Enabling Discoverable Trusted Services for Highly Dynamic Decentralized Workflows

Iain Barclay*, Chris Simpkin*, Graham Bent†, Tom La Porta‡, Declan Millar†,
Alun Preece*, Ian Taylor*, Dinesh Verma§
*School of Computer Science and Informatics, Cardiff University, UK
Email: BarclayIS@cardiff.ac.uk
†IBM Research, Europe
‡The Pennsylvania State University, USA
§IBM Research, USA

***Abstract*—Fifth generation (5G) mobile networks will revolutionize edge-based computing by providing fast and reliable network capabilities to remote sensors, devices and microservices. This heralds new opportunities for researchers, allowing remote instrumentation and analytic capabilities to be as accessible as local resources. The increased availability of remote data and services presents new opportunities for collaboration, yet introduces challenges for workflow orchestration, which will need to adapt to consider an increased choice of available services, including those from trusted partners and the wider community. In this paper we outline a workflow approach that provides decentralized discovery and orchestration of verifiably trustable services in support of multi-party operations. We base this work on the adoption of standardised data models and protocols emerging from hypermedia research, which has demonstrated success in using combinations of Linked Data, Web of Things (WoT) and semantic technologies to provide mechanisms for autonomous goal-directed agents to discover, execute and reuse new heterogeneous resources and behaviours in large-scale, dynamic environments. We adopt Verifiable Credentials (VCs) to securely share information amongst peers based on prior service usage in a cryptographically secure and tamperproof way, providing a trust-based framework for ratifying service qualities. Collating these new service description channels and integrating with existing decentralized workflow research based on vector symbolic architecture (VSA) provides an enhanced semantic search space for efficient and trusted service discovery that will be necessary for 5G edge-computing environments.**

I. Introduction

Emerging fifth generation (5G) mobile networks offer the promise of low latency, high capacity, and increased bandwidth, providing a completely new communications infrastructure organized as Multi-access Edge Computing (MEC), enabling edge resources to enjoy network connectivity on a par with that experienced in data centers today. As a result, the quantity and capabilities of devices and remote microservices deployed as part of the Internet of Things (IoT) will increase dramatically. In support, 5G network slices provide a new way to virtualize infrastructure using dynamic software defined provisioning technologies and service aggregation, allowing consortia groups and coalitions to create dynamic virtual private networks optimised for different use cases.

Workflow orchestrations of services assembled within these network slices and given access to the MEC infrastructure will require increasing levels of trust in the capabilities and integrity of services and IoT devices they select to use. Resources will often be provided by consortium partners, third parties such as city authorities, or the open source community, and it will often be critical to have strong authentication, confidentiality, availability and privacy guarantees. In the case of a cyberattack, for example, the consequences could be costly or dramatic e.g. an IoT based vehicle may provide severe damage reports if erroneous sensor information is provided. To counter such threats, service configurations need to establish trust in the IoT device itself, and the data it transmits, in a potentially trustless environment.

As a result, dynamic workflow configuration needs to be able to provide rapid and autonomous discovery and configuration of *suitable* resources from a plethora of available devices and services, based on factors such as capabilities and quality of service, and augmented by networks of trust. The author's previous work has shown that decentralized workflow mechanisms, such as those based on Vector Symbolic Architectures (VSA), can be used to facilitate efficient service discovery and workflow formulation in decentralized collaborative environments [1], [2], [3], without requiring central control [4]. Such decentralized mechanisms are well suited to the low latency, high capacity, and high bandwidth connectivity afforded to edge computing resources as part of 5G networks, and the business and research ecosystems that they will support.

In this paper we describe the use of a VSA to architect a mechanism that can be used for distributed discovery and orchestration of edge devices and microservices, where device and service descriptions are derived from interoperable linked data, semantic web technologies, and emerging open web standards, such that pre-existing descriptive resources can be utilised as far as possible. In addition, a dynamic layer of trust is added to service descriptions through the employment of certified credential documents (from organisations including World Wide Web Consortium (W3C)[1]) which are used to provide assurance on service qualities as experienced by trusted peers. The service descriptions and credentials are collated to

[1]https://w3.org

seed service and sensor descriptions in the following way:

- Interfaces (APIs) to microservices are provisioned with SPARQL [5] service descriptions, which offer self-describing mechanisms for interactions, exploiting a capability for "Continuous Acquisition of Behaviors" [6], to make key features prominent, based on prior successful usage [7]. RDF triples underlying the SPARQL descriptions can be further augmented by service descriptions provided as Web of Things (WoT) 'Thing Descriptions' (TD) [8] or other JSON-LD [9] format Linked Data descriptions.
- W3C standard Verifiable Credential proofs, in the form of JSON-LD documents, enable service providers to attest to the *specifications* of their service offerings, whilst peers can digitally sign assertions about their service *experiences*, based on their use of the service. E.g., "The AI service was effective at classifying tanks." As a result, JSON-LD documents are created describing both service specifications and service experiences.
- SPARQL and JSON-LD service descriptions, service specification JSON-LD VCs, and experiential JSON-LD VCs are encoded using VSA techniques, such that they can be integrated and used by a VSA workflow system, and provide a capability for efficient searching for services based on the semantic properties of VSA.

