Component-based Software Development*. World Scientific, Singapore, 1st edition, 2017.
- [60] Michael Bell. *SOA Modeling Patterns for Service Oriented Discovery and Analysis*. Wiley Publishing, Hoboken, NJ, USA, 1st edition, 2010.
- [61] Chandrashekar Jatoth, G. R. Gangadharan, and Rajkumar Buyya. Computational Intelligence Based QoS-Aware Web Service Composition: A Systematic Literature Review. *IEEE Transactions on Services Computing*, 10(3):475–492, 2017.
- [62] Dominique Guinard. Towards the Web of Things: Web Mashups for Embedded Devices. In *International World Wide Web Conference (WWW)*, pages 1–8. ACM, 2009.
- [63] Martin Bauer, Nicola Bui, Jourik De Loof, Carsten Magerkurth,

- Andreas Nettsträter, Julinda Stefa, and Joachim W. Walewski. IoT Reference Model. In Alessandro Bassi, Martin Bauer, Martin Fiedler, Thorsten Kramp, Rob van Kranenburg, Sebastian Lange, and Stefan Meissner, editors, *Enabling Things to Talk: Designing IoT solutions with the IoT Architectural Reference Model*, pages 113–162. Springer, Berlin, Germany, 2013.
- [64] João Rufino, Muhammad Alam, Joaquim Ferreira, Abdur Rehman, and Kim Fung Tsang. Orchestration of containerized microservices for IIoT using Docker. In *International Conference on Industrial Technology (ICIT)*, pages 1532–1536. IEEE, 2017.
- [65] Cesare Pautasso. RESTful Web service composition with BPEL for REST. *Data & Knowledge Engineering*, 68(9):851–866, 2009.
- [66] Roy Fielding. *Architectural Styles and the Design of Network-based Software Architectures*. PhD Thesis, University of California, Irvine, 2000.
- [67] Cesare Pautasso, Olaf Zimmermann, and Frank Leymann. Restful Web Services vs. "Big" Web Services: Making the Right Architectural Decision. In *International Conference on World Wide Web (WWW)*, pages 805–814. ACM, 2008.
- [68] Dominique Guinard, Iulia Ion, and Simon Mayer. In Search of an Internet of Things Service Architecture: REST or WS-*? A Developers' Perspective. In Alessandro Puiatti and Tao Gu, editors, *Mobile and Ubiquitous Systems: Computing, Networking, and Services*, volume 104 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pages 326–337. Berlin, Germany, 2012. Springer.
- [69] Choonhwa Lee, Chengyang Wang, Eunsam Kim, and Sumi Helal. Blueprint Flow: A Declarative Service Composition Framework for Cloud Applications. *IEEE Access*, 5:17634–17643, 2017.
- [70] Rebeca P. Diaz Redondo, Ana Fernandez Vilas, Manuel Ramos Cabrer, Jose J. Pazos Arias, and Marta Rey Lopez. Enhancing Residential Gateways: OSGi Service Composition. *IEEE Transactions on Consumer Electronics*, 53(1):87–95, 2007.
- [71] Hye-Young Paik, Angel Lagares Lemos, Moshe Chai Barukh, Boualem Benatallah, and Aarthi Natarajan. *Web Service Implementation and Composition Techniques*. Springer, Cham, Switzerland, 1st edition, 2017.
- [72] Angel Lagares Lemos, Florian Daniel, and Boualem Benatallah. Web Service Composition: A Survey of Techniques and Tools. *ACM Computing Surveys*, 48(3):1–41, 2015.
- [73] Remco Dijkman and Marlon Dumas. Service-oriented design: a multi-viewpoint approach. *International Journal of Cooperative Information Systems*, 13(4):337–368, 2004.
- [74] Damian Arellanes and Kung-Kiu Lau. Analysis and Classification of Service Interactions for the Scalability of the Internet of Things. In *International Congress on Internet of Things (ICIOT)*, pages 80–87. IEEE, 2018.
- [75] Matthew Shields. Control- Versus Data-Driven Workflows. In Ian J. Taylor, Ewa Deelman, Dennis B. Gannon, and Matthew Shields, editors, *Workflows for e-Science: Scientific Workflows for Grids*, pages 167–173. Springer, London, UK, 2007.
- [76] Amazon. Amazon Simple Workflow Service (SWF), 2019. URL <https://aws.amazon.com/swf/>.
- [77] Bertram Ludäscher, Ilkay Altintas, Chad Berkley, Dan Higgins, Efrat Jaeger, Matthew Jones, Edward A. Lee, Jing Tao, and Yang Zhao. Scientific workflow management and the Kepler system. *Concurrency and Computation: Practice and Experience*, 18(10):1039–1065, 2006.
- [78] Carl Adam Petri. *Kommunikation mit Automaten*. PhD Thesis, Institut für Instrumentelle Mathematik, 1962.
- [79] Yong Zhao, Michael Wilde, and Ian Foster. Virtual Data Language: A Typed Workflow Notation for Diversely Structured Scientific Data. In Ian J. Taylor, Ewa Deelman, Dennis B. Gannon, and Matthew Shields, editors, *Workflows for e-Science: Scientific Workflows for Grids*, pages 258–275. Springer, London, UK, 2007.
- [80] Tom Goodale, Gabrielle Allen, Gerd Lanfermann, Joan Massó, Thomas Radke, Edward Seidel, and John Shalf. The Cactus Framework and Toolkit: Design and Applications. In José M. L. M. Palma, A. Augusto Sousa, Jack Dongarra, and Vicente Hernández, editors, *High Performance Computing for Computational Science — VEC- PAR 2002*, volume 2565 of *Lecture Notes in Computer Science*, pages 197–227. Berlin, Germany, 2003. Springer.
- [81] OASIS. Web Services Business Process Execution Language Version 2.0, 2007. URL <http://docs.oasis-open.org/wsbpel/2.0/OS/wsbpel-v2.0-05.html>.
