



**Technical report for
Collaborative Research Center
SFB 876**

**Providing Information by Resource-
Constrained Data Analysis**

December 2020

Part of the work on this report has been supported by Deutsche Forschungsgemeinschaft (DFG) within the Collaborative Research Center SFB 876 "Providing Information by Resource-Constrained Analysis".

Speaker: Prof. Dr. Katharina Morik
Address: Technische Universität Dortmund
Fachbereich Informatik
Lehrstuhl für Künstliche Intelligenz, LS VIII
D-44221 Dortmund

Contents

1	Subproject A1	2
1.1	Sebastian Buschjäger	3
1.2	Lukas Pfahler	7
2	Subproject A2	12
2.1	Henning Funke	13
2.2	Amer Krivošija	17
2.3	Andreas Lang	21
2.4	Lars Lenssen	25
2.5	Lea Schönberger	29
3	Subproject A3	34
3.1	Andrea Bommert	35
3.2	Junjie Shi	39
4	Subproject A4	44
4.1	Stefan Böcker	45
4.2	Robert Falkenberg	49
4.3	Marcus Haferkamp	53
4.4	Pascal Jörke	57
4.5	Karthik Aswin Venkatapathy Ramachandran	61
5	Subproject A6	66
5.1	Nico Bertram	67
5.2	Jonas Ellert	71
5.3	Matthias Fey	75
6	Subproject B2	80
6.1	Konstantin Wüstefeld	81

7	Subproject B3	86
7.1	Felix Finkeldey	87
7.2	Amal Saadallah	91
8	Subproject B4	96
8.1	Fabian Eckermann	97
8.2	Karsten Heimann	101
8.3	Petros Polichronidis	105
8.4	Benjamin Sliwa	109
8.5	Tim Vranken	113
9	Subproject C1	118
9.1	Till Hartmann	119
9.2	Alicia Tüns	123
10	Subproject C3	128
10.1	Jean-Marco Alameddine	129
10.2	Dominik Baack	133
10.3	Noah Biederbeck	137
10.4	Mirko Bunse	141
10.5	Alicia Fattorini	145
10.6	Danny Heinrich	149
10.7	Mirco Hünnefeld	153
10.8	Leonora Kardum	157
10.9	Lena Linhoff	161
10.10	Lukas Nickel	165
10.11	Kevin Schmidt	169
10.12	Jan Soedingrekso	173
10.13	Johannes Werthebach	177

11 Subproject C4	182
11.1 Simon Omlor	183
11.2 Julian Riehl	187
12 Subproject C5	192
12.1 Maximilian Berens	193
12.2 Kevin Heinicke	197
12.3 Vukan Jevtić	201
12.4 Thomas Lindemann	205
12.5 Gerwin Meier	209
12.6 Antje Mödden	213
12.7 Holger Stevens	217



Subproject A1
Data Mining for Ubiquitous System Software

Katharina Morik Jian Jia Chen

Machine Learning on Embedded Systems

Sebastian Buschjäger
Artificial Intelligence Group, Chair 8
Technical University Dortmund
sebastian.buschjaeger@tu-dortmund.de

With increasing volumes of data and more sophisticated machine learning algorithms, the demand for fast and energy efficient computation systems is growing. To meet this demand, two approaches are possible: First, machine learning algorithms can be tailored specifically for the hardware at hand. Second, instead of changing the algorithm, we can change the hardware to suit the machine learning algorithms better. This report briefly discusses my last years' work which focused largely on the first approach and quickly outlines some ideas for future research.

1 Introduction

To make machine learning universally applicable, we need to bring its algorithms to small and embedded devices including both - the training and the application of models. From a computer architectural point of view, we may optimize these two aspects separately. In model application, we rapidly apply an already trained model for predictions and thus focus on the optimization of *inference*. In model training, however, we would like to *train* models on small devices directly, so that these devices dynamically adjust their prediction rules for new data.

2 Machine learning for Embedded Devices

In the previous year, I started to play around with Binarized Neural Networks and their application on small devices. This research led to several papers throughout the year. I started to implement BNNs in popular frameworks such as PyTorch and continued to use this code for my experiments. In my initial experiments, I found that BNNs are more difficult to train compared to their float-point counterparts, but often do not overfit as severely. Together with Lukas Pfahler and Jens Buss we prepared and analyzed the FACT data with BNNs. This work is published at the ECML PKDD 2020 conference [1]. In the corresponding paper we train a BNN model to distinguish gamma and hadron particles from raw telescope data. Our approach outperforms state of the art approaches utilizing hand-crafted features and Random Forest. Moreover, we validated our model using real-world data from the Crab Nebula and deployed the trained model on an FPGA for on-the-fly data analysis using the FastInference tool which was developed in 2019.

Closely related to this line of work, I also started to work together with Mikail Yayla from the embedded systems group at TU Dortmund as part of our SFB876-A1 collaboration. In this work, we studied the bit error tolerance of BNNs when executed on unreliable hardware. From previous research, we knew that BNNs are quite error-tolerant which means that they retain some performance when bits are randomly flipped e.g. due to a hardware error. To enhance the bit error tolerance further literature proposes to inject random errors during training. After implementing this type of training for reference, we investigated various hypotheses. We found, that the bit error tolerance is closely connected to the classification margin of the network which can be used to train more bit-error tolerant networks. This work has been accepted at the DATE 2021 [2].

3 Machine learning on Embedded Devices

Last year I did some experiments with the 'Rule of Three' in the context of submodular function maximization, which resulted in a paper. This paper has been met with mixed reviews. While some reviewers really liked the general idea, other reviews were completely against it. Thus, I was not able to publish this work by now. The main reason why reviewers are against it is, that my method does not offer a worst-case guarantee but a guarantee in high probability. This results in a slightly weaker guarantee, but uses fewer resources. This is a strong contrast to established literature, which exclusively offers worst-case bounds without looking much at the overall resource consumption. I, however, argue that this worst case is often a pathological case during the analysis and most practical applications do not fall in this category. I enhanced the paper [3] and code¹ accordingly and hope to publish it early next year.

¹<https://github.com/sbuschjaeger/SubmodularStreamingMaximization>

During my work on outlier detection with Isolation Forest (see below), I noticed that *any* decision tree with enough nodes can approximate *any* distribution. This insight led me back towards my first research topic: The training of decision trees on streaming data. I devised a simple algorithm that randomly samples multiple trees and learns which of these trees performs best. This procedure is faster than traditional approaches, because technically there is no DT-training involved, but just random sampling as well as model application. In initial experiments, my approach showed a competitive performance to traditional approaches, but so far the theory behind this algorithm is not fully explored.

4 Beyond Embedded Systems

At the end of 2019 and the beginning of 2020, I worked on a Generalized Bias Variance Decomposition for ensemble learning. My goal was to generalize the existing Bias-Variance Decomposition for the MSE error to other loss functions and use this to derive new algorithmic insights. While I had some initial success in deriving a new decomposition, its algorithmic insights were limited and did not yield any new algorithms. Thus, I stopped working on this project for most of the year. Interestingly, a very similar paper was published at the ECML PKDD 2020 and after discussing it with the author, I decided to come back to this topic. Our experimental results now look much better and the code² as well as the preliminary paper is available [4]. I hope to publish it early next year.

During the year I supervised the master thesis of Philipp Honysz on the topic of outlier detection. Due to his excellent practical results and some theoretical analysis from my side, we were able to publish this work [5] and presented it at the DSAA conference [6]. In this paper, we analyze Isolation-based outlier detection methods from a distributional point of view. We show, that *any* tree-based algorithm accurately approximates the underlying probability distribution given the trees are large enough. Furthermore, we show how existing methods approximate the mixture coefficients of a mixture distribution by using the average path length in the tree. Building on these insights, we devised a Generalized Isolation Forest which outperforms other tree-based algorithms in the majority of cases.

5 Future research

First, I plan to publish the papers mentioned above which are not yet published. As an extension of the ECML paper, Lukas, Jens and I are currently preparing more experiments which also include the origin reconstruction as well as energy estimation of. Finally, I plan to write my PhD thesis in 2021.

²<https://github.com/sbuschjaeger/gnc1/>

References

- [1] S. Buschjäger, L. Pfahler, J. Buss, K. Morik, and W. Rhode, "On-site gamma-hadron separation with deep learning on fpgas," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2020.
- [2] S. Buschjäger, J.-J. Chen, K.-H. Chen, M. Günzel, C. Hakert, K. Morik, R. Novkin, L. Pfahler, and M. Yayla, "Margin-maximization in binarized neural networks for optimizing bit error tolerance," in *Proceedings of DATE 2021*, 2021.
- [3] S. Buschjäger, P.-J. Honysz, and K. Morik, "Very fast streaming submodular function maximization," 2020. [Online]. Available: <https://arxiv.org/abs/2010.10059>
- [4] S. Buschjäger, L. Pfahler, and K. Morik, "Generalized negative correlation learning for deep ensembling," 2020. [Online]. Available: <https://arxiv.org/abs/2011.02952>
- [5] S. Buschjäger, P.-j. Honysz, and K. Morik, "Randomized outlier detection with trees," *International Journal of Data Science and Analytics*, 2020. [Online]. Available: <https://doi.org/10.1007/s41060-020-00238-w>
- [6] S. Buschjäger, P.-J. Honysz, and K. Morik, "Generalized isolation forest: Some theory and more applications – extended abstract," in *Proceedings 2020 IEEE 7th International Conference on Data Science and Advanced Analytics (DSAA 2020)*. IEEE, 2020.

Unsupervised Representation Learning

Lukas Pfahler

February 5, 2021

We review a number of applications and models for unsupervised representation learning that occurred within the context of SFB876.

1 Introduction

An abundance of labeled training data has fuelled the explosive rise in deep-learning methods in domains like computer vision with its infamous image-net dataset. However, in many scenarios we do not have the luxury of millions of labeled data points. On the contrary, in many situations we are presented with a large collection of data points without any annotations to learn from. Still these datasets have value, and unsupervised or self-supervised machine learning methods can yield useful models.

2 Applications and Learning Tasks

In any given application where we seek to apply machine learning models, it is critical to frame the problem as an appropriate learning task.

Take for instance the following application: We want to learn to identify similar mathematical expressions in a large collection of formulas. The collection is extracted from <http://arxiv.org> and contains more than 900,000 scientific preprints with a total of over 29 million mathematical expressions that appear in the LaTeX sources of the preprints. We only use formulas that appear in maths-environments and ignore all inline expressions. We can consider different approaches for learning representations that capture semantic similarity. Our initial approach follows the contextual similarity principle popularized by methods like word2vec [6]. We conjecture that two expressions are similar when they appear in a shared context, like in the same section or in the same paper [8, 10]. Hence the machine learning task becomes a similarity learning task where the similarity of two formulas from a shared context must

be higher than the similarity with a random third formula. Many different loss functions for these kinds of similarity learning problems have been proposed, in our work we find that the Histogram loss [12] performs better than the triplet margin loss [1]. In a second approach, we automatically identify equalities and inequalities in the training data using a heuristic that splits formulas at predefined operators ($=$, \leq , \geq , $<$, $>$) and filters the results such that the majority of the outputs looks like meaningful equalities, rather than e.g. definitions. We apply as contrastive learning task where the model has to identify the true partner in an equality or inequality from a set of distractor formulas. When training these models in a minibatch SGD-like algorithm, we can use all the other formulas in the minibatch as negative, contrastive examples. To evaluate the performance of models trained this way, we measure Precision@k for identifying the right partner in all equality-parts of the test data.

In another application, fair news article recommendation we are interested in an representation of news paper articles that does not allow the generation of highly biased or polarized news feeds. To this end we train a vectorial representation that is unpredictable of the news outlet that published the news article, while still capturing the content or topic of the article. To this end we design an adversarial learning tasks comprised of a generator that proposes a vectorial representation and a discriminator that tries to guess the right news outlet [7].

3 Models

The second important question for learning representation is the underlying model architecture. The architecture specifies the inductive biases for learning representations, e.g. for computer vision tasks, convolutional network architectures with their emphasis on local patterns seem to be the architecture of choice in the literature. We have applies convolutional neural networks for working on classification and regression tasks in Gamma ray astronomy, where the input data are images captured by a telescope's camera [2].

For working with mathematical expressions, we have chosen Graph Convolutaional Neural Networks, which allows us to represent the mathematical expressions as a tree structure. We use an XML layout format, MathML, to represent our data. This can be fed into any GCN, which we use as a feature extractor that outputs a vectorial representation [8, 9]. Additionally we have considered transformer models as made popular by BERT [4] for processing of mathematical expressions. Rather than working on a sequential LaTeX representation, we again favour the tree-based MathML representation. In order to feed it into a sequential BERT-like model, we use the in-order tree

traversal algorithm.

Transformer architectures are a popular research area giving us a flood of models to choose from. Their main design choice is the choice of key-value storage. The original transformer model uses dot-product self attention and computing this attention matrix quickly becomes a bottleneck with its quadratic time- and memory-complexity (with respect to the input length). Many alternatives have been proposed, including the Linear Transformer [5] that focus on speeding up the self-attention mechanism. We instead investigate the use of another form of fully-differentiable key-value storage: holographic reduced representations. Holographic reduced representations [11] were already proposed in gated recurrent networks [3], we test their performance in modern transformer architectures. This has become computationally feasible with the availability of fast implementations of Fourier transformations in modern deep learning frameworks.

4 Outlook

References

- [1] Vassileios Balntas, Edgar Riba, Daniel Ponsa, and Krystian Mikolajczyk. Learning local feature descriptors with triplets and shallow convolutional neural networks. In Richard C. Wilson, Edwin R. Hancock, and William A. P. Smith, editors, *Proceedings of the British Machine Vision Conference (BMVC)*, pages 119.1–119.11. BMVA Press, 2016.
- [2] Sebastian Buschjäger, Lukas Pfahler, Jens Buß, Wolfgang Rhode, and Katharina Morik. On-Site Gamma-Hadron Separation with Deep Learning on FPAGs. In *ECML PKDD, 2020*.
- [3] Ivo Danihelka, Greg Wayne, Benigno Uria, Nal Kalchbrenner, and Alex Graves. Associative Long Short-Term Memory. *arXiv*, 48, 2016.
- [4] Jacob Devlin, Ming-Wei Chang, Lee Kenton, and Kristina Toutanova. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of NAACL-HLT 2019*, pages 4171–4186. Association for Computational Linguistics, 2019.
- [5] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and Francois Fleuret. Transformers are RNNs : Fast Autoregressive Transformers with Linear Attention. 2020.
- [6] Tomas Mikolov, Kai Chen, Greg Corrado, Jeffrey Dean, Kai Chen, and Jeffrey Dean. Efficient Estimation of Word Representations in Vector Space. In *1st International Conference on Learning Representations*, 2013.

- [7] Lukas Pfahler and Katharina Morik. Fighting Filterbubbles with Adversarial Training. (July):2–5, 2020.
- [8] Lukas Pfahler and Katharina Morik. Semantic Search in Millions of Equations. In *KDD 2020*, 2020.
- [9] Lukas Pfahler and Jan Richter. Interpretable Nearest Neighbor Queries for Tree-Structured Data in Vector Databases of Graph-Neural Network Embeddings. In *Workshop Proceedings of the EDBT/ICDT 2020 Joint Conference*, pages 2–6.
- [10] Lukas Pfahler, Jonathan Schill, and Katharina Morik. The Search for Equations - Learning to Identify Similarities between Mathematical Expressions. In *Machine Learning and Knowledge Discovery in Databases. ECML PKDD 2019*, 2019.
- [11] Tony A. Plate. Holographic Reduced Representations. *IEEE Transactions on Neural Networks*, 6(3):623–641, 1995.
- [12] Evgeniya Ustinova and Victor Lempitsky. Learning Deep Embeddings with Histogram Loss. In *Advances In Neural Information Precessing Systems 2016*, 2016.



Subproject A2
Algorithmic aspect of learning methods in
embedded systems

Erich Schubert Jens Teubner

Hardware-accelerated Query Processing

Henning Funke

DBIS Group

Technische Universität Dortmund

henning.funke@tu-dortmund.de

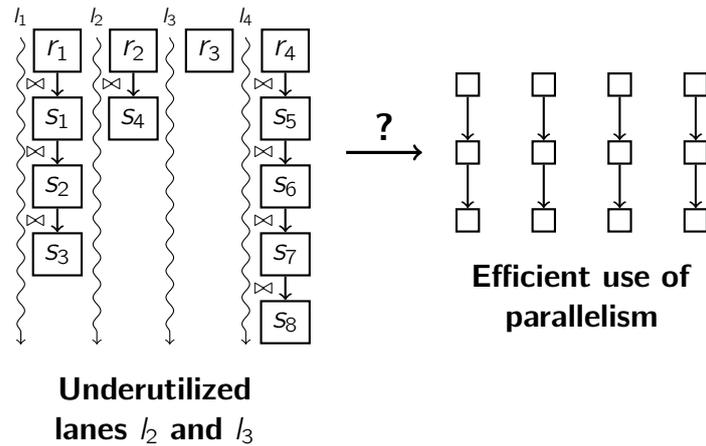
Query compilation is a processing technique for database queries that achieves very high resource-efficiency and throughput. It uses JIT (Just-in-Time)-compilation to generate machine code that is native to the target hardware for each query. This reduces overheads of traditional processing techniques, which stem from interpreting schemas and query plans *during* processing. Query compilation removes such interpretation work from the processing phase. At the time a query is compiled, the query plan and the schema are constants and can be evaluated before processing.

This article describes improvements to query compilation that were recently published. We first address parallelism in compiled queries. We show how GPUs with high parallelism degrees can be utilized effectively, even for skewed computations. Then we show improvements to the compilation speed of query compilers. We describe a new intermediate representation that enables very low compilation times.

1 Data-Parallel Query Compilation

GPUs have very high compute throughput and are a promising approach to overcome limitations of CPU-based processors. Being designed for data-parallelism, GPUs natively operate on several elements at a time. This allows GPUs to dedicate more chip resources to processing than CPUs.

Leveraging data-parallelism for query compilation in a beneficial way can be challenging. While uniform data can be processed naturally, irregular data and computation patterns may compromise the benefits. The following figure illustrates the problem for a database join operation. While rows r_1 and r_4 find three/four join partners, there is only a single



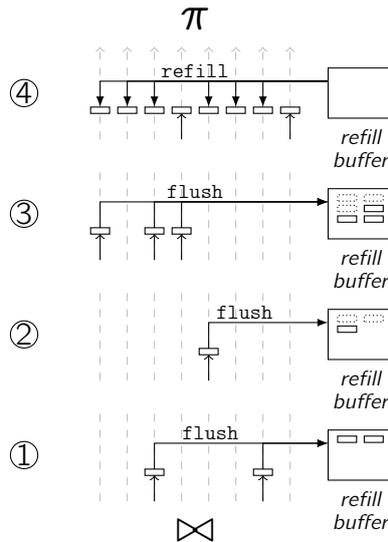
join partner for r_2 and none for r_3 . A naive data-parallel execution, therefore, will leave execution lanes l_2 and l_3 underutilized. This problem is called *divergence* because the work in some execution lanes differs from the work in others, i.e. it *diverges*.

Divergence is particularly harmful to query compilation techniques because query compilation links sequences of coherent operators into so-called *pipelines*. These pipelines can be executed together with improved efficiency. However, if some operators of the pipeline introduce divergence, the harmful effects are amplified by the execution of the successive operators. This is because the divergent processing state carries over from operator to operator.

We identify two different kinds of divergence and provide techniques to counter them. *Filter divergence* stems from processes that may eliminate some processing elements. For instance selections or semi-joins. *Expansion divergence* stems from processes that iterate through varying numbers of elements, e.g. joins or string operations. In our article [2], we analyze the negative performance effect introduced by divergence and present techniques that restore balanced parallelism when processing divergent workloads. We evaluate the benefit of divergence balancing with our prototype query compiler DogQC¹. The benefit of divergence balancing has also been shown as part of a demonstration [3]. In the following, we give an overview of one of the divergence balancing techniques: *Lane Refill*.

Lane refill can be used to restore parallelism in filter divergence situations. To this end, a lane refill operator is placed between two relational operators of a query pipeline. The lane refill operator has the ability to postpone processing of tuples by writing them to a buffer. Later the tuples from the buffer can be read to reactivate empty lanes. The following figure illustrates the application of lane refill in an example.

¹DogQC query compiler available at <https://github.com/Henning1/dogqc>

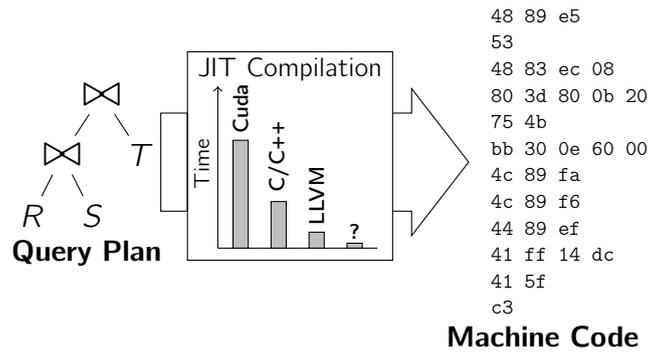


The figure contains two operators that are part of a JIT-compiled query pipeline. The join operator (bottom) precedes the projection operator (top) and introduces divergence. To restore processing efficiency, a lane refill operator is placed between both operators and performs the following operations in four consecutive iterations. In iterations 1, 2, and 3, only few lanes have a tuple. Most lanes are empty because their tuple was eliminated due to filter divergence. Lane refill suspends processing of the few remaining tuples by writing them to a buffer. In iteration 4 the buffer has sufficient elements to restore a high processing efficiency by loading tuples from the buffer into empty lanes. The elements are loaded and processing continues in iteration 4 without divergence. By placing such mechanisms between divergent operators, we can effectively restore balanced parallelism and improve processing efficiency and throughput.

2 Fast JIT-Compilation with Flounder IR

Query compilation makes query execution extremely efficient, but it introduces additional compilation time for translating the query before execution. When processing queries on smaller datasets, the relative cost of compilation time compared to the processing time increases. For small queries, a compiling query engine will spend most of the time during compilation before entering execution only for a very short time.

The intermediate representation (IR) that is used for compilation is an important design choice. During query compilation, queries are first translated to the IR and then the IR is translated to machine code. The following figure illustrates the effect of the IR choice on JIT compile times. Higher-level IRs, such as C++ or Cuda, have compilation times from hundreds of milliseconds to seconds. Lower level IRs, such as LLVM have compilation times of tens of milliseconds.



In our article [1], we propose *Flounder IR*, a lightweight intermediate representation for query compilation. Flounder IR is close to machine assembly and adds just that set of features that is necessary for efficient query compilation: virtual registers and function calls ease the construction of the compiler front-end; database-specific extensions enable efficient pipelining in query plans; more elaborate IR features are intentionally left out to maximize compilation speed.

By simplifying the IR and the translation process, it becomes feasible to achieve much shorter JIT compile times than those of LLVM. We show that Flounder IR accelerates query compilation by up to 24.6x over LLVM with 00-level optimizations and by up to 933x over LLVM with 03-level optimizations. The performance of the resulting machine code is similar to LLVM with 00-optimizations.

References

- [1] Henning Funke, Jan Mühlig, and Jens Teubner. Efficient generation of machine code for query compilers. *Damon@Sigmod*, 2020.
- [2] Henning Funke and Jens Teubner. Data-parallel query processing on non-uniform data. *PVLDB*, 13(6), 2020.
- [3] Henning Funke and Jens Teubner. Like water and oil: with a proper emulsifier, query compilation and data parallelism will mix well. *PVLDB*, 13(6), 2020.

On clustering and related problems on curves under the Fréchet distance

Amer Krivošija

Lehrstuhl für Datenbanken und Informationssysteme

Technische Universität Dortmund

amer.krivosija@tu-dortmund.de

An extended abstract of the dissertation submitted to the TU Dortmund.

Introduction

Sensor measurements are omnipresent. They can be represented as points in \mathbb{R}^d . Ordered by the time-stamps of these measurements, these points yield a time series in \mathbb{R}^d . By joining the points in this order by the straight lines, the time-series can be interpreted as a polygonal curve in the d -dimensional ambient space. Note that such representation ignores the explicit time stamps. This can be seen as a loss of information, but it concentrates on the shape of the curve. Such a choice is supported by the real-world examples of the tracking of airplanes, ships, bird migrations, but also in signature authentication and in the analysis of whale singing patterns. We studied several fundamental computational tasks on curves: clustering, simplification, and embedding.

The Fréchet distance is a popular distance measure for curves, in its continuous and discrete version. It is a distance measure of choice should the inner structure of the curves be observed, opposed to e.g. Hausdorff distance. Both variants of the Fréchet distance are (near)-metric. Related to the discrete Fréchet distance is the dynamic time warping (DTW) distance, that is very popular in the data mining community, but is not a metric. There exist algorithms to compute all three distance measures between two curves with m vertices each (called *complexity* of the curve), with (roughly) quadratic running time in m . However, it is widely believed that no algorithms exist to compute either the discrete or the continuous Fréchet, or the DTW distance between two curves of complexity m , in the subquadratic running time in m . The Fréchet distance lends itself

naturally to the computational tasks we investigate, in the corresponding metric spaces. Yet, the inherent cost of distance computation between the input curves puts a constraint on design of efficient algorithms, even in the one-dimensional ambient space.

Results in the one-dimensional ambient space

In our work we focused on curves in the one-dimensional ambient space \mathbb{R} (defined by e.g. stock market values, or weather phenomena). The main problem we studied is clustering curves in one-dimensional ambient space under the (continuous and discrete) Fréchet distance, in particular, the following variations of the well-known k -center and k -median problems. Given is a set P of n curves in one-dimensional ambient space, each of complexity at most m . Our goal is to find k one-dimensional curves, not necessarily from P , which we call *cluster centers* and each has complexity at most ℓ . In the (k, ℓ) -center problem, the maximum distance of an element of P to its nearest cluster center is minimized. In the (k, ℓ) -median problem, the sum of these distances is minimized. The need to introduce an additional parameter ℓ is justified, as by doing this we avoid of the overfitting of the cluster centers. We showed that both (k, ℓ) -clustering problems are NP-hard under both the discrete and the continuous Fréchet distance, if k is the part of the input. It was later shown by Buchin *et al.* [1] that both problems remain NP-hard if ℓ is the part of the input, and even if $k = 1$.

Under the continuous Fréchet distance, we gave $(1+\varepsilon)$ -approximation algorithms for both (k, ℓ) -center and (k, ℓ) -median problem, with running time near-linear in the input size for constant ε , k and ℓ . Our techniques yield constant-factor approximation algorithms for the observed problems under the discrete Fréchet distance. Our algorithms initiated a series of publications that address the problem in multi-dimensional ambient spaces \mathbb{R}^d , $d \geq 1$. We give a brief overview of the best results.

For the (k, ℓ) -center problem under the discrete Fréchet distance there is a $(1 + \varepsilon)$ -approximation by Buchin, Driemel and Struijs [2]. In the continuous case, a constant factor approximation algorithm by Buchin *et al.* [1] is known. For the (k, ℓ) -median problem under the discrete Fréchet distance, a $(1 + \varepsilon)$ -approximation was given by Nath and Taylor [6]. In the continuous case there was recently given a bicriteria $(1 + \varepsilon)$ -approximation by Buchin, Driemel and Rohde [3], but with the complexity $2\ell - 2$ of the cluster centers. A $(1 + \varepsilon)$ -approximation algorithm to (k, ℓ) -clustering for $d \geq 2$ for the continuous Fréchet distance is still an open problem. It is not much known on the related problems of (k, ℓ) -means and on the clustering under the DTW distance.

To obtain the $(1+\varepsilon)$ -approximation algorithms for the clustering problems under the continuous Fréchet distance, we develop a new simplification technique on one-dimensional curve. Our simplifications, called δ -*signatures*, provide the “shape” of the curve. The parameter δ relates to the minimum length of the edges of the simplified curve. The

main technical property of the δ -signatures is that for two given curves τ and π in \mathbb{R} at Fréchet distance at most δ , the vertices of π that are not close to the vertices of the δ -signature of τ can be omitted. Such obtained new curve π' has the continuous Fréchet distance to τ not larger than δ . This property implies that, during the search for the clustering centers, we can concentrate only to the area that is close to the vertices of the signatures of the input curves.

The signatures always exist, and we can compute them efficiently, provided that both the complexity and the distance to the original curve are given. Our signatures were recently used by Driemel and Psarros [4] to construct a $(2 + \varepsilon)$ -approximation to the approximate near neighbor problem for the one-dimensional polygonal curves. One limitation of the signatures is that they heavily utilize the properties of \mathbb{R} , and thus is not clear how to extend them into \mathbb{R}^d , for $d > 1$.

We also studied the problem of embedding of the (discrete and continuous) Fréchet distance (as well as the DTW) into one-dimensional ambient space. More precisely, we study distortion of the probabilistic embedding that results from projecting the curves onto a randomly chosen line. If the distortion of such embedding under the continuous Fréchet distance would be bounded, then it could be combined with our signatures.

We show that, in the worst case and under reasonable assumptions (that the curves are c -packed), the discrete Fréchet distance between two polygonal curves of complexity m in \mathbb{R}^d , where $d \in \{2, 3, 4, 5, 6, 7\}$, degrades by a factor linear in m with constant probability. The c -packedness of the curve τ intuitively says that the total length of subcurves of τ contained in any ball is bounded by a constant c times radius of the ball. We also showed the lower bounds on the distortion of the discrete and the continuous Fréchet distance, as well as for the DTW. There are no known results yet on upper bounds for the continuous Fréchet and the DTW distances.

Results on probabilistic points

Sometimes it is not the order of the measured values of interest but their frequency. Such an example is the position of the cellular phone user relative to the fixed antenna(s), where the needed strength of the antenna signal is to be minimized while providing a reliable service. Sensor measurements can also define a discrete probability distribution over possible locations of a point in \mathbb{R}^d . Then, the input consists of n probabilistic points. We study the probabilistic 1-center problem in Euclidean space \mathbb{R}^d , also known as the probabilistic smallest enclosing ball (pSEB) problem. Our main objective is to improve the best existing algorithm for the pSEB problem by Munteanu, Sohler and Feldman [5], by reducing its exponential dependence on the dimension d to linear.

To do so, we study the deterministic *set median* problem, which is a variant of the median problem for a collection of N point sets, each containing n deterministic points, in high dimensions. The set median problem generalizes the 1-median as well as the (probabilistic) 1-center problems. We present a $(1 + \epsilon)$ -approximation algorithm for the set median problem, using a novel combination of sampling techniques for clustering problems in metric spaces with the framework of stochastic subgradient descent. Our algorithm has running time $O((dn/\epsilon^4) \cdot \log^2 1/\epsilon)$. The running time of our algorithm does not depend on N and is only linear in d and n . However, we showed that the linear dependence on n cannot be removed, unless we lose the approximation factor of $\sqrt{2}$, or have the exponential dependence on d . This is not an option in high-dimensional case.

By incorporating our result for the set median problem into the framework of Munteanu, Sohler and Feldman [5] for the pSEB problem, we avoid the exponential dependence on d contained in their result. Our $(1 + \epsilon)$ -approximation algorithm for the pSEB problem takes $O((dn/\epsilon^4) \cdot \log^2 1/\epsilon)$ time. As a result, the pSEB algorithm becomes applicable to shape fitting problems in Hilbert spaces of unbounded dimension using kernel functions. We present an exemplary application by extending the support vector data description (SVDD) shape fitting method to the probabilistic case. This is done by simulating the pSEB algorithm implicitly in the feature space induced by the kernel function. Our probabilistic SVDD $(1 + \epsilon)$ -approximation algorithm has running time $O((dn/\epsilon^8) \cdot \log^2 1/\epsilon)$. It is an open question if our approach can be extended to other shape fitting methods, such as support vector machines (SVM).

References

- [1] K. Buchin, A. Driemel, J. Gudmundsson, M. Horton, I. Kostitsyna, M. Löffler, and M. Struijs. Approximating (k, ℓ) -center clustering for curves. In *SODA*, pages 2922–2938, 2019.
- [2] K. Buchin, A. Driemel, and M. Struijs. On the hardness of computing an average curve. In *SWAT*, volume 162, pages 19:1–19:19, 2020.
- [3] M. Buchin, A. Driemel, and D. Rohde. Approximating (k, ℓ) -median clustering for polygonal curves. In *SODA*, to appear, 2021.
- [4] A. Driemel and I. Psarros. $(2+\epsilon)$ -ANN for time series under the Fréchet distance. *CoRR*, abs/2008.09406, 2020.
- [5] A. Munteanu, C. Sohler, and D. Feldman. Smallest enclosing ball for probabilistic data. In *SoCG*, pages 214–223, 2014.
- [6] A. Nath and E. Taylor. k -Median Clustering Under Discrete Fréchet and Hausdorff Distances. In *SoCG*, volume 164, pages 58:1–58:15, 2020.

GMM with Numerically Stable BIRCH CF-Trees

Andreas Lang

Lehrstuhl für Informatik VIII

Technische Universität Dortmund

andreas.lang@tu-dortmund.de

BIRCH Clustering is a widely known approach for clustering, that has influenced much subsequent research and commercial products. The key contribution of BIRCH is the Clustering Feature tree (CF-Tree), which is a compressed representation of the input data. The idea has already been adopted to k -means, data-stream and density-based clustering. The Clustering Features that are used by BIRCH are simple summary statistics that can easily be upgraded with new data. Unfortunately, how the sum of squares is used in BIRCH is prone to catastrophic cancellation. We introduce a replacement cluster feature without this numerical problem, that could also be used in other work derived from BIRCH. With the improved stability, BIRCH can be used to accelerate Gaussian Mixture Modeling with the Expectation Maximization Algorithm.

1 Introduction

Cluster Analysis is an unsupervised learning task used in data mining and data analysis, it is also called clustering and is aimed at discovering potential structure in a data set when no labeled data or pattern examples are available. The BIRCH algorithm [4] is a widely known cluster analysis approach, that scales well to big data even with limited resources, because it processes the data as a stream and aggregates it into a compact summary of the data. It inspired various subsequent work, for example BICO [1] which combines BIRCH with the concept of coresets to implement approximate k -means clustering. Because many of the underlying problems (e.g., k -means clustering) are NP-hard, we often use approximation techniques and great concern is directed at the scalability.

Scalability is where the BIRCH algorithm shines. It is a multi-step procedure for numerical data that first aggregates the data into a tree based data structure much smaller than the original data. This condensed representation is then fed into a clustering method, that now only has to cluster the reduced data set. When examining the BIRCH algorithm we noticed that it is susceptible to a numerical problem known as “catastrophic cancellation”. This emerges when two large and similar floating point values are subtracted. With numerically stable Clustering Features it is possible to use the BIRCH CF-Tree to accelerate Gaussian Mixture Modeling by reducing the input data size.

2 Numerically Stable Clustering Features

The core concept of BIRCH are the Clustering Features $CF=(\vec{LS}, SS, N)$. Those are a summary data structure where each Clustering Feature represents N data points of dimension d , with the linear sum $\vec{LS} \in \mathbb{R}^d$ (with $LS_i = \sum_x x_i$) and the sum of squares $SS \in \mathbb{R}$. Clustering Features can be combined by using the additivity theorem of BIRCH $CF_{AB}=(\vec{LS}_A+\vec{LS}_B, SS_A+SS_B, N_A+N_B)$. Calculating the center of a Clustering Feature is done by computing LS/N . The variance of a Clustering Feature can also be calculated easily by using $Var(X)=E[X^2]-E[X]^2$, which can be computed, with data stored in a Clustering Features, as $Var(X)=\frac{1}{N}SS-(\frac{1}{N}\sum_i LS_i)^2$. This equation for variance computation is known to be numerically unstable because a numerical problem called “catastrophic cancellation”. It can return zero for non-constant values or even negative values because of rounding errors. To insert new data samples into a BIRCH tree it is necessary to calculate the distance from the new data sample to the existing Clustering Features. There are different distance metrics used in BIRCH (D0 - D4, R). Most of them are based on the variance:

$$D2(A, B) = \sqrt{\frac{1}{n_A n_B} \sum_{x \in A} \sum_{y \in B} \|x - y\|^2} \quad (1)$$

$$D3(A, B) = \sqrt{\frac{1}{n_{AB}(n_{AB}-1)} \sum_{x, y \in AB} \|x - y\|^2} \quad (2)$$

$$D4(A, B) = \sqrt{\sum_{x \in AB} \|x - \mu_{AB}\|^2 - \sum_{x \in A} \|x - \mu_A\|^2 - \sum_{x \in B} \|x - \mu_B\|^2} \quad (3)$$

$$R(A, B) = \sqrt{\frac{1}{n_{AB}} \sum_{x \in AB} \|x - \mu_{AB}\|^2} \quad (4)$$

For this reason we introduce the BETULA [2] cluster feature. This stores the summary of n data points as the mean μ and the sum of squared deviations $S := \sum_x n_x \|x - \mu\|^2$ in a cluster feature $CF=(\mu, S, n)$. By directly aggregating the sum of squared deviations instead of the sum of squares we can avoid these numerical problems as discussed by Schubert and Gertz in [3]. The variance calculation here is trivial: $Var(X)=\frac{S}{n}$. We can combine BETULA cluster features efficiently too. The number of data samples can be added $n_{AB}=n_A + n_B$ while the other two measures need slightly more effort. The

combined mean can be computed by $\mu_{AB} = \mu_A + \frac{n_A}{n_{AB}}(\mu_B - \mu_A)$ and the sum of squared deviations by $S_{AB} = S_A + S_B + n_B(\mu_A - \mu_B)(\mu_{AB} - \mu_B)$.

A BIRCH tree is built by sequential insertion of data points into the tree. The best fitting leaf is found by choosing the nearest node (Clustering Feature) on each level. If at the leaf level the new data point is within a set threshold distance of the nearest Clustering Feature, then the data point is added to this entry. If not, a new Clustering Feature with the data point is added to the leaf. In each case the path to the leaf entry is updated by adding the data to all Clustering Features belonging to it. Each node has a maximum capacity, when an insertion exceeds this, the node is split, which can propagate to higher levels. If the maximum capacity of the tree is reached the tree is rebuilt by increasing the threshold and reinserting all leaf entries in the same way a data point is inserted. With the BETULA cluster features the same tree can be built numerically stable without any performance disadvantages.

Clustering with BIRCH is a multi-phase approach where first the data is aggregated into the BIRCH tree and after that the global clustering takes place. For the global clustering phase a multitude of clustering algorithms can be used. We here introduce GMM with the EM algorithm on BIRCH clustering features.

3 GMM on CF-Features

Gaussian Mixture Modeling(GMM) with the Expectation Maximization(EM) algorithm is a popular but computationally expensive algorithm. It iteratively optimizes a GMM with two steps per iteration. First the probability density function (pdf) of each Gaussian with every data point is evaluated, then the distribution parameters are updated based on all points weighted by their probabilities. The algorithm is linear in the number of data points, but the number of iterations and the costly evaluation of the many pdfs makes it expensive. To scale this algorithm to large data sets and many clusters it is beneficial to use the data summarization of BIRCH / BETULA.

There are various variations of GMM, that differ in the parameters and shape of the clusters. Here we concentrate on some of the more popular variants. The spherical model, with varying weight per cluster and identical volume in each dimension and the diagonal model, that also allows for varying volume in each dimension. For BIRCH where we only have a scalar for SS , the spherical model is the only suitable model. When considering the new BETULA Clustering Features, we can also use the diagonal model.

When using GMM with BIRCH, the tree structure is discarded and only the Clustering Features in the leaves are kept. Then an initialization algorithm for example kmeans++ is used on them. Afterward the GMM with EM is executed. While normally each data

point is preprocessed separately, each Clustering Feature represents multiple data points, therefore the runtime of GMM with Clustering Features scales with their number and not with the number of data points. To improve the quality of the clustering we not only use the mean of the Clustering Features, but also the Gaussian distribution of the data in the CF to calculate the pdfs and update the clusters.

By utilizing BETULA cluster features and EM-GMM it is possible to cluster big data sets with limited memory and high numerical stability as shown in [2]

4 Outlook

Next, we want to make the BETULA cluster features more modular with the objective of adapting the stored data in the cluster features to the used GMM to optimize runtime and memory usage even further for a given model. Additionally, the cluster features will be expanded to be able to store the covariance instead of the variance which allows the usage of fully variable GMM that models covariance. To improve the performance we plan to explore the possibility of utilizing the tree structure to exclude or assign whole sub-trees to clusters to reduce the number of pdf computations necessary. After these improvements, we will focus on adding data stream clustering functionality to the algorithm.

References

- [1] Hendrik Fichtenberger, Marc Gillé, Melanie Schmidt, Chris Schwiegelshohn, and Christian Sohler. BICO: BIRCH meets coresets for k-means clustering. In *ESA*, 2013.
- [2] Andreas Lang and Erich Schubert. BETULA: Numerically Stable CF-Trees for BIRCH Clustering. In *Similarity Search and Applications*, Lecture Notes in Computer Science, 2020.
- [3] Erich Schubert and Michael Gertz. Numerically stable parallel computation of (co-)variance. In *SSDBM*, 2018.
- [4] Tian Zhang, Raghu Ramakrishnan, and Miron Livny. BIRCH: an efficient data clustering method for very large databases. In *SIGMOD*, 1996.

Fast Clustering under Application-specific Constraints

Lars Lenssen
Artificial Intelligence Group
Technical University Dortmund
lars.lenssen@tu-dortmund.de

Clustering of Euclidean and non-Euclidean data is a well researched but still relevant topic in machine learning. Many long-established and proven algorithms and implementations are expensive in run time costs. Depending on the application area, additional constraints arise that require specific adaptations to existing algorithms and increase sensitivity in terms of resource requirements.

In this report, approaches to reduce the runtime of k-medoids clustering for solving the facility location problem in generating distribution network models based on public data are presented. This use case in electrical engineering places high demands on the efficiency of the underlying clustering implementation due to large size of processed data. Constraints such as capacities and maximum distances results in additional requirements for the implementation.

1 Introduction

Clustering is basically the partitioning of a multidimensional unclassified data set. Similar data should be in the same cluster, while different data should remain in separate clusters. Two of the most common clustering methods are k-means and k-medoids. Both seek k centers of the surrounding data points, where the centers in k-medoids can only be set on existing data points. Thus, k-medoids is not limited to Euclidean data and can be applied to graph distances, which is relevant in the Distribution Grid Generation use case [3].

The objective function to be minimized for k-medoids

$$OF = \sum_{i \in Data} d(i, m(i)) \quad (1)$$

with $d(i, m(i))$ as the distance between data point i and the nearest medoid contains the sum of pairwise dissimilarities. The k-medoids problem is NP-hard to solve exactly. A common method to solving k-medoids is the Partitioning Around Medoids (PAM) [4]. PAM essentially goes through two phases: The BUILD phase and the SWAP phase. The BUILD phase consists of initializing the central points of the clusters (medoids) selected at the beginning. In the following SWAP phase, iterative medoid and non-medoid are swapped to minimize the objective function (1).

There are several additions to PAM that address the problem of high runtime PAMs on large datasets. Prominent examples are CLARA [4] with the extension CLARANS [5]. CLARA executes the PAM algorithm on a random subset of the initial data points and thus usually generates a significantly higher normalized loss than PAM. With CLARANS we can also save run time costs compared to PAM, but risking local minimas being computed as the result. Schubert and Rousseeuw [2] established FastPAM, another extension of original PAM, but also of CLARA and CLARANS. The adaptation essentially takes place in the SWAP phase. By an adjustment in the SWAP phase, it is possible to reduce the cost change computed $O(|D|^2)$ times in constant time with only one distance matrix lookup. Table 1 roughly outlines the results from FastPAM, where EagerPAM makes changes to the initialization in addition to the changes from FastPAM1. Finally, FasterPAM is a combination of both changes.

		Algorithm	Norm. Loss	Run Time	
100plants k = 100	Initialization		7.4%	2732 ms	100.0%
	PAM		0.5%	53813 ms	100.0%
	EagerPAM		0.7%	8910 ms	16.6%
	FastPAM1		0.5%	3645 ms	6.8%
	FasterPAM		0.5%	2992 ms	5.6%
optdigits k = 10	Initialization		6.6%	6726 ms	100.0%
	PAM		0.0%	29861 ms	100.0%
	EagerPAM		0.0%	17167 ms	57.5%
	FastPAM1		0.0%	11493 ms	38.5%
	FasterPAM		0.0%	9384 ms	31.4%

Table 1: Runtime and normalized loss on 100plants with $k = 100$ and optdigits with $k = 10$. The standard BUILD step from PAM was used for initialization [2].

2 Generating Distribution Network Models

Following to the Distribution Network Generation of Seack et al [3] the framework OS-MOGrid was developed, which generates energy grids in the form of graphs on the basis of public data. On the basis of this graph data structure, there are different computational tasks in which resource-efficient clustering methods are necessary. One of these tasks is the calculation of optimal power substation locations within the graph. This can theoretically be described as a facility location problem, which originally comes from urban, factory and public service planning. The objective function for facilities and demandpoints

$$FL = \sum_{i \in \text{Demands}} d(i, m(i)) + \sum_{j \in \text{Centers}} c(j) \quad (2)$$

with $c(j)$ as the cost of a facility has strong similarities to the objective function of k -medoids. In the concrete example, the transformers are the facilities and the consumers are the demandpoints. We need to focus on resource efficiency in clustering cause of the big data size and complexity of the data structures. Figure 1 shows results of the optimal power substation distribution for Hagen Eelsey-Nord [1]. Details from FastPAM were considered during implementation to ensure scalability to larger areas. The underlying OSM planet file contains about 1.2 TB of data.

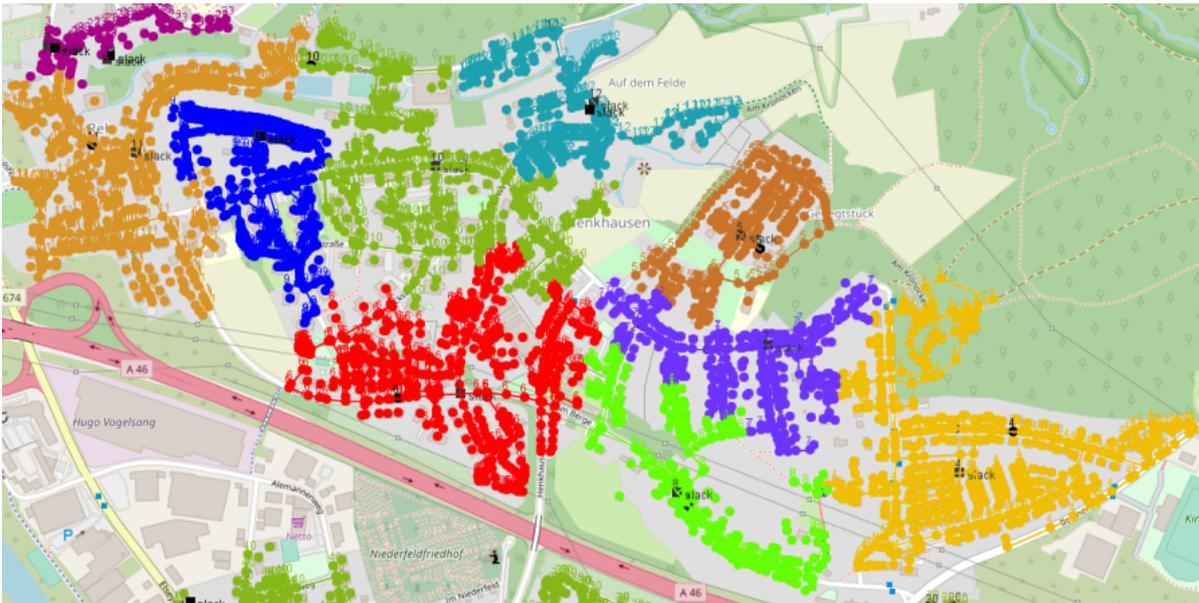


Figure 1: Clustering the demandpoints of the generated graph structure according to optimal power substation locations. [1]

3 Future Research

We plan to consider realistic constraints on the connection between consumers and substation for the concrete facility problem. Not all consumers can be serviced by all substation locations in reality, since facts like voltage drop over distance are physical limitations. In order to consider this sparse interface in the clustering algorithm, we need an adjustment of the implementation. This adaptation should focus on using the sparse data structure in such a way that runtime can be reduced. The given constraints lead to missing edges in the graph, which we represent in the distance matrix as infinite or very large constants. In case the number of edges is smaller than the product of the number of demandpoints and the number of facility locations, this leads to a reduction of the runtime. As a challenge, we still have to solve the problem that due to the missing edges in the graph, the task can be unsolvable for small k . Thus, the initialization is crucial. If necessary, we have to choose a larger k in the beginning, which we reduce dynamically at runtime.

References

- [1] Alexander Ostrop. Entwicklung eines Greedy-Algorithmus zur Generierung elektrischer Niederspannungsnetze auf Basis öffentlich verfügbarer Daten. *Bachelorarbeit*, TU Dortmund, 2019.
- [2] Erich Schubert and Peter J. Rousseeuw. Faster k-medoids Clustering: Improving the PAM, CLARA, and CLARANS Algorithms. *SISAP 2019. Lecture Notes in Computer Science*, 11807, 2019.
- [3] Jan Kays, André Seack, Theo Smirek, Fabian Westkamp, and Christian Rehtanz. The Generation of Distribution Grid Models on the Basis of Public Available Data. *IEEE Transactions on Power Systems*, 32:2346–2353, 2016.
- [4] Peter J. Rousseeuw and Leonard Kaufmann. *Finding Groups in Data: An Introduction to Cluster Analysis*. John Wiley Sons, 2009.
- [5] Raymond Ng and Jiawei Han. CLARANS: A method for clustering objects for spatial data mining. *IEEE Transactions on Knowledge and Data Engineering*, 14:1003–1016, 2002.

Towards End-to-End Latency Guarantees in Distributed Embedded Systems

Lea Schönberger

Lehrstuhl für Datenbanken und Informationssysteme

Technische Universität Dortmund

lea.schoenberger@tu-dortmund.de

Modern embedded systems are becoming more and more connected and can therefore be considered as distributed systems. Resulting from this, tasks of applications are executed in a distributed fashion and thus transform into time-critical cause-effect chains consuming a variety of heterogeneous communication as well as computation resources. Against this background, satisfying the timing requirements of an application in such distributed embedded systems, i.e., guaranteeing the so-called end-to-end latency, has become a major challenge. In this report, we give a short overview about our achievements and ongoing work related to this.

Many modern embedded systems are no longer designed as isolated stand-alone systems, but rather as highly dynamic distributed systems such as, e.g., robots, drones, and autonomous vehicles, being connected to external execution platforms via various communication solutions. As a consequence of this trend, applications are not executed on a local system only, but transform into time-critical cause-effect chains consuming a variety of heterogeneous communication as well as computation resources. Satisfying the timing requirements of an application in such distributed embedded systems, i.e., guaranteeing the so-called end-to-end latency, has become a major challenge [1]. In the following, we sketch our efforts made in order to address this challenge and point out our current research topics.

From a simplified local perspective, the sketched system can be reduced to a local execution platform and a communication channel, which serves as a connection to the

remaining resources, henceforth summarized under the term *remote system*. By the local system, the execution of a task of an application is perceived as the performance of an offloading operation, in the course of which the task interrupts its execution, transmits data to the remote system via the communication channel, and receives a response after a certain amount of time (denoted *suspension time*). Having received a response, the task resumes its local execution. Assuming an upper-bound on the suspension time, such systems can be analyzed with respect to their timing characteristics using any suitable method from the literature, e.g., our approach proposed in [4]. With respect to the suspension time, however, an upper-bound can only be assumed if the communication channel is reliable and the time required for processing the task's data on the remote system can be bounded as well.

The reliability of a communication channel strongly depends on the actual solution employed. Considering, for instance, 4G/5G or IEEE802.11p-based [3] wireless connections, the channel conditions can be highly variable due to various environmental and technological factors [6]. To soften the impact of such unreliability in terms of communication and to allow safe offloading, i.e., to guarantee end-to-end latencies, for tasks of safety- and mission-critical applications, we proposed two protocols in [5] applicable for systems executing so-called segmented self-suspending tasks with one suspension interval, i.e., tasks which (as perceived from a local system's perspective) are offloaded only once. For both protocols, we performed a formal timing analysis and provided a safe schedulability test. By means of comprehensive simulations and a case study, we verified the applicability of the protocols and demonstrated that the system behavior is strongly dependent on the quality of the communication channel, the percentage of critical tasks in the system, and the amount of offloaded workload (for detailed information please refer to [5]).

While assuming an upper-bound on a task's processing time on the remote system is inevitable for some approaches (including the above-mentioned), obtaining and guaranteeing this upper-bound is not trivial. In order to achieve this, it is necessary to switch from the local to a global view and to contemplate the distributed system as a whole. De facto, especially due to technologies such as cloud- and edge-computing, i.e., due to resources which can be used temporary and by to some extent arbitrary consumers or applications, distributed systems exhibit a high level of dynamicity. More precisely, these systems are not only dynamic in terms of applications, which may join and start their execution or terminate and leave, but also with respect to the resources pertaining to the system, which may join and leave as well, e.g., if edge-computing is considered and mobility is involved. So far, some effort has already been made by the scientific community to bound and guarantee end-to-end latencies in such dynamic distributed (embedded) systems (e.g., [2]).

