Opening the Loops - Towards Semantic, Information-centric Networking in the Internet of Things

Hanno Wirtz, Klaus Wehrle
Chair of Communication and Distributed Systems, RWTH Aachen
{wirtz, wehrle}@comsys.rwth-aachen.de

Abstract—**The advent of the Internet of Things (IoT) paradigm in increasing deployments promises a pervasive proliferation of smart things, capable of sensing, actuating, and processing information. In typical designs, however, application of each thing is restricted to a dedicated use case in a single network of connected devices, resulting in a closed loop of information flow. We argue that, given the envisioned diversity, capabilities, and sheer number of smart things, this obstructs the possibility of creating diverse and exciting applications that benefit of the generated information in public, global usage scenarios.**

In this paper, we thus aim to initiate the discussion of creating a true *Internet* of Things, i.e., interconnected IoT networks, based on provision and requests of generated information in a public infrastructure. We highlight the challenges in designing this infrastructure for feasible integration in the current Internet and IoT designs, comprehensive provision and retrieval of information, and versatile derivation of higher information contexts from single information sources. Assessing advantages and shortcomings of existing approaches, we propose a suitable approach and discuss both a centralized and distributed implementation of the proposed infrastructure.

I. INTRODUCTION

The vision of an *Internet of Things* (IoT) comprises a diversity of use cases, deployment and application scenarios. Popular examples are active building [1] and factory automation [2], health [3] and energy metering [4] but also control of public infrastructure [5], street traffic [6], and environmental changes [7]. Especially the general use case of smart cities and smart building envisions heterogeneous use cases for networked "things" that both sense and report data as well as process data and act on the results.

In this vision, a number of smart things thus form a tailor-made *control loop* that is deployed and operated towards a specific, single purpose. As this vision is about to become a reality, triggering the installation of numerous such networks for widely heterogeneous purposes, we argue that the single-purpose use of single networks unnecessarily restricts the benefits of an existing Internet of Things. This is because information gathered in a specific context may prove valuable in another context, either as secondary information to supplement the primary information or in combination with other information to defer higher-layer contexts. For example, information gathered from monitoring street traffic and traffic lights by a municipality may aid steering the barriers of private and commercial parking lots to prevent accidents. Similar, private noise measurements, in combination with environmental pollution sensors, may serve as input to publicly available route planning based on current traffic conditions. Leveraging the sensing capabilities of mobile devices, such as smartphones, participatory sensing [8] will further increase the number and proliferation of ubiquitous information sources. The combination of such diverse information requires users, services, and even networks to be able to query sources *across specific deployments and administrative domains*, i.e., private, public, and commercial.

Currently however, IoT designs revolve around a *closed loop* network architecture, i.e., an isolated network of (wirelessly) communicating things (wireless sensor nodes, RFID-equipped devices, etc.) that is delimited by a dedicated, closed gateway device. Emphasis is placed on connecting devices and on establishing the use case-specific information flow within each loop. Accordingly, information, i.e., sensed and/or processed data, in these deployments flows from things to the gateway and an enclosed database or web service that is neither visible nor reachable from the public Internet[1]. While supporting the designated use case, e.g., to monitor the temperature or transportation of items in a factory, this restricts usage of the information to the operator of the network.

For a public, versatile, and wide-spread use, closed loop-designs restrict the potential of an IoT. This is because already single information, if publicly reachable, would be beneficial for more consumers than just the operating (public) institution. For example, things measuring the pressure on a bridge's car lanes allow the bridge operator to judge the stress imposed on the bridge infrastructure but would also allow travelers to judge the current traffic flow on the bridge when planning a route. Combined with data from similar application-specific, closed deployments such as noise, pressure, and pollution and (at night) light sensors in light poles in the city streets, such data would enable the creation of an up-to-date, public traffic report of the city. Extending single information by the combination of multiple specific sources, higher-context services can be tailored to, e.g., report the status of a route in timely fashion. Similar, such combinations offer the possibility of equivalent information derived from different sources, exploiting the multitude and diversity of sources for redundancy, accuracy, and comparison.

