

HYPERSTREAM: a Workflow Engine for Streaming Data

Tom Diethe

Amazon, Cambridge, CB1 2GA, UK

TDIETHE@AMAZON.COM

Meelis Kull

University of Tartu, 50090 Tartu, Estonia

MEELIS.KULL@UT.EE

Niall Twomey

Intelligent Systems Laboratory, University of Bristol, BS8 1UB, UK

NIALL.TWOMEY@BRISTOL.AC.UK

Kacper Sokol

Intelligent Systems Laboratory, University of Bristol, BS8 1UB, UK

K.SOKOL@BRISTOL.AC.UK

Hao Song

Intelligent Systems Laboratory, University of Bristol, BS8 1UB, UK

HAO.SONG@BRISTOL.AC.UK

Miquel Perelló-Nieto

Intelligent Systems Laboratory, University of Bristol, BS8 1UB, UK

MP15688@BRISTOL.AC.UK

Emma Tonkin

Intelligent Systems Laboratory, University of Bristol, BS8 1UB, UK

EMMA.TONKIN@BRISTOL.AC.UK

Peter Flach

Intelligent Systems Laboratory, University of Bristol, BS8 1UB, UK

PETER.FLACH@BRISTOL.AC.UK

Editor: A. N. Other

Abstract

This paper describes HYPERSTREAM, a large-scale, flexible and robust software package, written in the Python language, for processing streaming data with workflow creation capabilities. HYPERSTREAM overcomes the limitations of other computational engines and provides high-level interfaces to execute complex nesting, fusion, and prediction both in online and offline forms in streaming environments. HYPERSTREAM is a general purpose tool that is well-suited for the design, development, and deployment of Machine Learning algorithms and predictive models in a wide space of sequential predictive problems.

Source code, installation instructions, examples, and documentation can be found at: <https://github.com/IRC-SPHERE/HyperStream>.

Keywords: Stream Processing, Workflow Engine, Compute Engine, Streaming Data

1. Introduction

Scientific workflow systems are designed to compose and execute a series of computational or data manipulation operations (workflow) (Deelman et al., 2009). Workflows simplify the process of sharing and reusing such operations, and enable scientists to track the provenance of execution results and the workflow creation steps. Generally workflow managers are designed to work in offline (batch) mode. Well known examples are Kepler¹ and Taverna².

-
1. <https://kepler-project.org/>
 2. <https://taverna.incubator.apache.org/>

In streaming data scenarios, common to most industry segments and big data use cases, dynamic data is generated on a continual basis. Stream processing solutions have been receiving increasing interest (Garofalakis et al., 2016), popular examples including Apache Spark™ Streaming³ and Microsoft® Azure Stream Analytics⁴.

HYPERSTREAM harnesses the rich environment provided by the Python language to provide both stream processing and workflow engine capabilities, while maintaining an easy-to-use Application Programming Interface (API). This answers a growing need for scientific streaming data analysis in both academic and industrial data intensive research, as well as in fields outside of core computer science, such as healthcare and smart environments. HYPERSTREAM differs from other related toolboxes in various respects: i) it is distributed under the permissive MIT license, encouraging its use in both academic and commercial settings; ii) it depends only on a small set of requirements to ease deployment; iii) it focuses on streaming data sources unlike most workflow engines; is suitable for limited resource environments such as found in Internet of Things (IoT) and Fog computing scenarios (Bonomi et al., 2012); and iv) it allows both online and offline computational modes unlike most streaming solutions.

2. Features

This software has been designed from the outset to be domain-independent, in order to provide maximum value to the wider community. Source code, issue tracking, installation instructions, examples, and documentation can be found on GitHub⁵, as well as a discussion room⁶. HYPERSTREAM is currently supported with Python 2.7 and 3.6 on *ix platforms (*e.g.* linux, OS-X) and Microsoft Windows. For ease of installation, Docker containers are provided. HYPERSTREAM also makes use of continuous integration using Travis-CI.

