Decentralized Stream Reasoning Agents

Gözde A. Tataroğlu Özbulak
gozde.tatarogluozbulak@hevs.ch
University of Applied Sciences and Arts Western Switzerland HES-SO
Sierre, Switzerland
Supervised by: Jean-Paul Calbimonte

ABSTRACT

This PhD project proposes the theoretical and technological foundations of an approach for decentralized processing of streaming knowledge graphs, where autonomous reasoners may combine individual and collective processing of continuous data. These decentralized stream processors shall be capable of sharing not only data stream knowledge, but also processing duties, using collaboration and negotiation protocols. Moreover, commonly agreed semantic vocabularies will be used to address the high dynamicity of reasoners' knowledge and goals. The approach proposed in this project goes beyond previous works on stream reasoning, enabling the self-organization and coordination among distributed stream reasoners, based on techniques and principles inspired by Multi-Agent systems. On the one hand, it adds the ability to explicate processing goals, capabilities and knowledge, while on the other it exploits potential ways of interconnecting them in ways that expand their combined capacity/efficacy for managing highly dynamic flows of streaming knowledge. Through this approach, efficient local stream processors can establish cooperative processing schemes, respecting data privacy restrictions and data locality requirements through the exchange of streaming Knowledge Graphs.

CCS CONCEPTS

• **Computer systems organization → Embedded and cyber-physical systems**; **Real-time system architecture**; • **Information systems → Semantic web description languages**; • **Software and its engineering → Interoperability**.

KEYWORDS

stream reasoning, RDF streams, Web Agents, Decentralized reasoning

ACM Reference Format:
Gözde A. Tataroğlu Özbulak. 2023. Decentralized Stream Reasoning Agents. In *Proceedings of 17TH ACM International Conference on Distributed and Event-Based Systems Doctoral Symposium (DEBS 2023).* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3583678.3603286

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
DEBS '23, June 27–30, 2023, Neuchatel, Switzerland
© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0122-1/23/06…$15.00
https://doi.org/10.1145/3583678.3603286

1 PROBLEM STATEMENT

Data streams are increasingly used to represent highly dynamic information in different domains, including eHealth, automation, transportation, and energy management [18]. In these environments, the challenge of structural heterogeneity is combined with the necessity of providing instantaneous processing and continuous outcomes for high data volume and velocity. In particular, we identify the challenge of handling data streams from virtual and IoT devices, which are inherently decentralized and may be distributed over geographic and organizational boundaries. Indeed, streaming data access is often restricted due to privacy concerns (e.g., in the healthcare domain), while local processing can also be needed in edge devices with real-time response requirements. Knowledge management techniques have been long used to address heterogeneity, even though they generally lack the capability of managing dynamic data flows in decentralized settings. Although the area of stream reasoning has advanced towards continuous processing of semantic streams in the last decade [17], there is an important research gap regarding decentralized reasoning over data streams, and even more so for organization and cooperation mechanisms among different stream processors. Even though there have been recent efforts towards semantization of streaming data in the form of dynamic Knowledge Graphs [8, 21], current models and implemented systems lack the ability to manage and reason over rapidly changing knowledge in a highly distributed and fully decentralized environment. Furthermore, stream reasoners can hardly be combined in order to share computing duties or aggregate streaming data results, given the diversity of their underlying processing models, and their lack of orchestration capabilities.

In this PhD project we will address these challenges, and we expect the results to open new research chapters in stream reasoning for knowledge graphs, with multiple applications specially in the domains of IoT and eHealth. Given the necessity of isolating data and computation due to risks in data privacy and protection, as well as considering the decentralized nature of the Web, the approach taken in this project is part of an increasing trend in autonomous computation and reasoning.

2 RESEARCH APPROACH/METHODOLOGY

To address the identified problem, we will adopt an iterative research approach, including the phases of planning, acting, observing, and reflecting. More specifically, we propose addressing the project challenges from the conceptual/theoretical, architectural, and application perspectives. The overall idea is to start extending existing stream reasoning models with the decentralization capabilities and their formalization. In particular, we will base the model on the definitions specified by the RSP Community Group, and the formalization of RDF-star [22] for streams. Afterwards, we will define the interaction model and possible cooperation strategies, formalized through agent-based protocols, such as multi-objective negotiation. Then, the architectural aspects can be addressed, leading to potential implementations that inherently include the distributed nature of streaming Knowledge Graphs. Finally, an evaluation and validation of this approach will be performed, considering performance, expressiveness, and applicability in a real-life scenario. Given the overwhelming importance of streaming data in the IoT domain, we will implement a proof-of-concept in this area, more specifically for supporting physical therapists through activity monitoring devices. In this scenario, the data privacy and data locality requirements will need to be addressed by the decentralized stream processing solution.