As a result, mechanisms can be provided which are able to take advantage of light-weight, standardized and interoperable service description technologies to enable efficient service discovery and orchestration via VSA across dynamic decentralized environments. VSA enables service discovery in a semantic space, which will be seeded by the linked data descriptions provided by service providers and users. Using a semantic search provides service discovery that is decoupled from the precise language used by service operators and users, whilst experience reports from users can be used to refine service selection, prioritising services which have been successfully used by trusted peers.

Our specific use case is being conducted in the context of the International Technology Alliance in Distributed Analytics and Information Sciences (DAIS ITA) [10] project, which is aiming to enable the creation of a distributed *cognitive computer system* that can perform analytics on demand across heterogeneous networks of interconnected devices in support of coalition operations, where multiple partners share sensing and information processing assets. In such an environment, clients want to be able to ensure services are ratified and trusted, and they also would like to make selections based on the degree of trust those services exhibit e.g. have they been ratified by a member of our defense force or by coalition member?

The rest of the paper is organized as follows. Section II describes related work in the service discovery area and Section III provides a military scenario within an extreme environment for the workflow orchestration example. Section IV considers technologies and specifications for service interfaces and drills down into how the use of SPARQL, RDF, JSON-LD, DIDs and Verifiable Credentials can help to address such requirements. Section V shows how such an approach can evolve in time to provide more trusted service selections and deeper semantic information. Section VI provides details on how service descriptions can be encoded into VSA vectors, which are used in Section VII to provide workflow orchestration. The paper concludes with Section VIII.

II. Related work

Service discovery is a key component in a transient distributed networked environment, but is a difficult problem even when services are hosted in centralized repositories, mainly because services are developed and deployed independently or developed by loosely cooperating developers in open environments. This has led to a complex mix of disparate service architectures employing different methodologies for the description of their inputs, outputs, and configurations. In support of such situations, we are employing vector based representations as a means of encoding service descriptions that can be semantically compared within particular contexts, in an extremely resource efficient way. Using such vectors, semantically rich queries in the form of vectors, can be sent out to the network, using protocols such as multicast for efficient querying in a complex space.

In order to seed vector representations of services, we propose to leverage a number of semantically rich Linked Data [11] service description standards, as might plausibly be developed by providers of IoT devices and published microservices. Linked Data principles provide mechanisms which are regarded by Mayer et al. [12] as having the ability "to underpin systems that integrate multiagent planning and acting with semantic technologies and with interoperable mixed reality interfaces", enabling "the creation of highly augmented environments... where physical and digital things coexist and interact with one another." Suri et al. [13] provide an analysis of the applicability of these 'physical and digital things' in a decentralized environment, and conclude that technical challenges enumerated by Zheng [14] in regards to connectivity, digital analytics, and interoperability of assets in decentralized environments can be addressed through the use of semantic web [15] technologies, which are identified as providing "(I) Open integration standards; (II) Reasoning support; (III) Support for data provenance management." and state that "one of the many applications of IoT would be shared sensing among mobile devices / sensor nodes in an area of interest. Sensing — e.g., of the environment, people, and devices — is at the core of IoT... Combined with robust short range communication, IoT devices would be able to utilize placement or sensing modality of other sensors to supplement their own sensing methods" [13].

Zschorn et al. [16] reflect that "Information requirements of Defence operational staff are ... many, varied and changing, and sometimes unpredictable. These various operations require at times communication and coordination with coalition military partners, federal and state police forces, other govern-

ment agencies, and non-government agencies." As such, it is important to be able to develop trust in providers and sources of data, which is raised as a concern by Michaelis et al. [17]: "...information derived from IoT sources may have varying degrees of integrity, possibly making it unfit for analyst/commander usage" who identify a need for "methods to associate records of provenance with information, sufficiently detailed for a collection of (possibly unforeseen) assessment tasks".

III. A Dynamic Service Configuration Use Case

Military scenarios provide an extreme environment for the application of workflow configuration architectures, as they consist of partner or peer organisations with varying and fluid levels of trust, and need to be deployed in fragile environments, often in transient mobile ad hoc networks (MANETs) which are typical of battlefield network scenarios, when devices are coming in and out of range, and network fragmentation occurs frequently.

As such, a mechanism is needed that can locate and orchestrate the required workflow in the face of these challenges, without central control since it in not possible to rely on centralised registries, or even to know the IP location of objects and devices as they come in and out of service. As described in [1], [2], [3], the VSA architecture can be used for peer-to-peer (P2P) discovery of appropriate devices and functional micro-services without service and device registries because the VSA representation is able to act as both the object description and the address of the object.

The architecture presented herein is enumerated with reference to a military use case scenario depicting a sequence of events, initiated by receipt of an intelligence report regarding a possible insurgent rendezvous at a given location. The intelligence report is analysed, and it is determined that assets available in the area should be located and connected into a workflow that can provide "anomaly detection covering location 58.145, 7.998".

As a result, a VSA workflow is encoded specifying the devices and functional analytic objects required to perform analysis and detection in the area. In response to the VSA multicast, capable services are located and configured. For example, a monitoring service covering the location could be configured from an acoustic sensor, a camera unit, and an AI service which collectively responded to the VSA multicast.