However, a number of challenges remain, on which we focus in our currently ongoing research, such as:

- Developing and implementing dynamic resource reservation mechanisms that can cope with the dynamicity of distributed embedded systems, especially in view of highly heterogeneous resources;
- designing an admission control for applications;
- investigating the interplay between dynamic resource reservation and admission control in order to control end-to-end latencies of applications;
- guaranteeing end-to-end latencies for applications with variable resource demands;
- and providing a portable solution allowing to guarantee end-to-end latencies in already existing systems.

References

- [1] A. Hamann, S. Saidi, D. Ginthoer, C. Wietfeld, and D. Ziegenbein. Building end-to-end iot applications with qos guarantees. In *2020 57th ACM/IEEE Design Automation Conference (DAC)*, pages 1–6, 2020.
- [2] Victor Millnert, Johan Eker, and Enrico Bini. End-To-End Deadlines over Dynamic Topologies. In Sophie Quinton, editor, *31st Euromicro Conference on Real-Time Systems (ECRTS 2019)*, volume 133 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 10:1–10:22, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [3] Institute of Electrical and Electronic Engineers. 802.11p-2010 - IEEE standard for information technology– local and metropolitan area networks– specific requirements– part 11: Wireless LAN medium access control (MAC) and physical layer (PHY) specifications amendment 6: Wireless access in vehicular environments. https://standards.ieee.org/standard/802_11p-2010.html, 2017.
- [4] Lea Schönberger, Wen-Hung Huang, Georg von der Brüggen, Kuan-Hsun Chen, and Jian-Jia Chen. Schedulability analysis and priority assignment for segmented self-suspending tasks. In *The 24th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA)*, Hakodate, Japan, August 28-31 2018.
- [5] Lea Schönberger, Georg von der Brüggen, Kuan-Hsun Chen, Benjamin Sliwa, Hazem Youssef, Aswin Karthik Ramachandran Venkatapathy, Christian Wietfeld, Michael ten Hompel, and Jian-Jia Chen. Offloading Safety- and Mission-Critical Tasks via

Unreliable Connections. In Marcus Völz, editor, *32nd Euromicro Conference on Real-Time Systems (ECRTS 2020)*, volume 165 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 18:1–18:22, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik.

- [6] Benjamin Sliwa, Robert Falkenberg, Thomas Liebig, Nico Piatkowski, and Christian Wietfeld. Boosting vehicle-to-cloud communication by machine learning-enabled context prediction. *IEEE Transactions on Intelligent Transportation Systems*, Jul 2019.



Subproject A3
Methods for Efficient Resource Utilization in
Machine Learning Algorithms

Jian Jia Chen Jörg Rahnenführer

Integration of Feature Selection Stability in Model Fitting

Andrea Bommert
Faculty of Statistics
TU Dortmund University
andrea.bommert@tu-dortmund.de

This report summarizes the four aspects connected to feature selection analyzed in the PhD thesis “Integration of Feature Selection Stability in Model Fitting”. Firstly, a benchmark of filter methods for feature selection was conducted. Secondly, measures for the assessment of feature selection stability were compared both theoretically and empirically. Thirdly, a multi-criteria approach for obtaining desirable models with respect to predictive accuracy, feature selection stability, and sparsity was proposed and evaluated. Fourthly, an approach for finding desirable models for data sets with many similar features was suggested and evaluated.

For the benchmark of filter methods, 20 filter methods were analyzed. The analyses were based on the analyses in [4]. First, the filter methods were compared with respect to the order in which they rank the features and with respect to their scaling behavior in order to identify groups of similar filter methods. Especially for the data sets that contain a large number of features, there were four groups of filter methods that rank the features in a similar way and some filter methods that were not similar to any other method. For the other data sets, most filter methods were very similar. The filters that were similar to each other mostly came from the same toolboxes. Regarding the scaling behavior of the filter methods, there were two groups of filters with similar behaviors: the filter methods that calculate all scores at the same time and the methods that calculate the scores iteratively. Next, the predictive accuracy of the filter methods when combined with a predictive model and the run time needed for feature selection as well as for building a good predictive model based on the selected features were analyzed. There was no subset of filter methods that performed better than the rest of the filter methods on all data sets. Instead, the best filter methods differed between the data sets. Also,

the differences in performance between the filter methods on the same data set were not large for many data sets, compared to the observed variation. Regarding the aim of choosing a subset of filter methods that perform well for most data sets, the mutual information based filter methods *JMI* and *MIM* as well as the random forest importance filter *impurity* seem advisable. For *JMI*, the best predictive accuracies and for *MIM*, the lowest run times among all filter methods were observed across data sets. When searching for a good filter method for a new data set, it appears reasonable to limit the search space to *JMI*, *MIM*, and *impurity*.

To find suitable measures for stability assessment, 20 stability measures were compared based on both theoretical properties and on their empirical behavior. Five of the measures were newly proposed by us. Regarding the theoretical comparison, it was assessed which of the desirable criteria for stability measures defined in [5] are fulfilled by the stability measures. Additionally, new desirable properties were defined and investigated in this thesis. For the empirical comparison, two scenarios were considered: In the first scenario, all possible combinations of sets of selected features were analyzed for a very small number of features. In the second scenario, feature sets selected from real data sets were investigated. In both scenarios, groups of similar stability measures were identified and the impact of the number of selected features on the stability values was studied. Additionally, in the scenario with the real feature sets, the run times for calculating the stability measures were analyzed. These analyses are extensions of the analyses in [2] and [3]. When analyzing the stability assessment behavior of the stability measures, it was investigated whether the measures consider the same feature sets as stable or unstable. Based on this behavior, four homogeneous groups of stability measures could be identified. One group was formed by stability measures that do not possess the theoretical property "correction for chance". These measures take the higher values, indicating a more stable feature selection, the more features are selected. The second group consisted of stability measures that fulfill this property but that do not take into account similarities between the features. The third group was created by stability measures that both are corrected for chance and consider similarities between features. These measures were named "adjusted" (for similar features). The fourth group existed only for the analysis on real feature sets and only contained one uncorrected stability measure that - in contrast to the other uncorrected stability measures - is not able to take small values for small numbers of selected features. The stability measures that are corrected for chance do not have restrictions on attainable values. The run times for most stability measures were almost neglectable. For the adjusted and corrected measures however, the run times were quite long, especially for large data sets with many similar features. For stability assessment in future analyses, the stability measures SMU, SMA-Count, SMJ, and SMH were chosen as representatives of the four groups, based on their properties and on run time considerations. SMU and SMA-Count are suitable measures for stability assessment. When employing the measures SMJ and SMH, it should be kept in mind that they do not fulfill the important theoretical property "correction for chance". All stability measures considered in this thesis have been implemented in the R package *stabm* [1].

When searching for a good predictive model for a given data set, the predictive accuracy is usually the only criterion considered in the model finding process. In this thesis, the benefits of not only considering the predictive accuracy but also the feature selection stability and the number of selected features were investigated. To find desirable configurations with respect to all three performance criteria, the hyperparameter tuning of combined methods consisting of a filter and a classification method was performed in a multi-criteria fashion: the subset of Pareto optimal configurations, whose predictive accuracy is at most 0.05 below the accuracy of the configuration with highest accuracy obtained for the same data set, was determined. This way, it could be investigated whether it is possible to find configurations that perform a more stable selection of fewer features without losing much predictive accuracy compared to model fitting only considering the predictive performance. These analyses are extensions of the analyses in [3]. It could be concluded that with multi-criteria tuning, it is possible to choose such configurations. The resulting Pareto fronts gave several options to choose sparse and stable configurations, which at the same time had high predictive accuracy. The losses in predictive accuracy that went along with the gains in stability and sparsity were different between the data sets and ranged from noticeable over neglectable to gains in predictive accuracy. Also, with multi-criteria tuning, models were obtained that over-fit the training data less than the models obtained with single-criteria tuning only with respect to predictive accuracy. The configurations that were obtained with multi-criteria tuning differed a bit when different stability measures were used. Nevertheless, the conclusions about the benefits of multi-criteria tuning applied to all of the considered stability measures. All in all, the multi-criteria approach of considering predictive accuracy, stability, and sparsity jointly during hyperparameter tuning proved to be advantageous.

For data sets with many similar, for example highly correlated features, feature selection is especially challenging. Most established feature selection methods are not able to select only one feature out of a group of similar features. They select either several or none of them. L_0 -regularized logistic regression is a method that is able to perform such a feature selection. Therefore, for data sets with similar features, we proposed the approach of using L_0 -regularized logistic regression and tuning its hyperparameter in a multi-criteria fashion with respect to both predictive accuracy and feature selection stability. The idea was that considering also the feature selection stability avoids the selection of irrelevant or redundant features. We suggested to assess the stability with an adjusted stability measure, that is, a stability measure that takes into account similarities between features. For selecting the best configuration, the proposed algorithm ϵ -constraint selection was employed. The approach was evaluated based on both simulated and real data sets. It was compared to the standard approach of L_0 -regularized logistic regression with hyperparameter tuning only with respect to predictive accuracy, to L_0 -regularized logistic regression with hyperparameter tuning with respect to predictive accuracy and stability assessed with an unadjusted stability measure, and to the state-of-the-art method stability selection with L_0 -regularized logistic regression as feature selection method. Based on simulated data, it was observed that the proposed approach achieved the same or

better predictive performance compared to the two established approaches. Considering the stability during tuning did not decrease the predictive accuracy of the resulting models. The proposed approach succeeded at selecting the relevant features while avoiding irrelevant or redundant features. With the proposed approach, much fewer irrelevant or redundant features were selected than with the other approaches. Especially in situations with many similar features, the proposed approach outperformed the competing approaches. The single-criteria approach failed at avoiding irrelevant or redundant features and the stability selection approach failed at selecting enough relevant features for achieving acceptable predictive accuracy. On real data, performing hyperparameter tuning with respect to both predictive accuracy and stability was beneficial for fitting models with fewer features without losing predictive accuracy compared to single-criteria tuning. For data sets with many similar features, it was necessary to use an adjusted stability measure for obtaining this benefit. For data sets with very few similar features, it was still beneficial to consider the stability during tuning, but for these data sets, an unadjusted stability measure was sufficient and faster to compute. By additionally considering the stability, no predictive accuracy was lost and for many data sets, the resulting models were more sparse. The stability selection approach was outperformed by the proposed approach on many of the data sets. With the stability selection approach, comparably sparse models were fitted. These models, however, often did not include enough relevant features and therefore obtained a comparably low predictive accuracy. Also, the proposed approach could easily be applied to all of the considered data sets while the stability selection approach would have been infeasible without a high performance compute cluster for the larger data sets because of its enormous run time.

References

- [1] Andrea Bommert and Michel Lang. *stabm: Stability Measures for Feature Selection*, 2020. R package version 1.1.3.
- [2] Andrea Bommert and Jörg Rahnenführer. Adjusted measures for feature selection stability for data sets with similar features. In *International Conference on Machine Learning, Optimization, and Data Science*, 2020. preprint: arXiv:2009.12075.
- [3] Andrea Bommert, Jörg Rahnenführer, and Michel Lang. A multicriteria approach to find predictive and sparse models with stable feature selection for high-dimensional data. *Computational and Mathematical Methods in Medicine*, 2017, 2017.
- [4] Andrea Bommert, Xudong Sun, Bernd Bischl, Jörg Rahnenführer, and Michel Lang. Benchmark for filter methods for feature selection in high-dimensional classification data. *Computational Statistics & Data Analysis*, 143, 2020.
- [5] Sarah Nogueira. *Quantifying the Stability of Feature Selection*. PhD thesis, University of Manchester, United Kingdom, 2018.

Scheduling of Real-Time Tasks with Multiple Critical Sections in Multiprocessor Systems

Junjie Shi

Lehrstuhl für Informatik XII

Technische Universität Dortmund

junjie.shi@tu-dortmund.de

The performance of multiprocessor synchronization and locking protocols is a key factor to utilize the computation power of multiprocessor systems under real-time constraints. While multiple protocols have been developed in the past decades, their performance highly depends on the task partition and prioritization. The recently proposed Dependency Graph Approach showed its advantages and attracted a lot of interest. It is, however, restricted to task sets where each task has at most one critical section. In our recent work, we remove this restriction and demonstrate how to utilize algorithms for the classical job shop scheduling problem to construct a dependency graph for tasks with multiple critical sections. To show the applicability, we provide extensive numerical evaluations under different configurations, which in many situations show significant improvement compared to the state-of-the-art.

For several decades, the primary focus when considering multiprocessor synchronization and locking in real-time systems has been the design and analysis of resource sharing protocols, where the protocols decide the order in which the new incoming requests access the shared resources dynamically. Contrarily, the Dependency Graph Approaches (DGA), that was proposed by Chen et al. [1] in 2018, pre-computes the order in which tasks are allowed to access resources. Chen et al. [1] showed significant improvement against existing protocol-based approaches from the empirical as well as from the theoretical perspective, and demonstrated the practical applicability of the DGA by implementation. However, the original dependency graph approaches presented in [1] has two strong limitations: 1) the construction in the first step allows only one critical section per

task, and 2) the presented algorithms can only be applied for frame-based real-time task systems, i.e., all tasks have the same period and release their jobs always at the same time. The latter has been recently removed by Shi et al. [4], who applied the DGA after unrolling the jobs in the hyper-period. However, the former remains open and is a fundamental obstacle which limits the generality of the dependency graph approaches.

In the original DGA, the assumption that each task has only one non-nested critical section allows the algorithm to partition the tasks according to their shared resources in the first step. However, when a task accesses multiple shared resources, such a partitioning is no longer possible. Therefore, to enable the DGA for tasks with multiple critical sections, an exploration of effective construction mechanisms for a dependency graph that considers the interactions of the shared resources is needed.

We consider a set \mathbf{T} of n recurrent tasks to be scheduled on M identical (homogeneous) processors. All tasks can have multiple (non-nested) critical sections and may access several of the Z shared resources. Each task τ_i is described by $\tau_i = ((\eta_i, C_i), T_i, D_i)$:

- η_i is the number of computation segments in task τ_i .
- C_i is the total worst-case execution time (WCET) of task τ_i .
- T_i is the period of τ_i .
- D_i is the relative deadline of τ_i .

We consider constrained deadlines, i.e., $\forall \tau_i \in \mathbf{T}, D_i \leq T_i$. For the j -th segment of task τ_i , denoted as $\theta_{i,j} = (C_{i,j}, \lambda_{i,j})$:

- $C_{i,j} \geq 0$ is the WCET of computation segment $\theta_{i,j}$ with $C_i = \sum_{j=1}^{\eta_i} C_{i,j}$.
- $\lambda_{i,j}$ indicates whether the corresponding segment is a non-critical section or a critical section. If $\theta_{i,j}$ is a critical section, $\lambda_{i,j}$ is 1; otherwise, $\lambda_{i,j}$ is 0.
- If $\theta_{i,j}$ is a non-critical section, then $\theta_{i,j-1}$ and $\theta_{i,j+1}$ must be critical sections (if they exist). That is, $\theta_{i,j}$ and $\theta_{i,j+1}$ cannot be both non-critical sections.
- If $\theta_{i,j}$ is a critical section, it starts from the lock of a mutex lock (or *wait* for a binary semaphore), denoted by $\sigma_{i,j}$, and ends at the unlock of the same mutex lock (or *signal* to the same binary semaphore).

We detail the DGA for tasks with multiple critical sections, based on job shop scheduling to construct a dependency graph.

- In the first step, we construct a directed *acyclic* graph $G = (V, E)$. For each sub-job $\theta_{i,j}$ of task τ_i in \mathbf{T} , we create a vertex in V . The sub-job $\theta_{i,j}$ is a predecessor of $\theta_{i,j+1}$ for $j = 1, 2, \dots, \eta_i - 1$. Suppose that Θ^z is the set of the computation segments that are critical sections guarded by mutex lock z , i.e., $\Theta^z \leftarrow \theta_{i,j} \lambda_{i,j} = 1$ and $\sigma_{i,j} = z$. For each $z = 1, 2, \dots, Z$, the subgraph of the computation segments in Θ^z is a directed chain, which represents the total execution order of these computation segments.
- In the second step, we construct a schedule of G on M processors either globally or partitioned, either preemptive or non-preemptive.

We now explain how to reduce from an input instance I^{MS} of the makespan problem to an input instance I^{JS} of the job shop scheduling problem $J_{Z+n}||C_{\max}$.

- We create $Z + n$ shops:
 - Shop $z \in \{1, 2, \dots, Z\}$ is exclusively used to execute critical sections guarded by mutex lock z . That is, only critical sections $\theta_{i,j}$ with $\lambda_{i,j} = 1$ and $\sigma_{i,j} = z$ (i.e., $\theta_{i,j} \in \Theta^z$) can be executed on shop z .
 - Shop $Z + i$ is exclusively used to execute non-critical sections of task τ_i . That is, only non-critical sections $\theta_{i,j}$ with $\lambda_{i,j} = 0$ can be executed on shop $Z + i$.
- The operation of each computation segment $\theta_{i,j}$ is transformed to the corresponding shop, and the processing time is the same as the segment's execution time.

Suppose that ρ^{JS} is a feasible job shop schedule for I^{JS} . Since ρ^{JS} is non-preemptive, the operations on a shop are executed sequentially in ρ^{JS} . The construction of the dependency graph G sets the precedence constraints of Θ^z by following the total order of the execution of the operations on shop z , i.e., the shop dedicated for Θ^z in ρ^{JS} .

Once the dependency graph G is constructed, a schedule ρ^{MS} of the original input instance I^{MS} can be generated by applying any scheduling algorithms to schedule G , as already detailed in [1, 4]. Additionally, its partitioned extension in [5] (P-EDF) can be applied to generate the partitioned schedule.

We conducted evaluations on $M = 4, 8,$ and 16 processors. Based on M , we generated 100 synthetic task sets with $10M$ tasks each, using the RandomFixedSum method [2]. We set the utilization level, i.e., $\sum_{\tau_i \in T} U_i$ from 0 to $100\% \times M$ in steps of 5% . and enforced $U_i \leq 0.5$ for each task τ_i , where $U_i = \frac{C_i}{T_i}$ is the utilization of a task. The number of shared resources (binary semaphores) Z was either $4, 8,$ or 16 . Each task τ_i accesses the available shared resource randomly between 2 and 5 times, i.e., $\sum \lambda_{i,j} \in [2, 5]$. The total length of the critical sections $\sum_{\lambda_{i,j}=1} C_{i,j}$ is a fraction of the total execution time C_i of task τ_i , depended on $H \in \{[5\% - 10\%], [10\% - 40\%], [40\% - 50\%]\}$. When considering shared resources in real-time systems, the utilization of critical sections for each task in classical settings is relatively low. However, with the increasing computation demand in real-time systems (e.g., for machine learning algorithms), adopted accelerators, like GPUs, behave like classical shared resources (i.e., they are non-preemptive and mutually exclusive), but have a relatively high utilization. Hence, we chose possible settings of H that cover the complete spectrum. The total length of critical sections and non-critical sections are split into dedicated segments by applying UUniFast [2] separately. For task τ_i , the number of critical sections Num_{cs} equals to $\sum \lambda_{i,j}$, and the number of non-critical sections $Num_{ncs} = Num_{cs} + 1$. In the end, the generated non-critical sections and critical sections are combined in pairs, and the last segment is the last non-critical section. We evaluated all resulting 27 combinations of $M, Z,$ and H .

The evaluation results show that our approach outperforms the other non-DGA based methods significantly for all evaluated frame-based task settings settings, and performs

slightly better than the methods using *PRP*.

We applied constraint programming to solve the job shop problem $J_Z|r_j, l_j|L_{\max}$ and construct the dependency graph. We extended the settings for frame-based task sets to periodic task systems by choosing the period T_i randomly from a set of semi-harmonic periods, i.e., $T_i \in \{1, 2, 5, 10\}$, which is a subset of the periods used in automotive systems [3]. We used a small range of periods to generate reasonable task sets with high utilization of the critical sections, which are otherwise by default not schedulable. The evaluation results show that our approach is very effective for frame-based real-time task systems. The evaluation results show that our approach has significant improvements, compared to existing protocols, in most evaluated cases except light shared resource utilization. This paper significantly improves the applicability of the DGA by allowing arbitrary configurations of the number of non-nested critical sections per task.

In this paper, we focus on the long-standing problem of resource sharing of periodic tasks and on providing a good solution for this most adopted real-time task model. As a result, we achieved a solution that outperforms the methods in the literature which can be applicable to this task model. In the future, we plan to explore the possibility to apply the dependency graph approach on sporadic task systems, which do not have predefined arrival times of jobs.

References

- [1] Jian-Jia Chen, Georg von der Bruggen, Junjie Shi, and Niklas Ueter. Dependency graph approach for multiprocessor real-time synchronization. In *IEEE Real-Time Systems Symposium, RTSS*, pages 434–446, 2018.
- [2] Paul Emberson, Roger Stafford, and Robert I Davis. Techniques for the synthesis of multiprocessor tasksets. In *WATERS*, pages 6–11, 2010.
- [3] Simon Kramer, Dirk Ziegenbein, and Arne Hamann. Real world automotive benchmark for free. In *6th International Workshop on Analysis Tools and Methodologies for Embedded and Real-time Systems (WATERS)*, 2015.
- [4] Junjie Shi, Niklas Ueter, Georg von der Brüggen, and Jian-jia Chen. Multiprocessor synchronization of periodic real-time tasks using dependency graphs. In *2019 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 279–292, 2019.
- [5] Junjie Shi, Niklas Ueter, Georg von der Brüggen, and Jian-Jia Chen. Partitioned scheduling for dependency graphs in multiprocessor real-time systems. In *Proceedings of the 25th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications, RTCSA*, 2019.



Subproject A4
Resource efficient and distributed platforms
for integrative data analysis

Michael ten Hompel

Christian Wietfeld

Data-driven measures to improve reliability of Mission-Critical Machine Type Communications

Stefan Böcker

Lehrstuhl für Kommunikationsnetze
Technische Universität Dortmund
stefan.boecker@tu-dortmund.de

The ongoing digitalization and the associated steadily increasing number of distributed sensor devices and Internet-of-Things (IoT) systems implies a massive increase of subscribers. At the same time, the amount of available frequency spectrum resources remains static. While current 5G networks are already aiming for large-scale connectivity with an ambitious node density of 1,000,000 devices per square kilometer mainly in the area of massive 5G machine-type communications (mMTC), the ongoing evolution of MTC will primarily focus on mission-critical MTC (cMTC). In this regard, this work highlights diverse challenges and future aspects of mission-critical MTC on 6G enabling technologies. First, the need for new service classes and its characterization is discussed, highlighting the need for future 6G systems to leverage application-domain information about the predictability of resource requirements and conditions. Subsequently, by the dynamic radio resource management, as well as model-predictive communication approaches, two building blocks for data-driven Resource-Efficient and reliable cMTC networks are presented.

1 Data-Driven Building Blocks for Resource-Efficient MC-MTC Networks

While 5G has already introduced mMTC for many IoT applications such as Smart City and Smart Home applications, cMTC will be the primary focus of MTC in 6G [1]. Future applications tend to require dependable service quality characteristics in terms of latency

and error rates, for example, in the context of life-critical alarming and control, practically equivalent to wired communications. To address such increased requirements, future MTC devices will need to anticipate their environment. Among others, this includes that future devices need to be able to identify available and estimated resources across different link, network or band options, which enable approaches to take final decisions about the resource choices. The following sections examine an overview about two main building blocks for data-driven resource-efficient and reliable cMTC networks.

1.1 Dynamic Radio Resource Management Function

To further increase the reliability of MTC networks, especially for critical services, different approaches to increase spectral efficiency are discussed. In addition to pure scheduling based approaches, a data-driven analysis of the spectral power density to predict and avoid technology-independent interferences promise optimization potential. The idea is that the allocation of resources will require proactive monitoring of available resources and prediction of future resources for distributed user equipment as well as centralized network parts. Resource awareness should be supported by ML and new network quality parameters delivered by the various networks (such as their current load level, which is an essential criterion for resource allocation, especially in distributed MTC networks).

For this purpose, the Energy Detection sensing method is implemented based on a software-defined radio for the SRD frequency band. Based on the identified information of interfered frequencies in the SRD band, the Autoregressive Integrated Moving Average (ARIMA) model enables a data-driven channel quality and interference prediction. The idea is to increase the robustness of LPWAN systems by centrally deriving communication profiles that address and bypass the predicted interference characteristics. The updated communication behavior are transmitted in beacons to all decentralized cell participants, which in turn adapt their own communication decisions in favor of lower interference potentials. Based on a model to evaluate the scalability limits of LPWAN technologies [2], three different scheduling approaches are presented for the dynamic enhancement of resource utilization. For this purpose, an evaluation of the statistically expected latency is conducted from the predicted channel activity. Determined areas that meet this requirement are identified as potential transmission areas while maintaining an Age of Information (AoI) shift potential. On this basis the three following procedures with increasing complexity are designed and evaluated in an exemplary scenario (see Fig. 1).

First, in Restricted Access Window (RAW) approach a daily profile is divided into No-Go Areas, in which excessive latency is expected, and Random Access Areas, in which compliance with the latency requirements is given. Transmissions take place only in the Random Access Areas, at a randomly determined time within these intervals. In the case of weighted experience, a subdivision into favorable and unfavorable areas also takes place. Within the random access areas, an additional weighting is performed in favor of periods with very low activity, which results in an average improvement of latency of about 60%. Finally, in the Coordinated RAW process, an optimised transmission

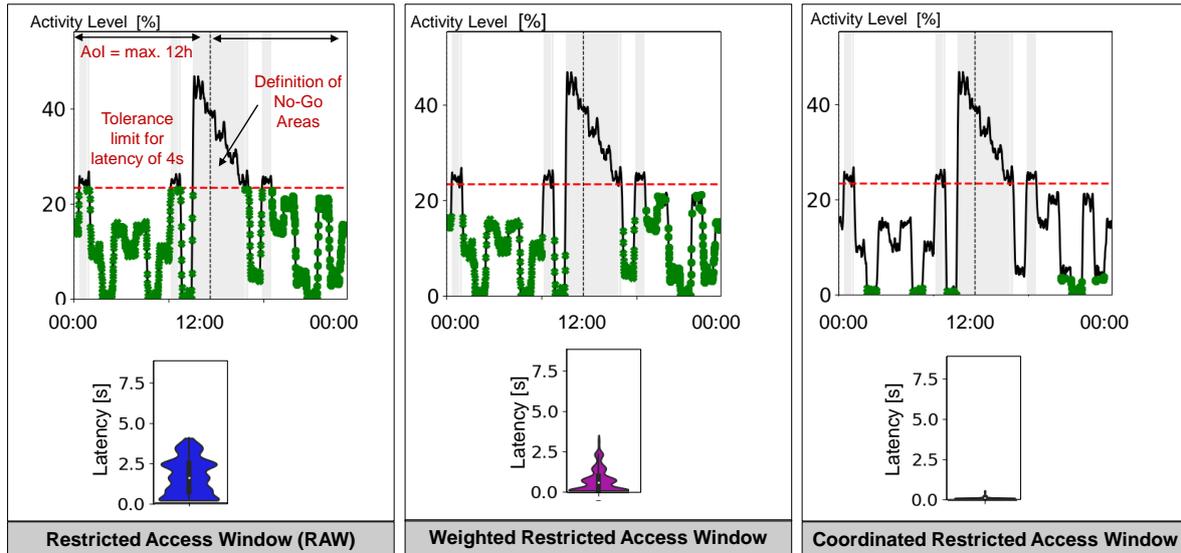


Figure 1: Dynamic Spectrum-Management to improve Scalability of mission-critical MTC Applications

time is iteratively determined for each participant, taking into account the activity of the entire fleet. The Coordinated RAW method exhibits the strongest performance in terms of expected latency due to the tight allocation of suitable transmission time points and achieves a gain of 95% compared to the simple RAW method. However, this is accompanied by an increased computational cost, which must be accounted for a real system.

1.2 Data-Driven Model-Predictive Communication

Apart from intelligently scheduling data transmission, another way for increasing efficiency is to reduce the amount of data that has to be transferred in the first place. Nowadays, initial generations of connected IoT devices and applications enabled by Cellular-IoT (CIoT) and LPWAN technologies are deliberately kept simple and based on equidistant, regular communication intervals. In contrast, we propose an Artificial Intelligence (AI) based model-predictive communication approach taking advantage of knowledge about the underlying data of a sensor system. Therefore an ARIMA model is used in order to depict the behavior of an applications sensor data, only leaving values deviating from the model to be transmitted. Analysis of different environmental sensor systems show a high potential in reducing communication effort using this approach. It can be shown that a data-driven model prediction based on ARIMA can reduce communication effort significantly. Compared to a regular communication interval, the presented approach achieves a reduction of communication events by 87% for a sensor data tolerance of 0.5°C up to 100% for tolerance ranges higher 1°C (see Fig. 2 [3]).

2 Conclusion and Further Research

In future work, the dynamic spectrum-management approach to improve the scalability of mission-critical sensor applications will be further enhanced and evaluated in the lab and real-world environments. In addition, the reliability of mcMTC applications will be

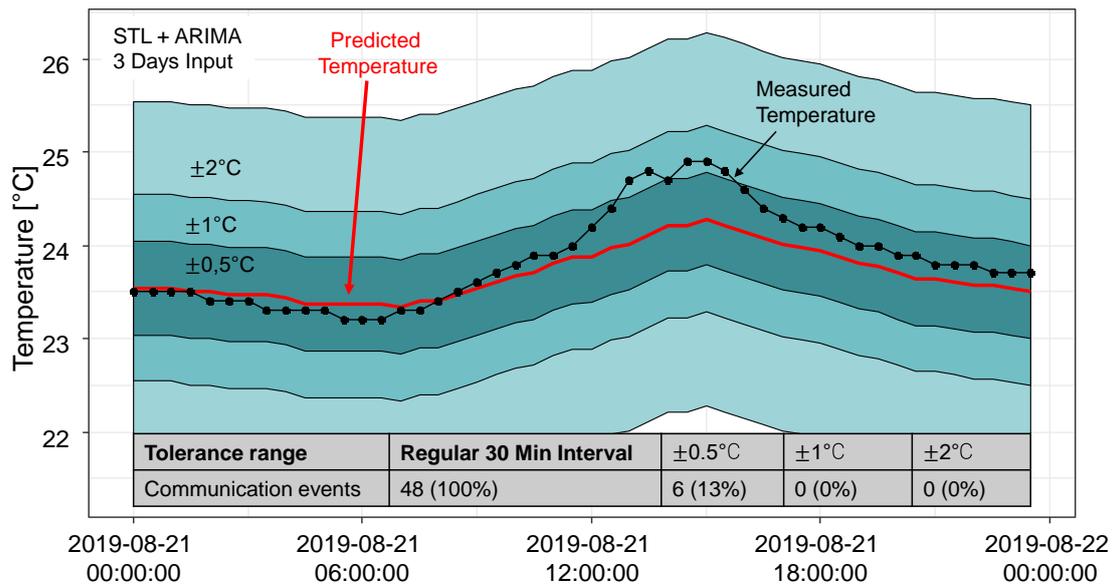


Figure 2: Exemplary forecast using ARIMA shows high potential for reduction of communication effort [3]

increased through ongoing improvements in measures to increase availability, as well as the consideration of future technology measurements, such as mmWave frequency ranges [4].

References

- [1] N. H. Mahmood, S. Böcker, A. Munari, F. Clazzer, I. Moerman, K. Mikhaylov, O. Lopez, O.-S. Park, E. Mercier, H. Bartz, R. Jüntti, R. Pragada, Y. Ma, E. Annanperä, C. Wietfeld, M. Andraud, G. Liva, Y. Chen, E. Garro, F. Burkhardt, H. Alves, C.-F. Liu, Y. Sadi, J.-B. Dore, E. Kim, J. Shin, G.-Y. Park, S.-K. Kim, C. Yoon, K. Anwar, and P. Seppänen, "White paper on critical and massive machine type communication towards 6G," Apr 2020. [Online]. Available: <http://arxiv.org/abs/2004.14146>
- [2] S. Böcker, C. Arendt, P. Jörke, and C. Wietfeld, "LPWAN in the context of 5G: Capability of LoRaWAN to contribute to mMTC," in *2019 IEEE World Forum on Internet of Things (WF-IoT 2019)*, Apr. 2019.
- [3] C. Arendt, S. Böcker, and C. Wietfeld, "Data-driven model-predictive communication for resource-efficient IoT networks," in *2020 IEEE 6th World Forum on Internet of Things (WF-IoT)*, Apr. 2020.
- [4] K. Heimann, J. Tiemann, S. Böcker, and C. Wietfeld, "Cross-bearing based positioning as a feature of 5G millimeter wave beam alignment," in *IEEE 91st Vehicular Technology Conference (VTC-Spring)*, Antwerp, Belgium, May 2020. [Online]. Available: <https://ieeexplore.ieee.org/document/9129182>

Predictive Mobile Network Selection for Resource-Efficient Communication in Crowded Scenarios

Robert Falkenberg
Lehrstuhl für Kommunikationsnetze
Technische Universität Dortmund
robert.falkenberg@tu-dortmund.de

1 Introduction

Increasing data traffic in public mobile networks leads to growing competition for available spectral resources. This is counteracted by opening up additional bands in higher frequency ranges, e.g. in the millimeter-wave range, which, however, is only economically attractive in densely populated areas due to its short range and the resulting smaller cell size. Regions outside this coverage area are therefore left to leverage available resources in the long-range sub-6 GHz frequency range ever more efficiently. As the available spectrum is divided among several network operators, the research question arises whether and how cross-network communication strategies can be used to increase the transmission rate in high-load scenarios. At the same time, the approach itself shall not induce any additional network load, as it would be the case for redundant transmissions or additional speed measurements, for example. Instead, the most promising network is to be selected by a local prediction of the expected data rate, which is estimated by passive observation of the channel quality and the instantaneous network load. Prior work has developed methods for network observation and prediction in controlled environments [1], and prepared them for use in the field with machine learning [2]. With the specially developed monitoring instrument Fast Analysis of LTE Control Channels (FALCON), the observation accuracy could be significantly increased by a faster user detection and the reduction of false detections [3].

This report presents some methods and key results of predictive network selection that emerged from an extensive evaluation under high-load conditions.

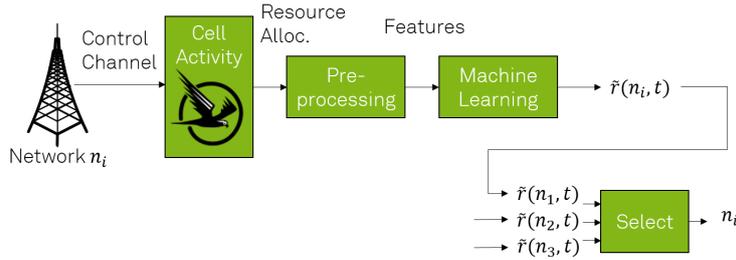


Figure 1: Schematic overview of the predictive network selection approach.

2 Methodology and Measurements

Predictive network selection aims to conduct upcoming transmissions over the network that promises the highest data rate by observing N available networks $\{A, B, C\}$ in parallel. In the saved transmission time, the User Equipment (UE) can switch to energy-saving mode to increase the battery life.

Fig. 1 shows a schematic representation of the processing pipeline of predictive network selection. Each available network is monitored by a FALCON instance, which analyzes the Physical Downlink Control Channel (PDCCH) and records the resource allocations of all active cell users in real-time and with a resolution of 1 ms. A preprocessing step calculates characteristic features from the captured assignments and feeds them to a regression model that is trained on the respective cell using machine learning. Finally, the network with the highest predicted data rate is selected for the transmission.

A large-scale event with extremely high user and data volumes is chosen for the evaluation of the approach, in which participants play an interactive augmented reality game with their smartphones on the grounds of Dortmund's Westfalenpark. A measurement system installed at the park boundary performs observations and transmissions in parallel on three mobile networks during the entire event period. At regular intervals, cell activity is monitored simultaneously in all networks for 5 s using FALCON, followed by a 5 MB data transfer. Transmissions are performed alternately in the downlink and uplink to evaluate both transmission directions independently. The parallel transmissions in all networks allow a subsequent evaluation of the selection made. Since the networks operate in dedicated frequency bands, mutual interference is excluded.

3 Prediction Accuracy

Fundamental for the network selection is the reliable data rate prediction. It is considered as a regression task to train a model f on observed feature vectors \vec{x}_i labeled by outcomes y_i such that $f : \vec{x}_i \rightarrow y_i$ for $i \in \{1, \dots, N\}$ with N being the size of the training set. The model is then applied to perform predictions \hat{y}_j on unlabeled and unseen features \vec{x}_j . Training objective is the minimization of the loss function $L(\hat{\vec{y}}, \vec{y})$ with $\vec{y} = (y_i)_{i \in \{1, \dots, N\}}$ and $\hat{\vec{y}} = (\hat{y}_j)_{j \in \{1, \dots, N\}}$. In this case, Mean Squared Error (MSE) is used as a loss function. Both Artificial Neural Networks (ANNs) with two hidden layers à 5 and 10 nodes and a learning rate $\eta = 0.001$ and Random Forest (RF) with 100 trees, $m_{\text{try}} = 1/3$ and unlimited depth are trained. The dataset comprises 250 to 750 samples per network and direction with \vec{x}_i being a 27-dimensional vector.

Fig. 2 shows the RMSE of the ANN and RF models for each network and direction averaged over all runs of a 10-fold Cross-Validation (CV). The error bars represent the standard deviation of the RMSE across the CV runs.

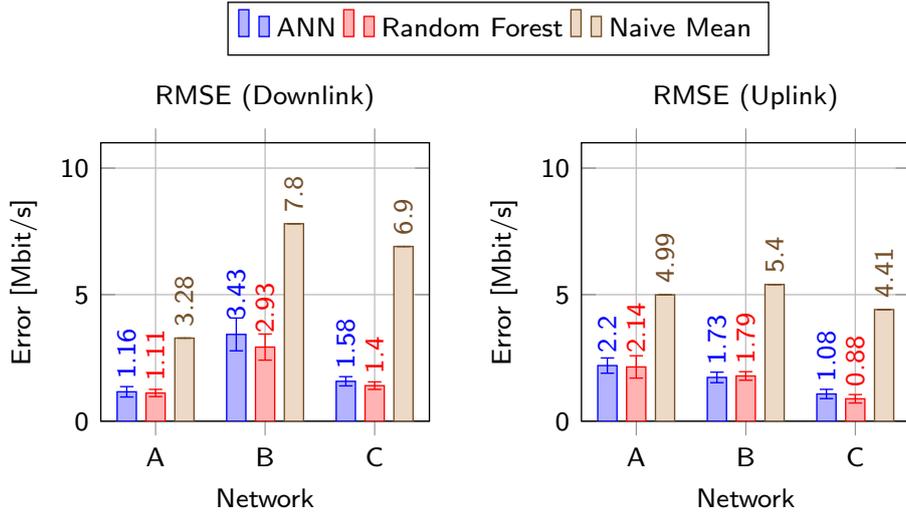


Figure 2: Prediction error in terms of RMSE for ANN and RF models in comparison to naively guessing the sample mean.

The results show that the prediction yields significantly higher accuracy compared to the naïve mean estimation. This is also visible in Fig. 3 which plots the measured vs. the predicted values for network C, and where the predicted values closely follow the first diagonal. The Coefficient of Determination (CoD) spans 0.8 to 0.96 across the scenarios. Moreover, Fig. 2 shows that both methods achieve comparable accuracy and dispersion, with RF having the lowest error in almost all scenarios. Moreover, RF has significantly shorter training times making this method the perfect candidate for further analysis. In particular, methods for dimension reduction, like Principal Components Analysis (PCA) and wrapper-based feature selection have negligible effects on the model accuracy and dispersion, hence the 27-dimensional RF model is not degraded by the presence of numerous collinear features.

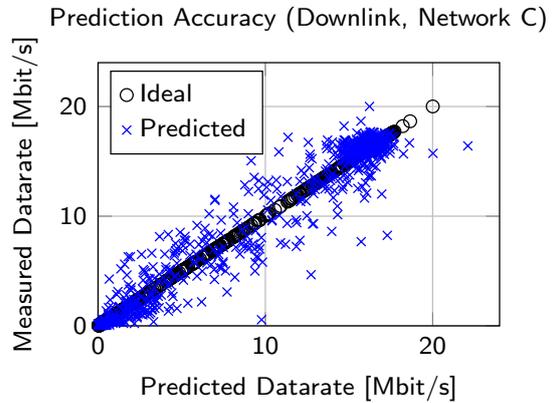


Figure 3: Predicted vs. measured data rate of the ANN for Network C.

4 Evaluation of the Predictive Network Selection

The predictive network selection approach takes the network with the highest predicted data rate for each (pending) transmission. To evaluate this approach based on the conducted experiments, the predicted data rate is computed for each transmission as part of the CV process. In other words, the prediction is always made on unseen data. Fig. 4 (left) highlights the data rate improvement of predictive network selection over a static mapping of the network. It also includes the upper bound of the network selection, by selecting the true maximum of each parallel transmission, which represents a perfect and error-free prediction. While single networks achieve in average 4.58 Mbit/s to 4.96 Mbit/s in downlink direction, the predictive approach reaches 6.66 Mbit/s which lays very close to the achievable maximum of 7.09 Mbit/s. In uplink direction, the data

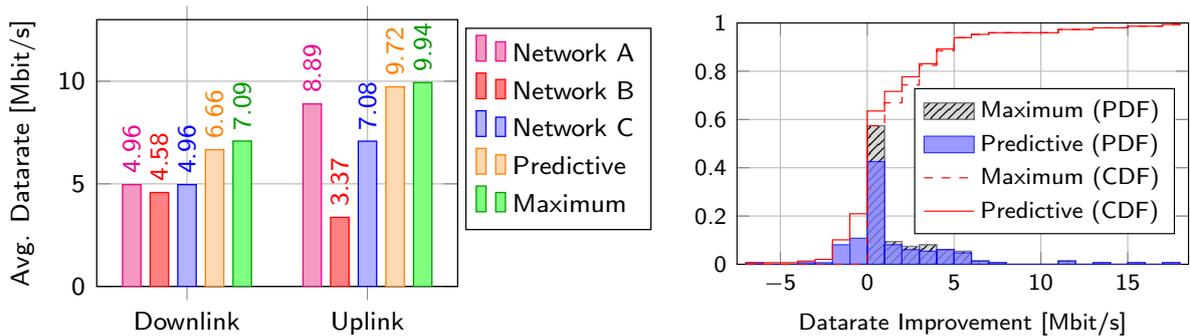


Figure 4: Left: Average data rate with predictive network selection in comparison to the single-network data rate (networks A, B, and C) and in comparison to the retrospective selection of the true maximum, i.e. the upper bound of the approach. Right: Distribution of downlink data rate improvement by predictive network selection versus static transmission over network A only.

rate is dominated by network A, hence the relative improvement by network selection is smaller. Nevertheless, the predictive selection is very close to the maximum.

The proximity to the upper bound by maximum selection is also reflected in the distribution functions. Fig. 4 (right) shows the Probability Density Function (PDF) and Cumulative Distribution Function (CDF) of data rate gain by the predictive approach in comparison to network A. The distribution of the ideal case (by maximum selection) is shown in the background as dashed/hatched elements. Despite occasional mispredictions characterized by a negative gain in the predictive approach, most of the opportunities for data rate increase are taken correctly.

5 Conclusion and Outlook

The results show that passive data rate prediction by monitoring cell-wide resource utilization enables reliable predictions even in public mobile networks. The predictive network selection implemented with this approach leads to a significantly increased average data rate that is close to the optimum. Mobile use cases with opportunistic transmission time selection also benefit to a large extent from the information about the instantaneous network load [4]. Subsequent work will extend the approach to 5G mobile networks.

References

- [1] Robert Falkenberg, Christoph Ide, and Christian Wietfeld. Client-based control channel analysis for connectivity estimation in LTE networks. In *2016 IEEE 84th Vehicular Technology Conference (VTC-Fall)*, Montréal, Canada, Sep 2016. IEEE.
- [2] Robert Falkenberg, Benjamin Sliwa, and Christian Wietfeld. Rushing full speed with LTE-Advanced is economical - a power consumption analysis. In *2017 IEEE 85th Vehicular Technology Conference (VTC Spring)*, Sydney, Australia, Jun 2017.
- [3] Robert Falkenberg and Christian Wietfeld. FALCON: An accurate real-time monitor for client-based mobile network data analytics. In *2019 IEEE Global Communications Conference (GLOBECOM)*, Waikoloa, Hawaii, USA, Dec 2019. IEEE.
- [4] Benjamin Sliwa, Robert Falkenberg, and Christian Wietfeld. Towards cooperative data rate prediction for future mobile and vehicular 6G networks. In *2nd 6G Wireless Summit (6G SUMMIT)*, Levi, Finland, Mar 2020.

Low-Cost and Modular Radio Tomography System for Vehicle Detection and Classification

Marcus Haferkamp
Lehrstuhl für Kommunikationsnetze
Technische Universität Dortmund
marcus.haferkamp@tu-dortmund.de

This report presents a low-cost modular detection and classification system utilizing radio tomography as a novel approach for analyzing heterogeneous road traffic in urban scenarios. In contrast to existing solutions, the proposed system approach allows for a cost-efficient and privacy-preserving ad-hoc deployment in road infrastructures while providing high classification accuracy.

1 Road Traffic Analysis

Precise detection and classification of different road users are required for correctly balancing heterogeneous road traffic flows. Starting from a complex radio system, which allows for reliable binary classification of cars and trucks [1, 3] (cf. Fig. 1, left), we enhance this approach by developing a miniaturized, cost-effective, and modular system setup. This improved system outperforms its predecessor by exploiting multiple in-depth channel parameters for individual links utilizing the radio technologies Ultra-Wideband (UWB) and WLAN and a Machine Learning (ML)-aided classification. The central assumption is that each road user influences specific radio channel parameters resulting in characteristic patterns (*fingerprints*), which are induced, among others, by the shape and materials of passing road users. By evaluating the temporal variance of radio links and their channel parameters, differently demanding applications can be implemented, including simple threshold-based detection (cf. Fig. 1, right) as well as reliable multi-type vehicle classification.

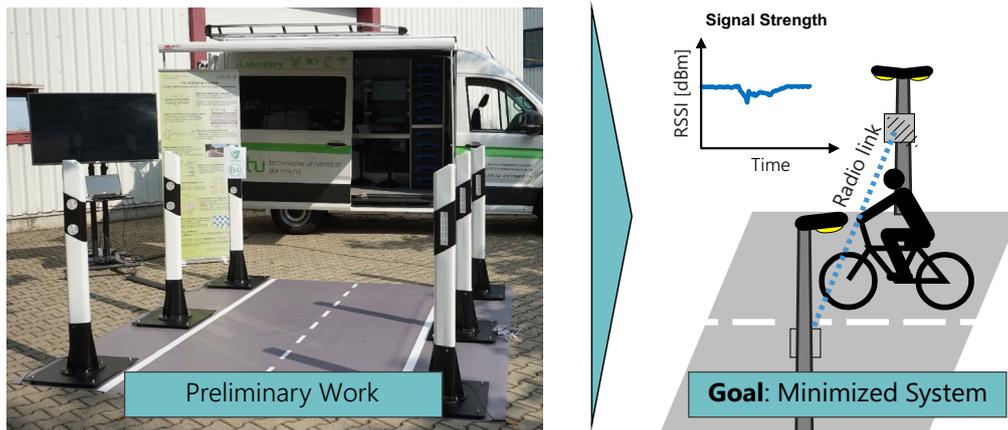


Figure 1: Overview of preliminary work and the project goal. While the previous system contains multiple nodes for vehicle detection and classification, the proposed approach aims at a low-cost and minimized system setup.

The main advantages of the radio-based approach compared to existing alternative solutions are high robustness against adverse weather conditions, low installation and maintenance costs predestining it for comprehensive deployments, and a preservation of the road users' privacy. To achieve this goal, we conduct a cross-method evaluation process containing the following steps:

- **Raytracing Simulations:** As a first step, we evaluate different system configurations respecting the characteristics of the radio technologies UWB and WLAN for different traffic scenarios. We use the resulting insights for the following field measurements.
- **Field Measurements:** Subsequently, we conduct comprehensive field measurements for gathering fingerprints of different road users required for performing an ML-aided detection and classification. Initially, we test and optimize the radio-based system using fingerprints of cyclists measured on a cycle path.
- **Data Processing:** Hereafter, we prepare the captured fingerprints represented as many channel parameters for the ML task. Specifically, we smooth the raw data using Gaussian filters with different kernels and derive multiple statistical means as features.
- **ML-aided Detection and Classification:** Lastly, we feed the extracted features as input for different well-established ML algorithms provided by the framework Lightweight Machine learning for IoT Systems (LIMITS) [2], including Artificial Neural Network (ANN), Random Forest (RF), and Support Vector Machine (SVM).

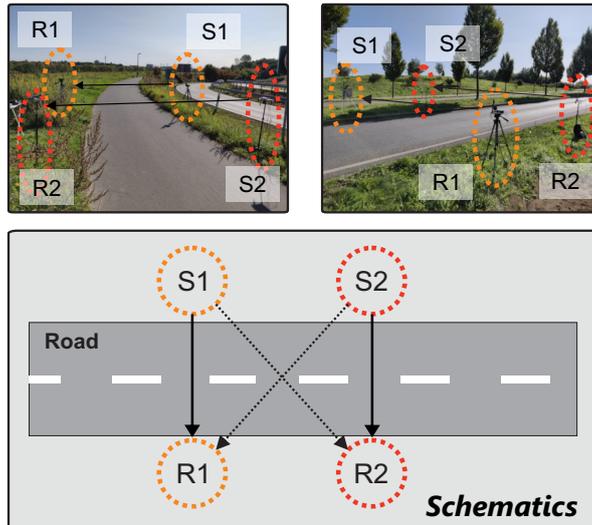


Figure 2: Field deployment setup for measuring cyclists (top left) and passenger cars (top right) utilizing WLAN and UWB transceiver nodes.

2 Real World Validation

To validate the proposed radio tomography-based vehicle detection and classification system, we conduct extensive field measurements gathering fingerprints for both cyclists and passenger cars. Fig. 2 depicts the schematics of the used field deployment setup, which consists of multiple nodes. As sensor nodes, we use Microcontroller Units (MCUs) equipped with WLAN or UWB transceiver modules providing high-resolution channel parameters, i. e., Channel State Information (CSI) or Channel Impulse Response (CIR) data, respectively. The receiver nodes utilize this information for channel estimation, allowing them to successfully decode the originally sent data despite a potential interference of the radio channel. While the usage of Received Signal Strength Indicator (RSSI) data only allows for a coarse-grained channel analysis, the Subcarriers (SCs) provided by WLAN CSI and UWB CIR qualify for an in-depth investigation. Finally, Fig. 3 shows the classification results for detecting cyclists (binary classification) and performing a multi-type vehicle classification (cyclists and passenger cars) using RF. For each task, we have evaluated data representing the passage of specific vehicle types and Line-of-Sight (LOS) measurements, i. e., no obstacle interfering with the radio link. For detecting cyclists, the usage of UWB and the corresponding quotient of the First Path Power (FPP) and the accumulated CIR power results in an accuracy of up to 100%. Regarding the multi-type classification task, an accuracy of more than 98% is achieved using amplitude information of the SCs provided by the *STBC-HT-LTF* training field (WLAN CSI). We note that we have evaluated the channel parameters of a single radio link separately. By correlating multiple radio links or evaluating parameter combinations, we may further improve the overall classification accuracy.

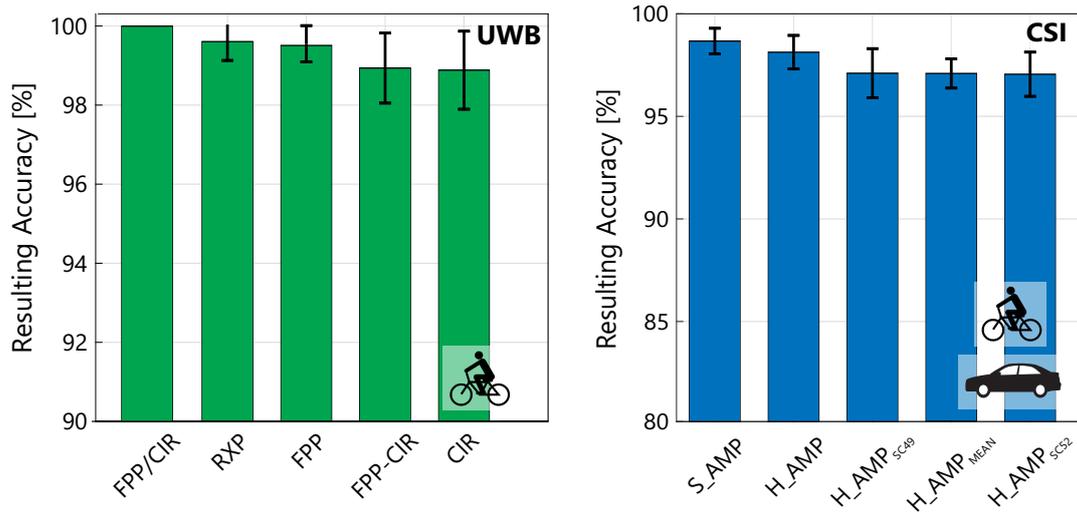


Figure 3: Classification results for detecting cyclists (UWB, *left*) and performing a multi-type vehicle classification (WLAN CSI, *right*) using RF.