In this notion, however, the intuitive solution of accessing the digital representation of every thing directly, e.g., in the

---

[1] While a notable exception, the Microsoft SenseWeb project [9] mainly provides visualization of sensor data without providing an open infrastructure.

cloud, to then construct the requested information, is infeasible due to three main reasons. First, single data points, as delivered by, e.g., a pressure sensor, only provide very limited and fluctuating information, thus requiring a local and repeated selection of numerous things and iteration over the delivered data points, e.g., all the pressure sensors in a given car lane. Second, addressing things directly, e.g., via IPv6, only establishes connectivity but does not yield any information about the context and type of the provided information. Similar to current Internet structures, this requires an additional discovery layer that allows representing services and information as well as queries to formulate requests. Third, directly accessing each thing to query information breaches the integrity to the original network and entails the inherent risk of misuse.

In this paper, we thus raise the discussion on a infrastructure that supports the full potential of (re-)enabling IoT-generated information for higher contexts while retaining control over the information at the original information provider. To this end, we illustrate a typical urban scenario (Section II) and discuss the requirements and existing approaches for a public infrastructure that allows provision, retrieval, and combination of IoT information sources (Section III). We build on this discussion to propose the Internet Indirection Infrastructure (i3) [10] as a suitable candidate and sketch our design and additional functionality (Section IV) and implementation issues (Section V). Concluding our paper, we highlight future work possibilities (Section VI).

II. SCENARIO

In this paper, we depart from the problem space of creating and operating single networks of connected things operated by a specific, public or private, entity. Instead, we focus on the scenario of (creating) a true *Internet or Things*, i.e., an arbitrary large set of interconnected networks of smart things being operated by arbitrary entities. In this, we follow the prevalent vision of billions of smart things being produced and deployed in the next 5 – 10 years [11], [12]. Specifically, we emphasize the question of *how to fully exploit the benefits possible in an existing Internet of Things*.

To this end, we assume a (large) number of existing networks of things being operated in a distributed fashion. We further assume these networks to follow the current closed-loop network paradigm of a separated *network of things* (NoT), communicating, e.g., over 6LowPAN [13], [14], that is attached to one or more gateway devices. Gateway devices provide connectivity to the Internet, thereby establishing the principal possibility of communication between NoTs, and serve as a proxy between possibly differing protocol and routing suites in the NoT and the Internet. note that, albeit connected to the Internet via a gateway, we still regard this setting as a closed loop, since uninvolved users are not able to query the information created in the network due to missing credentials, communication parameters, or even knowledge of its existence.

Figure 1 shows a number of networks (NoTs) in an example scenario. We will use this (artificial) scenario as an illustrating setting throughout the paper. Please note that we only use a limited local example scenario to facilitate our use case illustration but otherwise discuss a global infrastructure. In our example, a bridge connecting two city parts is equipped with pressure sensors in the car lanes to measure the structural integrity and stress. The street leading from the bridge to one of the city parts is monitored by noise sensors to judge the stress level imposed on citizens, while the other part of the city employs air quality and temperature sensors to measure pollution.

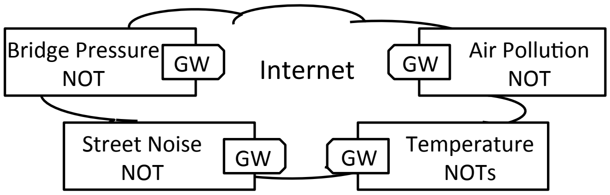

Fig. 1. Example scenario of separated Networks of Things (NoT) that could be combined to achieve higher context information in the presented scenario. Typically, each NoT constitutes a closed loop of information flow, prohibiting the exchange of information between NoTs or the combination of separated information sources outside of each loop, i.e., the creation of higher context loops.

For the sake of a general "big picture" discussion, we thereby make a number of abstractions and assumptions:

- **Internal network technology** We assume every single NoT to function using its own communication infrastructure. Things may thereby communicate wirelessly or be connected via some cable-bound network, using a standardized protocol or adaption layer, such as 6LowPAN [13], [14], a custom (sensor) network protocol, such as the Collection Tree Protocol (CTP) [15], or a proprietary protocol. We furthermore abstract from the number of gateway devices or sinks but assume that information provided by each things is reachable over the Internet, e.g., via the gateway device.
- **Network use case** We assume each network to be installed for one or more dedicated use cases, e.g., air pollution and noise monitoring, as mandated by the operator. However, we abstract from this specific use case and only regard and value the created information with regard to their semantic and context.
- **Internet routing protocol** We regard the Internet as a general routing substrate that allows communication between NoTs and arbitrary users but do not regard the specific routing or addressing mechanism, i.e., IPv4 or IPv6. Given the high degree of standardization, we assume gateway devices to provide protocol translation and proxy functionality between the existing routing variants.

