

# Polyglot Data Management: State-of-the-Art & Open Challenges

Tutorial proposal – 1.5 hours

Felix Kiehn<sup>†</sup>, Mareike Schmidt<sup>†</sup>, Daniel Glake<sup>†◊</sup>, Fabian Panse<sup>†</sup>, Wolfram Wingerath<sup>‡\*</sup>,  
Benjamin Wollmer<sup>†\*</sup>, Martin Poppinga<sup>†</sup>, Norbert Ritter<sup>†</sup>

<sup>†</sup>dbis-research@informatik.uni-hamburg.de, <sup>‡</sup>data-science@uni-oldenburg.de

<sup>†</sup>University of Hamburg (UHH), Germany; <sup>◊</sup>Hamburg University of Applied Science (HAW), Germany;

<sup>‡</sup>University of Oldenburg, Germany; <sup>\*</sup>Baqend, Hamburg, Germany

## ABSTRACT

Due to the increasing variety of the current database landscape, polyglot data management has become a hot research topic in recent years. The underlying idea is to combine the benefits of different data stores behind a predefined set of common interfaces and thus address use cases that individual stores cannot meet. This can be accomplished using different approaches which vary greatly in terms of capabilities, functionality, and architectural concepts. This tutorial provides a detailed overview of the current state of research in polyglot data management. We motivate its use by showing the high diversity of existing data stores and discussing three use cases in which individual stores are insufficient. Thereafter, we present different taxonomies for classifying polyglot data systems and give a detailed review of a number of selected systems. Finally, we compare these systems based on their features and discuss open challenges that still need to be addressed in future research.

## PVLDB Reference Format:

Felix Kiehn<sup>†</sup>, Mareike Schmidt<sup>†</sup>, Daniel Glake<sup>†◊</sup>, Fabian Panse<sup>†</sup>, Wolfram Wingerath<sup>‡\*</sup>, Benjamin Wollmer<sup>†\*</sup>, Martin Poppinga<sup>†</sup>, Norbert Ritter<sup>†</sup>.

Polyglot Data Management: State-of-the-Art & Open Challenges. PVLDB, 14(1): XXX-XXX, 2022.

doi:XX.XX/XXX.XX

## 1 INTRODUCTION

Data is nowadays processed and managed in many different contexts with often completely different requirements (e.g., consistency vs. scalability, real-time vs. batch, write heavy vs. read heavy workloads, key-based access vs. analytical queries). As a logical consequence, many new data management and (distributed) data processing systems<sup>1</sup> have been developed over the last two decades, and have established themselves alongside the still widely used relational DBMSs. Examples of such systems are key-value stores (e.g., Redis), document stores (e.g., MongoDB), wide-column stores (e.g., Cassandra), graph databases (e.g., Neo4J), time-series databases (e.g., Influx), real-time databases (e.g., Firebase), array-databases (e.g., SciDB), and distributed computation frameworks (e.g., Spark). All these systems have been designed with different applications in

This work is licensed under the Creative Commons BY-NC-ND 4.0 International License. Visit <https://creativecommons.org/licenses/by-nc-nd/4.0/> to view a copy of this license. For any use beyond those covered by this license, obtain permission by emailing [info@vldb.org](mailto:info@vldb.org). Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

Proceedings of the VLDB Endowment, Vol. 14, No. 1 ISSN 2150-8097.  
doi:XX.XX/XXX.XX

<sup>1</sup>In the following we summarize both kinds of systems under the term *data store*.

mind and it is up to the users to decide which system is best suited for their specific use cases. However, the sheer volume of potential candidates and the everlasting evolution of use cases poses several challenges to such a selection process:

- **Confusing System Space:** Even as an expert, it is hard to keep an overview of all these systems and their respective sweet spots and limitations (even if some guidelines [10] are available).
- **Complex Trade-Offs:** Complex applications may have sub-components with contradictory requirements so that using a single data store always ends up in a trade-off between them.
- **In-Flux Requirements:** As applications change and evolve, so do their requirements. Thus, the initially selected data store may not remain the best fitting one over time. However, those changes are often not immediately visible to the users and may therefore remain undetected for a long time.