The provisioned service monitors the target location, and subsequently detects a suspected anomaly (in this case, a possible shooting). The original VSA request can contain triggering thresholds, along with instructions for subsequent actions to take if triggering thresholds are met or not met. For example, if the trigger threshold was not met, monitoring could be restarted (a looping action) or if the threshold was met, further instructions could be carried out, for example a new multicast request could be made to request further monitoring services to be configured over a wider area.

IV. Service Descriptions

Use of interoperable semantic web technologies and emerging standards for describing microservices such as APIs, datasets and sensor devices, provides an opportunity to leverage semantic mechanisms and tooling from a wider community, and provides a route for services to ingest open source resources, including video infrastructure in a city [18], for example, as well as those provided within a collaboration. Where APIs and microservices from the open web are used as part of solution configurations, provider-supplied descriptions can be enriched with overlays of additional learning from service usage and reported service experiences, which can be used to seed or augment user's knowledge of the services. Web technologies can be adopted to allow peers to share service specifications and actual experiences in such a way that trust relationships can be used as a factor in service selection. At this stage, three streams of potential service description metadata have been identified, which can be collated to provide an overall service description which is semantically searchable via VSA.

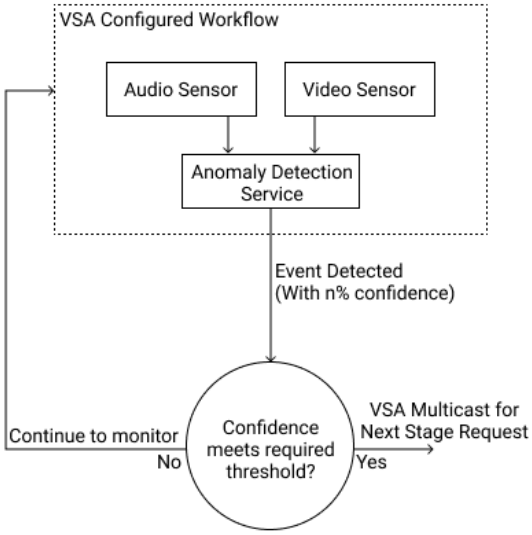

Fig. 1: Workflow Configuration

A. SPARQL, RDF and JSON-LD

Recent literature has described experiences in bridging web service API's with semantically searchable interfaces through the provision of SPARQL service descriptions. Michel et al. [19] developed an interface to a web-hosted dataset in which "a SPARQL micro-service evaluates a query against an RDF graph built at run-time from data obtained from the Web API", which was later extended [20] to provide a SPARQL front-end to a service which aggregated results of the previous experiment with those from a REST API. A PHP server application bridges a SPARQL client and the web APIs, converting SPARQL queries into search parameters. Such a service could be provisioned to provide access an interface to a web service, an ML model, a sensor or actuator, or a dataset [19]. As such, SPARQL interfaces can be created for a range of resources offered by service providers, third parties or open source, providing adaptable semantic search interfaces where none currently exists.

For some services, linked data service descriptions will be made available by the service provider. This might be the case for a physical asset, where a Web of Things 'Thing Description' document could be provided on the asset itself or via a proxy to give a JSON-LD Linked Data description of the device's affordances, or its capabilities and how it is to be used, "in order to increase interoperability between connected devices and develop arbitrarily complex mash-ups" [21]. Similarly, a dataset might be accompanied by a Linked Data description. Where such information is supplied by a service provider, this can be rendered as RDF triples to provide the SPARQL micro-service's description. Bienz et al. [22] have demonstrated the use of SPARQL endpoints to provide approximate search capabilities for physical devices, seeded from published WoT TDs.

B. Verifiable Credentials

The term Self-Sovereign Identity (SSI) [23] is used to describe the ability of an individual to take ownership of their personal data and to control who has access to that data, without the need for a centralised infrastructure, or any control or authorization being required by any third party. SSI has been the subject of research and ambition for several years, but has reached an inflection point in interest from industry and the research community as a result of the availability of distributed ledger and blockchain-based technologies, combined with an increased focus on individual's data privacy as they interact with web-based and social networking services [24].

SSI is decentralized, and is built upon well-established cryptographic techniques whereby a user holds a private and shares a public key [25]. The private key is used to sign documents, whilst the public key can be used by anybody with access to it to verify that the document has indeed been signed, and has not been tampered with. SSI uses a system built on decentralized identifiers (DIDs) to identify parties involved, with the DIDs resolving to documents which explain, in machine-readable format, how to locate the public key needed to validate claims made about that DID, in the same way as web addresses resolve to provide web pages. The SSI research community has developed data models and protocols [26] that provide mechanisms for any party identified by a DID to issue cryptographically verifiable sets of credentials to any subject entity, also identified by a DID. In this way, a party which believes something to be true about another party can declare this in a standardised way using a JSON-LD formatted document, and sign this attestation using asymmetric cryptography techniques, based on the DIDs used being able to be resolvable to validate the assertions made. This cryptographically signed document is known as a Verifiable Credential (VC), and will be held by the subject of the credential, or in the case of a child, or dataset or physical asset, by an authorised Holder.