- [82] Ian Taylor, Matthew Shields, Ian Wang, and Andrew Harrison. The Triana Workflow Environment: Architecture and Applications. In Ian J. Taylor, Ewa Deelman, Dennis B. Gannon, and Matthew Shields, editors, *Workflows for e-Science: Scientific Workflows for Grids*, pages 320–339. Springer, London, UK, 2007.
- [83] OpenJS Foundation. Node-RED: Flow-based programming for the Internet of Things, 2019. URL <https://nodered.org/>.
- [84] Li Da Xu and Wattana Viriyasitavat. A Novel Architecture for Requirement-Oriented Participation Decision in Service Workflows. *IEEE Transactions on Industrial Informatics*, 10(2):1478–1485, 2014.
- [85] Nathaniel Palmer, Swenson Keith, Reddy Surendra, Fingar Peter, Setrag Khoshafian, and Larry Hawes. *BPM Everywhere: Internet of Things, Process of Everything*. Future Strategies Inc., Lighthouse Point, FL, USA, 1st edition, 2015.
- [86] Jakob Mass, Chii Chang, and Satish N. Srirama. Workflow Model Distribution or Code Distribution? Ideal Approach for Service Composition of the Internet of Things. In *International Conference on Services Computing (SCC)*, pages 649–656. IEEE, 2016.
- [87] Ronny Seiger, Christine Keller, Florian Niebling, and Thomas Schlegel. Modelling complex and flexible processes for smart cyber-physical environments. *Journal of Computational Science*, 10:137–148, 2015.
- [88] Ronny Seiger, Steffen Huber, Peter Heisig, and Uwe Aßmann. Toward a framework for self-adaptive workflows in cyber-physical systems. *Software & Systems Modeling*, 18(2):1117–1134, 2019.
- [89] Shangguang Wang, Ao Zhou, Mingzhe Yang, Lei Sun, Ching-Hsien Hsu, and Fangchun Yang. Service Composition in Cyber-Physical-Social Systems. *IEEE Transactions on Emerging Topics in Computing*, pages 1–11, 2019. Advance online publication.
- [90] Alexander Jungmann and Felix Mohr. An approach towards adaptive service composition in markets of composed services. *Journal of Internet Services and Applications*, 6(1):1–18, 2015.
- [91] Mohamed Essaid Khanouche, Yacine Amirat, Abdelghani Chibani, Moussa Kerkar, and Ali Yachir. Energy-Centered and QoS-Aware Services Selection for Internet of Things. *IEEE Transactions on Automation Science and Engineering*, 13(3):1256–1269, 2016.
- [92] Shuiguang Deng, Zhengzhe Xiang, Jianwei Yin, Javid Taheri, and Albert Y. Zomaya. Composition-Driven IoT Service Provisioning in Distributed Edges. *IEEE Access*, 6:54258–54269, 2018.
- [93] Rong Yang, Bing Li, and Can Cheng. A Petri Net-Based Approach to Service Composition and Monitoring in the IOT. In *Asia-Pacific Services Computing Conference (APSCC)*, pages 16–22. IEEE, 2014.
- [94] Mengyu Sun, Zhensheng Shi, Shengjun Chen, Zhangbing Zhou, and Yucong Duan. Energy-Efficient Composition of Configurable Internet of Things Services. *IEEE Access*, 5:25609–25622, 2017.
- [95] Osama Alsaryrah, Ibrahim Mashal, and Tein-Yaw Chung. Bi-Objective Optimization for Energy Aware Internet of Things Service Composition. *IEEE Access*, 6:26809–26819, 2018.
- [96] Lei Huo and Zhiliang Wang. Service composition instantiation based on cross-modified artificial Bee Colony algorithm. *China Communications*, 13(10):233–244, 2016.
- [97] Zhangbing Zhou, Deng Zhao, Lu Liu, and Patrick C. K. Hung. Energy-aware composition for wireless sensor networks as a service. *Future Generation Computer Systems*, 80:299–310, 2018.
- [98] Qian Li, Runliang Dou, Fuzan Chen, and Guofang Nan. A QoS-oriented Web service composition approach based on multi-population genetic algorithm for Internet of things. *International Journal of Computational Intelligence Systems*, 7(sup2):26–34, 2014.
- [99] Safina Showkat Ara, Zia Ush Shamszaman, and Ilyoung Chong. Web-of-Objects Based User-Centric Semantic Service Composition

- Methodology in the Internet of Things. *International Journal of Distributed Sensor Networks*, 10(5):1–11, 2014.
- [100] Mahmoud M. Badawy, Zainab H. Ali, and Hesham A. Ali. QoS provisioning framework for service-oriented internet of things (IoT). *Cluster Computing*, pages 1–17, 2019. Advance online publication.
- [101] Samir Berrani, Ali Yachir, Badis Djemaa, and Mohamed Aissani. Extended multi-agent system based service composition in the Internet of things. In *International Conference on Pattern Analysis and Intelligent Systems (PAIS)*, pages 1–8. IEEE, 2018.
- [102] A. Urbieto, A. González-Beltrán, S. Ben Mokhtar, M. Anwar Hossain, and L. Capra. Adaptive and context-aware service composition for IoT-based smart cities. *Future Generation Computer Systems*, 76: 262–274, 2017.
- [103] Tomasz Szydio, Robert Brzoza-Woch, Joanna Sendorek, Mateusz Windak, and Chris Gniady. Flow-Based Programming for IoT Leveraging Fog Computing. In *International Conference on Enabling Technologies: Infrastructure for Collaborative Enterprises (WET-ICE)*, pages 74–79. IEEE, 2017.
- [104] Jonathan Lee, Shin-Jie Lee, and Ping-Feng Wang. A Framework for Composing SOAP, Non-SOAP and Non-Web Services. *IEEE Transactions on Services Computing*, 8(2):240–250, 2015.
- [105] Farzad Khodadadi, Amir Vahid Dastjerdi, and Rajkumar Buyya. Simurgh: A framework for effective discovery, programming, and integration of services exposed in IoT. In *International Conference on Recent Advances in Internet of Things (RIoT)*, pages 1–6. IEEE, 2015.