The approach of *polyglot data management* [26] aims to solve these problems by combining the benefits of several data stores (and their underlying techniques) without adopting their drawbacks and to hide the automatic coordination across these stores behind single or multiple interfaces. This has several advantages:

- **Unified Access:** Users only need to acquire knowledge about a few interfaces, which in the best case are also standardized.
- **Transparent System Evolution:** One can react internally – preferably automatically – to changing requirements by migrating data between data stores or even exchanging entire stores without users being affected.
- **Ease-of-Extensibility:** Newly developed data stores can be easily integrated into the existing polyglot data system, without having to completely rebuild existing structures.

Over the last decade, this basic idea has been implemented in the form of various systems, such as HadoopDB [1], BigIntegrator [35], PolyBase [7], Myria [32], FORWARD [24], BigDAWG [9], RHEEM [2], CloudMdsQL [18], or ESTOCADA [3]. Although most of these systems are based on the mediator-wrapper architecture [8], there are major differences between their implementations and limitations. In this tutorial, we give an overview of the current state-of-the-art in polyglot data management. We will (i) introduce key terminology, taxonomies, and techniques, (ii) discuss the features, advantages, and disadvantages of different existing systems, and (iii) present open challenges.

## 2 TUTORIAL OUTLINE

This tutorial is split into six parts:

- (1) **Motivation & Use Cases. (~10 min.)** We introduce the underlying idea of polyglot data management and motivate work

in this area by presenting three real-life use cases where individual data management solutions are not capable of meeting all of the given requirements.

- (2) **Terminology & Taxonomies. (~10 min.)** We introduce pivotal terms and present various taxonomies according to which polyglot data systems can be classified and compared.
- (3) **Overview Database Landscape. (~10 min.)** We give an overview of the different classes of currently used data stores. In addition, we delineate their respective advantages and disadvantages, and discuss how their individual design characteristics are linked to concrete functional and non-functional application requirements.
- (4) **Basic Techniques & Concepts. (~15 min.)** We discuss several concepts and techniques from distributed data management [25] and data integration [8] that are reused in solutions for polyglot data management, such as the mediator-wrapper architecture and operators for distributed data computation.
- (5) **Current Systems. (~30 min.)** In the main part of this tutorial, we present the key concepts of several existing polyglot data systems, such as PolyBase [7] and RHEEM [2]. We provide an overview of their respective architecture, functionality, and capabilities. We also analyze their respective strengths and limitations, from which we then derive potential use cases.
- (6) **Open Challenges. (~15 min.)** We close with a set of desiderata for polyglot data management, evaluate existing systems based on these desiderata, and discuss open problems based on the results of this evaluation.

### 3 GOALS AND OBJECTIVES

This section provides details about the tutorial parts.

#### 3.1 Part 1: Motivation & Uses Cases

We first motivate polyglot data management based on three real-world use cases, each covering a different type of application:

**E-Commerce and Customer Management:** An online store provides users with products using text descriptions, ratings, short films, and pictures, with a high availability requirement to serve purchasing services constantly. The different components require different levels of consistency (e.g., while product ratings tolerate low consistency in favor of partition tolerance and high availability, payments require high consistency). The high read throughput created by users requires a read-friendly distribution scheme. In addition to the typical OLTP traffic, analytical tasks (data mining & OLAP) are required to compute product recommendations and calculate annual business numbers.

**Agent-Based Simulation:** A mobility simulation for cities needs to process different kinds of agents (e.g., citizens or cars) which move in a spatial environment or along road networks, and whose decisions are affected by stationary objects, events, and other agents. Thus, it combines data about different agent types with graph and grid data, and requires complex data operations, such as routing or spatial filtering. Depending on the selected time-scale and spatial extension, the simulation produces an extensive set of diverse simulation results and thus a high write throughput. In addition, to display aggregated indicators and visualize moving agents on a map, data objects are streamed through the system at runtime.

**Healthcare Data Management:** A hospital needs to store and query detailed information on patients, laboratory, and radiology results, as well as doctors' and nurses' notes regarding observations and treatments. Thus, medical data (e.g., MIMIC II [20]) ranges from structured and unstructured data (patient data, lab results) to images and waveform data (CT and MRI images, electrocardiograms). The analysis of images and waveform data requires complex operations such as Fourier-Transformations (e.g., to find regions containing abnormal tissues). In addition, text search and graph algorithms are required (e.g., to find similar examination reports or analyze correlations between symptoms). To keep patients in an intensive care unit under continuous surveillance, real-time monitoring is essential. In case of emergencies, it is necessary that vital data, e.g., blood types, are available at any time.