3 Conclusion and Further Research

The classification results for the field measurement data show that the proposed modular and radio tomography-based system qualifies as an accurate method for detecting and classifying road users. In future work, we will further enhance the system approach by evaluating the combinations of radio links and high-precision channel parameters.

References

- [1] Marcus Haferkamp, Manar Al-Askary, Dennis Dorn, Benjamin Sliwa, Lars Habel, Michael Schreckenberger, and Christian Wietfeld. Radio-based traffic flow detection and vehicle classification for future smart cities. In *IEEE Vehicular Technology Conference (VTC-Spring) Workshop on Connecting All Things for Enabling Smart Cities (CONTEST)*, Sydney, June 2017.
- [2] Benjamin Sliwa, Nico Piatkowski, and Christian Wietfeld. LIMITS: Lightweight machine learning for IoT systems with resource limitations. In *2020 IEEE International Conference on Communications (ICC)*, Dublin, Ireland, June 2020. Best Paper Award.
- [3] Benjamin Sliwa, Niko Piatkowski, and Christian Wietfeld. The channel as a traffic sensor: Vehicle detection and classification based on radio fingerprinting. *IEEE Internet of Things Journal*, March 2020.

Performance of Application Protocols over NB-IoT: CoAP vs. MQTT for Sensor Networks

Pascal Jörke

Lehrstuhl für Kommunikationsnetze
Technische Universität Dortmund
pascal.joerke@tu-dortmund.de

The Internet of Things (IoT) will enable a variety of new use cases by integrating billions of IoT devices into the internet. Therefore new communication networks have to face upcoming challenges such as high communication ranges, small data transmission efficiency and large scalability. Besides the communication networks itself, application and transport protocols have to perform more efficiently as well. This work addresses the overhead of different application protocols used in cellular IoT environments such as NB-IoT by analyzing the latency, application data rate, battery life and spectral efficiency. The results show, that the UDP-based CoAP protocol clearly outperforms the TCP-based MQTT protocol. When transmitting 10 bytes of application data with 164 dB of coupling loss, acknowledged CoAP transmissions result in a latency that is three to four times lower than transmissions using MQTT with QoS level 2. In terms of power consumption CoAP can provide up to 4.5 times the battery lifetime of MQTT, depending on the use case. MQTT-SN fills the gap between MQTT and CoAP in all scenarios examined in this work. All in all, CoAP is well-suited for transmitting sensor data in challenging environments while MQTT is not recommended due to its large overhead. MQTT-SN enables a better performance than MQTT while still providing a broker-based infrastructure.

1 Application Protocols for the IoT

IoT-relevant communication technologies are designed to be highly efficient for small data transmissions. In order to decrease latency and power consumption, NB-IoT uses techniques like the Radio Resource Control (RRC) resume procedure, which reduces the overall overhead. While NB-IoT defines high-efficient communication in lower layers, all remaining layers of the Open Systems Interconnection (OSI) model need to be redesigned for efficient communication as well. This work focuses on the following application layer protocols:

1. *MQTT*: a TCP-based Publish-Subscribe based Protocol including three Quality of Service (QoS) levels [1]
2. *MQTT-SN*: derived from MQTT, but UDP-based, and with an additional, more lightweight QoS-1 level [4]
3. *CoAP*: also being UDP-based, but using an Request/Response model with acknowledged and unacknowledged data transmission, called confirmable and non-confirmable [3]

Figure 1 gives an overview on the overhead coming with different QoS options:

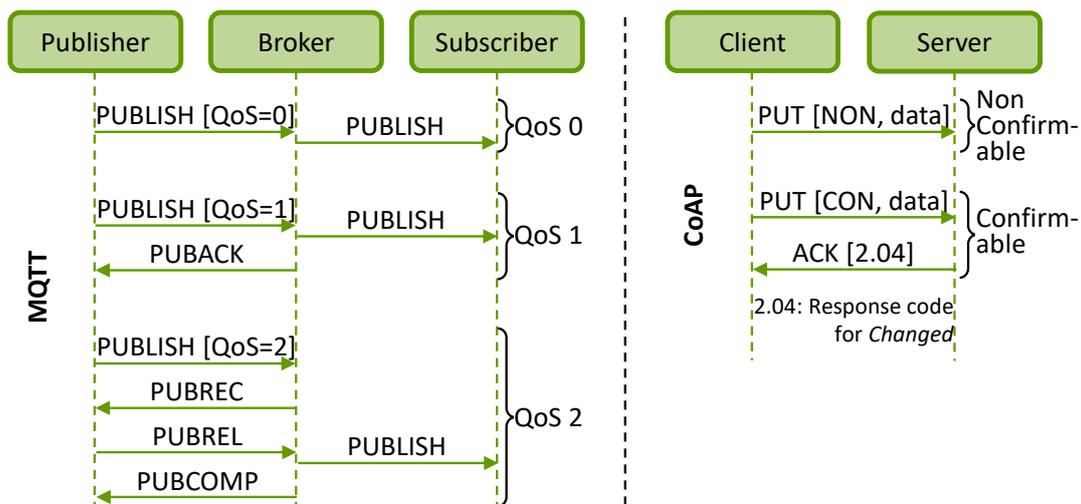


Figure 1: Overview of MQTT and CoAP data transmissions

Based on the power consumption and transmission model introduced in [2], NB-IoT data transmissions are modeled for all three IoT application protocols. Figure 2 shows the resulting messages on the lowest NB-IoT layer, using QoS levels that provide an acknowledgement for the successful data reception at the destination. As expected, MQTT comes with clearly more overhead based on multiple application protocol related messages as well as TCP overhead. When using MQTT-SN instead of MQTT, the overall number of messages is reduced from 65 messages to 33 messages. CoAP reduces the number of NB-IoT messages to only 23 messages.

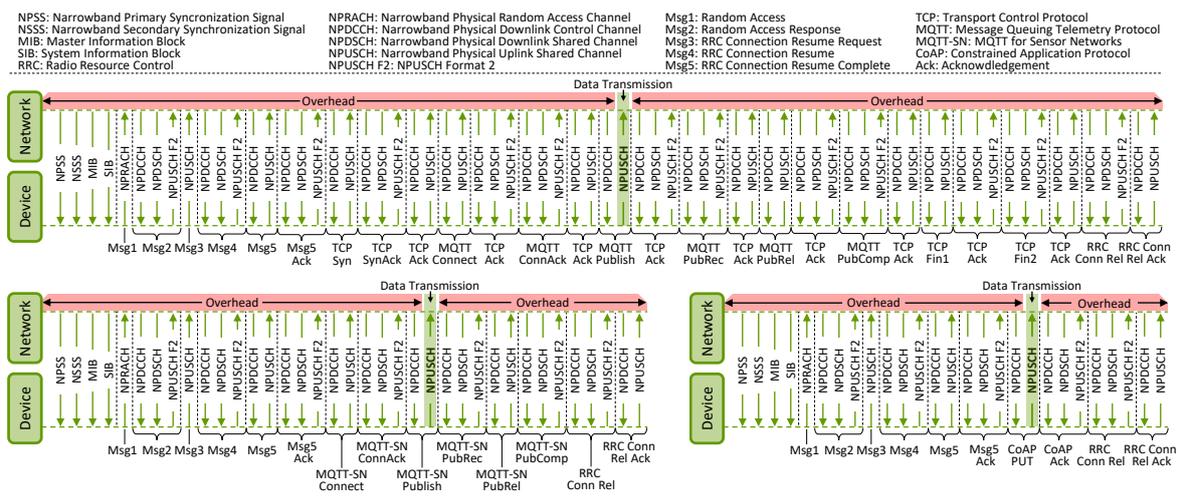


Figure 2: Protocol overhead for MQTT (top), MQTT-SN (bottom left) and CoAP (bottom right) transmissions

2 Evaluation of Application Protocol Performance

With MQTT, MQTT-SN, and CoAP, three IoT-suitable application protocols can be used to transmit data from sensor devices. The major differences between these application protocols are the underlying transport protocols as well as the different messaging systems publish/subscribe and request/response, which directly affect the overhead of each transmission. In this section, all three application protocols will be evaluated in terms of latency, data rate, battery life and spectral efficiency.

Figure 3 gives an overview on the simulation-based results. The results show, that the UDP-based CoAP protocol clearly outperforms MQTT and MQTT-SN. Since MQTT-SN is optimized by using the UDP protocol, it is still the better choice in terms of efficiency compared to MQTT. When transmitting 10 bytes of application data with 164 dB of coupling loss, acknowledged CoAP transmissions result in a latency two to three times lower than transmissions using MQTT with QoS level 2. In terms of power consumption CoAP can provide up to five times the battery life compared to MQTT and up to two times compared to MQTT-SN, respectively.

Since the latency of MQTT is many times over the latency of MQTT-SN and CoAP, its spectral efficiency performs badly as well, while CoAP performs better by a factor of 2.5. All in all, CoAP provides better performance not only for the user's point of view but also for the mobile network operator, as less frequency and time resources are required to transmit the same amount of data. This, in turn, benefits the scalability of the networks.

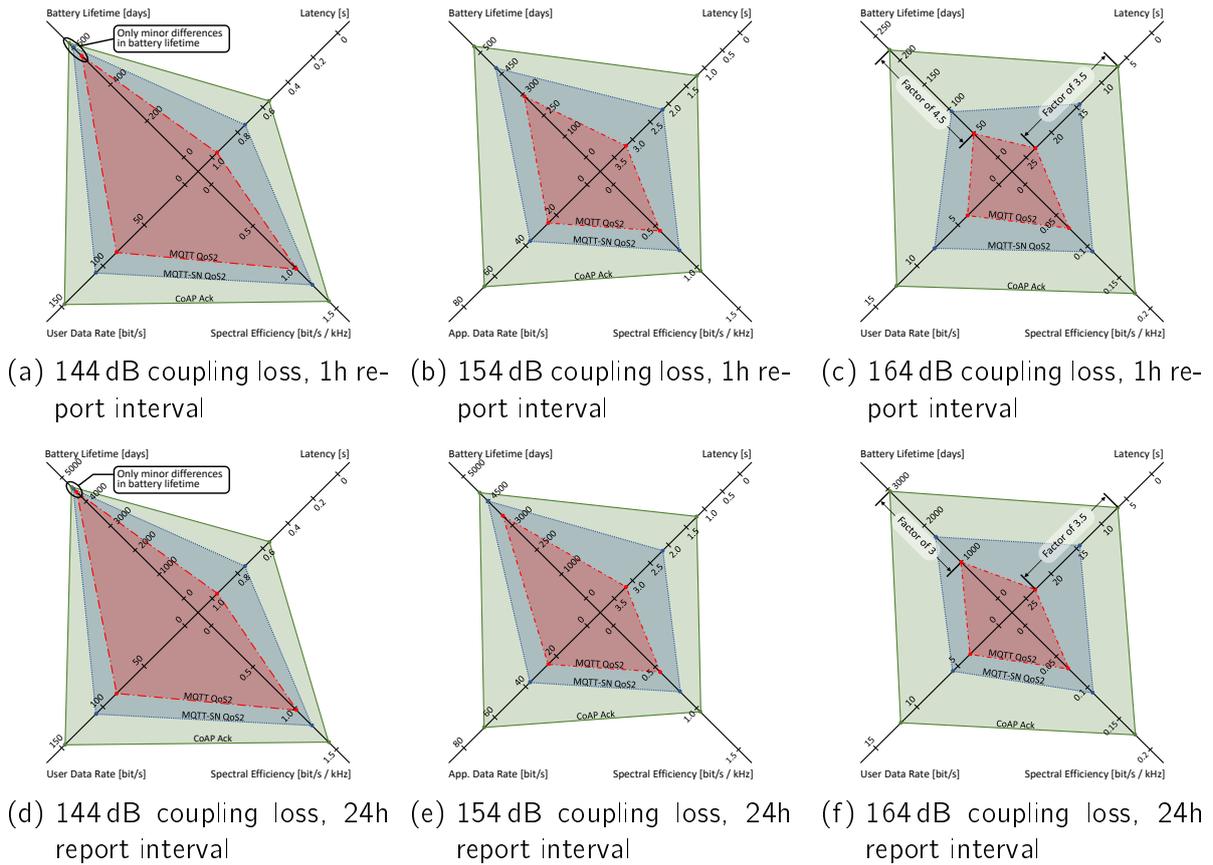


Figure 3: Comparison of app. protocols for 10 Byte data in diff. coverage conditions

3 Conclusion

The analysis has shown that picking the right application protocol is fundamental to fulfil the aim of an energy-efficient transmission for the user and a high-scalable network for the operator. Therefore the focus should not be narrowed down to the lower layers, but all layers, including transport protocols as well as application protocols.

References

- [1] Andrew Banks et al. MQTT Version 3.1.1. *OASIS standard*, 2014.
- [2] P. Jörke et al. Power Consumption Analysis of NB-IoT and eMTC in Challenging Smart City Environments. In *2018 IEEE Globecom Workshops*, pages 1–6, Dec 2018.
- [3] Z. Shelby et al. The Constrained Application Protocol. *Internet Engineering Task Force (IETF)*, 2014.
- [4] Andy Stanford-Clark et al. MQTT For Sensor Networks (MQTT-SN) Protocol Specification Version 1.2. *OASIS standard*, 2013.

The Sensor Floor: An IoT test bed for resource-constraint machine learning applications

Aswin Karthik Ramachandran Venkatapathy
Lehrstuhl für Förder- und Lagerwesen
Technische Universität Dortmund
aswinkarthik.ramachandran@tu-dortmund.de

In this report, the development of the Sensor Floor and its use in the 2020 Summer School hackathon is presented. The Sensor Floor is used in developing heterogeneous IoT applications and deploying machine learning models across aggregated streams of data. The Sensor Floor architecture and hardware components along with the data pipeline are presented. The data pipeline integrates heterogeneous sources that can be post-processed at the developer end. The use of the Sensor Floor and the type of data delivered and the processing of the data as validation for use in the hackathon are presented in this tech report.

1 Introduction

The PhyNetLab [3] was developed as a application-based experimentation platform ultra-low power industrial applications [3]. Due to the limitations provided by the test bed along with diverse requirements presented by IoT applications, the Sensor Floor is conceptualized and developed. It is an experiment infrastructure developed with the main reason for developing and rapid prototyping of physical layer characteristics on a large-scale. The Sensor Floor was conceptualized as a test bed and developed as an iteratively programmable test bed for dual-band application-based research. One such tested application is the Decentralized Brains that was developed for multiple parallel control of heterogeneous swarm systems [2]. It is an application-based testbed where the architecture design supports rapid prototyping of wireless network scenarios with a high

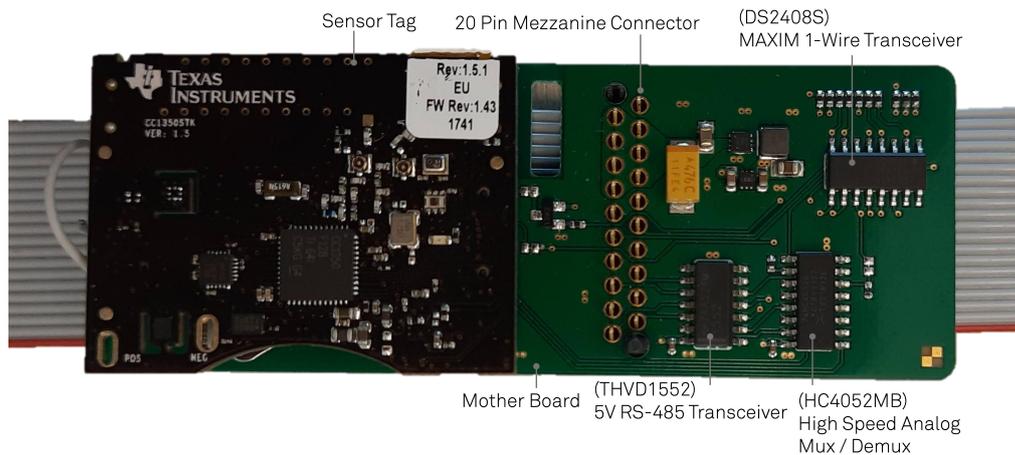


Figure 1: Sensor node with a sensor tag mounter with a mother board and the 20 PIN cable for data, power and control.

node density. The Sensor Floor is called so because 345 nodes are embedded under the research facility's wooden floor. This Sensor Floor infrastructure was used in the development and execution of the 2020 Summer School hackathon. The underlying concept of the hackathon was to develop learning models from large-scale sensors and to develop strategies in such tasks. IoT applications always provide a local measurement. These local measurements can be aggregated along with heterogeneous data streams to provide a global view of various integrated systems. As a first step the concept of the hackathon was developed to tackle the initial issues arising during such heterogeneous application development.

2 The Sensor Floor

The Sensor Floor has three components: the CC1350 STK, a carrier board where the communication bus and power lines are connected, and finally, a sink computer for data acquisition and firmware updates. The Sensor Floor is deployed in a 30-meter long hall with only 23 meter long usable areas. On the other side, it is 15 meters long, where we chose to deploy a cable of 15 meters with every meter for a sensor node. Every 15-meter length is called a strip composing 15 sensor nodes connected with power, 1-wire, and RS422 communication lines. One end of the 15-meter line terminates with two USB dongles for the 1-wire bus and the RS422 bus. The power supply is 12 volts powering each strip. In the end, there were 345 nodes in 23 strips and 23 internet-connected RPis. Each node can be flashed individually, and the whole floor can be flashed in parallel

using a command-line tool developed exclusively for this deployment [1]. The nodes' synchronization messages are delivered every 4 seconds on average due to the design choices of using a 1-wire bus to apply changes to the communication bus. Each column is called a strip with 15 nodes attached to a flat cable strip equidistantly 1 meter apart. The cable connects the 15 nodes to power and a data line. Each node is a sensor tag CC1350 STK with on-board sensors for measuring different physical phenomena. Some of the sensors used for this event are the nine-axis IMU with accelerometer, gyroscope, and magnetometer. The radio interface is used to measure the received signal strength indicator (RSSI).

3 2020 Summer School Hackathon

In the hackathon, we will be using the 9 axis IMU and the RSSI from each of the 345 sensor nodes to localize a robot driving on the floor. There are two tasks in this hackathon, where the first one is about predicting the robot's location using the provided training data set. Once a model is developed, a localization service is then used in task 2 to develop a robot controller using the live data to drive the robot through a predefined path in the hall. This first task was organized in Kaggle. The outcome of the first task provides a localization algorithm for controlling the robot, which will be used in the next task where documentation for controlling a robot's movement is provided and access to live sensor floor data as well. Using both, the participant controlled the robot to drive through predefined waypoints. Fig. 2, shows the inferences that were used for validating the Sensor Floor. The sensors used were the RSSI which is the received signal strength to add dimension to the data. The IMU was also used where a 9 axis data is received. From the fig. 2 it is evident that when a robot moves above a sensor, the sensor reacts according to the changes in the physical phenomena. The magnetometer has spikes while the robot moves above which is due to the changes in the magnetic field of the earth as sensed by the magnetometer. The vibration data is illustrated by the accelerometer and gyroscope data. As there is overall vibration that is recorded, it is possible to see certain peaks in co-relation with the magnetometer. Along with the RSSI data and Vicon data as a benchmark, it was possible to create a learning-based localization algorithm.

4 Future Work

The 2020 summer school successfully used the live data from the Sensor Floor. The participants developed localization algorithms using machine learning and deep neural networks. The various approaches for developing such an algorithm will be summarized for publication. Followed by the publication, the constraints on live data will be investigated. There was a delay of 4 seconds per frame update, this delay can be mitigated

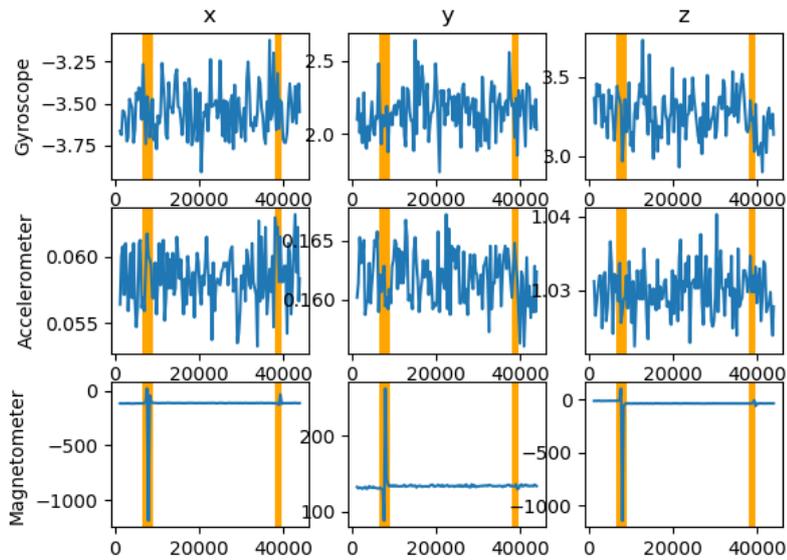


Figure 2: Sensor data interpretation of a single node with all three axes x,y and z from the Gyroscope, Accelerometer, and Magnetometer

by deploying a threshold-based filter. The filter will reduce communication frequency and send data only when the value reaches a certain threshold. This will also allow the sensor to communicate using the wireless medium making it easier to deploy in new environments.

References

- [1] Aswin Karthik Ramachandran Venkatapathy. Sensor floor management software. <https://github.com/akrv/sensorfloor>. Accessed: 2020-12-30.
- [2] Aswin Karthik Ramachandran Venkatapathy, Ariel Ekblaw, Michael ten Hompel, and Joseph Paradiso. Decentralized brain in low data-rate, low power networks for collaborative manoeuvres in space. In *2018 6th IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*, pages 83–88. IEEE, 2018.
- [3] Aswin Karthik Ramachandran Venkatapathy, Moritz Roidl, Andreas Riesner, Jan Emmerich, and Michael ten Hompel. Phynetlab: Architecture design of ultra-low power wireless sensor network testbed. In *2015 IEEE 16th International Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, pages 1–6. IEEE, 2015.



Subproject A6
Resource-efficient Graph Mining

Johannes Fischer

Frank Weichert

A Parallel Framework For Computing a Directed Maximum Cut

Nico Bertram

Lehrstuhl für Algorithm Engineering

Technische Universität Dortmund

nico.bertram@tu-dortmund.de

A Max-Dicut is a partitioning of the nodes of a graph into two subsets which maximizes the sum of the edge-weights between the two subsets. This problem has applications in compressing a string with a contextfree grammar, i.e. a Max-Dicut with high quality leads to a higher compression rates for strings. Unfortunately Max-Dicut is NP-hard so we try to develop an approach which computes the Max-Dicut approximately. We present a parallel framework which partitions a graph with a *graph partitioner* into k subgraphs. On each subgraph we run in parallel a sequential algorithm to compute a local solution. In a final step we try to improve the global solution by merging the locally computed solutions.

1 Directed Max-Cut

A *cut* in a directed and weighted graph $G = (V, E, w)$ is a partitioning of V into the subsets S and T so that the sum of the edge-weights, which cross the subsets S and T , is maximized (an edge is only counted if it points from S to T). The value of a cut with respect to S and T is defined by $W(S, T) = \sum_{i \in S, j \in T} w(i, j)$. The maximum cut is then defined by $W_{max} = \max_{S, T} W(S, T)$. In Fig. 1 we see an example for a cut in a directed graph.

We will denote the problem to find a maximum cut in a directed graph with Max-Dicut. Max-Dicut is NP-hard, which means there is probably no polynomial time algorithm to compute an optimal solution. One common approach in theory and praxis is to solve this kind of problem approximately. That means we develop algorithms which allow for

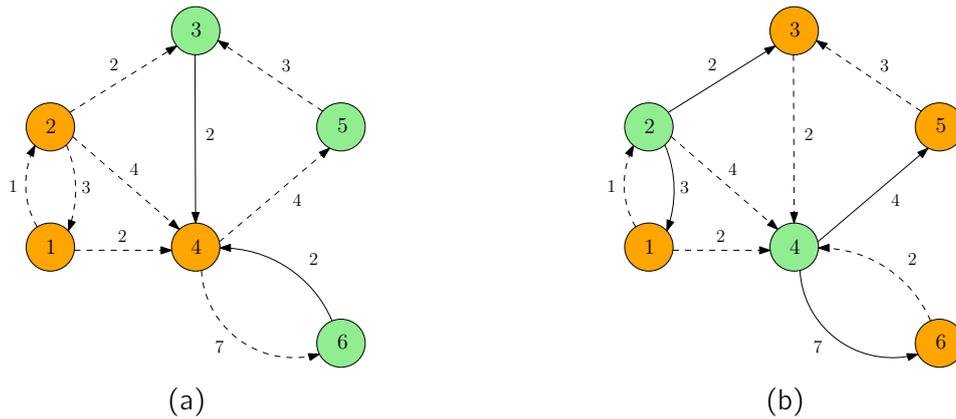


Figure 1: A graph with two example cuts. The nodes which are in S are colored in green and The nodes which are in T are colored in orange. The edges which are not counted for the cut are dashed. The cut in Fig. 1a has the value 4. The cut in Fig. 1b has the value 16, which is the optimal value.

a multiplicative error α with $0 < \alpha < 1$, so that the computed value is in the worst case by the factor α worse than the optimal solution. We call this factor α the *performance guarantee* of an algorithm.

One simple randomized algorithm is to assign each node with probability $\frac{1}{2}$ either to S or T which outputs a solution with an expected performance guarantee of $\frac{1}{4}$. This algorithm can be derandomized with the method of *conditional expectations*. A linear time randomized algorithm with performance guarantee of $\frac{1}{3}$ was described in [2], which can also be randomized to achieve an expected performance guarantee of $\frac{1}{2}$.

The currently best known performance guarantee uses the approach described in [3]. This approach uses a formulation of Max-Dicut as an Integer Program which is then relaxed into a *Semidefinite Program*. By rounding the solution with a random hyperplane the expected performance guarantee of 0.79607 is achieved. This algorithm can be derandomized as well which is described in [8]. The performance guarantee was further improved to 0.859 in [11]. The currently best known performance guarantee of 0.874 was achieved by further improving this approach in [7]. In case that the *Unique Games Conjecture* [6] is true the performance guarantee can only be improved up to 0.878.

2 Application in String Compression

Max-Dicut is used in building a succinct data structure over strings to answer *LCE queries* efficiently. A LCE query over a string s asks for two positions i and j what the longest common prefix of the suffixes starting at position i and j is.

To answer such queries efficiently one could use the *recompression* technique which was described in [5]. With this technique a string s is compressed into a contextfree grammar which generates exactly s . Then on top of the grammar we build a LCE data structure [4]. The memory usage is $\mathcal{O}(z \log(\frac{n}{z}))$ and the query time is $\mathcal{O}(\log(n))$ where z is the size of the Lempel-Ziv 77 factorization [10] and n is the size of s .

During the compression of s we try to find pairs ab and build a rule $X \rightarrow ab$ so that as many pairs are covered by a rule. To do that we build a directed graph G in which each node represents a character of s and we insert an edge from a to b if the pair ab appears in s . Then a cut in G represents a partition of the characters into two subsets S and T , so that we can compress as many pairs ab with $a \in S$ and $b \in T$ as possible without overlapping pairs. So there is a direct correlation between the quality of the computed cut and the compression rate of s .

3 Framework

We are currently working on a parallel framework which computes a Max-Dicut with good quality in Shared and Distributed Memory. Our approach is to partition an input graph G into k subgraphs G_i of roughly equal size, so that the dependency between the subgraphs is minimized, i.e. the edge-weights between the subgraphs is minimized. In recent years many algorithms to compute such graph partitions were developed. In our framework we use the graph partitioner KaHIP [1]. But we want to integrate other graph partitioners such as B-GRAP [9] into our framework as well.

On each computed subgraph G_i we run in parallel a sequential Max-Dicut algorithm to compute a local cut consisting of the subsets S_i and T_i . There are many approaches to merge the computed local cuts to a global cut. One approach is to union each subset S_i to S and each subset T_i to T . The problem with this approach is that we did not consider the edges between the subgraphs. It might be possible that it is more advantageous to swap the subsets S_i and T_i in the global graph. To consider each possible combination of swaps we reduce the problem of merging the local solutions to another Max-Dicut instance. We build a complete Graph H with $2k$ nodes in which each node represents a locally computed subset S_i or T_i . Let X and Y be two nodes of H . We add an edge (X, Y) to H with weight $\sum_{i \in X, j \in Y} w(i, j)$. Then we can run a Max-Dicut algorithm on H . Since the graph has only $2k$ nodes we can use an expensive algorithm to compute an exact solution. In our framework we implemented a simple bruteforce algorithm and an algorithm which solves the Formulation of Max-Dicut as an Integer Program.

References

- [1] Y. Akhremtsev, P. Sanders, and C. Schulz. High-quality shared-memory graph partitioning. *IEEE Transactions on Parallel and Distributed Systems*, 31(11):2710–2722, 2020.
- [2] N. Buchbinder, M. Feldman, J. Naor, and R. Schwartz. A tight linear time $(1/2)$ -approximation for unconstrained submodular maximization. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 649–658, 2012.
- [3] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42(6):1115–1145, 1995.
- [4] Tomohiro I. Longest Common Extensions with Recompression. In J. Kärkkäinen, J. Radoszewski, and W. Rytter, editors, *28th Annual Symposium on Combinatorial Pattern Matching (CPM 2017)*, volume 78 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 18:1–18:15, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.
- [5] A. Jež. Faster fully compressed pattern matching by recompression. In A. Czumaj, K. Mehlhorn, A. Pitts, and R. Wattenhofer, editors, *Automata, Languages, and Programming*, pages 533–544, Berlin, Heidelberg, 2012. Springer Berlin Heidelberg.
- [6] S. Khot, G. Kindler, E. Mossel, and R. O’Donnell. Optimal inapproximability results for max-cut and other 2-variable csps? *SIAM J. Comput.*, 37(1):319–357, 2007.
- [7] M. Lewin, D. Livnat, and U. Zwick. Improved rounding techniques for the max 2-sat and max di-cut problems. In W. J. Cook and A. S. Schulz, editors, *Integer Programming and Combinatorial Optimization*, pages 67–82, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg.
- [8] S. Mahajan and H. Ramesh. Derandomizing approximation algorithms based on semidefinite programming. *SIAM Journal on Computing*, 28(5):1641–1663, 1999.
- [9] A. Moussawi, N. B. Seghouani, and F. Bugiotti. A graph partitioning algorithm for edge or vertex balance. In S. Hartmann, J. Küng, G. Kotsis, A. M. Tjoa, and I. Khalil, editors, *Database and Expert Systems Applications*, pages 23–37. Springer International Publishing, 2020.
- [10] J. Ziv and A. Lempel. A universal algorithm for sequential data compression. *IEEE Transactions on Information Theory*, 23(3):337–343, 1977.
- [11] U. Zwick. Analyzing the max 2-sat and max di-cut approximation algorithms of feige and goemans. 2000.

Computing Runs over General Ordered Alphabets

Jonas Ellert

Lehrstuhl für Algorithm Engineering

Technische Universität Dortmund

jonas.ellert@tu-dortmund.de

A run in a string is a maximal periodic substring. For example, the string `bananatree` contains the run `anana` = $(an)^{3/2}$. There are less than n runs in any length- n string, and computing all runs in a string over a polynomial alphabet takes $\mathcal{O}(n)$ time (Bannai et al., SODA 2015). For general ordered alphabets, the best known result is $\mathcal{O}(n\alpha(n))$ time (where α is the extremely slowly growing inverse Ackermann function, Crochemore et al., SPIRE 2016). We improve this bound to $\mathcal{O}(n)$ time by exploiting combinatorial properties of the Lyndon array. Our new technique results in an algorithm that is not only theoretically fast, but also easy to implement and highly efficient in practice.

A run in a string is a maximal periodic substring. For example, the string `bananatree` contains exactly the runs `anana` = $(an)^{3/2}$ and `ee` = e^2 . Finding all runs in a string is of relevance for applications like compression and text indexing, where identifying repetitive structures is oftentimes crucial for efficient algorithms. Bannai et al. proved that any length- n string contains less than n runs [1], and also provided a new linear time algorithm to compute the runs (even though other linear time algorithms had previously been known [5]). However, all existing algorithms only achieve linear time if the alphabet is *polynomial*, i.e. if the symbols of the string are from the set $\{1, \dots, n^{\mathcal{O}(1)}\}$. We introduce the first linear time runs algorithm that also works for *general ordered alphabets*, i.e. alphabets that are totally ordered, and where the order of any two symbols can be tested in constant time. The best previous algorithm for this problem requires $\mathcal{O}(n\alpha(n))$ time (where α is the extremely slowly growing inverse Ackermann function) [3]. In this report, we give a preview of the techniques that we use in order to achieve linear time. The full description is expected to be published next year.

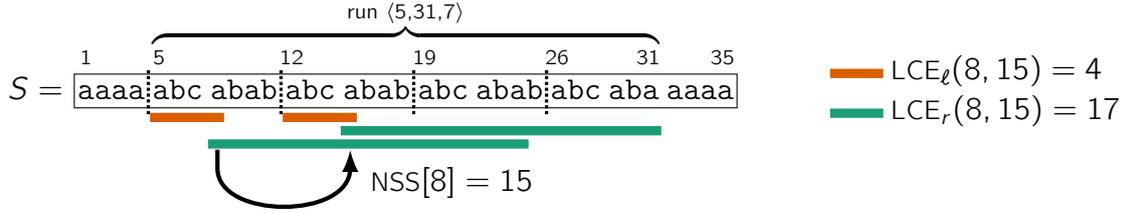


Figure 1: The substring $S[5..31] = (\text{abcabab})^{27/7}$ is a run with period 7. The run is rooted in index $x = 8$ with $y = \text{NSS}[8] = 8 + 7 = 15$. The left border of the run can be computed as $x - \text{LCE}_\ell(x, y) + 1 = 8 - 4 + 1 = 5$. Analogously, the right border is $y + \text{LCE}_r(x, y) - 1 = 15 + 17 - 1 = 31$.

First, we introduce the necessary definitions and notations. We write $[a, b] = (a - 1, b] = [a, b + 1)$ with $a \leq b$ to denote the integer interval $\{a, a + 1, \dots, b\}$. A string S of length $|S| = n$ is a sequence $S[1]S[2] \dots S[n]$ of symbols from a totally ordered alphabet Σ . We write $S[i..j] = S(i - 1..j) = S(i..j + 1)$ with $i \leq j$ to denote the substring $S[i] \dots S[j]$ of length $j - i + 1$. If $i > j$, then $S[i..j]$ is the empty string. The i -th suffix S_i is defined as $S[i..n]$. Given two strings S and T , we say that S is lexicographically smaller than T and write $S \prec T$, if either $|S| < |T|$ and $S = T[1..|S|]$, or if there is some positive integer $\ell < \min(|S|, |T|)$ with $S[1..\ell] = T[1..\ell]$ and $S[\ell + 1] < T[\ell + 1]$.

We say that $p \in \mathbb{N}^+$ is the *shortest period* of a substring $S[i..j]$, if and only if $S[i..j - p] = S[i + p..j]$ and there is no $p' < p$ with $S[i..j - p'] = S[i + p'..j]$. A repetition is a triple $\langle i, j, p \rangle$ where p is the shortest period of $S[i..j]$, and the length $j - i + 1$ of the substring is at least $2p$. A run is a repetition that cannot be extended to either side, i.e. it satisfies $i = 0 \vee S[i - 1] \neq S[i + p - 1]$ and $j = n \vee S[j + 1] \neq S[j - p + 1]$.

Computing Runs from Next Smaller Suffixes. Our solution for computing all runs is an improved version of the algorithm by Bannai et al., which uses the concept of nearest smaller suffixes. For any index $i \in [1, n]$, its so called *next smaller suffix (NSS)* is located at index $\text{NSS}[i] = \min\{j \mid j = n + 1 \vee (j \in (i, n] \wedge S_i \succ S_j)\}$. Less formally, the NSS of S_i is the first suffix that begins to the right of S_i , and is lexicographically smaller than S_i . Analogously, we can define the position $\text{NLS}[i]$ of the next *larger* suffix (NLS) of S_i . The arrays NSS and NLS are also known under the term *Lyndon array* (with a slightly different but equivalent definition). A key property shown by Bannai et al. is that every run is *rooted* in either an NSS or NLS relation. Particularly, for any run $\langle i, j, p \rangle$, there is exactly one index $x \in [i, i + p)$ such that $\text{NSS}[x] = x + p \vee \text{NLS}[x] = x + p$. The index x is then called *root* of the run. An example is provided in Figure 1. Using this insight, we can compute all runs from the arrays NSS and NLS. We simply investigate each index $x \in [1, n]$, and check if x is the root of a run with period $\text{NSS}[x] - x$ or $\text{NLS}[x] - x$.

Bannai et al. compute the arrays NSS and NLS from the inverse suffix array, which takes $\mathcal{O}(n)$ time for polynomial alphabets, and superlinear time for general ordered alphabets.

We use our new algorithm [2] for computing the Lyndon array instead, which runs in linear time even for general ordered alphabets. It remains to be shown how to check whether an index x is the root of a run with period $p = \text{NSS}[x] - x$ (the procedure for period $\text{NLS}[x] - x$ works analogously). Let $y = \text{NSS}[x]$. If x is the root of a run $\langle i, j, p \rangle$ with period $p = y - x$, then clearly $i \leq x$ and $j \geq y$. It is easy to see that $S[i..y]$ is exactly the longest p -periodic substring that ends in position y (otherwise the run could be extended to the left). Similarly, $S[x..j]$ is exactly the longest p -periodic substring that starts at position x . We can find these substrings using the *left-hand* and *right-hand longest common extensions* (*L-* and *R-LCE*), which are defined as

$$\begin{aligned} \text{LCE}_\ell(x, y) &= \max\{\ell \mid x - \ell \geq 0 \quad \wedge \quad S(x - \ell..x) = S(y - \ell..y)\}, \text{ and} \\ \text{LCE}_r(x, y) &= \max\{\ell \mid y + \ell \leq n + 1 \wedge S[x..x + \ell] = S[y..y + \ell]\}. \end{aligned}$$

The left border of the potential run is $i = x - \text{LCE}_\ell(x, y) + 1$, while the right border is $j = y + \text{LCE}_r(x, y) - 1$. Figure 1 visualizes this relation between LCEs and runs. If the LCEs are known, then computing i and j takes constant time. In order to determine whether $\langle i, j, p \rangle$ is actually a run, we only have to check whether its length $j - i + 1$ is at least $2p$, which also takes constant time. Since we consider each index x as the root of exactly two runs (one with period $\text{NSS}[x] - x$ and one with period $\text{NLS}[x] - x$), the time needed to process all indices (apart from the computation time of the LCEs) is $\mathcal{O}(n)$.

Computing the LCE Values. It remains to be shown how to efficiently compute the LCE values $\text{LCE}_\ell(x, \text{NSS}[x])$, $\text{LCE}_r(x, \text{NSS}[x])$, $\text{LCE}_\ell(x, \text{NLS}[x])$, and $\text{LCE}_r(x, \text{NLS}[x])$ for all indices $x \in [1, n]$. Bannai et al. use a data structure that can be precomputed in $\mathcal{O}(n)$ time and answers arbitrary LCE queries in constant time [1, 4]. However, the data structure requires a polynomial alphabet. The best known result for general ordered alphabets exploits a special non-crossing property of the LCEs, and achieves $\mathcal{O}(n\alpha(n))$ time (where α is the inverse Ackermann function) [3]. Our new technique is even more specialized; we achieve linear time for general ordered alphabets by using combinatorial properties of the arrays NSS and NLS . In the remainder of this report, we show an example of how these properties accelerate LCE computations.

Generally, computing $\ell = \text{LCE}_r(x, y)$ for some index x with $y = \text{NSS}[x]$ can be realized in $\mathcal{O}(\ell)$ time by performing $\ell + 1$ individual character comparisons. We simultaneously scan the suffixes S_x and S_y from left to right, searching for the first mismatching character. If we were using this naive strategy to compute all the LCE values, we would not achieve linear time. However, by computing the LCEs in a specific order, we can oftentimes reuse previously computed LCEs. For example, assume that we already computed the aforementioned R-LCE value $\ell = \text{LCE}_r(x, y)$, and now we want to compute another R-LCE value $\ell' = \text{LCE}_r(x', y')$ with $x' < x$ and $y' = \text{NSS}[x'] > y$. In a moment, we will show that (as visualized in Figure 2) it holds $\ell' \geq \ell - (y' - y)$. Therefore, we do not need to compute ℓ' by naively scanning the suffixes $S_{x'}$ and $S_{y'}$ until we find a mismatch. Instead, we can skip the first $\ell - (y' - y)$ character comparisons and compute the LCE value as

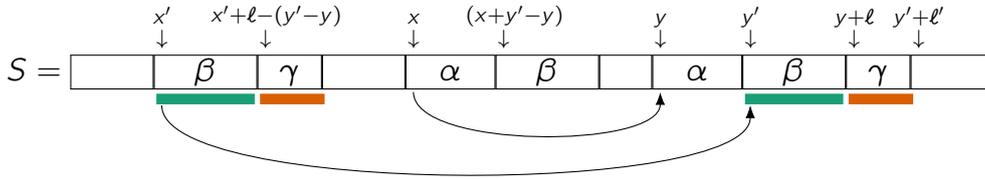


Figure 2: Exploiting previously computed LCE values.

$\ell' = \ell - (y' - y) + \text{LCE}_r(x' + \ell - (y' - y), y + \ell)$. In terms of Figure 2, this means that we only need to scan the text areas that are marked orange, while we skip the teal areas.

Lastly, we show that $\ell' \geq \ell - (y' - y)$. First, note that $x + y' - y$ lies in between x' and $y' = \text{NSS}[x']$. Thus, by definition of next smaller suffixes it holds $S_{y'} \prec S_{x'} \prec S_{x+y'-y}$. From $\ell = \text{LCE}_r(x, y)$ follows $\text{LCE}_r(x + y' - y, y') = \ell - (y' - y)$, which means that $S_{x+y'-y}$ and $S_{y'}$ share a prefix β of length $\ell - (y' - y)$. Now assume that $S_{x'}$ does not start with β , then it holds $S_{x'} \prec S_{x+y'-y} \iff S_{x'} \prec S_{y'}$, which contradicts our previous observation $S_{y'} \prec S_{x'} \prec S_{x+y'-y}$. Thus $S_{x'}$ and $S_{y'}$ share the prefix β , and we have $\ell' = \text{LCE}_r(x', y') \geq |\beta| = \ell - (y' - y)$.

Conclusion and Future Work. We have seen one example of how to accelerate LCE computations by using the results of previous computations. In our upcoming paper, we show that by computing the LCEs in a specific order and carefully distinguishing between a number of different cases, we can compute all LCEs required by the runs algorithm in linear time. In future work, we want to generalize our techniques to also work for the runs computation on tries.

References

- [1] Hideo Bannai, Tomohiro I, Shunsuke Inenaga, Yuto Nakashima, Masayuki Takeda, and Kazuya Tsuruta. The “runs” theorem. *SIAM J. Comput.*, 46(5):1501–1514, 2017.
- [2] Philip Bille, Jonas Ellert, Johannes Fischer, Inge Li Gørtz, Florian Kurpicz, J. Ian Munro, and Eva Rotenberg. Space efficient construction of Lyndon arrays in linear time. In *ICALP 2020*, pages 14:1–14:18, 2020.
- [3] Maxime Crochemore, Costas S. Iliopoulos, Tomasz Kociumaka, Ritu Kundu, Solon P. Pissis, Jakub Radoszewski, Wojciech Rytter, and Tomasz Waleń. Near-optimal computation of runs over general alphabet via non-crossing lce queries. In *SPIRE 2016*, pages 22–34, 2016.
- [4] Johannes Fischer and Volker Heun. Theoretical and practical improvements on the RMQ-problem, with applications to LCA and LCE. In *CPM 2006*, pages 36–48, 2006.
- [5] R. Kolpakov and G. Kucherov. Finding maximal repetitions in a word in linear time. In *FOCS 1999*, pages 596–604, 1999.

Scalable Graph Neural Networks

Matthias Fey

Department of Computer Graphics

TU Dortmund University

matthias.fey@tu-dortmund.de

We have presented the Open Graph Benchmark (OGB), a diverse set of challenging and realistic benchmark datasets to facilitate scalable, robust, and reproducible graph machine learning (ML) research. OGB datasets are large-scale, encompass multiple important graph ML tasks, and cover a diverse range of domains, ranging from social and information networks to biological networks, molecular graphs, source code ASTs, and knowledge graphs. We further present a general framework for scaling arbitrary Graph Neural Networks (GNNs) to large graphs by utilizing historical node activations.

1 Introduction

Graphs are widely used for abstracting complex systems of interacting objects, such as social networks, knowledge graphs, molecular graphs, and biological networks, as well as for modeling 3D objects, manifolds, and source code. Machine learning (ML), especially deep learning, on graphs is an emerging field with significant methodological advances which have produced promising results in applications from diverse domains.

How can we further advance research in graph ML? Historically, high-quality and large-scale datasets have played significant roles in advancing research, as exemplified by, *e.g.*, ImageNet in computer vision. However, in graph ML research, commonly-used datasets and evaluation procedures present issues that may negatively impact future progress.

For example, most of the frequently-used graph datasets are extremely small compared to graphs found in real applications. As models are extensively developed on these small datasets, the majority of them turn out to be not scalable to larger graphs. In fact, the performance of GNNs on these datasets is often unstable and nearly statistically identical

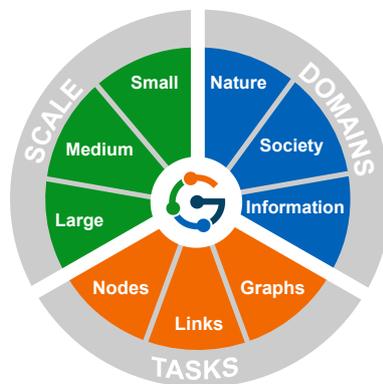
to each other, due to the small number of samples the models are trained and evaluated on.

Despite the indisputable potential of Graph Neural Networks in learning node representations, one of the challenges that have so far precluded their wide adoption in industrial and social applications is the difficulty to scale them to large graphs. In particular, the embedding of a given node depends recursively on all its neighbor’s embeddings, leading to high inter-dependency between nodes that grows exponentially with respect to the number of layers. Hence, the scalability of Graph Neural Networks is challenging, and demands for approximations to the underlying computation.

2 OGB: A Benchmark Suite of Large-scale Graphs

There is an urgent need for a comprehensive suite of real-world benchmarks that combine a diverse set of datasets of various sizes coming from different domains. Data splits as well as evaluation metrics are important so that progress can be measured in a consistent and reproducible way.

Towards this goal, we have presented the Open Graph Benchmark (OGB) in order to facilitate scalable, robust, and reproducible graph ML research [2]. The premise of OGB is to develop a diverse set of challenging and realistic benchmark datasets that can empower the rigorous advancements in graph ML. The OGB datasets are designed to have the following three characteristics:



1. *Large-scale*: The OGB datasets are orders-of-magnitude larger than existing benchmarks and can be categorized into three different scales (small, medium, and large). Even the “small” OGB graphs have more than 100 thousand nodes or more than 1 million edges, but are small enough to fit into the memory of a single GPU, making them suitable for testing computationally intensive algorithms. Additionally, OGB introduces “medium” (1 million nodes or 10 million edges) and “large” (100 million nodes or 1 billion edges) datasets, which can facilitate the development of scalable models based on mini-batching and distributed training.
2. *Diverse domains*: The OGB datasets aim to include graphs that are representative of a wide range of domains. The broad coverage of domains in OGB empowers the development and demonstration of general-purpose models, and can be used to distinguish them from domain-specific techniques. Furthermore, for each dataset, OGB adopts domain-specific data splits (*e.g.*, based on time, species, molecular

structure, GitHub project, etc.) that are more realistic and meaningful than conventional random splits.

3. *Multiple task categories*: Besides data diversity, OGB supports three categories of fundamental graph ML tasks, *i.e.*, node, link, and graph property predictions.

Currently, OGB includes 15 diverse graph datasets, with at least 4 datasets for each task category. Thorough extensive benchmark experiments, we have highlighted research challenges and opportunities provided by each dataset, especially on (1) scaling models to large graphs, and (2) improving out-of-distribution generalization performance under the realistic data split scenarios. The OGB website (<https://ogb.stanford.edu>) provides the documentation, example scripts, and public leaderboards.

3 Scaling-up Graph Neural Networks

While the full-gradient in GNNs is straightforward to compute since we have access to *all* hidden node representations of *all* layers, this is not feasible in large-scale graphs due to memory limitations and slow convergence [3]. Therefore, given a loss function ϕ , it is desirable to approximate its full-batch gradient stochastically via

$$\nabla \mathcal{L}(\theta) = \frac{1}{|\mathcal{V}|} \sum_{v \in \mathcal{V}} \nabla \phi(\mathbf{h}_v^{(L)}, y_v) \approx \frac{1}{|\mathcal{B}|} \sum_{v \in \mathcal{B} \subseteq \mathcal{V}} \nabla \phi(\mathbf{h}_v^{(L)}, y_v), \quad (1)$$

which considers only a mini-batch $\mathcal{B} \subseteq \mathcal{V}$ of nodes for loss computation. However, this stochastic gradient is still expensive to compute due to the recursive inter-dependency of nodes, a phenomenon framed as *neighbor explosion*. Specifically, the representation of a given node depends recursively on all its neighbor’s representations, and the number of dependencies grows exponentially with respect to the number of layers. Recent works try to alleviate this problem by proposing various different sampling techniques [3]. *Node-wise sampling* recursively samples a fixed number k of 1-hop neighbors, *i.e.*, $|\tilde{\mathcal{N}}^{(\ell)}(v)| = k$. In *subgraph sampling*, a full GNN is run on an entire subgraph $\mathcal{G}[\mathcal{B}]$ induced by a sampled batch of nodes $\mathcal{B} \subseteq \mathcal{V}$. These approaches get completely rid of the neighbor explosion problem, but challenges arise to preserve most of the original edges in order to still present a meaningful topological structure.

An alternative to the aforementioned approaches is to utilize the historical activations of nodes. Specifically, the idea is to maintain a history $\bar{\mathbf{h}}_v^{(\ell)}$ of activations $\mathbf{h}_v^{(\ell)}$ for each node $v \in \mathcal{V}$ and each layer $\ell \in \{1, \dots, L-1\}$ as an affordable approximation, given that the model weights do not change too fast during training. Let $\mathbf{h}_v^{(\ell)}$ denote the exact node embedding in layer ℓ of a node $v \in \mathcal{B}$ in a mini-batch $\mathcal{B} \subseteq \mathcal{V}$, and let $\tilde{\mathbf{h}}_v^{(\ell)}$ denote its

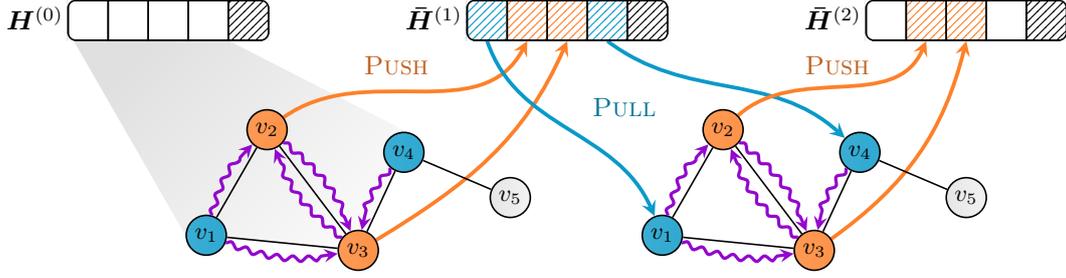


Figure 1: **High-level illustration of the mini-batch processing of GNNs via historical activations.** Here, ■ denote the nodes in the current mini-batch and ■ represent their direct neighbors. Nodes in the current mini-batch *push* their updated activations to the history $\bar{\mathbf{H}}^{(\ell)}$, while their direct neighbors *pull* their most recent historical activations from $\bar{\mathbf{H}}^{(\ell)}$ for further processing.

estimated node embedding. Then, we can approximate $\mathbf{h}_v^{(\ell)}$ via

$$\begin{aligned}
 \mathbf{h}_v^{(\ell)} &= \mathbf{f}_\theta^{(\ell)} \left(\mathbf{h}_v^{(\ell-1)}, \left\{ \left\{ \mathbf{h}_w^{(\ell-1)} : w \in \mathcal{N}(v) \right\} \right\} \right) \\
 &= \mathbf{f}_\theta^{(\ell)} \left(\mathbf{h}_v^{(\ell-1)}, \left\{ \left\{ \mathbf{h}_w^{(\ell-1)} : w \in \mathcal{N}(v) \cap \mathcal{B} \right\} \cup \left\{ \left\{ \mathbf{h}_w^{(\ell-1)} : w \in \mathcal{N}(v) \setminus \mathcal{B} \right\} \right\} \right) \\
 &\approx \mathbf{f}_\theta^{(\ell)} \left(\underbrace{\left\{ \tilde{\mathbf{h}}_v^{(\ell-1)} \right\}}_{\text{Current activations}}, \underbrace{\left\{ \left\{ \tilde{\mathbf{h}}_w^{(\ell-1)} : w \in \mathcal{N}(v) \cap \mathcal{B} \right\} \cup \left\{ \left\{ \tilde{\mathbf{h}}_w^{(\ell-1)} : w \in \mathcal{N}(v) \setminus \mathcal{B} \right\} \right\} \right)}_{\text{Historical activations}} \right) =: \tilde{\mathbf{h}}_v^{(\ell)}, \quad (2)
 \end{aligned}$$

where we substitute estimated activations $\tilde{\mathbf{h}}_w^{(\ell-1)}$ by historical activations $\bar{\mathbf{h}}_w^{(\ell-1)}$ for neighboring nodes outside of the batch. This scheme is a generalization from the specific versions formulated by [1]. A high-level illustration of its computation flow is visualized in Figure 1. Despite reducing the variance of estimation in comparison to other sampling approaches, historical activations provide several advantages: (1) They allow the GNN to make use of all available neighborhood information, *i.e.* *no* edges are dropped. (2) By restricting the GNN to the 1-hop neighborhood around each node $v \in \mathcal{B}$, this scheme does not need to maintain recursive layer-wise computation graphs. (3) Computing node embeddings via historical activations provides theoretical guarantees. In particular, if the model weights are kept fixed, $\tilde{\mathbf{h}}_v^{(\ell)}$ eventually equals $\mathbf{h}_v^{(\ell)}$.

