KTH Royal Institute of Technology June 8, 2023

CONTINUOUS DEEP ANALYTICS

Paris Carbone, Seif Haridi

Final Report

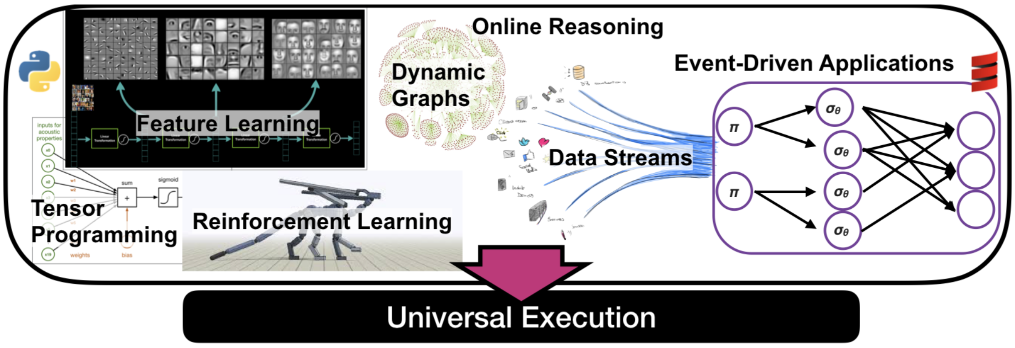

Executive Summary

The Continuous Deep Analytics SSF project addresses a set of core challenges in distributed computing systems for declarative data management and AI in support of critical decision making processes.

> **Research Statement:** *Modern end-to-end data pipelines are highly complex and unoptimised. They combine code from different frontends (e.g., SQL, Beam, Keras), declared in different programming languages (e.g., Python, Scala) and execute across many backend runtimes (e.g., Spark, Flink, Tensorflow). Data and intermediate results take a long and slow path through excessive materialization and conversions down to different partially supported hardware accelerators. This hinders the prospect of reliable actuation and data-driven critical decision making. The Continuous Deep Analytics (CDA) project aims to shape the next-generation systems for scalable, data-driven applications and pipelines for critical decision making. Our work aims to combine state of the art mechanisms in compiler and database technology together with hardware-accelerated machine learning and modern methods in reliable online ML.*

Our areas of focus are defined as follows:

Uncertainty, Dynamicity and Interpretability in Continuous Deep Analysis: Contemporary ML algorithms and methods produce models that output point predictions, without proper uncertainty quantification while assuming that training and test instances are drawn from the same (static) underlying distribution. The CDA project applies and further develops recent frameworks to obtain guarantees on the prediction error and well-calibrated probability distributions, respectively, and employs these frameworks for detecting and adapting to changes in the underlying distributions. The project also investigates auxiliary techniques for explaining and complementing predictions of popular black-box models (e.g., CNNs, GCNs,) and strives to provide human-understandable knowledge for critical decision-making.

Declarative Data Programming: We aim for simple and clear user-facing data programming models suitable for a seamless declaration of dynamic ML and event-based applications while enabling optimisation and unified execution through exposing special properties in computation, state and time via novel type systems.

Compilation and Code Generation for Continuous Deep Analytics: In particular, we explore the prospect and benefits of multi-stage compilation techniques that allow multiple levels of optimization, common in AI applications, to be effectively employed. In addition, we exploit the latest advances in tools for optimization and code generation (e.g., Google's MLIR - part of the LLVM library) on diverse hardware in support of various types of data such as tensors (matrices and vectors), streams, and relational tables.

Distributed Runtime Support: Distributed execution is one of the key enablers of democratizing deep analysis. We investigate two fronts of distributed execution: 1) distributed execution algorithms/protocols and graph-driven analysis. First, we explore foundational data replication protocols and programming models for building decentralized, high-performance dataflow applications. Furthermore, we tackle decentralization via graph/network analysis methods, applying existing but also evolving further graph algorithms that generalize and also unify different models of computation in a parallelized framework.

1 Project Background

The CDA project investigated the foundations for continuous deep analytics, starting from the problem domain to the specification, code generation, and distributed execution of diverse tasks on heterogeneous computing platforms. The following sections detail all the interdisciplinary areas, the specific problems, and the milestones we achieved throughout the course of the project.

1.1 BACKGROUND, MOTIVATION, AND ACHIEVEMENTS

The CDA project was launched in Q3 2017 by a group of ML and computer system researchers at KTH Royal Institute of Technology and RISE Digital Systems (former Swedish Institute of Computer Science) in Stockholm, Sweden. With a firm vision to establish the foundations of Continuous Deep Analytics, the project birthed a unique research group comprising active data science practitioners, theoretical ML researchers, programming language experts, and distributed computing researchers today.

Over the last years, the necessity for continuous deep analytics became more pronounced. Large-scale machine learning methods, a prime focus of research and industry, have been optimized to scale the training of black box models to large amounts of input data in a nearly linear fashion with respect to accuracy. GPT-4 by OpenAI is a significant landmark in language prediction "black-box" models, housing a capacity of hundreds of billions of parameters. However, relatively less effort has been committed to addressing the core limitations of large-scale machine learning, such as lack of support for interpretability, programmability, and seamless integration with critical decision-making processes that are important at a societal scale.

Building on these challenges, our project achieved in several ways to fill these gaps. We succeeded in creating reliable continuous processing technologies, via research studies and associated open-source software that can automate and enrich data-driven decision-making processes. We could further claim that the CDA project has significantly pushed the envelope of simplifying the complexities associated with data science, programming, and configuration tasks, which traditionally required complex engineering and rare domain expertise.

1.1.1 The Declarative Data Programming Challenge

Existing individual data processing frameworks served the sole purpose of providing a set of operations on specific types of data. However, our project saw the necessity and fulfilled the creation of a unified programming model that seamlessly integrates all these types of workloads under a single program. We accomplished this primarily in the Arcon sub-project. Arcon combines our original unified language called ArcLang which we designed from scratch, as well as a prototype dataflow engine, Arcon, and a storage solution for continuous deep analytics called WheelDB.

1.1.2 The Distributed Continuous System Execution Challenge

An integral component of scalable data processing is computational distribution. As of most frameworks available today, computation is distributed and scheduled to run on multiple compute nodes by each runtime in parallel while aiming for a specific workload such as "stream execution" on live data or "bulk iterative execution" on historic data employed e.g., by streaming systems or batch processing systems respectively. In addition, each runtime provides its own unique set of processing guarantees, fault tolerance and support for hardware accelerators according to its needs. The execution of data-driven continuous applications requires a combination of different runtime technologies to incorporate offline deep analytics (e.g., ML model training) and online preprocessing and inference (e.g., continuous stream computation). This leads to data pipelines consisting of fragmented codebases that run separately across system "silos", with no shared utilisation of hardware, causing excessive materialization of intermediate results (e.g., large ML models communicated through big files). Therefore, modern data pipelines are complex and dominated by unnecessary data synchronisation latencies that make critical decision-making impossible.

In the CDA project, we were addressed the distributed execution challenges in several different fronts.

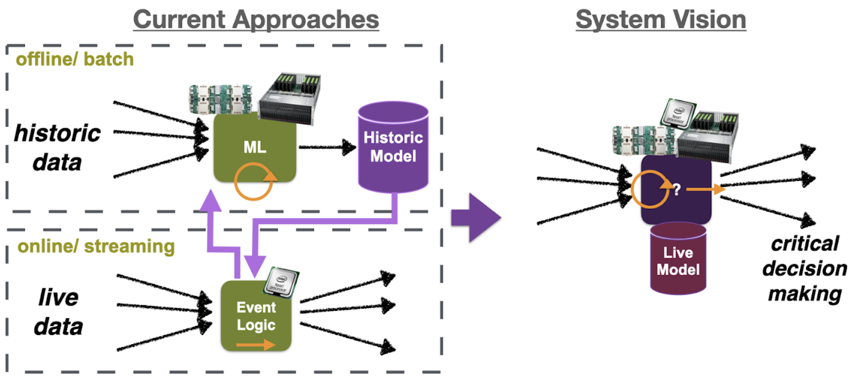

Figure 1: Approaches to Critical Decision Making.

1) Our Arcon runtime and ArcLang compiler employed a method called Multi-Stage compilation which allowed us to support heterogeneous resources and hardware on demand at compile time. 2) We investigated ways to optimise replicated data systems which led to the creation of the popular OmniPaxos protocol and library as well as UniCache, a novel technique that allows for data compression of such systems with up to 60% of data transmission reductions. 3) We conducted numerous studies exploring the use of graph representation learning, with significant improvements of the state of the art (Sarunas Girdzijauskas' work) as well as applying it to make smarter scalable data systems as we showed in Orb DB and GCNSplit.

1.1.3 The Reliable Online Learning Challenge

As machine learning is increasingly used not only for decision support but also automated decision-making, trust in the resulting decisions or recommendations becomes vital. Consequently, how to make machine learning solutions reliable is today a key question addressed by researchers from many disciplines. Being able to explain the logic behind the predictions of a machine-learning model is widely considered to be one of the cornerstones for enabling trust. However interpretable models that fulfill this criterion, can only rarely be used as direct replacements for the strongest black-box models. Instead, attention has shifted towards approaches to bridging the gap between (strong) black-box models and (weak) interpretable models, which allow predictions of the black-box models to be explained using (local or global) interpretable approximations. Our new novel methodologies developed in CDA addressed this issue through the use of conformal prediction which places statistically sound error estimates on every ML prediction. This line of work was supported by Henrik Boström's pioneering work in the project and our applications in the ORB inference-driven database design that was published and patented during the project.

1.2 CONCRETE GOALS AND OBJECTIVES

The CDA project aimed to address all aspects that define continuous deep analytics, from conceptual models and algorithms to concrete system specifications, from the programming model to the execution of data computing tasks. Figure 2 (not provided here) illustrated how we organized our objectives hierarchically to provide the necessary requirements of CDA from the bottom up. From system execution to programming abstractions and applications, we aimed to fulfill objectives within the proposed system implementation. The achievements of each underlying area and domain are summarized below.

1.2.1 Algorithms and ML Applications for CDA

Traditional ML and data engineering methods were designed to fit specific computation models. Recognizing the limitations of these existing methods, a core objective of the project was to identify emerging

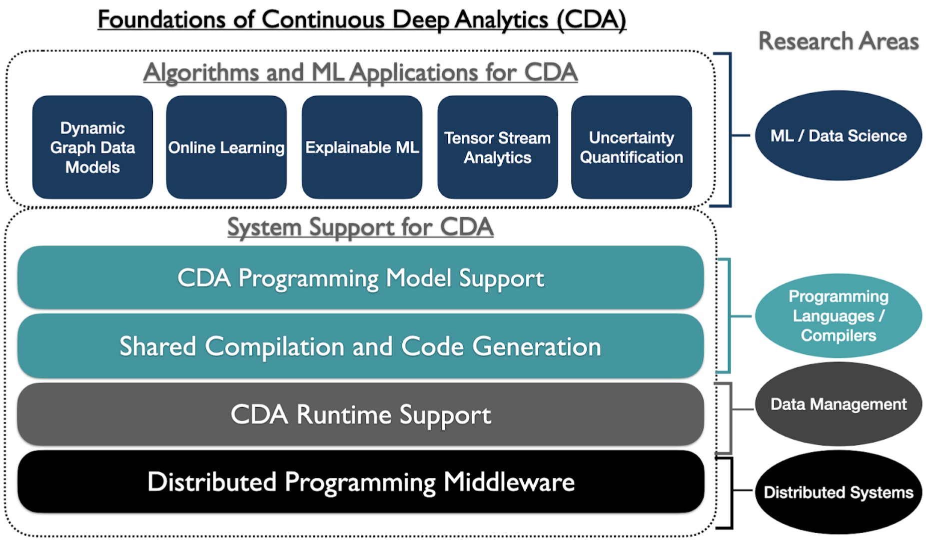

Figure 2: An Overview of Goals and Objectives for Continuous Deep Analytics.

problems that called for new computational models or a combination of existing ones. One of the primary domains we addressed was dynamic graphs. Utilized today in a multitude of applications such as social network analysis, recommender systems, traffic networks in large cities, 5G telecommunication performance monitoring, and medical data modeling, graph data has become increasingly significant.

Moreover, the emerging field of "Graph Embedding Models" proved promising for creating highly accurate predictive models from massive graphs, e.g., using Convolutional Neural Networks (CNNs). However, these models faced limitations in their applicability for critical decision-making due to continuous concept drift in the data, which invalidated their accuracy, coupled with slow training times.

In response, the CDA project aimed to address these challenges by establishing a formal specification of dynamic graph embeddings and the necessary system requirements to support it. Our achievements included online learning, efficient learning from streaming data, uncertainty quantification, ensuring the uncertainty of the predictions were properly quantified, concept drift detection, where the learning system was able to detect and adapt to changes in the underlying distributions, and explainable machine learning, which made the logic behind the predictions understandable to humans.

Overall, these achievements marked a significant advance in the field of continuous deep analytics, opening up new possibilities for ML and data engineering methods.

1.2.2 Data Programming for CDA

The CDA project investigated the case of a unified programming model for continuous and deep analytics. This included data types and operators that facilitated data transformations, serving, and querying (e.g., relational and stream operators), as well as model creation and reasoning (e.g., tensors, uncertainty quantification). To address the dynamic properties of data, our programming model integrated concept drift detection and adaptation, as well as reliable discretization (e.g., stream windows), while making runtime concerns transparent to the user.

Beyond language/model specification challenges, we strived to design a sophisticated compiler infrastructure for optimal code generation across operators, each with distinct workload and semantic differences, yet with a common need for their combined execution on shared hardware. This led us to utilize "Intermediate Representations" (IRs) as a common ground, while also exploiting active open-source libraries for code compilation, including LLVM and MLIR, which are used internally in Tensorflow, among other systems.

1.2.3 Runtime Support for CDA

The CDA project required fresh thinking in terms of data management and runtime execution. We questioned the purpose of current data processing technologies and opened new directions in computational systems design. Systems designed for continuous execution traditionally struggle with supporting heavy ML workloads, like the bulk iterative workflows of Tensorflow, and systems specialising in batch workloads often fail to meet the real-time processing requirements of critical applications.

Recognizing these challenges, we embarked on a new runtime design that could exploit heterogeneous hardware for mixed workloads. This led us to explore novel approaches to sharing active application state across different representations. We made significant strides in utilizing the latest advances in embedded databases, such as Log-Structure-Merge storage and adaptive indexing, aligning with the CDA project's vision.

1.2.4 Middleware Support for Scalable Execution in CDA

Creating correct and efficient distributed software is a challenging task. While existing languages and frameworks, such as Erlang or Akka, often focus on high-level application development abstractions, they sometimes compromise performance for convenience.

Our middleware aims included the design and development of programming abstractions that empower runtime developers to use the best communication strategy for each individual algorithm in a system, rather than conforming to a one-size-fits-all solution. Our developed middleware has been a major effort of the CDA project with numerous accomplishments and new frameworks (e.g., Kompact, OmniPaxos, UniCache, Portals). This allowed our platform to be efficient, correct, and maintainable in the long run, thereby meeting one of the key objectives of the project.

1.3 BASIC ORGANISATION, LEADERSHIP, GRANTS

The project composition consists of a series of work-packages organised based on the concrete area of their respective objectives. Work Packages 3 and 4 that are centred around algorithms and ML applications have been coalesced into a single work-package. The main project coordination is conducted by Prof. Seif Haridi while the work-package Principal Investigators (PIs) are shown below:

WP1 Dataflow Runtime with Elasticity in Time, Resources and Uncertainty (Christian Schulte, *Paris Carbone since 2019*).

WP2 Declarative Programming Models (Seif Haridi)

WP3 (previously WP3+WP4) Algorithms and ML Applications for CDA (Daniel Gillblad, Anders Holst, *Henrik Boström since 2019*)

The CDA group had regular (weekly) reports from all the sub-projects as well as adjacent projects from related grants seen below.