3 RESEARCH SETTING

The identified research problem has led to the formulation of the following research questions, which will be studied during the PhD. This refers first to **stream reasoning models**, and *How these can be extended in order to support decentralized continuous processing of entailments?*, and *What are the semantics of decentralized reasoning performed over streaming Knowledge Graphs?*. From this point it is also pertinent to understand *What level of expressiveness can be supported in decentralized stream reasoning tasks?*. Once such a model is defined, it is fundamental to know how **stream reasoners can interact**, and understand *What languages and interaction protocols are needed in order to specify goals, behaviors, and affordances of decentralized stream processors?*. Moreover, we can study if *stream reasoners can be interconnected in order to provide coordinated reasoning capabilities?*, and *What strategies can decentralized reasoners implement for cooperation and negotiation in stream reasoning tasks?*

From a more technological perspective, it is key to understand **decentralized stream reasoning architectures**, and *What types of architectures can be used to implement a framework for streaming Knowledge Graph reasoning?*, if *existing stream reasoning engines be adapted to be integrated into a distributed and decentralized reasoning framework?*, and *How can decentralized stream reasoners be incorporated to autonomous IoT devices, to support the generation of dynamic knowledge from personal sensor data?*

Based on these research questions, we formulate the following general objective: *Study, conceive and implement a model for collaborative and decentralized stream processing considering both ontology-based and AI-based stream processors interacting through streaming Knowledge Graphs.* The specific research objectives derived form the general one are:

- **O1** Study and conceive a model for decentralized stream reasoning over dynamic Knowledge Graphs, including the formalization of its underlying semantics.
- **O2** Study and define a language and protocol for establishing interactions among decentralized stream processors.
- **O3** Study and develop cooperation strategies and algorithms that enable negotiation among stream reasoners to perform coordinated tasks.
- **O4** Design, implement and validate an architecture for the deployment of decentralized stream reasoners based on the model and interaction protocols defined previously.
- **O5** Evaluate the application of the proposed architecture for decentralized streaming Knowledge Graphs in a Web of Things environment, and in particular for supporting physical therapy professionals as a use-case.

The fundamental novelty of our approach is embedded in the streaming nature of highly dynamic Knowledge Graphs, which are increasingly made available on the Web via IoT devices and social networks. Unlike other types of Knowledge Graphs, these streams are potentially unbounded, time-dependent, and may be produced (and consumed) continuously at high velocity, requiring different querying, reasoning, and interaction strategies.

4 RELATED WORK

The project topic lies at the intersection of different areas, including Knowledge Graphs, semantic data management, stream reasoning, and distributed reasoning.

Knowledge Graphs. Knowledge Graphs emerged in recent years as a flexible yet powerful paradigm for representation of information, including rich semantics and complex relationships among conceptual entities [19]. Different modelling approaches have been proposed in the literature, such as property graphs, or direct edge-labelled graphs as RDF, with a considerable level of adoption in the industry (e.g. implementations by Google, Facebook, Amazon, LinkedIn, among others [29]). Although most knowledge graph technologies focus on stored and static data, approaches for handling streaming graphs (e.g. RDF streams) have also been proposed, typically limited to continuous querying [3, 10, 30].

Distributed and decentralized knowledge graphs. The availability of Knowledge Graphs on the Web has opened a window for enhanced interoperability and the potential for information exchange across disciplines and organizational boundaries. Nevertheless, the acknowledgment of the decentralized nature of the Web often clashes with the typically top-down approaches for data federation [24] and integration in Knowledge Graphs [28, 34]. Up to now, even though knowledge graphs on the Web are published independently, there is still a lack of models that allow going beyond a traditional top-down approach to query processing and reasoning, and even less for the case of streaming knowledge.