Furthermore, we expect users searching for data and information to only know and subsequently specify *what* they search for, but not necessarily *where* or *how* they expect it to be provided. This assumption reflects on two prevalent observations from the current way of discovering and using information in Internet-based systems. First, analogous to the use of search engines in the Internet, memorization of exact location and addresses of sources is hard for human users. Second, the internals of providing an information, i.e., the exact logic, structure, sequence, and dependencies of processing and aggregating data sources, is irrelevant for users, which mainly care about information quality. From this arise both potentially fuzzy specifications of information queries, e.g., "traffic jam in <city_name>", as well as the possibility to gather results over multiple venues and processes.

III. INFRASTRUCTURE REQUIREMENTS AND DESIGN TRAITS

We argue that a suitable infrastructure needs to support three distinct functionalities: *Provision* of information needs to keep the single network operation and control over the information untouched to foster adoption and preserve the integrity of the respective network use cases. Furthermore, provision needs to reflect the information context to allow distinguishing multiple sources with varying degrees of separation. *Retrieval* of information then bases on the ability to match provided information context to the query specification as given by the user. Lastly, to alleviate the limited information content and specific use case of single sources, *aggregation and combination* of sources, i.e., the provided information, inside the infrastructure enables the creation of volatile or persistent higher context information.

However, we first address the question of how to address and model information (sources) in the envisioned scenario and infrastructure.

A. Information Modeling

Given the expected proliferation, number, diversity, and flexibility of networked things as information sources, describing or specifying information in such an infrastructure closely resembles the task of *context modeling*, as originally considered in ubiquitous computing and communication [16].In this, context comprises all relevant parameters describing an entity, in our case an information element, such as location, time, provider, etc. that are of relevance in providing and retrieving information. While concerned with context modeling in mobile, location- and time-dependent ubiquitous computing systems, uniformly representing context and specifying an associated query language are highly desirable properties in a more distributed, abstract infrastructure as required for information handling in the Internet of Things.

In [16] the respective authors evaluate existing modeling approaches with regard to different aspects required to support context modeling in ubiquitous computing. We argue that the regarded aspects[2] directly map to the requirements we identified for a suitable infrastructure in a versatile, public IoT infrastructure. Specifically, support for possible *incompleteness and ambiguity* as well as varying *richness and quality* of information provided by information sources (or NoTs) in context modeling approaches already accounts for the diversity of devices in information *provision*. Similar, building on adjustable *formality* in the representation, well-defined interfaces simplifies parsing of human-readable and potentially fuzzy *retrieval* queries. Subsequently, *partial validation* of queries and results against a given or expected context allow incomplete results as well as deriving a measure of result quality. Last, support by the modeling approach to incorporate *distributed composition* of information directly implements the envisioned possibility to *aggregate and combine* information sources.

In their evaluation [16] the authors identify *object oriented* and *ontology based* modeling approaches as the most suitable approaches to model, compare, and validate heterogeneous contexts[3]. Object oriented approaches follow the design traits of object oriented programming and relational databases and are therefore suited to encapsulate context information and structure access to content. Ontology based approaches provide equivalent modeling capabilities, as the instruments of concepts and facts in ontologies can be mapped to classes and instances. However, due to more fine-grained modeling of relations between concepts, ontology based approaches provide better support for reasoning on the exchangeability or similarity of contexts. Following this evaluation, we will discuss the applicability of these two general approaches with regard to provision, retrieval, and aggregation of information sources in the respective following sections.

B. Provision

In providing information sources to a public infrastructure, we expect the first priority of operators of NoTs to be the preservation of i) the security and integrity of the actual network and its resources, ii) the implementation of the original network use case, and iii) control over the respective information as well as external factors, e.g., network traffic. Coupled with the inherent resource constraints of smart things, this implies that the provision of information sources occurs on an Internet-bound host inside the operator's domain that is not part or crucial to the operation of the respective NoT. In current NoT designs, the gateway device offers a natural fit, albeit depending on the available resources in terms of processing and storage. Already connected to the NoT for protocol translation and, more importantly, managing and forwarding of sensed or actuated data, events, and notifications, processing the respective contained information for provision is a feasible task. In this, the transition from *data sources* in a closed network to *information sources* in a public infrastructure (figuratively) opens the loops of information flows.

Providing information sources in a way that allows for a comprehensive description of the parameters of an information sources again highlights the connection to the task of context modeling. This is because an information source that, for example, provides a temperature measurement loses its meaning unless augmented with context information, such as the location in GPS coordinates and time of the measurement. With regard to the proposed context modeling techniques [16], we argue that an object oriented representation of information is best suited for the task of provision.