#### 3.2 Part 2: Terminology & Taxonomies

Like federated databases [25] (and in contrast to multi-model data stores [21]), polyglot data systems provide a unified access on multiple data stores. However, in contrast to federated databases, they are designed to deal with data-store-heterogeneity on technical (e.g., architecture) and conceptual levels (e.g., data model, query language). Despite this common ground, the term *polyglot data management* has often been used with different meanings. To get a clear overview of potential interpretations, we present three taxonomies that assign existing systems to clear classes.

The first taxonomy distinguishes them based on the number of query interfaces they provide:

- **Multistores** provide access through a single query language,
- **Polystores** allow users to submit queries in different languages (e.g., the native languages of the individual stores).

The second taxonomy takes the internal architecture of these systems into account. Here we distinguish between:

- **Loosely-coupled** systems correspond to a loose network of already existing and often independently managed data stores where the mediator has read rights, but no rights to write data or reconfigure the individual stores. In general, they bear a strong resemblance to virtual integration systems.
- **Tightly-coupled** systems were often designed from the beginning as one overall system. The mediator has write permissions on the individual stores and is therefore also able to migrate data between them. In this setting, the individual stores are typically not autonomous, so that the mediator is able to adapt their configurations to its own requirements.
- **Hybrid** systems are those in which some stores are tightly coupled and others are loosely coupled.

Both taxonomies are used to classify existing systems in Table 1. The last taxonomy, not included in Table 1 due to space limitations, classifies these systems based on the degrees of (i) **data-store-heterogeneity**, (ii) **data-store-autonomy**, (iii) **transparency**, and (iv) **flexibility** as well as the **cost models** they support [29].

#### 3.3 Part 3: Overview Database Landscape

Since the emergence of various application areas with unique requirements has led to the development of many highly specialized systems, the current data management landscape is more diverse and heterogeneous than ever before. The focus of these individual

systems can vary wildly: While some abandon consistency and querying power in favor of schema-freeness and scalability, others handle specific data types (e.g., spatio-temporal), data models (e.g., graphs), or combine transactional with analytic processing (HTAP). This potpourri of different systems provides the possibility to adapt data management to the requirements of the own application. However, it also complicates the selection of a suitable system, since an overview of advantages and disadvantages of these systems is hard to procure due to their sheer mass. In addition, there are – as our examples in Section 3.1 show – numerous use cases in which requirements of different subareas are combined and thus none of these systems can be used in these use cases standalone.

In this part of the tutorial, we give an overview of the current data management landscape, including different classes of data stores (NoSQL, real-time, etc.) and data processing frameworks (e.g., Hadoop, Spark, Flink). Moreover, we present the strengths and weaknesses of the different classes of NoSQL data stores regarding different user requirements (functional and non-functional) [10].

### 3.4 Part 4: Basic Techniques & Concepts

The idea of polyglot data management is usually implemented using the mediator-wrapper architecture, which is well-known from virtual data integration systems [8]. Thus, the schemas of the individual data stores are typically mapped to a (virtual) mediated schema, e.g., via tuple-generating dependencies [30] or views [16]. These mappings are used to decompose user queries into several subqueries, each forwarded to another data store, and to combine their results to a single (consistent) set of query answers. The result of such a query rewriting is one or more *query execution plans* (QEPs), each following predefined optimization goals and determining the most efficient join-order to integrate the relevant data of the respective stores. When building the QEPs, the mediator must decide where to execute the individual operators based on the store-specific location of the processed data, also known as the *operator-placement problem* [28]. Since this problem is NP-hard, the optimal placement needs to be approximated [19].

In this part of the tutorial, we present basic knowledge on (i) the mediator-wrapper architecture, (ii) schema mapping languages, (iii) operators that allow an efficient combination of data across stores, such as bind and skew-aware joins [14], (iv) typical cost models used for optimization, and (v) several heuristics to solve the operator-placement problem during query planning.