Related Grants

- EU-horizon 2020 StreamLine project (Dec. 2015 to Nov. 2018)
- Swedish Governmental Agency for Innovation Systems (VINNOVA), Predictive Models with Interpretability and Concept Drift Analytics (Oct. 2018 to March 2021)
- Swedish Governmental Agency for Innovation Systems (VINNOVA), Resultaten i staten (Oct. 2019 to Sep. 2021)
- Wallenberg AI NEST Data-Bound Computing Project (2022-2026)
- SSF Industrial Doctorand - Resilient Stateful Serverless Computing at the Edge(2022-2026)
- Digital Futures Pairs Project - PODS (2020-2022)
- Digital Futures Consolidation Project - Portals (2023-2024)
- Vinnova Avancerad Digitalisering - FASTER AI (2023-2024)

1.4 PROJECT CHANGES

No changes were made to the project objectives. The only notable change has been the extension of the project expiration to Dec 2022 for the research activities and until May 2023 for the exploitation activities.

1.5 STEERING GROUP

Since there was no official steering group claimed in the proposal of the project we omitted Appendices 1 and 2 from this report. However, we discuss several minor changes to the project management here. Professor Seif Haridi holds the overall management of the project as it was set in its initial proposal and vision. Professor Christian Schulte who was originally responsible for WP1 passed away in March 2020 while Professor Henrik Boström (KTH) effectively replaced Anders Holst (RISE) and Daniel Gillblad (RISE) in 2019 as the responsible for the ML research work in WP3-WP4. Furthermore, Lars Kroll (RISE) and Paris Carbone (RISE) became senior contributors to the project following their PhD graduation at KTH in Jan-2020 and Nov-2018 respectively. The latter also maintains an Assistant Professor position at KTH and is currently leading the efforts of the CDA project at RISE. None of the aforementioned changes in the steering group has affected the progress of the project.

2 The Research of CDA

Research in CDA has been driven by the vision of designing and creating novel software technologies from the ground up that can enable continuous and deep data analytics. The rest of this section summarises our published research work as well as the corresponding system contributions that have been released for open-source.

2.1 SCIENTIFIC RESULTS

The CDA project obtained several dozens of results addressing all of its objectives, as well as expanding in cross-cutting subjects that combined several of the objectives of our fields and materialized them in full-fletched software systems.

2.1.1 Contributions in Machine Learning Models and Methods

This part of the project has investigated major challenges in reliable machine learning to form the core system requirements for online ML and continuous deep analytics. In this line of work, we have looked into the following directions: i) quantifying prediction uncertainty, ii) online learning and scaling out algorithms, and iii) explainable machine learning.

Uncertainty Quantification

The project has contributed with a novel approach for uncertainty quantification, called Mondrian Conformal Regressors (Boström Johansson, 2020), which overcomes two weaknesses of the original conformal regression approach; i) predicted intervals may be several times larger or smaller than any previously observed error, and ii) the variance of the sizes of the predicted intervals is not positively correlated to the available information on the relative difficulty (prediction quality) of the predicted instances. In contrast, Mondrian conformal regressors can never produce intervals that are larger than twice the largest observed error. Moreover, a large-scale empirical investigation was conducted, showing that Mondrian conformal regressors have the desired property that the variance of the size of the predicted intervals is positively correlated with the accuracy of the function that is used to estimate prediction quality. As an alternative venue, (Johansson et al, 2021) investigated the effect of a crucial parameter (beta) of normalized conformal regressors on the efficiency (interval size) and studied the relation between the size of output intervals and error rate.

A Mondrian approach was proposed also for conformal predictive systems (CPS), called Mondrian conformal predictive systems (Boström Johansson, 2021). Such systems output cumulative distribution functions for real-valued labels of test examples, rather than point predictions (as output by regular regression models) or confidence intervals (as output by conformal regressors). By forming Mondrian categories from the predictions of the underlying point estimator, the proposed technique was demonstrated, using regression forests as underlying models, to outperform not only standard CPS with respect to continuous ranked probability score, but also to produce as tight prediction intervals as produced by normalized conformal regressors, while improving upon the point predictions of the underlying regression forest. In an earlier paper (Werner et al, 2020), a large-scale investigation of CPS showed that by using either variance or the k-nearest-neighbour method for estimating prediction quality, a significant increase in performance, as measured by the continuous ranked probability score, can be obtained for regression forests, compared to omitting the quality (difficulty) estimation. The results furthermore showed that the use of out-of-bag examples for calibration is competitive with the most effective way of splitting training data into a proper training set and a calibration set, without requiring tuning of the calibration set size.

A Python package, called crepes, for developing conformal regressors and predictive systems on top of any regression model has been developed (Boström, 2022). The package implements standard, normalized and Mondrian conformal regressors and predictive systems, allowing the user to employ own difficulty estimates and Mondrian categories, as well as some provided standard options for these. The documentation is found at https://crepes.readthedocs.io, while source code is found at the GitHub repository: https://github.com/henrikbostrom/crepes

The application of Venn-Abers predictors to the task of calibrating multi-class models was also investigated in (Johansson et al, 2021), which received the Alexey Chervonenkis award for the best paper at the Tenth Symposium on Conformal and Probabilistic Prediction and Applications. The paper proposed to estimate the probability that the class label predicted by the underlying model is correct, rather than providing a probability distribution over all labels. An extensive empirical investigation showed that the suggested approach, when applied to both Platt scaling and Venn-Abers, is able to improve the probability estimates from decision trees, random forests and extreme gradient boosting.

In an award-winning paper (Karlsson et al, 2020), inductive conformal prediction was applied to a dataset of laboratory-generated aerosol particles, i.e., small airborne particles suspended in air affecting the climate and human health, consisting of ten particle subclasses that were grouped into four parent classes. Different types of particles come from different sources and impact the environment in different ways, which is why a reliable particle classification is of interest. The performance of the inductive conformal predictor (ICP) was evaluated on particle subclasses that were not included in training or calibration and was shown to give accurate predictions in some cases, namely if the unknown particle is similar to the known ones in the parent class. The precision of the underlying model was not high enough to reject all unknown particles for any subclass at the chosen significance levels, but the ICP managed to reject them at a higher rate if they were sufficiently different from the training and calibration samples.

Scaling Out Algorithms for Online ML

Following a collaboration with Amazon AWS (Seattle), one focus of the project has been on the scalability of distributed gradient boosting trees. Methods that allow scaling along both the data and feature dimension (block-distribution), improving the scalability characteristics, were proposed and, through proper use of sparsity, were shown to reduce the communication cost of the algorithm by orders of magnitude for highly sparse data. The corresponding paper (Vasiloudis et al, SIGIR, 2019) received the best short paper award at SIGIR, which is the top conference on Information Retrieval.

Uncertainty quantification for online random forests was investigated in (Vasiloudis et al, 2019), where two algorithms were proposed for forming prediction intervals; one based on conformal prediction, and the other based on quantile regression. The approaches were demonstrated to be able to maintain specified error rates, with constant computational cost per example and bounded memory usage.

The above work on scaling out gradient boosting trees and online uncertainty quantification were included in the Ph.D. thesis Scalable Machine Learning through Approximation and Distributed Computing (Vasiloudis, 2019).

In (Werner et al, 2021), different strategies for updating split conformal predictive systems in an online (streaming) setting were evaluated. The updating strategies varied in the extent and frequency of retraining as well as in how training data was split into proper training and calibration sets. An empirical evaluation was presented, considering passenger booking data from a ferry company, which stretches over a number of years. The passenger volumes changed drastically during 2020 due to COVID-19 and part of the evaluation focused on which updating strategies work best under such circumstances. Some strategies were observed to outperform others with respect to continuous ranked probability score and validity, highlighting the potential value of choosing a proper strategy.

Explainable Machine Learning

In (Boström et al, 2020), techniques for explaining forecasting models for multivariate time-series were investigated. Various approaches to explaining predictions of black box models have been proposed in the past, including model-agnostic techniques that measure feature importance (or effect) by presenting modified test instances to the underlying black-box model. These modifications typically rely on choosing feature values from the complete range of observed values. However, when applying machine learning algorithms to the task of forecasting from multivariate time-series, the temporal aspect should be taken into account when analysing the feature effects. In (Boström et al, 2020), we proposed a modification of individual conditional expectation (ICE) plots, called ICE-T plots, which displays the prediction change for temporally ordered feature values, and demonstrated, in collaboration with the Swedish National Financial Management Authority, its use through a case study on predicting the Swedish gross domestic product (GDP) based on a comprehensive set of indicator and prognostic variables.

In (Amr et al, 2022), a novel explanation technique, called CEGA (Characteristic Explanatory General Association rules), was proposed. The technique employs association rule mining to aggregate multiple explanations generated by any standard local explanation technique into a set of characteristic rules. The technique was compared to two state-of-the-art methods, Anchors and GLocalX, for producing local and aggregated explanations in the form of discriminative rules. The results suggest that the proposed approach provides a better trade-off between fidelity and complexity compared to the two state-of-the-art approaches; CEGA and Anchors significantly outperform GLocalX with respect to fidelity, while CEGA and GLocalX significantly outperform Anchors with respect to the number of generated rules.

Dynamic Graph Representation Learning

This part of the project focused on developing new techniques in the field of graph representation learning (GRL). GRL involves extracting meaningful and useful features from graph data, such as social networks, biological networks, and citation networks, and utilizing them in existing machine learning architectures for tasks such as node classification, link prediction, and personalized recommendation. In particular, we investigated and made contributions within several aspects of GRL, including contributions on novel random-walk based techniques, as well as on new approaches using Graph Neural Networks (GNNs). Furthermore, we investigated the use of self-supervised learning techniques in GNNs. Self-supervised learning is a promising approach to machine learning that allows models to learn from unlabeled data, and we applied it to GNNs with the purpose of bridging the gap between supervised and unsupervised learning. Random-Walk based GRL

Most existing GRL methods explore the structure and metadata associated with the graph using random walks and employ an unsupervised or semi-supervised learning schemes. Learning in these methods is context-free, resulting in only a single representation per node. However, recent studies have argued on the adequacy of a single representation and proposed context-sensitive approaches, which are capable of extracting multiple node representations for different contexts. This proved to be highly effective in applications such as link prediction and ranking. In our projecgt we focused on GRL techniques that result in multiple node representations depending on the context. In particular, in our work [Kefato et al. 2020] we show that in order to extract high-quality context-sensitive node representations it is not needed to rely on supplementary node features, nor to employ computationally heavy and complex models. We proposed GOAT, a context-sensitive algorithm inspired by gossip communication and a mutual attention mechanism simply over the structure of the graph. We investigated the efficacy of GOAT on link prediction and node clustering tasks and showed that it consistently outperforms SOTA on real-world datasets.

Furthermore, in a follow-up work [Samy et al. 2022] we have investigated how random-walk-based GRL techniques can be applied for semantically rich heterogeneous networks consisting of multiple node and edges types. Most HIN embedding methods exploit meta-paths to retain high-order structures, yet, their performance is conditioned on the quality of the (generated/manually-defined) meta-paths and their suitability for the specific label set. Through our work we proposed SchemaWalk - an approach that utilizes network schema as a unique blueprint of HIN, which is capable of uniformly sampling all edge types within the network schema during the random walk. In contrast to SOTA, our approach does not need to have any prior knowledge of the domain, nor need to hold any assumptions on heterogeneous graph structure (i.e. edge types). Our experiments on node classification show the robustness of the method regardless of the graph structure, as compared to the SOTA baselines.

Recommender Systems

Recommender systems (RSs) are often modeled as temporal interaction networks between two sets of entities (users and items). The terms cover a variety of notions, e.g. users could be customers in an e-commerce system, or accounts on Reddit, YouTube, or Spotify; items could be products, posts, or media produced or consumed by users. Predicting the next item that a user interacts with is critical for any successful recommender system. Often RSs are centered around the user, who is modeled using her recent sequence of activities. Recent studies have shown the effectiveness of modeling the mutual interactions between users and items using separate user and item embeddings.

In our work [Kefato et al. WWW2021] we proposed a novel method called DeePRed (Dynamic Embeddings for Interaction Prediction) that predicts the next items that users interact with, based on their previous history of interaction. We model our problem through bipartite temporal interaction networks, as they can naturally and effectively represent user-item interactions over time. DeePRed computes separate embeddings from the point of view of both: users and items. The key idea behind the effectiveness of DeePRed is that, each time a user interacts with an item, it is modeled using a sequence of k recent items she interacted with, which reflects a context of interaction. Experiments show that DeePRed significantly outperforms the best state-of-the-art approaches on the next item prediction task, while gaining more than an order of magnitude speedup over the best-performing baselines. Although this study is mainly concerned with temporal interaction networks, we also show the power and flexibility of DeePRed by adapting it to the case of static interaction networks, substituting the short- and long-term aspects with local and global ones.

Graph Neural Networks

We continued our line of work on link prediction/recommendation using Graph Neural Networks, which have gained popularity in recent years due to their effectiveness in a variety of tasks, including social network analysis, recommender systems, molecule property prediction etc. Most approaches are based on variants of graph neural networks (GNNs) that focus on transductive link prediction and have high inference latency. However, many real-world applications require fast inference over new nodes in inductive settings where no information on connectivity is available for these nodes. Thereby, node features provide an inevitable alternative in the latter scenario. To that end, we proposed Graph2Feat [Samy et al. 2023], which enables inductive link prediction by exploiting knowledge distillation (KD) through the Student-Teacher learning framework. In particular, we show that Graph2Feat learns to match the representations of a lightweight student multi-layer perceptron (MLP) with a more expressive teacher GNN while learning to predict missing links based on the node features, thus attaining both GNN's expressiveness and MLP's fast inference. Our proposed approach is general; it is suitable for transductive and inductive link predictions on different types of graphs regardless of them being homogeneous or heterogeneous, directed or undirected. Our experiments demonstrated that Graph2Feat significantly outperformed SOTA methods in terms of AUC and average precision in homogeneous and heterogeneous graphs as well as showed the minimum inference time compared to the SOTA methods, and 100x acceleration compared to GNNs.

Self-supervised learning for GNNs

We continued our work on GNNs, and in particular, investigated the use of self-supervised learning techniques for GNNs. Self-supervised learning is a promising approach in circumstances where ML workloads have unlabeled (or sparsely labeled) data due to the high costs of (manually) labeling data. In this project, we have investigated unsupervised learning techniques that are powerful enough to achieve comparable results as semi-supervised/supervised techniques. In [Kefato et al. SSL2021] we propose, SelfGNN, a novel contrastive self-supervised graph neural network (GNN) without relying on explicit contrastive terms. We leverage Batch Normalization, which introduces implicit contrastive terms, without sacrificing performance. Furthermore, as data augmentation is key in contrastive learning, we introduce four feature augmentation (FA) techniques for graphs. Though graph topological augmentation (TA) is commonly used, our empirical findings show that FA performs as good as TA. Our empirical evaluation real-world data shows that SelfGNN is powerful and leads to a performance comparable with SOTA supervised GNNs and is always better than SOTA semi-supervised and unsupervised GNNs.

However, often SSL approaches rely on heuristically crafted data augmentation techniques. Our follow-up work [Kefato et al. 2022] addresses this issue and proposed GRAPHSURGEON (self-supervised GNN that jointly learns to augment), a novel SSL method for GNNs, which, unlike most existing methods, requires neither data augmentation using heuristics nor explicit negative samples. We design an SSL architecture in such a way that data augmentation is jointly learned with the graph representation. Moreover, by using a principled constrained optimization objective, we avoid the need for explicit negative samples. In particular, we take advantage of the flexibility of the learnable data augmentation and introduce a new strategy that augments in the embedding space, called post-augmentation. This strategy has a significantly lower memory overhead and run-time cost. Furthermore, as it is difficult to sample truly contrastive terms, we avoid explicit negative sampling, and instead of relying on engineering tricks, we use a scalable constrained optimization objective motivated by Laplacian Eigenmaps to avoid trivial solutions. To validate the practical use of GRAPHSURGEON, we perform an empirical evaluation using 14 public datasets across a number of domains, ranging from small to large-scale graphs with hundreds of millions of edges. Our finding shows that GRAPHSURGEON is comparable to six SOTA semisupervised and on par with five SOTA self-supervised baselines in node classification tasks.