A gateway device would thus gather data from things in the respective NoT, performing in this its unmodified current functionality. In addition to the respective use case, however, it would encapsulate the gathered data in *information objects* that carry the data and the associated context. Then, given a specific user request, this context information, e.g., "time = xx:yy" and "location = aa.bb cc.dd", can be compared to the request specification using well-defined object interfaces.

C. Retrieval

We regard the efficient search for relevant information sources as the main challenge in providing the envisioned infrastructure. While information objects allow an efficient comparison with the parameters set in the search request, the sheer number of envisioned information sources renders a brute-force iteration over all objects infeasible. For example,

---

[2] Aspect/requirement terminology as in [16].

[3] In a more recent survey paper [17], the authors derive similar conclusions.

a user constructing a temperature heat map of a given city would specify his request for temperature values along with area boundaries defined by GPS coordinates and a reference time frame. Processing of this request then requires finding the information objects that match the specified parameters of type, location, and time. In the following, we sketch a small number of possibilities to enable retrieval for information objects and identify associated problems.

A central entity, such as Google in the current Internet, that indexes retrievable information objects and processes user requests thereby is both the simplest and most expensive solution. While being able to provide a simple interface for retrieval as well as a single point of contact, this requires exhaustive resources to index and manage the envisioned number of devices and thereby transitively information objects. In addition, comprehensively disclosing information sources to a single entity entails the traditional risks of misuse, unlawful disclosure to others, and malicious data gathering.

Direct, human-readable addressing of resources, i.e., in this case information objects, has been proposed and realized in RESTful deployments [18], [19]. A careful combination of context parameters in URIs could thereby reflect the context diversity and depth of information objects, given a host URI that is known in advance or retrievable via a separate look-up structure. While rather static in terms of retrieval, the large body of existing RESTful functionality promises applicability and integration.

Catering to the emphasis on information rather than endpoint-based communication, data-oriented [20] and content-centric [21] networking offer direct addressing of information objects. Similar to RESTful approaches, derivable host URIs in content-centric networking [21], i.e., globally routable names, that allow a look-up of the providing device, e.g., via DNS, are required on top of a base routing layer such as IP. In contrast, data-oriented networking [20] alleviates the difficulties of host look-ups by using a route-by-name paradigm to *find* the relevant (and closest) *registered* copies of a requested item, denoted by the managing principal and a label describing the item. As the authors state, realization of the required resolution handlers that resolve such requests is possible in multiple ways, e.g., in a hierarchy similar to the current DNS system or distributed in a peer-to-peer fashion. In this, the inherent inclusion of locality in the retrieval of information objects as well as a publish-subscribe design promises a good match of the requirements we identified. When requesting information objects, specifying their labels but leaving the principal as a wildcard would allow a semantic search within specified parameters such as location and time.

Still, realization of the data-oriented (routing) architecture as well as resolution handlers requires, at least in parts, a "fork-lift upgrade" of the current Internet and routing structure [22], which has been proven to be the prohibitive point of such clean-slate approaches. In addition, designed for single pieces of content, no inherent support for the combination or aggregation of information objects exists *inside* the architecture. In Section IV, we sketch the proposal of a suitable approach.

D. Aggregation and Combination

Given the ability to provide (publish) an information source and to query it publicly, an interesting question is whether information retrieval needs to be limited to the original context of strictly specified sources, with derivation of higher contexts happening at end hosts. This question is motivated by three main aspects. First, the information information typically is *static* and *volatile*, i.e., a temperature or pressure sensor periodically provide a single, new measurement value, just as a camera provides a (still) picture. As such, information objects can not (and should not, given the need for integrity protection) represent the specific higher context of user interests, such as a temperature gradient over a given timespan. Second, end hosts deriving higher contexts, using local resources, have no means of publishing these contexts, i.e., the means of deriving them, to the public for others to benefit from, similar to closed information loops in NoTs. As an example, the combination of multiple information sources to derive the temperature heat map of a city or the detection of traffic jams, as established by a user, may directly benefit other users. Third, information sources may provide redundant information. This allows substitution of sources in case a specified source fails or context may be derived more efficiently or suitably via another set of sources, e.g., replacing camera stills to judge a bridge's traffic amount by a combination of pressure sensors.