At a later date, when the holder seeks to enter into a transaction, a service provider may request proof of status or entitlements. The Verifiable Credential document provides a means for this proof to be provided, as the holder of the credential can generate a Verifiable Presentation containing assertions from the VC document. By processing the Verifiable Presentation document, the Verifier can use the accessible public keys to check that i) the presented proof pertains to the subject it is being presented on behalf of, ii) the presented proof contains assertions signed by the original Issuer, and finally iii) that the presented documents have not been tampered with. As such, triangles of trust [27] can be leveraged to enable parties to issue, hold and verify credentials without reliance on any central authority.

To date, the focus of effort in the SSI community has been on personal identity and data privacy for individuals [28], however the underlying computer science techniques can be applied to any type of entity, including digital assets such as datasets [29], and devices [30]. Systems based on the paradigm of the self-sovereignty of human participants, data resources and devices are inherently decentralized, with attributes held at the edges of the ecosystem.

1) Specification Credentials: To provide further context on services, we can use linked data documents, in the form of VCs, to augment information provided by service description documents or RDF triples with further contextual information about the service which can be used for reasoning on its suitability for use in a workflow.

In the first instance, solution vendors, systems integrators or representatives of service providers can issue attestations relating to the specifications or qualities of the service, which can be stored and made available for inspection. These are cryptographically signed JSON-LD documents (VC_S), which can be verified against a public key known to be owned by the signer and credential issuer, which can prove that the document was signed by the issuer and has not been tampered with. As an example, a service provider deploying a video camera at a certain location could issue a signed VC asserting the coordinates of the location of the camera, such that any parties interested in using the camera could inspect the credential and (provided they trusted the signing party) could be assured of the location of the camera.

2) Experience Credentials: Additionally, it is possible for any service users to issue a signed credential relating to their own experiences in using the service, which would be held as part of the service's metadata corpus. For example, if an AI service offers vehicle identification[31] and a user has had good success using the service for identifying Sports Utility Vehicles (SUV), then they are able to create a JSON-LD document attesting to this and sign it, resulting in an experience credential (VC_XP) being held by the service. Subsequent parties seeking SUV identification would be able to inspect this document, and determine its suitability as additional information, based on any knowledge and trust they have in the party providing the document, as identified by the decentralized identifier (DID) of the signing party.

As a result, there are three channels which collectively describing the service (Figure 2), its specifications and usage experiences - service descriptions, comprised of RDF triples derived from a SPARQL interface and which can be serialised

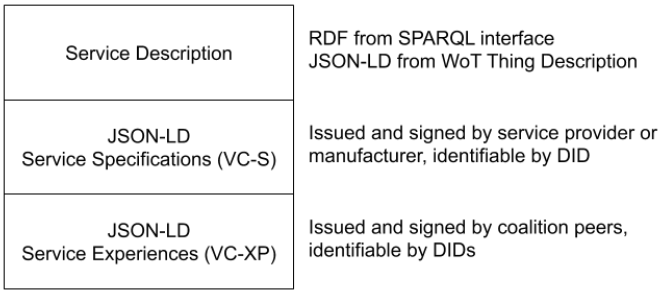

Fig. 2: Web Standards are used to provide service descriptions

and combined with any JSON-LD device or service descriptions, JSON-LD specification credentials (VC_S) provided by service owners, and the JSON-LD experiential reports from service users (VC_XP). Each information channel can be VSA encoded, resulting in a vector describing the service's capabilities, as will be described further in Section VI.

In the context of the vignette introduced in Section III, services under consideration for configuration as part of an anomaly detection service could include *Input sensors*, for example, an audio listening device and a camera, with data which can be used to build knowledge of the service coming from:

- A JSON-LD Web of Things TD interface for each device
- RDF triples, derived from observed use of the service via a SPARQL interface.
- Specifications for the service deployment (e.g., the sensor's location) issued by the service provider and held as VC_S.
- Experiences of service use. E.g., descriptions of landmarks or points of interest that can be seen from the camera, held by the service as VC_XP.

And an *AI service* requiring audio and video inputs and capable of triggering on detection of anomalies, which can be described in terms of:

- RDF Triples derived from a SPARQL interface to a Web API.
- Specification credentials from the service provider as VC_S.
- Experiences from past users of the service, as VC_XP.

V. Improvements with Time

As microservices and devices are deployed and used in different workflows, users are able to craft experience reports about their own usage of the service or device. These reports, structured as JSON-LD documents, can be cryptographically signed and protocols for issuing Verifiable Credentials can be used to assign them to the service. As such, new information about actual experiences of services becomes increasingly available as services are used (Figure 3), and is able to provide deeper semantic information about service qualities. Furthermore, as cryptographic signatures based on decentralized identifier properties are used to sign these experience attestations, they can be traced back to pseudonymous identities who may be known by potential service users. As a result, service selection can begin to include the presence of trusted experience reports as a selection criteria in workflow configuration – reverting back to a military context, a UK commander would be more likely to trust a service ratified by their US counterpart, than one not ratified at all, or indeed one ratified by a more transient ally.