2.1.2 Contributions in Computer Systems Research

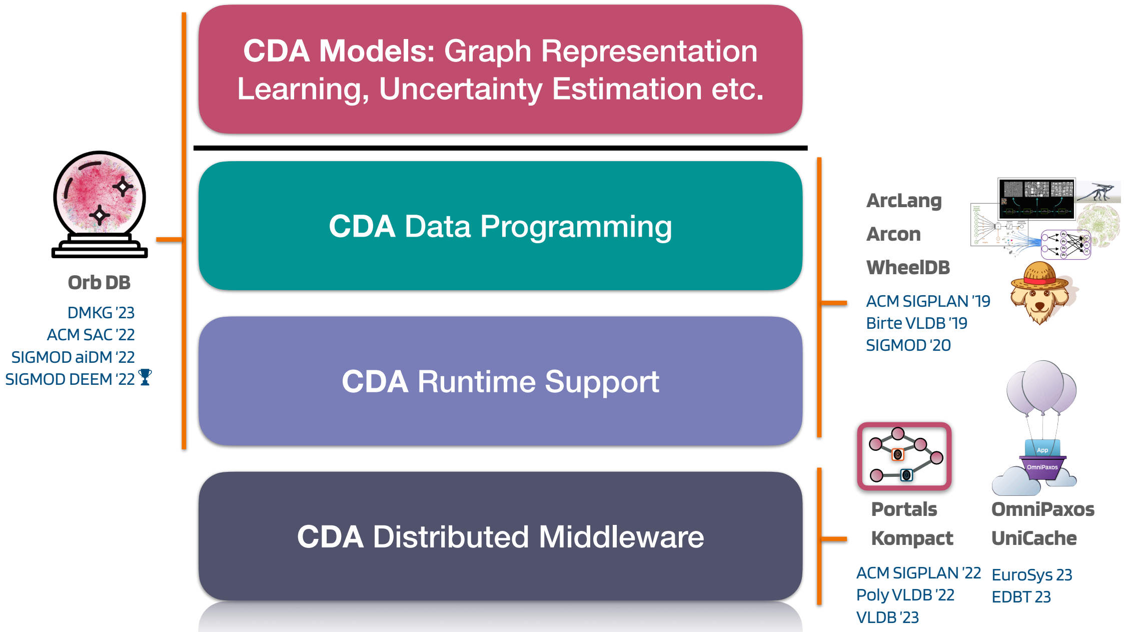

Figure 3: An Overview of System Contributions in CDA.

The project contributions in systems research have touched on open problems in compiler technology as well as declarative models and data management. Figure 3 depicts our core contributions within and across our different problem domains.

Below we summarise the components and systems research of CDA starting from a top-down fashion: *Orb DB*, *ArcLang* (data programming) and its LLVM-based compiler, including our *Arcon* and *WheelDB* runtime, the *Portals* serverless middleware for data applications and finally, the resilient execution results of *OmniPaxos*.

Orb DB

Orb DB is a special form of database which adopts many of our findings in uncertainty estimation as well as graph representation learning and offers next-generation system support for data management and advanced reasoning. Reasoning in knowledge graphs can take various forms including (1) pointwise queries (e.g., identifying a person with a specific personal number), (2) multi-hop queries (e.g., locating all family members of actors in films directed by friends of Tom Cruise), (3) monitoring queries (e.g., identifying drugs with excessive side effects), (4) exploratory searches (e.g., determining the top 10 relevant activities for an individual), and (5) similarity searches (e.g., finding painters comparable to Carl Larsson). Although knowledge graphs (KGs) play a crucial role as enablers for applications of significant societal importance, recent studies reveal that progress has stagnated due to the absence of a suitable system architecture capable of efficiently storing large dynamic graph data or facilitating complex queries involving new, unknown entities and connections. Orb DB [Horchidan, Carbone '23] materializes this idea, by applying the power of graph neural networks and generative AI.

Generative AI for Data Systems During the first half of the CDA project we investigated a new series of studies in "AI-Assisted" data systems a series of data management research problems that eventually led to the conception of Orb DB. Our initial goal was to try and replace monolithic and costly, yet state of the art algorithms (e.g., heuristics, optimizers) with low-cost AI models. The first results have been inspiring and encouraging to continue on this avenue of research. More concretely we were able to solve the NP-hard stream graph partitioning problem using Graph Neural Networks combined with Heuristics [Abbas et al. '22]. Figure 4 showcases the approach. With GCNSplit we investigated the prospect of "training a partitioning algorithm" by fusing the multiple objectives of stream partitioning in the process of model training. The outcome was a powerful model, capable of satisfying all objectives on unseen streams, generalizing its applicability to different graphs, many of which varied from the original graph that the model was trained on.

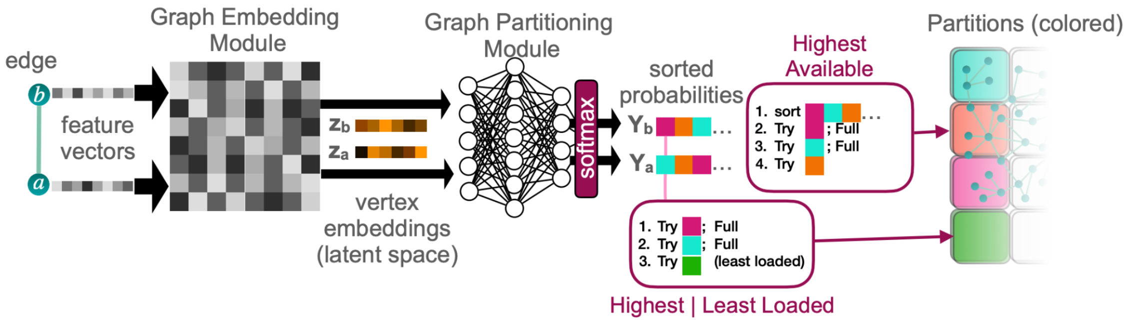

Figure 4: GCNSplit, a showcase of solving stream graph partitioning with AI.

The core idea of Orb The potential of inference-based query execution was investigated within a feasibility study with doctoral student Sonia Horchidan which resulted in a vision paper accepted for publication at the first DMKG symposium(Data Management in Knowledge Graphs) 2023. Among other highlights, we compared the performance scalability of Neo4j (state-of-the-art graph database) in comparison to Query2Box (AI-based query execution method) for multi-hop queries (1 to 3 traversals/hops long) in terms of execution latency. Our results showed a distinct ability in AI-based models in accommodating varying levels of complexity without significant performance overheads.

Orb DB and CDA As summarized in [Horchidan, Carbone '23], Orb materializes an important trade-off and moves the decision to the hands of the user. At one extreme we can let Orb DB execute a query completely in the embedding space at the highest uncertainty level. This would include many generated results, connections, and labels and execute fast. At another extreme, we can restrict Orb execute every computation using graph traversals that are evidently slow and incomplete in skewed graphs. Therefore, Orb's model allows users to set a maximum uncertainty level and run queries at a hybrid mode, tuning for less uncertainty while also gaining speed. Orb is an example of a continuous deep analytics system built from the ground up with the help of generative AI. Several of our underlying studies and techniques that let to Orb were published and discussed in SIGMOD aiDB amd DEEM workshops where we also got the best paper and best presentation awards.

ArcLang, Arcon and the MLIR Compiler

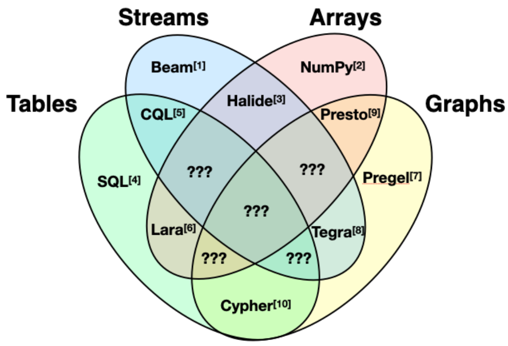

Figure 5: Open Domains in Data Programming of Continuous Deep Analytics.

Dataflow systems have surfaced as a premier choice for executing continuous analytics applications in cloud environments, boasting scalability, quick response times, and robust processing assurances. However, their potential is often curtailed by restrictions set by their programming interfaces, be it dataflow libraries with unfettered generality or query languages with confining data models. To counteract these obstacles, we developed Arc-Lang, a programming language tailored for continuous and deep analytics, which embraces flexibility, safety, efficiency, and succinctness and provides the missing ground in expressivity for continuous and deep models of computation as shown in Figure 5.

ArcLang contains a feature complete set of semantics and also makes use of a compiler implemented on top of the MLIR project (Multi-Level Intermediate Representation), an active effort for supporting intermediate languages and common compiler infrastructure for data analytics. MLIR allows expressing computations across different problem domains. It was initially contributed by Google to the LLVM ecosystem. An MLIR extension is called a “dialect” and multiple dialects in MLIR can coexist in the same program. Using MLIR further allows us to reuse functionality and optimizations provided by, for example the TensorFlow dialect, and their optimisations.

Data generation is happening today at an unparalleled scale, originating from a myriad of sources such as social media interactions to automated industrial sensors. These sources tirelessly produce vast quantities of raw and semi-structured data. This explosive growth has given rise to a new realm of continuous analytics applications, demanding instantaneous analysis of data as it is produced to facilitate swift and precise decision-making. Instances of such continuous analytics applications encompass financial fraud detection, coordination of ride-sharing services, stock market forecasting, real-time medical diagnosis, meteorological predictions, cyber security measures, and speech recognition technologies. Continuous analytics applications are generally constructed as a chain of stages that progressively extract meaningful insights from data. The raw data first has to be processed using application-specific logic to clean it, following which it is transformed into a structured format. To make it suitable for analysis, the data is further refined with historical information. The analysis can vary from simple summarization calculations to sophisticated machine learning inference. The ArcLang project evolved from Arc [Segeljat et al SIGPLAN 2019] which was built on the Weld intermediate representation and is as of today today a standalone programming language targeting dataflow execution semantics.

The Arcon Dataflow Runtime

Arcon [Meldrum et al 2019] is a data processing runtime built from the ground up to support hybrid batch and stream pipelines. At its current state it provides support for stream operators such as stream windows and a novel state management layer using the latest advances in embedded database technology. Already at its early stage, Arcon has attracted the attention of data engineering groups as well as top academic communities due to its novel design and promising performance. Our group presented Arcon at BIRTE the premier real time-analytics venue at VLDB in Los Angeles 2019 as well as the research track of Flink Forward 2019, the most popular event in stream processing technology. Arcon has also gained visibility in the Swedish ICT at the Data Science Summit 2019, CASTOR events organized by KTH for the Swedish industry as well as the Chaos Engineering conference in Stockholm in 2019. Industrial partners such as King and Ericsson have also expressed interest into putting Arcon into practice and experimenting with its capabilities once the first production-ready release is out. Arcon can run code generated by Arc and orchestrate it to utilize scalable state and computational resources available in a cluster.

WheelDB, is one of the extensions of the Arcon runtime that became an independent software library. The WheelDB research work behind the system is highly novel and currently under submission at a top-venue in data management. In essence, WheelDB is the first data structure (embedded database) that targets continuous and deep analytics. It can be used as a persistence engine for dataflow state, allowing efficient real-time operations on streams as well as fast analytical lookup queries. The key behind WheelDB is its ability to pre-materialize results on a wheel-like data structure and to exploit implicit application time ordering.

Portals

The Portals framework is an effort within CDA to provide a complete programming and execution environment that allows for arbitrary continuous and deep applications and services to be written and interact in a more organic and trustful way. In a sense, Portals address the distributed middleware challenges in continuous deep analytics. If we look into the currently available programming frameworks they are built around silos: Tensorflow is used for training, Flink for streaming, Ray for reinforcement learning, yet, different applications across frameworks cannot trivially be used together and interact in a service-oriented manner.

The Portals framework [Spenger et al ’22] as shown in Figure 6 abstracts a common programming layer that is flexible for building applications, similar to actor programming. At the same time, Portals enable data-parallel execution while automating fault tolerance and data-privacy (GDPR) across data applications that grow organically. The methodologies we combined to build Portals include static program analysis, serverless execution with the portals abstraction and atomic streams at runtime to enable these strong guarantees. The Portals framework will have its official release in August 2023 and it will

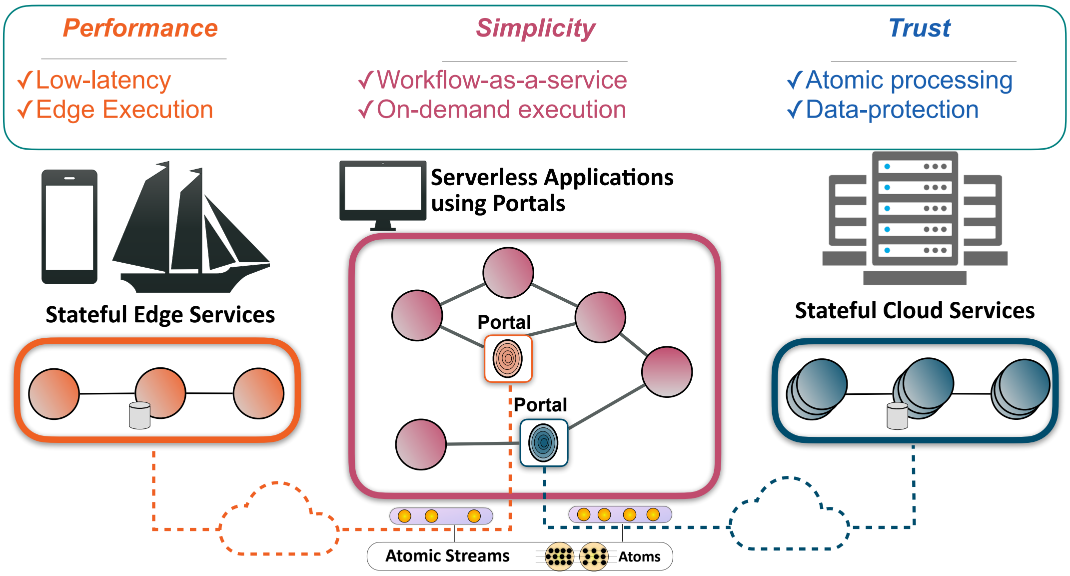

Figure 6: An Overview of the Portals Framework Achievements.

be demonstrated for the first time at the VLDB conference [Spenger et al '23].

Kompact

Kompact [Kroll '20] is a programming middleware we created that is exposing a lower abstraction level compared to Portals. Yet, the primer focus of Kompact is performance and not flexibility and trust. Portals contributes a novel hybrid message-passing programming model, combining the strengths of the Actor model and component models. It does so by providing developers with the choice of three different messaging semantics at a fine grained level, so that every algorithm can utilise the best available guarantees and avoid inefficiencies arising from semantic mismatches. Additionally, Kompact's embedding in the Rust language produces efficient, statically typed, native code, produced and optimised for a variety of target platforms by the popular LLVM compiler backend.

At a high level, Kompact provides a light-weight process abstraction with exclusive internal state—called a component—that can communicate with other components by exchanging discrete pieces of information, called messages or events. As in other message-passing models, such as the Actor model or the Kompics component model, this design allows services written in Kompact to scale with load. It does so by efficiently utilising modern multi-core architectures, without introducing bottlenecks such as lock contention, which are common in shared-memory programming models. However, Kompact's reliance on static typing and the benefits of Rust's powerful memory management and compiler optimisations, have allowed it to show performance benefits over other state-of-the-art message-passing frameworks of up to 27×.

Not only does the component-based design of Kompact avoid scalability bottlenecks, it also provides the means for modular composition of abstractions and subservices into larger services and subsystems. This approach allows the Arcon runtime, which is built on top of Kompact, to remain both maintainable in the long term and flexible enough to allow future research on different subsystems by simply exchanging them as desired. Additionally, modularity allows testing and verification of the code at different levels of abstraction. This benefit, together with the compile-time guarantees of static typing, helped us to develop and maintain correct implementations of the set of reliable distributed services that make up unified stream and batch runtime, such as Arcon.