The core requirements for HYPERSTREAM are summarised as follows:

1. the capability to create complex interlinked workflows
2. a computational engine that is designed to be “compute-on-request”
3. to be capable of storing the history of computation
4. a plugin system for user extensions
5. to be able to operate in online and offline mode
6. to be lightweight and have minimal requirements

(2) and (3) reduce unnecessary repeated computation, and enable full data pipeline provenance. One of the main motivating factors for (6) was that computations should be capable of being performed on minimal hardware, such as found in IoT settings (see Section 4.1 below).

3. Design

HYPERSTREAM is written in Python, and uses MONGODB for the back-end. This means that all system configuration and persistence is in MONGODB, although this does not mean

3. <http://spark.apache.org/streaming/>

4. <https://azure.microsoft.com/en-us/services/stream-analytics/>

5. <https://github.com/IRC-SPHERE/HyperStream>

6. <https://gitter.im/IRC-SPHERE-HyperStream>

that HYPERSTREAM is limited to MONGODB for stream storage (see below). The system consists of two layers: the stream layer and the workflow layer, as described below.

3.1 Stream Layer

At the stream layer there exist only *streams* and *tools*. Tools operate on streams to produce new streams, hence creating a chain of operation. A simple example of this can be seen in Figure 1⁷. Here the data originates from a comma-separated value file, is imported using the *csv_import* tool into the *sea_ice* memory stream, and then processed by the *sum_list* tool into a database stream called *sea_ice_sum*.

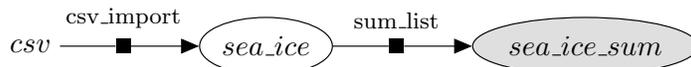


Figure 1: Example chain of computations. The filled (grey) node indicates that the *sea_ice_sum* stream is stored in the database rather than memory.

We treat all data as *streams* of *documents*, where a *document* can contain most Python object types, as long as they can be converted to Binary JavaScript Object Notation (BSON).

Tools are the computation elements, with fixed parameters and filters defined that can reduce the amount of data that needs to be read from the database. They take input data in a standard format (an iterator over stream documents) and output data via a generator in the same standard format. Tools are agnostic to where the data actually lives (*i.e.* memory/files/database).

Channels define the manifestation of streams for time ranges that have been computed, along with any specific processing required to read and write the streams, which abstracts away the specifics of interacting with different data sources. The built-in channels are the memory, database (MONGODB), file, assets, (Python) module, and tool channels. The tool channel is a subclass of the module channel, which in turn is a subclass of the file channel, which means that the tools themselves are stored in streams. The HYPERSTREAM plugin system allows users to define their own channels, in order to work with custom databases, file-based storage with custom formats or locations, or to modify the default capabilities of existing channels. An example machine learning plugin can be found at ⁸. This wraps Scikit-learn (Pedregosa et al., 2011) linear models into a HyperStream plugin, and provides examples for how this would be used for online learning and anomaly detection, and can for example be extended to the continual learning setting (Diethe et al., 2019).

3.2 Workflow Layer

Taking inspiration from factor graph notation for probabilistic graphical models (Buntine, 1994), workflows define a graph of “nodes” connected by “factors”, which can be surrounded by “plates”. Workflows can have multiple time ranges, which will cause the streams contained in the nodes to be computed on all of the ranges given. Workflows can be defined

7. https://github.com/IRC-SPHERE/HyperStream/blob/master/examples/tutorial_03.ipynb

8. <https://github.com/IRC-SPHERE/HyperStreamOnlineLearning>

to be operable in offline-only mode, or also available to the HYPERSTREAM online engine, which will cause the workflow to be executed continuously. Workflows are serialised to MONGODB by HYPERSTREAM for ease of deployment.

3.2.1 PLATES

Plates can be thought of as a “for loop” over parts of the computational graph contained within them. This is conceptually similar to the notion of plates in factor graphs. Both nodes and factors can be contained inside a plate.