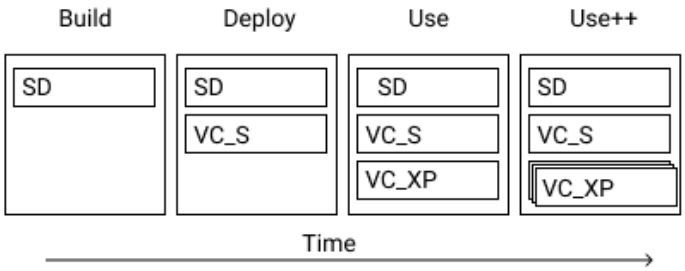

Fig. 3: Experiences are added as service is used

VI. Encoding Sensor Device and Service Objects

Vector Symbolic Architectures are a family of bio-inspired methods for representing and manipulating concepts and their meanings in a high dimensional vector space [32]. They are a form of 'brain like' distributed representation that enables large volumes of data to be combined into a fixed size feature vector such that the semantic meaning of the data and relationships that they represent is preserved. Such vector representations were originally proposed by Hinton [33] who identified that they have recursive binding properties that allow for higher level semantic vector representations to be formulated from, and in the same format as, their lower level semantic vector components. Eliasmith, in his book 'How to Build a Brain' [34], shows how these vector representations can be used to perform 'brain like' neuromorphic cognitive processing. He coined the phrase 'semantic pointer' for such a vector since it acts as both a 'semantic' description of the concept, which can be manipulated directly and a 'pointer' to the concept. As such they are said to be semantically self-describing. VSAs are also capable of supporting a large range of cognitive tasks such as: Semantic composition and matching; Representing meaning and order; Analogical mapping; and Logical reasoning [35]. Consequentially they have been used in natural language processing [35], [36] and cognitive modelling [34], [37].

Our approach for creating semantically rich representation of services and workflows is to represent them as high level semantic concept vectors that are themselves constructed from semantic vectors representing their sub features in a recursive manner using vector binding and superposition operations as described in [3], [2]. Reviewing that X_r represents a role vector and Y_v a value vector, role vectors can be bound to filler vectors using the binding operation $X_r \ . \ Y_v$ (where '.' is the binding operation).

$$\boldsymbol{Z_v = SD_r \cdot SD_v + VC\_S_r \cdot VC\_S_v + VC\_XP_r \cdot VC\_XP_v} \quad (1)$$

Thus, $\boldsymbol{Z_v}$, the high-level semantic vector representation of the sensor/device or service object, is made up of a nested superposition of its sub-feature vectors, $\boldsymbol{SD_v}$, $\boldsymbol{VC\_S_v}$, $\boldsymbol{VC\_XP_v}$ which are themselves high-level concept vectors built from RDF triple sets or parsed JSON-LD.

As an example, Listing 1 shows a Web of Things 'Thing Description' for a web camera [38], encoded in JSON-LD, as $\boldsymbol{SD_v}$ of the $\boldsymbol{Z_v}$ object description. This in turn is converted to a flattened collection of sub-features as described in [2].

Listing 1: WoT Thing Description for Camera Sensor

```
{
 "@context": "https://iot.mozilla.org/schemas/",
 "@type": ["Camera", "VideoCamera"],
 "name": "Web Camera",
 "description": "Mobile web camera",
 "properties": {
   "video": {
     "@type": "VideoProperty",
     "title": "Stream",
     "links": [{
       "href": "rtsp://eg.com/video.mp4",
       "mediaType": "video/mp4"
     }]
   },
   "image": {
     "@type": "ImageProperty",
     "title": "Snapshot",
     "links": [{
       "href": "http://eg.com/image.jpg",
       "mediaType": "image/jpg"
     }]
   }
 }
}
```

The $\boldsymbol{VC\_S_v}$ and $\boldsymbol{VC\_XP_v}$ components can be encoded in a similar manner. Listing 2 shows an example of a Verifiable Credential, used as a VC_S, issued by a service provider when deploying a camera device. The credential stores information about the deployed location of the camera, along with a cryptographic proof which can be used by relying parties to verify that the credential document has not been tampered with. The credential could be deleted or revoked when the camera is moved to a new location.

Listing 2: An example VC_S for a deployed Camera Sensor

```
{
 "@context": [
   "https://www.w3.org/2018/credentials/v1",
   "https://schema.org/"
 ],
 "id": "http://eg.com/credentials/e94a16cb",
 "type": [
     "VerifiableCredential",
     "DeployedDeviceCredential"
 ],
 "name": "Camera Deployment",
 "description": "Roadside camera deployed.",
 "issuer": "did:v1:nym:z6Mk..63oP39k",
 "issuanceDate": "2020-04-09T21:13:13Z",
 "credentialSubject": {
   "deviceIdentifier": "3a185b8f",
   "deployedLocation": {
     "address": "Kirkegata, Anglova",
     "latitude": "58.145",
     "longitude": "7.998"
   }
   ...
 },
 "proof": {
   "type": "Ed25519Signature2018",
   "created": "2020-04-09T21:13:28Z",
   "proofPurpose": "assertionMethod",
   "verificationMethod": "did:v1::nym:z6Mkhdm"
 }
}
```

All sensor, device or microservice descriptions are encoded into VSA service descriptions, directly on the edge resource, or on a proxy capable of representing the resource. As such, each entity becomes VSA aware, and is able to support a VSA cognitive wrapper service executed on the device object or on its proxy, facilitating peer-to-peer service discovery and orchestration.