OmniPaxos

Continuous deep applications are data applications and it is typical for data applications to be replicated in cloud environments using consensus and state machine replication. In the last five years, during the duration of CDA we noticed that major cloud service disruptions (e.g., cloudflare) were caused by the inability of existing consensus protocols to withstand network partitions. To that end, we created, tested and designed OmniPaxos, a new protocol and system that quickly became the most robust method for data replication to date. Furthermore, OmniPaxos was presented at EuroSys '23 [Ng, Haridi, Carbone '23], a premier systems venue with an acceptance rate of 16%.

OmniPaxos achieves strong resilience through a decoupled design that distinctly separates the execution and state of leader election (primary server) from log replication. The leader election component of OmniPaxos is built upon the concept of quorum-connected servers, with a primary focus on ensuring connectivity. This decoupled design allows for a more flexible and efficient system, as it separates the concerns of leader election and log replication, each with its unique challenges and requirements.

Furthermore, OmniPaxos introduces a decoupling of reconfiguration from log replication. This innovative approach allows for flexible and parallel log migration, significantly improving the performance and robustness of reconfiguration. The decoupled design enables the system to handle reconfiguration and log replication independently, thereby enhancing the overall efficiency and resilience of the system.

Our evaluation of OmniPaxos highlighted two key advantages over existing state-of-the-art protocols. First, it guarantees recovery in at most four election timeouts under extreme partial network partitions. This feature ensures the system's robustness and reliability, even in challenging network conditions. Second, OmniPaxos achieves up to 8x shorter reconfiguration periods with 46% less I/O at the leader. This improvement significantly enhances the performance of the system, reducing the time and resources required for reconfiguration.

UniCache

In the context of improving data replicated services we also investigated the use of AI for adaptive data compression between primary and backup servers. Our investigation led to the creation of UniCache [Ng, Wu, Carbone '23], an innnovative method and software which was presented at EDBT '23.

UniCache leverages the fact that each replica has access to a consistent prefix of the replicated log. This access enables them to construct a uniform lookup cache, which is used for consistently compressing and decompressing commands. This innovative approach significantly reduces the data transfer load on the primary server, thereby improving the overall efficiency of the system.

In application workloads with a skewed data distribution, UniCache achieves effective speedups and reduces the primary load. Our experimental studies demonstrate a low pre-processing overhead for UniCache, indicating its efficiency and practicality for real-world applications.

Moreover, UniCache exhibits the highest performance gains in cross-data center deployments over wide area networks. This finding suggests that UniCache is particularly beneficial in distributed systems where data must be transferred across different geographical locations.

In conclusion, UniCache offers a promising solution to the data transfer burdens associated with replicated state machines. By introducing a learned replicated cache, UniCache enhances the efficiency and performance of cloud data service workloads, particularly in cross-data center deployments.

Scalable File Systems and Serverless Streaming

The project also contributed to the design of scalable distributed storage/file systems that now runs as a managed system in cloud infrastructures including multi-datacenter clouds. The main work is on the HopsFS distributed file system for storing and managing large data sets. In particular one work in the project is to extend HopsFS to efficiently manage and store files of different varying sizes. Traditional cloud-based storage systems only efficiently manage files consisting of large block sizes. This is not optimal as a backend to ML training systems that for example in image recognition training where an image is a moderately small file. The work results in publications at ACM Middleware 2018 [Niazi et al 18], and a Ph.D. thesis by Salman Niazi. The work is integrated in HopsFS currently part of product portfolio of the start-up LogicalClocks[1].

Another joint research direction we conducted together with Hopsworks was the investigation of externalizing state in streaming applications to scalable and fast NoSQL databases. A series of four MSc thesis projects was conducted which led to a new contributed state backend for Flink based on RonDB. Furthermore, the concept was further exploited to research the use of external state for low-latency feature store databases that have become prevalent today.

[1] `https://logicalclocks.com`

2.2 PARTICIPATING RESEARCHERS

The CDA project has funded a large group of researchers in computer systems, programming languages and machine learning focusing on continuous and deep analysis. The complete list of core researchers that participated in the project can be found in Appendix A3. The resources of the project were shared between KTH and RISE researchers. Non-senior researchers and engineers had been at RISE whereas senior researchers were employed at KTH. RISE consumed slightly higher than KTH due to the surplass of senior researchers as well as employing all the non-doctoral engineers which played a crucial role in the project. Several employees were re-located from RISE to KTH when they graduated, e.g., Paris Carbone, Lars Kroll, Max Meldrum, Klas Segeljakt and Harald Ng.

The recruitment process for graduate students was staled during the COVID period (2022-2021) and as a result most doctoral students initiated their involvement late in the project (which justified a deadline extension). The competition has also been fierce in the recruitment process, however, the reputation and success of the group managed to attract top candidates. One example is Sonia-Florina Horchidan who left her position at Google to join the group as a doctoral researcher in the hot field of continuous deep analytics.

Internally, our onboarding process involves familiarizing potential MSc candidates through internships and thesis supervision from related programs at KTH (e.g., distributed systems engineering and data science). We strived to achieve a balanced gender ratio at our best effort which was at 30% at a MSc level and 16% at a doctoral level. To compare, on average the female student ratio at our respective MSc programs is consistently below 20% over the last 5 years.

2.3 EVENTS AND ACTIVITIES

The CDA project led to the organisation of a plethora of events and activities. We refer to Appendix A5 for more information.

3 Strategic Relevance

The CDA project established the creation of a new joint group between RISE and KTH that eventually became one of the most strategic hubs today in Sweden and Northern Europe in the field of data systems, spanning the creation of a dozen production-grade software systems as well as numerous high-impact publications and awards.

3.1 INDUSTRIAL UTILIZATION

The CDA research met unprecedented outreach through various strategic collaborations and partnerships with industry leaders and stakeholders. These partnerships have facilitated the transfer of knowledge and technology from the research project to practical, real-world applications.

Firstly, we have ensured that the research findings and the developed technologies are accessible to the industry. This has been achieved through regular presentations at industry conferences, publication of our findings in widely-read industry journals (e.g., VLDB), and direct communication with industry partners in Sweden and beyond.

Seif Haridi, led the adoption of several of our systems in the Hopsworks company that specializes in AI software infrastructure and collaborates with several dozens of customers from around the world, including healthcare institutes (e.g., Karolinska).

Henrik Boström led several collaborations on applying ML methodologies to real time data both at Karolinska as well as at Scania, primarily looking into the application of ML inference in the context of optimising routing and battery efficiency in transportation.

Paris Carbone led the adoption effort in four avenues: 1) The Digital Futures network which hosts as of 2023 a wide range of strategic industry partnerships include both big (e.g., Ericsson, SAAB, Astrazeneca) and smaller companies which resulted in collaborations for the Portals project. 2) Wallenberg AI network with active participation at the Wallenberg Research Arena activities which resulted in several active collaborations for the Orb and OmniPaxos projects, 3) The CASTOR network which resulted into collaborations with critical product development at Ericsson and SAAB in the context of the ArcLang and Arcon projects and finally, 4) The KTH Centre for Data Driven Health (CDDH) which spiralled new collaborations with healthcare institutes around the globe in the context of Orb and graph representation learning.

Secondly, we have also actively sought opportunities for collaboration with industry partners. These collaborations have taken various forms, including joint research projects and internships for PhD students. These collaborations not only provided opportunities for our research to be applied in an industrial context, but also allowed our group members to gain valuable experience in the industry.

Finally, we have also explored the potential for commercialization of our research results. This has involved conducting market research to identify potential applications for our technologies, and seeking advice from business development experts to develop a commercialization strategy as well as a recent patent application for the Orb DB project.

3.2 COLLABORATION STRATEGY

Our Wallenberg network participation had been fruitful, primarily for the projects conducted by our WASP-affiliated PhDs Sonia Horhidan and Harald Ng. The OmniPaxos project attracted the SAAB-owned company Combitech which asked our team to adapt the OmniPaxos algorithm to their vehicle state management problems and boost their reliability. The project will enter its demonstration phase later in 2023. Furthermore, the Orb DB project attracted the interest of several WASP-affiliated companies such as Adledge and Scandinavian Airlines which discussed use-cases that would benefit from the inference power of the Orb database. The Orb team is currently working on addressing these use-cases in their upcoming release.

During 2022 we discussed the prospect of using our multi-stage compilation technology developed in CDA to solve safety- and time- critical system problems at industries of core infrastructure such as

Swedish telecommunications and defence. Paris Carbone discussed these avenues together with Ericsson and SAAB in the context of his leadership role in the CASTOR area of “data-driven development”. The results of those discussions led to a new strategic collaboration under the Vinnova FASTER-AI project, led by KTH (Paris) which aims to apply our expertise in order to infuse AI-driven decision making in critical hardware of specialized use. The project was granted a total of 16MSEK to satisfy this goal by 2026.

The vision of Orb DB attracted healthcare professionals at KTH CDDH as well as Karolinska and Massachussets General Hospital and led us to form new collaborations which took the form of submitted EU and other proposals. The aim is to apply our newly developed graph reasoning technology to optimise healthcare management as well as to provide more insights for the use of advanced clinical trials that are rooted on the patient data.

During 2022, Henrik Boström collaborated actively with Scania CV AB, through the Vinnova-funded project RAPIDS (ref. no. 2021-02522) on the task of predicting with confidence. Although the project, which is lead by Scania, focuses on partly different questions to CDA, such as learning models for predictive maintenance from censored data, the underlying algorithms and package for conformal prediction developed in CDA, have been a starting point for the RAPIDS project.

Henrik Boström has collaborated with the Swedish National Financial Management Authority, through the Vinnova-funded project ”Datalabb Resultaten i staten” (ref. no. 2019-02252), during 2019-2021. Henrik Boström’s work focused on explainable machine learning for time-series forecasting, with an application to predicting the gross domestic product of Sweden. Apart from a publication, which is summarized in section 2.1 as it is also a contribution of CDA, he contributed with the Webb application ”Förutse BNP” (”Predict GDP”), available at https://datalabb.esv.se/bnp/

Since 2022, Henrik Boström is collaborating with Karolinska University hospital and Karolinska Institute, through the project ”CLINOMICS - Data-driven precision medicine: from discovery to clinical decision support systems” (HMT, ref. no. FoUI-966658), where techniques for explainable machine learning will be developed and applied to dicover risk factors for unsuccessful prescription of antibiotics for sepsis, so called ”drug-bug mismatch”.

During 2021, Henrik Boström was responsible for a commissioned research project on using explainable machine learning for characterizing customers that are about to leave (churn prediction) with funding from the research and development fund of the insurance company Länsförsäkringar (ref. no. P1.21). From 2022, he is leading a new commissioned research project for the same agency (ref. no. P3.22) on developing techniques for identifying risk factors for insurance claims to allow customers to take preventive actions.

3.3 EXPLOITABLE RESULTS

As listed in Appendix A6 the CDA project has been highly fruitful. Driven by our effort to make actionable research, most of our research findings are accessible in the form of standalone software libraries ranging from TRL-2 to TRL-6. More concretely, we have 1) *Crepes*, a python data science library that democratizes the use of conformal regressors. 2) *Orb DB* (TRL-2) is a database system soon to be released which makes a novel combination of generative AI and graph traversals to answer complex deep questions fast and continuously. Orb, from its inception has attracted investor interest and we proceeded with a patent application before its release to protect its novelty from external exploitation. 3) *OmniPaxos* (TRL-6) has become a highly influential, production grade library for high availability and one out of handful systems available today for building replicated services. Around OmniPaxos we build an ecosystem of novel methods which can eventually evolve into a managed cloud product, depending on the interest. 4) *ArcLang* (TRL-4) as well as our Arcon and WheelDB software libraries are highly novel and the first released software for continuous deep analytics. They cover the needs of declaring versatile data applications combining real-time and analytical workloads and could potentially evolve into a commercial ecosystem on their own in the future. Finally, 5) *Portals* is a novel serverless open source framework that aims for high adoption following its autumn 2023 release which was jointly built with Philipp Haller (KTH), one of the co-creators of the Scala language. Portals combines several novel features of stream processing technologies we created in the past (i.e., Apache Flink) with actor models (i.e., Scala Actors by Philipp Haller).

Finally, our close collaboration with Hopsworks in adapting stream processing technology with noSQL databases (MSc theses by Haseeb and Kumar) led to a new prototype of a managed streaming service based on Flink which provides zero downtime when adjusting its scalability or recovering from failures.

3.4 INTELLECTUAL ASSETS AND PROPERTY RIGHTS

As listed in Appendix A7 the CDA project led to a patent application in spring 2023 that aimed to protect the intellectual assets of Orb DB, including its hybrid query execution and optimisation design.

While no other IP was registered in the context of CDA, we expect several of our software libraries to increase in adoption and justify the potential for more trademarking. Those include: Crepes, ArcLang, Portals and OmniPaxos among others. To that end, we are in talks with several innovation offices (KTH, WASP, AI Sweeden, RISE) in Sweden as well as VC funds in the USA to assess the readiness of the products, their teams and market for such initiatives.

3.5 INDUSTRIAL SHORT-TERM ADOPTION

Several of our systems and libraries are already in use today or planned to be in the following months:

Crepes is currently in use as well as being further developed at Scania, in the context of the RAPIDS project (see section 3.2).

The OmniPaxos team currently finalizes an integration of the library with file replication protocols used at Combitech, that is a SAAB / Wallenberg-funded company.

Our scalable file system inventions (2.1.2) are currently in use in the last three years in production at Hopsworks, supporting their feature store products.

An pre-alpha version (TRL-3) of Orb DB is planned to be developed within a year in support of use-cases discussed with several WARA-affiliated companies.

3.6 PUBLIC OUTREACH

As outlined in our activities in Appendix A5, CDA had a wide outreach spanning across different audiences of varying background.

Several tech talks given at influential industrial events such as NDSML, Flink Forward, FOSDEM and Scala Days aimed to attract the attention of practitioners as well as expanding our network and incresing impact. This way we also started a conversation with several interested companies that want to try early prototypes of our systems and provide us with feedback and use-cases.

In addition, we used our affiliated research centers and respective channels (Digital Futures Days, CASTOR Days, WARA Events) to reach leading Swedish ICT and retail companies and disseminate our work primarily to management and senior R&D levels.

Our doctoral were encouraged to participate in doctoral workshops hosted at top tier conferences. They had several opportunities to submit their poster presentations and participate in big events while learning how to promote their research and seeking for use-case partners as well as strategic international partners. Overall we presented posters at SIGMOD, VLDB, EuroSys and more conferences with high expertise and insights in data systems.

Finally, we organised our own workshops and hackathons, where we provided access to KTH students to participate, discuss interesting problems with guest researchers and get more engaged with applied systems research problems.

3.7 RESULT UTILIZATION AND EXPLOITATION - NYTTIGGÖRANDE

Our exploitation plan gave a great boost to the CDA projects which were in highest need, helping us develop competitive proof-of-principle (PoP) software out of our research as well as identifying commercialization potential for several of our systems. Both types of these important activities were funded directly from our nyttigörande CDA sub-project. In detail, we made the following utilizations:

We managed to make a PoP implementation of the ArcLang and Arcon projects. This required specialized expertise in multi-stage compilation and additional efforts to integrate diverse software libraries together into a final product. More information and documentation of the resulting proof of principle can be found here: https://cda-group.github.io/arc/

For the Portals project we achieved a market analysis and a preparation of a unique interactive demo experience which will be presented live in September 2023 at VLDB as well as other tech events we plan to use towards increasing the commercialization potential of Portals. Parts of the "Portals Playground" project are currently being released on Github : https://github.com/portals-project/playground.

Figure 7: A view of live visitors of omnipaxos.com during its first release.

Finally, we launched an effort to boost the visibility and readiness of the OmniPaxos project following its launch in May 2023. We then analyzed the demographics of visitors of the OmniPaxos project which contributed greatly to our understanding of potential users and outreach needs. For example, after noticing the first day in our live user feed (Figure 7) that China did not contribute with visitors we included a Chinese version of the website which greatly increased the outreach in Asia.

The following table summarises all utilized nyttigörande activities in our exploitation phase.