References

- [1] J. Chen, J. Zhu, and L. Song. Stochastic training of graph convolutional networks with variance reduction. In *ICML*, 2018.
- [2] W. Hu, M. Fey, M. Zitnik, Y. Dong, H. Ren, B. Liu, M. Catasta, and J. Leskovec. Open Graph Benchmark: Datasets for machine learning on graphs. In *NeurIPS*, 2020.
- [3] Y. Ma and J. Tang. *Deep Learning on Graphs*. Cambridge University Press, 2020.



Subproject B2

Resource optimizing real time analysis of
artifactual image sequences for the
detection of nano objects

Roland Hergenröder

Frank Weichert

Towards Robust Spatiotemporal Detection of Nanoparticles

Konstantin Wüstefeld
Department of Computer Graphics
TU Dortmund University
konstantin.wuestefeld@tu-dortmund.de

The ongoing Covid-19 pandemic stresses the importance of medical testing and in particular the availability of rapid tests for virus infections. The PAMONO (Plasmon-Assisted Microscopy of Nano-Objects) sensor being the basis of this research project is capable of detecting viruses and virus-like particles in nanometer scale. This report briefly describes earlier work, recent work and observations regarding the spatiotemporal detection of nanoparticles in artifactitious image streams. At the end ideas for future research are outlined.

1 Previous Work on Nanoparticle Detection

Earlier works related to the PAMONO sensor have shown that it is capable of recording the signals of nanometer sized viruses and virus-like particles [4, 5]. While classic, non-deep learning approaches can achieve high detection rates, their need for manual adjustments of parameters to fit a certain characteristic is a drawback for the usage outside a controlled laboratory environment. Imagining a usage as a rapid test on-site, it is infeasible to manually determine the needed parameters frequently in order to compensate dynamically changing external influence.

An earlier approach [3] based on convolutional neural networks reduced user interaction for the detection part while relying on noise reduction methods and shape based checks from a former detection pipeline [4]. These parts were manually adapted to a specific characteristic of particle signals. The detection itself is done by classifying fixed size tiles based on Fourier respectively Wavelet features for each tile separately [3]. While

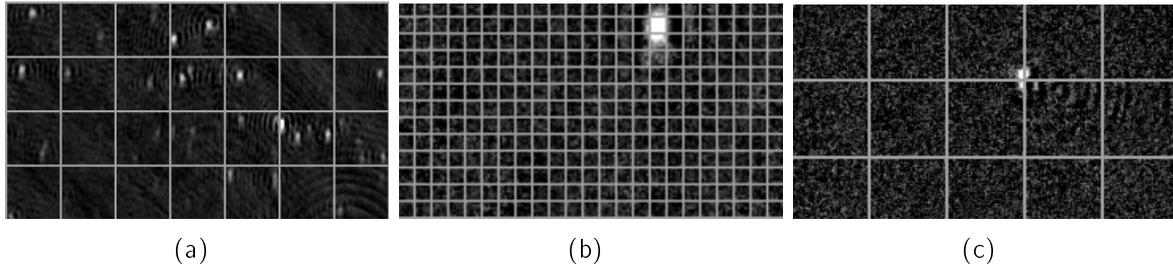


Figure 1: Examples of problematic cases with fixed size tiling: (a) Multiple particles in one cell, (b) cell size smaller than particle size, (c) cell borders distributing a particle over multiple cells.

achieving good results on similar datasets to those which were used for training, detection rates are decreasing significantly for changing characteristics [3].

2 Difficulties with Classifications on Fixed Size Tiles

While approaches based on the classification of fixed size tiles can yield good results coupled with low execution times [3], the principle of tiling is sensitive to the choice of the tile size. A too small tile size as well as higher particle concentration in a sample can make the particle signal indistinguishable from the background signal by spatial features as there is no delimitation against a reference region in one tile. A higher tile size increases the probability of multiple particle signals lying in the same tile. When multiple particles are located in one tile, a maximum of one particle can be found in the binary classification. An issue which also has to be handled with fixed size tiling is cutting a particle into multiple parts distributing it over neighboring classification regions. This can lead to duplicate as well as missing detections. Figure 1 shows visualizations for the mentioned problem cases. A possible solution for this is to analyze the same image multiple times with different tile sizes which at the same time increases the need for additional memory and calculation time. Another solution is presented in the next section.

3 Spatiotemporal Object Detection Based on Deep Learning

The recently presented spatiotemporal pipeline [6] shown in Figure 2 was developed to overcome the problems originating from dividing the original image into fixed size tiles and to take regional information around particle candidates stronger into account than it would be possible with the classic tiling strategy. Using the Mask R-CNN [2] network

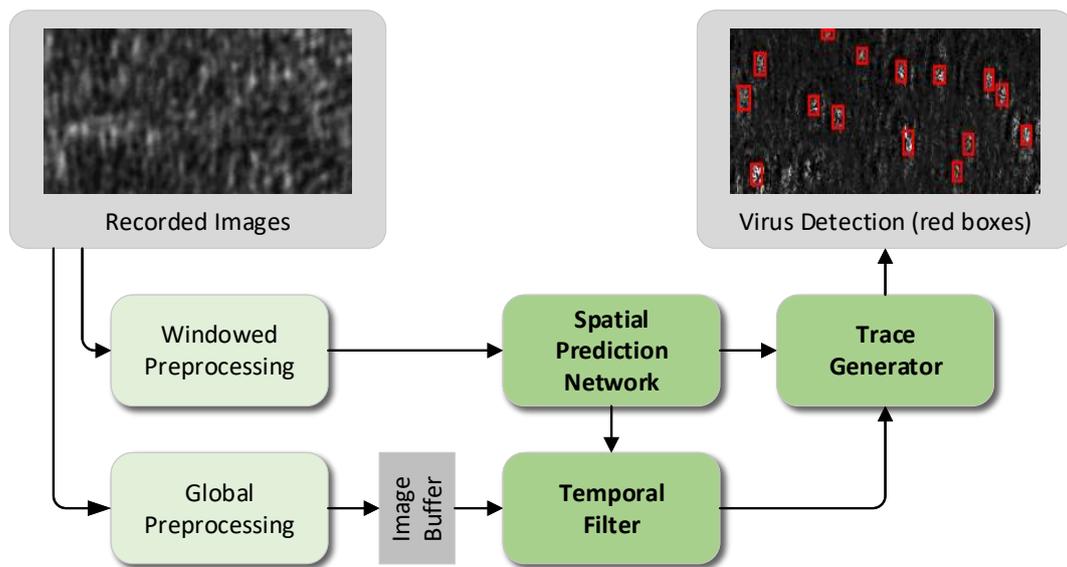


Figure 2: Architecture of the spatiotemporal pipeline for (virus) particle detection [6].

for spatial predictions, candidate regions of different sizes can be proposed so that the need for choosing a fixed tile size is removed. For merging candidates of different sizes at overlapping regions the method of non-maximum suppression [1] is used. It selects only the region with the highest confidence score from a set of overlapping regions determined by an intersection over union value exceeding a given threshold. After the regions are proposed purely locally they are checked from a temporal point of view in order to distinguish signals originating from particles of interest from those of disturbances like dust particles or air bubbles in the sample. This is done in a fully connected network of three layers with time features relativized by the spatial vicinity of the candidate region and normalized over a time window. The effect is that the classification network has to take the characteristic shape of the time series into account instead of overfitting to absolute values. The presented pipeline was able to achieve a count exactness of over 80% on synthetic particles of 80 nm, 100 nm and 200 nm contained in ten recorded samples in total. These results were achieved by training only 200 nm sized particles which shows the generalization capability of the pipeline. Despite satisfying results on the tested datasets recorded in a laboratory it is expected that lower values will be reached when analyzing samples with stronger disturbances and less intense particle signals.

4 Current and Future Research

One of the challenges in detecting the signals of nanoparticles is the ability to detect particles of smaller sizes. With decreasing particle sizes resulting in weaker signals of interest, handling disturbances becomes increasingly important. This is especially true

for the use as a rapid test on-site where external influences such as tremors in the vicinity of the device and unclean samples contaminated with dust particles have to be taken into account. Most of the disturbance characteristics can be sorted into one of the categories of random noise, global lightness differences over time, locally restricted lightness differences over time, wave-like and line-like structures. All of these are expected to become stronger when using the device outside the laboratory where it is placed on an air-decoupled desk [4]. For that reason the current focus lies on methods that either remove artifacts directly before the actual detection takes place or make the detection itself more robust against disturbances by training it with augmented data. At the same time the used methods should require as few manual interactions as possible to enable also non-experts to use the system in a short time.

References

- [1] Ross B. Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. *CoRR*, abs/1311.2524, 2013.
- [2] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross B. Girshick. Mask R-CNN. *CoRR*, abs/1703.06870, 2017.
- [3] Jan Eric Lenssen, Anas Toma, Albert Seebold, Victoria Shpacovitch, Pascal Libuschewski, Frank Weichert, Jian-Jia Chen, and Roland Hergenröder. Real-time low SNR signal processing for nanoparticle analysis with deep neural networks. In *International Conference on Bio-inspired Systems and Signal Processing (BIOSIGNALS)*, 2018.
- [4] Pascal Libuschewski. *Exploration of cyber-physical systems for GPGPU computer vision-based detection of biological viruses*. PhD thesis, Technical University Dortmund, 2017.
- [5] Dominic Siedhoff. *A parameter-optimizing model-based approach to the analysis of low-SNR image sequences for biological virus detection*. PhD thesis, Technical University Dortmund, 2016.
- [6] Konstantin Wüstefeld and Frank Weichert. An automated rapid test for viral nanoparticles based on spatiotemporal deep learning. In *IEEE Sensors 2020 (accepted for publication)*, 2020.



Subproject B3
Data Mining on Sensor Data of Automated
Processes

Jochen Deuse Katharina Morik
Petra Wiederkehr

Detection of unstable milling processes using wavelets

Felix Finkeldey
Chair 14 for Software Engineering
Department of Computer Science
TU Dortmund University
felix.finkeldey@tu-dortmund.de

Different challenges arise for process optimizations of milling processes, e.g., minimizing tool load or avoiding chatter vibrations. Tool vibrations can deteriorate the machined workpiece surface and lead to scrap parts. Detecting unstable vibrations during milling operations can be crucial for process monitoring, in order to adjust process parameter values accordingly. This work presents a methodology to identify unstable process states based on an analysis of acoustic emission signals using the wavelet transform. To validate the proposed approach, different process configurations were investigated.

1 Introduction

In the aerospace industry, where long and slender milling tools are often used to reach deep cavities of structural components, avoiding chatter vibrations is a challenging task [1]. In addition, when difficult-to-cut materials are machined, tool wear, which can slowly progress but also change abruptly, can alter the dynamic behavior of the system consisting of the machine tool, the spindle and the milling tool. Thus, previously identified stable process configurations can still lead to instabilities for long-running processes. The field of process monitoring deals with the observation of characteristics, e.g., process forces, tool vibrations or tool conditions, parallel to process conduction in order to identify critical conditions and adapt process parameter values accordingly. However, using external sensors, which can measure forces or vibrations, is often not feasible in an industrial manufacturing environment. In contrast, acoustic emission sensors are easy to position

near the machining area. This report presents an approach to analyze acoustic emission signals acquired for milling processes, using the wavelet transform, in order to identify unstable process behaviors. The proposed approach was applied for the identification of stability limits in milling processes for different inclination angles of the tool [2]. In addition, the results of an application of this methodology could also be used for a following classification learning objective, whereby instabilities could be predicted with a specific prediction horizon. This would allow predicting critical process conditions based on in-process measurements before they would occur. Subsequently, a prevention of these conditions would be possible by adjusting the process parameter values to ensure stability for the current engagement situation.

2 Detection of chatter vibrations

A discretized version of the continuous wavelet transform [3]

$$W(a, b) = \sum_{i=1}^N \xi^{(i)} \left(\frac{\delta_t}{a} \right)^{\frac{1}{2}} \Psi^* \left[\frac{(i-b)\delta_t}{a} \right] \quad (1)$$

was used to analyze each acoustic emission signal, whereby $\xi^{(i)}$ is the i -th sample of the acoustic emission signal, N is the total number of samples, Ψ^* is the complex conjugate of a mother wavelet Ψ , δ_t is the time difference between two samples and a and b are the scaling and translation variables, respectively. In contrast to using spectrograms, there is no trade-off between time and frequency resolutions when using the wavelet transform. Each process section of an index

$$i_{\text{crit}} = i \in \left\{ \mathbb{N} \mid \bigcap_a T(W^{(i)}(a), N_T) > \overline{W}(a) + \tau \cdot W_s(a), \right. \quad (2)$$

$$\left. \forall i \in \{1, \dots, N\}, \forall a \in S \right\}, \quad (3)$$

$$\overline{W}(a) = \frac{1}{N} \sum_{i=1}^N W^{(i)}(a), \quad (4)$$

$$W_s(a) = \sqrt{\frac{\sum_{i=1}^N (W^{(i)}(a) - \overline{W}(a))^2}{N-1}} \quad (5)$$

of a considered acoustic emission signal was classified as unstable, whereby $W(a)$ is the wavelet transform of the signal, τ is a user-defined threshold and S is the set of investigated scales, which correspond to a set of natural frequencies of a measured frequency response function (FRF) of a specific system consisting of the machine tool, the spindle and the milling tool. When identifying FRFs by impact hammer tests, natural

frequencies were measured for a non-rotating spindle. Since these frequencies may differ from natural frequencies of a rotating system, the function $T(W^{(i)}(a), N_T)$ calculates a weighted average of the wavelet intensities, incorporating N_T neighboring frequencies of the natural frequency, which corresponds to the scaling a . For the weighting, the Blackman window function [4]

$$w(j) = \alpha_0 - \alpha_1 \cos\left(\frac{2\pi j}{N_T}\right) + \alpha_2 \cos\left(\frac{4\pi j}{N_T}\right) \quad (6)$$

with

$$\alpha_0 = 0.42, \quad \alpha_1 = 0.5, \quad \alpha_2 = 0.08 \quad (7)$$

was used, whereby $W^{(i)}(a)$ was weighted with $w(N_T/2)$.

In the following, an exemplary application of the proposed approach is presented for slot milling processes using different inclination angles of the tool. Each slot was machined using a fixed width of cut and an increasing depth of cut along the cutting path, starting at $a_{p,\text{start}} = 0$ mm. The depth of cut at the end of each slot was $a_{p,\text{end}} = 0.8$ mm. For the mother wavelet, the complex Morlet wavelet [5]

$$\Psi(\eta) = \pi^{-\frac{1}{4}} e^{i\omega_0\eta} e^{-\frac{\eta^2}{2}} \quad (8)$$

was used and $\tau \in [2, 2.5]$. A number of $N_T = 10$ neighboring frequencies were incorporated for each of the three dominant natural frequencies of each FRF. Figure 1 visualizes the results of the wavelet-based identification of stability limits. The three dominant natural frequencies were 1438 Hz, 1526 Hz and 1569 Hz for an inclination angle of $B = -30^\circ$ and 1401 Hz, 1521 Hz and 1591 Hz for $B = -45^\circ$. It can be seen that the identified stability limits correspond to visibly high intensities of the wavelet transform. A limit of $a_{p,\text{crit}} = 0.28$ mm was identified for an inclination angle of $B = -30^\circ$. For $B = -45^\circ$, a limit of $a_{p,\text{crit}} = 0.57$ mm was estimated.

3 Further work

Further research activities could include a combination of different sensor technologies to improve the detection accuracy without the necessity of adjusting τ for each process configuration. In this context, also different mother wavelets have to be investigated, since the accuracy of the frequency information, which results from an application of the wavelet transform, is highly dependent on how similar the mother wavelet is to the investigated signal patterns. In addition, resulting stability limits could be used for further classification learning tasks in order to realize a process monitoring system based on real-time predictions of upcoming unstable process behaviors.

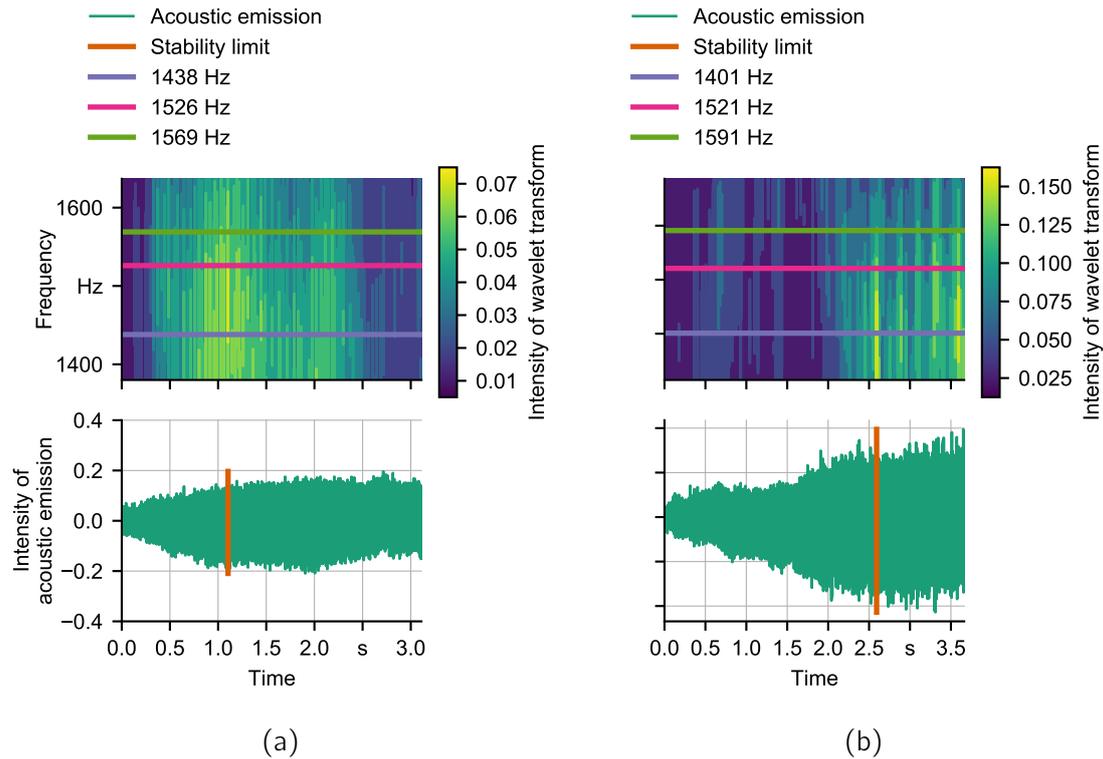


Figure 1: Identification of stability based on acoustic emission signals for inclination angles of (a) -30° and (b) -45° .

References

- [1] P. Wiederkehr and T. Siebrecht, "Virtual machining: capabilities and challenges of process simulations in the aerospace industry," *Procedia Manufacturing*, vol. 6, no. Supplement C, pp. 80–87, 2016.
- [2] F. Finkeldey, A. Wirtz, T. Merhofe, and P. Wiederkehr, "Learning-based prediction of pose-dependent dynamics," *Journal of Manufacturing and Materials Processing*, vol. 4, no. 3, p. 85, 2020.
- [3] C. Torrence and G. P. Compo, "A practical guide to wavelet analysis," *Bulletin of the American Meteorological Society*, vol. 79, no. 1, pp. 61–78, 1998.
- [4] R. B. Blackman and J. W. Tukey, "The measurement of power spectra from the point of view of communications engineering – part i," *Bell System Technical Journal*, vol. 37, no. 1, pp. 185–282, 1958.
- [5] W. K. Ngui, M. S. Leong, L. M. Hee, and A. M. Abdelrhman, "Wavelet analysis: Mother wavelet selection methods," in *Advances in Manufacturing and Mechanical Engineering*, vol. 393, pp. 953–958, 11 2013.

Dynamic Ensemble Aggregation using Deep Reinforcement Learning for Time-Series Forecasting

Amal SAADALLAH

Lehrstuhl für Künstliche Intelligenz, LS 8

Technische Universität Dortmund

amal.saadallah@tu-dortmund.de

An ensemble model is defined as a collection of several models that addresses the same learning task as the individual models to improve the overall performance. Ensemble combination can be coped via meta-learning. We exploit the paradigm of Reinforcement Learning (RL) to learn a combination policy as a meta-learning technique. Some approaches address the opposite task of Meta-RL, which denotes meta-learning on RL problems. However, to the best of our knowledge, we are the first to employ RL for ensemble aggregation that can be used for adaptive re-weighting of ensembles.

1 Problem Formulation

A univariate time-series X_t is a temporal sequence of values up to time t , $X_t = \{x_1, x_2, \dots, x_t\}$, where x_i represents the value of X at time i . Let $\mathbb{M} = \{f_1, f_2, \dots, f_m\}$ be a pool of m base models trained to approximate a true unknown function f that generated X_t . An ensemble model \bar{f} of \mathbb{M} at a future data-point $t + j$ ($j \geq 1$) can be formally expressed as a convex combination of predictions of the individual base models in \mathbb{M} : $\bar{f}(x_{t+j}) = \sum_{i=1}^m w_i^{t+j} f_i(x_{t+j})$, where $w_i^{t+j}, i \in \{1, \dots, m\}$, determine the weights of the ensemble for the time $t + j$. For notational simplicity, assume in the following that our objective is to predict x_{t+1} (i.e., $j = 1$) without loss of generality. We thus aim to find a set of weights for the ensemble that minimizes the expected prediction error for the next forecast: $\operatorname{argmin}_{w_i^{t+1}, \forall i} \mathbb{E}[(f(x_{t+1}) - \bar{f}(x_{t+1}))^2 | X_t]$, s.t. $w_i^{t+1} \geq 0, \forall i \in \{1, \dots, m\}, \sum_{i=1}^m w_i^{t+1} = 1$.

2 Methodology

2.1 The MDP Framework

An RL problem is mathematically formulated via a Markov Decision Process (MDP) [?]. An MDP is defined by a five tuple $(\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{P}, \gamma)$, in which \mathcal{S} are the states, \mathcal{A} the actions, $\mathcal{R} : \mathcal{S} \rightarrow \mathbb{R}$ is the reward function, $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ is the transition function, and $\gamma \in [0, 1)$ is the discount factor. The goal of an MDP is to learn a policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ that maximizes the total obtained reward, and decides about what actions to take at every state. In this section, we characterize our meta-learning task for finding the optimal weights of ensembles for time-series forecasting in an MDP framework.

Actions. An action is interpreted as a decision made at time t to be executed at $t + 1$. In our setting, a decision at every time-step is to determine the set of weights for the ensemble that minimizes the expected prediction error for the next forecast. Therefore, the action $\mathbf{a}_t \in \mathcal{A}$ is defined as an m -dimensional vector, $\mathbf{a}_t = (a_{t,1}, a_{t,2}, \dots, a_{t,m})^\top$, taken at time t that corresponds to the vector of ensemble weights $\mathbf{w}_{t+1} = (w_1^{t+1}, w_2^{t+1}, \dots, w_m^{t+1})^\top$. These weights are attributed to each of the m single models in the ensemble to predict the future value of x_{t+1} . This definition of actions leads to a continuous m -dimensional action space.

States. Let ω be a provided window size of a validation set X^ω that corresponds to the previous ω values of the time-series until time t , $X^\omega = \{x_{t-\omega+1}, x_{t-\omega+2}, \dots, x_t\}$. We consider state s_t to be the current window of time-series that is used for forecasting the next value, i.e., x_{t+1} . However, the next state $s_{t+1} \in \mathcal{S}$ should be devised in a way that it reflects the result of a taken action \mathbf{a}_t at state s_t . Therefore, we take the output of the ensemble \bar{f} on X^ω as a state, instead of X^ω itself, since it reflects the result of the internal combination policy (i.e., \mathbf{a}_t), and also, captures the characteristic traits of X^ω . That means, $s_t = \{\hat{x}_{t-\omega+1}, \hat{x}_{t-\omega+2}, \dots, \hat{x}_t\}$, where \hat{x}_i is the output of ensemble at time t and using the weight vector \mathbf{a}_t , $\hat{x}_i = \bar{f}(x_i)$.

Transition function. The transition function is deterministic in our problem as selecting an action only leads to one possible next state: $\mathcal{P}(s_{t+1}|s_t, \mathbf{a}_t) = 1$ when $s_{t+1} = \{\hat{x}_{t-\omega+2}, \hat{x}_{t-\omega+3}, \dots, \hat{x}_{t+1}\}$ and is zero otherwise.

Reward function. We further determine the reward of taking an action \mathbf{a}_t in state s_t as a function $\mathcal{R}(s_t, \mathbf{a}_t)$, which for brevity in notation we denote by r_t .

Assume that the set of m models in \mathbb{M} together with the ensemble model form a total of $m + 1$ models. Subsequently, we compile a ranked list of all the models (including the ensemble) using their corresponding forecasting error, in which, ρ^{f_i} indicates the rank of model f_i . The lower the rank is in the obtained ranking, the more accurate the model is (i.e., rank 1 means the model performs the best). The reward r_t is thus defined by

$r_t = m + 1 - \rho^{\bar{f}}$, where $\rho^{\bar{f}}$ is the rank of the ensemble among all the models, and the lower it is, the more accurate the ensemble is and the higher the reward value will be.

2.2 Learning the Combination Policy

Once the meta-learning task is phrased in an MDP framework, the policy π is learned in favor of maximizing the reward which correlates to an inverse measure of the ensemble performance. We employ the deep actor-critic approach presented in [2] to learn an optimal combination policy in a continuous action space. This approach is selected since it is well-suited for both continuous and high-dimensional action and state spaces (for large m and ω). In this architecture, the actor is accountable for selecting an action given the current state, and the critic estimates a value function which provides adequate evaluation for the actor. Both parts are represented by (deep) neural networks that can be optimized by gradient descent-based methods. As a result, the actor and the critic networks are called the *policy network* and the *value network*, respectively. The value network predicts the value of an action \mathbf{a}_t in state s_t via $Q(s_t, \mathbf{a}_t|\phi)$, where ϕ is the parameter vector of the value network; see [2]: Eq. (3)-(5). On the other hand, the policy network learns a policy $\pi(s_t|\theta)$ which yields a deterministic policy in state s_t given the network parameters θ ; see [2]: Eq. (6). During the learning, the actor takes the gradients derived from the policy gradient theorem and adjusts the policy parameters θ , and the critic network estimates the approximate value function for the current policy π via Bellman equation.

After the policy network $\pi(s|\theta)$ is learned, we apply the model for predicting the weights of the ensemble (i.e., actions) that will be used for predicting the future values of time-series in an online manner. Let state s be X^ω , the predicted weights via a are used to predict x_{t+1} . Afterwards, the ω -length vector of time-series X^ω (i.e., the state s) is moving forward by one value. That means, the oldest value is removed and the predicted value \hat{x}_{t+1}^E is added to the current window. The new state s' and $\pi(s|\theta)$ are employed to predict the weights of the ensemble to forecast the next value of the time-series. The procedure is repeated until N_f desired values of the time-series are forecasted.

3 Experiments

The results are evaluated using the RMSE and ranks of our method compared to s.o.a methods for ensemble aggregation are computed. A rank of 1 means the model was the best performing on all time series.

The approaches that are based on combining individual forecasters, e.g., SE, SWE, etc., and common ensemble methods, such as RF, GBM, Stacked LSTM and Stacking, show

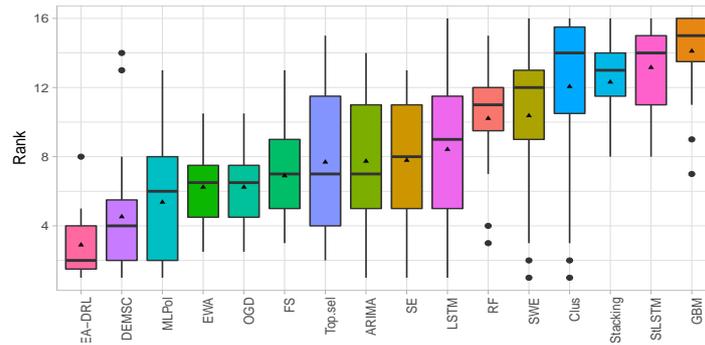


Figure 1: Distribution of the ranks of ensemble methods across the different time series

inferior performance compared to EA-DRL. ARIMA and LSTM, state-of-the-art methods for forecasting, have a considerable difference in the average rank as well. The two competitive approaches to our method are DEMSC and MLPOL that perform well in the pairwise comparison, nevertheless, both attain a higher average rank in Figure 1. DEMSC is based on real-time update of meta-learning strategy behind, while EA-DRL is devised offline and only predictions are computed in real time. Looking at the distribution of ranks in Figure 1, we see that EA-DRL is within the range of the first 4 ranks and has a lower variance in the rank compared to DEMSC and MLPOL. In addition, MLPOL have a relatively higher rank compared to DEMSC and EA-DRL. More details about the method and the evaluation can be found in [1].

4 Future work

Our approach relies on a policy which is learned offline. One potential future research direction would be to investigate the impact of an online update of the policy, for instance in a periodic manner, or in an informed fashion following a drift-detection mechanism in the data and/or the performance of the ensemble.

References

- [1] A. Saadallah, M. Tavakol, K. Morik, "An Actor-Critic Ensemble Aggregation Model for Time-Series Forecasting," 2021 IEEE 37th International Conference on Data Engineering (ICDE).
- [2] Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., ... Wierstra, D. (2015). Continuous control with deep reinforcement learning. arXiv preprint arXiv:1509.02971.



Subproject B4
Analysis and Communication for Dynamic
Traffic Prognosis

Thomas Liebig Michael Schreckenberg
Christian Wietfeld

Validation of CAM Position Information in C-V2X Communication

Fabian Eckermann

Lehrstuhl für Kommunikationsnetze

Technische Universität Dortmund

fabian.eckermann@tu-dortmund.de

Cellular-Vehicle-to-Everything (C-V2X) communication is an essential component of future automated traffic systems. These connected vehicles exchange location data obtained by on-board satellite receivers. Such information is vital for cooperative applications, so its validation is an important safety feature.

In this report, a method to repurpose available information from C-V2X to estimate the distance between sender and receiver and therefore validate Cooperative Awareness Messages (CAM) position information is presented.

1 Range Estimation and Validation Using C-V2X

The automotive sector has standardized cooperative awareness messages (CAM) to exchange Global Navigation Satellite System (GNSS) positions and further status information such as speed or heading. The C-V2X specification was introduced in release 14 of the fourth generation of mobile communication systems (4G), standardized by 3GPP. It introduced the sidelink, opposed to up- and downlink, as a direct link between C-V2X User Equipments (UEs), that is independent from cellular network coverage by using a distributed channel access scheme. C-V2X has been designed for periodic ad-hoc broadcasting of status messages, to enable cooperative awareness in intelligent transportation systems (ITS) for road safety and autonomous driving. Furthermore, the number of reference signals utilized for signal detection and channel estimation is increased, addressing high relative speeds of agents and decentralized time synchronization.

Timing of transmissions in C-V2X is organized in a fixed subframe grid of 1 ms. The

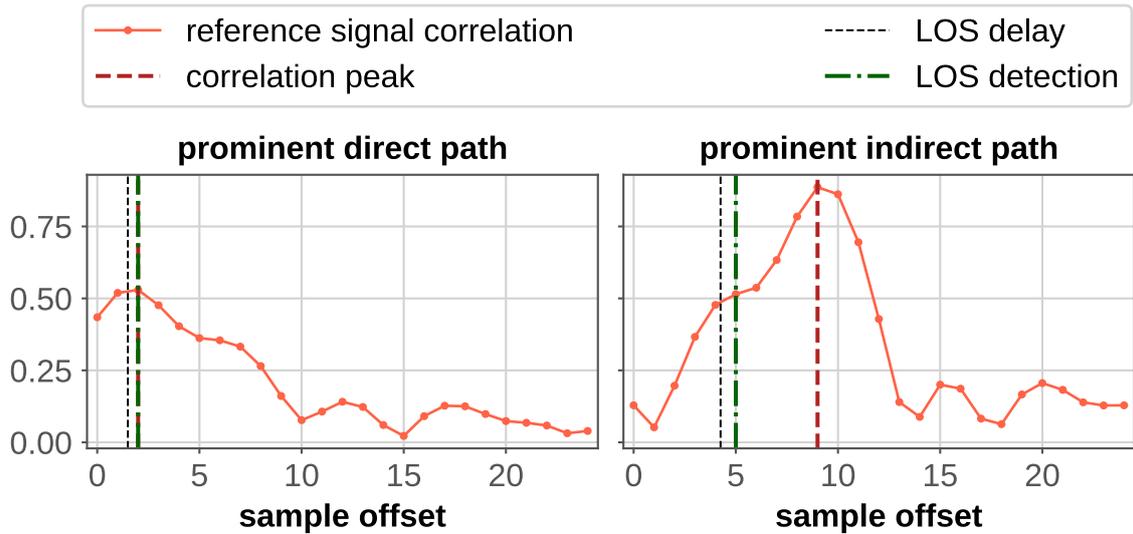


Figure 1: Examples of the cross correlation of expected and received reference signals. Due to multipath fading the prominent path does not always correspond to the LOS, which is vital for position estimation.

time between the local start of a subframe and the start of a received signal, the sample offset, can be used as an estimate for the signal propagation time that correlates to the distance between sender and transmitter. The accuracy of this estimate depends on the synchronization error and the signal detection method. At the receiver, pre-computed, expected Demodulation Reference Signals (DMRS) are usually cross-correlated with received signals to detect the start of a specific transmission. The maximum of this correlation indicates the instance of the reference signal that arrived with the highest power, increasing the chance of successful decoding. It is important to note that in a multipath environment, the strongest path might not always be the direct Line-of-Sight (LOS), making a correlation peak less than ideal for range estimation, see Figure 1.

2 Validation Results

When a message has been successfully decoded, the correlation function is analyzed to find a LOS delay, by extracting local maxima and saddle points. A comparison of the basic peak correlation and the extended LOS detection is shown by Figure 1. The trilateration of the sender's position requires knowledge of other receiver's positions and estimated ranges. Since positions are already exchanged through CAM, a straightforward approach could be to insert a list of past measured ranges into a CAM broadcast. The difference between the sender's position and the CAM position is the trilateration error. The analyzed trilateration approach starts with an initial linear approximation [2], followed by a Gauss-Newton optimization. Robustness against timing errors caused by multipath

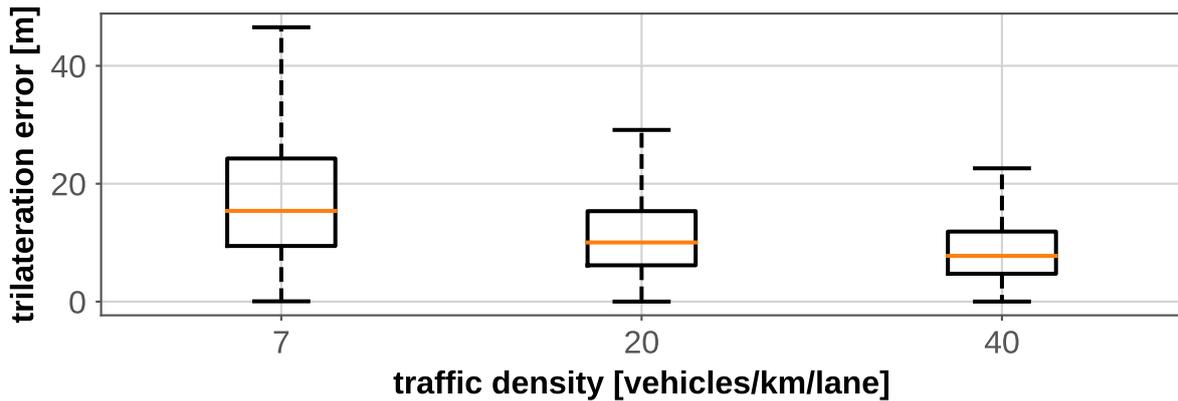


Figure 2: Trilateration error for different traffic densities.

propagation and sampling granularity is increased by an iterative outlier removal scheme as presented in [3]. Range measurements disagreeing with the current estimate are iteratively removed until all residuals lie within a one sample range. The performance of both approaches depends on the number of receivers whose estimates are combined. Hence three vehicular highway scenarios with different traffic densities are evaluated: free, stable and congested (7, 20, 40 vehicles/km/lane) based on [1].

Figure 2 depicts the trilateration error for different traffic densities. A higher traffic density decreases the maximum trilateration error from 50 m to 22 m. For a stable traffic density (20 vehicles/km/lane) a maximum trilateration error of 30 m and a median error of 10 m are achieved. For the classification of erroneous positions, sensitivity and specificity are used as the primary statistical performance measures. Sensitivity, also known as the true positive rate describes the proportion of erroneous positions that are correctly identified. The specificity (true negative rate) defines the proportion of valid positions, that are classified as such. These metrics are related to the position tolerance i.e. the difference between a CAM position and an estimated position above which the CAM position is classified as invalid.

The sensitivity and specificity of the triangulation based validation approach is shown by Figure 3. Invalid positions can be classified with approximately 99% confidence across all evaluated tolerances. The sensitivity slightly decreases when tolerance increases. This can be explained by an increase of the valid area and together with its circumference, making it likelier for simulated positions to lie close to this decision boundary. These positions are harder to classify and slightly more of them are falsely classified as valid. The specificity of the classification highly correlates to the distribution of the trilateration error. The saturation of the specificity is reached if the position tolerance matches the maximum of the trilateration error (shown in Figure 2). Starting from this point, all correct positions can be classified as such and the classification is almost perfect. For a traffic density of 7 vehicles/km/lane a position tolerance of 50 m is necessary to reach a sensitivity of 99 % and a specificity of around 100 %. For a density of 40 vehicles/km/lane even better results are achieved for half the position tolerance.

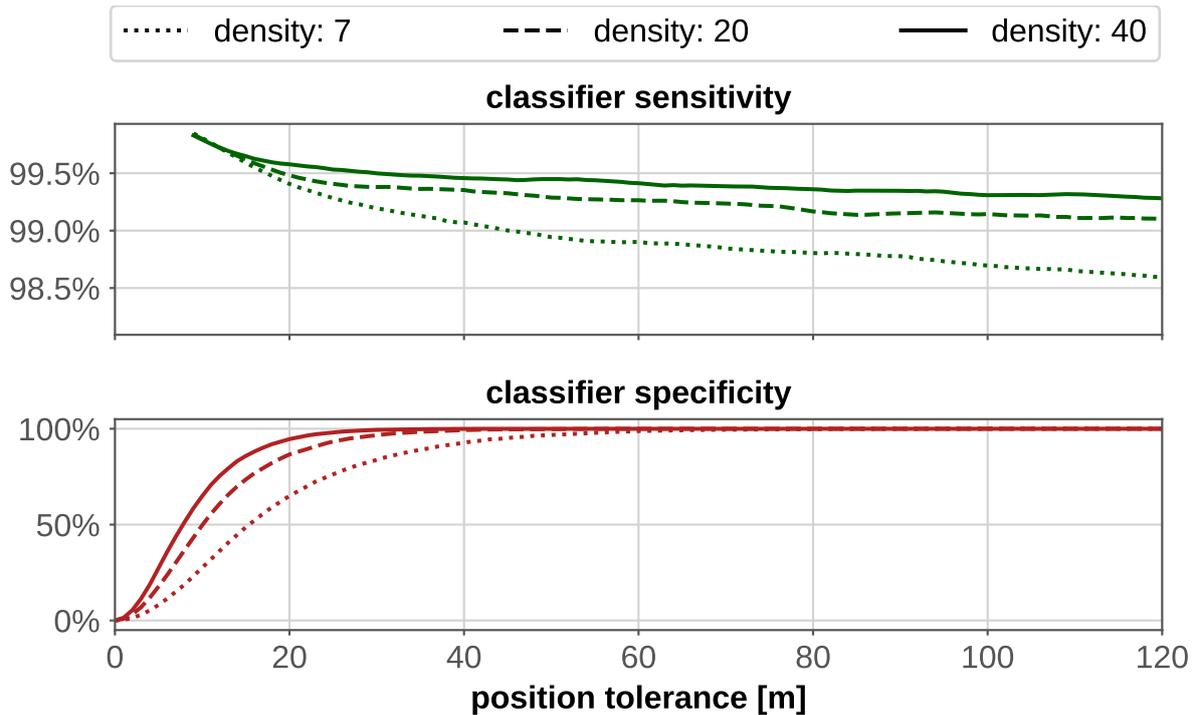


Figure 3: Classification performance using trilateration on offsets obtained by LOS detection for different traffic densities (in [vehicles/km/lane]).

3 Conclusion

In this report, it is presented how existing measurements from the C-V2X physical layer can be used to validate positioning data. Using a trilateration algorithm, and a position tolerance of 22-50 m, errors are classified with a confidence of 99% even in low traffic densities while at the same time being robust against false positives.

References

- [1] Transportation Research Board. *Highway Capacity Manual*. Transportation Research Board, National Research Council, 2000.
- [2] Yue Wang. Linear least squares localization in sensor networks. *EURASIP Journal on Wireless Communications and Networking*, 2015, 2015.
- [3] Andrew Martin Robert Ward. *Sensor driven computing*. PhD thesis, University of Cambridge, 1998.

Leveraging 5G Millimeter Wave Beam Alignment to Provide a Positioning Service in Vehicular Networks

Karsten Heimann

Lehrstuhl für Kommunikationsnetze
Technische Universität Dortmund
karsten.heimann@tu-dortmund.de

Several vehicular applications necessitate a positioning service. Besides well-known, dedicated infrastructures like a global navigation satellite system (GNSS), mobile radio networks are also able to determine the network subscriber's position based on radio link metrics like the receive power or the propagation delay. Introducing electrically steerable, directional communications at the millimeter wave (mmWave) spectrum, high gain antennas allow for a targeted signal transmission. In doing so, the adjustable main lobe direction embodies a promising new metric for a positioning service in 5G mmWave networks based on direction finding (DF). In a cross-bearing manner, the DFs to a mobile device from base stations with known poses are merged to estimate the device's position in addition to the usual mobile communication. The experimental evaluation based on our extensive laboratory system proves that information on the base station's configured directivity bear the potential to provide a precise positioning service to mobile devices.

1 Positioning in Vehicular Communications

The millimeter wave (mmWave) mobile networks of the fifth generation of mobile communication (5G) are designated to satisfy the growing demand for high throughput. Due to the inevitable need for directional communication at this frequency spectrum, this direction knowledge qualifies for being utilized as RF-based positioning service. Conveniently, especially vehicular applications, which may benefit from an increased network

performance, often rely on positioning services e.g. within the road traffic for navigation or remote/autonomous driving. While designated positioning systems like global navigation satellite system (GNSS) or ultra-wideband (UWB) require the deployment of infrastructure like satellites or anchor nodes, positioning as a feature of wireless communications (like wi-fi and cellular) is already valued as a compromise between efficiency and precision nowadays. In terms of an experimental proof-of-concept evaluation, the feasibility of a position estimate based on mmWave beam alignment is analyzed in [1].

2 Beam Alignment for Position Estimation

As discussed in [2], the mmWave mobile network access of a subscriber in motion may be provided and maintained by means of a signal quality-based beam alignment. This means the base station antenna's main lobe or *beam* is electrically steered towards the designated mobile device based on signal quality measurements. According to that, this beam pointing direction of the base station can be leveraged as direction finding (DF).

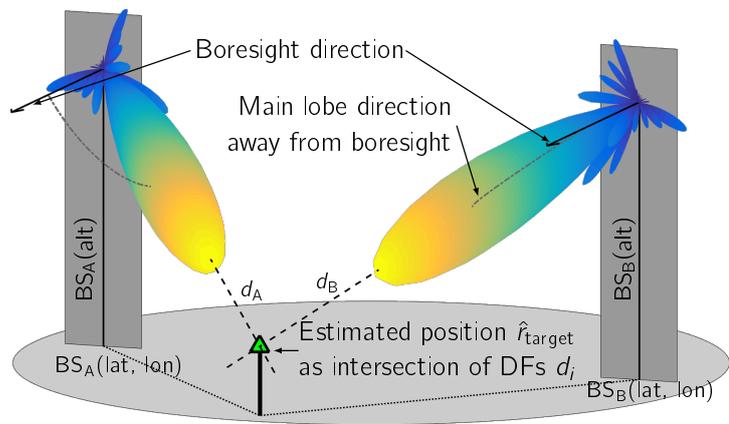


Figure 1: Concept of determining the position estimate as intersection of given beam directions.

A DF of a base station may even be available independently of data transmissions, since the network needs to broadcast its presence for both attached and appearing devices to update and initiate their connections, respectively. However, during a data transmission, a more pronounced directivity with a high gain and thus a narrow beam width is required due to its high throughput character. This not only leads to a more challenging beam alignment, but also provides fine-grained information of the mobile device direction. On the contrary, the steady broadcast signals like from the synchronization signal/physical broadcast channel (SS/PBCH) may use wide beams to efficiently serve the base station coverage area, so that only a coarse-grained DF may be derived from this signal. Additionally, the base stations' positions and orientations need to be known as reference for subsequent mobile device positioning intentions.

Although a single base station might suffice for a position estimation, if its beam direction and range (or pseudo-range) was measured, this approach would rather represent a hybrid positioning method by also taking the propagation delay into account. In [1], nevertheless, a solely DF-based approach is analyzed by determining suited beam pointing directions of at least two base stations as DFs and estimating the mobile device position at their intersection as depicted in Figure 1.

3 Experimental Evaluation

The methodology of the conducted laboratory experiment is depicted in Figure 3. While beam tracking approaches like in [2] would reduce the required time for DF at the cost of the achievable resolution, exhaustive sweeps exploiting the antenna’s whole coverable angular domain are performed with a step size of 1° in azimuthal and elevation plane to get a comprehensive overview on the beam direction dependent signal quality. This leads to heatmap representations like depicted in the second block of the figure.

Based on these measurements, the uplink angle of arrival (UL–AoA) or downlink angle of departure (DL–AoD) at the base station or the mobile station, respectively, may be determined by analyzing the signal quality distribution and assess a proper DF. Since a positioning service tend to be offered by the network, the UL–AoA variant is measured in this work, but a self–positioning at the mobile device may be taken into consideration as well, if the base station reference poses are known at the mobile device.

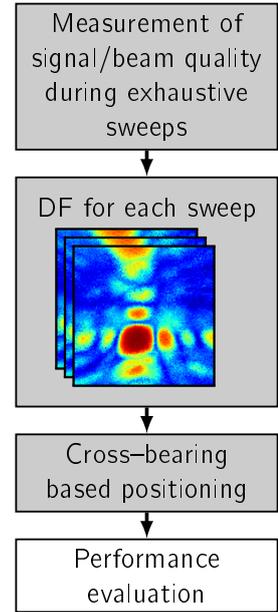


Figure 3: Methodology of the experimental evaluation.

However, from the heatmap illustration, the DF is derived as centroid of the region containing the best signal quality (red area). The measurements are repeated one hundred times per base station and mobile station position using two base stations and two reference positions of the mobile device. For each evaluated mobile device position, each permutation of DF pairs of the respective base stations is used for distinct positioning runs by computing a least squares approximation of the intersection of the two bearings. Finally, the three dimensional Euclidean positioning error is evaluated as depicted in Figure 2.

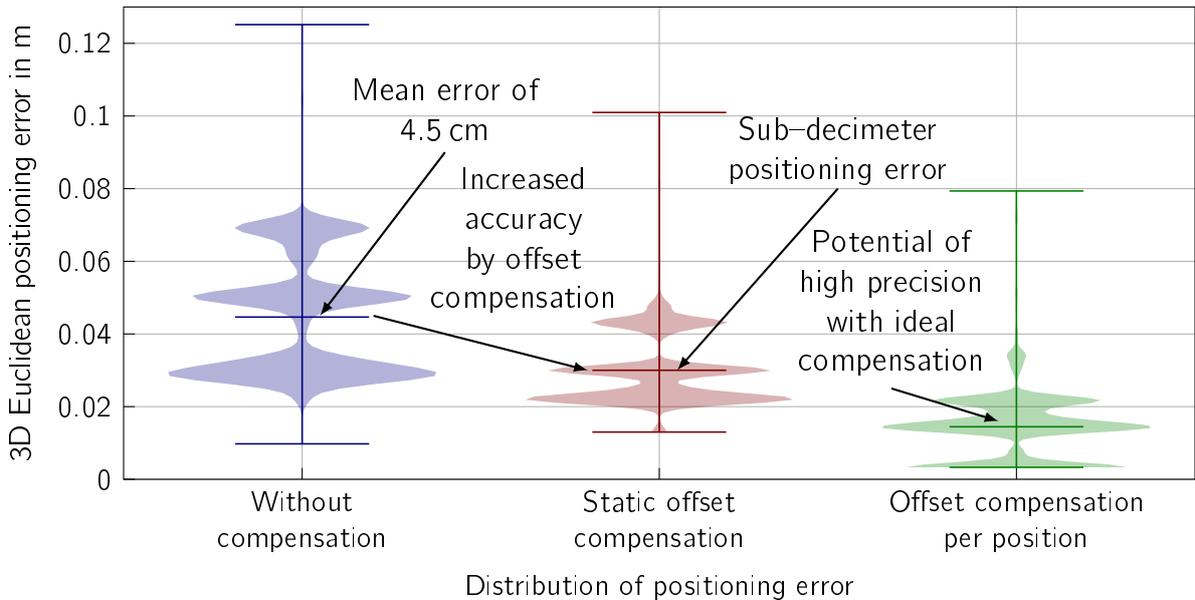


Figure 2: Statistical evaluation of the positioning performance. The applied compensation accounts for the DFs’ mean angular offset in total or per mobile device position.

The mean error of about 4.5 cm as well as the variance can be reduced in post-processing by applying a static offset to the DFs. This is achieved by subtracting the angular mean deviation to the reference direction per base station. A further enhanced precision is reached by an individual compensation per base station and position. At this laboratory evaluation, a mean position error of 1.4 cm with a standard deviation of 0.9 cm constitutes a reasonable and promising performance in terms of positioning accuracy and precision.

4 Conclusion and Further Research

Positioning services are being offered as an additional feature of current mobile networks. With 5G utilizing the mmWave spectrum, the performance of a communication network based positioning can even be improved by base stations applying beamforming by means of phased array antennas. In doing so, the antenna's directivity is electrically steerable forming narrow beams towards particular mobile users. This allows for information about the directions of the mobile devices, while the base station poses ought to be known as well. The presented UL-AoA-based approach achieves a promising accuracy and precision especially with ideal compensation and thus points out the potential of a mmWave positioning service leveraging directional knowledge of the antennas' beam alignment.

While this study relies on a line-of-sight, so-called reconfigurable intelligent surfaces (RISs) are believed to facilitate an efficient utilization of the mmWave spectrum even in obstructed areas. In our initial simulation work [3], their deployment in an example vehicular scenario leads to a comprehensive network coverage thanks to purposefully adjusted reflection paths. Since the angle of reflection is again related to a bearing/DF, the presented positioning method might be extended to non-line-of-sight conditions in RIS-enhanced future mobile networks.

References

- [1] K. Heimann, J. Tiemann, S. Böcker, and C. Wietfeld, "Cross-bearing based positioning as a feature of 5G millimeter wave beam alignment," in *IEEE 91st Vehicular Technology Conference (VTC-Spring)*, May 2020.
- [2] K. Heimann, J. Tiemann, D. Yolchyan, and C. Wietfeld, "Experimental 5G mmWave beam tracking testbed for evaluation of vehicular communications," in *IEEE 2nd 5G World Forum (5GWF) (WF-5G'19)*, Sept. 2019.
- [3] K. Heimann, A. Marsch, B. Sliwa, and C. Wietfeld, "Reflecting surfaces for beyond line-of-sight coverage in millimeter wave vehicular networks," in *IEEE Vehicular Networking Conference (VNC)*, Dec. 2020.

Kinematic phases in one-dimensional, heterogeneous procession traffic with open boundary conditions on the example of the Cologne Rose Monday Parade

Petros Polichronidis
Physics of Transport and Traffic
University of Duisburg-Essen
petros.polichronidis@uni-due.de

In the context of this paper, the analyses of the kinematics of the Cologne Rose Monday Parade are extended. The paradoxical velocity conditions observed there could be explained by a constant preceding but velocity restricting participant in the lead [1]. In the now extended analyses of the data set from 2014, it could be shown that the exit of the parade head manifests itself in an acceleration wave that propagates through the entire parade. The following pageant part changes into a different kinematic phase by adopting higher velocities. On the basis of different velocity ranges these kinematic phases could be characterized (stagnant, leisurely, rapid). They can be of temporary or persistent nature, but depend on the presence and pre-set velocity of the parade leader and the nature of the boundary condition of the parade itself. The results extend the understanding of the movement of processions and can be used to extend existing models describing procession traffic and their predictive capabilities.