### 3.5 Part 5: Current Systems

Polyglot data management is a subject of intense research, and various systems have been developed in recent years. Table 1 lists some popular systems classified based on the first two taxonomies we have introduced in Section 3.2. In this part, we give a deeper insight into the architecture and functionality of some of these systems. The selected systems implement highly diverse concepts and cover different categories from Table 1. These systems are:

- **PolyBase**: a loosely-coupled multistore that allows the integration of unstructured (HDFS) data from Hadoop clusters with relational and NoSQL data by using SQL [7]. Polybase is now an integral part of MS SQL Server<sup>2</sup>.

<sup>2</sup><https://docs.microsoft.com/en-us/sql/relational-databases/polybase/>

	<i>multistore</i>	<i>polystore</i>
<i>loosly-coupled</i>	<b>PolyBase</b> [7] BigIntegrator [35] FORWARD [24] Apache Drill (Calcite) [4] TATOOINE [6] QoX [27] QUEPA [23] Odyssey [15]	Myria [32]
<i>tightly-coupled</i>	<b>RHEEM</b> [2] MuSQL [12] HadoopDB [1]	<b>ESTOCADA</b> [3] Polypheny-DB [31]
<i>hybrid</i>	<b>CloudMdsQL</b> [18] SparkSQL [5]	<b>BigDAWG</b> [9]

**Table 1: Existing systems for polyglot data management. The selected representatives are in bold.**

- **RHEEM** (now **Apache Wayang**<sup>3</sup>): a tightly-coupled multistore that utilizes ML techniques to learn costs for data shipping operators from already existing execution logs [2, 17, 19].
- **ESTOCADA**: a tightly coupled polystore that automatically distributes data across multiple non-autonomous data stores (relational and NoSQL) to optimize query performance [3].
- **CloudMdsQL**: a hybrid multistore whose query interface uses a nesting approach to combine several native query languages within a single one [18].
- **BigDAWG**: A hybrid polystore with support for several data models (including their respective query languages) that is able to migrate data between individual stores [9].

After describing the individual systems, we assess for which use cases they are particularly well suited and for which rather less.

### 3.6 Part 6: Open Challenges

In the final part of the tutorial, we describe a set of properties desirable for polyglot data systems and illustrate them using the three use cases from Section 3.1. This includes (i) **preserving the non-functional capabilities** (e.g., scalability) of each underlying store despite unified access across multiple stores, (ii) support of **ad-hoc data manipulation** through a global interface, (iii) **automatic detection of changes** in system requirements or workloads, and corresponding **adjustments** of cross-store data distribution (permanent data migration between stores) and store-specific **reconfigurations**, (iv) efficient cross-system **query planning and execution** including cost-efficient **data shipping** and **optimal operator placement** under consideration of functional and non-functional system properties, (v) the ability to process data in **real-time** in a polyglot setup, and (vi) **multi-model schema management**.

Thereafter, we compare the systems of Section 3.5 based on these characteristics and highlight common deficits. Finally, we summarize the current status of polyglot data management and discuss the open challenges we have identified in this tutorial.

## 4 RELATIONSHIP TO PRIOR TUTORIALS

The tutorial has not been given before. The only past tutorial addressing a similar subject was held by Lu et al. 2018 [22]. However, while we discuss polyglot data management from the perspective

<sup>3</sup><https://wayang.apache.org/>

of several functional and non-functional properties, Lu et al. consider it in comparison to multi-model databases and thus focus on the collective use of different data models. Moreover, we present basic knowledge that has not been covered by Lu et al., such as the operator-placement problem and cross-system join operators. Finally, we present five systems in more detail (see Section 3.5). Of these, two were barely (BigDAWG) or not at all (CloudMdsQL) addressed by Lu et al. In addition and like polyglot data management in general, the other three systems (PolyBase, RHEEM, ESTOCADA) have been further developed in the meantime, and thus contain many novel aspects of interest to the audience.

## 5 INTENDED AUDIENCE AND MATERIAL

Our goal is to give VLDB attendees an overview of the current state-of-the-art and open challenges in both research and practice on polyglot data management. We will not assume any background in NoSQL systems, data integration, and distributed data management, making our tutorial appropriate for researchers, practitioners, and graduate students.