Project	**Type**
A Virtual Log service for replicated reliable applications (OmniPaxos)	Bed. av komm.
Serverless data analysis for decentralized systems (Portals)	Bed. av komm.
Live Data Exploration on Stream State (Arcon)	PoP
Standardizing a new language and code generator for analytics (ArcLang)	PoP

3.8 PATENTS

A single patent was filed individually via a granted VFT-1 Vinnova/KTH funding to protect Orb DB's distinct architecture. We further refer to section 3.4 and Appendix A7 for further information.

3.9 SPIN-OFF COMPANIES

Currently there has been one spin-off company named **Hopsworks**[2] co-founded by CDA members Seif Haridi and Salman Niazi among others. The company is currently entering a round B of investment capital and is one of the leading tech companies originating in the EU that offer AI-platform and infrastructure services.

[2] https://hopsworks.ai

3.10 STRATEGIC SIGNIFICANCE

The strategic significance of the CDA project lies in its potential to revolutionize the way we approach data analysis and management. The project's innovative research and development efforts showcased that it is possible to think of deep and continuous computation seamlessly. This brings far-reaching implications for various sectors of society.

Orb DB, with its unique AI graph database technology, carries significant strategic value for both Sweden and the EU. This technology enhances Sweden's capacity to sustain its domestic tech solutions, ensuring resilience against potential disruptions in geopolitical or value chain scenarios. Orb's ability to produce dependable data with modifiable accuracy and resilience provides a degree of stability that could be vital during crises. By relying on homegrown, reliable deep tech like Orb DB, Sweden can secure its technological infrastructure, lessening reliance on foreign systems such as OpenAI or Google, and minimizing potential risks. Furthermore, Orb's capability to autocomplete knowledge in graph form could stimulate innovation, supporting Sweden's strategic objectives of promoting its digital economy and competitiveness in the AI market.

OmniPaxos and **Unicache** are another example of software with strategic significance. OmniPaxos, with its resilience to partial network partitions, offers a level of stability and reliability that is crucial in today's interconnected world. Its decoupled design enhances the efficiency of state machine replication, a fundamental component of many distributed systems. By reducing the impact of network disruptions, OmniPaxos can contribute to the robustness of national technological infrastructures, reducing dependence on potentially unstable foreign systems and mitigating potential threats. UniCache, on the other hand, addresses a significant challenge in the realm of cloud data services - the heavy data transfer duties imposed on primary servers. By introducing a learned replicated cache, UniCache offers a solution that not only enhances the efficiency of data transfer but also reduces the load on primary servers. This can lead to improved performance of cloud data services, a critical aspect of the digital economy.

Finally, our core CDA data technologies such as **ArcLang, Arcon and Portals** offer a robust and flexible solution for data service development that was proven to be strategically relevant for high-tech companies like SAAB and Ericsson since it attracted their interests. These systems' contributions to enhancing the efficiency and reliability of service development are poised to make a lasting impact on these companies and hopefully the tech industry as a whole.

4 The Graduate Training of the Project

The CDA project incubated the creation of an active “data systems” research lab that quickly gained traction within and beyond KTH and RISE for its graduate training.

4.1 GRADUATE DEVELOPMENT AND TRAINING

Within the focus of the CDA project we organised a series of courses, international events and refined previous courses to incorporate our CDA research interests. Our complete training activities are included in Appendices A8 and A5.

Graduate Courses: From the kickoff event of the CDA project in 2017 we initiated a weekly seminar, reading group and discussion series on **“Continuous Deep Analytics”** that is still held until today on every Friday. The Weekly seminar participation list expanded greatly throughout the years from 5 members to 54 active members. The seminars had the form of a short lecture and discussion on a research topic overview or a paper discussion focusing on our CDA topics. These events were hosted by a different researcher each week, where we prioritized graduate students to present in order to develop their communication and analysis skills. The seminars contributed to 7.5 credits per graduate student for their efforts and developed skills.

In addition to the main CDA seminar series where everybody participated, we also initiated a set of 3rd cycle courses which were more focused on concrete topics. This includes a course on “advanced databases” and a course on “ML for Graphs”. These two courses were based on online material available by leading researchers and teachers at CMU and Stanford University. To justify a full-year 3rd cycle course, we also enriched these course materials with papers and a research project each, all of which granted 7.5 credits to each student. The courses were used as a vehicle to bring our graduate students up to date with the state of the art in these technologies early in their studies. This also helped the overall expertise and interests in the group to grow, get familiar with the data management and systems community and familiarize with the mindset and methodologies used in these research fields.

Development Meetings: In addition to the CDA meeting we also hosted another weekly discussion series focusing primarily on software development. There, all graduate students had the chance to 1) present their weekly progress, 2) bring up technical issues they faced and discuss all together for solutions and 3) ask for feedback and criticism of new ideas or tools that came up during their development process. We started this series in 2021 and we noticed several positive results. 1) We separated software methodology concerns from our research meetings (e.g., CDA seminars), 2) Boosted collaboration and team spirit between graduate students and 3) Managed to overcome software development hurdles faster than before. In Figure 8 we show an example of a development meeting with both student and their mentors.

Summer Schools: In the context of CDA we also organised summer schools. One example is the summer school on “Linked Data Analytics and Big Data Management” which was organised by Sarunas Girdizjauskas , 12 – 15 September, 2022, Stockholm, Sweden. This included seminars, projects and student competitions on our core CDA topics with a focus on the graph representation learning problems and solutions.

KTH Courses Redesigned: We redesigned several of the KTH courses that we are responsible for in order to fuse CDA-related research interests and state-of-the-art. This included two highly popular KTH courses : 1st cycle IV1351 Data Storage Paradigms, and 2nd and 3rd cycle ID2203 Advanced Course in Distributed Systems. The content we added corresponded to lessons learned from our system development, and insights into data technologies such as dealing with data skew, or data replication more efficiently to compensate for deep analysis without sacrificing latency. We received a highly positive feedback from the students and great enthusiasm, more than ever before, to join our research team for their MSc theses in the CDA topics.

MSc theses Supervision: Several of our core studies in CDA were conducted by MSc thesis student interns at RISE and KTH. The doctoral students had the chance to test several hypotheses early in their studies via the help of their MSc interns and also learned valuable lessons on how to properly supervise diverse groups of students with different talents and interests. PhDs also had the chance to participate

Figure 8: Weekly CDA Development Meeting with students and mentors.

on several student fair events (Figure 13) with CDA posters that grasped the attention of hundreds of thesis students looking into an interesting topic and appropriate group with industry impact.

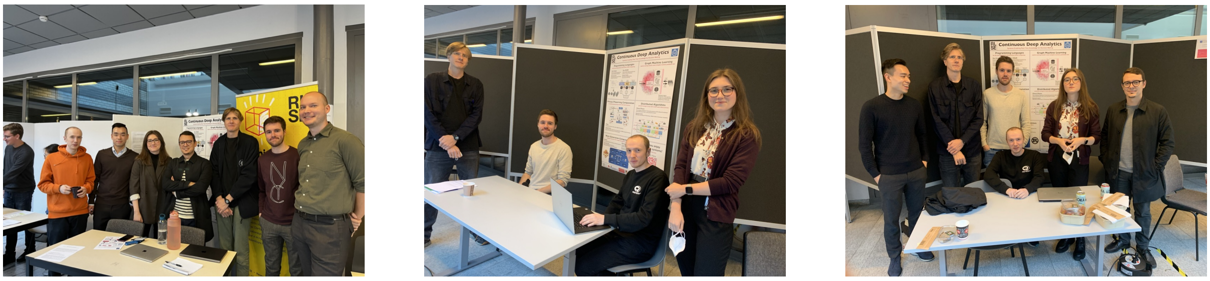

Figure 9: Highlights from our CDA doctoral student representatives at KTH student fairs

MooC Development: in January 2020 we also had the chance to collaborate on the development of a high-influential MOOC series on data stream processing technology .[3] implemented with TUDelft named “Taming Big Data Streams”. Paris Carbone from CDA contributed to several video seminars and content on Stream State Management and Reliability for continuous data processing.

4.2 THE CDA GRADUATES

As listed in the Appendix, the CDA project had overall five graduates: Paris Carbone, Lars Kroll, Jingna Zheng, Salman Niazi, and Theodore Vasiloudis who defended their theses between 2018 and 2020. Our initial plan was to hire the next generation of graduate students early in the project, however, we completed our hiring as late as 2021. This means that, despite the large number of affiliated doctoral students in CDA, only half have so far graduated. According to our table in Appendix A9 as well as A11 six out of eleven doctoral students are still completing their doctoral studies. One student (Klas Segeljakt) is expected to finish soon in 2023, whereas the rest are planning to defend until 2026.

Graduate Independence: Among the complete graduates we have good examples of leadership with

[3]https://online-learning.tudelft.nl/courses/taming-big-data-streams-real-time-data-processing-at-scale/

distinctions in their post-graduate career following their involvement in CDA. Lars Kroll continued one more year as a postdoc researcher and engineer in the project at RISE with crucial contributions to the project's middleware. Eventually, Lars Kroll moved to Databricks, a leading company in data technologies where he serves as a senior software engineer and architect within Databrick's cloud storage platform. Salman Niazi, is one of the founding members of Hopsworks, a leading AI company in the EU and a core engineer of their platform. Furthermore, Theodore Vasiloudis is leading a research team under Amazon Research, residing in Seattle USA with a focus on graph representation learning and data management. Jingna Zheng, is helping Volvo Cars with cutting-edge software technologies and expertise she acquired during the CDA project. Finally, Paris Carbone, is leading the Data Systems Lab at KTH and RISE which expands and develops further all distinct technologies and research we were able to achieve during the CDA project. Furthermore, as an independent researcher and Asst Professor Paris contributes to the development of CASTOR, Digital Futures and other newly founded research centers with a focus on industrial and societal research impact.

5 Collaborations

5.1 INTERNAL COLLABORATION

The CDA group helped creating a cluster of expertise across four distinct groups between KTH and RISE. Collaborations within CDA have been fruitful in the following ways:

- Helped speeding up knowledge transfer within joint discussions in weekly meetings.
- Direct communication with a highly active dedicated mailist (~ 60 members) and Slack channel (~ 100 members).
- Joint supervision of M.Sc. and Ph.D. theses across RISE and KTH.
- Joint publications of CDA researchers (e.g., work in Block-distributed Gradient Boosted Trees with Theodore Vasiloudis and Henrik Boström).
- Joint workshops with all CDA researchers.

The fusion of ML and Systems experts across KTH and RISE created a new perspective and a series of research collaborations in the new trending field of Systems for ML (or ML for Systems). This would not have been possible without a close interaction achieved in the CDA project. Apart from knowledge transfer between the two groups, significant results were also achieved through joint supervision e.g., Sonia Horchidan co-supervised by Paris Carbone and Henrik Boström and Dr. Theodore Vasiloudis' Ph.D. work that was supervised by Seif Haridi and Henrik Boström.

5.2 CROSS-DISCIPLINE COLLABORATION

The CDA project also showcased the cross-discipline nature and importance of continuous deep analysis and its applications.

CDA in Healthcare: Our work with Orb DB quickly became popular with the Centre of data driven health (CDDH) that is is a multi-disciplinary arena for research and development in health data. Their mission is to create an infrastructure solution to some major unsolved issues for health care, research, citizens, and society. After several discussions with the CDDH team we decided to form a wider expert group at KTH and apply the inference capabilities of Orb DB developed in the CDA project in order to help with data integration across hospitals, a highly important issue. To that end, we discussed use-cases and wrote two pending project proposals, one EU proposal together with US partners and hospitals alongside Karolinska, as well as a MedTechLabs call of Region Stockholm. Furthermore, we started new collaborations at the level of PhD students, exchanging ideas in graph representation learning and data integration problems. Another example of CDA in healthcare is the project "CLINOMICS - Data-driven precision medicine: from discovery to clinical decision support systems" (HMT, ref. no. FoUI-966658). This has been a joint effort, during 2022-2023, between Karolinska University hospital/Karolinska Institute and Henrik Boström at KTH based on applied results from the CDA project in conformal prediction topics.

CDA in Robotics and Edge Computing: A good example of direct collaboration with industry was our collaboration with Combitech. During one of our poster events at WASP, Harald Ng's presentation of his work on OmniPaxos grasped the attention of Combitech, a modern company working with drones, decentralized vehicular control and robotics. They were in high need of a reliable middleware to replicate data commands across remote devices and were willing to try the OmniPaxos library for such synchronization purposes. Our team made a tailored prototype to solve this problem based on the OmniPaxos protocol.

CDA in Critical Infrastructures: Our advanced programming language and compiler technologies in ArcLang and Arcon are currently in use in the Vinnova FASTER-AI project, led by KTH (Paris) which embeds AI-model logic within critical hardware of specialized use at SAAB and Ericsson. The project was granted a total of 16MSEK to satisfy this goal between 2023 and 2025.

CDA in Sustainability Research: The work presented in (Karlsson et al, 2020) on classification of aerosol particles using inductive conformal prediction was a collaboration between the Department of Environmental Science at Stockholm University and the School of Electrical Engineering and Computer Science KTH Royal Institute of Technology.

CDA in Economics: The work presented in (Boström et al, 2020) showcases explainable multivariate time series financial forecasts which were researched within the CDA project. This was a collaboration with researchers/practitioners in economics at the Swedish National Financial Management Authority.

(a) KTH and Boston University collaborators at SIGMOD 22 - Philadelphia USA.

(b) Jonas Spenger discussing the Portals vision with prog. language experts at SPLASH 22 in Oakland NZ

Figure 10: International collaboration highlights in CDA

5.3 INTERNATIONAL COLLABORATION

Several of the results of the project would not have been possible without the feedback and close collaboration with other research institutes internationally. Below we highlight a few collaborations that were established or reinforced within the scope of the CDA project.

Stanford University: Our work on ArcLang's intermediate representation (IR) support for batch and stream analytics began as a collaboration with Stanford University and Matei Zaharia's group (creator of Apache Spark). Via SSF's support, our researcher Lars Kroll from KTH spent two weeks at Matei's research group understanding and extending Weld, Stanford's novel IR for optimising data pipelines with stream capabilities.
TU Delft: Several of the goals of CDA align with the vision of Web Information Systems (WIS) research group at TUDelft. The expertise at TUDelft has complemented ours with use cases and ideas in problems related to data management. Several results of our collaboration with Asterios Katsifodimos' group include a joint MOOC, a newly published survey (on Stream Processing Systems) and a Tutorial as SIGMOD 2020 on steam processing while two joint papers are also currently under submission.
Boston University: Our continued collaboration with Vasiliki Kalavri (currently Assistant Professor at BU) has led to significant results in the intersection of stream technology and graph analytics. Via a series of joint MSc projects co-supervised by our two groups we have managed to invent novel techniques for graph analytics including partitioning techniques (published at VLDB 2018) and AI for graphs (two papers accepted and three under submission).
Imperial College London: The Large Scale Data and Systems group at ICL (Peter Pietzuch, Holger Pirk) has a long history and expertise in systems for hardware accelerated analytics. From the early beginnings, we co-organised joint workshops and established new collaborations between PhD students of which many results have come to materialise today. A core example of this collaboration is the design of a low-level representation for incremental analytics for which we also plan to get EU (ERC) funding in the upcoming months along joint publications.
TUBerlin: Our collaboration history with TUBerlin (DIMA group led by Volker Markl) has its roots in our joint efforts in the Stratosphere and Flink system (SSF E2E project) as well as EU projects and EIT programs that we jointly participated in. To this day we still maintain close collaboration, especially in the prospect of combining our efforts between the Nebula Streams project (IoT streaming at TUBerlin) and CDA via the use of our work on IRs and code generation in the scope of sensor networks. In the

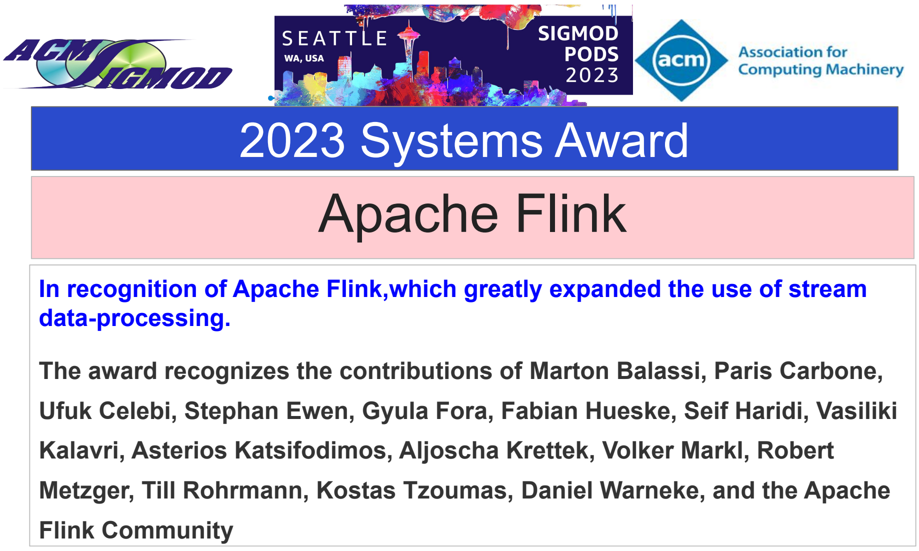

Figure 11: A highlight from the planned award ceremony at SIGMOD 2023.

latest results from the DIMA group (SIGMOD 2020) it is already mentioned that ArcLang's IR will be a potential candidate in the future to boost performance in the Nebula Streams platform.
University of Edinburgh: Our showcase application on online graph embeddings is done in collaboration with PhD student Massimo Perini from UoE under the co-supervision of Milos Nicolic's group that focuses on algorithms for stream analytics and ML. This work is planned to be submitted and published in 2021. In addition, our upcoming work on iterative dataflow models is also done under the co-supervision of Pramod Bhatotia from UoE.