VII. Service Construction and Orchestration

Once sensor and functional micro-services have been encoded into VSA vectors, we can construct a VSA workflow vector capable of discovering, connecting, and orchestrating such services, as described in [3, pages 28-31] and [3]. On injection of a workflow vector into a MANET, for example, the VSA architecture will automatically locate appropriate devices and service objects and assemble them into a sensor chain arrangement. Each such sensor chain would be encoded into a VSA vector in a similar manner. Equation (2) shows a generalised sensor chain encoding using the hierarchical binding notation defined in [2], [3], for example.

$$\begin{aligned} Sensor\_Chain\_n = {} & p0^0 \cdot Sensor\_Z_v^1 \\ & + p0^0 \cdot p1^0 \cdot Sensor\_Analysis\_Z_v^2 \\ & + p0^0 \cdot p1^0 \cdot p2^0 \cdot Stream\_until\_Triggered_r^3 \\ & + p0^0 \cdot p1^0 \cdot p2^0 \cdot p3^0 \cdot p4^0 \cdot HCW\_Collector_r^4 \\ & + p0^0 \cdot p1^0 \cdot p2^0 \cdot p3^0 \cdot p4^0 \cdot p5^0 \cdot Results\_To\_UAV^5 \end{aligned} \quad (2)$$

Where $Sensor\_Z_v$ and $Sensor\_Analysis\_Z_v$ are VSA encodings of the particular object descriptions built as described in section VI. Multiple sensor chains, can be initiated by creating a VSA as shown in Eq. (3) and multi-casting it to a listening network.

$$\begin{aligned} Start\_Sensor\_Chains = {} & p0^0 \cdot (Senor\_Chain\_01 \\ & + Sensor\_Chain\_02 \\ & + Senor\_Chain\_03 + \cdots)^1 \end{aligned} \quad (3)$$

The multiple sensor chains continue independently until an anomaly is detected when each will unbind and activate its data collector. The $HCW\_Collector_r$ uses a bully algorithm [39, page 330] to merge multiple parallel streams into a single

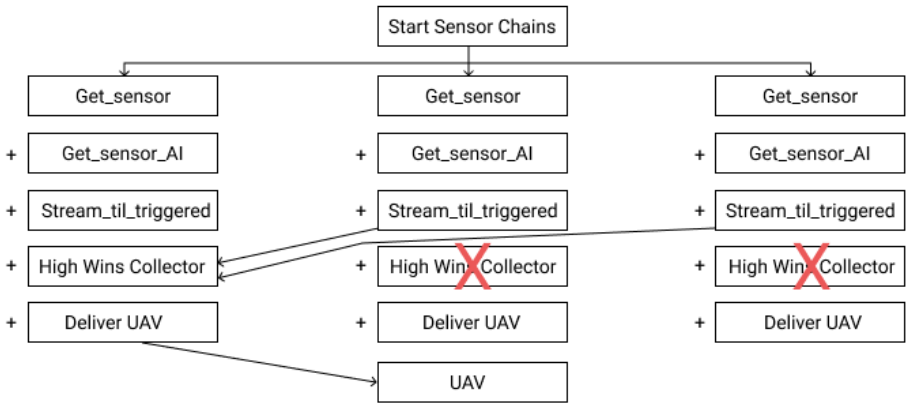

Fig. 4: VSA Workflow Graph

stream and subsequently only one $Results\_To\_UAV$ will be activated. Note that, during sensor chain recruitment each node participates in *local arbitration* as described in [3, Section 7.2, Page 79]. This enables each node to inspect the VC_S and VC_XP credentials of its partner nodes and select the best partner service with which to connect. Figure 4 shows a sketch of the VSA workflow graph.

Verifiable Credentials provide a mechanism for service users to attach reports about service use experiences to the services themselves, through VC_XP 'Experience Credentials'. These reports, which are expressed in JSON-LD formatted documents, are semantically searchable through the VSA mechanisms detailed above, and also provide additional contextual information as they are signed by the party that is making the claim. The identity of the signing party is expressed only as a decentralized identifier. This DID may or may not be known to parties checking the experience reports, and as such, a varying amount of regard can be given to it.

The use of VC_XPs allow services to be categorised and selected with a degree of granularity, where an absence of VC_XPs, means that no other party has left a signed report for the service, and where VC_XPs are available, configurations can be made such that priority is given to VC_XPs signed by close associates ahead of those signed by unknown parties. It can be envisaged that an ordering of service selection might be made where priority is given such that services with known VC_XPs are selected ahead of services with unknown VC_XPs, with services with zero VC_XPs chosen as a last resort. The selection field could vary depending on the urgency of the situation, and the quantity of resource required versus resource availability.