1 Introduction

GPS data sets of the Cologne Rose Monday Parade provided the first scientific insight into the kinematics of a carnival procession [1]. The starting point of these investigations was a paradoxical observation in the kinematic behavior of the Cologne Rose Monday procession. It was observed that despite considerable standstills and jam waves, the travel times of the participants shortened and depended primarily on the position in the pageant. For example, the last participant in 2014 had a shorter parading duration than the first participant by about 1 hour. This could be explained by the speed behavior of the pageant participant in the lead. It is set to about 2 km/h in the case of the Cologne Rose Monday procession. Since this is much lower than typical pedestrian speeds (4 km/h - 5 km/h), it acts as an artificial velocity restriction for the following procession. On the one hand, this is to minimize the creation of gaps between participants, and on the other hand, to take into account the physical condition and velocity capabilities of young participants. Further analyses showed that the movement of the pageant is characterized by different traffic states or kinematic phases, which occur primarily due to the speed-limiting pageant leader and the open boundary conditions. Based on the GPS data set of the Cologne Rose Monday procession from 2014 the phases are explained in the following.

2 Kinematic Phases in the Cologne Rose Monday Parade 2014

The individual velocity and distance behavior of the participants and the slow-moving pageant leader acting as a speed barrier mean that the pageant can exhibit different kinematic *phases* during the parade. These can be classified by three velocity ranges (stagnant, leisurely, rapid) and will be explained below.

In general, a distinction can be made between the *stagnant phase* and the two other phases in motion. The stagnant phase is equivalent to a standstill, similar to a traffic jam ("stop-and-go traffic") and is characterized by standstill times of different lengths. Kinematic phases in motion are characterized by different velocity ranges (leisurely and rapid). These can originate from a kind of stop-and-go traffic, comparable to a synchronized traffic phase, but also from a quasi-free traffic flow.

The *leisurely phase* is defined by the average speed of the pageant leader, since its speed is often smaller than that of the typical pedestrian movement. As long as there are no large gaps between the following participants, the participants are within the range of influence of the pageant leader and can only adopt speeds that are equal to or smaller than that of the pageant leader. Thus, this part of the procession is accordingly in the leisurely phase. But even beyond the presence of the pageant leader, dawdling

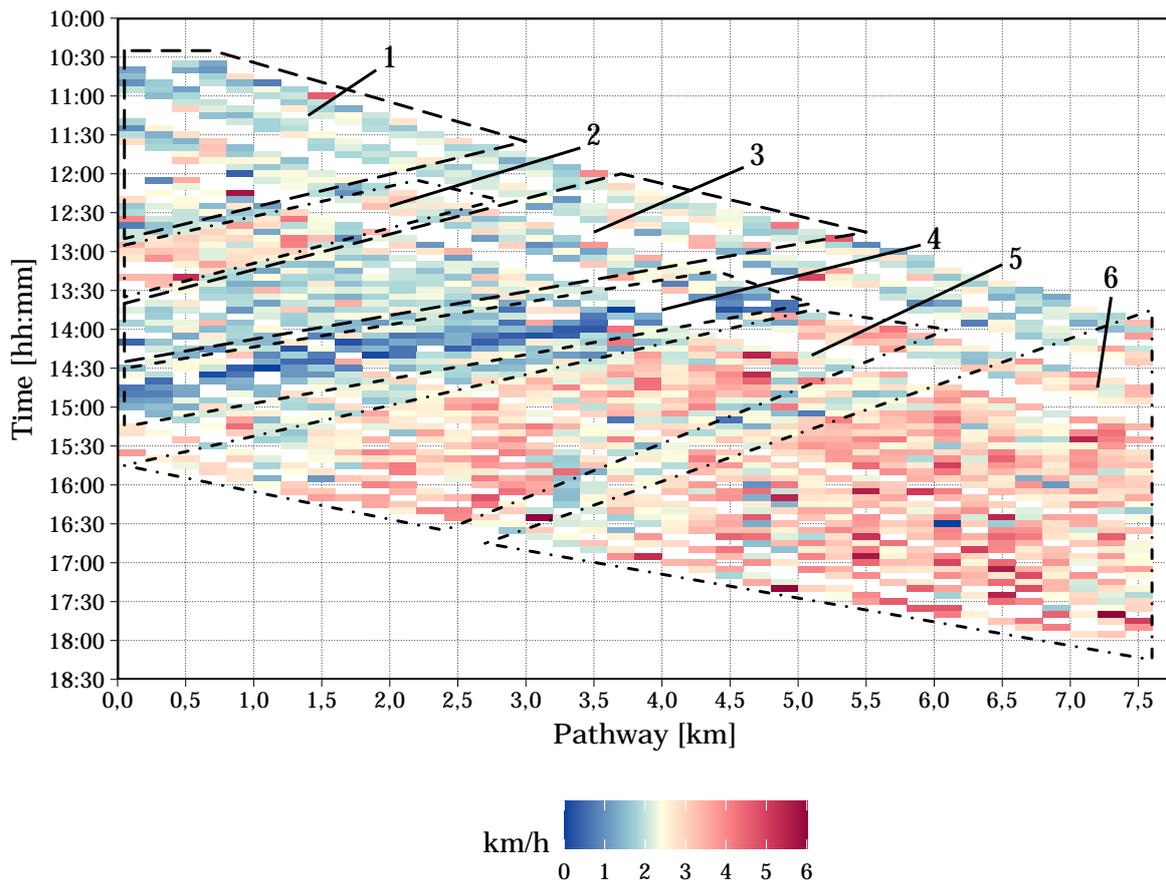


Figure 1: Quantized path-time diagram of the Cologne Rose Monday Parade from 2014: Within spatial and temporal intervals (200 m and 5 min, respectively), the arithmetic mean of all occurring speeds is calculated and color-coded. The coloring corresponds to stagnant (blue), leisurely (yellow) and rapid phases (red).

participants can behave in such a way that the procession gets into the leisurely phase in various sections.

A pageant section in the *rapid phase* is defined by a speed range that lies above the average speed of the pageant leader. A rapid phase can be of short or long duration. If the speed-limiting parade leader exits the pathway, the subsequent pageant as a whole may transition to a *prolonged rapid phase*. *Temporary rapid phases* occur especially during the closing of gaps created by excessively high speeds after a stagnant phase caused by various reasons. Due to the distance behavior of the following participants, this can also cause a certain part of the following pageant to transition into a rapid phase. Thus, different kinematic phases may coexist in different parts of the procession. In the further course, the different kinematic phases are illustrated in quantized path-time diagrams. The time is divided into 5 minute intervals and the pathway into 200 m sections, within which the arithmetic mean of all velocities occurring there is formed and color-coded according

to the three phases. The coloring corresponds to stagnant (blue), leisurely (yellow) and rapid phases (red). In this representation, the spatiotemporal expression of the different phases can be illustrated.

In figure 1, 6 areas are marked in which the different traffic states of the pageant can be clearly seen. In area 1 and area 3, the participants mainly show velocities that are in the range of the participant in the lead of the pageant. It can therefore be assumed that the participants are all holding on to each other. Occasional red intervals originate from relatively short dawdling or acceleration processes. In area 2, the participants show similar behavior. The foremost participant in this area dawdles or allows a large gap to develop, which is then closed excessively fast. Since the following participants are supposed to not let gaps occur, they all exhibit excessively high velocities. The part of the pageant in area 3 is in a temporary rapid phase. Area 4 is characterized by an extended wave of congestion with duration of standstill of about a quarter of an hour. By definition, the pageant part in area 4 is in the stagnant phase. As the participants not affected by the jam front continue to move, a considerable gap is created between them and the frontmost participant affected by the standstill. This gap is made up at excessively high speeds, so that the pageant part in area 6 is in a temporary rapid phase until it catches up with the part ahead. Between area 5 and area 6, the participants closed up the gap and are thus in the leisurely phase. In area 6, the pageant enters the continuous rapid phase, which is created by the absence of the speed-limiting participant in the lead. The following pageant is not subject to any speed restriction and an acceleration wave spreads through the entire pageant in the course of which velocities are adopted that are in the range of typical pedestrian speeds.

3 Outlook

The described kinematic phases should be typical for the kind of boundary condition of the Rose Monday Parade in Cologne. This is a phenomenon critical of the presence of a speed-limiting pageant leader. It would therefore be interesting to investigate the kinematics of other pageants with open boundary conditions for universal behavior. This would improve existing models for describing procession traffic and their predictive capabilities.

References

- [1] Petros Polichronidis, Dominik Wegerle, Alexander Dieper, and Michael Schreckenberg. Traffic dynamics of carnival processions. *EPL (Europhysics Letters)*, 121(6):68003, 2018.

Lightweight Machine Learning for IoT Systems with Resource Limitations

Benjamin Sliwa

Lehrstuhl für Kommunikationsnetze

Technische Universität Dortmund

benjamin.sliwa@tu-dortmund.de

In this report, an overview about the novel Lightweight Machine learning for IoT Systems (LIMITS) framework for high-level automation of machine learning tasks is provided. LIMITS implements a hybrid approach which brings together the convenience features of high-level automation with the resource-awareness of low-level programming. For this purpose, a platform-in-the-loop approach is applied which explicitly integrates the resource requirements of the targeted Internet of Things (IoT) platform into the model optimization process.

1 Platform-in-the-Loop Data Analysis with LIMITS

It is expected that future mobile communication systems will heavily rely on predictive Quality of Service (QoS) mechanisms e.g., adaptive data transfer methods that rely on predicting the achievable data rate along vehicular trajectories.

Machine learning has been recognized as a catalyst for developing these mechanisms as it is implicitly able to learn the hidden interdependencies between measurable indicators which are often too complex to bring together in an analytical description. As a consequence, a clear trend of replacing detailed mathematical models with learned representations can be identified. However, there exists a methodological gap between the *insight* and *exploitation* domains. While data analysis is often performed using high-level programming languages (e.g., `python`, `R`, `MATLAB`), the targeted IoT platforms operate on low-level

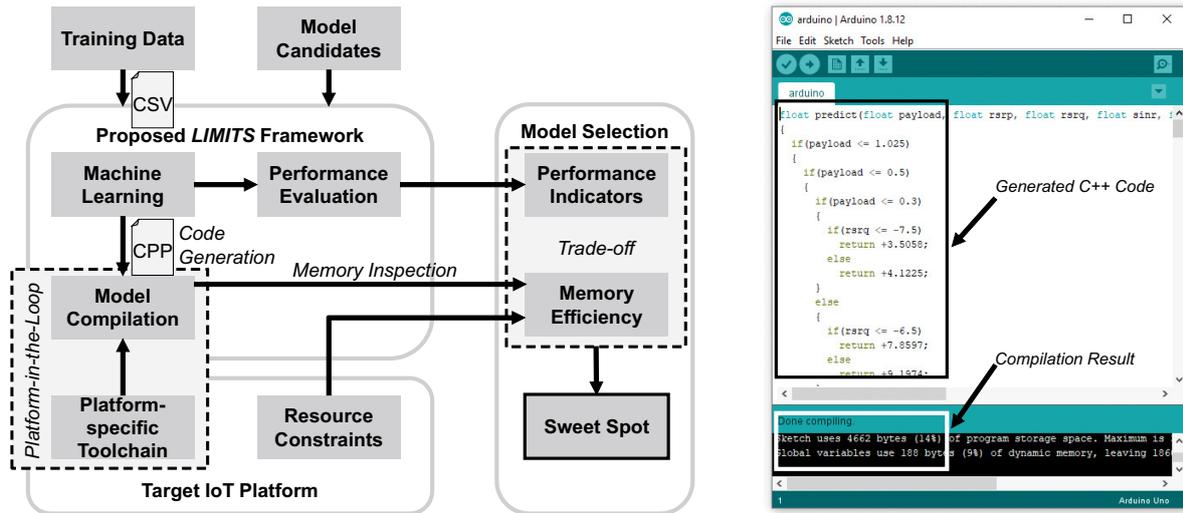


Figure 1: System architecture model of the proposed LIMITS framework and example for platform-specific code generation and compilation.

code such as C/C++. Moreover, the real world devices are impacted by resource constraints such as memory, computation speed, and energy limitations.

The novel LIMITS [3] framework aims to close this gap by acting as a facilitator between the high-level and low-level domains. It represents the methodological distillate of previous attempts to automatize machine learning-based data analysis targeted at resource-constrained IoT systems. Its corresponding system architecture model is shown in Fig. 1. LIMITS relies on intuitive `python` methods to automatize even complex data analysis tasks. In addition, a Command Line Interface (CLI) can be utilized for rapid result analysis as shown in Fig. 2. However, the actual low-level machine learning processes are performed based on validated models of the well-known Waikato Environment for Knowledge Analysis (WEKA) framework. The achieved outputs – the parameters of the trained models – of the latter are then parsed and converted into abstract models which support C/C++ code generation.

In order to optimize machine learning models for highly resource-constrained platforms such as Microcontroller Units (MCUs), a key challenge is the determination of the required platform resources for a trained model. While the memory footprint of certain model classes can be described analytically, these estimations are often *optimistic* as they do not consider platform-dependent aspects such as external libraries for math functions and floating point emulation. For overcoming this issue, LIMITS applies a *platform-in-the-loop* approach which explicitly integrates the targeted hardware platform into the model selection/optimization process. Through compiling the automatically generated C/C++ code of the trained model with the actual compilation toolchain of the IoT platform, the final on-device resource consumption of the model can be precisely assessed.

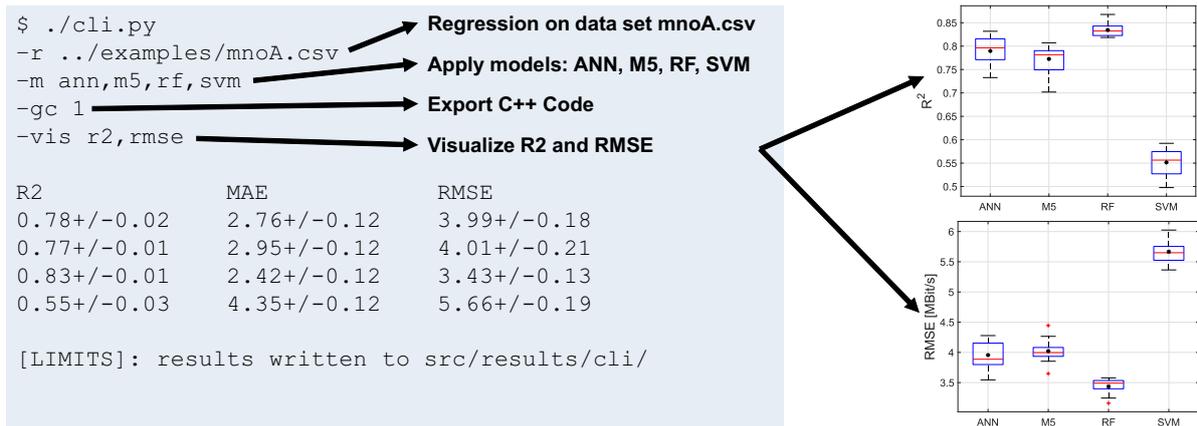


Figure 2: Example for the CLI usage for data analysis, visualization, and code generation.

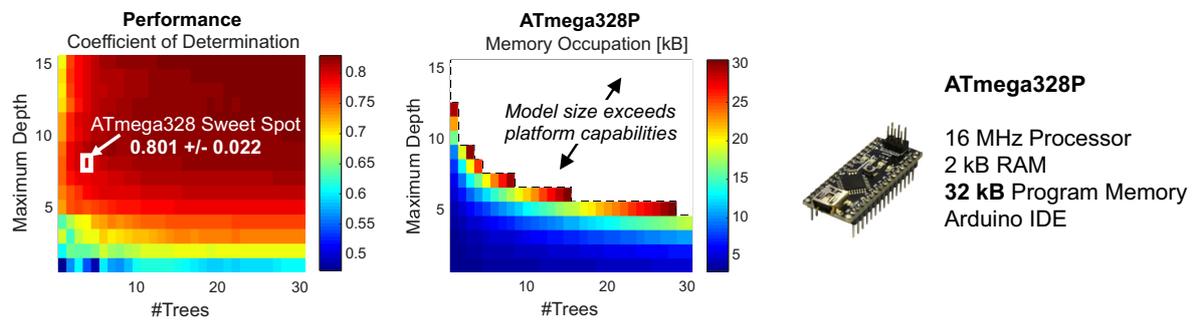


Figure 3: Example for the platform-specific sweet spot model parameterization process.

Moreover, it allows to determine the *sweet spot* model parameterization – as a trade-off between platform resources and model performance – in order to optimize a prediction model based on the given bounds of the hardware platform. Fig. 3 shows an example of such a process. Hereby, different parameter variants for the number of trees and the maximum tree depth of an Random Forest (RF) regression model are explored. While in this case higher parameter values generally lead to a better prediction performance, the memory size of the targeted ATmega328P only allows to implement a subset of the analyzed model variants. However, the fine-grained analysis of the memory requirements allows to select the parameter values such that the maximum achievable prediction performance for the given platform can be utilized.

The current version of the proposed framework supports different model classes such as Artificial Neural Networks (ANNs), M5 Regression Trees (M5s), RFs, and Support Vector Machines (SVMs). The practical usage has been demonstrate in different research works focusing on resource-efficient vehicular data transfer [1] as well as for Data-Driven Network Simulation (DDNS) [2] and radio-based vehicle classification [4].

2 Conclusion and Further Research

LIMITS is a novel high-level framework for automating data analysis tasks and for deploying trained prediction models to resource-constrained real world IoT platforms. In future work, LIMITS will be extended with additional machine learning models. In addition, it might be possible to integrate methods for automatic Worst Case Execution Time (WCET) analysis which would be utilized to anticipate the on-device inference time of trained models.

References

- [1] Benjamin Sliwa, Rick Adam, and Christian Wietfeld. Acting selfish for the good of all: Contextual bandits for resource-efficient transmission of vehicular sensor data. In *Proceedings of the ACM MobiHoc Workshop on Cooperative Data Dissemination in Future Vehicular Networks (D2VNet)*, Online, Oct 2020.
- [2] Benjamin Sliwa, Manuel Patchou, and Christian Wietfeld. The best of both worlds: Hybrid data-driven and model-based vehicular network simulation. In *2020 IEEE Global Communications Conference (GLOBECOM)*, Taipei, Taiwan, Dec 2020.
- [3] Benjamin Sliwa, Nico Piatkowski, and Christian Wietfeld. LIMITS: Lightweight machine learning for IoT systems with resource limitations. In *2020 IEEE International Conference on Communications (ICC)*, Dublin, Ireland, Jun 2020. Best paper award.
- [4] Benjamin Sliwa, Niko Piatkowski, and Christian Wietfeld. The channel as a traffic sensor: Vehicle detection and classification based on radio fingerprinting. *IEEE Internet of Things Journal*, 7(8):7392–7406, 2020.

Heterogeneous vehicle traffic with human driven, automated, and communicating automated vehicles

Tim Vranken
Physik von Transport und Verkehr
Universität Duisburg-Essen
tim.vranken@uni-due.de

This report sums up our work on creating a cellular automata model that is able to simulate heterogeneous traffic where human driven vehicles mix with automated and communicating automated vehicles. The aim of this model was to introduce a cellular automata model that uses a time step length below 1 s without allowing human drivers to react more often than in previous cellular automata models. Furthermore, different behaviours between automated and non automated vehicles besides reaction times were defined and implemented.

1 Introduction

Automated vehicles are expected to improve traffic flow while reducing the risk of accidents, fuel consumption, and travel time [2]. However, the situation in heterogeneous traffic, where automated and human driven vehicles mix, is still unclear and because there is next to no empirical data on it yet [5] a lot of heterogeneous traffic models were developed to simulate and analyse it (see [1] for a overview). Most of the cellular automata models in this regard however are not able to simulate the reduced reaction time of automated vehicles due to their 1 s time step length or differences between the behaviours of automated and non automated vehicles.

The Lee-model [4] has introduced two different states for human drivers that correspond to optimistic driving (where they follow the leading vehicle with less than the necessary

safety distance) and pessimistic driving (where they follow with slightly more than the necessary safety distance). Furthermore, it operates with a finite breaking capability. Later, Pottmeier et al. [7] slightly changed the model so that even through it operates with a finite breaking capability and agents can follow their leader with less than the safety distance, no accidents occur.

This makes it a perfect candidate to adapt for heterogeneous traffic simulations. To this end, we will first create a model based on Pottmeiers version of the Lee-model with a 0.1 second time step length in section 2 before we introduce automated and communicating vehicles in 3.

2 0.1 second time step length model

To change the time step length of the original model to a 0.1 s length, the length of one cell also has to be reduced. Because the acceleration of $1.5 \frac{m}{s^2}$ should be retained, the cell length has to be reduced by the square of the reduction of the time step length. This means that in the new model, one cell equals $\frac{1}{100}$ of the original model, or 0.015 m. Next, human agents are only allowed to calculate the velocity with that they can drive safely without risking an accident every 10 time steps to prevent a reaction time of less than one second. Furthermore, the dawdling of agents has to be strengthened because now they only dawdle for $\frac{1}{10}$ of what they are supposed to. To compensate for this, an agent that dawdled in one time step will not only dawdle in the following 9 time steps but also not accelerate against it in these time steps, which means that they reduce their velocity by a total of 10 if they dawdle. After these changes, among a few more details changes done to the equations of the model, the model is able to simulate homogeneous vehicle traffic, as shown in figure 1, after the parameters are changed accordingly. Based on this homogeneous traffic model, automated and communicating vehicles can now be introduced in the following chapter.

3 Automated and communicating vehicles

There are a few main differences between automated and human driven vehicles to point out. Firstly, automated vehicles always behave the same in all traffic situations which means that the different driving behaviours introduced in the original Lee-model should not be used by them. Instead they should always drive while keeping the safety distance. Secondly, automated vehicle do not have an attention span and so should not dawdle in any situation. Thirdly, they have a reduced reaction time of 0.5 for automated or 0.1 for communicating and automated vehicles. Communicating vehicles inform their following vehicle of their current position and velocity. This means that the following vehicle does not have to gather these information by lidar, radar, or through a camera and so it can

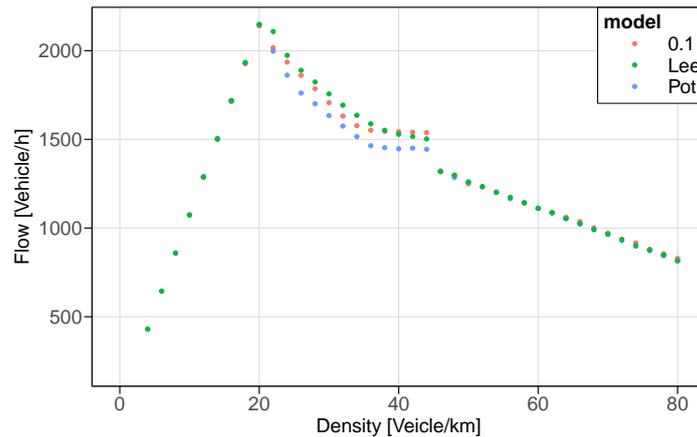


Figure 1: The fundamental diagram of the Lee-model (green), after Pottmeiers accident prevention correction (blue), as well as the new 0.1 second time-step length model (red).

operate quicker. After implementing these changes, homogeneous automated vehicle traffic can be simulated.

To also simulate heterogeneous vehicle traffic where automated and human driven vehicle mix, different driving behaviours between the vehicles have to be applied. For example, humans could always drive passively if they follow automated vehicles. However, current research suggest that this is unlikely and instead the automated vehicles should operate more passively if they follow a human driven vehicle by increasing its following distance [3]. Furthermore, they should not follow leading human drivers too closely [6] because tailing can increase the accident risk.

After these changes are implemented heterogenous traffic can be simulated and as one can see in figure 2, the traffic capacity of the road $C_m(x)$ correlates with the number of automated vehicles.

References

- [1] Wooseok Do, Omid M. Rouhani, and Luis Miranda-Moreno. Simulation-based connected and automated vehicle models on highway sections: A literature review. *Journal of Advanced Transportation*, 2019, 2019.
- [2] Daniel J. Fagnant and Kara Kockelman. Preparing a nation for autonomous vehicles: opportunities, barriers and policy recommendations. *Transportation Research Part A: Policy and Practice*, 77:167–181, 2015.
- [3] Bernhard Friedrich. The effect of autonomous vehicles on traffic. In *Autonomous driving*, pages 317–334. Springer Open, Berlin, 2016.

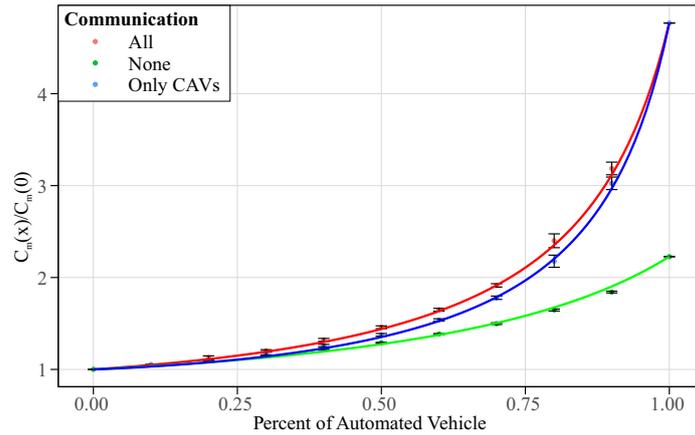


Figure 2: A example intersection for the cellular automaton model. The arrows show the directions a agent is allow to take, based on the lane he passes into the intersection from.

- [4] Hyun Keun Lee, Robert Barlovic, Michael Schreckenberg, and Doochul Kim. Mechanical restriction versus human overreaction triggering congested traffic states. *Phys. Rev. Lett.*, 92:238702, Jun 2004.
- [5] Jiaqi Ma, Fang Zhou, Zhitong Huang, Christopher L. Melson, Rachel James, and Xiaoxiao Zhang. Hardware-in-the-loop testing of connected and automated vehicle applications: A use case for queue-aware signalized intersection approach and departure. *Transportation Research Record: Journal of the Transportation Research Board*, 2672(22):36–46, 2018.
- [6] Christopher Nowakowski, Jessica O’Connell, Steven E. Shladover, and Delphine Cody. Cooperative adaptive cruise control: Driver acceptance of following gap settings less than one second. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 54(24):2033–2037, 2010.
- [7] Andreas Pottmeier, Christian Thiemann, Andreas Schadschneider, and Michael Schreckenberg. Mechanical restriction versus human overreaction: Accident avoidance and two-lane traffic simulations. In Springer, editor, *Traffic And Granular Flow*, pages 503–508. 2005.



Subproject C1

Feature selection in high dimensional data
for risk prognosis in oncology

Sven Rahmann

Alexander Schramm

Detection of circular reads in sequencing data

Till Hartmann
Lehrstuhl für Genominformatik
Universität Duisburg-Essen
till.hartmann@uni-due.de

While Oxford nanopore sequencing data has higher error rates than Illumina sequencing data, its longer reads are better suited to identify large-scale structural variation, such as gene fusions or circular DNA. Extending the standard PairHMM model by introducing "Hop" states, it is possible to express homopolymer runs and obtain statistically sound and comprehensible probabilities for sequence similarity of nanopore sequencing data.

1 Homopolymer aware PairHMM model

Pair Hidden Markov Models are a way of comparing two sequences, enabling a probabilistic assessment of the relatedness of these two sequences (with respect to the underlying model). The standard PairHMM [1] for sequence alignment has 3 relevant states¹, `Match`, `Gapx` and `Gapy`. However, nanopore reads have a higher rate of homopolymer errors, which are not captured by the standard PairHMM model. Our preliminary results show that it is a reasonable choice to modify the model in the following way while still keeping it relatively simple, under the assumption that "extending" a homopolymer run has a constant probability: First, split the singular match state `Match` into one match state `MS` for each possible nucleotide $S \in \{A, C, G, T\}$, then define two additional "hop" states `HSX` and `HSY` (where X and Y correspond to read and reference sequence, respectively) with incoming and outgoing edges only to the respective match state `MS` and one edge to itself.

¹omitting start and end states

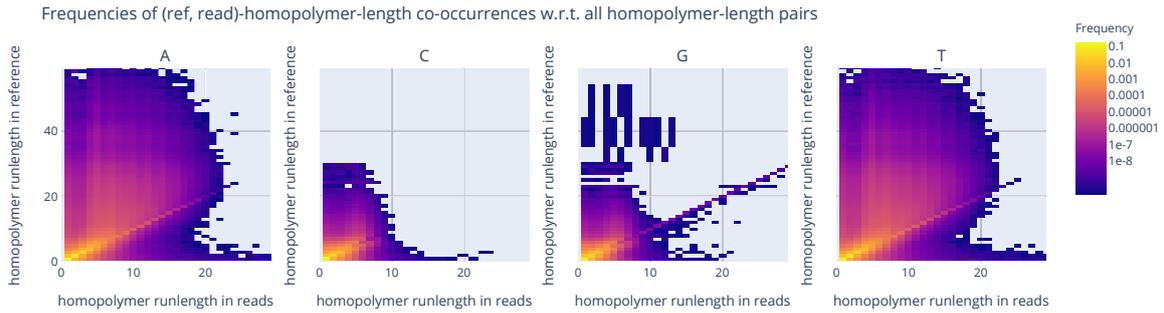


Figure 1: Relative frequencies of occurrences of (run-length in read, run-length in reference) homopolymer pairs in Oxford nanopore sequencing data, listed separately by nucleotide.

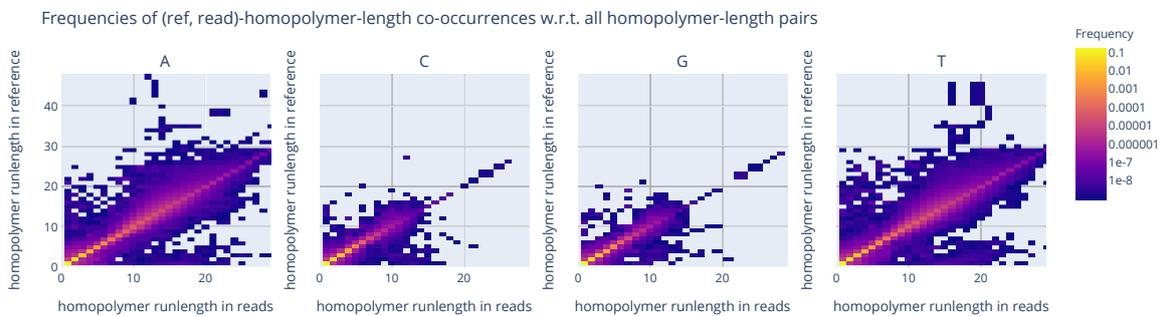


Figure 2: Relative frequencies of occurrences of (run-length in read, run-length in reference) homopolymer pairs in Illumina sequencing data, listed separately by nucleotide.

These modifications allow us to do uncertainty aware variant calling for Oxford nanopore sequencing data in *varlociraptor* [3].

In order to use this model, the probabilities of extending a homopolymer run (with respect to either the reference or the read) need to be estimated from data with known (i.e. verified) variations; all in addition to established error rates for substitutions, insertions and deletions [2,4]. The Genome In A Bottle (GIAB) project has both Illumina and Oxford nanopore data publicly available for the NA12878 sample, with accompanying verified variants [5]. Using this data, we first verify that nanopore data indeed exhibits more homopolymer errors than Illumina data: Figure 1 details the relative frequencies of occurrences of (run-length in read, run-length in reference) pairs obtained from the nanopore data, while Figure 2 does the same for Illumina data. It is obvious that homopolymer errors are way more likely in nanopore data and that homopolymer runs corresponding to deletions are more likely than those corresponding to insertions.

2 Detecting circular DNA in Oxford nanopore sequencing data

Extrachromosomal circular DNA is – as the name suggests – DNA that is both circular in nature and is found "outside" of chromosomes. As it has been associated with cancer, it is of interest to be able to identify reads that stem from such DNA.

Using Oxford nanopore sequencing technology, which can span tens to hundreds of kilobases [2], it is possible to produce reads which "traverse" circular DNA more than once. However, the way nanopore sequencing works makes reads error-prone to so called *homopolymer errors*: during base-calling, the length of runs of identical bases in a sequence are misjudged to be either larger or smaller than the actual run-lengths. To cope with these errors, we propose to collapse any homopolymer run in both nanopore reads and target sequence. While this increases ambiguity for possible alignments of the reads (or kmers thereof) to the target sequence, it eliminates the need to address homopolymer errors specifically.

Building a kmer index of the collapsed target sequence allows quick lookup of the positions of each kmer of a collapsed read in the target. Using such an index to map a read to the target, a *circular* read is characterised as follows:

1. Diagonal hit counts are high. A diagonal hit is defined as position of kmer in target – position of kmer in read. Intuitively, as long as target and read match exactly, there will only be one diagonal hit with a count of exactly the read's length.
2. Periodicity of diagonal hits. For example, the human mitochondrial genome is about 16.6kbp long. Assuming a read that traverses this sequence exactly two times, there are two diagonal hits both with a count of 16.6k and which are 16.6kbp apart.

However, there are some issues that make identifying circular reads more difficult:

1. Since reads seldom map perfectly to the target, there will be insertions, deletions and homopolymer errors which lead to "fuzzy" diagonal hit peaks.
2. It is difficult to distinguish between linear repeats (stemming from repetitive DNA) and actual circles (stemming from circular DNA).
3. Circular DNA may be composed of different parts of different chromosomes, i.e. the repeat structure can become quite complex.
4. Reads may not be long enough to traverse a circle twice.

References

- [1] Richard Durbin, Sean R. Eddy, Anders Krogh, and Graeme Mitchison. *Biological Sequence Analysis: Probabilistic Models of Proteins and Nucleic Acids*. Cambridge University Press, 1998.
- [2] Miten Jain, Hugh E Olsen, Benedict Paten, and Mark Akeson. The oxford nanopore minion: delivery of nanopore sequencing to the genomics community. *Genome biology*, 17(1):239, 2016.
- [3] Johannes Köster, Louis J Dijkstra, Tobias Marschall, and Alexander Schönhuth. Enhancing sensitivity and controlling false discovery rate in somatic indel discovery. *BioRxiv*, page 741256, 2019.
- [4] Melanie Schirmer, Rosalinda D’Amore, Umer Z. Ijaz, Neil Hall, and Christopher Quince. Illumina error profiles: resolving fine-scale variation in metagenomic sequencing data. *BMC Bioinformatics*, 17(1):125, 2016.
- [5] Justin M. Zook, David Catoe, Jennifer McDaniel, Lindsay Vang, Noah Spies, Arend Sidow, Ziming Weng, Yuling Liu, Christopher E. Mason, Noah Alexander, Elizabeth Henaff, Alexa B.R. McIntyre, Dhruva Chandramohan, Feng Chen, Erich Jaeger, Ali Moshrefi, Khoa Pham, William Stedman, Tiffany Liang, Michael Saghbini, Zeljko Dzakula, Alex Hastie, Han Cao, Gintaras Deikus, Eric Schadt, Robert Sebra, Ali Bashir, Rebecca M. Truty, Christopher C. Chang, Natali Gulbahce, Keyan Zhao, Srinka Ghosh, Fiona Hyland, Yutao Fu, Mark Chaisson, Chunlin Xiao, Jonathan Trow, Stephen T. Sherry, Alexander W. Zaranek, Madeleine Ball, Jason Bobe, Preston Estep, George M. Church, Patrick Marks, Sofia Kyriazopoulou-Panagiotopoulou, Grace X.Y. Zheng, Michael Schnall-Levin, Heather S. Ordonez, Patrice A. Mudivarti, Kristina Giorda, Ying Sheng, Karoline Bjarnesdatter Rypdal, and Marc Salit. Extensive sequencing of seven human genomes to characterize benchmark reference materials. *Scientific Data*, 3(1):160025, Jun 2016.

Detecting Extrachromosomal Circular DNA in Lung Cancer by Nanopore Sequencing

Alicia Isabell Tüns
Molekulare Onkologie
Innere Klinik/Tumorforschung
Universitätsklinikum Essen
Alicia.Tuens@uk-essen.de

Lung cancer (LC) is the leading cause of cancer-related death and five-year survival rates are below 20% due to advanced stage at diagnosis and therapy resistance. Early detection of tumor progression and resistance to therapy is therefore an unmet medical need. Studies have shown that extrachromosomal circularization of DNA (ecDNA) is a genetic marker of aggressive LC. Since nanopore sequencing processes long reads, it is an emerging technology potentially allowing the detection of circularized DNA with high sensitivity and specificity. In the current project we are establishing a workflow for detecting circularized DNA in lung cancer cell lines.

1 Introduction

Lung cancer (LC) is the most common cancer type and the leading cause of cancer-related death. LC can be categorized into small-cell lung cancer (SCLC) and non-small-cell lung cancer (NSCLC) that account for 15% and 85%, respectively [1]. Molecular profiling of aggressive LC revealed various genetic alterations such as extrachromosomal circularization of DNA (ecDNA) [2]. ecDNAs originate from the linear genome and encode whole genes or just regulatory elements [3]. The process of extrachromosomal circularization in cancer is a powerful driver of intratumoral heterogeneity. Since ecDNAs replicate during cell division but lack centromeres, they undergo non-mendelian

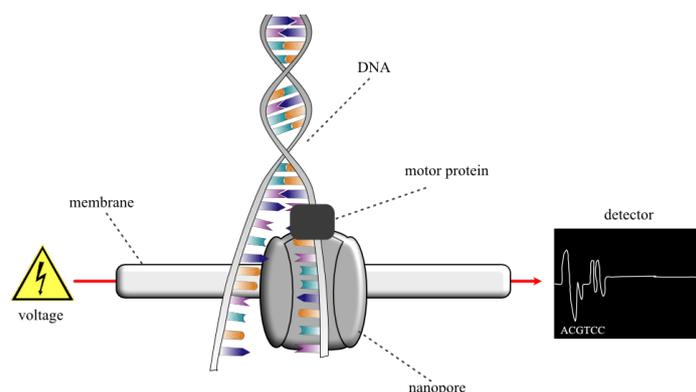


Figure 1: When a DNA strand approaches the nanopore, a motor protein unzips the DNA and pulls the single-strand through the nanopore. The nucleotides entering the nanopore cause characteristic changes in the ion current that can be measured and characterized.

inheritance which enables tumors to achieve high oncogene copy numbers. Additionally, ecDNAs contribute to genome remodeling through reintegration into the linear genome at a different site [2,4]. Therefore, extrachromosomal circularization possibly enables tumors to evolve rapidly in response to changing conditions such as treatments, which could reveal a new therapy resistance mechanism [5]. To assess the clinical relevance of ecDNAs in lung cancer, first their complex structures must be analyzed by conventional and new sequencing strategies. Nanopore sequencing is a relatively new method that allows real-time sequencing of long reads by measuring ion current changes when a DNA/RNA strand passes through a nanopore (figure 1). The extended read length potentially facilitates the detection of ecDNAs and their structural analysis.

2 Detecting ecDNA in lung cancer cell lines using the nanopore sequencer MinION

High molecular weight DNAs from nine lung cancer cell lines were enriched for circular structures by exonuclease digestion of linear chromosomal DNA (figure 2A). Afterwards, circular DNAs were amplified and sequenced using the MinION nanopore sequencer and the *Rapid Barcoding Kit* (SQK-RBK004, Oxford Nanopore Technologies). Reads that passed the quality control were aligned against the human Ensembl DNA database (GRCh38) using minimap2. When dealing with long nanopore reads, two categories of reads confirming circularity were expected. Firstly, the amplification of circular structures produces long DNA strands with repetitive sequences, so that nanopore reads containing repetitive motifs indicate a circular origin. Secondly, reads covering the point of fusion, which appear as split reads in a linear alignment, suggest circularity (figure 2B). To test



Figure 2: **A**, ecDNA enrichment. **B**, Types of reads produced by nanopore sequencing that suggest circular structures. **C**, Circular mitochondrial chromosome. **D**, Representative reads, which map to MT genome and support ecDNA detection strategy. **E**, Number of sequenced reads and their average length. **F**, Principle of de-novo assembly to reconstruct long circles. **G**, Number of reconstructed circles and their size. **H**, MT genome de-novo assembly of H1975 and PF611.

our ecDNA detection strategy, the mitochondrial chromosome served as a positive control (figure 2C). By visual inspection, both categories of reads were observed, however only one read with a MT genome coverage $> 1x$ was identified throughout all samples (figure 2D). Considering this observation and the average read length of 1835 bp across all samples, looking for repetitive motifs in reads just allows for the detection of small circles. To identify larger circles, de-novo assembly of all reads was performed using flye (figure 3F). After the assembly, the average size of circles was 11841 bp (figure 3G). Next, the proper assembly of the MT genome was checked. The accurate MT assembly of sample H1975 was successful, however failed for the rest of the samples (figure 2H).

3 Conclusion and Outlook

Although nanopore sequencing processes longer reads than conventional sequencing methods, the detection and structural analysis of ecDNAs is still challenging. Long reads with repetitive motifs mainly allow for identification of small circles. Larger circles must be reconstructed by de-novo assembly using long read de-novo assemblers such as flye. Although flye reconstructed various circular structures, proper MT genome assembly failed for most of the samples. These technical issues are currently being addressed by adjusting the settings or using alternative assemblers like raven or canu.

4 References

- [1] HERBST, R.; HEYMACH, John V.; LIPPMAN, S. Molecular origins of cancer. *N Engl J Med*, 2008, 359. Jg., Nr. 13, S. 1367-80.
- [2] KIM, Hoon, et al. Extrachromosomal DNA is associated with oncogene amplification and poor outcome across multiple cancers. *Nature Genetics*, 2020, 52. Jg., Nr. 9, S. 891-897.
- [3] VERHAAK, Roel GW; BAFNA, Vineet; MISCHER, Paul S. Extrachromosomal oncogene amplification in tumour pathogenesis and evolution. *Nature Reviews Cancer*, 2019, 19. Jg., Nr. 5, S. 283-288.]
- [4] KOCHER, Richard P., et al. Extrachromosomal circular DNA drives oncogenic genome remodeling in neuroblastoma. *Nature Genetics*, 2020, 52. Jg., Nr. 1, S. 29-34.
- [5] TURNER, Kristen M., et al. Extrachromosomal oncogene amplification drives tumour evolution and genetic heterogeneity. *Nature*, 2017, 543. Jg., Nr. 7643, S. 122-125.



Subproject C3

Multi-level statistical analysis of
high-frequency spatio-temporal process data

Katharina Morik Wolfgang Rhode
Tim Ruhe

Recent developments on the lepton propagator PROPOSAL

Jean-Marco Alameddine
Experimentelle Physik 5
Technische Universität Dortmund
jean-marco.alameddine@tu-dortmund.de

The Monte Carlo simulation library PROPOSAL has been developed to simulate the propagation of muon and tau leptons, for example for the Ice-Cube Neutrino Observatory. Currently, this library is being extended to be able to simulate electron, positron and photon interactions. This offers the possibility to use PROPOSAL as an electromagnetic shower model in the currently developed new version of the air simulation software CORSIKA. This report provides a summary of the recent developments in PROPOSAL as well as an overview of resulting future research possibilities.

1 The Monte Carlo simulation software PROPOSAL

PROPOSAL is a Monte Carlo software library to simulate the propagation of high-energy particles¹. To predict properties of unlabeled, real experimental data sets, these Monte Carlo simulations are necessary since they provide labeled datasets used to train and validate the underlying data analysis methods. Hence, a Monte Carlo simulation software such as PROPOSAL needs to be optimized for both physical accuracy and computational performance.

Originally, PROPOSAL has been developed [2] and enhanced [6, 7] in Dortmund, mainly to be used in the simulation chain of the IceCube Neutrino Observatory. Here, PROPOSAL simulates the interactions of high-energy muon and tau leptons, which create characteristic detector signatures in the ice [9].

¹The current development of PROPOSAL is publicly available under: <https://github.com/tudo-astroparticlephysics/PROPOSAL>.

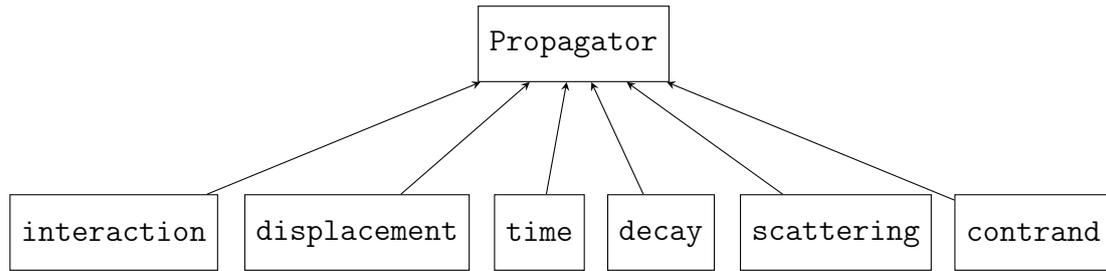


Figure 1: Modular structure of the physical calculation components of PROPOSAL. The user can either use each module individually or a combined version of all modules called `Propagator`, providing a full particle simulation. The individual modules are described in more detail in [5].

2 Inclusion of PROPOSAL in the shower simulation framework CORSIKA

CORSIKA is Monte Carlo computer software that is widely used for the simulation of extensive air showers [4]. Accurate Monte Carlo simulations from these air showers are important for all experiments using showering effects for their observations, are directly interested in the detection of cosmic rays or need to consider them as a background.

Currently, a new version of CORSIKA is under development, replacing the old FORTRAN code base with a modern and modular code structure with the goal to turn CORSIKA 8 into a flexible particle cascade simulation framework [10]. For CORSIKA 8, PROPOSAL will be used as an external library to provide the description of electromagnetic interactions for electrons, positrons and photons as well as muon and tau leptons.

To prepare PROPOSAL for the integration into CORSIKA 8, a description of photon propagation as well as an improved description of electron and positron propagation has been implemented in PROPOSAL [1]. This includes the description of Compton scattering and electron-positron pair production for photons and the description of annihilation for positrons. Furthermore, a restructuring of the PROPOSAL code structure, splitting all physical calculations necessary for the propagation process into six individual modules, as shown in Figure 1, has been conducted [5]. This allows CORSIKA, where the transport code is interested in single propagation step suggestions, to use the individual physical modules from PROPOSAL. Additionally, all modules are still combined to a `Propagator` object, providing a full particle propagation simulation. This is necessary for analysis interested in energy losses of a single particle track, for example the IceCube simulation chain.

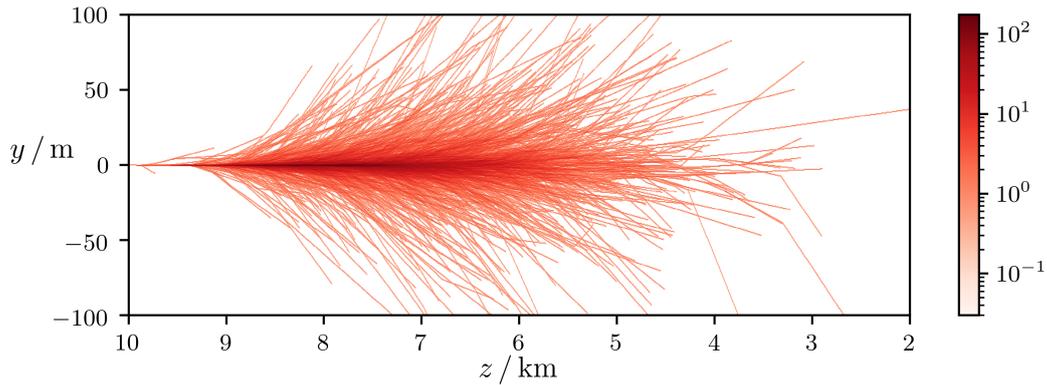


Figure 2: Particle density of an electromagnetic shower, induced by a photon with an initial energy of 10^6 MeV at a height of $z = 10$ km. Particles with an energy below 50 MeV have been omitted. Adapted from [1].

3 Results and outlook

As a proof of concept and as a first validation of the physical description of electron, positron and photon propagation, PROPOSAL-internal electromagnetic showers have been successfully simulated. One example of an electromagnetic shower, induced by a photon, is shown in Figure 2.

A first version of an interface, combining CORSIKA 8 and PROPOSAL with its improved code structure, has been implemented². Using this interface, it was possible to produce first CORSIKA 8 showers, with the electromagnetic component being simulated by PROPOSAL.

Currently, a new major release of PROPOSAL, including the new code structure written for the CORSIKA 8 interface, is about to be finalized and will be released soon. Preliminary results show that this release will either decrease or, at the very least, not significantly increase, the runtime of PROPOSAL simulations. Further physical improvements of this release will include the deflection of particles during stochastic interactions, the possibility to resimulate particles after the initial propagation process and minor bug fixes.

After the release, the interfaces for the IceCube simulation chain as well as CORSIKA 8 need to be updated to benefit from these improvements. Further studies in CORSIKA 8 need to be conducted to validate the results provided by PROPOSAL. This may include comparisons with the previous version of CORSIKA, CORSIKA 7, which used EGS4 as an electromagnetic shower model [8].

²The current development of CORSIKA 8 is publicly available under: <https://gitlab.ikp.kit.edu/AirShowerPhysics/corsika>.

PROPOSAL provides different, modern parametrizations for each physical interaction type that can be individually changed for each analysis. This provides the possibility to investigate how changes in the theoretical description of electromagnetic processes affect the whole shower development. Yet unresolved problems in observations of air showers, for example the Muon Puzzle, a significant discrepancy between observations and simulations reporting different muon numbers, may benefit from these new possibilities [3].

References

- [1] J. Alameddine. Simulation of rare processes and electromagnetic shower components within the Monte Carlo propagation library PROPOSAL. Master's thesis, Technische Universität Dortmund, February 2020.
- [2] D. Chirkin and W. Rhode. Propagating leptons through matter with Muon Monte Carlo (MMC). July 2004.
- [3] H.P. Dembinski. The Muon Puzzle in High-Energy Air Showers. *Phys. Atom. Nucl.*, 82(6):644–648, 2020.
- [4] D. Heck et al. CORSIKA: A Monte Carlo code to simulate extensive air showers. February 1998.
- [5] J. Alameddine et al. PROPOSAL: A library to propagate leptons and high energy photons. *Journal of Physics: Conference Series*, 1690:012021, December 2020.
- [6] J.-H. Koehne et al. Proposal: A tool for propagation of charged leptons. *Computer Physics Communications*, 184(9):2070–2090, 2013.
- [7] M. Dunsch et al. Recent improvements for the lepton propagator proposal. *Computer Physics Communications*, 242:132–144, 2019.
- [8] W. R. Nelson et al. The EGS4 Code System. December 1985.
- [9] M. G. Aartsen et al. (IceCube Collaboration). Neutrinos below 100 TeV from the southern sky employing refined veto techniques to IceCube data. February 2019.
- [10] M. Reininghaus and R. Ulrich. CORSIKA 8 - Towards a modern framework for the simulation of extensive air showers. In *European Physical Journal Web of Conferences*, volume 210, page 02011, October 2019.

GPU Accelerated IACT/Fluorescence Simulation

Dominik Baack
E5b Physic
Technical University Dortmund
dominik.baack@tu-dortmund.de

The physics-driven and widely used astroparticle simulation CORSIKA [3] is an ideal candidate for the application and further development of the techniques developed within SFB876 "Providing Information by Resource-Constrained Data Analysis". With this focus on computational efficiency, new simulations that are currently limited by runtime can be enabled. On the other hand, using fewer environmental resources, such as electrical power, can reduce computational costs. As several new or expanded cosmic ray experiments are conducted in the coming years, the need for simulation increases equally. In addition, the long-standing CORSIKA 7 Fortran codebase will be completely rewritten to a "state of the art" C++ simulation framework [2], this will allow the use of new techniques that were not possible in previous work [1].

The methods described here are specifically tailored to the field of optical light propagation through the atmosphere. Used in Imaging Atmospheric Cherenkov Telescopes (IACT) and fluorescence telescopes, light propagation currently contributes to over 80 percent of the total runtime. With the application of early cuts optimized by machine learning and specific hardware tailored for parallel execution, such as GPUs, the runtime can be greatly reduced.

1 Architecture

To use the developed acceleration methods across a variety of computing cluster infrastructures, an additional layer is required that separates specialized hardware (e.g.

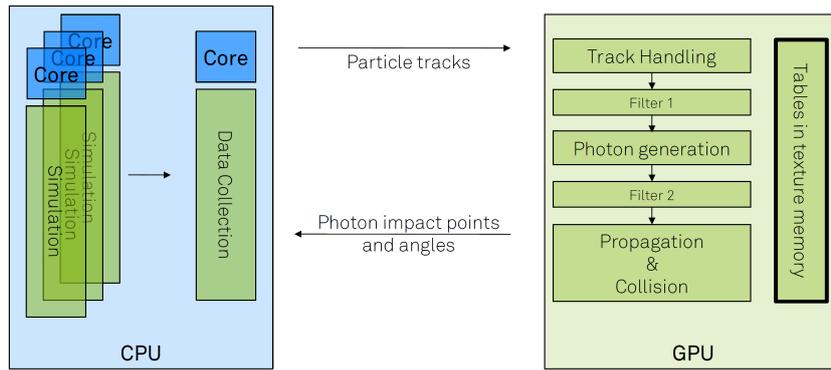


Figure 1: Schematic presentation of the program flow used for the simulation of optical photons inside the context of the particle cascade simulation framework CORSIKA.