- [106] Panagiotis Vlacheas, Raffaele Giaffreda, Vera Stavroulaki, Dimitris Kelaidonis, Vassilis Foteinos, George Poullos, Panagiotis Demestichas, Andrey Somov, Abdur Rahim Biswas, and Klaus Moessner. Enabling smart cities through a cognitive management framework for the internet of things. *IEEE Communications Magazine*, 51(6): 102–111, 2013.
- [107] Hongming Cai, Yizhi Gu, Athanasios V. Vasilakos, Boyi Xu, and Jianzhong Zhou. Model-Driven Development Patterns for Mobile Services in Cloud of Things. *IEEE Transactions on Cloud Computing*, 6(3):771–784, 2018.
- [108] In-Young Ko, Han-Gyu Ko, Angel Jimenez Molina, and Jung-Hyun Kwon. SoIoT: Toward A User-Centric IoT-Based Service Framework. *ACM Transactions on Internet Technology*, 16(2):1–21, 2016.
- [109] Alfred Åkesson, Görel Hedin, Mattias Nordahl, and Boris Magnusson. ComPOS: Composing Oblivious Services. In *International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*, pages 132–138. IEEE, 2019.
- [110] Fabrizio Montesi, Claudio Guidi, and Gianluigi Zavattaro. Service-Oriented Programming with Jolie. In Athman Bouguettaya, Quan Z. Sheng, and Florian Daniel, editors, *Web Services Foundations*, pages 81–107. Springer, New York, NY, USA, 2014.
- [111] Le Kim-Hung, Soumya Kanti Datta, Christian Bonnet, François Hamon, and Alexandre Boudonne. A scalable IoT framework to design logical data flow using virtual sensor. In *International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, pages 1–7. IEEE, 2017.
- [112] Anne H. Ngu, Mario Gutierrez, Vangelis Metsis, Surya Nepal, and Quan Z. Sheng. IoT Middleware: A Survey on Issues and Enabling Technologies. *IEEE Internet of Things Journal*, 4(1):1–20, 2017.
- [113] Federico Montori, Luca Bedogni, and Luciano Bononi. A Collaborative Internet of Things Architecture for Smart Cities and Environmental Monitoring. *IEEE Internet of Things Journal*, 5(2):592–605, 2018.
- [114] Antonio Marcos Alberti, Gabriel Dias Scarpioni, Vaner Magalhães, Arismar Cerqueira Sodre, Joel Rodrigues, and Rodrigo da Rosa Righi. Advancing NovaGenesis Architecture Towards Future Internet of Things. *IEEE Internet of Things Journal*, 6(1):215–229, 2019.
- [115] Sandra Rodríguez-Valenzuela, Juan A. Holgado-Terriza, José M. Gutiérrez-Guerrero, and Jesús L. Muros-Cobos. Distributed service-based approach for sensor data fusion in IoT environments. *Sensors*, 14(10):19200–19228, 2014.
- [116] Robert Kleinfeld, Stephan Steglich, Lukasz Radziwonowicz, and Charalampos Doukas. Glue.Things: A Mashup Platform for Wiring the Internet of Things with the Internet of Services. In *International Workshop on Web of Things (WoT)*, pages 16–21. ACM, 2014.
- [117] Sylvain Cherrier, Ismail Salhi, Yacine M. Ghamri-Doudane, Stéphane Lohier, and Philippe Valembois. BeC 3: Behaviour Crowd Centric Composition for IoT applications. *Mobile Networks and Applications*, 19(1):18–32, 2014.
- [118] Alexis Huf and Frank Siqueira. Composition of heterogeneous web services: A systematic review. *Journal of Network and Computer Applications*, 143:89–110, 2019.
- [119] Parvaneh Asghari, Amir Masoud Rahmani, and Hamid Haj Seyyed Javadi. Internet of Things applications: A systematic review. *Computer Networks*, 148:241–261, 2019.
- [120] Xabier Larrucea, Izaskun Santamaria, Ricardo Colomo-Palacios, and Christof Ebert. Microservices. *IEEE Software*, 35(3):96–100, 2018.
- [121] Cloves Carneiro and Tim Schmelmer. Microservices: The What and the Why. In Cloves Carneiro and Tim Schmelmer, editors, *Microservices From Day One*, pages 3–18. Apress, Berkeley, CA, USA, 2016.
- [122] Sam Newman. *Building Microservices*. O’Reilly Media, Sebastopol, CA, USA, 1st edition, 2015.
- [123] Ion-Dorinel Filip, Florin Pop, Cristina Serbanescu, and Chang Choi. Microservices Scheduling Model Over Heterogeneous Cloud-Edge Environments As Support for IoT Applications. *IEEE Internet of Things Journal*, 5(4):2672–2681, 2018.
- [124] Ahmed E. Khaled, Abdelsalam Helal, Wyatt Lindquist, and Choonhwa Lee. IoT-DDL—Device Description Language for the “T” in IoT. *IEEE Access*, 6:24048–24063, 2018.
- [125] Lorenzo Carnevale, Antonio Celesti, Antonino Galletta, Schahram Dustdar, and Massimo Villari. From the Cloud to Edge and IoT: a Smart Orchestration Architecture for Enabling Osmotic Computing. In *International Conference on Advanced Information Networking and Applications Workshops (WAINA)*, pages 419–424. IEEE, 2018.
- [126] Yuansong Qiao, Robert Nolani, Saul Gill, Guiming Fang, and Brian Lee. ThingNet: A micro-service based IoT macro-programming platform over edges and cloud. In *Conference on Innovation in Clouds, Internet and Networks and Workshops (ICIN)*, pages 1–4. IEEE, 2018.
- [127] Kleantes Thramboulidis, Danai C. Vachtsevanou, and Alexandros Solanos. Cyber-physical microservices: An IoT-based framework for manufacturing systems. In *Industrial Cyber-Physical Systems (ICPS)*, pages 232–239. IEEE, 2018.
- [128] Olaf Zimmermann. Microservices tenets. *Computer Science - Research and Development*, 32(3):301–310, 2017.
- [129] J. Paul Morrison. *Flow-Based Programming: A New Approach to Application Development*. CreateSpace, Paramount, CA, USA, 2nd edition, 2010.