VIII. Conclusions and Future Research

Utilising semantic web technologies and open web standards allows service providers to describe their service offerings using interoperable data structures, which has the potential to improve service discovery and orchestration, whilst leveraging open source or municipal microservices where beneficial. Providing a means for service users to describe their actual experience with services provides an opportunity for trustworthy metadata to be added to service descriptions, backed by cryptographic assurance of the faithfulness of the party leaving it. Pseudonymous decentralized identifiers introduce an opportunity for parties to build networks of trust, such that they can develop policies to influence service selection based on fluid relationships with their peers. Converting these self-describing linked data structures into VSA vectors and building on previous work in semantic-based service discovery and orchestration via multicast service requests provides a pathway for efficient and flexible workflow construction, based on decentralized constructs. Potentially suitable service matches can be identified as a result of a semantic search, bringing a wider pool of services into consideration. The field can then be narrowed by policies which prioritise selection based on the availability of experience or quality of service reports from trusted partners, resulting in selection of the most suitable service to perform a particular task.

To provide a pathway to further research, the use case scenario presented in Section III will be enumerated through service descriptions based on JSON-LD documents, including WoT TD documents, and augmented with sample VCs to add context to service deployments and illustrative experience reports from peers. The linked data descriptions will be encoded as VSA vectors, and experiments will be conducted to determine the effectiveness of semantic service discovery based on linked data descriptions. If this proves successful, further experimental work will be designed to introduce service selection policies based on trust networks, and to understand the requirements for service configurations using this method, suitable for use on MANETs and 5G networks.

IX. Acknowledgements

This research was sponsored by the U.S. Army Research Laboratory and the U.K. Ministry of Defence under Agreement Number W911NF-16-3-0001. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Army Research Laboratory, the U.S. Government, the U.K. Ministry of Defence or the U.K. Government. The U.S. and U.K. Governments are au-

thorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation hereon.

REFERENCES

[1] C. Simpkin, I. Taylor, G. A. Bent, G. de Mel, and R. K. Ganti, "A scalable vector symbolic architecture approach for decentralized workflows," in *COLLA 2018 The Eighth International Conference on Advanced Collaborative Networks, Systems and Applications*. IARIA, 2018, pp. 21–27.

[2] C. Simpkin, I. Taylor, D. Harborne, G. Bent, A. Preece, and R. K. Ganti, "Dynamic distributed orchestration of node-red iot workflows using a vector symbolic architecture," in *2018 IEEE/ACM Workflows in Support of Large-Scale Science (WORKS)*. IEEE, 2018, pp. 52–63.

[3] C. Simpkin, I. Taylor, G. A. Bent, G. de Mel, S. Rallapalli, L. Ma, and M. Srivatsa, "Constructing distributed time-critical applications using cognitive enabled services," *Future Generation Computer Systems*, vol. 100, pp. 70–85, 2019.

[4] D. Verma, G. Bent, and I. Taylor, "Towards a distributed federated brain architecture using cognitive iot devices," in *9th International Conference on Advanced Cognitive Technologies and Applications (COGNITIVE 17)*, 2017.

[5] W. W. W. Consortium *et al.*, "Sparql 1.1 overview," 2013.

[6] D. Vachtsevanou, P. Junker, A. Ciortea, I. Mizutani, and S. Mayer, "Long-lived agents on the web: Continuous acquisition of behaviors in hypermedia environments," in *Companion Proceedings of the Web Conference 2020*, 2020, pp. 185–189.

[7] A. Ciortea, O. Boissier, and A. Ricci, "Engineering world-wide multi-agent systems with hypermedia," in *International Workshop on Engineering Multi-Agent Systems*. Springer, 2018, pp. 285–301.

[8] L. Sciullo, C. Aguzzi, M. Di Felice, and T. S. Cinotti, "Wot store: Enabling things and applications discovery for the w3c web of things," in *2019 16th IEEE Annual Consumer Communications & Networking Conference (CCNC)*. IEEE, 2019, pp. 1–8.

[9] M. Lanthaler and C. Gütl, "On using json-ld to create evolvable restful services," in *Proceedings of the Third International Workshop on RESTful Design*, 2012, pp. 25–32.

[10] T. Pham, G. Cirincione, A. Swami, G. Pearson, and C. Williams, "Distributed analytics and information science," in *In IEEE International Conference on Information Fusion (Fusion),*, 2015.

[11] C. Bizer, T. Heath, and T. Berners-Lee, "Linked data: The story so far," in *Semantic services, interoperability and web applications: emerging concepts*. IGI Global, 2011, pp. 205–227.

[12] S. Mayer, A. Ciortea, A. Ricci, M. I. Robles, M. Kovatsch, and A. Croatti, "Hypermedia to connect them all: Autonomous hypermedia agents and socio-technical interactions," *Internet Technology Letters*, vol. 1, no. 4, p. e50, 2018.

[13] N. Suri, M. Tortonesi, J. Michaelis, P. Budulas, G. Benincasa, S. Russell, C. Stefanelli, and R. Winkler, "Analyzing the applicability of internet of things to the battlefield environment," in *2016 international conference on military communications and information systems (ICMCIS)*. IEEE, 2016, pp. 1–8.

[14] D. E. Zheng and W. A. Carter, *Leveraging the internet of things for a more efficient and effective military*. Rowman & Littlefield, 2015.

[15] T. Berners-Lee, J. Hendler, and O. Lassila, "The semantic web," *Scientific american*, vol. 284, no. 5, pp. 34–43, 2001.

[16] A. Zschorn, H.-W. Kwok, and W. Mayer, "Microservice api design to support c2 semantic integration," 2019.