## 6 PRESENTER BIOGRAPHIES

Our team has a solid background in data management through more than 20 years of research at the Universities of Hamburg (UHH & HAW) and Oldenburg. We address the problems of polyglot data management in both research and practice, both in the multi-agent system MARS<sup>4</sup> and in the area of web technologies in the Backend-as-a-Service company Baqend<sup>5</sup>. Our knowledge on NoSQL, real-time, stream, and polyglot data management has been published in several surveys [10, 13], tutorials [33], and books [11, 34].

**Felix Kiehn** is a Ph.D. student at the UHH. In his Ph.D. research, he is exploring potential strategies for scalable and consistent ad-hoc data manipulation of data sharded and replicated across different data stores through a global interface.

**Mareike Schmidt** is also a Ph.D. student at the UHH. In her Ph.D. thesis, she is studying potential strategies for identifying changes in application requirements and workloads that require to migrate data between different data stores.

**Daniel Glake** is a Ph.D. student at the UHH and HAW. His Ph.D. research focuses on polyglot data management for multi-agent simulation systems where he has to work on different categories of data (e.g., spatio-temporal, array, graph).

**Wolfram “Wolle” Wingerath** is a professor for data science at the University of Oldenburg and a scientific adviser at Baqend. His current research focuses on data management in web applications, but he also has a strong background in NoSQL data stores, real-time databases, and stream processing.

**Norbert Ritter** is a full professor at the University of Hamburg. His research interests include distributed and federated database systems, transaction processing, caching, cloud data management, information integration, and autonomous database systems. He has been teaching NoSQL topics in various courses for several years.

<sup>4</sup><https://www.mars-group.org/>

<sup>5</sup><https://www.baqend.com>

## REFERENCES

[1] Azza Abouzeid, Kamil Bajda-Pawlikowski, Daniel J. Abadi, et al. 2009. HadoopDB: An architectural hybrid of MapReduce and DBMS technologies for analytical workloads. *Proc. VLDB Endow.* 2, 1 (2009), 922–933.

[2] Divy Agrawal, Sanjay Chawla, Bertty Contreras-Rojas, et al. 2018. RHEEM: Enabling cross-platform data processing - may the big data be with you! -. *Proc. VLDB Endow.* 11, 11 (2018), 1414–1427.

[3] Rana Alotaibi, Damian Bursztyn, Alin Deutsch, et al. 2019. Towards scalable hybrid stores: Constraint-based rewriting to the rescue. In *SIGMOD*. ACM, 1660–1677.

[4] Apache Software Foundation. 2022. *Apache Drill*.

[5] Michael Armbrust, Reynold S. Xin, Cheng Lian, et al. 2015. Spark SQL: Relational data processing in Spark. In *SIGMOD*. ACM, 1383–1394.

[6] Raphaël Bonaque, Tien Duc Cao, Bogdan Cautis, et al. 2016. Mixed-instance querying: A lightweight integration architecture for data journalism. *Proc. VLDB Endow.* 9, 13 (2016), 1513–1516.

[7] David J. DeWitt, Alan Halverson, Rimma V. Nehme, et al. 2013. Split query processing in polybase. In *SIGMOD*. ACM, 1255–1266.

[8] AnHai Doan, Alon Y. Halevy, and Zachary G. Ives. 2012. *Principles of Data Integration*. Morgan Kaufmann.

[9] Jennie Duggan, Aaron J. Elmore, Michael Stonebraker, et al. 2015. The BigDAWG polystore system. *SIGMOD Rec.* 44, 2 (2015), 11–16.

[10] Felix Gessert, Wolfram Wingerath, Steffen Friedrich, et al. 2017. NoSQL database systems: A survey and decision guidance. *Comput. Sci. Res. Dev.* 32, 3-4 (2017), 353–365.

[11] Felix Gessert, Wolfram Wingerath, and Norbert Ritter. 2020. *Fast and Scalable Cloud Data Management*. Springer.

[12] Victor Giannakouris, Nikolaos Papailiou, Dimitrios Tsoumakos, et al. 2016. MuSQL: Distributed SQL query execution over multiple engine environments. In *Big Data*. IEEE, 452–461.

[13] Daniel Glake, Felix Kiehn, Mareike Schmidt, et al. 2022. Towards polyglot data stores: Overview and open research questions. *CoRR abs/2204.05779* (2022).

[14] Ankush M. Gupta, Vijay Gadepally, and Michael Stonebraker. 2016. Cross-engine query execution in federated database systems. In *HPEC*. IEEE, 1–6.