- [130] Wesley M. Johnston, J. R. Paul Hanna, and Richard J. Millar. Advances in Dataflow Programming Languages. *ACM Computing Surveys*, 36(1):1–34, 2004.
- [131] J. Paul Morrison. Data Stream Linkage Mechanism. *IBM Systems Journal*, 17(4):383–408, 1978.
- [132] Gustavo Alonso, Fabio Casati, Harumi Kuno, and Vijay Machiraju. Web Services. In Gustavo Alonso, Fabio Casati, Harumi Kuno, and Vijay Machiraju, editors, *Web Services: Concepts, Architectures and Applications*, pages 123–149. Springer, Berlin, Germany, 2004.
- [133] Nam Ky Giang, Michael Blackstock, Rodger Lea, and Victor C. M. Leung. Developing IoT applications in the Fog: A Distributed Dataflow approach. In *International Conference on the Internet of Things (IOT)*, pages 155–162. IEEE, 2015.
- [134] Joseph Noor, Hsiao-Yun Tseng, Luis Garcia, and Mani Srivastava. DDFlow: Visualized Declarative Programming for Heterogeneous IoT Networks. In *International Conference on Internet of Things Design and Implementation (IoTDI)*, pages 172–177. ACM, 2019.
- [135] National Instruments. What Is LabView?. 2019. URL <http://www.ni.com/en-gb/shop/labview.html>.
- [136] Chris Peltz. Web Services Orchestration and Choreography. *Computer*, 36(10):46–52, 2003.

- [137] Alistair Barros, Marlon Dumas, and Phillipa Oaks. Standards for Web Service Choreography and Orchestration: Status and Perspectives. In Christop J. Bussler and Armin Haller, editors, *Business Process Management Workshops*, volume 3812 of *Lecture Notes in Computer Science*, pages 61–74, Berlin, Germany, 2005. Springer.
- [138] Quan Sheng, Xiaoqiang Qiao, Athanasios Vasilakos, Claudia Szabo, Scott Bourne, and Xiaofei Xu. Web services composition: A decade's overview. *Information Sciences*, 280:218–238, 2014.
- [139] Pieter Hens, Monique Snoeck, Geert Poels, and Manu De Backer. Process fragmentation, distribution and execution using an event-based interaction scheme. *Journal of Systems and Software*, 89:170–192, 2014.
- [140] Ward Jaradat, Alan Dearle, and Adam Barker. Towards an autonomous decentralized orchestration system. *Concurrency and Computation: Practice and Experience*, 28(11):3164–3179, 2016.
- [141] Mangala Gowri Nanda, Satish Chandra, and Vivek Sarkar. Decentralizing Execution of Composite Web Services. In *ACM SIGPLAN Conference on Object-oriented Programming, Systems, Languages, and Applications (OOPSLA)*, pages 170–187. ACM, 2004.
- [142] Girish Chafle, Sunil Chandra, Vijay Mann, and Mangala Gowri Nanda. Decentralized orchestration of composite web services. In *International World Wide Web conference (WWW)*, pages 134–143. ACM, 2004.
- [143] Michael Pantazoglou, Ioannis Pogkas, and Aphrodite Tsalgaidou. Decentralized Enactment of BPEL Processes. *IEEE Transactions on Services Computing*, 7(2):184–197, 2014.
- [144] Daniel Wutke, Daniel Martin, and Frank Leymann. Model and infrastructure for decentralized workflow enactment. In *ACM Symposium on Applied Computing (SAC)*, pages 90–94. ACM, 2008.
- [145] Daniel Martin, Daniel Wutke, and Frank Leymann. Tuplespace middleware for Petri net-based workflow execution. *International Journal of Web and Grid Services*, 6(1):35–57, 2010.
- [146] Mirko Sonntag, Katharina Gorchach, Dimka Karastoyanova, Frank Leymann, and Michael Reiter. Process space-based scientific workflow enactment. *International Journal of Business Process Integration and Management*, 5(1):32–44, 2010.
- [147] Xitong Kang, Xudong Liu, Hailong Sun, Yanjiu Huang, and Chao Zhou. Improving Performance for Decentralized Execution of Composite Web Services. In *World Congress on Services (SERVICES)*, pages 582–589. IEEE, 2010.
- [148] Walid Fdhila, Marlon Dumas, Claude Godart, and Luciano García-Bañuelos. Heuristics for composite Web service decentralization. *Software & Systems Modeling*, 13(2):599–619, 2014.
- [149] Ronny Seiger, Steffen Huber, and Thomas Schlegel. PROtEUS: An Integrated System for Process Execution in Cyber-Physical Systems. In Khaled Gaaloul, Rainer Schmidt, Selmin Nurcan, Sérgio Guerreiro, and Qin Ma, editors, *Enterprise, Business-Process and Information Systems Modeling*, volume 214 of *Lecture Notes in Business Information Processing*, pages 265–280, Cham, Switzerland, 2015. Springer.
- [150] Yongyang Cheng, Shuai Zhao, Bo Cheng, and Junliang Chen. A Service-Based Fog Execution Environment for the IoT-Aware Business Process Applications. In *International Conference on Web Services (ICWS)*, pages 323–326. IEEE, 2018.
- [151] Amazon. AWS Step Functions, 2019. URL <http://aws.amazon.com/step-functions/>.
- [152] Adam Barker, Christopher D. Walton, and David Robertson. Choreographing Web Services. *IEEE Transactions on Services Computing*, 2(2):152–166, 2009.
- [153] Gero Decker, Oliver Kopp, and Alistair Barros. An Introduction to Service Choreographies. *Information Technology*, 52(2):122–127, 2008.
- [154] OW2 Consortium. CHOReVOLUTION. URL <http://www.chorevolution.eu>.
- [155] S. Cherrier and R. Langar. Services organisation in IoT: mixing Orchestration and Choreography. In *Global Information Infrastructure and Networking Symposium (GIIS)*, pages 1–4. IEEE, 2018.
- [156] ActnConnect. The choreographic platform Actorsphere, 2019. URL http://actnconnect.de/actorsphere_en.