Stream Reasoning. The rich semantics exposed by ontologies underlying Knowledge Graphs naturally brings the possibility of computing entailments, following reasoning mechanisms explored, for example in the context of Description Logics and Datalog[5]. However, when data and knowledge are highly dynamic, i.e., when produced or generated in the form of data streams, traditional approaches like incremental reasoning need to be adapted [35]. Stream reasoning emerged in the last decade to bridge this gap, producing a number of contributions including continuous data querying for RDF streams [30], semantic complex event processing [1], incremental maintenance of materialization [31], concurrent query processing [26], answer-set-programming for streams [4], or online inductive analysis [2].

Decentralized reasoning for rapidly changing information. Decentralized data management is a key aspect of the Web that has not been sufficiently explored, even if it has been present since its introduction by Tim Berners Lee [7], and further described in his Semantic Web vision [6]. In the case of data streams, the need for local decisions is even more relevant, as IoT and smart devices are capable of performing stream analytics, while also requiring to establish cooperation mechanisms with other devices. In the Web of Things (WoT) domain, new specifications such as the Things Description W3C Recommendation [13, 20] constitute a first step towards achieving this goal, although they do not address the question of negotiation, cooperation or self-organization, nor the implications related to reasoning over decentralized streaming knowledge.

5 RESEARCH APPROACH

To address the research problem and the derived research questions the general approach will be conducted as follows.

To achieve **O1** we will start with the formalization of the decentralized stream reasoning model for knowledge graphs. As seen in the literature, most of these works target streaming query federation at most, and do not address the problem of modeling reasoning among autonomous streaming engines. The proposed novel model will consider autonomous reasoners that may operate over knowledge graphs, or even streaming ML machines exposed through virtual RDF streams. The model is expected to extend the formalization of RDF-star for streams [22], incorporating elements from previous work on operational semantics [15], correctness [16]. This model is expected to cover point-in time semantics for RDF graphs, excluding temporal range semantics. Semantics will be specified for virtual RDF stream graphs, following the principles of ontology-based data access [10, 23]. Moreover, the model will define the possible entailment regimes, which might differ if it is a native or virtual RDF stream.

The definition of this model will serve then as a basis for defining interaction protocols, following **O2**. Interaction protocols will specify which directives the stream reasoning agents will use to communicate among them. For instance, they will allow the discovery of stream according to their metadata (e.g., using vocabularies for data streams [33]) or according to their affordances (e.g., Things description [13]). Protocols will also allows indicating stream reasoners to join a cooperative network, or offering reasoning and continuous query capabilities. These interaction protocols will be inspired by agent-based primitives [9, 11], with extensions indicating stream reasoners' affordances and capabilities [14], as well as goals and (streaming) data needs. The protocol will also specify if a processor produces materialized streaming entailments, or if it produces virtual streams, generated through Streaming ML. The entailment regime, ontology complexity, and expressiveness may also be indicated, and the protocol will provide the necessary primitives for reasoners to initiate coordinated actions.

Different strategies for stream reasoners' interaction and cooperation will be studied and designed, in order to fulfill **O3**. In particular, these protocols will use agent-based languages for specifying communication primitives or BSPL [32]. These primitives will allow establishing multi-agent negotiation patterns, e.g., using cost and goal prediction models in order to compute potential benefits [25]. Moreover, strategies may encompass negotiations regarding the stream velocity, the reasoning expressiveness, the notification of results, etc. Also, processors may include computing costs and requirements as part of the negotiation information (e.g., Streaming ML training overhead, or incremental materialization complexity). Considering the eHealth/IoT use case scenario of the project, we will focus on negotiation criteria relevant to this domain, in which data quality, computational and memory constraints, data privacy, and responsiveness are of particular interest, and may guide the self-organization decisions of stream reasoners.

6 EVALUATION PLAN

The implementation and evaluation of the aforementioned streaming model, as well as the interaction protocols and strategies will contribute to attain **O4**. A system for decentralized stream reasoning will be developed as a result, designed to cope with discovery of new stream reasoners within a network (e.g. using semantic vocabularies like Vocals [33]), It will also use the interaction protocols to declare capabilities, affordances, goals, etc., and initiate negotiation and coordination exchanges, thus enabling streaming knowledge sharing, distributed processing, or establishing reasoning pipelines. This implementation will be specifically targeted towards Web of Things scenarios where reasoners ingest streams from smart devices through virtual or materialized knowledge graphs [27]. The implementation will include the technical evaluation and study of optimizations for the agent-based interactions and negotiations. Specifically, we will consider metrics relative to stream reasoning and querying throughput, as well as agent-level goal fulfillment after negotiation.