[15] Hakan Hacigümüs, Jagan Sankaranarayanan, Jun’ichi Tatemura, et al. 2013. Odyssey: A multi-store system for evolutionary analytics. *Proc. VLDB Endow.* 6, 11 (2013), 1180–1181.

[16] Alon Y. Halevy. 2001. Answering queries using views: A survey. *VLDB J.* 10, 4 (2001), 270–294.

[17] Zoi Kaoudi, Jorge-Arnulfo Quiané-Ruiz, Bertty Contreras-Rojas, et al. 2020. ML-based cross-platform query optimization. In *ICDE*. IEEE, 1489–1500.

[18] Boyan Kolev, Patrick Valduriez, Carlyna Bondimouy, et al. 2016. CloudMdsQL: Querying heterogeneous cloud data stores with a common language. *Distributed Parallel Databases* 34, 4 (2016), 463–503.

[19] Sebastian Kruse, Zoi Kaoudi, Bertty Contreras-Rojas, et al. 2020. RHEEMix in the data jungle: A cost-based optimizer for cross-platform systems. *VLDB J.* 29, 6 (2020), 1287–1310.

[20] Joon Lee, Daniel J. Scott, Mauricio Villarreal, et al. 2011. Open-access MIMIC-II database for intensive care research. In *EMBC*. IEEE, 8315–8318.

[21] Jiaheng Lu and Irena Holubová. 2019. Multi-model databases: A new journey to handle the variety of data. *ACM Comput. Surv.* 52, 3 (2019), 55:1–55:38.

[22] Jiaheng Lu, Irena Holubová, and Bogdan Cautis. 2018. Multi-model databases and tightly integrated polystores: Current practices, comparisons, and open challenges. In *CIKM*. ACM, 2301–2302.

[23] Antonio Maccioni, Edoardo Basili, and Riccardo Torlone. 2016. QUEPA: QUerying and Exploring a Polystore by Augmentation. In *SIGMOD*. ACM, 2133–2136.

[24] Kian Win Ong, Yannis Papakonstantinou, and Romain Vernoux. 2014. The SQL++ semi-structured data model and query language: A capabilities survey of SQL-on-Hadoop, NoSQL and NewSQL databases. *CoRR abs/1405.3631* (2014).

[25] M. Tamer Özsu and Patrick Valduriez. 2020. *Principles of Distributed Database Systems, 4th Edition*. Springer.

[26] Pramod J. Sadalage and Martin Fowler. 2012. *NoSQL Distilled: A Brief Guide to the Emerging World of Polyglot Persistence*. Pearson Education.

[27] Alkis Simitsis, Kevin Wilkinson, Malú Castellanos, et al. 2012. Optimizing analytic data flows for multiple execution engines. In *SIGMOD*. ACM, 829–840.

[28] Utkarsh Srivastava, Kamesh Munagala, and Jennifer Widom. 2005. Operator placement for in-network stream query processing. In *PODS*. ACM, 250–258.

[29] Ran Tan, Rada Chirkova, Vijay Gadepally, et al. 2017. Enabling query processing across heterogeneous data models: A survey. In *BigData*. IEEE, 3211–3220.

[30] Balder ten Cate and Phokion G. Kolaitis. 2010. Structural characterizations of schema-mapping languages. *Commun. ACM* 53, 1 (2010), 101–110.

[31] Marco Vogt, Alexander Stiemer, and Heiko Schuldt. 2018. Polypheny-DB: Towards a distributed and self-adaptive polystore. In *Big Data*. IEEE, 3364–3373.

[32] Jingjing Wang, Tobin Baker, Magdalena Balazinska, et al. 2017. The Myria big data management and analytics system and cloud services. In *CIDR*.

[33] Wolfram Wingerath, Felix Gessert, Erik Witt, et al. 2018. Real-time data management for big data. In *EDBT*. 524–527.

[34] Wolfram Wingerath, Norbert Ritter, and Felix Gessert. 2019. *Real-Time & Stream Data Management - Push-Based Data in Research & Practice*. Springer.

[35] Minpeng Zhu and Tore Risch. 2011. Querying combined cloud-based and relational databases. In *CSC*. IEEE, 330–335.