- [157] Damian Arellanes and Kung-Kiu Lau. Workflow Variability for Autonomic IoT Systems. In *International Conference on Autonomic Computing (ICAC)*. IEEE, 2019.
- [158] Damian Arellanes and Kung-Kiu Lau. Algebraic Service Composition for User-Centric IoT Applications. In Dimitrios Georgakopoulos and Liang-Jie Zhang, editors, *Internet of Things – ICIOT 2018*, volume 10972 of *Lecture Notes in Computer Science*, pages 56–69, Cham, Switzerland, 2018. Springer.
- [159] Damian Arellanes and Kung-Kiu Lau. Exogenous Connectors for Hierarchical Service Composition. In *International Conference on Service-Oriented Computing and Applications (SOCA)*, pages 125–132. IEEE, 2017.
- [160] Kung-Kiu Lau, Perla Velasco Elizondo, and Zheng Wang. Exogenous Connectors for Software Components. In George T. Heineman, Ivica Crnkovic, Heinz W. Schmidt, Judith A. Stafford, Clemens Szyperski, and Kurt Wallnau, editors, *Component-Based Software Engineering*, volume 3489 of *Lecture Notes in Computer Science*, pages 90–106, Berlin, Germany, 2005. Springer.
- [161] Perla Velasco Elizondo and Kung-Kiu Lau. A catalogue of component connectors to support development with reuse. *Journal of Systems and Software*, 83(7):1165–1178, 2010.
- [162] Damian Arellanes and Kung-Kiu Lau. Decentralized Data Flows in Algebraic Service Compositions for the Scalability of IoT Systems. In *World Forum on Internet of Things (WF-IoT)*, pages 668–673. IEEE, 2019.
- [163] Damian Arellanes and Kung-Kiu Lau. D-XMAN: A Platform For Total Compositionality in Service-Oriented Architectures. In *International Symposium on Cloud and Service Computing (SC2)*, pages 283–286. IEEE, 2017.
- [164] T. Lin, H. Rivano, and F. Le Mouél. A Survey of Smart Parking Solutions. *IEEE Transactions on Intelligent Transportation Systems*, 18(12):3229–3253, 2017.
- [165] Tullio Giuffrè, Sabato Marco Siniscalchi, and Giovanni Tesoriere. A Novel Architecture of Parking Management for Smart Cities. *Procedia - Social and Behavioral Sciences*, 53:16–28, 2012.
- [166] Jan Höller, Vlasios Tsiatsis, Catherine Mulligan, Stamatios Karnouskos, Stefan Avesand, and David Boyle. *From Machine-To-Machine to the Internet of Things: Introduction to a New Age of Intelligence*. Academic Press, Oxford, UK, 1st edition, 2014.
- [167] Tariq Samad and Anuradha Annaswamy. The Impact of Control Technology: Overview, Success Stories, and Research Challenges. Technical report, IEEE Control Systems Society, 2011.
- [168] Jasmine Sekhon and Cody Fleming. Towards improved testing for deep learning. In *International Conference on Software Engineering (ICSE)*. IEEE, 2019.
- [169] P. Barros, R. Just, S. Millstein, P. Vines, W. Dietl, M. dAmorim, and M. D. Ernst. Static Analysis of Implicit Control Flow: Resolving Java Reflection and Android Intents (T). In *International Conference on Automated Software Engineering (ASE)*, pages 669–679. IEEE, 2015.
- [170] Animesh Nandi, Atri Mandal, Shubham Atreja, Gargi B. Dasgupta, and Subhrajit Bhattacharya. Anomaly Detection Using Program Control Flow Graph Mining From Execution Logs. In *International ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD)*, pages 215–224. ACM, 2016.
- [171] Fábio Bezerra, Jacques Wainer, and W. M. P. van der Aalst. Anomaly Detection Using Process Mining. In Terry Halpin, John Krogstie, Selmin Nurcan, Erik Proper, Rainer Schmidt, Pnina Soffer, and Roland Ukör, editors, *Enterprise, Business-Process and Information Systems Modeling*, volume 29 of *Lecture Notes in Business Information Processing*, pages 149–161, Berlin, Germany, 2009. Springer.
- [172] Tigist Abera, N. Asokan, Lucas Davi, Jan-Erik Ekberg, Thomas Nyman, Andrew Paverd, Ahmad-Reza Sadeghi, and Gene Tsudik. C-FLAT: Control-Flow Attestation for Embedded Systems Software. In *ACM SIGSAC Conference on Computer and Communications Security*, pages 743–754. ACM, 2016.
- [173] Christian Collberg, Clark Thomborson, and Douglas Low. Manufacturing Cheap, Resilient, and Stealthy Opaque Constructs. In

- ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL), pages 184–196. ACM, 1998.
- [174] Babak Yadegari, Jon Stephens, and Saumya Debray. Analysis of Exception-Based Control Transfers. In *ACM Conference on Data and Application Security and Privacy (CODASPY)*, pages 205–216. ACM, 2017.
- [175] Zipkin. OpenZipkin · A distributed tracing system, 2019. URL <https://zipkin.io/>.
- [176] Alexandre Bergel, Felipe Bañados, Romain Robbes, and Walter Binder. Execution profiling blueprints. *Software: Practice and Experience*, 42(9):1165–1192, 2012.
- [177] Joshua Garcia, Daniel Popescu, Gholamreza Safi, William G. J. Halfond, and Nenad Medvidovic. Identifying Message Flow in Distributed Event-based Systems. In *Joint Meeting on Foundations of Software Engineering (ESEC/FSE)*, pages 367–377. ACM, 2013.
- [178] A. Burattin, A. Sperduti, and W. M. P. van der Aalst. Control-flow discovery from event streams. In *IEEE Congress on Evolutionary Computation (CEC)*, pages 2420–2427. IEEE, 2014.
- [179] Wil van der Aalst. *Process Mining: Data Science in Action*. Springer, Berlin, Germany, second edition, 2016.
- [180] Volodymyr Leno, Abel Armas-Cervantes, Marlon Dumas, Marcello La Rosa, and Fabrizio M. Maggi. Discovering Process Maps from Event Streams. In *International Conference on Software and System Process (ICSSP)*, pages 86–95. ACM, 2018.