This motivates the idea of *virtual sources*, information objects that build on a number ≥ 1 of information objects with a physical representation, by incorporating their information according to a specified functionality. The incorporated information objects may thereby contain diverse information and be provided by different entities. Similar, the functionality may be a near arbitrary piece of processing logic, with likely candidates being aggregation and filter functions to capture (relevant) values from (multiple) sources. Each "higher context" virtual source then exists in the infrastructure in the same way a "single context" source does, i.e., providing information along with a set of parameters describing its semantic context. This then allows the combination of virtual sources in other virtual sources in hierarchies.

Especially, higher context in virtual sources could allow the processing of fuzzy user requests, such as the aforementioned general request "Is there a traffic jam in <city_name>?". This is because multiple, diverse methods of providing information towards this request exist, from traffic camera pictures during the day and pressure sensors at night to a combination of air quality and noise information. To this end, multiple virtual sources could exist that serve a common purpose, achieving redundancy and quick failure recovery.

Such processing and management of higher layer context in virtual sources, instead of static, singular information objects, motivates *contextual reasoning* and *consistency checking* on the modeled context and information [16]. With regard to the possible modeling approaches, both object-oriented and ontology-based models provide such capabilities. In object-oriented approaches, inheritance relation between classes of information objects allow reasoning about the relation between actual, available instances as well as their consistency. Similar, the interrelation of concepts allows reasoning about facts in ontology-based approaches, which in addition may inherit the

information ID	*content_type:value*	···	*content_type:value*

(a) Generic structure.

*h(*temperature*)*	*location:gps_co*$_1$	*time:2013-03-30 12:00*

(b) Example request for temperature information from location <$gps\_coordinate_1$> at time $timestamp_1$.

Fig. 2. Proposed identifier layout in generic form and with example values.

inference capabilities of the ontology's underlying description language.

IV. i3 – A SUITABLE APPROACH?

In the previous sections, we identified the requirements and shortcoming of existing approaches. In this section, we thus propose to adapt the Internet Indirection Infrastructure (i3) [10] as an approach that could provide the required characteristics and serve as a viable basis for possible extensions. i3 provides a look-up and rendezvous structure that enables publishing abstract *(identifier, host)* tuples and subsequent communication by sending packets towards the respective identifier or *subscribing* to it. This indirection mitigates the difficulty of requiring a destination host look-up through the encoding of request semantics in identifiers. In the following, we illustrate how we envision i3 to enable the aforementioned functionalities and refer to the original paper [10] for a detailed design description. Figure 4 illustrates the overall system layout, in the course of this section, we detail the functionality we envision for each component.

i3 supports *provision* of information objects through publishing the respective identifier, with the host entry pointing to the IP of the providing device. In this, the objects and their content remain on the providing devices, allowing it to retain control over access to, updates of, and eventual removal of the information itself. To derive an identifier, we propose to encode the type of the information, i.e., the object class or concept, into an *information ID* field, e.g., using a hash function, which serves as the identifier. This information ID thereby represents a categorization of the information object with scalable details, e.g., "temperature" or "temperature <gps_coordinate>" (single context) or "traffic" (aggregated context). Context parameters, such as location or time, are then concatenated behind the information ID in a generic *(content_type:value)* encoding, where content type are again well defined classes, types, or facts in the model. Figure 2 illustrates this structure[4] while Figure 3(a) illustrates the provision of multiple information sources by a gateway.

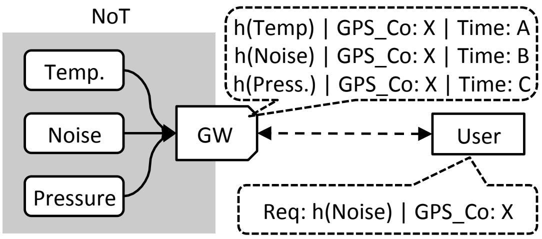

(a) Provision and retrieval of single information sources.

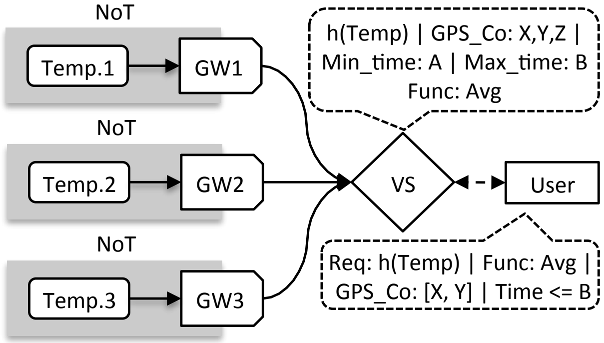

(b) Provision and retrieval of an aggregated, higher context information through a virtual source.