An example is given below, where we construct an outer plate ‘C’ that loops over continents, and then an inner plate that loops over countries within each continent:

```

1 countries_dict = {
2     'Asia': ['Bangkok', 'HongKong', 'KualaLumpur', 'NewDelhi', 'Tokyo'],
3     'Australia': ['Brisbane', 'Canberra', 'GoldCoast', 'Melbourne', 'Sydney'],
4     'NZ': ['Auckland', 'Christchurch', 'Dunedin', 'Hamilton', 'Wellington'],
5     'USA': ['Chicago', 'Houston', 'LosAngeles', 'NY', 'Seattle']
6 }
7
8 # delete_plate requires the deletion to be first childs and then parents
9 for plate_id in ['C.C', 'C']:
10     if plate_id in [plate[0] for plate in hs.plate_manager.plates.items()]:
11         hs.plate_manager.delete_plate(plate_id=plate_id, delete_meta_data=
12             True)
13 for country in countries_dict:
14     id_country = 'country_' + country
15     if not hs.plate_manager.meta_data_manager.contains(identifier=id_country
16 ):
17         hs.plate_manager.meta_data_manager.insert(
18             parent='root', data=country, tag='country', identifier=
19             id_country)
20     for city in countries_dict[country]:
21         id_city = id_country + '.' + 'city_' + city
22         if not hs.plate_manager.meta_data_manager.contains(identifier=
23             id_city):
24             hs.plate_manager.meta_data_manager.insert(
25                 parent=id_country, data=city, tag='city', identifier=id_city
26 )
27 C = hs.plate_manager.create_plate(plate_id="C", description="Countries",
28     values=[], complement=True,
29     parent_plate=None, meta_data_id="country")
30 CC = hs.plate_manager.create_plate(plate_id="C.C", description="Cities",
31     values=[], complement=True,
32     parent_plate="C", meta_data_id="city")

```

3.2.2 NODES

A *node* is a collection of streams that live on the same plate, *i.e.* they have shared meta-data keys, and are connected in the computational graph by factors.

3.2.3 FACTORS

Factors are the workflow implementations of tools. A factor defines the element of computation: the tool along with the source and sink nodes. Basic factors will take input streams on (a) given plate(s), execute the tool on these streams, and output a stream that is one same plate(s). Multi-output factors are able to take streams from a plate and output streams on a sub-plate of that plate (*e.g.* by splitting).

Usually, the first factor in a workflow will be a special “raw” factor that uses a tool with no input streams that pulls in data from a custom data source outside of HYPERSTREAM.

4. Domain Specific Languages

HYPERSTREAM workflows are defined in terms of a Domain Specific Language (DSL).

4.1 Case Study: The SPHERE Project

The Sensor Platform for HEalthcare in a Residential Environment (SPHERE) project (Zhu et al., 2015; Woznowski et al., 2017; Diethe et al., 2018) uses multiple heterogeneous sensors for the purpose of health monitoring within the home environment. Part of the project involves the deployment of sensor systems to the homes of up to 100 volunteer families within the Bristol area of the UK. Each house has a limited computational budget due to the inconvenience of off-site hardware installation. In this setting a HYPERSTREAM instance runs in each house in *online* mode on an Intel® NUC, a 4 × 4 inch mini PC, alongside other services such as MONGODB, Apache ActiveMQ™, and Apache HTTP Server™. Here HYPERSTREAM is used to provide pseudo-real-time predictions using trained Machine Learning models and to perform online processing and summarising of the sensor data. In addition, when data is retrieved from the houses, computations are performed on a centralised database to perform aggregate computations and further “meta-summaries”. The SPHERE project has made heavy use of the workflow capabilities and the plugin architecture of HYPERSTREAM.

Figure 2 depicts an example workflow for prediction of sleep, showing nested plates, nodes and factors. Here the raw data comes from the SPHERE deployment houses, which are on the **H** plate. The wearable data is then split by its unique identifier (since there is more than one wearable per house) onto the **W** plate, which is nested inside the **H** plate. Two `sliding_apply` tools are then executed for each wearable in each house with differing length sliding windows (5s and 300s) to first compute windowed arm angles and then a windowed inactivity estimate, which is stored in the database channel and subsequently used as part of a sleep prediction algorithm.