- [181] Olin Grigsby Shivers. *Control-flow Analysis of Higher-order Languages of Taming Lambda*. PhD Thesis, Carnegie Mellon University, 1991.
- [182] Ryan Cunningham and Eddie Kohler. Making Events Less Slippery with Eel. In *Conference on Hot Topics in Operating Systems (HOTOS)*, pages 1–6. USENIX Association, 2005.
- [183] Josh Evans. Mastering Chaos - A Netflix Guide to Microservices, 2016. URL <https://www.infoq.com/presentations/netflix-chaos-microservices>.
- [184] Payam Barnaghi and Amit Sheth. On Searching the Internet of Things: Requirements and Challenges. *IEEE Intelligent Systems*, 31(6):71–75, 2016.
- [185] Zhenyu Wen, Renyu Yang, Peter Garraghan, Tao Lin, Jie Xu, and Michael Rovatsos. Fog Orchestration for Internet of Things Services. *IEEE Internet Computing*, 21(2):16–24, 2017.
- [186] Michele Zorzi, Alexander Gluhak, Sebastian Lange, and Alessandro Bassi. From today’s INTRANet of things to a future INTERNet of things: a wireless- and mobility-related view. *IEEE Wireless Communications*, 17(6):44–51, 2010.
- [187] Jörg Heuer, Johannes Hund, and Oliver Pfaff. Toward the Web of Things: Applying Web Technologies to the Physical World. *Computer*, 48(5):34–42, 2015.
- [188] Samir Tata, Rakesh Jain, Heiko Ludwig, and Sandeep Gopisetty. Living in the Cloud or on the Edge: Opportunities and Challenges of IOT Application Architecture. In *International Conference on Services Computing (SCC)*, pages 220–224. IEEE, 2017.
- [189] Adam Barker and Rajkumar Buyya. Decentralised Orchestration of Service-oriented Scientific Workflows. In *International Conference on Cloud Computing and Services Science (CLOSER)*, pages 222–231. SciTePress, 2011.
- [190] Tanveer Ahmed, Abhinav Tripathi, and Abhishek Srivastava. Rain4service: An Approach towards Decentralized Web Service Composition. In *International Conference on Services Computing (SCC)*, pages 267–274, Anchorage, AK, USA, 2014. IEEE.
- [191] T. Ahmed, M. Mrissa, and A. Srivastava. MagEl: A Magneto-Electric Effect-Inspired Approach for Web Service Composition. In *International Conference on Web Services (ICWS)*, pages 455–462. IEEE, 2014.
- [192] Walid Fdhila, Ustun Yildiz, and Claude Godart. A Flexible Approach for Automatic Process Decentralization Using Dependency Tables. In *International Conference on Web Services (ICWS)*, pages 847–855. IEEE, 2009.
- [193] M. Schmidt, B. Hutchison, P. Lambros, and R. Phippen. The Enterprise Service Bus: Making service-oriented architecture real. *IBM Systems Journal*, 44(4):781–797, 2005.
- [194] Nicolai Josuttis. *Soa in Practice: The Art of Distributed System Design*. O’Reilly Media, Sebastopol, CA, USA, 1st edition, 2007.
- [195] Walter Binder, Ion Constantinescu, and Boi Faltings. Service invocation triggers: a lightweight routing infrastructure for decentralised workflow orchestration. *International Journal of High Performance Computing and Networking*, 6(1):81–90, 2009.
- [196] Janggwan Im, Seonghoon Kim, and Daeyoung Kim. IoT Mashup as a Service: Cloud-Based Mashup Service for the Internet of Things. In *International Conference on Services Computing (SCC)*, pages 462–469. IEEE, 2013.
- [197] Felipe Pontes Guimaraes, Eduardo Hideo Kuroda, and Daniel Macedo Batista. Performance Evaluation of Choreographies and Orchestration with a New Simulator for Service Compositions. In *International Workshop on Computer Aided Modeling and Design of Communication Links and Networks (CAMAD)*, pages 140–144. IEEE, 2012.
- [198] David Liu. Data-flow Distribution in FICAS Service Composition Infrastructure. In *International Conference on Parallel and Distributed Computing Systems (PDCAT)*, pages 1–6, 2002.
- [199] Jan Seeger, Rohit A. Deshmukh, Vasil Sarafov, and Arne Bröring. Dynamic IoT Choreographies. *IEEE Pervasive Computing*, 18(1):19–27, 2019.
- [200] Mahda Noura, Mohammed Atiquzzaman, and Martin Gaedke. Interoperability in Internet of Things: Taxonomies and Open Challenges. *Mobile Networks and Applications*, 24(3):796–809, 2019.
- [201] Frank Leymann. Web Services: Distributed Applications Without Limits - An Outline. In *Database Systems for Business, Technology and Web (BTW)*, pages 1–22, 2003.
- [202] Kung-Kiu Lau, Vladyslav Ukis, Perla Velasco, and Zheng Wang. A Component Model for Separation of Control Flow from Computation in Component-Based Systems. In *International Workshop on Aspect-Based and Model-Based Separation of Concerns in Software Systems (ABMB)*, pages 57–69, 2006.
- [203] Farhad Arbab. Composition of Interacting Computations. In Dina Goldin, Scott A. Smolka, and Peter Wegner, editors, *Interactive Computation*, pages 277–321. Springer, Berlin, Germany, 2006.
- [204] Ahmed Safwat and M. B. Senousy. Addressing Challenges of Ultra Large Scale System on Requirements Engineering. In *International Conference on Communications, Management, and Information Technology (ICCMIT)*, International Conference on Communications, management, and Information technology (ICCMIT’2015), pages 442–449, 2015.
- [205] Mehdi Mirakhorli, Amir Azim Sharifloo, and Fereidoon Shams. Architectural challenges of ultra large scale systems. In *International Workshop on Ultra-Large-Scale Software-Intensive Systems (ULSSIS)*, pages 45–48. ACM, 2008.