Fig. 3. Example interaction between a requesting user, information sources inside a NoT, the providing gateway device (GW) and a virtual source (VS).

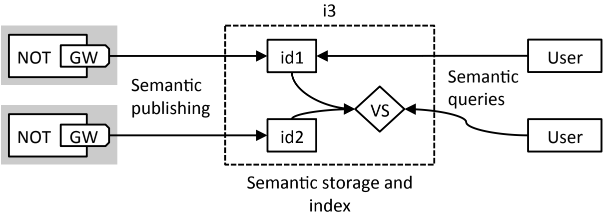

Fig. 4. Example layout of using i3 as an indexing and storage approach. Previously closed loops (NOTs) are able to semantically publish (and remove) identifiers (id$_i$) that point to information sources without compromising the original NOT use case or network integrity. Within the indexing structure, single sources may be combined to higher-context virtual sources (VS). Users are then able to query either single sources directly or query aggregated information sources in a VS.

To match a request to a provided object for *retrieval*, we propose to follow the authors' original proposal [10] of *longest prefix matching*, while allowing fuzzy user requests through wildcards. A user would thereby specify a request and derive an information ID that all information objects need to match. Given further context parameters in the request, e.g., a GPS coordinate and a time, these can be compared against the available objects' context parameters. Figure 3(a) illustrates the case of an exact match with regard to the location, while the time value is disregarded. Wildcards thereby cover fuzzy or incomplete specifications in the request, we assume information providers to set all available context parameters. Since i3 envisions tuples with the same identifier, i.e., information ID in our case, to be stored at the same host, no wide dissemination of the request is necessary, saving routing effort. The best matching tuple is subsequently returned to the requester, enabling him to retrieve the information object. To enable parametrization and optimization of comparisons in searches, context parameters may be augmented by priority and exactness scores that specify to which extent each parameter needs to be matched by an object, e.g., in case of a location denoted by a GPS coordinate.

To *aggregate and combine* information objects in a virtual source, we envision to use i3's *service composition* and *multicast* functionality. Service composition allows to specify the

---

[4] The original dimensioning of identifiers $\leq 256\,$bit would need to be expanded to accommodate context parameters etc.

identifier of (multiple) services that need to be called before the result is delivered. In our case, a user would request a virtual source by its information ID, which in turn points to the identifier(s) of (multiple) information objects that need to be retrieved in order to process their information according to the virtual source's functionality. Figure 3(b) illustrates the aggregation of temperature values from multiple sources in a virtual source. Please note that i3, in its original design, does not envisage the association of functionality with tuples, i.e., virtual sources, and will need to be extended in this regard. Furthermore, service composition is designed to occur sequentially, requiring a virtual source to specify a chain of objects and the transport of subsequent information. In order to support service composition in parallel, we envision to extend service composition by adapting i3's multicast functionality to allow parallel requests along the multicast tree.

Based on the information ID of a request, along with context parameters and priority and exactness scores, *reasoning* about the similarity and relation of information objects and virtual sources becomes possible. Requested objects and sources may thus be exchanged in case of failure, given an appropriate similarity, or combined with other sources to achieve adequate similarity.

V. IMPLEMENTATION ISSUES OR IS THERE A BUSINESS MODEL?

The discussed infrastructure would, next to providing the basic indirection infrastructure proposed by i3 [10], allow for the creation, processing, and management of virtual sources of varying complexity. Similar to content delivery networks (CDNs), this raises questions about the cost, performance, availability, and reliability of such an infrastructure. In this section, we thus briefly compare the two apparent venues of implementing the envisioned infrastructure, namely *centralized* or *distributed*.

In their original design and evaluation, the authors of i3 assumed a distributed hash table (DHT), namely Chord [23], as the substrate "below" i3. The DHT thereby is responsible for managing the distribution of identifiers, and thus storage of tuples, among the nodes in the DHT as well as performing routing in the identifier name space. To construction and maintain the DHT, all providing devices, i.e., the gateway devices, would manage a share of the name space as well as the associated tuples. Cooperative operation of the infrastructure would distributes the costs and effort among all participants, while redundancy mechanisms in the underlying DHT, e.g., as outlined in [23], would afford a measure of redundancy and thereby reliability and availability. However, given the heterogeneous resources of providing devices as well as their primary use case in the original network, partially poor performance and node churn as well as high latencies have to be expected.