Finally, **O5** will be achieved through a validation of our system implementation in the context of a physical therapy environment. To perform this validation, the team will collaborate with the PhysioLab of the Institute of Health of HES-SO Valais-Wallis. Following previous joint work in the context of physical exercise for at-risk people at home [12], we will use a combination of wearable sensor data streams with live IR video recording data that extracts human skeleton movements. These measurements are used of r exercise assessment, automatic assessment and feedback, and thus requiring real-time responses, continuous processing and decentralized management of IoT data with privacy guarantees. In this scenario, individual stream processors may jointly coordinate based on data observed on a single person (e.g. a balance and strength sensors), while at higher level certain outputs may be streamed back to another reasoners under control of the therapist. We will evaluate the feasibility of the entire approach, and at a specific level evaluate response time and feedback, as well as efficacy of the composition of decentralized stream processors.

7 CONCLUSIONS AND REFLECTIONS

As this PhD project has just started, we are currently only at the beginning of the planned work. adaptation may be needed depending on the findings of the first iterations, as explained in the methodology. In summary, we will continue with the described tasks, starting with the formalization of a model for decentralized stream processing agents, capable of coordinated reasoning over streaming

Knowledge Graphs. Then we will study and propose an interaction language and protocols to implement negotiation and coordination among stream processors. Afterwards we will conceive, design and implement a framework that follows the proposed decentralized stream processing model, incorporating Knowledge Graphs for both the negotiation and cooperation protocols. Moreover, we will validate the approach and its implementation through experimental evaluation, including both synthetic Knowledge Graph streams and existing benchmarks for IoT and stream processing. Finally, we will explore the impact of this approach through a proof-of-concept implementation in the domain of eHealth for activity assessment.

ACKNOWLEDGMENTS

This PhD project is supported by the Swiss National Science foundation project StreamDKG.

REFERENCES

[1] Darko Anicic, Paul Fodor, Sebastian Rudolph, and Nenad Stojanovic. 2011. EP-SPARQL: a unified language for event processing and stream reasoning. In *Proceedings of the 20th International Conference on World Wide Web, WWW 2011, Hyderabad, India, March 28 - April 1, 2011*. ACM, 635–644. https://doi.org/10.1145/1963405.1963495

[2] Davide Barbieri, Daniele Braga, Stefano Ceri, Emanuele Della Valle, Yi Huang, Volker Tresp, Achim Rettinger, and Hendrik Wermser. 2010. Deductive and Inductive Stream Reasoning for Semantic Social Media Analytics. *IEEE Intelligent Systems* 25, 6 (Nov. 2010), 32–41. https://doi.org/10.1109/MIS.2010.142

[3] Davide Francesco Barbieri, Daniele Braga, Stefano Ceri, Emanuele Della Valle, and Michael Grossniklaus. 2009. C-SPARQL. (2009). https://doi.org/10.1145/1526709.1526856

[4] Hamid R. Bazoobandi, Harald Beck, and Jacopo Urbani. 2017. Expressive Stream Reasoning with Laser. In *The Semantic Web – ISWC 2017 (Lecture Notes in Computer Science)*. Springer International Publishing, Cham, 87–103. https://doi.org/10.1007/978-3-319-68288-4_6

[5] Luigi Bellomarini, Davide Benedetto, Georg Gottlob, and Emanuel Sallinger. 2020. Vadalog: A modern architecture for automated reasoning with large knowledge graphs. *Information Systems* (may 2020), 101528. https://doi.org/10.1016/j.is.2020.101528

[6] Tim Berners-Lee, James Hendler, and Ora Lassila. 2001. The Ssemantic Web. *Scientific American* 284, 5 (2001), 34–43. https://www.jstor.org/stable/26059207

[7] T. J. Berners-Lee. 1992. The world-wide web. *Computer Networks and ISDN Systems* 25, 4 (Nov. 1992), 454–459. https://doi.org/10.1016/0169-7552(92)90039-S

[8] Pieter Bonte, Riccardo Tommasini, Filip De Turck, Femke Ongenae, and Emanuele Della Valle. 2019. C-Sprite. (jun 2019). https://doi.org/10.1145/3328905.3329502

[9] Jean-Paul Calbimonte, Davide Calvaresi, and Michael Schumacher. 2018. Multi-agent Interactions on the Web Through Linked Data Notifications. In *Multi-Agent Systems and Agreement Technologies*. Springer International Publishing, Cham, 44–53. https://doi.org/10.1007/978-3-030-01713-2_4