5.4 COLLABORATIONS

As mentioned in Section 3.10 we have disseminated our efforts in CDA to CASTOR and Digital Futures. This, in combination with strategic partnerships at RISE helped us reach data engineers and analysts in the industry and get exposed to real needs early in the project. As of today, there have been numerous industrial collaborations in the form of M.Sc. thesis co-supervision for the following projects:

- Reliable External State Management for Dataflow Backends – Co-supervised by **Hopsworks** and co-funded by **Google Research** (Research Cloud Grant) – Hasseb Asif
- Externalising Dataflow State to NoSQL Databases – Co-supervised by **Hopsworks** and co-funded by **Google Research** (Research Cloud Grant) – Sruthi Kumar
- Real-time Financial Anomaly Detection – Co-supervised by **Swedbank** – Anna Martignano
- Window Aggregation Algorithm Repository[4] - A cooperation with **IBM (Watson) Research**.

Finally, to avoid duplication please refer to section 3.2 on the collaboration details with Ericsson, SAAB, Combitech, Scania CV AB, the Swedish National Financial Management Authority, Länsförsäkringar, Karolinska University Hospital and Karolinska Institute.

5.5 NEW RESEARCH DIRECTIONS

The CDA project helped our group get a thematic identity that is both interesting and relevant. In essence, it fused ML expertise into our systems knowledge and vice versa, benefiting both ML and systems groups.

With CDA we exposed our researchers to hot topics early on, before they eventually became popular such as graph representation learning, AI-for-Systems, and multi-stage compilation.

Through CDA we established deeper and long-term collaborations with the industry such as Scania

[4] https://github.com/IBM/sliding-window-aggregators

(Henrik Boström) and Ericsson and SAAB (Paris Carbone), among others. Furthermore, we also discovered novel and exciting new areas that did not exist before such as inference-based databases (Orb) and the new area of meta-coordination systems which was inspired by our studies on OmniPaxos and Unicache. We believe that these areas will establish more demand and academic interest in the near future.

5.6 AWARDS

As listed in Appendix A13, the CDA project attracted several prestigious awards. The SIGMOD Systems Award [5] which was given to Paris Carbone and Seif Haridi from the CDA project is a rare achievement. This award was granted (Figure 11) for all contributions to Apache Flink from its conception until today. We are the first non-US researchers that receive this big price which was previously given to the creators of the widely used Postgres and Spark systems, among others.

Best paper awards were granted to three works of Henrik Boström as well as a paper and presentation award to Sonia Horchidan and Paris Carbone.

[5] https://sigmod.org/2023-sigmod-systems-award/

6 Project Continuation

6.1 MAIN LESSONS LEARNED

The achievements of the CDA project are unprecedented for our group. Within 5 years we were able to set up an impressive portfolio of innovative projects, start new collaborations and strengthen our partnerships. Several takeaways from the CDA project include:

- Interdisciplinary collaborations are the most fruitful. Having complementary expertise such as ML and Systems in CDA was an important success-factor.
- The importance of industry partnerships cannot be overstated. These collaborations not only provided practical insights but also facilitated the application of CDA in real-world scenarios.
- The increased focus on PhD training in CDA was beneficial for both the students and the project. The students gained valuable experience and contributed fresh perspectives, while the project benefited from their research efforts.
- Despite our successes, we faced challenges primarily during the COVID19 period. In terms of hiring people, we experienced a 1.5 year delay which in-turn delayed further our efforts. We plan to complete all hires in year 1 in future projects.

6.2 WHAT WILL HAPPEN TO CDA

There are two types of results in the CDA project, 1) Prototypes and Methods and 2) Production-Ready Open-Source Systems. For the latter, we foresee the continuation of further development for at least three to five years after the CDA project ends, given the potential for applying for incubation, e.g., in the Apache Foundation as well as commercialization. More concretely we plan to continue the development of the Kompact and further extend its capabilities towards edge computing and sensor networks as well as building higher-level services such as decentralized serverless capabilities. Arcon is already the first system of its kind when it comes to hybrid batch and stream analytics and it is expected to serve hundreds of industrial use cases in data analytics that aim for high performance. Thus, due to the development demands we plan to build a user group and an active development community that can take over our efforts and expand Arcon's ecosystem such as new connectors to external systems and message queues as well as code generators for new hardware. As we move forward, we plan to build upon this foundation and continue to advance our research. Here are some of the steps we plan to take:

- **Continued Research:** The research conducted under the CDA project has opened up several new avenues for exploration. We plan to continue our research efforts, delving deeper into the topics we've already explored and branching out into new areas of interest.
- **Collaborations and Partnerships:** We will continue to strengthen our existing collaborations and seek out new partnerships, both within academia and industry. These collaborations will provide opportunities for knowledge exchange and potential funding sources.
- **Commercialization:** Where applicable, we will explore the potential for commercializing our research results. This could provide a sustainable source of funding for our research activities and contribute to the development of new technologies in the field of data analysis. Primarily we look into Orb DB and its related IP, however, we also plan to extend our activity based on the market analyses we could make within CDA.
- **Education and Training:** We will continue to involve students in our research activities, providing them with valuable experience and contributing to the development of the next generation of researchers in our field. We will also seek opportunities to incorporate our research findings into educational programs and training courses at KTH.

6.3 FIELD LONG-TERM PERSPECTIVE

Continuous Deep Analytics is currently not just a KTH project but a global trend. Already in 2023, with the rise of generative AI and LLMs we can observe that society is catching up at a much faster pace than ever before. Among related trends:

Graph representation learning technologies are becoming more popular than ever. Gartner estimates that 80% of data technologies will use graphs by 2025 (10% in 2021). Insight Partners say that "The Graph DB market is expected to grow to 7.22Bil by 2028 (1.85B in 2021)". In fact, the Graph Market could subsume Relational DBs. LLMs have already shown their fundamental limitations in reasoning based on ground truth. Yet, GNNs circumvent that through a richer mapping of ML representations and ground truth.

Uncertainty estimation has been gaining traction in research due to the increasing recognition of its importance in the deployment of AI systems in real-world applications. As AI models are increasingly used in decision-making processes, the need for understanding and quantifying the uncertainty associated with their predictions has become more apparent. This is particularly relevant in high-stakes domains such as healthcare, finance, and autonomous vehicles, where incorrect predictions can have significant consequences.

Uncertainty estimation provides a measure of the model's confidence in its predictions, which can be used to assess the reliability of the model's outputs and make informed decisions about when to trust the model and when to defer to human judgment. This can lead to safer and more reliable AI systems. Moreover, uncertainty estimation can also be used to identify areas where the model lacks sufficient training data, guiding data collection efforts to improve the model's performance. It can also be used to detect out-of-distribution inputs, which are inputs that are significantly different from the data the model was trained on and for which the model's predictions are likely to be unreliable.

Language support for diverse data analysis has been a hot open topic. Our ArcLang language combines a Python-level intuitive programming semantics with the performance of finely tuned C programs. Recently, a group of tech professionals such as the creator of the Swift language and contributors of the LLVM infrastructure joined forces to create a new AI application language, called Mojo [6]. Despite this being a language of a different purpose, the methodologies behind Mojo are almost identical to those of ArcLang, both have high level, user-intuitive semantics and the compilation to high-performance binaries adopts the same techniques we used in ArcLang such as multi-stage compilation. This is proof that ArcLang is on the right trajectory which just now started to unravel.

[6] https://www.modular.com/mojo

7 Budget of the Project

	Year 2017	Year 2018	Year 2019	Year 2020	Year 2021	Year 2022	SUM
PI salary	174 125	378 941	342 875	343 257	572 086	147 752	1 959 036
Seniors salary	360 475	710 127	673 716	1 452 631	1 897 806	1 522 450	6 617 205
PostDocs salary		694 953		787 769			1482722
PhDs salary	658 117	2 028 714	1 066 054	558 797	2 132 439	2 458 177	8 902 298
Other salary			793 239	1 490 923	752 094	24 952	3061208
Equipment	240 944		21 427	26 043	59 266	40 777	3061208
Material		600	26 816	13 647			41063
travel	67 414	254 883	483 536	38 678	7 700	369 206	1 221 417
Administration							
Information		1 671	2 001				3672
Consultants		52 396		59 000	20 000		131396
Other costs		13 409	2 306	6 300	12 041	16 091	50147
Sum costs	1 501 075	4 135 694	3 411 970	4 777 045	5 453 432	4 579 405	23 858 621
Overhead	441046	1447493	1 006 349	1 642 201	1 880 958	1 588 520	8006567
% OH	29,4	35,44	29,68	35	34,99		33,56
Sum incl. OH	1 942 121	5 583 187	4 418 319	6 419 246	7 334 390	6 167 925	31 865 188

Figure 12: Summary of the annual economic CDA reports

The table in Figure 14 summarises the budget consumed until 31/12/2022 which corresponds to the full length of the project, excluding its exploitation. The budget increases per year are attributed to new hires and senior-member promotions in the group. For example, in 2018 we had one additional Ph.D. hire and one promotion, while in 2019 we had two graduations and one more Ph.D. hire. In 2021 we reached the highest number of active PhDs supported by the project with another two new hires.

As shown we reserved twice the estimated budget for the second half of the project ($20\,066\,373$ SEK), since according to our previous experience most research result exploitation and critical development efforts are typically required towards the end of the project, close to its maturity.

For the exploitation (nyttig.) budget information please refer to Section 3.7 and the budgets in the table below. We did not utilize the total 3% since we decided to concentrate on these four utilizations and skip the Hopsworks exploitation project in favor of time and focus prioritization.

Project	**Type**	**Amount**	**Host**
A Virtual Log service for replicated reliable applications (OmniPaxos)	Bed. av komm.	150000	RISE
Serverless data analysis for decentralized systems (Portals)	Bed. av komm.	150000	RISE
Live Data Exploration on Stream State (Arcon)	PoP	240000	RISE
Standardizing a new language and code generator for analytics (ArcLang)	PoP	300000	RISE

8 External Information and Activities

Alongside conference presentations, we have also focused major efforts into communicating the benefits that the CDA project will bring to society through a series of workshops, tutorials and technical seminars in popular research and industrial venues.

International Tutorials: In 2020 we got granted by the ACM SIGMOD committee a 5-hour tutorial session [7] to talk about the evolution of stream processing and related data technologies. The event was co-organized together with our collaborators Asterios Katsifodimos from TUDelft and Vasiliki Kalavri from Boston University. The tutorial included an overview of the state of the art, yet, we spent two hours discussing future directions of particular interest to continuous deep analysis. The discussions also included a hosted panel session with leading researchers and industry practitioners who were invited to spark new novel ideas. The tutorial has been fruitful for the CDA project, first to promote our CDA project and second, to get feedback and further discuss the best ways to move forward in this fast-evolving field of data systems.

(a) SIGMOD-DEEM 22 Awarded Presentation by Sonia Horchidan

(b) OmniPaxos Presentation in big aula by Harald Ng (EuroSys 22)

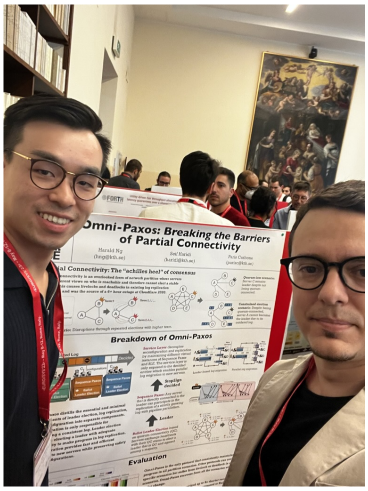

(c) Poster Presentation of Omni-Paxos at the EuroSys conference

Figure 13: Highlights from top conferences in USA and EU.

International Participation and Networking: During the CDA years we had the chance to expose graduate students to leading research networks by encouraging them to participate in international top conferences. Graduate students were able to physically attend and present posters, make paper presentations and mingle with the world's leading experts. Our senior researcher team (Seif, Henrik, Sarunas, Paris) were responsible of introducing students to leading faculty around the world and hosting several joint activities. This also included several group visits such as a full CDA Group Research Visit to Imperial College London for a Joint Workshop with Peter Pietzuch's group in Dec. 2018.

Tech Talks and Events

For a complete list of dissemination events please refer to Appendix A5.

Press

Through the duration of the CDA project we had two articles hosted at the Framtidens Forskning magazine. One was a full-page article (Figure 14) in 2019 and one half page article in 2021.

CDA Web

Project Website: `https://cda-group.github.io/papers.html`

[7] https://streaming-research.github.io/Tutorial-SIGMOD-2020/

38 © NextMedia

Realtidsintelligens för det smarta samhället

Vi står på tröskeln till framtidens smarta samhälle. Men för att de mängder data som ständigt genereras ska göra verklig nytta som kritiskt beslutsstöd måste de kunna bearbetas och analyseras i realtid. Ett forskningssamarbete mellan KTH och RISE uppfinner lösningar för att skapa nytta av framtidens strömmade data.

Lars Kroll, senior forskare vid RISE, Seif Haridi, professor i datasystem vid KTH och Paris Carbone, senior forskare och ansvarig för projektet vid RISE. Foto: Johan Marklund

Den snabbt ökande digitaliseringen genererar enorma mängder data från olika enheter. Idag lagras dessa data framför allt i stora molndatacentra, där de senare kan bearbetas och analyseras. Denna metod har emellertid stora brister. Data blir snabbt gammal och därmed har lagrad data begränsat värde som stöd för kunskapsinhämtning och tidskritiska affärsbeslut. I dagens nuvarande lösningar sammanlänkas olika hårdvaror och system av ett komplicerat lapptäcke av olika teknologier, koder och programspråk. Detta omöjliggör effektiv konstruktion av sådana system.

I framtidens snabbrörliga, digitala samhälle är miljontals olika enheter sammankopplade, inte minst genom 5G, edge-moln och sakernas internet, IoT. För detta krävs ny teknologi som gör det möjligt att fatta beslut i realtid, eller "on the fly", som det engelska uttrycket lyder. Kontinuerlig Djup Analys, CDA, är ett stort forskningsprojekt som adresserar detta. Projektet drivs i samverkan mellan KTH och RISE (Research Institutes of Sweden).