GPUs) from the base simulation. This allows independent handling of optical photons on nodes where acceleration hardware is not available. Aggregating particle data from multiple simulation instances has the added benefit of accumulating enough particle traces to maximize GPU utilization. Figure 1 shows the basic structure of the program flow schematically. To offload as much work as possible and reduce the amount of data to be transferred, photon generation from the individual particle traces is performed on the GPU. This includes two filter stages that apply cuts to particles and later photons that are very unlikely to be measurable by experiment.

2 Photon handling

In the propagation of optical photons through the atmosphere, two main effects must be considered. These effects are refraction and absorption. Absorption is a superposition of several individual effects that lead to the removal of photons from the measurement result. Optical refraction is due to the fact that as the distance from the ground increases, the density and, to a lesser extent, the composition of the atmosphere changes. Consequently, the local refractive index also changes. As a result, photons cannot propagate in a straight line. This relationship is illustrated in figure 2.

2.1 Absorption

In the propagation of optical photons, several effects lead to the extinction of individual photons. Some examples are listed here:

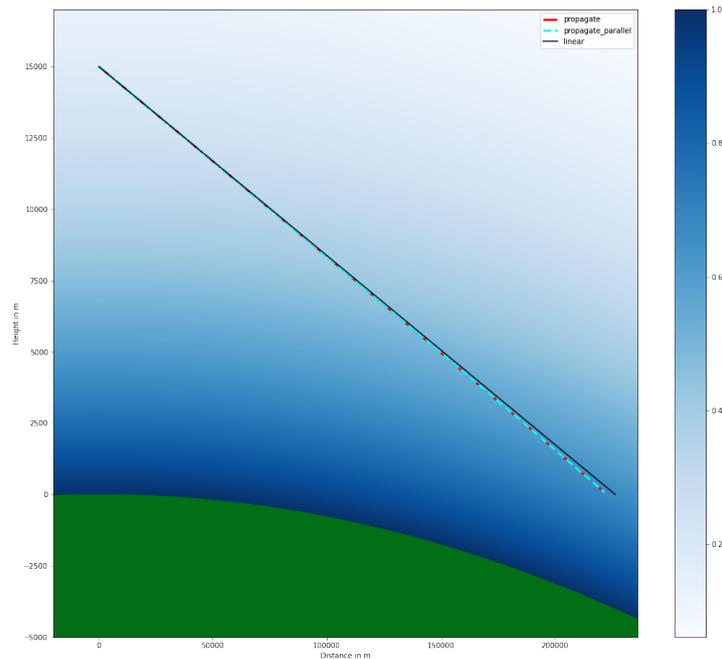


Figure 2: Shown is the deviation of the straight line propagation of an optical photon due to atmospheric refraction. The curved lines were calculated with two different methods, the red line (propagate) with a precise numerical integration and the blue line (propagate_parallel) with a fast approximation method.

- Atmospheric absorption as a function of propagation distance, wavelength and position.
- specular reflectivity as a function of wavelength
- Sensor acceptance

These effects can lead to an overall reduction of individual Cherenkov photons by orders of magnitude. By folding all absorption effects that may occur during propagation into the Cherenkov spectrum emitted by the particles, the number of optical photons that must be calculated can be drastically reduced. The price for this is some loss of accuracy for more complex effects such as angle-dependent reflectivity, but even these could be applied subsequently in an additional step.

2.2 Atmospheric refraction

In modern experiments, even the comparatively small effect of refraction can be measured. The two figures 4 and 3 show the difference between the physically correct simulation using the measured refractive index (in this case the US standard atmosphere [4]) and the

simple straight line propagation. For real atmospheric conditions, there is no analytical expression for calculating the point of impact. Therefore, one possibility is to approximate the change in refractive index using one or more exponential distributions. Comparison with a computationally intensive but accurate numerical integration along the path shows that for real atmospheres the error introduced is higher than acceptable for modern telescopes. Therefore, an interpolation method based on linear interpolation including correction factors was developed. The use of a specialized texture memory on the GPU leads to a fast and scalable method for the calculation of photon impact points and angles.

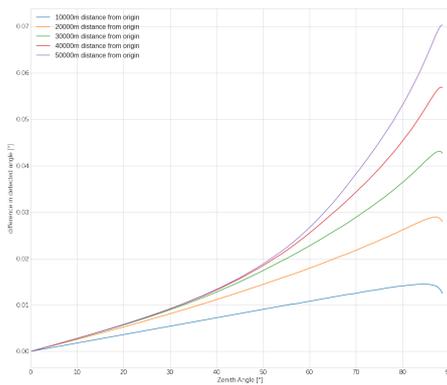


Figure 3:

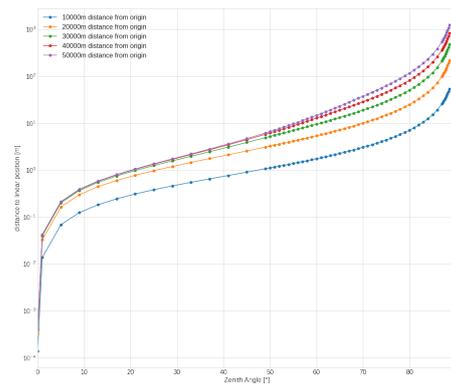


Figure 4:

References

- [1] Dominik Baack. Data reduction for CORSIKA. Technical Report 2, E5b, Faculty Physic, TU Dortmund, 06 2016.
- [2] Ralph Engel, Dieter Heck, Tim Huege, Tanguy Pierog, Maximilian Reininghaus, Felix Riehn, Ralf Ulrich, Michael Unger, and Darko Veberič. Towards a next generation of corsika: a framework for the simulation of particle cascades in astroparticle physics. *Computing and Software for Big Science*, 3(1):2, 2019.
- [3] Dieter Heck, J Knapp, JN Capdevielle, G Schatz, T Thouw, et al. Corsika: A monte carlo code to simulate extensive air showers. *Report fzka*, 6019(11), 1998.
- [4] Arlin J Krueger and Raymond A Minzner. A mid-latitude ozone model for the 1976 us standard atmosphere. *Journal of Geophysical Research*, 81(24):4477–4481, 1976.

Deep Learning based Likelihood Reconstruction of CTA Events

Noah Biederbeck

Experimentelle Physik 5

Technische Universität Dortmund

noah.biederbeck@tu-dortmund.de

The prototype of the LST1 is the first of over one hundred Imaging Air Cherenkov Telescopes (IACTs) for the new Cherenkov Telescope Array (CTA). It is located at the Roque de Los Muchachos Observatory on La Palma, Spain, right next to the MAGIC Telescopes, which it exceeds in mirror area by a factor of two.

The current state-of-the-art event reconstruction is based on the classical approach publicly known as *Hillas Analysis*.

In this report I will show the feasibility of a likelihood approach which is supported by deep learning. This technique makes use of an analysis paradigm that, without further ado, plays out its full power when applied to high-multiplicity events.

Classical Hillas Analysis

Since the first Large-Sized Telescope (LST1) is under commissioning, little data has been taken. Although, not to forget, the Crab Nebula, depicted as *standard candle* for Cherenkov astronomy has been detected in November 2019 [1].

One has to use Monte Carlo simulations to train machine learning algorithms that predict event properties on measured data. In this case, I use the Prod3 La Palma Benchmark dataset. It consists of over 10 000 simulated runs, each containing a multitude of events where at least a single telescope triggered.

In a classical analysis, the obtained images get cleaned (using some hard values) and parameterized. On these parameters, classical machine learning algorithms like a random

forest get trained. This approach is more or less standard procedure since the 1990s [2]. A big plus of this approach is the usage of battle-tested algorithms and parameterizations. Also, it's fast and robust.

Drawbacks of the classical approach are that most reconstruction algorithms get more complex the more telescopes triggered a single event, and the missing possibility of giving uncertainties to the estimations.

Deep Learning facilitated Likelihood Predictions

A simple neural network architecture predicts the Cherenkov photon content of a pixel from a given set of parameters: energy, direction as source position in the field-of-view (2), pointing information (2), and impact position relative to the telescope (2). These parameters get scaled in a way that the input for the neural network are numbers in the range of $[-1, 1]$.

The architecture is defined as five fully connected layers, with increasing and again decreasing amount of nodes. I use $\tanh(x)$ as activation function for the layers and $\exp(x)$ as activation for the output layer. A single batch consist of a full camera, for the LST consisting of 1855 individual pixels. The loss function is a negative logarithmic Poisson likelihood loss:

$$L_{\text{pixel}} = \hat{y} - y \ln(\hat{y} + \epsilon) + y \ln(y + \epsilon) - y \quad (1)$$

This makes it easy to scale to the full batch (a full camera) and beyond:

$$L_{\text{telescope}} = \sum L_{\text{pixel}} \quad (2)$$

$$L_{\text{array}} = \sum L_{\text{telescope}} \quad (3)$$

Training is done in a standard fashion for up to 1000 Epochs. See the training loss in figure 1.

In figure 2 I show for a selected camera image the performance of the model.

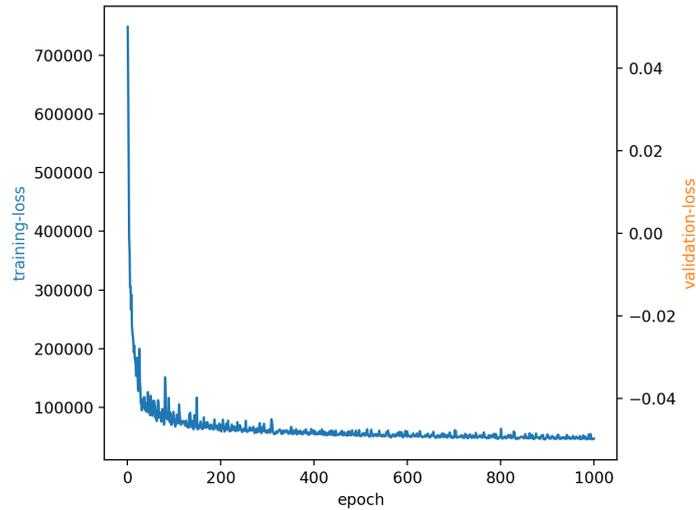


Figure 1: Training loss.

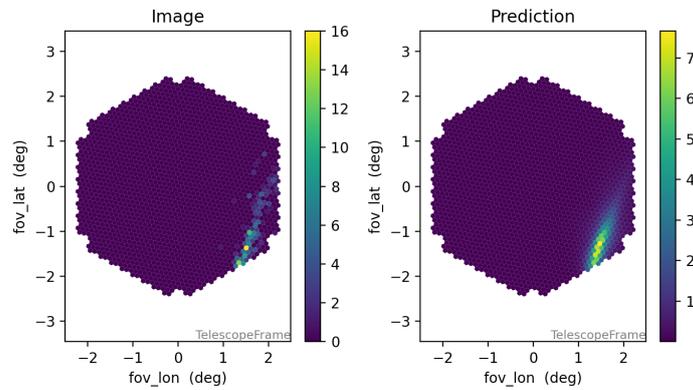


Figure 2: Truth and model prediction for a selected camera image.

Reconstruction of Event Properties

Using another very simple neural network architecture I implemented a minimizer for the above network. The trainable parameters are the above mentioned event properties. These get trained in a similar fashion to that of the original network. Basically, for each event this network will over-fit into the minimum for the estimation of the original parameters.

In figure 3 I show the energy migration using this minimization approach.

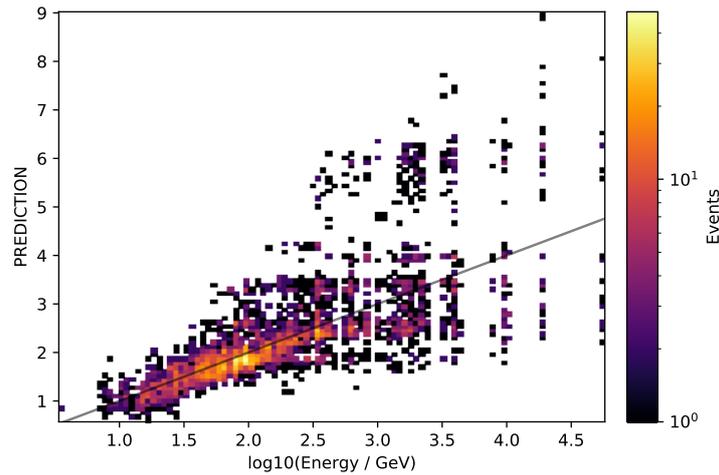


Figure 3: Energy estimation of the minimization procedure. Minimizer trained (i.e. over-fitted) for 100 epochs.

Conclusion and Future Work

As seen in figure 3, the model overestimates the higher energies. This is due to the low number of training events used in this report. The reason for other deviations is that the model predicts not individual photon counts per pixel, but regresses floating-point values that describe the photon count. Additionally, the shower blobs in the camera look like a 2D-Gaussian, and not at all hacked into pieces like the true image. The images in figure 2 show both points. This needs to be fixed, likely with a different network architecture. Predicting integer-values is another approach that could fix this problem, but the automatic differentiation packages of popular deep learning frameworks (e.g. pytorch [3]) do not implement this.

References

- [1] CTA Observatory. The LST-1 detects its first gamma-ray signal, 2019.
- [2] A. M. Hillas. Cerenkov light images of EAS produced by primary gamma. *Proceedings from the 19th International Cosmic Ray Conference*, 3:445–448, August 1985.
- [3] Adam Paszke et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.

Optimality in Active Class Selection

Mirko Bunse

Lehrstuhl für Künstliche Intelligenz

Technische Universität Dortmund

mirko.bunse@tu-dortmund.de

Active class selection tries to optimize the class proportions in newly acquired data, so that a classifier trained from this data achieves maximum performance during deployment. In the last year, we have presented an information-theoretic examination of this task, from which we can draw at least two conclusions: First, the more data is acquired, the less beneficial are class proportions which are different from those that naturally occur during deployment. Second, the degree to which other class proportions are eligible depends on the correlation between features and labels. In the long run, we intend to employ the active selection of classes as a resource-aware control mechanism for astro-particle simulations.

1 Introduction

Astro-particle simulations are label-dependent data generators [2]. That is, they produce features X based on the latent quantities Y which are to be predicted, e.g. based on the type of the primary particle, its direction and energy. It is therefore necessary to decide for the proportions $\vec{q} \in \mathbb{R}^{|\mathcal{Y}|}$ of classes before any training data can be simulated. The idea behind active class selection (ACS, see Fig. 1) is to base this decision iteratively on information about the training progress of the classifier [4, 5].

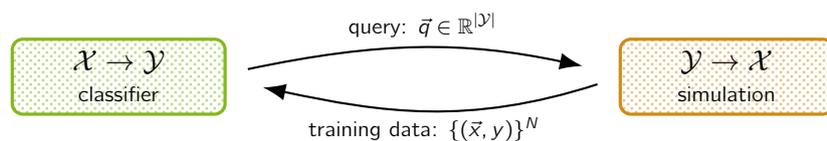


Figure 1: N additional samples are produced from desired label proportions $\vec{q} \in \mathbb{R}^{|\mathcal{Y}|}$.

2 Theoretic Analysis for Probabilistic Classifiers

Previous publications on ACS have focussed on the empirical evaluation of heuristic data acquisition strategies. Our ICDM-paper [3] adds an information-theoretic contribution to this debate, which is based on Kullback-Leibler divergences $\hat{d}_{(\cdot)}$ and $\bar{d}_{(\cdot)}$. Here, a vertical bar relates to divergences between the training set and the deployment data; a hat relates to divergences between the deployment data and the classifier; subscripts denote the random variables that are addressed by these divergences. ACS intends to minimize $\hat{d}_{Y|X}$, the error of the predictor. We are interested in how $\hat{d}_{Y|X}$ is affected by \bar{d}_Y , the difference between training and test data in terms of class proportions.

We assume that the data-generating mechanism $\mathcal{Y} \rightarrow \mathcal{X}$ is identical among the training and test sets. This assumption allows us to study the effect of mistaken class proportions in isolation; yet, it is weaker than assuming i.i.d. data. In terms of divergences, assuming an identical mechanism means to assume that $\bar{d}_{X|Y} = 0$. Moreover, we assume that a Bayesian classifier learns a consistent generator, i.e. for any $\varepsilon > 0$ there is an $N \in \mathbb{N}$ such that $d(\bar{\mathbb{P}}(X | Y) || \hat{\mathbb{P}}_n(X | Y)) \leq \varepsilon$ for all $n > N$.

We analyze $\hat{d}_{Y|X}$ in two steps (see Fig. 2). First, we study the effect of ACS on the distribution of the training data, ignoring the potential effects of estimating $\mathbb{P}(Y | X)$. Second, we introduce a Bayesian classifier and finite amounts of data into our analysis. Each analysis step yields one theorem which we can study in experiments.

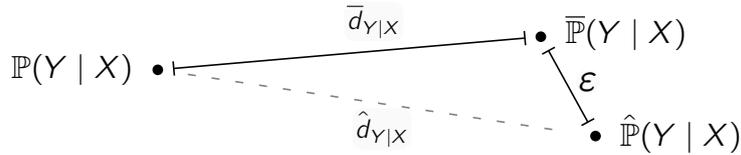


Figure 2: The two steps of our analysis: $\hat{d}_{Y|X}$ is bounded through $\bar{d}_{Y|X}$ and ε .

Theorem 1: The amount of dis-information in the training data is bounded above by the divergence between the natural and the ACS-chosen class proportions [3]:

$$\bar{d}_{Y|X} = \bar{d}_Y - \bar{d}_X \leq \bar{d}_Y$$

Consequently, if we want to minimize $\bar{d}_{Y|X}$, we can instead minimize its upper bound \bar{d}_Y . The beneficial influence of \bar{d}_X stems from the correlation between features and labels. For instance, a full correlation would yield $\bar{d}_X = \bar{d}_Y$ and thus always minimize $\bar{d}_{Y|X}$.

Theorem 2: The error of a Bayesian classifier with a consistent generative estimator $\mathcal{Y} \rightarrow \mathcal{X}$ is bounded above by the error of the assumed class proportions and by the error $\varepsilon_{X|Y}$ of the generative part. Eventually, $\varepsilon_{X|Y}$ will become small [3]:

$$\hat{d}_{Y|X} \leq \hat{d}_Y + \varepsilon_{X|Y}$$

3 Experiments

We check the tightness of Theorem 1 by altering the probability $\bar{\mathbb{P}}(Y = y)$ of each class $y \in \{0, 1, 2\}$ separately, keeping the relative probability of the other classes intact. The results in Fig. 3 indicate that—due to the correlation between features and labels—several class proportions other than the natural ones can yield accurate classifiers.

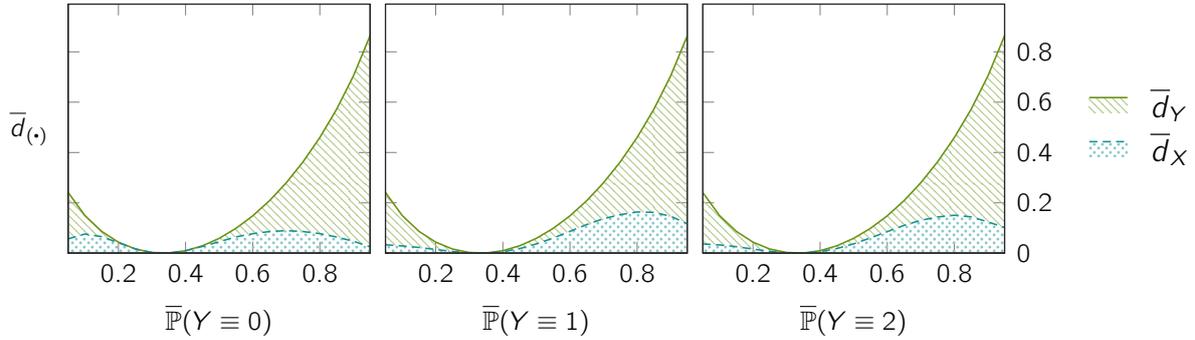


Figure 3: The divergence $\bar{d}_{Y|X} = \bar{d}_Y - \bar{d}_X$ for the 3clusters data set. Class proportions with $\bar{d}_X \approx \bar{d}_Y$ will not impair the performance of the classifier [3].

In another experiment we confirm the suggestion of Theorem 2, that the natural class proportions will become better during data acquisition, due to consistency. To this end, we rank several class proportions $\vec{q} \in \mathbb{R}^{|\mathcal{Y}|}$ with respect to their classifier accuracy. Fig. 4 displays the decrease of the average rank of the natural proportions while data is acquired. Any \vec{q} that is not natural is losing its potential benefit over time. ACS strategies should transition towards the natural proportions as more data is being acquired.

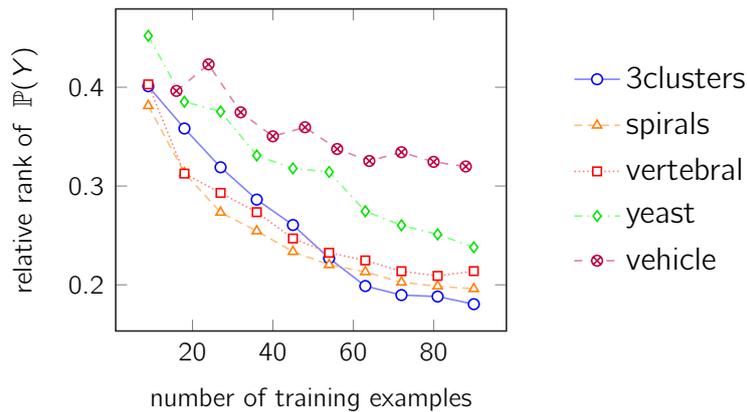


Figure 4: The natural proportions gain superiority over their competitors, indicated by decreasing relative ranks in terms of accuracy, while the training set grows.

4 Conclusion

Our theoretical findings confirm a common empirical observation, that even sophisticated ACS strategies often cannot outperform the class proportions that also occur in the test data. Moreover, our findings enrich this observation by revealing that the severeness of using other, non-natural proportions depends i) on the size of the training set; and ii) on the underlying conditional distribution $\mathbb{P}(X | Y)$.

Since simulated data is the basis for every analysis in astro-particle physics, optimizing simulations for resource efficiency will improve the resource efficiency of all work packages in the C3 project. To get there, we still have to extend our analysis towards two particularities of the astro-particle use case. First, the true proportions of classes are highly imbalanced. Second, they are not precisely known. Class ratios are usually reported to be roughly between $1:10^3$ and $1:10^4$ [1]. The latest information about our efforts, now also featuring videos, is hosted at <https://sfb876.tu-dortmund.de/simulation-data-mining>.

References

- [1] Christian Bockermann, Kai Brügge, Jens Buss, Alexey Egorov, Katharina Morik, Wolfgang Rhode, and Tim Ruhe. Online analysis of high-volume data streams in astroparticle physics. In *ECML-PKDD*, volume 9286 of *LNCS*, pages 100–115. Springer, 2015.
- [2] Mirko Bunse, Amal Saadallah, and Katharina Morik. Towards active simulation data mining. In *Proc. of the 3rd Int. Workshop and Tutorial on Interactive Adaptive Learning at ECML-PKDD 2019*, volume 2444 of *CEUR Workshop Proceedings*, pages 104–107, 2019.
- [3] Mirko Bunse, Dorina Weichert, Alexander Kister, and Katharina Morik. Optimal probabilistic classification in active class selection. In *International Conference on Data Mining (ICDM)*. IEEE, 2020. In press.
- [4] Daniel Kottke, Georg Krempl, Marianne Stecklina, Cornelius Styp von Rekowski, Tim Sabsch, Tuan Pham Minh, Matthias Deliano, Myra Spiliopoulou, and Bernhard Sick. Probabilistic active learning for active class selection. In *Proc. of the NIPS Workshop on the Future of Interactive Learning Machines*, 2016.
- [5] Rachel Lomasky, Carla E. Brodley, M. Aernecke, D. Walt, and Mark A. Friedl. Active class selection. In *Proc. of the ECML 2007*, volume 4701 of *LNCS*, pages 640–647. Springer, 2007.

Particle Classification of Cosmic Rays with a Machine Learning Approach using IACT Data

Alicia Fattorini
Experimentelle Physik 5
Technische Universität Dortmund
alicia.fattorini@tu-dortmund.de

different The MAGIC telescopes, two Air Imaging Cherenkov Telescopes at La Palma are sensitive in the energy regime from the GeV to the TeV range and are extremely capable instruments for studying gamma-ray sources in the Universe. With a ratio of photons to hadrons up to 1:10000 for the detected showers, the background offers large statistics. The measured hadrons are protons and heavier nuclei such as helium and iron. The main challenge is to estimate the class of the primary particle, since particle Identification can yield interesting results both for background rejection and scientific studies of the Cosmic Ray spectrum. The approach is a random forest classifier trained and tested on monte-carlo simulations of protons, helium and iron nuclei. The separation of the particles with IACT data is quite challenging, therefore an approach for a correction of the data sets follows.

The MAGIC Telescopes

MAGIC [1] is a system of two Imaging Air Cherenkov Telescopes at the Roque de los Muchachos, a mountain of 2200 m height at La Palma, Canary Islands. Originally it was built for the detection and investigation of gamma-ray sources. Besides high energetic photons, cosmic rays also cause showers entering the atmosphere. The ratio of photons to hadrons is up to 1:10000 for the showers detected in the MAGIC's energy detection range. Hence the measurements of IACTs offer besides the analysis of gamma-ray sources a great opportunity to investigate the cosmic ray spectrum in the certain energy regime.

The Cosmic Ray Spectrum

Unlike photons or neutrinos, hadrons pose the challenge that their origin is difficult to reconstruct because they can be deflected by intergalactic magnetic fields on the way from their source to earth. This leads to the fact, that a diffuse cosmic ray flux is measured on earth. Although the cosmic ray spectrum is now being investigated in many experiments, the energy range of MAGIC is relatively unexplored. Therefore it is reasonable to create hadronic simulations and build models to classify the hadronic particles detected by MAGIC.

Analysis and Results

This approach uses a random forest to classify the hadronic particle. The framework named 'aict-tools' [3] is used, especially written for the analysis of CTA and FACT data. The random forest is trained on simulated data for proton, helium and iron showers above a total charge of 250 phe in the camera pixels identified as part of the shower image. In this work the classification is done in two steps. The first step is the separation of iron particles from the lighter particles, assuming that the the air shower projection in the MAGIC camera clearly changes with the massnumber of the particle. In Figure 1 the results of the classification is shown.

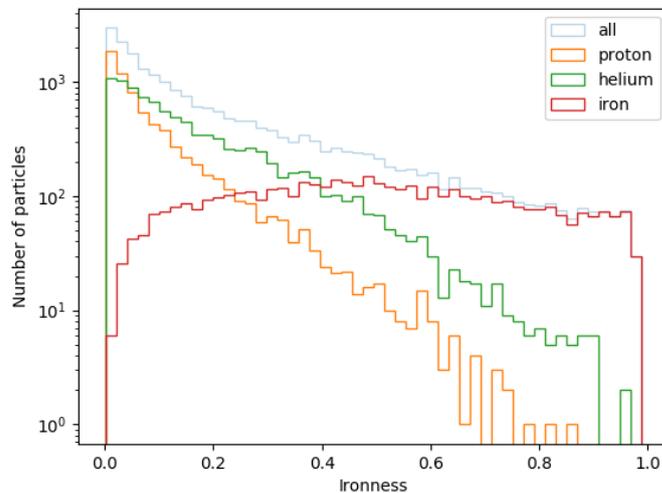


Figure 1: Distribution of Ironness of the different particles.

The classifier trained on a data set of protons, helium and iron identifies most of the lighter particles as non-iron nuclei. For proton and helium the classification with the

random forest looks promising. On the other hand the identification of iron is challenging. This can be explained by the fact that the shower images of iron can be similar to those of light particles especially when many subshowers are outside the camera.

The cut of 0.4 is done in the Ironness parameter to separate the heavier from the lighter particles.

The second step of the particle classification is the separation of helium and protons. The random forest is trained on a data set consisting of simulated protons and helium. The resulting classification of the remaining particles after the previous cut in the Ironness is shown in Figure 2.

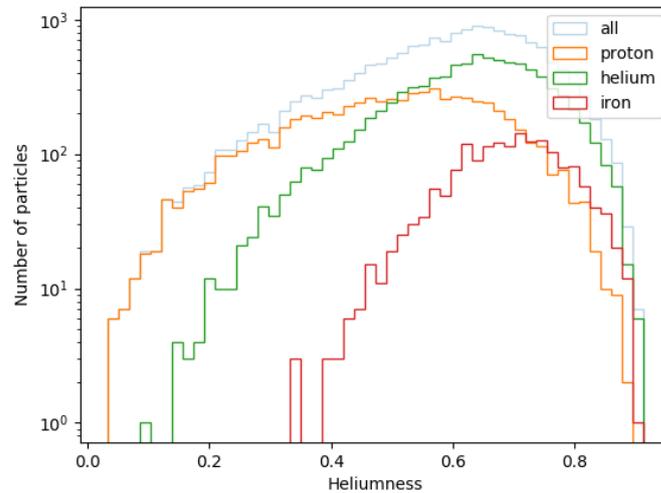


Figure 2: Distribution of Heliumness of the remaining particles after the Ironness cut.

The separation of protons and helium is more challenging than the separation of heavy and light particles, since the shower images can look very similar. The classifier identifies iron more like helium than proton particles. Since the aim of this work is a proton sample, this result is, even if not surprising, pleasing.

For the separation of proton and helium, as well as the surviving iron particles, a cut of 0.4 is made in the Heliumness. In Table 1 the number of particles surviving the cuts is shown. The majority of the simulated proton sample of 6993 particles survives the cut in the Ironness. A percentage of 97.28 % were estimated by equal or higher than 0.4 in the Ironness. For helium the classification leads to 8759 particles surviving the Ironness cut, which is 90.28 % of the simulated data set. A high amount of iron particles (66.64 %) is discarded after the classification with the first random forest. 37.36 % of the iron nuclei survive the cut but most of them are discarded by the following cut in the Heliumness. In total 26.84 % of the protons, 5.29 % of the helium and 0.11 % survive the cuts of the analysis.

	simulated particles	particles surviving Ironness cut	particles surviving Heliumness cut
proton	6993	6803 (SI97.28%)	1877 (SI26.84%)
helium	9656	8759 (SI90.71%)	511 (SI5.29%)
iron	4601	1719 (SI37.36%)	5 (SI0.11%)

Table 1: Particles surviving the first cut (in Ironness) and both cuts, the first cut (in Ironness) and the second one (in Heliumness). Percentage of simulated particles after the cuts.

Correction of the Data Samples

As shown in the previous section the classification of particles does not lead to pure proton sample and the random forest cannot be significantly improved by hyper parameter tuning. Anyhow, a correction of the data samples can be done by the following formula exemplified by the dataset of events classified as protons.

$$N_{\text{protons}} = \frac{(p_{2,2} - p_{2,3})(\epsilon_1 - p_{1,3}) - (p_{1,2} - p_{1,3})(\epsilon_2 - p_{2,3})}{(p_{2,2} - p_{2,3})(p_{1,1} - p_{1,3}) - (p_{1,2} - p_{1,3})(p_{2,1} - p_{2,3})} \frac{N}{\epsilon}$$

p_1, i : Probability to classify a particle i as proton with proton-helium-classifier

p_2, i : Probability to classify a particle i as proton with proton-iron-classifier

$i = 1, 2, 3$: proton, helium, iron

The correction must be done dependently on the energy, the zenith angle of the telescope pointing and the cosine of the proton's arriving angle. The error of the number of protons N_{protons} can be propagated. In this case it is limited by the monte carlo statistics.

References

- [1] Aleksic, J, et al., The major upgrade of the MAGIC telescopes, Part II: A performance study using observations of the Crab Nebula. *Astroparticle Physics* 72, pp. 76-94.
- [2] Nöthe, M., Brügge, K. and Buß, J., *aict-tools* - Reproducible Artificial Intelligence for Cherenkov Telescopes. <https://github.com/fact-project/aict-tools>.
- [3] Particle Data Group, Patrignani, C. et al., Review of Particle Physics - 29. Cosmic Rays. *Chin. Phys. C*, 40, 100001 (2016) and 2017 update. <http://pdg.lbl.gov>.

Towards Applied Federated and Distributed Learning

Danny Heinrich
Artificial Intelligence Group, Chair 8
Technische Universität Dortmund
danny.heinrich@tu-dortmund.de

With an increasing number of small and smart devices, data collection and processing becomes an increasingly challenging task of growing importance. Distributed data sources continue to accelerate the rate at which the data is collected, while also generating more diverse data. It is becoming more and more difficult to extract knowledge as limited bandwidth and increased privacy concerns limit our access to the raw data. To alleviate this issue we employ distributed and federated machine learning techniques on the devices themselves. Instead of computing a central machine learning model, which is deployed on each device, we learn machine learning models on the edge without moving the data from the device. This report discusses and briefly outlines my research and progress on the theoretical and practical applications of distributed learning, model aggregation, and resource-efficient communication as part of the Collaborative Research Center (CRC) 876.

1 Introduction

Smart and small devices equipped with multiple sensors have become more commonly available in our private lives as well as in the industry. The amount of data collected by these devices has increased tremendously over the past few years and is expected to grow even further in the years to come. As the number of devices grows, decreasing the upkeep costs for these devices has become a crucial part of the development process. While individual devices do not contribute much to a total budget, multiple such devices quickly accumulate high upkeep costs if not carefully selected and maintained. Hence,

ultra-low-power devices and energy efficiency are at the forefront for developing new and more efficient algorithms for smart devices.

Large-scale scientific applications include the Cherenkov Telescope Array (CTA) for high energy gamma ray detection, with two construction sites in Chile and La Palma. In logistics applications distributed learning can be used for automated warehousing, e.g., for controlling shipping processes such as moving payloads, using drones and robots. We investigate how distributed and federated learning applications in these scenarios are able to efficiently reduce communication cost and energy footprint.

2 Distributed and Federated Machine Learning

Distributed learning poses significant additional challenges on top of the already demanding machine learning tasks. Given a set of distributed devices we are typically facing additional constraints such as limited bandwidth and communication costs. Reducing the amount of bandwidth used has two major advantages. It reduces the overall network load and therefore the energy footprint of the overall system as less communication operations are frequently related to an overall less active system. Furthermore, it is crucial to consider privacy constraints in applications where data is gathered from private and personal devices. If the collected data contains private and sensitive information, we require the data to remain on these devices. However, we still may want to incorporate some of this information in our models to optimize processes on the devices themselves such as battery usage, process speedup or information exchange between devices in the network.

During my Master's Thesis we investigated and developed aggregation mechanisms for distributed models using only model parameters or meta-information without explicit requirement to transfer local data [1]. We explored different solutions for model aggregation in a federated learning environment. Federated learning was recently introduced by McMahan et.al. [3, 4] as a variant of distributed learning where devices communicate exclusively with a central coordinator node. The coordinator usually handles requests and exchanges between devices and updates local parameters if necessary. Furthermore, when interested in a global model, the coordinator aggregates the local models received from messages to a single global model.

Instead of transmitting the full data to a central server, we only communicate the model or its parameters, which already reduces communication costs significantly. In some cases enabling communication between distributed devices leads to better results as local information can be used to augment the models of other devices. Hence, we explored the possibility of allowing limited communication between distributed devices, exchanging parameters to and evaluating the model on previously unseen data. The feedback is then

again just the performance metric used to score each model, eliminating the requirement of data transmission.

Overall, our findings suggest that due to the asymptotic properties of exponential family parameters, the arithmetic mean is the true maximum likelihood estimator. More specifically, each parameter vector from a distributed device is just a sample from a normal Gaussian distribution around the true parameter vector with some variance, given that the data is identically and independently distributed.

Furthermore, we not only explored aggregation techniques, but we also investigated reducing the amount of communication required in a federated learning environment by actively choosing when to aggregate models. Instead of sending the model parameters at every update we tied communication to a set of conditions. We compared different communication strategies [2] based on local deviation from the global standard or sufficiently large updates on the local devices. By storing parameters as fixed-width integers [5] we were able to decrease the communication cost further. Naturally, using fixed-width integers in an approximation of the optimal floating point model parameters. However, we have shown that this error is bounded based on the maximum distance between optimal floating point and integer model parameters. Limiting models to be integer only additionally eliminates the requirement for floating-point architecture decreasing the overall architecture cost. Our results suggest a performance and bandwidth tradeoff, i.e., updating less often and using integer parameters decrease performance but reduces communication costs by several orders of magnitude.

3 Distributed Learning in Automated Warehousing

On the basis of our theoretical findings and experimental results, we want to further explore practical applications of distributed and federated machine learning. Together with members of the A4 project and the Fraunhofer IML, we started preparing a machine learning framework for a distributed Wireless Sensor Network (WSN) for indoor warehouse localization and distributed sensing. `SENSOR FLOOR` is a 345 node WSN primarily used for distributed machine learning and distributed sensing application prototyping. The full hardware and software stack includes data collection, data processing, feature exploration, feature extraction and distributed machine learning on the derived features. The sensors provide updates in an asynchronous fashion allowing for more flexible and independent data collection, while posing an additional challenges of data and time-series matching. As the sensor data does not contain positional information about the robots we use an external 3D-tracking system to gather the labels for our machine learning problem.

Moving towards an application for distributed machine learning we started by performing preliminary experiments on the full data collected onto a centralized server. Our initial

results show that we can indeed reliably localize robot positions based on data obtained from distributed data sources.

4 Future Research

Our research focus opens up a variety of possible applications for distributed learning on edge devices. First, we want to investigate possible methods to further decrease the amount of communication required in distributed learning. Reducing communication cost usually comes at the cost of decreased performance. However, initial results show that model parameters in federated learning may exhibit structural redundancies, which can be exploited to further reduce the communication cost required for local and global updates. If our first experiments are promising, we plan to apply our methods to practical applications within the CRC. First, our cooperation with project A4 resulted in a novel dataset for distributed sensing and localization, which is an inherently distributed task. Within the scope of this project our goal is to accelerate the real-time localization performance on the distributed sensing network, while preserving optimal performance. Second, within project C3 we have access to simulation data from the CTA project, where the telescopes, which are simultaneously the data sources, are distributed. Here, an on-site distributed algorithm for data-reduction or possibly preliminary inference could help reduce the amount of data transferred significantly.

References

- [1] Danny Heinrich. Aggregation mechanisms for distributed exponential family models. Master's thesis, TU Dortmund University, 2020.
- [2] Lukas Heppe, Michael Kamp, Linara Adilova, Nico Piatkowski, Danny Heinrich, and Katharina Morik. Resource-constrained on-device learning by dynamic averaging. 2020.
- [3] Jakub Konečný, H Brendan McMahan, Felix X Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated learning: Strategies for improving communication efficiency. *arXiv preprint arXiv:1610.05492*, 2016.
- [4] H Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, et al. Communication-efficient learning of deep networks from decentralized data. *arXiv preprint arXiv:1602.05629*, 2016.
- [5] Nico Philipp Piatkowski. *Exponential families on resource-constrained systems*. PhD thesis, 2018.

Neural Network based Cascade Event Selection

Mirco Hünnefeld
Lehrstuhl für Experimentelle Physik E5b
Technische Universität Dortmund
mirco.huennefeld@tu-dortmund.de

IceCube is a neutrino detector located at the geographic South Pole, instrumenting a cubic kilometer of glacial ice. A major goal of IceCube is the detection of astrophysical neutrino sources. Neutrino interactions are detected via Cherenkov radiation of charged secondary particles. The two main detection channels consist of charged muon-neutrino interactions, so called tracks, and cascade events, which result from all other neutrino interactions. Although the selection and angular reconstruction of cascades is challenging, these events enhance IceCube's capabilities to probe the southern neutrino sky. In this paper, a machine learning-based cascade event selection is presented. The event selection utilizes a series of convolutional neural networks and gradient boosted decision trees. The presented event selection improves upon the performance of previous selections, while greatly reducing the necessary computation time, enabling the application in real-time.

1 Cascade Events in IceCube

The IceCube Neutrino Observatory [1] indirectly measures neutrino interactions through the emitted Cherenkov radiation of charged secondary particles. Depending on the interaction type, there are two main detection channels consisting of tracks, induced by charged current muon-neutrino interactions, and cascade events, which are almost spherical energy depositions. Cascade events are difficult to reconstruct, as the resulting particle shower appears almost point-like, given IceCube's inter-string spacing of 125 m [2]. Nevertheless, these events improve IceCube's sensitivity in the southern neutrino sky [3].

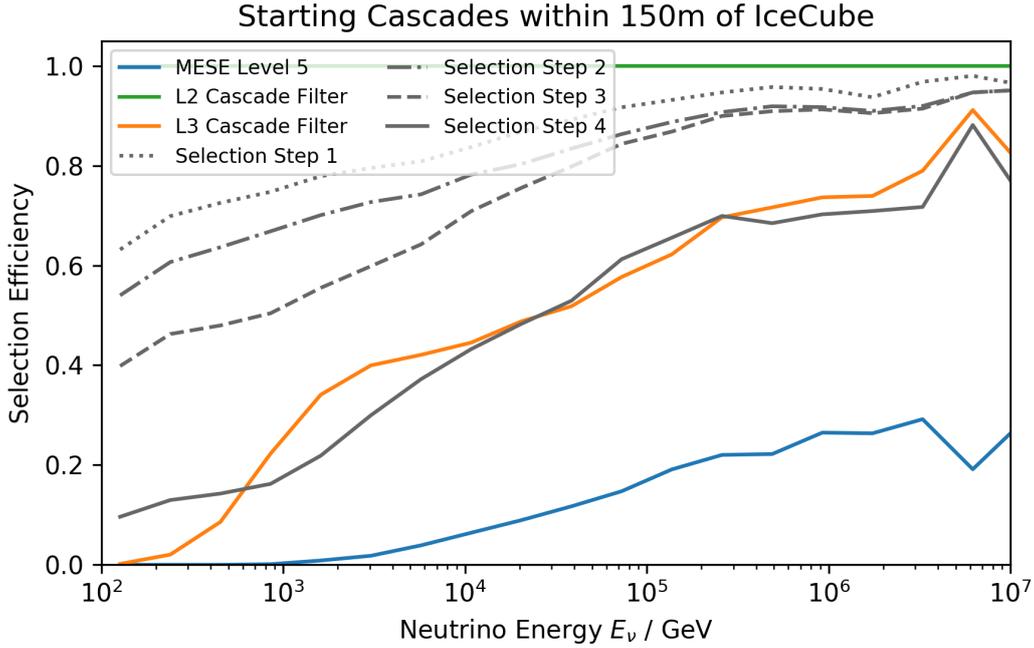


Figure 1: The signal efficiency of the four selection steps are compared to the initial data level (L2 Cascade Filter), to an intermediate step (L3 Cascade Filter) and final level (MESE Level 5) of the previous cascade selection utilized in Ref. [3].

A more efficient event selection in combination with an improved cascade reconstruction [4] may further enhance IceCube’s sensitivity towards potential neutrino emission from the galactic plane. In the following section, a new cascade event selection is presented that heavily relies on machine learning to increase the signal efficiency while reducing atmospheric muon background contamination.

2 Neural Network based Event Selection

In order to increase the signal efficiency of the cascade event selection, a series of convolutional neural networks (CNNs) [5] is applied. Due to the high class imbalance of signal versus background on the order of $1 : 10^6$, it is beneficial to separate out the classification task into multiple steps. Therefore, a variety of CNNs is trained at different selection levels. This allows initial CNNs to classify and discard the dominant background events. Subsequent CNNs can then focus on the more challenging events. In addition to an improved performance, this staged selection procedure also reduces computation time, because events are discarded as early as possible. The utilized CNNs are trained with a binary cross entropy loss. Further details are provided in Refs. [6, 7].

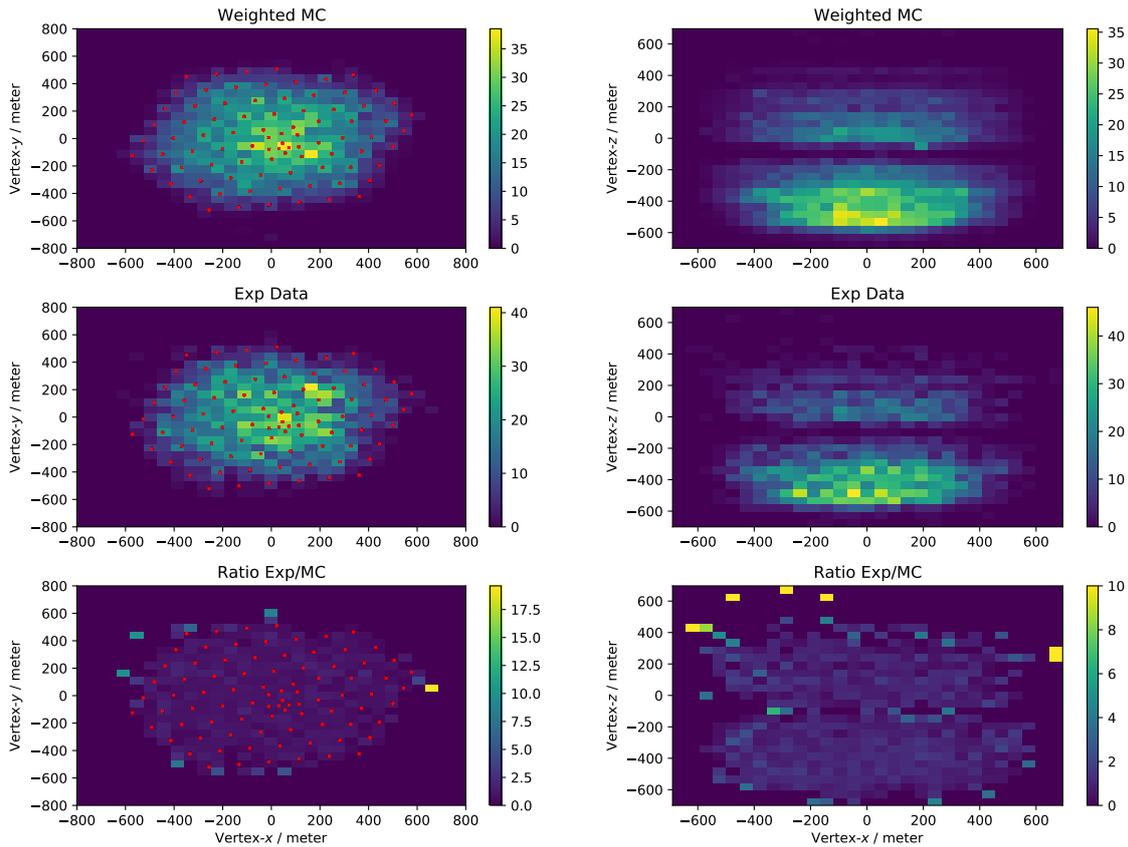


Figure 2: The distribution of neutrino interaction vertices of selected events are shown for the x - y -plane on the left and for the x - z -plane on the right. From top to bottom, the illustrated distributions show the simulated Monte Carlo (MC) events, the experimental data and the ratio of experimental data and MC events. Detector strings are indicated by red dots.

A total of four selection steps are defined. The first two selection steps consist of a total of three small CNNs that are used to filter out most of the dominant atmospheric muon background. These CNNs have a runtime of about 1 ms per event and are therefore well suited for the initial data reduction steps. After selection step 2, the data rate has been decreased by about 98 %. In selection step 3, further CNNs with larger network architectures are applied. Final level cuts are performed with the help of gradient boosted decision trees from the XGBoost library¹. The BDTs obtain the classification scores of the CNNs as input features, as well as reconstructed cascade quantities such as the neutrino interaction vertex, direction and energy.

The described event selection is able to retain a larger fraction of signal events, while providing a similar purity to existing event selections. The signal efficiency of the selection

¹<https://xgboost.ai/>

steps are illustrated in [Figure 1](#). In comparison to the previous cascade selection employed in Ref. [3], cascades with an interaction vertex close to or slightly outside of the IceCube detector are also included in this selection. The interaction vertices of selected events are shown in [Figure 2](#) for simulated and experimental data.

3 Conclusion

An improved event selection is presented that leverages machine learning to enhance the signal efficiency of cascade events in IceCube. This leads to an increased sensitivity to neutrino emission from the southern sky. Future IceCube neutrino source searches will benefit from the presented cascade selection.

References

- [1] IceCube collaboration, M. Aartsen et al., *The IceCube Neutrino Observatory: Instrumentation and Online Systems*, *JINST* **12** (2017) P03012 [[1612.05093](#)].
- [2] IceCube collaboration, M. Aartsen et al., *Energy Reconstruction Methods in the IceCube Neutrino Telescope*, *JINST* **9** (2014) P03009 [[1311.4767](#)].
- [3] IceCube collaboration, M. Aartsen et al., *Search for Sources of Astrophysical Neutrinos Using Seven Years of IceCube Cascade Events*, *Astrophys. J.* **886** (2019) [12](#) [[1907.06714](#)].
- [4] IceCube collaboration, M. Huennefeld, *Reconstruction Techniques in IceCube using Convolutional and Generative Neural Networks*, *EPJ Web Conf.* **207** (2019) 05005.
- [5] Y. LeCun et al., *Handwritten digit recognition with a back-propagation network*, in *Advances in Neural Information Processing Systems 2*, pp. 396–404, Morgan-Kaufmann (1990), <http://papers.nips.cc/paper/293-handwritten-digit-recognition-with-a-back-propagation-network.pdf>.
- [6] IceCube collaboration, M. Huennefeld, *Deep Learning in Physics exemplified by the Reconstruction of Muon-Neutrino Events in IceCube*, *PoS ICRC2017* (2018) 1057.
- [7] IceCube collaboration, R. Abbasi et al., *A Convolutional Neural Network based Cascade Reconstruction for the IceCube Neutrino Observatory*, [2101.11589](#).

Approaches to Unfolding the Neutrino Energy Spectrum

Leonora Kardum
Experimentelle Physik 5
Technische Universität Dortmund
leonora.kardum@tu-dortmund.de

The IceCube Neutrino Observatory is a cubic kilometer detector embedded in ice at the geographic South Pole, utilizing the properties of particles' passage through a medium to its detection. Due to its architecture and reconstruction algorithms, IceCube is capable of detecting particles up to PeV energies. A precise analysis of the high energy region is desirable with the initiative to differ among three predicted neutrino components. As the nature of neutrinos allows only detection through secondary interactions, the need for highly clean samples and accurate unfolding approaches arises.

1 Neutrino Energy Spectrum

Of high interest is the neutrino spectrum in the energy range between gigaelectronvolt (GeV) and petaelectronvolt (PeV) energies as it contains the atmospheric component, constituted of conventional and prompt, and the transition to an expectedly flatter astrophysical component. The conventional component arises from decays of kaons and pions and describes energies up to teraelectronvolt (TeV) while the higher energy regions are taken over by the prompt incurred by the decay of short lived charmed mesons. Both components are purely atmospheric whereas the analysis of PeV ranges aims at uncovering the astrophysical component predicted as a result of energetic sources in the cosmos. However, the component has not yet been fully described and was not connected to any known sources [1]. IceCube remains the only experiment to ever detect a neutrino in this range [2].

2 Unfolding approaches

Neutrinos interact with the ice of the detector producing secondary particles in the process. Resulting particles emit photons which create observable Cherenkov light if the particles move faster than photons in given medium. From the timestamps, intensity and position of the detected Cherenkov photons, the primary interaction can be reconstructed. Since we can only infer the energy of the primary particle from different observed and constructed features, the process is called deconvoluting the spectrum, or commonly named unfolding in astrophysics. Unfolding is done by training algorithms on faux Monte Carlo data to a certain response of the detector. From this, the response to an actual event can be described.

A model-dependant approach can be taken where the algorithm is trained to a previously known shape of the sought after distribution, called weighting, which greatly advances the accuracy of the prediction. However, this also mandates knowing the precise distribution in as much extent possible. A separate machine learning model was built in the sole purpose of further cleaning the data sample from background noise and muon particles that can disrupt the analysis. A purity of over 99.7 percent was achieved, higher than in previously known methods [3]. A seminal technique of treating features in a redistributed fashion enables a fair spread of statistics in the observable space. We take on a likelihood based approach in which the collective energy spectrum is derived from the values of measurements' distributions. The resulting spectra are sampled through Monte Carlo Markov Chains, in purpose of determining the most appropriate parameters that describe the given output. A strong regularization is required in high energy region sparsely populated by events. We have successfully unfolded the 125 GeV to 2 PeV energy distribution, in a singular zenith region covering the whole area of interest represented by angles from 86 to 180 degrees. Different regions of interest can be unfolded in the same manner. Nonetheless, the algorithm analyses energy spectra resulting from measurements, effectively erasing the information of single events' contributions.