[10] Jean-Paul Calbimonte, Oscar Corcho, and Alasdair J. G. Gray. 2010. Enabling Ontology-Based Access to Streaming Data Sources. In *The Semantic Web – ISWC 2010 (Lecture Notes in Computer Science)*. Springer, Berlin, Heidelberg, 96–111. https://doi.org/10.1007/978-3-642-17746-0_7

[11] Davide Calvaresi and Jean-Paul Calbimonte. 2020. Real-Time Compliant Stream Processing Agents for Physical Rehabilitation. *Sensors* 20, 3 (Jan. 2020), 746. https://doi.org/10.3390/s20030746

[12] Davide Calvaresi, Jean-Paul Calbimonte, Enrico Siboni, Stefan Eggenschwiler, Gaetano Manzo, Roger Hilfiker, and Michael Schumacher. 2021. EREBOTS: Privacy-Compliant Agent-Based Platform for Multi-Scenario Personalized Health-Assistant Chatbots. *Electronics* 10, 6 (Jan. 2021), 666. https://doi.org/10.3390/electronics10060666

[13] Victor Charpenay and Sebastian Käbisch. 2020. On Modeling the Physical World as a Collection of Things: The W3C Thing Description Ontology. In *The Semantic Web (Lecture Notes in Computer Science)*. Springer International Publishing, Cham, 599–615. https://doi.org/10.1007/978-3-030-49461-2_35

[14] Andrei Ciortea, Simon Mayer, Olivier Boissier, and Fabien Gandon. 2019. Exploiting Interaction Affordances: On Engineering Autonomous Systems for the Web of Things. In *Second W3C Workshop on the Web of Things The Open Web to Challenge IoT Fragmentation*. Munich, Germany.

[15] Daniele Dell'Aglio, Emanuele Della Valle, Jean-Paul Calbimonte, and Oscar Corcho. 2014. RSP-QL Semantics. *International Journal on Semantic Web and Information Systems* 10, 4 (oct 2014), 17–44. https://doi.org/10.4018/ijswis.2014100102

[16] Daniele Dell'Aglio, Jean-Paul Calbimonte, Marco Balduini, Oscar Corcho, and Emanuele Della Valle. [n. d.]. On Correctness in RDF Stream Processor Benchmarking. ([n. d.]), 326–342. https://doi.org/10.1007/978-3-642-41338-4_21

[17] Daniele Dell'Aglio, Emanuele Della Valle, Frank van Harmelen, and Abraham Bernstein. 2017. Stream reasoning: A survey and outlook. *Data Science* 1 (2017), 59–83. https://doi.org/10.3233/DS-170006

[18] Laurence Goasduff. 2020. Gartner Top 10 Trends in Data and Analytics for 2020. https://www.gartner.com/smarterwithgartner/gartner-top-10-trends-in-data-and-analytics-for-2020/. Gartner.

[19] Aidan Hogan, Eva Blomqvist, Michael Cochez, Claudia d'Amato, Gerard de Melo, Claudio Gutierrez, José Emilio Labra Gayo, Sabrina Kirrane, Sebastian Neumaier, Axel Polleres, Roberto Navigli, Axel-Cyrille Ngonga Ngomo, Sabbir M. Rashid, Anisa Rula, Lukas Schmelzeisen, Juan Sequeda, Steffen Staab, and Antoine Zimmermann. 2020. Knowledge Graphs. (March 2020). arXiv:2003.02320 [cs.AI] https://arxiv.org/abs/2003.02320

[20] S. Kaebisch, T. Kamiya, McCool M., V. Charpenay, and M. Kovatsch. 2020. Web of things (WoT) thing description. W3C Recommendation. https://www.w3.org/TR/wot-thing-description/. W3C.

[21] Seyed Mehran Kazemi, Rishab Goel, Kshitij Jain, Ivan Kobyzev, Akshay Sethi, Peter Forsyth, and Pascal Poupart. 2020. Representation Learning for Dynamic Graphs: A Survey. *Journal of Machine Learning Research* 21, 70 (2020), 1–73.