- [206] Andrea Ceccarelli, Andrea Bondavalli, Bernhard Froemel, Oliver Hoefberger, and Hermann Kopetz. Basic Concepts on Systems of Systems. In Andrea Bondavalli, Sara Bouchenak, and Hermann Kopetz, editors, *Cyber-Physical Systems of Systems: Foundations – A Conceptual Model and Some Derivations: The AMADEOS Legacy*, pages 1–39. Springer, Cham, Switzerland, 2016.
- [207] Amel Bennaceur, Ciaran McCormick, Jesús García-Galán, Charith Perera, Andrew Smith, Andrea Zisman, and Bashar Nuseibeh. Feed Me, Feed Me: An Exemplar for Engineering Adaptive Software. In *International Symposium on Software Engineering for Adaptive and Self-Managing Systems (SEAMS)*, pages 89–95. IEEE, 2016.
- [208] Marcello La Rosa, Wil M. P. Van Der Aalst, Marlon Dumas, and Fredrik P. Milani. Business Process Variability Modeling: A Survey. *ACM Computing Surveys*, 50(1):1–45, 2017.
- [209] Official government website for Northern Ireland citizens. Off-street and on-street parking, 2019. URL <https://www.nidirect.gov.uk/articles/street-and-street-parking>.
- [210] Wei Fan, Zhengyong Chen, Zhang Xiong, and Hui Chen. The Internet of data: a new idea to extend the IOT in the digital world. *Frontiers of Computer Science*, 6(6):660–667, 2012.
- [211] Richard N. Taylor, Nenad Medvidovic, and Eric M. Dashofy. *Soft-*

- ware Architecture: Foundations, Theory, and Practice. Wiley Publishing, Hoboken, NJ, USA, 1st edition, 2009.
- [212] Paolo Nesi, Gianni Pantaleo, Michela Paolucci, and Imad Zaza. Auditing and Assessment of Data Traffic Flows in an IoT Architecture. In *International Conference on Collaboration and Internet Computing (CIC)*, pages 388–391. IEEE, 2018.
- [213] Giacomo Ghidini, Sajal K. Das, and Vipul Gupta. Fuseviz: A framework for web-based data fusion and visualization in smart environments. In *International Conference on Mobile Ad-Hoc and Sensor Systems (MASS)*, pages 468–472. IEEE, 2012.
- [214] Juan Luis Pérez, Álvaro Villalba, David Carrera, Iker Larizgoitia, and Vlad Trifa. The COMPOSE API for the Internet of Things. In *International Conference on World Wide Web (WWW)*, pages 971–976. ACM, 2014.
- [215] Antonio Pintus, Davide Carboni, and Andrea Piras. Paraimpu: a platform for a social web of things. In *International Conference on World Wide Web (WWW)*, pages 401–404. ACM, 2012.
- [216] Spacebrew. What is Spacebrew?. 2019. URL spacebrew.cc.
- [217] Gennaro De Luca, Zhongtao Li, Sami Mian, and Yinong Chen. Visual programming language environment for different IoT and robotics platforms in computer science education. *CAAI Transactions on Intelligence Technology*, 3(2):119–130, 2018.
- [218] Per Persson and Ola Angelsmark. Calvin – Merging Cloud and IoT. In *International Conference on Ambient Systems, Networks and Technologies (ANT)*, pages 210–217. Elsevier, 2015.
- [219] Intel. IoT Services Orchestration Layer, 2019. URL <https://github.com/intel/intel-iot-services-orchestration-layer>.
- [220] NoFlo. NoFlo: Flow-Based Programming for JavaScript, 2019. URL noflojs.org.
- [221] Gábor Paller, Endri Bezati, Nebojša Taušan, Gábor Farkas, and Gábor Élő. Dataflow-based Heterogeneous Code Generator for IoT Applications. pages 428–434. IEEE, 2019.
- [222] Yuuichi Teranishi, Takashi Kimata, Hiroaki Yamanaka, Eiji Kawai, and Hiroaki Harai. Dynamic Data Flow Processing in Edge Computing Environments. In *Annual Computer Software and Applications Conference (COMPSAC)*, pages 935–944. IEEE, 2017.
- [223] Marco Mesiti, Stefano Valtolina, Luca Ferrari, Minh-Son Dao, and Koji Zettsu. An editable live ETL system for Ambient Intelligence environments. In *World Forum on Internet of Things (WF-IoT)*, pages 393–394. IEEE, 2015.
- [224] Mahmoud Hussein, Shuai Li, and Ansgar Radermacher. Model-driven development of adaptive IoT systems. In *International Workshop on Interplay of Model-Driven and Component-Based Software Engineering (ModComp)*, pages 20–27, 2017.
- [225] Christian Prehofer and Ilias Gerostathopoulos. Modeling RESTful Web of Things Services: Concepts and Tools. In Quan Z. Sheng, Yongrui Qin, Lina Yao, and Boualem Benatallah, editors, *Managing the Web of Things*, pages 73–104. Morgan Kaufmann, Boston, MA, USA, 2017.
- [226] Andrei Ciortea, Olivier Boissier, Antoine Zimmermann, and Adina Magda Florea. Responsive Decentralized Composition of Service Mashups for the Internet of Things. In *International Conference on the Internet of Things (IoT)*, IoT’16, pages 53–61. ACM, 2016.
- [227] Florian Daniel and Maristella Matera. *Mashups: Concepts, Models and Architectures*. Data-Centric Systems and Applications. Springer, Heidelberg, Germany, 1st edition, 2014.
- [228] Michael Blackstock and Rodger Lea. IoT mashups with the WoTKit. In *International Conference on the Internet of Things (IoT)*, pages 1–8. IEEE, 2012.
- [229] Florian Daniel and Barbara Pernici. Insights into Web Service Orchestration and Choreography. *International Journal of E-Business Research*, 2(1):58–77, 2006.
- [230] Patrik Spiess, Stamatis Karnouskos, Dominique Guinard, Domnic Savio, Oliver Baecker, Luciana Moreira Sá De Souza, and Vlad Trifa. SOA-based integration of the internet of things in enterprise services. In *International Conference on Web Services (ICWS)*, pages 968–975. IEEE, 2009.