In contrast, we further the development of a different approach realised through deconvolution of observed events respectively, from which the overall distribution can be derived. Here, a classifier of choice from a widely used library of tools is employed in the process, leaving the choice to the nature of the problem and the analysis wanted. Similarly, the algorithm is trained on the Monte Carlo data corresponding to a particular response. The unfolding of singular events depends on the chosen approach, whereas we mostly focused ours on probabilistic-based models due to their soundness in later physical interpretation. We treat the sought after quantity as as a set of unrelated classes representing different energy ranges that together constitute the full range of interest [4]. This reduces the problem to classification. When presented with the training sample from the Monte Carlo, the distribution is flattened out essentially erasing dependence on the weighting, which is then updated on subsequent steps in this iterative algorithm.

Updating density improves estimate with every step, but calls for a detailed treatment of its convergence and regularization. Regularization is also handled through iteration, and this effect is yet to be quantified. Each of different classification tools have to be treated individually since their predictions are based on different equations. We recently finished a thorough analysis of the probabilistic-based model and the effects of the iteration on it.

Current work focuses on the comparison of different approaches, and their advantages (or disadvantages) in practice. The likelihood based approach showed to be unparalleled in statistical analysis of the sample, and the meticulous treatment of regularization parameters. The second approach proved useful especially in density independent samples, usually with comparable statistics in all regions.

3 Outlook

Prospective work includes applying the developed models to real data detected with IceCube and evaluating the difference. The analysis of the event-based approach is to be extended to other models, possibly even non-probabilistic ones which require special manner. The focus is to ensure reproducibility on coming future data without the need of our severe intervention.

References

- [1] Aartsen, M. Abraham, K. Ackermann, M. Adams, J. Aguilar Sánchez, Juan Antonio Ahlers, M. Ahrens, M. Altmann, David Anderson, Travi Archinger, M. Arguelles, C. Arlen, Timothy Auffenberg, J. Bai, Xiaoqiong Barwick, S. Baum, V. Bay, Rabia Beatty, J. Tjus, J. (2015). The IceCube Neutrino Observatory - Contributions to ICRC 2015 Part II: Atmospheric and Astrophysical Diffuse Neutrino Searches of All Flavors. e-print. arXiv:1510.05223 [astro-ph.HE].
- [2] Dawn R. Williams. (2019) Results from IceCube. arXiv: 1909.05173 [astro-ph.HE]
- [3] Borner, M. (2018) Bestimmung des Energiespektrums von atmosphärischen Myon-neutrinos mit 3 Jahren Daten des IceCube-Detektors. doi: 10.17877/DE290R-19089
- [4] Bunse, M. (2018) DSEA Rock-Solid. PhD diss., Technische Universität Dortmund.

Multiwavelength Analysis of 3C 84/NGC 1275

Lena Linhoff

Experimentelle Physik 5

Technische Universität Dortmund

lena.linhoff@tu-dortmund.de

The elliptical galaxy 3C 84 ($z = 0.017$) is located at the Perseus cluster and one of the closest and brightest radio galaxies. In 2008, its counterpart in the gamma-ray regime, NGC 1275 was detected by *Fermi*-LAT [1]. Because of its proximity, it has been observed and studied quite well over years with different ground- and space-based detectors and telescopes. The fact, that we measure both gamma and radio emission and are able to distinguish different radio emission regions within the source opens unique opportunities. Since the acceleration mechanisms are unclear, TeV radio galaxies like 3C 84 are perfect candidates to study these mechanisms. To restrict the possible gamma-ray emission region, I calculate the optical absorption depth of high-energy photons in the broad-line region of 3C 84, depending on their energy and distance to the central black hole. Based on these calculations, a lower limit on the distance of the emission region from the central black hole can be derived. A paper on this research was published [6].

In a second project, I aim to provide the first long-term gamma-ray lightcurve of NGC 1275 at very-high energies. To achieve this goal, 10 years of data obtained by the Major Atmospheric Gamma-Ray Imaging Cherenkov Telescopes (MAGIC) have to be analyzed. This is not possible with the current status of the analysis software and workflow used by MAGIC. For this reason, I am developing an automatized framework to reduce human interaction and produce stable and reproducible analysis results.

Restricting the Gamma-Ray Emission Region in 3C 84

Based on calculation regarding the optical depth of the broad-line region (BLR) of an active galactic nucleus (AGN) by Finke [5], we are able to exclude some regions within the core of the AGN as gamma-ray emission regions. For this purpose, I combined data measured by *Fermi*-LAT and MAGIC to a joined spectral energy distribution (SED) and fitted this data with a log-parabola function modified with an absorption term, taking the optical depth of the broad-line region into account. The data were measured during two flux states of the source, since MAGIC detected a gamma-ray flare on the first of January 2017 [4]. The low flux state was measured from 2016, October to 2016, December, and the flux state in the first days of January 2017. Since the geometry of the BLR is not known, the model is obtained for two approaches: A spherical shell and a flattened ring geometry.

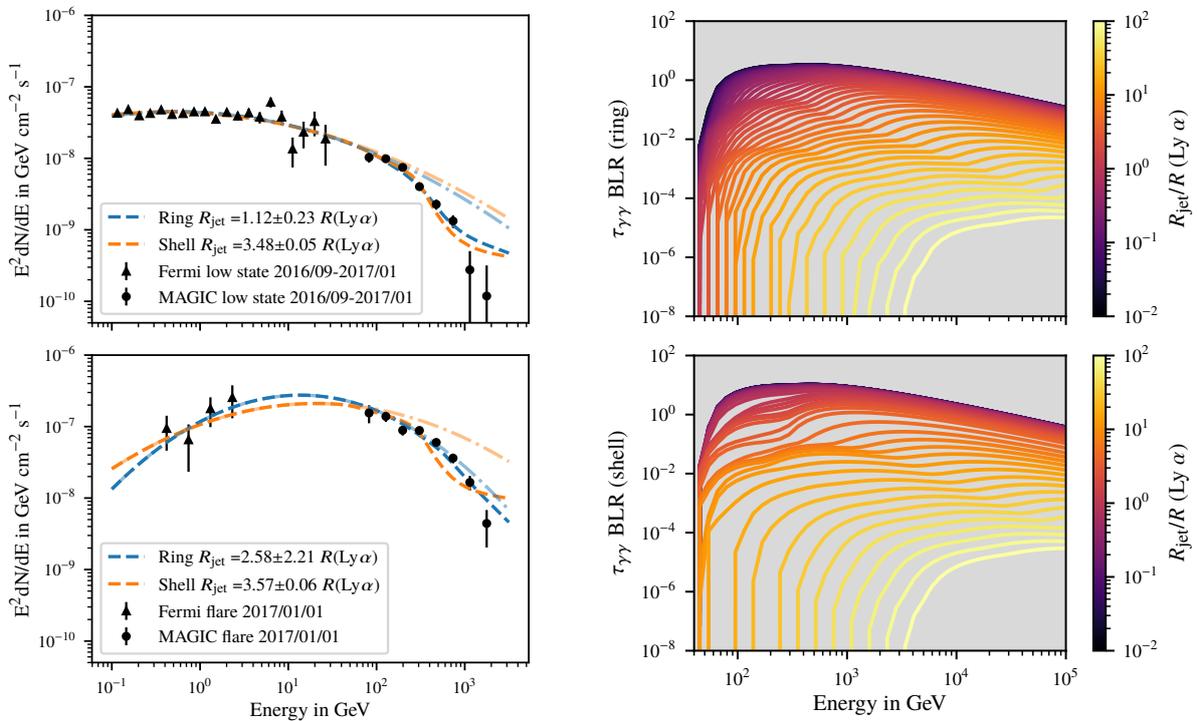


Figure 1: Left: Combined SED with *Fermi*-LAT and MAGIC data, fitted with a log-parabola function, modified with an absorption term. Right: Optical depth of the broad-line region, depending on particle energy and distance of the emission region from the center of the source.

With the results obtained from the SED fit, I can restrict the gamma-ray emission region. In the case of the shell geometry, the emission takes place outside the BLR. For the ring geometry, the emission region can be located at the outer edge of the BLR in the case of the low state. For the flaring state, no conclusive result was found due to the big

uncertainties.

This result challenges all theoretical models placing the gamma-ray emission region close to the black hole. Mechanisms producing very-high-energy gamma-ray emission further downstream an AGN's jet seems more likely.

autoMAGIC

MAGIC is a ground-based Cherenkov telescope, located at the Canary Island of La Palma. Since 2009, it observes NGC 1275 on irregular distributed time slots. The current standard analysis for MAGIC data deals with a huge amount of human interaction and is not suitable for large data sets [2]. Because NGC 1275 is a weak source, and another object (IC 310) appears in the field of view, the data analysis of this source needs special care. Furthermore, NGC 1275 was measured during moon nights, which makes the data reduction even harder. This is the reason, why no long-term lightcurve could be created with the MAGIC data in the last years. Having this information would open one a great deal of new analysis possibility, since long-term lightcurves exist for several other experiments and energy ranges, partly they span over decades. Correlating and comparing them with each other would give us new hints, what mechanisms take place in the source and which emission is produces in the same processes.

To create such a long-term light curve, one needs an analysis framework which can cope with the different observation conditions, data was taken under. Additionally, we aim for reproducible analysis results. For this reason, the autoMAGIC project was born, which aims to automatize and simplifies the data reduction for the whole MAGIC community and all data ever taken by the MAGIC telescopes. In the last years, a new standard data format for gamma-ray astronomy was developed, which is and will be used by all currently operating telescopes [3]. With our autoMAGIC project we can now produce this data format at so-called level 3, and make sure that our DL3 analysis results can be used as input for the current and future high-level analysis tools.

By the Port d'Informació Científica (PIC) in Barcelona, we are equipped with a cluster and a database system, which allows us to produce DL3 files in an automatized manner. For this task, a database was created with over 20 tables modeling the data reduction process. The database takes care of job monitoring, as well as the detector and simulations specific parameters. autoMAGIC is not restricted to my scientific topics and will be usable for the whole MAGIC collaboration.

In the near future, I will start to analyze data from NGC 1275 taken by MAGIC and combine my results with observations of other facilities to a multiwavelength study.

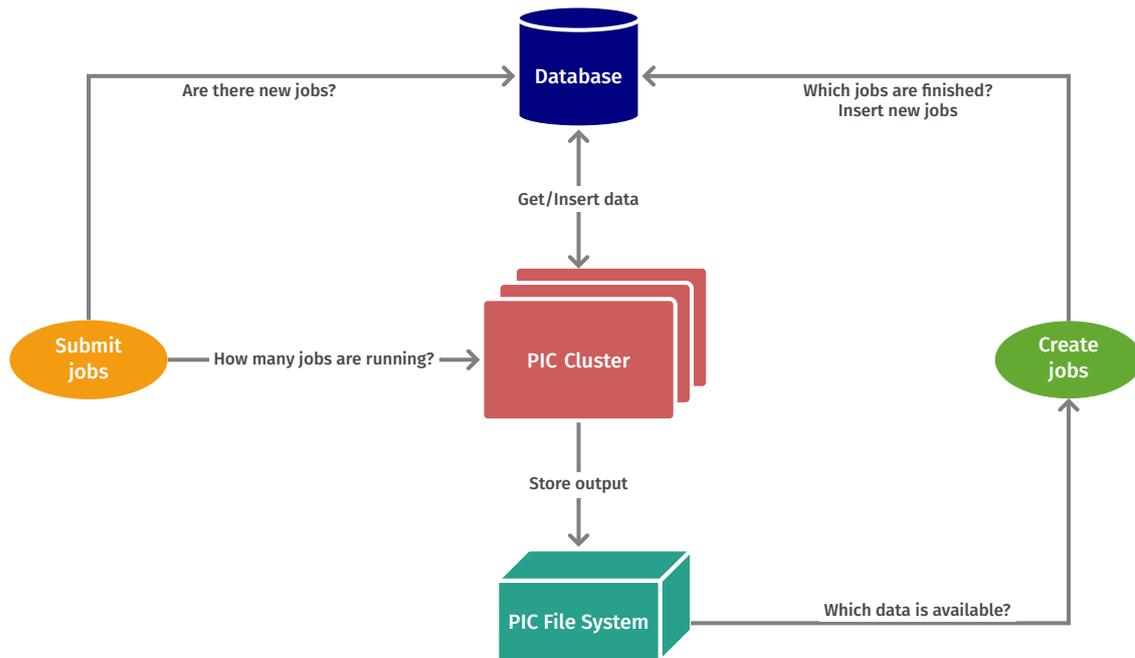


Figure 2: Production scheme of the autoMAGIC project.

References

- [1] Aous A Abdo, et al. Fermi discovery of gamma-ray emission from NGC 1275. *The Astrophysical Journal*, 699(1):31, 2009.
- [2] J Aleksí et al. The major upgrade of the MAGIC telescopes, Part II: A performance study using observations of the Crab Nebula. *Astroparticle Physics*, 72: 76-94, 2016.
- [3] C Deil et al. Open high-level data formats and software for gamma-ray astronomy. *AIP Conference Proceedings*. 1792(1), AIP Publishing LLC, 2017.
- [4] S Ansoldi, LA Antonelli, C Arcaro, D Baack, A Babić, B Banerjee, P Bangale, U Barres de Almeida, JA Barrio, J Becerra González, et al. Gamma-ray flaring activity of ngc 1275 in 2016-2017 measured by magic. *arXiv preprint arXiv:1806.01559*, 2018.
- [5] Justin D Finke. External compton scattering in blazar jets and the location of the gamma-ray emitting region. *The Astrophysical Journal*, 830(2):94, 2016.
- [6] L Linhoff, A Sandrock, M Kadler, D Elsässer, and W Rhode. Excluding possible sites of high-energy emission in 3C 84. *Monthly Notices of the Royal Astronomical Society*, 500(4):4671–4677, 2020.

Crab nebula observations with the Large Sized Telescope

Lukas Nickel

Experimentelle Physik 5

Technische Universität Dortmund

lukas.nickel@tu-dortmund.de

With the first prototype telescope of the Large Sized Telescope (LST-1) constructed for the Cherenkov Telescope Array (CTA), the current focus is on understanding every aspect of the new telescope and performing first studies. Mismatches between expectations and measurements have to be understood and incorporated into the new analysis pipelines in order to pass the critical design review and prepare for operation.

1 Status of CTA and the LST

The Cherenkov Telescope Array (CTA) is going to be the next major experiment in the field of ground-based gamma-ray astronomy. With two arrays composed of three different telescope types, it is expected to extend the observable energy range while at the same time improving on the sensitivity compared to existing experiments. Ultimately, an energy range of 20 GeV to 300 TeV is to be achieved [5].

The first constructed telescope, the LST-1, was inaugurated in October 2018 [3]. Starting with in November 2019 [2], multiple observations of the Crab Nebula and other sources have been performed, also detecting the crab pulsar [1].

Since the LST-1 is so far operating in monoscopic mode, the analysis methods differ slightly from the stereoscopic methods, that are likely to be used in later stages of observation.

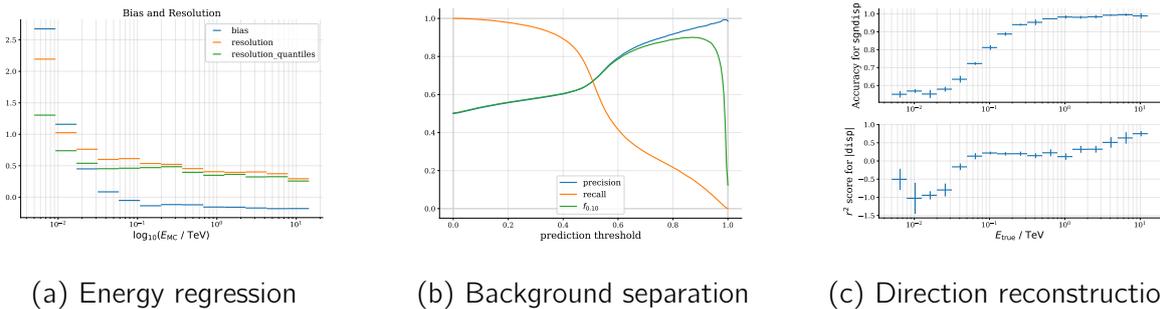
2 The Crab Nebula

The Crab Nebula is a supernova remnant located approximately 2000 parsecs away from the earth. It is the first detected very high energy source in gamma-ray astronomy [6] and, due to its high and steady flux at very high energies, is often used as a standard candle (e.g. [4]) for point source analyses.

3 Reconstruction on a per-event basis

Starting at the level of image parameters after raw data has been processed and camera images extracted, the next analysis steps consist of a background separation, energy regression and reconstruction of the event direction. This marks the step from data level 1 to data level 2. All of these tasks are performed using random forest models, that are trained on simulated data using the image parameters from data level 1.

Figure 1 shows performance metrics for each of the tasks obtained via cross-validation on the training datasets.



(a) Energy regression (b) Background separation (c) Direction reconstruction

Figure 1: Cross-validated metrics for the machine learning tasks at data level 2.

The next data level, data level 3, combines individual events to sets of observations including an estimation of the instrument response to draw conclusions on the real distributions given the measured data. Based on the predicted event properties, an event selection is performed with the goal to eliminate most of the hadronic background present in gamma-ray astronomy and achieve the highest possible sensitivity. The instrument response is commonly split into three separate contributions: The effective area as a measure of the portion of measured events, the energy migration as a measure of the difference between estimated and true event energy and the point spread function, which quantifies the spread of the direction reconstruction.

The instrument response after the optimized event selection is presented in figure 2.

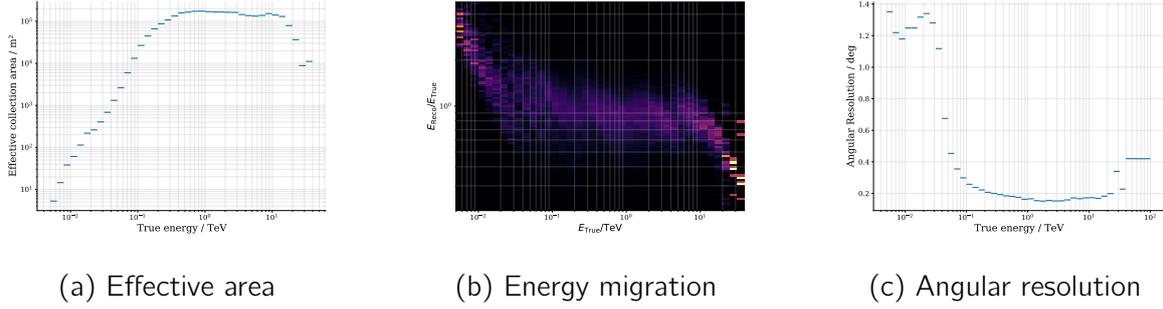


Figure 2: Instrument response functions

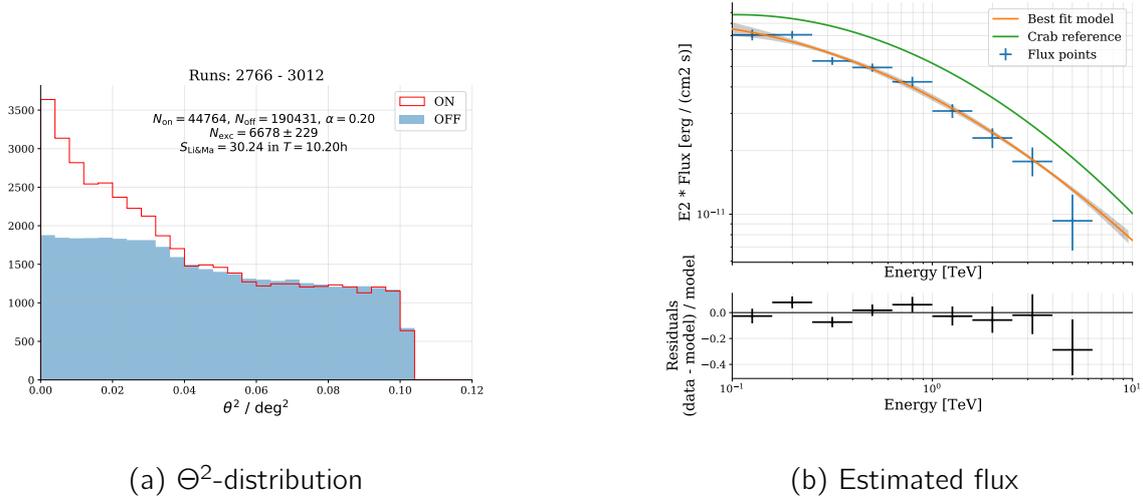


Figure 3: High level results for the observations

Analyzing the events, that pass all selection criteria, an excess at the assumed source position compared to background regions at equivalent positions in the field of view is present (Figure 3a) and an energy spectrum can be estimated from the remaining signal counts (Figure 3b). To estimate the spectrum, a log-parabola function, which is a widely used function for the flux of the crab nebula in the observed energy range, is used. The resulting flux is lower than than the reference of the MAGIC telescopes [7]. This behaviour has been observed with multiple, independent analyses in the working group and is likely connected to mismatches in the simulations regarding the optical efficiency of the telescope.

4 Outlook

As short-term goal the various efforts of analysing LST-1 data need to be brought in line with the main CTA development. New simulations are expected to solve some of the

issues in current analyses. Additional observations will be performed and observations of other sources are already present and can be analysed.

Apart of the LST, a wide range of longer term goals are getting investigated inside the CTA consortium, including, amongst others, the search for dark matter via possible gamma ray emission.

References

- [1] The Cherenkov Telescope Array Consortium. Cta prototype lst-1 detects very high-energy emission from the crab pulsar.
- [2] The Cherenkov Telescope Array Consortium. Large-sized telescope prototype records its first light.
- [3] Juan Cortina. Status of the large size telescopes of the cherenkov telescope array. *arXiv preprint arXiv:1907.10146*, 2019.
- [4] J. Aleksy et al. Performance of the magic stereo system obtained with crab nebula data. *Astroparticle Physics*, 35(7):435 – 448, 2012.
- [5] Daniel Mazin. The cherenkov telescope array. *arXiv preprint arXiv:1907.08530*, 2019.
- [6] T. C. Weekes, M. F. Cawley, D. J. Fegan, K. G. Gibbs, A. M. Hillas, P. W. Kowk, R. C. Lamb, D. A. Lewis, D. Macomb, N. A. Porter, P. T. Reynolds, and G. Vacanti. Observation of TeV Gamma Rays from the Crab Nebula Using the Atmospheric Cerenkov Imaging Technique. 342:379, July 1989.
- [7] Roberta Zanin, Daniel Mazin, Emiliano Carmona, Pierre Colin, Juan Cortina, Tobias Jogler, Stefan Klepser, Abelardo Moralejo, Julian Sitarek, Dieter Horns, et al. Magic measurement of the crab nebula spectrum over three decades in energy. *arXiv preprint arXiv:1110.2987*, 2011.

Uncertainty Estimations for Image Reconstruction in Radio Interferometry

Kevin Schmidt

Experimentelle Physik 5

Technische Universität Dortmund

kevin3.schmidt@tu-dortmund.de

Radio interferometry enables high-resolution imaging of astrophysical objects. For scientific analysis, cleaning of the observational data is necessary. The available software packages are suitable to provide reconstructed source images. Nevertheless, they are not capable of providing uncertainty estimations for the obtained images.

Using Deep Learning algorithms for the cleaning task permits fast and reproducible analysis strategies. The choice of the loss function can help to create additional uncertainty maps besides the reconstructed source images. This report illustrates how such an uncertainty estimation can be implemented using Convolutional Neural Networks.

Image Reconstruction and Uncertainty Estimation

In radio interferometry, correlated signals of antenna pairs are used to acquire information about astrophysical sources. Measured values correspond to complex visibilities in Fourier space. During observations, sparse antenna layouts lead to incomplete sampling. [2] Figure 1 illustrates this sampling with the help of a simulated observation. Only the information of colored pixels is available. The values of all grayed out pixels are lost in the sampling process during the observation. A direct application of the inverse Fourier transformation to these incomplete spectra results in noisy source images.

In this work, a Convolutional Neural Network reconstructs the missing values in the incomplete Fourier spectra. For model training, a simulated data set based on radio galaxies consisting of Gaussian components is created. The images of the radio galaxies

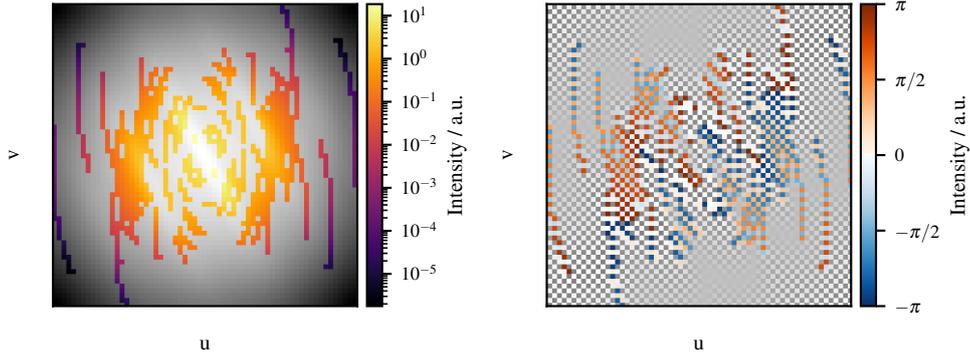


Figure 1: Amplitude (left) and phase (right) of a simulated radio galaxy. Grayed out pixels are lost in the sampling process during an observation with a radio interferometer.

are Fourier transformed. Afterward, a sampling comparable to an observation by a radio interferometer is performed. Sampled spectra serve as model input, while complete spectra serve as reconstruction targets. The utilized architecture is comparable to the one proposed in [1]. Here, residual blocks based on the layout described in [3] are applied. For simplification, missing pixel values are interpolated with nearest-neighbor interpolation before the input images get passed to the network.

A loss function based on a Gaussian likelihood is used to get an uncertainty estimation for each predicted pixel. Starting from a two-dimensional Gaussian distribution calculating the negative log-likelihood leads to

$$-\mathcal{L} = 2 \log(\sigma) + \frac{(x - \mu)^2}{\sigma^2}. \quad (1)$$

This function is minimized in the training process. Here, the network predicts two values for each pixel, a mean value μ and an uncertainty σ . In this work, amplitude and phase were passed to a single model using two different channels. The model was trained for 200 epochs with a learning rate of $1 \cdot 10^{-3}$.

Figure 2 visualizes the results for the amplitude reconstruction. It displays the reconstructed amplitude (top left), the true amplitude (top right), the difference between both (bottom left), and the predicted uncertainty (bottom right). Only small deviations between predicted and true spectrum are visible. Some reconstruction problems occur in the central part of the image. These parts are easy to identify because the predicted uncertainty for these pixels is higher than for the rest.

Figure 3 illustrates the results for the phase reconstruction. Again, the reconstructed image matches the true phase distribution well. Deviations appear for the structures in the central part. Also for the phase, the more problematic areas are easy to identify as the predicted uncertainties are higher for these pixels.

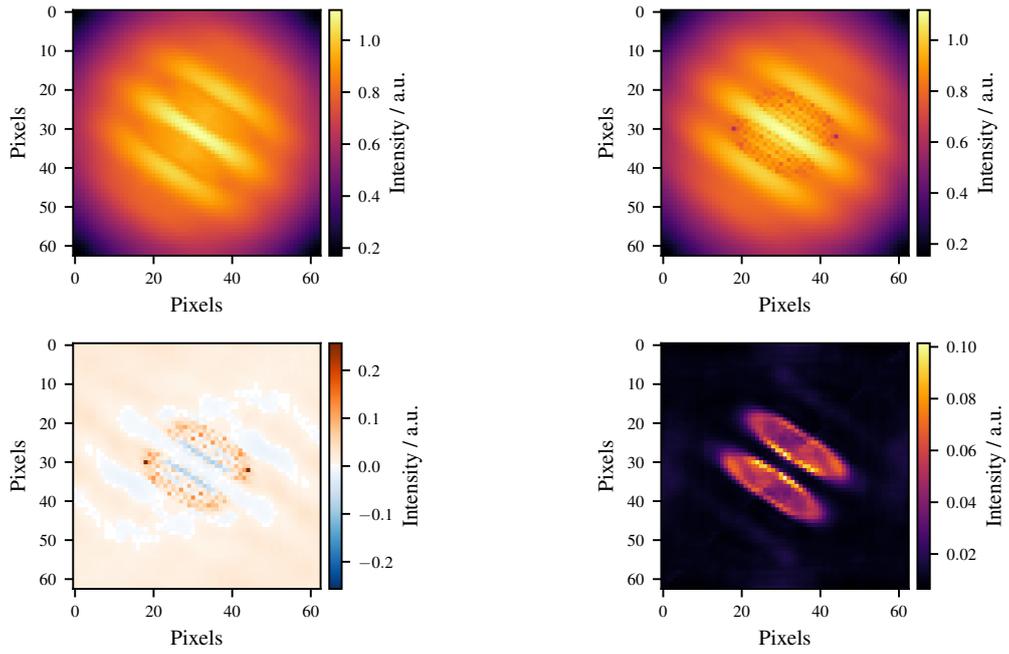


Figure 2: Reconstructed amplitude (top left), true amplitude (top right), difference between both (bottom left), and predicted uncertainties (bottom right).

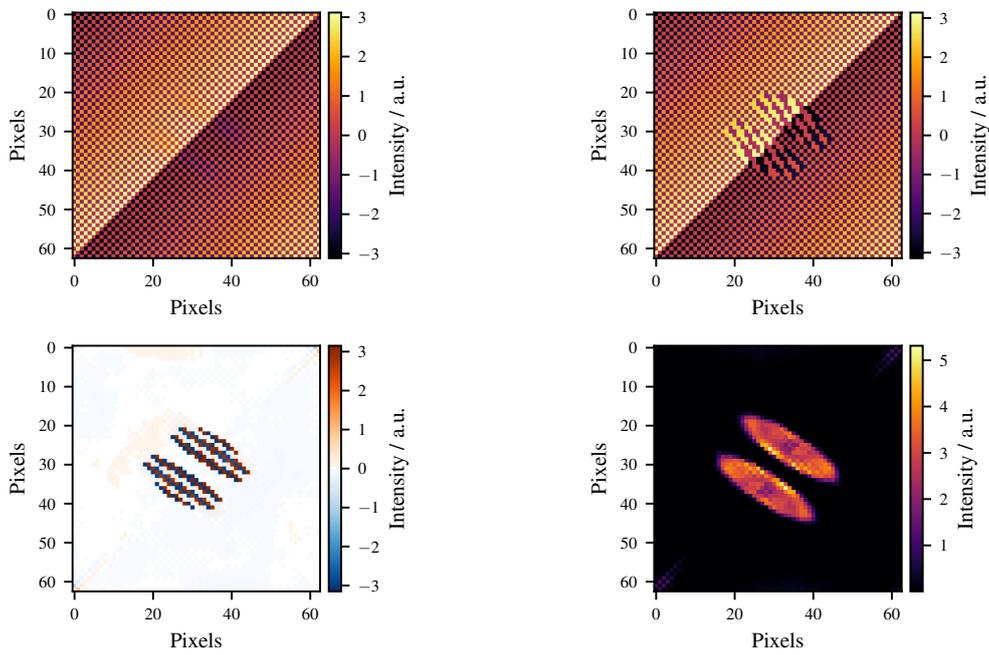


Figure 3: Reconstructed phase (top left), true phase (top right), difference between both (bottom left), and predicted uncertainties (bottom right).

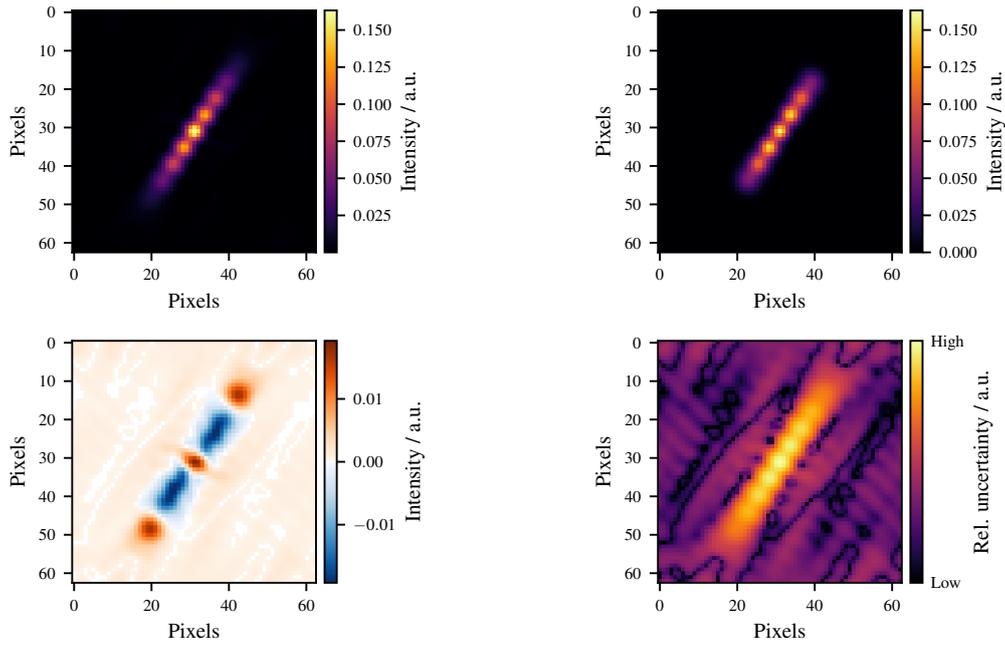


Figure 4: Reconstructed flux distribution (top left), true flux (top right), difference between both (bottom left), and calculated rel. uncertainty map (bottom right).

Predicted μ and σ enable creating Gaussian distributions for each pixel. These distributions are used to sample 10 000 versions of the amplitude and phase spectra. Afterward, these different versions are used to create 10 000 reconstructed source images utilizing the inverse Fourier transformation. Uncertainty maps are produced by calculating the relative standard deviation of the flux distribution samples. Figure 4 visualizes the source reconstruction for the example predictions shown above. It displays the reconstructed flux distribution (top left), the true flux distribution (top right), the difference between both (bottom left), and the calculated uncertainty map (bottom right). Areas with reconstruction problems can be identified as these areas stand out due to higher relative uncertainties. This allows an uncertainty estimation in cases where the true flux distribution is unknown. Until now, the method can only indicate the relative uncertainty. One plan is the quantification of the uncertainty for each pixel for better interpretation.

References

- [1] Ledig et al. Photo-realistic single image super-resolution using a generative adversarial network, 2016.
- [2] Thompson et al. *Interferometry and Synthesis in Radio Astronomy*. 01 1991.
- [3] Sam Gross and Michael Wilber. Training and investigating residual nets, 2016.

Feasibility study to measure the muon bremsstrahlung cross section using neutrino telescopes

Jan Soedingrekso
Experimentelle Physik 5b
Technische Universität Dortmund
jan.soedingrekso@tu-dortmund.de

Muons are the dominant event signature for neutrino telescopes like IceCube and they are the main background for neutrino searches. Furthermore, they are used to investigate extended air showers. In both cases, the stochasticity of the muon propagation plays a key role in the data extraction step and an accurate understanding, even of the edge cases, is crucial. The main process driving stochastic losses for TeV scale muons is bremsstrahlung.

A feasibility study is developed to measure the cross section of stochastic losses using neutrino-induced muons. The simulation study is based on the propagation of muons using the Monte-Carlo library PROPOSAL. For different reconstruction methods and resolutions, the energy loss distribution for different muon energies is used to estimate the sensitivity to measure the bremsstrahlung cross section. Two further systematic parameters, the detection efficiency, which scales the amount of detected light, and the spectral index are also estimated to analyze their correlation to the estimated bremsstrahlung normalization. The statistics of the simulated dataset correspond to 10 years of up-going muon neutrino data in IceCube.

1 Introduction

For neutrino telescopes and other deep underground experiments, muons are the main background, dominating the event rate by several orders of magnitude. Due to their

stochastic behavior during propagation, muons may remain undetected by the veto mechanism [5] while producing a single large energy loss inside the detector volume, mimicking neutrino-like signal events. Therefore, a precise description of the muon propagation in the Monte-Carlo simulation is crucial for an accurate background estimation, especially for the edge cases like the highest energetic losses.

The large, stochastic energy losses are mainly driven by bremsstrahlung interactions. The bremsstrahlung cross section has an overall uncertainty of a few percent [4], while recent calculations increase the accuracy below a percent [6]. However, both calculations assume that the outgoing particles are highly boosted in the forward direction, which does not hold for the highest energy losses resulting in increased uncertainties in these regions.

To answer the question whether neutrino telescopes are sensitive to these uncertainties, this simulation study was developed. The sensitivity to measure the normalization of the bremsstrahlung cross section is calculated using the energy loss distribution of muons propagated through the detector.

The event statistic in this simulation study is loosely based on the Northern track sample of the IceCube detector from the diffuse ν_μ analysis [7], which contains almost 10 years of data. According to the event spectrum of the aforementioned selection, there are around 245 000 events between the neutrino energies of a TeV and 100 TeV with a spectral index of 1.63. The resulting secondary muons have a steeper spectrum due to propagation losses. Their distribution is further shifted to lower energies because a part of the neutrino energy is transferred to the hadronic cascade at the vertex. The muon spectrum assumed here is a rough estimation, but nevertheless sufficient for this feasibility study.

2 Simulation and Measurement

The Monte-Carlo simulation of the muons is performed using the lepton propagator PROPOSAL [3, 2]. Muons with a flat energy spectrum between 100 GeV and 1 PeV are propagated through ice. To take into account the different propagation lengths inside the detector, the maximum propagation distance is varied between 100 m and 1 km. The produced secondaries through interactions or decays can be stacked to an energy loss distribution. In order to reduce resource requirements and to ensure applicability to neutrino telescopes in all generality, a simplified photon and detector simulation is applied. In this simplified approach, the secondaries produced in the muon simulation are smeared out to produce realistic and measurable energy loss segments along the muon track. In addition to the energy losses per track segment, the propagation length is smeared independently.

Based on the track segments, the energy is reconstructed using two independent methods. One method called *truncated energy* estimates the energy analogous to the truncated mean method from [1] using the correlation between the average energy loss and the muon energy. The second approach, a convolutional network trained on the reconstructed track segments, is employed as a cross-check to the *truncated energy* method, resulting in a similar performance. Furthermore, the same network architecture is used to estimate an uncertainty on the estimated energy. This estimate is used for both energy reconstruction methods to discard events with a high energy uncertainty.

An event selection is performed to discard mis-reconstructed events. First, events traveling less than 100 m are filtered out to avoid muons skimming the detector volume. The second cut is performed on the energy uncertainty and is chosen to strike a balance between selecting well reconstructed events and retaining enough statistics. Furthermore, only muon energies between 1 TeV and 100 TeV are selected.

For five muon energy ranges, an energy loss histogram is created based on the reconstructed muon track segments while accounting for the simulation weights for the chosen energy spectrum. Multiple simulation datasets are created, each with a different scaling of the bremsstrahlung cross section, or bremsstrahlung multiplier, varied by $\pm 10\%$, to interpolate the differences in the energy loss histograms. Therefore, the histograms are normalized by the propagated length.

Two further systematic parameters are considered, the DOM efficiency and the spectral index of the muon energy spectrum as they both affect the energy loss distribution. The DOM efficiency scales the energy losses per track segment, representing the uncertainty in the efficiency of the photomultiplier. It is varied by $\pm 10\%$. The spectral index specifies the power law of the assumed muon energy spectrum, which is varied between 1.5 and 1.9 with 1.7 as the baseline value. Their effects on the energy loss profile are also interpolated, resulting in a three dimensional interpolation function for each energy loss bin. With these interpolations, a Poisson Likelihood is defined to describe the bin contents of the energy loss distributions. An MCMC sampling is then applied to estimate the three parameters including their correlations.

To estimate the performance of the bremsstrahlung multiplier measurement for the different resolutions, simulation sets with random bremsstrahlung multipliers and fixed systematic parameters are produced. The correlations of the MCMC results are shown in Figure 1.

For a detector with an IceCube-like resolution, the normalization of the bremsstrahlung cross section can be measured with an uncertainty of $\pm 4\%$. As the bremsstrahlung uncertainties have detectable effects, they could be considered as a further systematic parameter in event selections which depend on the muon stochasticity.

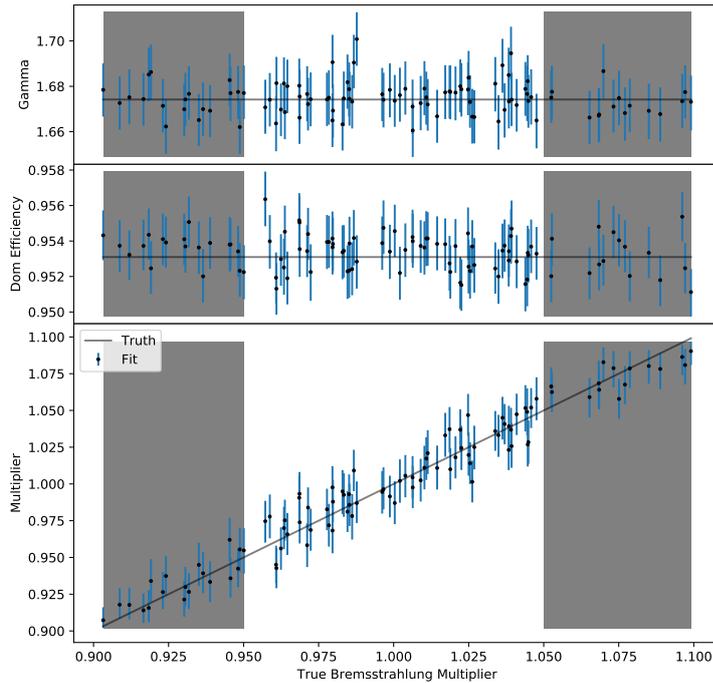


Figure 1: Results of the MCMC samplings with the high resolution settings for the Spectral index, the DOM efficiency and the bremsstrahlung multiplier. The muon energy is reconstructed using the truncated energy method. The region below 0.95 and above 1.05 is neglected for the performance to avoid boundary effects during the MCMC sampling at the edge of the allowed interpolation region.

References

- [1] M. G. Aartsen et al. "Energy Reconstruction Methods in the IceCube Neutrino Telescope". In: *JINST* 9 (2014), P03009. DOI: 10.1088/1748-0221/9/03/P03009. arXiv: 1311.4767.
- [2] M. Dunsch et al. "Recent Improvements for the Lepton Propagator PROPOSAL". In: *Comput. Phys. Commun.* 242 (2019), p. 132. DOI: 10.1016/j.cpc.2019.03.021. arXiv: 1809.07740.
- [3] J.-H. Koehne et al. "PROPOSAL: A tool for propagation of charged leptons". In: *Comput. Phys. Commun.* 184 (2013), p. 2070. DOI: 10.1016/j.cpc.2013.04.001.
- [4] R. P. Kokoulin. "Uncertainties in underground muon flux calculations". In: *Nucl. Phys. B Proc. Suppl.* 70 (1999), p. 475. DOI: 10.1016/S0920-5632(98)00475-7.
- [5] L. Lu, C Haack, and T Yuan. "A novel method of rejecting muon backgrounds for the detection of the highest energy neutrinos". In: *PoS(ICRC2019)*. 36th International Cosmic Ray Conference. 945. 2019.
- [6] J Soedingrekso, A Sandroek, and W Rhode. "The effect of improved high-energy muon cross sections". In: *PoS(ICRC2019)*. 36th International Cosmic Ray Conference. 429. 2019. arXiv: 1910.07050.
- [7] J. Stettner. "Measurement of the diffuse astrophysical muon-neutrino spectrum with ten years of IceCube data". In: *PoS(ICRC2019)*. 36th International Cosmic Ray Conference. 1017. 2019. arXiv: 1908.09551.

Characterization of the high energy muon flux using leading muons in the IceCube detector

Johannes Werthebach
Experimentelle Physik 5b
Technische Universität Dortmund
johannes.werthebach@tu-dortmund.de

Muons produced in cosmic ray showers are the most dominant part of events in large-volume underground neutrino detectors. Although these muons are considered as background for neutrino searches, they play an important role in investigating the extended air showers. Due to its size of one cubic kilometer, the IceCube detector is ideally suited to track high energy atmospheric events over a long distance and find events dominated by just one muon. The main approach to identify TeV scale muons is by looking for stochastic losses along the track. Such a dataset can be used to extract the single muon spectrum and constrain the contribution of prompt decays of short-lived hadrons.

In this work, a model is developed to characterize the energy deposition inside the detector and separate events with a high stochastic loss compared to the total energy deposition. The energy of the primary muon in the event is reconstructed. Using unfolding techniques the single muon energy spectrum up to several PeV is determined. Previous analysis show a strong indication of a non negligible component of prompt decays in the data. Using up to eight times more data from the IceCube detector will decrease statistical uncertainties, one of the major limitations of previous analysis.

1 Introduction

The IceCube detector located near the geographic South Pole is a particle detector with an instrumented volume of about one cubic kilometer [6]. The main part of the detector consists of a total of 5160 digital optical modules (DOMs) divided into 86 cables, each holding 60 of those DOMs at depths between 1450 m and 2450 m. The second part is located on the surface and called IceTop. This surface array is build of 162 frozen water tanks bundled into pairs and placed close to the 81 main cables of the in-ice detector. Each DOM contains a photomultiplier tube as well as readout electronics and communicates with a server responsible to handle data acquisition.

IceCube's main scientific goal - the search for astrophysical neutrinos - was achieved by the detection of the first neutrino with astrophysical origin in 2013 [3]. After several years of data acquisition, IceCube confirmed the detection of the astrophysical neutrino flux [7]. These analyses use high-energy starting events (HESE) separated from background events via a veto region inside the main in-ice detector. These background events consisting mainly of atmospheric muons and some atmospheric neutrinos. As this background can be filtered out using a detection channel selecting charged-current muon-neutrino events, the muons itself are of high interest for atmospheric muon studies [1], primary composition studies [2] or could be used to determine the contribution of short-lived hadrons to the atmospheric muon flux, see fig. 1 [5]. These short-lived particles, called *prompt*, decay before any energy loss and, therefore, preserve the same spectral index. Reconstructing the prompt muon flux provides knowledge about their production cross sections in forward direction. Such a measurement cannot be obtained by modern particle accelerators as the beam pipe blocks the most probable forward direction.

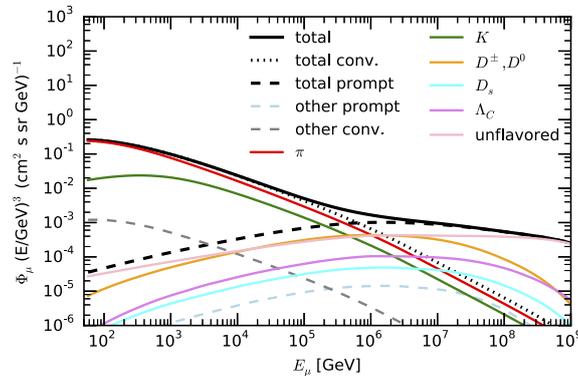


Figure 1: Contributions of the intermediate particles to the atmospheric muon flux calculated with SIBYLL-2.3 RC1. The prompt contribution to the flux is dominant at energies above several hundred TeV. Fig. from [4]

2 Analysis

The major task in this analysis is the creation of a data sample that contains the so called leading muons. Here a leading muon contains most of the energy deposited in the detector and can be interpreted as a single muon. This allows the reconstruction of the muon energy from the measured event. In general, muons arrive in bundles at the detector. Their multiplicity can range from just a single muon to several thousand. Due to the average distance of 125 m between the single strings it is not possible to resolve single muons in one event. Every measured value in one event is therefore a combination of all muons. To separate the leading muon from the bundle events the event topology is important. Figure 2 shows two exemplary event topologies of a leading muon on the left and a normal muon bundle on the right. The difference between these two topologies is the stochastic loss in the left picture indicated by the big orange bubble. The occurrence of a high stochastic loss is more likely if one muon carries most of the total bundle energy as the cross section for such a stochastic loss (eg. bremsstrahlung) increases with energy.

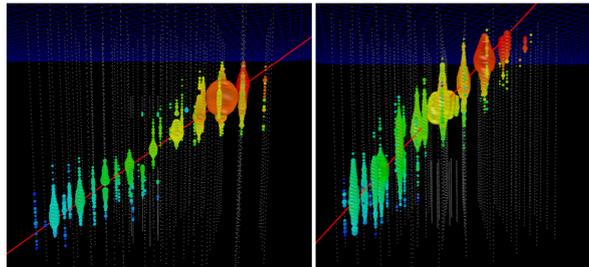


Figure 2: Exemplary topologies of a leading muon (left) and a muon bundle (right)

The definition of a leading muon is somewhat blurry. Previous analyses defined a leading muon by the fraction of energy it was carrying, namely 50 % of the total bundle energy. Here a different approach is chosen. The energy difference between the most and second most energetic muon is used. If this energy distance is more than 40 % of the total bundle energy, the most energetic muon is considered a leading muon and all the other events are considered background. In order to separate the leading muon events advanced data-mining techniques are applied. In a first step the topology is characterized and possible stochastic losses are identified. Depending on these characteristics a Boosted Decision Tree (BDT) is trained to identify the background events. A very loose cut in the confidence of the BDT removes a significant amount of background events. A second approach with a Deep Neural Net is currently tested to replace the BDT as a pre-cut to remove background events. Secondly, another BDT is trained to create the final data sample. The resulting data sample can be used to unfold a leading muon spectrum.

The spectrum in figure 3 from 2011 data is the latest measurement of high energy muons. At high energies it is limited by the high statistical errors caused by the low

amount of Monte-Carlo-Data available to reconstruct the spectrum. Nevertheless, a hint for a contribution of the prompt component was found. Newer Monte-Carlo-Data with approximately ten times the statistics are used in this analysis and an improvement of the spectrum at high energies is expected.

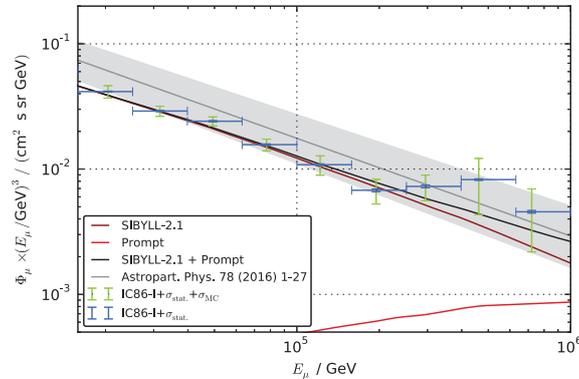


Figure 3: The leading muon spectrum with the IC86-I data (2011). Fig. from [5]

References

- [1] IceCube Collaboration, M. G. Aartsen et al. “Characterization of the atmospheric muon flux in IceCube”. In: *Astroparticle Physics* 78 (2016). DOI: 10.1016/j.astropartphys.2016.01.006.
- [2] IceCube Collaboration, M. G. Aartsen et al. “Cosmic ray spectrum and composition from PeV to EeV using 3 years of data from IceTop and IceCube”. In: *Phys. Rev. D* 100 (8 2019). DOI: 10.1103/PhysRevD.100.082002.
- [3] IceCube Collaboration, M. G. Aartsen et al. “Evidence for High-Energy Extraterrestrial Neutrinos at the IceCube Detector”. In: *Science* 342 (6161 2013). DOI: 10.1126/science.1242856.
- [4] A. Fedynitch et al. “Calculation of conventional and prompt lepton fluxes at very high energy”. In: *EPJ Web Conf.* 99 (2015). DOI: 10.1051/epjconf/20159908001.
- [5] T. Fuchs. “Charmante Myonen im Eis”. PhD thesis. TU Dortmund, 2016. DOI: 10.17877/DE290R-17241.
- [6] IceCube Collaboration, A. Karle. “IceCube”. In: *Proceedings of the 31st ICRC* (2009). URL: <http://arxiv.org/abs/1003.5715v1>.
- [7] IceCube Collaboration, C. Kopper. “Observation of Astrophysical Neutrinos in Six Years of IceCube Data”. In: *Proceedings of the 35th ICRC* (2017). DOI: 10.22323/1.301.0981.



Subproject C4
Regression approaches for large-scale
high-dimensional data

Katja Ickstadt

Alexander Munteanu

Coresets for logistic regression

Simon Omlor
Faculty of Statistics
TU Dortmund University
simon.omlor@tu-dortmund.de

In several areas of research and in the industry we are faced with massively large data sets, such that classical methods for regression or classification reach their limits of efficiency. Therefore, new methods need to be developed for handling massive data sets and their statistical analysis. To handle these kinds of problems it is helpful to first reduce the amount of data. Our current research aims towards reducing the size of the data via subsampling and reweighting the data or computing linear combinations of data points via sparse and fast linear sketches. On the data summary of small size we get the usual algorithms run efficiently. However sometimes it is impossible to reduce the amount of data without changing the optimal solution too much. Thus one goal of us is to determine properties of data sets that allow us to shrink the data and to determine instances where no small coresets exist. The problems we consider are generalized linear regression models.

Introduction and recent developments

We are considering the following type of problem: We are given a data set given by matrix $X \in \mathbb{R}^{n \times d}$ where each row represents a data point $x_i \in \mathbb{R}^d$ for $i \in \{1, \dots, n\}$. We want to find $\beta \in \mathbb{R}^d$ minimizing an objective function $f(X\beta)$. In the setting of logistic regression we have $f(X\beta) = \sum_{i=1}^n \ln(1 + \exp(x_i\beta))$. Our goal is to compute a coreset which is defined as follows:

Definition 1. Given $X \in \mathbb{R}^{n \times d}$ and a target function f a matrix $C \in \mathbb{R}^{k \times d}$ together with a weight vector $w \in \mathbb{R}^k$ is a (weighted) $(1 \pm \epsilon)$ -Coreset if for all $\beta \in \mathbb{R}^d$ it holds

$$(1 - \epsilon)f(X\beta) \leq f_w(C\beta) \leq (1 + \epsilon)f(X\beta)$$

where $f_w(C\beta) = \sum_{i=1}^k w_i \ln(1 + \exp(c_i\beta))$.