Nästa generations teknologi

Syftet med CDA är att utveckla nästa generations mjukvaruplattformar för skalbara, datadrivna tillämpningar, där information bearbetas och analyseras

realtid kunna nyttja data för beslut och strategier. Med dagens system får vi retrospektiva svar på våra frågor. Vårt mål är att skapa en teknologi för framtidens behov, med hållbara och uthålliga lösningar som är enkla att använda, förklarar projektledare Seif Haridi, professor i datasystem vid KTH.

Inkorporerar AI

CDA är en fortsättning på ett tidigare omfattande forskningsprojekt där Seif Haridi och hans team medverkade till att utveckla öppenkällkod-programvara, Apache-Flink, för analys av strömmande data, samt HOPS, som är den första kompletta europeiska plattformen för dataanalys och maskininlär-

av Uber och Lyft för dynamisk prissättning.

– CDA tar den här forskningen till nästa steg, där vi kombinerar tekniker för AI (artificiell intelligens) med traditionell dataanalys för att fatta komplexa beslut i realtid. En tillämpning som har stor relevans just nu kan exempelvis vara att på ett enkelt och smidigt sätt skapa en skalbar tjänst för avancerad och detaljerad spårning och analys av smittspridning av covid-19, berättar Paris Carbone, senior forskare och ansvarig för projektet vid RISE.

Anpassas till hårdvara

Ett viktigt fokus för CDA är att tekno-

– Det finns stora skillnader mellan olika typer av hårdvara och det är en utmaning att skapa lösningar som är agnostiska och oberoende av i vilken hårdvara de används. Vi skapar mjukvara med ett enkelt gränssnitt, som lätt anpassas till alla typer av hårdvara, säger Paris Carbone.

CDA finansieras av Stiftelsen för strategisk forskning inom ramen för programmet Big data och beräkningsvetenskap. Området har identifierats som strategiskt viktigt, där Sverige har goda förutsättningar att ta en ledande roll.

– Finansieringen från SSF

Figure 14: Highlight from Framtidens Forskning

CDA Slack Group: `https://cda-kth-sics.slack.com`
Project Mail-list: `https://groups.google.com/g/datasys-kth-rise`
Github Group: `https://github.com/cda-group`

9 SWOT Analysis

Our SWOT Analysis (mid-term) was as follows:

Strengths

- CDA itself defines a new and distinct problem domain for systems and ML alike with high significance and interesting challenges alike.
- Software design approach in CDA is decoupled to hardware and relevant advances which can be incorporated to the project fast.

Weaknesses

- Supporting $100\,\%$ existing functionality already present in the ecosystem of big data technologies in addition to CDA's novelties might add a strong development burden in the long-term.
- Research in "Systems for ML" is criticised by both ML and systems communities for its identity and that increases the challenge of publishing work in the right venue.

Opportunities

- New open source initiatives in compiler technology such as MLIR helps reduce the development overhead and boost usability of results in the project to a great extent.
- Increasing awareness of the limitations of "black box" models and transformers such as GPT-3 will fuel interest for reliable analytics which CDA aims to provide ahead of time.

Threats

- The continuous dependence on cloud-provided ML services could shift public focus from novel in-house data management such as CDA to general models under the centralised control of big powers in cloud computing such as Amazon, Google, OpenAI etc.
- It is challenging to keep talent and critical mass in Sweden given the potential opportunities and offers that are constantly given by big corporations in Silicon Valley and China.

9.1 COMMENTS

While our SWOT analysis was accurate, we also found it lacking since it did not incorporate sudden events that can occur and corresponding actions. One such event was the loss of Christian Schulte, one of the founding members of CDA, which left Christian's rare expertise in constrained programming inaccessible at a time that we needed it. However, this also created a new opportunity to look into AI methods for system optimization. Another unanticipated issue was the COVID-19 outbreak which stalled new hires at KTH and RISE, resulting in a multi-year delay of new valuable doctoral students.

A Appendix

A3. RESEARCHERS OF CDA

Researcher	Gender	University	Position	PhD Year
Seif Haridi	M	KTH/RISE	Professor	1981
Henrik Böstrom	M	KTH	Professor	1993
Paris Carbone	M	RISE/KTH	Senior Researcher/Asst Prof.	2018
Christian Schulte	M	KTH	Professor	2001
Sarunas Girdzijauskas	M	RISE	Senior Researcher	2009
Frej Drejhammar	M	RISE	Senior Researcher	2005 (Lic)
Salman Niazi	M	KTH	Doctoral Student	2018
Theodore Vasiloudis	M	RISE	Doctoral Student	2019
Jingna Zheng	F	KTH	Doctoral Student	2020
Lars Kroll	M	KTH/RISE	Doctoral Student/Researcher	2020
Max Meldrum	M	RISE/KTH	Engineer/Doctoral Student	2024*
Adam Hasselberg	M	KTH/RISE	Engineer/Doctoral Student	2026*
Jonas Spenger	M	RISE	Engineer/Doctoral Student	2026*
Harald Ng	M	RISE/KTH	Engineer/Doctoral Student	2026*
Sonia-Florina Horchidan	F	KTH	Doctoral Student	2026*
Klas Segeljakt	M	KTH	Doctoral Student	2023*
Mikolaj Robakowski	M	RISE	Research Engineer	
Sruthi Kumar	F	RISE	Research Engineer	
Haseeb Asif	M	RISE	Research Engineer	
Shahnur Isgandarli	M	RISE	Research Engineer	
Xin Long Han	M	RISE	Research Engineer	
Gabriele Morello	M	RISE	Research Engineer	
Abyel Tesfay	M	RISE	Research Engineer	
Alexandru Mindru	M	RISE	Research Engineer	
Jonathan Arns	M	RISE	Research Engineer	
Jun-Wei Liu	M	RISE	Research Engineer	
Kevin Harrison	M	RISE	Research Engineer	
Maximilian Georg Kurzawski	M	RISE	Research Engineer	

A4. PUBLICATIONS OF CDA

Doctoral Dissertations

Paris Carbone
Scalable and Reliable Data Stream Processing (2018) [Source]

Salman Niazi
Scaling Distributed Hierarchical File Systems Using NewSQL Databases (2018) [Source]

Theodore Vasiloudis
Scalable Machine Learning through Approximation and Distributed Computing (2019) [Source]

Lars Kroll
Compile-time Safety and Runtime Performance in Programming Frameworks for Distributed Systems (2020) [Source]

Jingna Zeng
Augmenting Transactional Memory with the Future Abstraction (2020) [Source]

Book Chapters / Technical Reports

Sherif Sakr, Tilmann Rabl, Martin Hirzel, Paris Carbone, Martin Strohbach
Report of Dagstuhl Seminar in Big Stream Processing
SIGMOD Record (2018)

P Carbone, A Katsifodimos, S Haridi
Stream Window Aggregation Semantics and Optimization
Encyclopedia of Big Data Technologies - Springer (2019)

S Haridi, L Kroll, P Carbone
Lecture Notes on Leader-based Sequence Paxos–An Understandable Sequence Consensus Algorithm
arXiv preprint arXiv:2008.13456 (2020)

M Fragkoulis, P Carbone, V Kalavri, A Katsifodimos
A Survey on the Evolution of Stream Processing Systems
arXiv preprint arXiv:2008.00842 (2020)

J.Spenger, P. Carbone, P.Haller *A Survey of Actor-Like Programming Models for Serverless Computing Active Object Languages 2023 (Review)*

Klas Segeljakt, Seif Haridi, Frej Drejhammar, Paris Carbone
Technical Report of the ArcLang language (2023)

Journals/Conferences

2017

Vasiloudis, Theodore, Foteini Beligianni, and Gianmarco De Francisci Morales *BoostVHT: Boosting Distributed Streaming Decision Trees Proceedings of the 2017 ACM Conference on Information and Knowledge Management* ACM, 2017

2018

Salman Niazi, Mikael Ronström, Seif Haridi, Jim Dowling *Size Matters: Improving the Performance of Small Files in Hadoop ACM Middleware 2018*

Z Abbas, V Kalavri, P Carbone, V Vlassov *Streaming graph partitioning: an experimental study (journal) PVLDB 2018*

Vasiliki Kalavri, Vladimir Vlassov, Seif Haridi *High-level programming abstractions for distributed graph processing IEEE Transactions in Knowledge Data Engineering (journal) TKDE 2018*

J Zeng, P Romano, J Barreto, L Rodrigues, S Haridi *Online Tuning of Parallelism Degree in Parallel Nesting Transactional Memory IEEE International Parallel and Distributed Processing Symposium (IPDPS 2018)*

2019

T Vasiloudis, H Cho, H Boström *Block-distributed Gradient Boosted Trees 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, 2019*

Theodore Vasiloudis, Hyunsu Cho (Amazon Web Services), Henrik Boström *Block-distributed Gradient Boosted Trees Workshop on Systems for ML at NeurIPS 2019*

T Vasiloudis, GDF Morales, H Boström *Quantifying Uncertainty in Online Regression Forests Journal of Machine Learning Research 20 (155), 2019*

M Meldrum, K Segeljakt, L Kroll, P Carbone, C Schulte, S Haridi *Arcon: Continuous and Deep Data Stream Analytics Proceedings of Real-Time Business Intelligence and Analytics, 1-3 @ VLDB 2019*

Lars Kroll, Klas Segeljakt, Paris Carbone, Christian Schulte, Seif Haridi *Arc: An IR for Batch and Stream Programming DBPL Workshop at ACM PLDI 2019*

M Ismail, M Ronström, S Haridi, J Dowling *ePipe: Near Real-Time Polyglot Persistence of HopsFS Metadata*
19th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing 2019

2020

P.Carbone, M. Fragkoulis, V.Kalavri, A.Katsifodimos *Beyond Analytics: the Evolution of Stream Processing Systems SIGMOD 2020*

H Bostrom, U Johansson *Mondrian Conformal Regressors COPA 2020*

H Bostrom, P Hoglund, S-O Junker, A-S Oberg, M Sparr *Explaining Multivariate Time Series Forecasts: an Application to Predicting the Swedish GDP XI-ML 2020*

L Karlsson, H Bostrom, P Zieger *Classification of Aerosol Particles using Inductive Conformal Prediction COPA 2020*

H Werner, L Carlsson, E Ahlberg, H Bostrom *Evaluating Different Approaches to Calibrating Conformal Predictive Systems COPA 2020*

Kefato, Girdzijauskas et al. *Context-sensitive node representations - Gossip and Attend: Context-Sensitive Graph Representation Learning ICWSM 2020*

Antaris, Girdzijauskas et al. *EGAD: Evolving Graph Representation Learning with Self-Attention and Knowledge Distillation for Live Video Streaming Events IEEE Big Data 2020*

Lee, Girdzijauskas et al. *Z-Embedding: A spectral representation of event intervals for efficient clustering and classification ECML/PKDD 2020*

2021

Zekarias Kefato, Sarunas Girdzijauskas, Nasrullah Sheikh, Alberto Montresor *Dynamic Embeddings for Interaction Prediction The Web Conference WWW 2021*

Z. T. Kefato and S. Girdzijauskas *Self-supervised Graph Neural Networks without explicit negative sampling The International Workshop on Self-Supervised Learning for the Web (SSL'21), at WWW'21, 2021*

Z. T. Kefato, S. Girdzijauskas and H. Stärk *Jointly Learnable Data Augmentations for Self-Supervised GNNs NeurIPS Workshop on Self-Supervised Learning: Theory and Practice, Virtual, 2021*

Zekarias T. Kefato, Sarunas Girdzijauskas *SelfGNN: Self-supervised Graph Neural Networks without explicit negative sampling International Workshop on Self-Supervised Learning for the Web (SSL'21) @ the Web Conference WWW 2021*

Henrik Boström, Ulf Johansson, Tuwe Löfström *Mondrian conformal predictive distributions Proceedings of the Tenth Symposium on Conformal and Probabilistic Prediction and Applications, PMLR 152:24-38 2021*

Ulf Johansson, Tuwe Löfström, Henrik Boström *Well-Calibrated and Sharp Interpretable Multi-Class Models Proceedings of the International Conference on Modeling Decisions for Artificial Intelligence (MDAI) 2021: 193-204*

Ulf Johansson, Tuwe Löfström, Henrik Boström *Calibrating multi-class models Proceedings of the Tenth Symposium on Conformal and Probabilistic Prediction and Applications, PMLR 2021 152:111-130* (Alexey Chervonenkis award for the best paper)

Ulf Johansson, Henrik Boström, Tuwe Löfström *Investigating Normalized Conformal Regressors IEEE Symposium Series on Computational Intelligence (SSCI) 2021: 1-8*

Hugo Werner, Lars Carlsson, Ernst Ahlberg, Henrik Boström *Evaluation of updating strategies for conformal predictive systems in the presence of extreme events Proceedings of the Tenth Symposium on Conformal and Probabilistic Prediction and Applications 2021, PMLR 152:229-242*

Spenger, Jonas, Paris Carbone, and Philipp Haller *Pods: Privacy Compliant Scalable Decentralized Data Services Heterogeneous Data Management, Polystores, and Analytics for Healthcare* Springer, Cham, 2021, 70-82.

Jingna Zeng, Shady Issa, Paolo Romano, Luis Rodrigues, Seif Haridi *Investigating the semantics of futures in transactional memory systems Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, 2021*

2022

M Perini, G Ramponi, P Carbone, V Kalavri *Learning on Streaming Graphs with Experience Replay Proceedings of the 37th Annual ACM Symposium on Applied Computing 2022*

Carbone, G Dan, J Gross, B Göransson, M Petrova *NeuroRAN: rethinking virtualization for AI-native radio access networks in 6G arXiv preprint arXiv:2104.08111*

S. Horchidan, E. Kritharakis, V. Kalavri, P. Carbone *Evaluating Model Serving Strategies over Streaming Data DEEM-SIGMOD 2022*

M. Zwolak, Z. Abbas, S. Horchidan, P. Carbone, V. Kalavri *GCNSplit: Bounding the State of Streaming Graph Partitioning aiDM-SIGMOD 2022*

J Spenger, P Carbone, P Haller *Portals: an Extension of Dataflow Streaming for Stateful Serverless ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software (Onward!@SPLASH'22)*

Filip Cornell, Jussi Karlgren, Sarunas Girdzijauskas *Symbolic Hyperdimensional Vectors with Sparse Graph Convolutional Neural Networks International Joint Conference on Neural Networks (IJCNN) 2022*

Yifei Jin, Marios Daoutis, Sarunas Girdzijauskas, Aristides Gionis *Open World Learning Graph Convolution for Latency Estimation in Routing Networks International Joint Conference on Neural Networks (IJCNN) 2022*

Susanna Pozzoli, Sarunas Girdzijauskas *Not Only Degree Matters: Diffusion-Driven Role Recognition 2nd International Workshop on Open Challenges in Online Social Networks, OASIS 2022*

Edvin Listo Zec, Ebba Ekblom, Martin Willbo, Olof Mogren, Sarunas Girdzijauskas *Decentralized adaptive clustering of deep nets is beneficial for client collaboration International Workshop on Trustworthy Federated Learning in Conjunction with IJCAI (FL-IJCAI), 2022*

Martin Isaksson, Edvin Listo Zec, Rickard Cöster, Daniel Gillblad, Sarunas Girdzijauskas *Adaptive Expert Models for Personalization in Federated Learning International Workshop on Trustworthy Federated Learning in Conjunction with IJCAI (FL-IJCAI), 2022*

Filip Cornell, Chenda Zhang, Jussi Karlgren, Sarunas Girdzijauskas *Challenging the Assumption of Structure-based embeddings in Few-and Zero-shot Knowledge Graph Completion Proceedings of the Thirteenth Language Resources and Evaluation Conference, 2022*

2023

Harald Ng, Seif Haridi, Paris Carbone *Omni-Paxos: Breaking the barriers of partial connectivity 18th European Conference on Computer Systems (EuroSys) 2023*

Harald Ng, Kun Wu, Paris Carbone ***Harald Ng, Kun Wu, Paris Carbone*** UniCache: Efficient Log Replication through Learning Workload Patterns 26th International Conference on Extending Database Technology (EDBT '23)