[22] Robin Keskisärkkä, Eva Blomqvist, Leili Lind, and Olaf Hartig. 2019. RSP-QL$$^{\star }\!$$: Enabling Statement-Level Annotations in RDF Streams. In *Semantic Systems. The Power of AI and Knowledge Graphs (Lecture Notes in Computer Science)*. Springer International Publishing, Cham, 140–155. https://doi.org/10.1007/978-3-030-33220-4_11

[23] Evgeny Kharlamov, Sebastian Brandt, Martin Giese, Ernesto Jiménez-Ruiz, Yannis Kotidis, Steffen Lamparter, Theofilos Mailis, Christian Neuenstadt, Ö Özçep, Christoph Pinkel, et al. 2016. Enabling semantic access to static and streaming distributed data with optique. In *Proceedings of the 10th ACM International Conference on Distributed and Event-based Systems*. 350–353.

[24] Hyeongsik Kim, Abhisha Bhattacharyya, and Kemafor Anyanwu. 2019. Semantic query transformations for increased parallelization in distributed knowledge graph query processing. (nov 2019). https://doi.org/10.1145/3295500.3356212

[25] Sarit Kraus. 1997. Negotiation and cooperation in multi-agent environments. *Artificial Intelligence* 94, 1 (July 1997), 79–97. https://doi.org/10.1016/S0004-3702(97)00025-8

[26] Chan Le Van, Feng Gao, and Muhammad Intizar Ali. 2017. Optimizing the Performance of Concurrent RDF Stream Processing Queries. In *The Semantic Web (Lecture Notes in Computer Science)*, Eva Blomqvist, Diana Maynard, Aldo Gangemi, Rinke Hoekstra, Pascal Hitzler, and Olaf Hartig (Eds.). Springer International Publishing, Cham, 238–253. https://doi.org/10.1007/978-3-319-58068-5_15

[27] Manh Nguyen-Duc, Anh Le-Tuan, Jean-Paul Calbimonte, Manfred Hauswirth, and Danh Le-Phuoc. 2020. Autonomous RDF Stream Processing for IoT Edge Devices. In *Semantic Technology*. Springer International Publishing, Cham, 304–319. https://doi.org/10.1007/978-3-030-41407-8_20

[28] Andriy Nikolov, Peter Haase, Johannes Trame, and Artem Kozlov. 2017. Ephedra: Efficiently Combining RDF Data and Services Using SPARQL Federation. (2017), 246–262. https://doi.org/10.1007/978-3-319-69548-8_17

[29] Natasha Noy, Yuqing Gao, Anshu Jain, Anant Narayanan, Alan Patterson, and Jamie Taylor. 2019. Industry-scale knowledge graphs. *Commun. ACM* 62, 8 (jul 2019), 36–43. https://doi.org/10.1145/3331166

[30] Danh Le Phuoc, Minh Dao-Tran, Josiane Xavier Parreira, and Manfred Hauswirth. 2011. A Native and Adaptive Approach for Unified Processing of Linked Streams and Linked Data. In *The Semantic Web - ISWC 2011 - 10th International Semantic Web Conference (Lecture Notes in Computer Science, Vol. 7031)*. Springer, 370–388. https://doi.org/10.1007/978-3-642-25073-6_24

[31] Yuan Ren and Jeff Z. Pan. 2011. Optimising ontology stream reasoning with truth maintenance system. (2011). https://doi.org/10.1145/2063576.2063696

[32] Munindar P Singh. 2012. Semantics and verification of information-based protocols.. In *AAMAS*. Citeseer, 1149–1156.

[33] Riccardo Tommasini, Yehia Abo Sedira, Daniele Dell'Aglio, Marco Balduini, Muhammad Intizar Ali, Danh Le Phuoc, Emanuele Della Valle, and Jean-Paul Calbimonte. 2018. VoCaLS: Vocabulary and Catalog of Linked Streams. In *The Semantic Web – ISWC 2018*. Springer International Publishing, Cham, 256–272. https://doi.org/10.1007/978-3-030-00668-6_16

[34] Ruben Verborgh, Miel Vander Sande, Olaf Hartig, Joachim Van Herwegen, Laurens De Vocht, Ben De Meester, Gerald Haesendonck, and Pieter Colpaert. 2016. Triple Pattern Fragments: A low-cost knowledge graph interface for the Web. *Journal of Web Semantics* 37-38 (mar 2016), 184–206. https://doi.org/10.1016/j.websem.2016.03.003

[35] Yifei Wang and Jie Luo. 2018. An Incremental Reasoning Algorithm for Large Scale Knowledge Graph. (2018), 503–513. https://doi.org/10.1007/978-3-319-99365-2_45