In contrast to a distributed operation, a centralized provider could follow a business model similar to the Apple App Store and Google Play. For example, the provider could manage the indirection infrastructure in a CDN-like network free of charge and offer commercial services for the creation, persistent storage, distribution, and advertisement of virtual sources, similar to current smartphone or desktop applications. The resource requirements of storage and processing could thus be met while guaranteeing high availability, reliability, and performance. Processing user requests in a centralized infrastructure could furthermore, through (machine) learning approaches and quality monitoring, enable a better processing of fuzzy user requests.

VI. CONCLUSION AND OUTLOOK

In this paper, we raised the discussion of networking single deployments in the ascending Internet of Things. As our main focus, we identified the benefits and challenges in providing and controlling information generated by arbitrary smart things to a public IoT outside of the closed loop of current network designs. In this, our proposal for a public provision, retrieval, and processing infrastructure revolves around i) modeling the respective information on a global scale of billions of information sources, emphasizing the respective information context, ii) finding and retrieving information given a user request, with regard to the advantages and shortcomings of existing approaches, and iii) processing singular information sources to derive higher context for public reuse and improvement.

We argue that the Internet Indirection Infrastructure meets the identified requirements and proposed functional extensions to support processing of (combined) information sources. Assessing both a distributed and a centralized implementation of the envisioned infrastructure, we highlighted the respective possibilities and advantages. We plan to specify our design and realize the described infrastructure in a local, distributed, and cloud-based implementation to assess both the feasibility and performance of our approach. In this, we are eager to incorporate feedback, suggestions, and hints arising from the reviews and possible discussions at the workshop, as we still are in the early stages of our design. Incorporating heterogeneous smart things, ranging from sensor nodes to smartphones, will thereby require a comprehensive design that accounts for the diversity of today's information sources.

ACKNOWLEDGEMENTS

This work was funded in part by the DFG Cluster of Excellence on Ultra High-Speed Mobile Information and Communication (UMIC) and by the NRW Ziel 2 programme and EU EFRE.

REFERENCES

[1] M. Jung, C. Reinisch, and W. Kastner, “Integrating building automation systems and ipv6 in the internet of things,” in *Innovative Mobile and Internet Services in Ubiquitous Computing (IMIS), 2012 Sixth International Conference on*, July, pp. 683–688.

[2] S. Karnouskos, O. Baecker, L. de Souza, and P. Spiess, “Integration of soa-ready networked embedded devices in enterprise systems via a cross-layered web service infrastructure,” in *Emerging Technologies and Factory Automation, 2007. ETFA. IEEE Conference on*, Sept., pp. 293–300.

[3] D. Niyato, E. Hossain, and S. Camorlinga, “Remote patient monitoring service using heterogeneous wireless access networks: architecture and optimization,” *IEEE J.Sel. A. Commun.*, vol. 27, no. 4, pp. 412–423, May 2009.

[4] C. Wei and Y. Li, “Design of energy consumption monitoring and energy-saving management system of intelligent building based on the internet of things,” in *Electronics, Communications and Control (ICECC), 2011 International Conference on*, Sept., pp. 3650–3652.

[5] R.-G. Lee, K.-C. Chen, S.-S. Chiang, C.-C. Lai, H.-S. Liu, and M.-S. Wei, "A backup routing with wireless sensor network for bridge monitoring system," in *Proceedings of the 4th Annual Communication Networks and Services Research Conference*, ser. CNSR '06. Washington, DC, USA: IEEE Computer Society, 2006, pp. 157–161. [Online]. Available: http://dx.doi.org/10.1109/CNSR.2006.5

[6] J. Yoon, B. Noble, and M. Liu, "Surface street traffic estimation," in *Proceedings of the 5th international conference on Mobile systems, applications and services*, ser. MobiSys '07. New York, NY, USA: ACM, 2007, pp. 220–232. [Online]. Available: http://doi.acm.org/10.1145/1247660.1247686

[7] A. Mainwaring, D. Culler, J. Polastre, R. Szewczyk, and J. Anderson, "Wireless sensor networks for habitat monitoring," in *Proceedings of the 1st ACM international workshop on Wireless sensor networks and applications*, ser. WSNA '02. New York, NY, USA: ACM, 2002, pp. 88–97. [Online]. Available: http://doi.acm.org/10.1145/570738.570751

[8] J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, and M. B. Srivastava, "Participatory sensing," in *Workshop on World-Sensor-Web (WSW)*, 2006.