Ahmed E Samy, Lodovico Giaretta, Zekarias T Kefato, Sarunas Girdzijauskas SchemaWalk: Schema Aware Random Walks for Heterogeneous Graph Embedding Companion Proceedings of the ACM Web Conference, 2023

S. Horhidan, P. Carbone ORB: Empowering Graph Queries through Inference DMKG 2023

J. Spenger, P. Haller, P. Carbone Portals: A Showcase of Multi-Dataflow Stateful Serverless PVLDB 2023

A5. EVENTS

[Tutorial] *- Organisation of VLDB BOSS Workshop Tutorial on Stream Processing with Apache Flink - Paris Carbone and Seif Haridi (August 2017)*

[Tutorial] *- Organisation of the Data Stream Systems Section at the 1st Dagstuhl Seminar on Big Stream Processing - Paris Carbone (November 2017)*

[Seminar] *Paris Carbone presented Asynchronous Epoch Commits for Fast and Consistent Stateful Streaming with Apache Flink* LADIS 2018, part of ACM PODC *(July 2018)*

[Tech Talk] *- Stream Loops on Flink: Reinventing the wheel for the streaming era Talk by Paris Carbone - Flink Forward 2018 Berlin https://www.ververica.com/flink-forward-berlin/resources/stream-loops-on-flink-reinventing-the-wheel-for-the-streaming-era (August 2018)*

[Workshop] *- Organisation of the first "Distributed Computing and Analytics Workshop" held in Stockholm, with many distinguished guests Website: https://discan18.github.io/ (September 2018)*

[Guest Seminar] *Distinguished speaker in the center of unstoppable computing Seif Haridi - University of Chicago (October 2018)*

[Tech Talk] *- Graph processing and cardinality approximation on massive streaming data with Apache Flink Theodore Vasiloudis (June 2018)*

[Tutorial] *Beyond Analytics: the Evolution of Stream Processing Systems SIGMOD (June 2020)*

[Workshop] *on Programming Languages and Distributed Systems (PLDS) https://plds.github.io/index.html (Nov 2020)*

[Tech Talk] *Frej Drejhammar - Extending Clang and LLVM for Interpreter Profiling Perfection (Euro-LLVM 2020)*

[Tech Talk] *- Klas Segeljakt - Experiences using LLVM to implement a custom language (Euro-LLVM 2020)*

[Seminar] *- Klas Segeljakt - Arc: An MLIR dialect for Data Analytics (PLDS Mar. 2020)*

[Seminar] *- Paris Carbone - Seamless Batch and Stream Computation on Heterogeneous Hardware with Arcon (PLDS Mar. 2020)*

[Seminar] *- Paris Carbone - Reliable Stream Processing at Scale with Apache Flink (Chaos Engineering Workshop 2019)*

[Tech Talk] *- Paris Carbone - NDSML Summit - Continuous Intelligence - Intersecting Event-Based Business Logic and ML Chaos Engineering Workshop - (Oct. 2019)*

[Seminar] *- Paris Carbone - Continuous Intelligence through Computation Sharing (CASTOR Software Days Oct. 2019)*

[Tech Talk] *- Massimo Perini - Deep Stream Dynamic Graph Analytics with Grapharis (Flink Forward Berlin 2019)*

[Tech Talk] *Klas Segeljakt, Max Meldrum - Introducing Arc Flink Forward Global Event (Oct. 2020)*

[Tech Talk] *Sruthi Kumar & Hasseb Asif – FlinkNDB: Skyrocketing Stateful Capabilities of Apache Flink, Flink Forward Global Event (Oct. 2020)*

[Tech Talk] Sruthi Kumar & Hasseb Asif – NEXMark-Beam: Your best companion for testing and benchmarking new core stream processing libraries, Beam Summit (Aug. 2020)

[Tech Talk] *- Max Meldrum - Presenting the Arcon Rust Streaming System (FOSDEM January 2021)*

[Industrial Research Talk] *- Paris Carbone - Research Overview of Data Systems Group (CASTOR Days 2022)*

[Summer School] *Sarunas Girdzijauskas was the main organizer of a Summer School on "Linked Data Analytics and Big Data Management", 12 – 15 September, 2022, Stockholm, Sweden*

[Workshop] *Paris Carbone was the main organizer of an invitation-only Data Computing workshop in September 2022 that took place in Kista, RISE and had invited guests from Ericsson, TUBerlin, and SAAB.*

A6. INNOVATIONS, PROTOTYPES, AND SPIN-OFF COMPANIES

During the CDA project we had one spinoff company; Hopsworks has been a company that Seif Haridi and Salman Niazi co-founded. Furthermore, CDA led to one patent issued, as well as several innovative open source systems, some of which yielding great potential to evolve into commercial products depending on the market demand. Overall, the project led to the following software products:

Crepes *A Python package, for developing conformal regressors and predictive systems on top of any regression model led by Henrik Boström. The documentation is found at https://crepes.readthedocs.io, while source code is found at the GitHub repository: https://github.com/henrikbostrom/crepes*

Orb DB *(TRL-2) is a software library and methodology for a new generation of graph databases which aims to combine generative AI with traditional processing to infuse inference in reasoning under user-defined uncertainty bounds. Orb has initiated several collaborations, a patent application and is currently entering a pre-seed funding phase to boost its development.*

OmniPaxos *(TRL-6) is a software protocol and Rust library that simplifies replication and fault-tolerance of distributed systems, which are challenging tasks when building highly-available services. With OmniPaxos, developers no longer need to worry about the complexities of replication and consensus; we provide an API that makes replication as simple as working with an append-only log. The website and documentation can be found at: https://omnipaxos.com/, whereas the source code is available on Github: https://github.com/haraldng/omnipaxos*

ArcLang *(TRL-4) is a language and compiler infrastructure for continuous deep analytics. The programming model of ArcLang is versatile allowing both event-based logic and query-driven comprehensions. The documentation and website of ArcLang can be found here: https://cda-group.github.io/arc/ , wherease its source code can be found here: https://github.com/cda-group/arc*

Portals *(TRL-3) is a framework which provides all building blocks necessary to construct complex scalable decentralized services that span across cloud and edge infrastructure. Portals automates serverless execution, fault tolerance and privacy (GDPR) in one system, making the process of digitalization accessible by everyone. Website: https://www.portals-project.org/*

A7. INTELLECTUAL ASSETS AND PROPERTIES

Patent Application *As part of our exploitation activities in CDA we filed in spring 2023 our first patent application to protect market-critical technologies developed within CDA for our ORB-graph database. The patent was funded in its majority by the VFK-1 grant from Vinnova and KTH Innovation Office and the project is entering a pre-incubation phase soon.*

Other CDA software projects *that are currently considered for trademarking in the future are:* ***ArcLang, OmniPaxos, and Portals***

A8. COURSES DEVELOPED WITHIN THE PROJECT

[summer school program] *Sarunas Girdizjauskas was the main organizer of a Summer School on "Linked Data Analytics and Big Data Management", 12 – 15 September, 2022, Stockholm, Sweden*

[SIGMOD Tutorial Course] *The evolution of Stream Processing Systems Tutorial was organized within the premier conference of Data Management and Databases (SIGMOD) was done in collaboration with TUDelft and Boston University and has been a big success. [Source]*

[KTH IV1351 and ID2203 Courses] *The IV1351 first cycle course on "Data Storage Paradigms" was redesigned and taken over completely by Paris Carbone at KTH since 2020 based on a CDA project 3-year reading course on data systems. Furthermore, half of the content of ID2203 course on "Distributed Systems" was remade to fit our research interests in CDA, and involve MSc students to the development of the various software research projects in CDA.*

[CDA Reading Group Course] *A third-cycle reading study group based on related technologies to CDA. The course took place in the form of seminars presented by senior members or students held on Fridays from 2017 to 2023. The content was heavily based on the studies conducted within CDA with invited talks from our area representatives (e.g., Henrik Boström, Sarunas Girdzijauskas) and many other external speakers from the USA. The course contributed wich 7.5 to each graduate student that successfully completed all of the required seminars and literature-review lectures.*

[CDA Course on Data Systems] *A third-cycle course for graduate students who want to focus on data technologies was organised using our resources and focus of the CDA group. The course included studying the content of three state-of-the-art 2nd and 3rd cycle courses by Assoc. Prof. Andy Pavlo at Carnegie Mellon University on advanced databases and hardware accelerated data processing. The students also had to complete a research project in advanced data systems to get all 7.5 credits.*

[CDA Course on ML for Graphs] *A third-cycle course for graduate students who want to focus on graph and ML technologies was jointly organised with Sarunas G. and Paris C. using our resources and focus of the CDA group. The course included studying the content of the widely available 2nd cycle course on "ML for Graphs" by Prof. Jure Leskovec at Stanford University. The students also had to complete a research project in advanced graph ML systems to get all 7.5 credits.*

[MooC] Joint Course with TUDelft on Stream Processing (2020)

A9. PHD EXAMS

Researcher	Gender	Deparment	Current Employer	PhD Year
Paris Carbone	M	KTH SCS	KTH	2018
Salman Niazi	M	KTH SCS	Hopsworks	2018
Theodore Vasiloudis	M	RISE/KTH SCS	Amazon	2019
Jingna Zheng	F	KTH	Volvo Cars	2020
Lars Kroll	M	KTH SCS	Databricks	2020

A10. LIC EXAMS

No Lic exams were obtained during the duration of the project.

A11. FUTURE EXAMS

Researcher	Gender	University	Reason	Expected Exam Year
Max Meldrum	M	KTH	1 year left in 5y program	2024*
Adam Hasselberg	M	RISE/KTH	3 years left in 4y program	2026*
Jonas Spenger	M	RISE/KTH	3 years left in 4y program	2026*
Harald Ng	M	KTH	3 years left in 5y program	2026*
Sonia-Florina Horchidan	F	KTH	3 years left in 5y program	2026*
Klas Segeljakt	M	KTH	doctoral report writing in progress	2023*

A12. NO EXAMS

No SSF-funded doctoral students are expected to not complete their doctoral exam.

A13. AWARDS

[ACM SIGMOD Systems Award 2013] *Paris Carbone, Seif Haridi et al. The prestigious 2023 SIGMOD Systems Award goes to the Apache Flink team: "Apache Flink greatly expanded the use of stream data-processing." https://sigmod.org/2023-sigmod-systems-award/*

[best paper & best presentation awards] S Horchidan*, E Kritharakis, V Kalavri, P Carbone - Evaluating Model Serving Strategies over Streaming Data DEEM - SIGMOD 2022*

[best student paper award] *Linn Karlsson, Henrik Boström, Paul Zieger. 2020. Classication of aerosol particles using inductive conformal prediction. Proceedings of the Ninth Symposium on Conformal and Probabilistic Prediction and Applications, PMLR 128:257-268*

[best paper award] *Theodore Vasiloudis, Hyunsu Cho, Henrik Boström: Block-distributed Gradient Boosted Trees. SIGIR 2019: 1025-1028 award]*

[Alexey Chervonenkis award for the best paper] *Ulf Johansson, Tuwe Löfström, Henrik Boström. 2021b. Calibrating multi-class models. Proceedings of the Tenth Symposium on Conformal and Probabilistic Prediction and Applications, PMLR 152:111-130*

A14. GRAND IMPACT STATEMENT

Impact Category	***Total number of Acknowledgements***
Original articles in refereed scientific journals	*6*
Original contributions to refereed scientific conferences	*41*
Review articles, book chapters, etc	*6*
Other scientific contributions	*2*
Patents and other IPR (such as trade marks)	*1*
Demos, software, copyrights	*7*
Popular science contributions and communication activities	*12*
PhD graduates	*5 (+6 candidates)*
Lic graduates	*0*
Spin out companies	*1*
Licensing deals	*0*
Person exchanges	*0*
Industry collaborations	*4*
Public sector collaborations	*1*
Research Institute collaborations	*4*
Courses developed due to the project	*8*

B Questions for the Project leader(s)

B.1 IF THE PROJECT HAD BEEN SET UP TODAY, WHAT CHANGES WOULD YOU HAVE MADE TO IT GIVEN EVERYTHING THAT YOU NOW KNOW [APART FROM THE RESEARCH RESULTS, OF COURSE]?

One of the core changes we would employ in similar strategic projects would be a more open approach to the development of the system technologies. The CDA vision, in the beginning, had a very well-defined software stack composed of interconnected software projects. This, at one end, eased the process of exploration with a more narrow design and made the work of doctoral students more straightforward, however, while the project evolved we realized several alternative architectures to explore. In the mid of the project we spinned, after hiring another batch of PhDs we spined of several alternative systems that ended up being more successful than we anticipated. To that end, we would maintain this agile approach in developing cutting-edge technology due to the speed and fast-pacing change of needs that occur in the field.

B.2 WHAT – IF ANYTHING – WILL ULTIMATELY BE THE MAIN IMPACT OF THE PROJECT ON SOCIETY AND ACADEMY?

The CDA project's main impact on society and the academy will be its transformation of data analysis and management processes. Its products, including Orb DB, OmniPaxos, UniCache, ArcLang, Arcon, and Portals, provide novel solutions to challenges in various tech sectors. Orb DB's AI graph database technology will boost Sweden's tech resilience and stimulate digital economy innovation. OmniPaxos and UniCache offer enhanced stability for technological infrastructures and improve the efficiency of cloud data services, respectively. Finally, core CDA technologies such as ArcLang, Arcon, and Portals, which have already garnered interest from major tech companies, will likely influence service development practices in the tech industry, driving efficiency and reliability.

B.3 WHAT DO YOU EXPECT WILL HAPPEN [WHAT HAS HAPPENED...] TO THE ACTIVITIES WITHIN THE PROJECT AFTER THE FOUNDATION FUNDING HAS EXPIRED?

Given the project's strategic significance and its interest from major tech companies like SAAB and Ericsson, it's reasonable to expect that activities within the project will continue after Foundation funding has expired. The developed technologies, like Orb DB, OmniPaxos, UniCache, ArcLang, Arcon, and Portals, could attract additional investment or commercial partnerships. Additionally, these technologies could be integrated into existing systems or further developed to foster innovation and resilience in the tech industry, particularly within Sweden and the EU. These activities should maintain, if not accelerate, their momentum post-funding.

B.4 WHAT WERE THE PROBLEMS OF THE PROJECT?

The project faced several challenges that could not be anticipated in the original SWOT analysis. The sudden loss of founding member Christian Schulte deprived the project of unique expertise in constraint programming at a crucial time, although it did open avenues for exploring AI methods for system optimization. Moreover, the unforeseen COVID-19 pandemic stalled the hiring process at KTH and RISE, causing significant delays in onboarding new doctoral students, a resource that would have added value to the project. This delay resulted in a approximately one-year setback for the project.

B.5 WHAT WAS THE MOST FUN OF THE PROJECT?

The project brought about many enjoyable aspects, particularly through team building and collaborative development. Working together to create innovative solutions like Orb DB, OmniPaxos, UniCache, ArcLang, Arcon, and Portals was not just a challenging but also a rewarding and enjoyable process. The sense of camaraderie developed in the process of tackling complex problems brought a fun element to the project. Brainstorming sessions and group activities also provided an avenue for exchanging creative ideas and fostered a lively and supportive team culture. Events, whether focused on knowledge sharing or celebrating project milestones, also contributed to the fun atmosphere, boosting morale and encouraging a vibrant and enthusiastic team dynamic.

B.6 YOUR MAIN APPRECIATION/COMPLAINTS OF THE FOUNDATION?

Our main appreciations of the Foundation revolve around their proactive approach to supporting strategic projects like ours. This has been instrumental in promoting innovation in Sweden, and we strongly believe such initiatives should continue. We also commend the SSF's flexibility during the COVID-19 pandemic, which was disruptive to many facets of the project. The Foundation's willingness to extend the project timeline in response was a significant relief and a testament to their commitment to the projects they support. We had no concerns regarding the SSF foundation throughout the project, however, we would recommend making the exploitation strategy a little more relaxed/flexible in terms of the host restrictions. KTH, today, for example, has the means to support and facilitate innovation activities within the KTH innovation office which also encourages external funding for such activities via KTH Holding AB with less overheads and support for IP management as well.