- [231] Antonio Pintus, Davide Carboni, Andrea Piras, and Alessandro Gior-dano. Connecting Smart Things through Web Services Orchestration. In Florian Daniel and Federico Michele Facca, editors, *Current Trends in Web Engineering*, volume 6385 of *Lecture Notes in Computer Science*, pages 431–441. Berlin, Germany, 2010. Springer.
- [232] Nils Glombitza, Sebastian Ebers, Dennis Pfisterer, and Stefan Fischer. Using BPEL to Realize Business Processes for an Internet of Things. In Hannes Frey, Xu Li, and Stefan Ruehrup, editors, *Ad-hoc, Mobile, and Wireless Networks*, volume 6811 of *Lecture Notes in Computer Science*, pages 294–307. Berlin, Germany, 2011. Springer.
- [233] Hagen Overdick. Towards Resource-Oriented BPEL. In Thomas Gschwind and Cesare Pautasso, editors, *Emerging Web Services Technology*, pages 129–140. Birkhäuser, Basel, Switzerland, 2008.
- [234] Feng Chen, Changrui Ren, Qinhua Wang, and Bing Shao. A process definition language for Internet of things. In *International Conference on Service Operations and Logistics, and Informatics (SOLI)*, pages 107–110. IEEE, 2012.
- [235] Dulce Domingos, Francisco Martins, Ricardo Martinho, and Mário Silva. Ad-hoc changes in IoT-aware business processes. In *Internet of Things (IOT)*, pages 1–7. IEEE, 2010.
- [236] OMG. Business Process Model And Notation (BPMN) Version 2.0, 2011. URL <https://www.omg.org/spec/BPMN/2.0/>.
- [237] Camunda. BPMN Workflow Engine, 2019. URL <https://camunda.com/products/bpmn-engine/>.
- [238] Tijs Rademakers. *Activiti in Action: Executable Business Processes in BPMN 2.0*. Manning Publications, Greenwich, CT, USA, 1st edition, 2012.
- [239] Minjae Park, Hyunah Kim, Hyun Ahn, and Kwanghoon Pio Kim. A process-aware IoT application execution environment design. In *International Conference on Advanced Communication Technology (ICACT)*, pages 724–727. IEEE, 2018.
- [240] Gero Decker, Oliver Kopp, Frank Leymann, and Mathias Weske. BPEL4chor: Extending BPEL for Modeling Choreographies. In *International Conference on Web Services (ICWS)*, pages 296–303. IEEE, 2007.
- [241] Steve Ross-Talbot and Tony Fletcher. *Web Services Choreography Description Language: Primer*, 2006. URL <https://www.w3.org/TR/ws-cdl-10-primer/>.
- [242] Johannes Maria Zaha, Alistair Barros, Marlon Dumas, and Arthur ter Hofstede. Let’s Dance: A Language for Service Behavior Modeling. In *On the Move to Meaningful Internet Systems 2006: CoopIS, DOA, GADA, and ODBASE*, volume 4275 of *Lecture Notes in Computer Science*, pages 145–162. Berlin, Germany, 2006. Springer.
- [243] Assaf Arkin, Sid Askary, and Scott Fordin. Web Service Choreography Interface (WSCI) 1.0, 2002. URL <https://www.w3.org/TR/2002/NOTE-wsci-20020808/>.
- [244] Gero Decker and Alistair Barros. Interaction Modeling Using BPMN. In Arthur ter Hofstede, Boualem Benatallah, and Hye-Young Paik, editors, *Business Process Management Workshops*, volume 4928 of *Lecture Notes in Computer Science*, pages 208–219. Berlin, Germany, 2007. Springer.
- [245] Michael Zimmermann, Uwe Breitenbücher, and Frank Leymann. A TOSCA-based Programming Model for Interacting Components of Automatically Deployed Cloud and IoT Applications. In *International Conference on Enterprise Information Systems (ICEIS)*, pages 121–131. SciTePress, 2017.
- [246] Andreas Weiß, Vasilios Andrikopoulos, Santiago Gómez Sáez, Michael Hahn, and Dimka Karastoyanova. ChorSystem: A Message-Based System for the Life Cycle Management of Choreographies. In Christophe Debruyne, Hervé Panetto, Robert Meersman, Tharam Dillon, eva Kühn, Declan O’Sullivan, and Claudio Agostino Ardagna, editors, *On the Move to Meaningful Internet Systems: OTM 2016 Conferences*, volume 10033 of *Lecture Notes in Computer Science*, pages 503–521. Cham, Switzerland, 2016. Springer.
- [247] F. L. Traversa and M. Di Ventra. Universal Memcomputing Machines. *IEEE Transactions on Neural Networks and Learning Systems*, 26(11):2702–2715, 2015.
- [248] A. Rajhans, A. Bhawe, I. Ruchkin, B. H. Krogh, D. Garlan, A. Platzler, and B. Schmerl. Supporting Heterogeneity in Cyber-Physical Systems

- Architectures. *IEEE Transactions on Automatic Control*, 59(12): 3178–3193, 2014.
- [249] Tomas Bures, Ilias Gerostathopoulos, Petr Hnetyka, Jaroslav Keznikl, Michal Kit, and Frantisek Plasil. DEECO: An Ensemble-based Component System. In *International ACM Sigsoft symposium on Component-based software engineering (CBSE)*, pages 81–90. ACM, 2013.
- [250] Ikuo Nakagawa, Masahiro Hiji, and Hiroshi Esaki. Dripcast - Architecture and Implementation of Server-less Java Programming Framework for Billions of IoT Devices. *Journal of Information Processing*, 23(4):458–464, 2015.
- [251] Y. Zhang, J. L. Chen, and B. Cheng. Integrating Events into SOA for IoT Services. *IEEE Communications Magazine*, 55(9):180–186, 2017.
- [252] F. Halili, E. Rufati, and I. Ninka. Styles of Service Composition – Analysis and Comparison Methods. In *International Conference on Computational Intelligence, Communication Systems and Networks (CICSyN)*, pages 302–307. IEEE, 2013.