[9] A. Kansal, S. Nath, J. Liu, and F. Zhao, "Senseweb: An infrastructure for shared sensing," *IEEE MultiMedia*, vol. 14, no. 4, pp. 8–13, Oct. 2007. [Online]. Available: http://dx.doi.org/10.1109/MMUL.2007.82

[10] I. Stoica, D. Adkins, S. Zhuang, S. Shenker, and S. Surana, "Internet indirection infrastructure," in *Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications*, ser. SIGCOMM '02. New York, NY, USA: ACM, 2002, pp. 73–86. [Online]. Available: http://doi.acm.org/10.1145/633025.633033

[11] Ericsson White Paper, "More than 50 billion connected devices," [Online] http://www.ericsson.com/res/docs/whitepapers/wp-50-billions.pdf, 2011.

[12] Cisco White Paper, "The Internet of Things – How the Next Evolution of the Internet Is Changing Everything," [Online] http://www.cisco.com/web/about/ac79/docs/innov/IoT_IBSG_0411FINAL.pdf, 2011.

[13] J. Hui and D. Culler, "Extending IP to Low-Power, Wireless Personal Area Networks," *Internet Computing, IEEE*, 2008.

[14] G. Montenegro, N. Kushalnagar, J. Hui, and D. Culler, "Transmission of IPv6 Packets over IEEE 802.15.4 Networks," RFC 4944, IETF, 2007.

[15] O. Gnawali, R. Fonseca, K. Jamieson, D. Moss, and P. Levis, "Collection tree protocol," in *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*, ser. SenSys '09. New York, NY, USA: ACM, 2009, pp. 1–14. [Online]. Available: http://doi.acm.org/10.1145/1644038.1644040

[16] T. Strang and C. Linnhoff-Popien, "A context modeling survey," in *In: Workshop on Advanced Context Modelling, Reasoning and Management, UbiComp 2004-The Sixth International Conference on Ubiquitous Computing*, 2004.

[17] C. Bettini, O. Brdiczka, K. Henricksen, J. Indulska, D. Nicklas, A. Ranganathan, and D. Riboni, "A survey of context modelling and reasoning techniques," *Pervasive and Mobile Computing*, vol. 6, no. 2, pp. 161–180, 2010.

[18] A. P. Castellani, N. Bui, P. Casari, M. Rossi, Z. Shelby, and M. Zorzi, "Architecture and protocols for the internet of things: A case study," in *Pervasive Computing and Communications Workshops (PERCOM Workshops), 2010 8th IEEE International Conference on*. IEEE, 2010, pp. 678–683.

[19] W. Colitti, K. Steenhaut, and N. De Caro, "Integrating wireless sensor networks with the web," in *Workshop on Extending the Internet to Low power and Lossy Networks (IP + SN)*, 2011.

[20] T. Koponen, M. Chawla, B.-G. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker, and I. Stoica, "A data-oriented (and beyond) network architecture," in *Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications*, ser. SIGCOMM '07. New York, NY, USA: ACM, 2007, pp. 181–192. [Online]. Available: http://doi.acm.org/10.1145/1282380.1282402

[21] V. Jacobson, D. K. Smetters, J. D. Thornton, M. F. Plass, N. H. Briggs, and R. L. Braynard, "Networking named content," in *Proceedings of the 5th international conference on Emerging networking experiments and technologies*, ser. CoNEXT '09. New York, NY, USA: ACM, 2009, pp. 1–12. [Online]. Available: http://doi.acm.org/10.1145/1658939.1658941

[22] P. Jokela, A. Zahemszky, C. Esteve Rothenberg, S. Arianfar, and P. Nikander, "Lipsin: line speed publish/subscribe inter-networking," in *Proceedings of the ACM SIGCOMM 2009 conference on Data communication*, ser. SIGCOMM '09. New York, NY, USA: ACM, 2009, pp. 195–206. [Online]. Available: http://doi.acm.org/10.1145/1592568.1592592

[23] I. Stoica, R. Morris, D. Karger, M. F. Kaashoek, and H. Balakrishnan, "Chord: A scalable peer-to-peer lookup service for internet applications," in *Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications*, ser. SIGCOMM '01. New York, NY, USA: ACM, 2001, pp. 149–160. [Online]. Available: http://doi.acm.org/10.1145/383059.383071