[17] J. R. Michaelis, M. Tortonesi, M. Baker, and N. Suri, "Applying semantics-aware services for military iot infrastructures," in *21st International Command and Control Research and Technology Symposium: C2 in a Complex Connected Battlespace*, 2016.

[18] T. S. Perry, "San diego's streetlights get smart," *IEEE Spectrum*, vol. 55, no. 1, pp. 30–31, 2018.

[19] F. Michel, C. Faron Zucker, O. Gargominy, and F. Gandon, "Integration of web apis and linked data using sparql micro-services—application to biodiversity use cases," *Information*, vol. 9, no. 12, p. 310, 2018.

[20] F. Michel, C. Faron-Zucker, O. Corby, and F. Gandon, "Enabling automatic discovery and querying of web apis at web scale using linked data standards," in *Companion Proceedings of The 2019 World Wide Web Conference*, 2019, pp. 883–892.

[21] V. Charpenay and S. Käbisch, "On modeling the physical world as a collection of things: The w3c thing description ontology," in *European Semantic Web Conference*. Springer, 2020, pp. 599–615.

[22] S. Bienz, A. Ciortea, S. Mayer, F. Gandon, and O. Corby, "Escaping the streetlight effect: Semantic hypermedia search enhances autonomous behavior in the web of things," in *Proceedings of the 9th International Conference on the Internet of Things*, 2019, pp. 1–8.

[23] C. Allen, "The path to self-sovereign identity," *URL: http://www.lifewithalacrity. com/previous/2016/04/the-path-to-selfsovereign-identity.html*, 2016.

[24] A. Tobin and D. Reed, "The inevitable rise of self-sovereign identity," *The Sovrin Foundation*, vol. 29, no. 2016, 2016.

[25] B. Preneel, "Cryptographic hash functions," *European Transactions on Telecommunications*, vol. 5, no. 4, pp. 431–448, 1994.

[26] M. Sporny, G. Noble, D. Longley, D. C. Burnett, and B. Zundel, "Verifiable credentials data model," November 2019. [Online]. Available: https://www.w3.org/TR/vc-data-model/

[27] M. Davie, D. Gisolfi, D. Hardman, J. Jordan, D. O'Donnell, and D. Reed, "The trust over ip stack," *IEEE Communications Standards Magazine*, vol. 3, no. 4, pp. 46–51, 2019.

[28] F. Wang and P. De Filippi, "Self-sovereign identity in a globalized world: Credentials-based identity systems as a driver for economic inclusion," *Frontiers in Blockchain*, vol. 2, p. 28, 2020.

[29] I. Barclay, S. Radha, A. Preece, I. Taylor, and J. Nabrzyski, "Certifying provenance of scientific datasets with self-sovereign identity and verifiable credentials," in *Proceedings of 12th International Workshop on Science Gateways*, 2020.

[30] P. C. Bartolomeu, E. Vieira, S. M. Hosseini, and J. Ferreira, "Self-sovereign identity: Use-cases, technologies, and challenges for industrial iot," in *2019 24th IEEE International Conference on Emerging Technologies and Factory Automation (ETFA)*. IEEE, 2019, pp. 1173–1180.

[31] I. S. Ahmad and B. Boufama, "Automatic vehicle identification through visual features," in *Proceedings of the 17th International Conference on Advances in Mobile Computing & Multimedia*, 2019, pp. 185–194.

[32] P. Kanerva, "Hyperdimensional computing: An introduction to computing in distributed representation with high-dimensional random vectors." *Cognitive Computation*, vol. 1, no. 2, pp. 139–159, 2009. [Online]. Available: http://dblp.uni-trier.de/db/journals/cogcom/cogcom1.html#Kanerva09

[33] G. E. Hinton, "Mapping part-whole hierarchies into connectionist networks," *Artificial Intelligence*, vol. 46, no. 1-2, pp. 47–75, 1990.

[34] C. Eliasmith, *How to build a brain: A neural architecture for biological cognition*. Oxford University Press, 2013.

[35] M. N. Jones and D. J. K. Mewhort, "Representing word meaning and order information in a composite holographic lexicon," *psychological Review*, vol. 114, no. 1, pp. 1–37, 2007.

[36] G. Recchia, M. Sahlgren, P. Kanerva, and M. N. Jones, "Encoding sequential information in semantic space models: comparing holographic reduced representation and random permutation," *Computational intelligence and neuroscience*, vol. 2015, p. 58, 2015.

[37] C. Eliasmith, T. C. Stewart, X. Choo, T. Bekolay, T. DeWolf, Y. Tang, and D. Rasmussen, "A large-scale model of the functioning brain," *Science*, vol. 338, no. 6111, pp. 1202–1205, Nov. 2012. [Online]. Available: http://www.sciencemag.org/content/338/6111/1202

[38] B. Francis, *Cameras, Sensors & What's Next For Mozilla'S Things Gateway*, 2019 (Accessed 11 August 2020). [Online]. Available: https://hacks.mozilla.org/2019/01/cameras-sensors-whats-next-for-mozillas-things-gateway

[39] M. Van Steen and A. S. Tanenbaum, *Distributed systems*. Maarten van Steen Leiden, The Netherlands, 2017.