5. Concluding remarks

We have presented HYPERSTREAM, a software package for processing streaming data with workflow creation capabilities, with a flexible plugin architecture. HYPERSTREAM is in active development, and contributions are actively welcomed (see <https://github.com/IRC-SPHERE/HyperStream/wiki/How-to-contribute>). Moreover, we would like to ac-

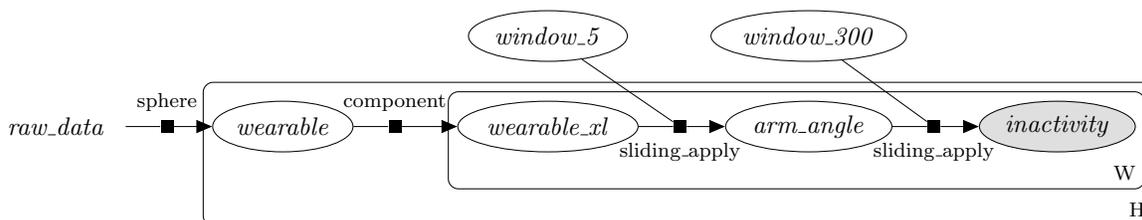


Figure 2: Example workflow.

knowledge all HYPERSTREAM contributors, who can be identified using the `git log` command.

The SPHERE Interdisciplinary Research Collaboration (IRC) is funded by the UK Engineering and Physical Sciences Research Council (EPSRC) under Grant EP/K031910/1.

References

- Flavio Bonomi, Rodolfo Milito, Jiang Zhu, and Sateesh Addepalli. Fog computing and its role in the internet of things. In *Proceedings of the first edition of the MCC workshop on Mobile cloud computing*, pages 13–16. ACM, 2012.
- Wray L Buntine. Operations for learning with graphical models. *Journal of artificial intelligence research*, 1994.
- Ewa Deelman, Dennis Gannon, Matthew Shields, and Ian Taylor. Workflows and e-science: An overview of workflow system features and capabilities. *Future generation computer systems*, 25(5):528–540, 2009.
- Tom Diethe, Mike Holmes, Meelis Kull, Miquel Perello Nieto, Kacper Sokol, Hao Song, Emma Tonkin, Niall Twomey, and Peter Flach. Releasing eHealth analytics into the wild: Lessons learnt from the SPHERE project. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '18*, pages 243–252, New York, NY, USA, 2018. ACM. ISBN 978-1-4503-5552-0. doi: 10.1145/3219819.3219883. URL <http://doi.acm.org/10.1145/3219819.3219883>.
- Tom Diethe, Tom Borchert, Eno Thereska, Borja de Balle Pigem, and Neil Lawrence. Continual learning in practice. *arXiv preprint arXiv:1903.05202*, 2019.
- Minos Garofalakis, Johannes Gehrke, and Rajeev Rastogi. *Data Stream Management: Processing High-Speed Data Streams*. Springer, 2016.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Przemyslaw Woznowski, Alison Burrows, Tom Diethe, Xenofon Fafoutis, Jake Hall, Sion Hannuna, Massimo Camplani, Niall Twomey, Michal Kozłowski, Bo Tan, Ni Zhu, Atis Elsts, Antonis Vafeas, Adeline Paiement, Lili Tao, Majid Mirmehdi, Tilo Burghardt, Dima Damen, Peter Flach, Robert Piechocki, Ian Craddock, and George Oikonomou. SPHERE: A sensor platform for healthcare in a residential environment. In *Designing, Developing, and Facilitating Smart Cities*, pages 315–333. Springer International Publishing, Cham, 2017.

Ni Zhu, Tom Diethe, Massimo Camplani, Lili Tao, Alison Burrows, Niall Twomey, Dritan Kaleshi, Majid Mirmehdi, Peter Flach, and Ian Craddock. Bridging e-health and the internet of things: The SPHERE project. *Intelligent Systems, IEEE*, 30(4):39–46, 2015.