In general there are instances for which no coresets of sublinear size exist. Thus finding sufficient conditions for the existence of small coresets is crucial if we want to use them. One example of a bad instance is an instance where every point is linearly separable from each other point. Here every point is important for some $\beta \in \mathbb{R}^d$. One sufficient condition for the existence of a small coreset is μ -complexity:

Definition 2. Let $\mu > 1$. We say that $X \in \mathbb{R}^{n \times d}$ is μ -complex if for all $\beta \in \mathbb{R}^d$ it holds

$$\frac{\sum_{x_i\beta > 0} x_i\beta}{\sum_{x_i\beta < 0} -x_i\beta} \leq \mu.$$

In μ -complex data sets there always exist coresets of size polynomial in \sqrt{n} , d and μ . Those can be found by combining sketching techniques and sensitivity sampling.

Theorem 1. [5] Let $X \in \mathbb{R}^{n \times d}$ be a μ -complex data set. Then there exists a $(1 \pm \varepsilon)$ -Coreset of size polynomial in \sqrt{n} , d , μ and ε^{-1} . The coreset can be constructed in time $\mathcal{O}(\text{nnz}(X) \log(n) + \text{poly}(d) \log(n))$ in 2 passes over the data and with success probability $1 - \frac{1}{n^c}$ for any absolute constant $c > 1$.

By repeating the procedure recursively we can reduce the size of our coreset even further:

Theorem 2. [5] Let $X \in \mathbb{R}^{n \times d}$ be a μ -complex data set. Then there exists a $(1 \pm \varepsilon)$ -Coreset of size polynomial in $\log(n)$, d , μ and ε^{-1} . The coreset can be constructed in time $\mathcal{O}(\text{nnz}(X) + \text{poly}(d) \log(n) \log \log(n))$ in $2 \log(\frac{1}{\eta})$ passes over the data for some small $\eta > 0$ and with success probability $1 - \frac{1}{n^c}$ for any absolute constant $c > 1$.

In the following we give a short description of the tools used to construct the coreset.

Sketching

The idea of sketching is to reduce the size of a matrix $X \in \mathbb{R}^{n \times d}$ with $n > d$. One example of a sketch is the count sketch. In the count sketch each row x_i gets hashed into a one of k buckets together with a sign which is chosen randomly. Then the entries in each bucket are summed up to a row of a new matrix X' . If k is large enough it holds $\|X\beta\|_2(1 - \varepsilon) \leq \|X'\beta\|_2 \leq \|X\beta\|_2(1 + \varepsilon)$ for all $\beta \in \mathbb{R}^d$ with high probability. For a more detailed explanation of sketching we refer to the work of Woodruff [6] and Clarkson et al [3]. Clarkson and Woodruff [2] developed a sketching algorithm that preserves not only

the ℓ_2 -norm but also functions ranging from the ℓ_1 -norm to the ℓ_2 -norm that is useful for our research in particular. However it comes with two downsides, we just get a constant approximation guarantee and it will not always preserve the target value of all $\beta \in \mathbb{R}^d$. Sketching can also be used to approximate the leverage scores of X .

Sensitivity sampling

Sensitivity is defined as follows:

Definition 3. [4] Consider a family of functions $\mathcal{F} = \{g_1, \dots, g_n\}$ mapping from \mathbb{R}^d to $[0, \infty)$ and weighted by $w \in \mathbb{R}_{>0}^n$. The sensitivity of g_i for $f_w(x) = \sum_{i=1}^n w_i g_i(x)$ is

$$\varsigma_i = \sup \frac{w_i g_i(x)}{f_w(x)} \quad (1)$$

where the sup is over all $x \in \mathbb{R}^d$ with $f_w(x) > 0$. If this set is empty then $\varsigma_i = 0$. The total sensitivity is $\mathfrak{S} = \sum_{i=1}^n \varsigma_i$.

The sensitivity of a point measures its worst-case importance for approximating the objective function on the entire input data set. Performing importance sampling proportional to the sensitivities of the input points thus yields a good approximation. In general sensitivities are hard to compute [1] but for logistic regression it is possible to get an upper bound for the sensitivity of each point x_i which depends on μ and the leverage score of x_i . As we want to get a good approximation for all $\beta \in \mathbb{R}^d$ we also need that the VC-dimension of the range space of \mathcal{F} is bounded by d . A framework combining the sensitivity scores with a theory on the VC dimension of range spaces was developed in [1].

Future research

In our future research we will focus on the following questions:

Since the data is sometimes saved external and slowly accessible or is given via a data stream one question is whether we can construct a small coresnet in just one pass over data? Our first goal here will be to develop a fast algorithm that can compute a constant factor approximation of the optimal β minimizing $f(X\beta)$ for logistic regression. The second goal will be to determine a class of instances where we can get $1+\epsilon$ -Approximation.

Another aim of us is to develop a coreset for the so called probit model where the objective function is given by

$$f(X\beta) = \sum_{i=1}^n -\ln(\Phi(x_i\beta))$$

where Φ is the cumulative distribution function of the normal distribution. Similar as in logistic regression our method is to sample a small number of data points from a carefully designed importance sampling distribution. This will guarantee that the likelihood is approximated for every possible parameter vector, not only for the "optimal" maximum likelihood estimator. As a consequence this allows us to approximate the distribution of parameter vectors in the Bayesian setting. Finally, we will carry out an empirical evaluation of the method in the frequentist as well as in the Bayesian case. This will be done using a Gibbs Sampler.

We also would like to extend our findings to more complex models functions.

References

- [1] Vladimir Braverman, Dan Feldman, and Harry Lang. New frameworks for offline and streaming coreset constructions. *arXiv preprint arXiv:1612.00889*, 2016.
- [2] Kenneth L Clarkson and David P Woodruff. Sketching for m-estimators: A unified approach to robust regression. In *Proceedings of the twenty-sixth annual ACM-SIAM symposium on Discrete algorithms*, pages 921–939. SIAM, 2014.
- [3] Kenneth L. Clarkson and David P. Woodruff. Low-rank approximation and regression in input sparsity time. *J. ACM*, 63(6), January 2017.
- [4] Michael Langberg and Leonard J Schulman. Universal ϵ -approximators for integrals. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pages 598–607. SIAM, 2010.
- [5] Alexander Munteanu, Chris Schwiegelshohn, Christian Sohler, and David P. Woodruff. On coresets for logistic regression. *CoRR*, abs/1805.08571, 2018.
- [6] David P Woodruff. Sketching as a tool for numerical linear algebra. *arXiv preprint arXiv:1411.4357*, 2014.

Regression approaches for high-dimensional data under maintaining of subgroups

Julian Riehl

Lehrstuhl für Mathematische Statistik und
biometrische Anwendungen

Technische Universität Dortmund

julian.riehl@tu-dortmund.de

The central theme of project C4 is the consideration of regression approaches for large-scale high-dimensional data. There are already theoretical guarantees for data reduction methods such that a regression model calculated on the smaller data set is close to the original model on the source data set. Since a complete statistical analysis also includes aspects such as model diagnostics and the further development of models, the goal of future research is to adapt the data reduction processes so that these points are also covered. A special focus will be on the maintenance of subgroups. The targeted steps to be taken on the way to such a result are explained in this technical report.

1 Introduction

In the context of linear regression, many approaches have been investigated in recent years that aim to reduce very large data sets in such a way that a regression model calculated on the smaller data set is close to the original model on the initial data set. The current approaches to reduce data sets in the presence of a large number of observations can be roughly divided into the categories "random projections" and "(weighted) samples". The theoretical guarantees from the perspective of computer science as well as the statistical properties of the resulting models have been widely studied and offer different

advantages. However, the studies have in common that they mainly refer to the regression model to be obtained. In the context of a statistical analysis, however, it is of great importance to check the model for its assumptions. Furthermore, it is important to decide on the basis of the model and data whether the current model describes the data well or whether variables should be transformed or interactions should be included, whether subgroups or outliers are adequately represented and whether there are indications for the absence of important variables. In this sense, each model is first of all a preliminary model, which aims to show improvements and modeling options - iterative until a "useful model" is found.

The goal of future research is now to adapt the methods for reducing very large data sets in such a way that the entire process of statistical analysis is covered, especially the steps of further development of models and diagnostics, with an increased focus on maintaining of subgroups and outliers. Which steps seem to be reasonable to reach this goal shall be described in the following.

2 Planned Procedure

In a first step, it is useful to describe the given situation. For this purpose it is necessary to define which steps are part of a typical statistical analysis. Furthermore, it is necessary to consider which scenarios can occur in an analysis after the determination of a first model. Possibly it turns out that an important variable or interaction term is missing. It is also conceivable that the transformation of a variable is necessary or that a subset of the data shows a very different behavior. After defining such scenarios, the question arises what happens if existing approaches to data reduction are applied in the presence of single scenarios or a combination of several scenarios. The effects of this approach on an analysis and the possible misconclusions it may lead to would then have to be checked by means of suitable empirical studies. Of particular interest is the case already mentioned, that there are subgroups that differ greatly. In such a situation it would of course be desirable that the subgroup structure of the original data set is also found in the reduced data set.

In a second step, the process is extended for a special model. For this purpose, a comparatively simple model class is selected, which is to be considered more closely for this step. The most obvious way would certainly be to consider a linear model. Depending on the results of the first step, the question arises whether it is possible to apply the known reduction methods to perform analysis, diagnostics and further development on one data set, or whether a second, reduced data set is required for this purpose. If the latter case occurs, it seems reasonable to consider approaches from sampling theory. For example, the use of stratified samples may possibly maintain in subgroups. A further option is to modify the existing procedures in order to pursue the goals of "diagnostics" and "further development" in a meaningful way. For this purpose a cooperation with

computer scientists from project C4 would be conceivable. Depending on the results of the previous investigations, it could now be checked whether there are differences between the scenarios from the first step. It would of course be desirable if one method could cover all scenarios in the same way. It seems more likely, however, that several methods adapted to the individual scenarios will be necessary.

If the described questions have been sufficiently investigated for a simply selected model class, in a third step the extension to other model classes is possible. It is questionable to what extent the findings from the second step can be transferred to other model classes. At best, the results are completely transformable. Otherwise, it has to be checked which restrictions have to be made to allow the applicability of the previous results. It may be sufficient to make minor changes to the reduction methods.

The third step involves greater imponderables. Thus, it is quite possible that the results from the second step are trivially valid for a large model class. In this case, the third step could possibly be completely absorbed by the second step. Conversely, it is also conceivable that a transfer would be very costly. Also imaginable is a situation that lies between the above two. The basic idea is to first focus on the scenarios and one or more suitable reduction methods and then to broaden the model class.

3 Outlook on the maintaining of subgroups

Finally, it is now a question of looking at existing research work that could be methodologically linked to achieve the goals described above. By way of example, the following explanations are limited to the question of how to deal with subgroups from the original data set that show behavior that deviates from the overall population.

In the described consideration of subgroups, a distinction is made between two approaches. If the goal is to estimate the effects of an already known subgroup on the basis of a very large data set, the problem is trivial. In this case, the data set is reduced in a first step so that it contains only the data points that belong to the desired subgroup. Then one of the known methods for data reduction is applied and the regression model is fitted to the resulting small data set, for which corresponding theoretical guarantees then apply. The same procedure is repeated for all required subgroups.

Much more interesting is the case where the goal is to find subgroups with behavior that deviates from the overall population. The same applies to a situation in which the subgroups change predictively, e.g. by merging or splitting. For these situations it is important to have identified a reduced data set that represents the entire original data set and still retains the subgroup structure. Research in the field of clusters could be helpful here, where the aim is to determine a coresets for a very large data set that retains a given cluster structure. There has already been extensive research in this area. For example, [1] is about training a statistical mixture model on a very large data set by means of coresets. In [2] a novel coresets construction for approximate posterior inference in

the nonparametric Dirichlet process Gaussian mixture model is developed. Also interesting is [3], where the existence of coresets for DP-Means, a prototypical nonparametric clustering problem, is shown and a practical construction algorithm is provided.

4 References

- [1] Mario Lucic, Matthew Faulkner, Andreas Krause and Dan Feldman. Training Gaussian Mixture Models at Scale via Coresets. *Journal of Machine Learning Research*, 18(160):1-25, 2018.
- [2] Zalán Borsos, Olivier Bachem and Andreas Krause. Variational Inference for DPGMM with Coresets. *Advances in Approximate Bayesian Inference*, 2017.
- [3] Olivier Bachem, Mario Lucic and Andreas Krause. Coresets for Nonparametric Estimation - the Case of DP-Means. *Proceedings of the 32nd International Conference on Machine Learning, PMLR*, 37:209-217, 2015.



Subproject C5

Real-Time Analysis and Storage of High-Volume Data in Particle Physics

Bernhard Spaan

Jens Teubner

Selecting Rare, High-Dimensional Objects with Compositional Indexing

Maximilian Berens

Lehrstuhl für Datenbanken und Informationssysteme

Technische Universität Dortmund

maximilian.berens@tu-dortmund.de

The search for rare objects in vast amounts of high-dimensional, scientific data sets is an important problem. Using inverted indices and quantization, we develop new methods to accelerate selective range queries. In the following, we will briefly describe our approach and important performance tradeoffs.

Finding rare objects in very large databases is a frequent objective in many scientific research fields, such as genetics, astronomy or particle physics. During the life time of a research project, the ever-developing knowledge usually invalidates previous assumptions and evolves into new expectations on the properties of the objects under investigation. This usually requires the means to (re-)process vast amounts of data from previous experiments in an offline fashion. In addition, high dimensionality and limited storage capacity prevents the use of traditional, exact database indexing techniques, such as B+-trees or R-trees, in order to accelerate searches with continuously evolving sets of queries.

Two general concepts are usually applied to address these problems: For high dimensionality, the performance of (often tree-based) indices deteriorates, since the additional cost of the index access offsets the cost savings of excluding parts of the data from the search. Even a plain scan outperforms traditional indices in high dimensions [5]. Scan accelerators, such as VA-Files [5], Column Sketches [1] or Column Imprints [3], therefore try to improve the inevitable scan by reducing the cost of evaluating individual tuples. These techniques are also said to defeat the curse of high dimensionality, as their cost only depends linearly on the number of dimensions. However, performance gains can be insufficient, because these techniques still require either the consideration of every single tuple, i.e. they do not scale sublinearly in the number of tuples, or a clustered storage order that is usually unavailable, especially when the set of relevant attributes in a query varies.

The second idea, that also finds application in many scan accelerators, is the admission of inaccuracy in an index's result. A common approach is to apply some form of quantization, where data is replaced with a set of representatives that is supposed to capture the general properties (or the distribution) of the data. Two well known examples for nearest neighbor search are the aforementioned VA-Files, where tuples are represented via quantiles, and Product Quantization [2], that is based on (k-means-) clustering. Column Sketches, also based on quantiles, can be used for ranged queries and additionally offer the property to dynamically select only columns relevant to a given query. Because some scientific queries can not be formulated in terms of a nearest neighbor problem, we will focus our attention solely on techniques that accelerate ranged queries. A bitmap-based, but (arguably) comparable technique to Column Sketches, is FastBit [6]. It utilizes "binning" to represent high cardinality attributes. When using so called equi-height binning, FastBit effectively also applies quantile-quantization. In contrast to Column Sketches, that are originally developed for the use in in-memory systems, the actual representation of a tuple differs and bitmap compression and encoding play a key role in the efficiency of FastBit, which is used to accelerate queries in large, disk-based scientific databases. To elaborate, when using b bins to represent values of a tuple, Column Sketches require $\log_2 b$ bits per value, whereas FastBit uses b bit, since b -bitmaps of length n are derived. FastBit therefore has to rely on compression to offset this additional storage cost. On the other hand, answering a query with FastBit can be done by considering only the bitmaps that actually intersect with the query region, ignoring the rest. As a result, selective predicates translate to reduced access costs, whereas Column Sketches always incur the same cost for all values. It is noteworthy that bitmap encoding techniques offer a way to reduce the number of relevant bitmaps, although at the expense of a decreased compressibility. In the following, we will argue how this cost can be decreased substantially and then briefly discuss the design decisions and their impact on performance.

In order to search through large amounts of high dimensional data, we begin by first group all attributes into sets the size d_{sub} . Then, a number of b quantiles are separately calculated over all attributes. For each group of attributes, the quantile-borders span a partitioning grid of $(b + 1)^{d_{\text{sub}}}$ cells over the data space, where each cell (or bucket) contains at most $\frac{n}{b}$ values. We then proceed by creating an inverted index as well as a multidimensional count-matrix for bucket cardinalities. These two structures combined will be called a sub-index in the following. Each tuple can be assigned to exactly one cell in each sub-index, according to it's position in the partitioned data space. Thus, for each tuple, we add it's id to the corresponding cell's posting list and increment the count for this list in the count-matrix. Eventually, when all tuples are assigned, all inverted indices/sub-indices will contain all n ids. The posting lists are then stored linearly, either as an in-memory array or in a file on disk (depending on the magnitude of n). The count-matrix can be used to derive the starting-address of a bucket, when it's counts are accumulated in the same order as the cells are stored/linearized. A simple linearization is the column-major order. As a result, every sub-index represents a set of (up to) d_{sub} attributes and can be used to filter tuples in the following way: To process a range query, posting lists of grid cells that intersect with

the query-region are fetched from each relevant sub-index. For every attribute involved in a query, we have to fetch (at least one of) it's associated sub-indices. Irrelevant sub-indices can be skipped, similar to columns in a Column Sketch. The next step involves the logical combination of the fetched buckets via set-intersection and union. Before evaluating the disjunctions and conjunctions of the query, the buckets of each sub-index are first merged. After performing all logical operations, the surviving ids form the result of the index access. This result can be used to retrieve the actual data from storage. Although the true result is guaranteed to be a subset of this answer, additional false positives can also be returned and increase the cost of this so called candidate refinement step. Note that this cost factor arises in all indexing techniques that use some form of quantization to approximate the actual result.

In order to assess the overall performance of our approach, multiple aspects have to be considered. As indicated above, indexing techniques, such as FastBit, make use of predicate selectivity in order to exclude a portion of the data, subsequently decreasing the index access costs. However, even when a query is expected to be very selective overall (i.e. lower than 1%), individual predicates are often not very selective on their own. We therefore advocate to group multiple attributes in the above mentioned way. This idea comes with a cost tradeoff: the number of cells in a sub-index increases exponentially with d , increasing the count of relevant posting-lists (that intersect with a query). For a fixed database size and increasing sub-index dimensionality, the average cardinality drastically decreases, deteriorating the storage media's read bandwidth due to random access. We argue, however, that this cost increase is sufficiently low für moderate configurations of $b^{d_{\text{sub}}}$. Furthermore, it is guaranteed that the ratio of required posting lists and subsequently the number of relevant ids, can not be larger than the selectivity of the most selective predicate. To provide an example, assume a database of 100 tuples and a conjunctive query of three predicates over attributes A, B and C , with selectivities $s_A = 0.9, s_B = 0.8$ and $s_C = 0.2$, respectively. For FastBit, with equal-cardinality binning and 10 bitmaps per attribute, this requires to consider approx. $90 + 80 + 20 = 190\%$ of the bitmaps. In contrast, when all three attributes are grouped together, only 20% of the ids of a single inverted index are considered. In general, it is impossible to find perfect groupings for all possible queries without exhaustively forming all possible attribute groupings. However, prior knowledge about frequently considered, selective attributes can be used to create favourable sets. Note that, while it is technically possible to apply attribute grouping strategies in FastBit, the high sparsity of high dimensional space makes bitmaps an inefficient data structure to represent sorted lists of integers [4]. Another advantage of representing multiple columns with a single sub-index is the reduction of the index size. More specifically, this applies regardless of the magnitude of b , i.e. at a higher index precision: Where FastBit linearly increases the storage costs in b , the number of required bits stays the same for our approach. Only the order of the ids in the linearized sub-index changes, when adding more cells (by increasing either the dimensionality or the number of quantiles). These advantages don't come entirely without cost, however. When values are distributed over an increasing amount of cells, the order to ids in the linearized array also becomes more random and the number of I/O requests rises. Also, the general

compressibility of the array deteriorates.

Preliminary tests with data from particle physics indicate that, for a 14 dimensional query with a high selectivity of 0.0014%, significant reductions in both index size and accessed data volume are possible. In comparison with Column Sketches, a decrease in total index size is readily achieved already by arbitrarily grouping sufficiently many attributes together. The access cost due to an increasing sub-index dimensionality (due to random access) has to be compensated by a large number of tuples, however. Moreover, another impactful design parameter is the actual combination of attributes. Because the accessed index-volume is small for low sub-index selectivities, a good grouping allows to consistently combine multiple predicates per sub-index, subsequently dropping the accessed volume. As a result, a compressed representation of just a few bits per value, as it is used by Column Sketches, has to be overcome by either high, average (sub-index) selectivities or, once again, a large amount of tuples. We therefore conclude that our approach works best for very large databases, because cost penalties, due to random access and the usage of ids (instead of compressed values), are eventually amortized by increased bucket cardinality. Furthermore, the break-even point of efficiency is influenced by the efficient grouping of attributes and subsequently high average sub-index selectivity and more efficient space partitioning strategies.

Our future research will encompass the development of a formal cost model, more efficient space partitionings as well as strategies for consistent attribute grouping.

References

- [1] B. Hentschel et al. Column sketches: A scan accelerator for rapid and robust predicate evaluation. In *Proceedings of the 2018 International Conference on Management of Data*, pages 857–872, 2018.
- [2] H. Jegou et al. Product quantization for nearest neighbor search. *IEEE transactions on pattern analysis and machine intelligence*, 33(1):117–128, 2010.
- [3] L. Sidirourgos and M. Kersten. Column imprints: a secondary index structure. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*, pages 893–904, 2013.
- [4] J. Wang et al. An experimental study of bitmap compression vs. inverted list compression. In *Proceedings of the 2017 ACM International Conference on Management of Data*, pages 993–1008, 2017.
- [5] R. Weber et al. A quantitative analysis and performance study for similarity-search methods in high-dimensional spaces. In *VLDB*, volume 98, pages 194–205, 1998.
- [6] K. Wu et al. Fastbit: interactively searching massive data. In *Journal of Physics: Conference Series*, volume 180, 2009.

Precision measurement of Particle-Antiparticle oscillation using Flavour Tagging at LHCb

Kevin Heinicke

Lehrstuhl für Experimentelle Physik 5

Technische Universität Dortmund

kevin.heinicke@tu-dortmund.de

The oscillating frequency Δm_s between B_s^0 and \bar{B}_s^0 mesons is measured with so far unprecedented statistical and systematic precision of approximately $\sigma_{\Delta m_s} = (\pm 0.005_{\text{stat.}} \pm 0.004_{\text{syst.}}) \text{ps}^{-1}$. A key ingredient for this measurement is the Flavour Tagging software, used to deduce the initial flavour of B mesons. It has been re-optimised until the end of 2017 and this measurement can be used as a performance baseline for newly developed Flavour Tagging algorithms.

1 Introduction to the LHCb Experiment

The LHCb experiment is one of the four big experiments located at the Large Hadron Collider near Geneva, Switzerland. Its main focus is the search for New Physics effects in CP -violating and rare decays of beauty and charm hadrons. In order to maximize the sensitivity with respect to these specialized targets the LHCb detector is built as a forward single arm spectrometer (see Figure 1).

In the collisions of protons inside the vertex locator (VELO), new particles are created and decay until they finally leave traces in the various subcomponents of the detector. These traces are hits in the tracking systems (VELO, TT, T1-T3, M1-M6), clusters in the calorimeters (ECAL, HCAL) and Cherenkov radiation in the Ring Imaging Cherenkov Detectors (RICH1, RICH2). To conclude on the presence of particles the information of these sub detectors needs to be reconstructed, e.g. by fits of trajectories to ensembles of hits and pattern recognition algorithms looking for clusters of energy deposition. Finally, particle candidates need to be combined to heavier particles in order to perform physics measurements on the same.

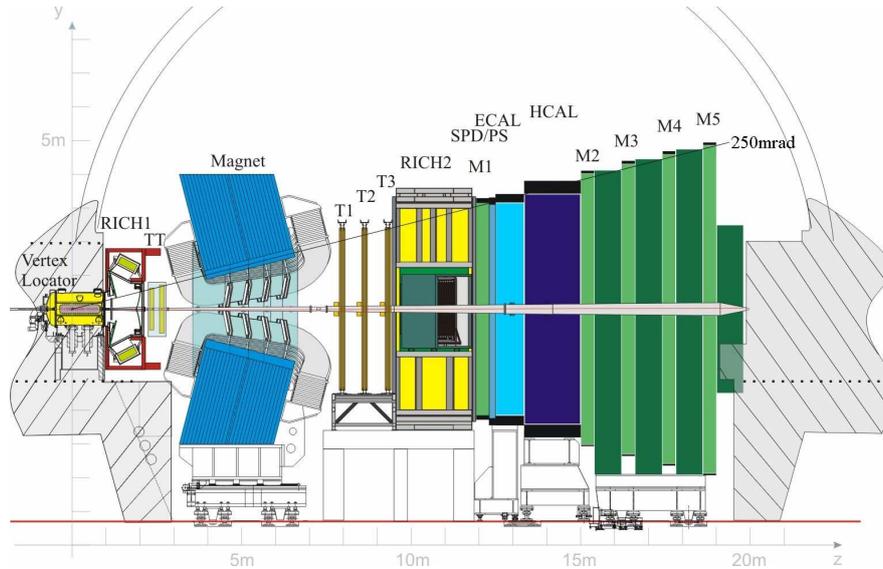


Figure 1: The LHCb detector with the various sub detectors for the identification of particles and reconstruction of their tracks [1].

The endeavour to find the particles of interest is hindered for two main reasons. Firstly, hundreds of particles are produced inside the angular acceptance which allows for a large number of combinations to be made in the reconstruction as well as the combination of particles. Secondly, the interaction rate of 50 ns / 25 ns together with the limitation on the bandwidth that can be written to disk enforces a fast reconstruction that leads to the selection of interesting events and the rejection of physically uninteresting ones. Both these points together set the frame for investigating these tasks in the context of resource limitation: The reconstruction and combination tasks can be parallelized and therefore performed faster.

2 Measurement of the B_s oscillation frequency Δm_s

A major subset of the LHCb analyses involve studies of CP violation, which can be used to test the Standard Model of particles physics (SM). These measurements include time-dependent decay rate measurements, many of which are subject to mixing of neutral B meson states. One of the key parameters of aforementioned measurements of CP violation in the B_s^0 system is the oscillation frequency Δm_s of neutral B_s^0 mesons.

To measure Δm_s , a data sample of a flavour-specific process is needed for which the decay $B_s^0 \rightarrow D_s^- \pi^+$ is a perfect candidate: The charge of the final state π^\pm determines the final flavour of the B_s meson. Combined with the flavour tagging information from previously mentioned flavour tagging algorithms, and the decay time, which can be measured with high precision in the VELO detector, the oscillation frequency Δm_s can be measured.

A statistically pure sample of 378 000 $B_s^0 \rightarrow D_s^- \pi^+$ decays candidates is extracted using the sPlot method [4]. The expected combined tagging power of OS combination and SS

Kaon tagger is estimated to be $\varepsilon_{\text{eff}} = (7.03 \pm 0.03_{\text{stat.}} \pm 0.32_{\text{syst.}}) \%$.

Together with a calibration of the decay time resolution, the expected systematic uncertainty on Δm_s can be determined. A blind fit to data yields $\sigma_{\Delta m_s, \text{stat}} = 0.005 \text{ ps}^{-1}$. Besides others, extensive studies of the data pre-selection, the sPlot parametrisation, the correlation among the decay time t and the reconstructed B_s^0 and D_s masses, and the VELO alignment, have been performed, leading to a systematic uncertainty of approximately $\sigma_{\Delta m_s, \text{syst}} = 0.004 \text{ ps}^{-1}$.

3 Study of the Inclusive Tagging Algorithm

The flavour of neutral B mesons, as used in the previously described oscillation measurement, is extracted with several Flavour Tagging algorithms, which are executed after the particle-decays have been fully reconstructed. The algorithms are designed to reconstruct particles on the same-side (SS) and opposite-side (OS) of the signal B candidate. The charge information of these particles is correlated with the initial flavour of the B candidate via different weak transitions. The different types of algorithms are depicted in figure 2.

In general each algorithm aims to identify a decay product either from the OS non-signal B meson, or from a SS hadron which hadronised with the signal B meson by applying different selection criteria. The selected particles are referred to as tagging particles. The algorithms differ in their specific decay product and selection strategies. Due to the large number of tracks which do not necessarily originate from the primary vertex, this process is error prone. The quality of the identification is therefore evaluated by calculating an overall tagging efficiency

$$\varepsilon_{\text{tag}} = \frac{N_{\text{tagged}}}{N_{\text{tagged}} + N_{\text{untagged}}}, \quad (1)$$

describing the ratio of B candidates for which a tagging particle has been found within all B candidates, and the mistag rate

$$\omega = \frac{N_{\text{incorrect}}}{N_{\text{tagged}}}, \quad (2)$$

which is the rate of incorrectly tagged events or false-positive rate. The overall performance of these algorithms is usually described in terms of the tagging power

$$\bar{\varepsilon}_{\text{eff}} = \varepsilon_{\text{tag}} (1 - 2\omega)^2. \quad (3)$$

To estimate the tagging power on untagged data, the mistag rate ω_i is predicted on a per-event basis by different multivariate analysis tools for each individual tagging algorithm, leading to the final figure of merit for Flavour Tagging, the per-event tagging power

$$\varepsilon_{\text{eff}} = \frac{1}{N_{\text{tagged}} + N_{\text{untagged}}} \sum (1 - 2\omega_i)^2. \quad (4)$$

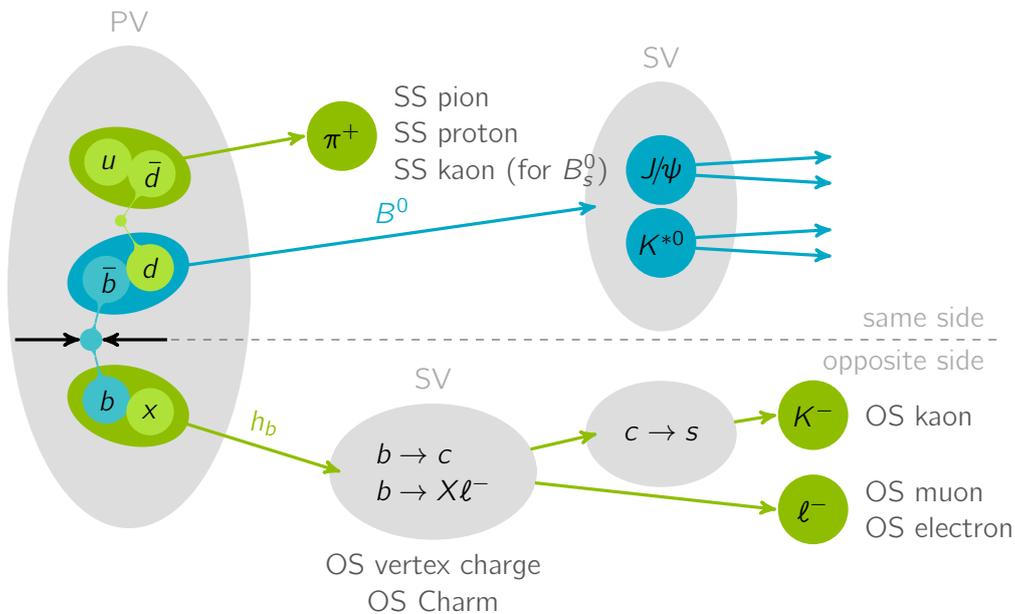


Figure 2: Schematic description of different Flavour Tagging algorithms. The same side algorithms infer the initial b flavour from pions and protons that hadronise alongside with the B meson. The opposite side algorithms infer this information from the non-signal B partner and its decay into leptons or $b \rightarrow c \rightarrow s$ transitions.

A newly proposed strategy to infer the initial B flavour has been implemented together with Vukan Jevtic and Quentin Fühling, using the previously reported, re-factored Flavour Tagging software package. The Inclusive Tagger aims to defer the flavour information via a Recurrent Neural Network, trained on simulated event samples, using `keras` and the `tensorflow` backend [2, 3]. It will replace the individual flavour tagging algorithms depicted in Figure 2 with a single algorithm, allowing to deduce flavour information from the correlation between the different physical processes. After a study of the interpretability of the new algorithm has been performed and reported in the previous period, the validation of the algorithm is prepared. Part of this is the calibration of the established flavour tagging algorithms on the full LHCb Run 2 data sample from 2015 to 2018 which is described in the upcoming section for a sample of $B_s^- \rightarrow D_s^- \pi^+$ decays.

References

- [1] A. A. Alves, Jr. et al. The LHCb detector at the LHC. *JINST*, 3:S08005, 2008.
- [2] Martín Abadi et al. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [3] François Chollet et al. Keras. <https://keras.io>, 2015.
- [4] Muriel Pivk and Francois R. Le Diberder. SPlot: A Statistical tool to unfold data distributions. *Nucl. Instrum. Meth. A*, 555:356–369, 2005.

Measurement of CP violation in $B^0 \rightarrow \psi(\rightarrow \ell\ell)K_S^0(\rightarrow \pi^+\pi^-)$ and $B_S^0 \rightarrow J/\psi(\rightarrow \mu\mu)K_S^0(\rightarrow \pi^+\pi^-)$ decays

Vukan Jevtić
Lehrstuhl für Experimentelle Physik 5
Technische Universität Dortmund
vukan.jevtic@tu-dortmund.de

The LHCb detector [1] is one of the major particle detectors at the LHC. The main research subjects at LHCb are measurements of CP violation and the search for rare decays in the decays of beauty -and charm mesons. Since B mesons are mainly created in the forward region (along the beam axis), the LHCb detector is constructed as a one-arm forward-spectrometer. The LHCb detector that recorded the Run 2 dataset (2015 to 2018) that is used in this analysis is shown in figure 1. While the LHC is running, proton bunches are brought to collision in the Vertex Locator (VELO) at a rate of 40 MHz. The Vertex Locator measures the primary and secondary vertices of a proton-proton collision. The particles that are created in these collisions then traverse the rest of the detector. For physics analyses, a measurement of the particles mass and energy and charge is crucial. For this reason the LHCb detector is equipped with multiple subdetectors that measure the track hits created by charged particles traversing the detector as well as the momentum of the tracks. The tracking components of the LHCb detector are the VELO, the Tracker Turicensis (TT), which is a silicon strip detector located upstream of the LHCb magnet, the tracking stations T1-T3 consisting of an inner silicon tracker and a straw tube outer tracker as well as muon chambers. If charged particles reach these detector components, electromagnetic interactions occur and the detectors produce signals. That way, tracks can be reconstructed from multiple space-time points inside the detector. The warm LHCb dipole magnet bends tracks of charged particles and therefore allows for a measurement of the track curvature from which the track charge can be inferred. By measuring the track curvature and the track

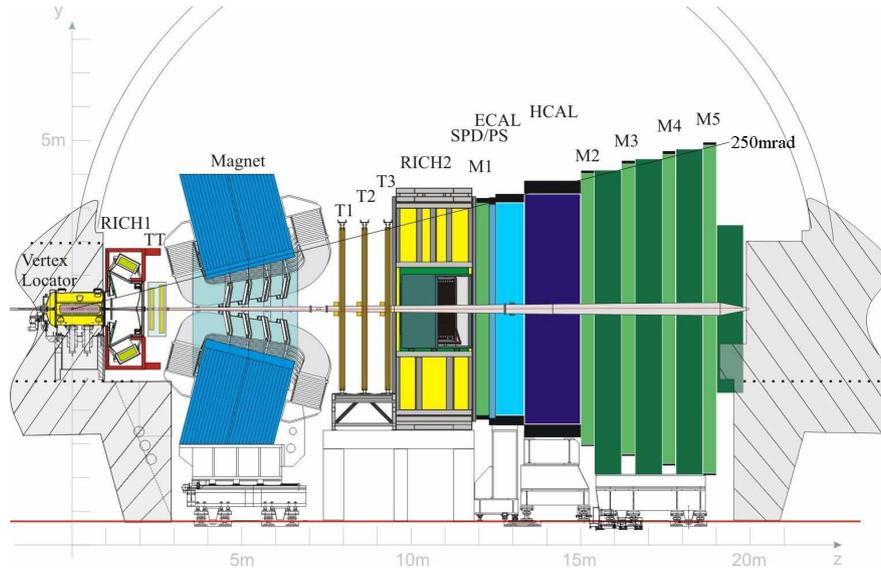


Figure 1: The LHCb detector with the subdetectors for the identification of particles and reconstruction of their tracks [1].

momentum, the particles mass can be computed which is a major indicator of a particles type. The track momentum is measured by the two Ring Imaging Cherenkov detectors RICH 1 and RICH 2, one is located upstream and the other is located downstream of the magnet, and they measure in two different momentum ranges. Inside the RICH detectors, particles are travelling faster than the speed of light of the RICH medium, and therefore a Cherenkov light cone is created at an angle that is directly related to the particles momentum. In the electromagnetic calorimeter (ECAL), the energy of interacting particles is measured. Particles like electrons and photons produce showers in the ECAL medium and lose all of their initial energy in this process. This energy loss is estimated by the amount of photons that are created in the ECALs scintillating layers. Muons and charged hadrons, like pions or protons traverse the ECAL without losing a significant fraction of their energy. The hadronic calorimeter (HCAL) is located after the ECAL and measures the energy of hadrons that produce showers of photons and gluons. Muons generally traverse the whole detector without losing a significant fraction of their initial energy. In the muon chambers M1 to M5 the tracks hits produced by muons are measured and allow in combination with the remaining LHCb tracking components, a very precise measurement of the muon momentum.

In this analysis, the CP -violation parameters S and C are measured in $B^0 \rightarrow J/\psi(\rightarrow \mu\mu)K_S^0$, $B^0 \rightarrow J/\psi(\rightarrow ee)K_S^0$ and $B^0 \rightarrow \psi(2S)(\rightarrow \mu\mu)K_S^0$ decays which are also referred to as *golden modes* for the measurement of indirect CP violation. This is due to the fact, that higher-order decay processes like penguin decays with non- CP violating phases only add small corrections to the tree-level decay [6]. Since the LHCb detector is reaching

the precision necessary to measure these effects, the CP -violation parameters ϕ_s^{eff} , C_s and $\mathcal{A}_{\Delta\Gamma_s}$ are measured in $B_s^0 \rightarrow J/\psi (\rightarrow \mu\mu) K_S^0$ in addition. These are closely related to the formerly mentioned decays and this measurement will help, in combination with other analyses, to constrain penguin effects for the measurement of S [5]. This allows for an exceptionally clean measurement of the CP -violation parameter S , a constant parameterizing, in part, the CP -violating nature of weak interactions. In these decays, CP violation occurs in the interference of mixing, i.e. the time dependent oscillation of B mesons between their matter and antimatter states B^0 and \bar{B}^0 and their decay into a CP -invariant final state. The CP -invariant final state analysed in the scope of this analysis is a charmonium state ψ consisting of two charm quarks and a strange meson K_S^0 consisting of a strange quark and a down quark. The charmonium state J/ψ is reconstructed from either two electrons or two muons, while the charmonium state $\psi(2S)$ is only reconstructed from decays into two muons. The K_S^0 meson is in all channels reconstructed from decays into two oppositely charged pions.

The CP -violation parameter S is determined by measuring the time dependent CP -asymmetry which is defined as

$$\mathcal{A}_{CP}(t) \equiv \frac{\Gamma(\bar{B}^0(t) \rightarrow \psi K_S^0) - \Gamma(B^0(t) \rightarrow \psi K_S^0)}{\Gamma(\bar{B}^0(t) \rightarrow \psi K_S^0) + \Gamma(B^0(t) \rightarrow \psi K_S^0)} = \frac{S \sin(\Delta m t) - C \cos(\Delta m t)}{\cosh(\frac{\Delta\Gamma t}{2}) + A_{\Delta\Gamma} \sinh(\frac{\Delta\Gamma t}{2})}. \quad (1)$$

Since the aforementioned higher-order effects are small, S is closely related to a fundamental physical constant β , which is one of the parameters needed to measure the size of CP violation in the universe: $S \approx \sin(2\beta)$.

According to eq. (1), the goal of this analysis is therefore to select as many signal decays as possible while achieving a good background rejection, and identifying the B flavour at production. Background is rejected, with the help of machine learning algorithms like boosted decisions trees. Boosted decision trees are supervised learning algorithms that are trained to separate a dataset into classes. In addition, the B flavour at production needs to be determined. This is achieved by measuring the charges of particles that are created in the hadronisation of either the signal B meson or the signal B 's partner-beauty hadron. This is possible since b quarks are created in pairs. A simplified schematic of a $b\bar{b}$ production is shown in figure 2. This is an especially challenging task which generally significantly decreases the effective size of the data sample. Therefore, a huge effort is made to utilize existing LHCb flavour tagging algorithms to their fullest potential. In addition, a new kind of flavour tagging algorithm is being developed that relies on deep neural networks and elements of long short-term memory units like gated recurrent units. This new method aims at interpreting a full proton-proton collision in order to interpret it as a whole instead of focusing on specific decay signatures in the event. The results that have been achieved so far reach the performance of the classical method and could supersede its performance in the near future.

The final goal of this analysis is to measure the CP violation parameter S in a simultaneous fit to all discussed modes. This measurement is going to result in the most precise single

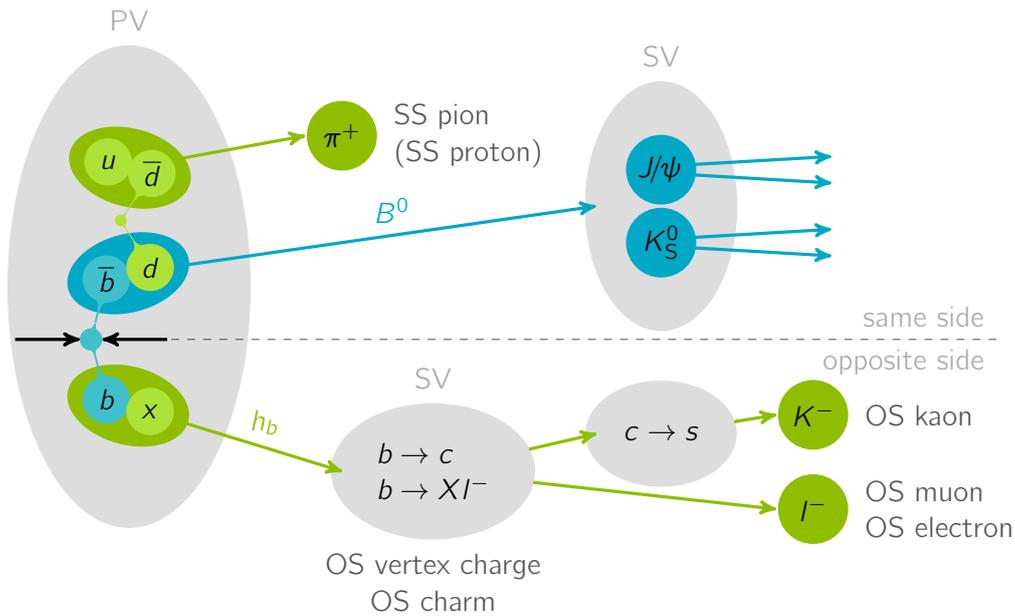


Figure 2: Illustration of the event topology of a $b\bar{b}$ production [7].

determination of the CP -violation parameter S to date and the measurement precision will supersede the current world average. Previous measurements by LHCb that have been performed using the Run 1 dataset [2–4] are likely going to be superseded in precision.

References

- [1] A. A. Alves, Jr. et al. The LHCb detector at the LHC. *JINST*, 3:S08005, 2008.
- [2] R. Aaij et al. Measurement of CP violation in $B^0 \rightarrow J/\psi K_S^0$ decays. *Phys. Rev. Lett.*, 115:031601, 2015.
- [3] R. Aaij et al. Measurement of the time-dependent CP asymmetries in $B_s^0 \rightarrow J/\psi K_S^0$. *JHEP*, 06:131, 2015.
- [4] R. Aaij et al. Measurement of CP violation in $B^0 \rightarrow J/\psi K_S^0$ and $B^0 \rightarrow \psi(2S)K_S^0$ decays. *JHEP*, 11:170, 2017.
- [5] Kristof De Bruyn and Robert Fleischer. A Roadmap to Control Penguin Effects in $B_d^0 \rightarrow J/\psi K_S^0$ and $B_s^0 \rightarrow J/\psi \phi$. *JHEP*, 03:145, 2015.
- [6] Philipp Frings, Ulrich Nierste, and Martin Wiebusch. Penguin contributions to cp phases in $B_{d,s}$ decays to charmonium. *Phys. Rev. Lett.*, 115:061802, Aug 2015.
- [7] J. T. Wishahi. <https://gitlab.cern.ch/lhcb-ft/TikZpics>, 2014.

Resource-aware Scientific Data Analysis on Energy-Efficient Storage Clusters

Thomas Lindemann
Databases and Information Systems Group (DBIS)
TU Dortmund University
thomas.lindemann@cs.tu-dortmund.de

The main objective of our research is to get over hardware restrictions in processing capabilities due to electrical power consumption or thermal discharge constraints. Our current work aims to handle high volume scientific data on energy-efficient modern hardware clusters.

To validate our approaches, we are processing real world scientific use cases on our custom-assembled experimental hardware clusters using modern hardware architectures. The use cases get provided by our collaboration partners of the particle physics department in the SFB876-C5 collaboration. We aim to improve the processing efficiency from different approaches by adjusting the hardware configuration, the algorithms and data storage patterns, while metering the execution time and energy consumption compared to existing solutions.

Since the last report, we worked on improving resource-aware data storage side for the scientific analyzing side. With our newly introduced data storage platform, we created an experimental approach of a hardware-software-co-designed storage solution for fast data analyzing. We set the focus especially on fast processing of multi user requests on data subsets with distributed storage, what covers the observed physicist's workflow.

The hard- and software of our design is matched to each other to achieve an excellent power profile, but it is not exclusively required. The newly built hardware cluster is and the distributed data storage software framework is also able to run on other cluster and processor architectures.

1 MAGPIE: Distributed Data Storage Analysis Approach

To validate our approach, we created a small hardware platform prototype in shape of a 2U 19" Blade with integrated embedded storage nodes, consisting of embedded x86_64 based systems, SSD storage devices, cooling system and power supply with a custom designed integrated power control and energy consumption metering.

The system is by design scalable by just adding more nodes or similar node blades connected by the network back-end. We have decided to use very power efficient embedded systems for our evaluation system, because based on previous experiments, we observed that high performance processors have little advantage for the scan- and filter operations in the storage nodes, assuming the operation is memory .

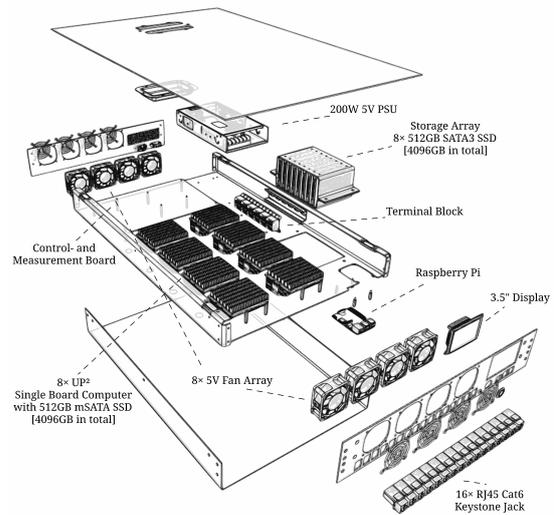


Figure 1: MAGPIE Hardware Assembly

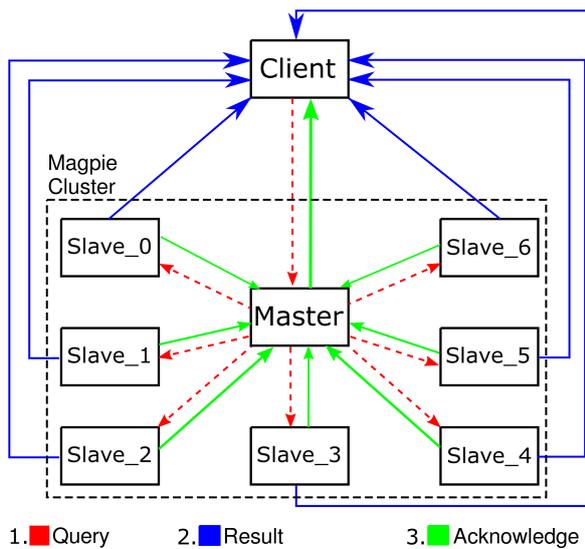


Figure 2: MAGPIE Software Architecture

To avoid a bottleneck on the master node, the data is sent directly to the requesting client, the master node acts only as a pool server for all active data slave nodes, it forwards the request and gets the acknowledge, redundancy or load distribution is possible. The currently used solution of data analysis in the C5 project is time intensive, as the users have to wait for stripping line creation several days and comes along with data redundancy what is increasing the overall needed amount of storage space.

Our new approach is different to the existing storage solutions. Instead of querying the data from a single storage device, the data is distributed as horizontally partitioned parts among a number of storage nodes and is received in parallel. Furthermore, the many embedded controllers on the storage nodes utilize pre-filtering properties. So that, a request is not only done on a specific data stream but combined with a predicate and a projection on needed columns of the data set. Moreover, Aggregation of the data on the storage nodes is supported as well. Figure 2 is visualizing the query processing schema of the MAGPIE storage cluster.

Instead of pre-processing the data analyst’s requests for a specific stripping line and materialize it for the requesting entity the final selection could be done on demand. This doesn’t restrict the idea of working on different subsets of the very large experiment data and buffer these locally, it is rather supported by the distributed architecture of the approach, only the subsets could be larger and more universal. Due to the distributed character of the system with parallel processing, for frequently requested data, there is also the option for creating replicates on several storage nodes in the system and balancing the load.

2 Experimental Evaluation

2.1 Data Query Evaluation

We evaluated the performance of MAGPIE and the comparison Postgres server system. With variants of tests with different selectivity levels, we are aiming to cover several data pre-filtering scenarios which are able to do a big part of the data pre-processing on the distributed slave nodes and less post-processing on the client. The plots in Figure 3 shows the results on MAGPIE and the comparison system.

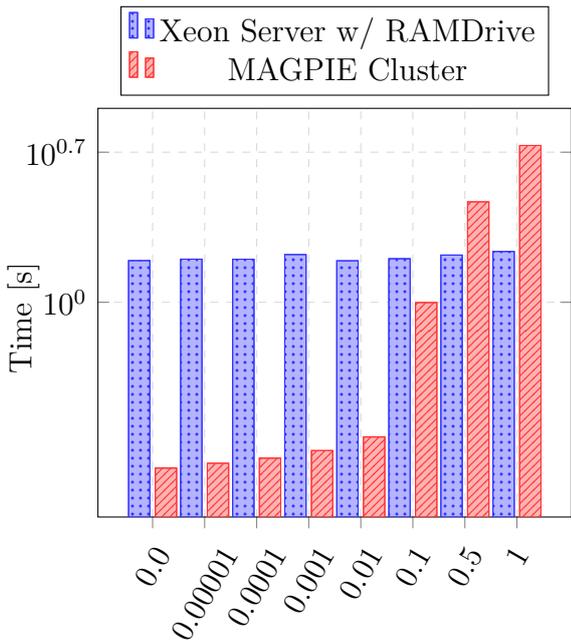


Figure 3: Execution Time of Queries Comparison on MAGPIE vs. ref. Server w/ PostgreSQL

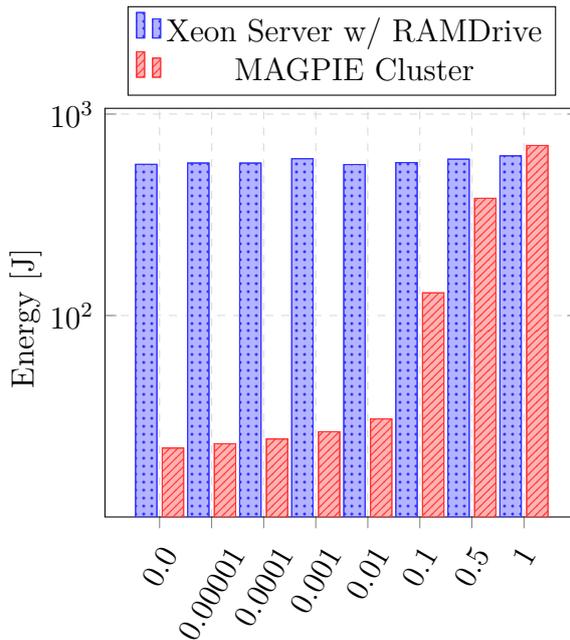


Figure 4: Energy Consumption of Queries Comparison on MAGPIE vs. ref. Server w/ PostgreSQL

Our experiments confirm the approach, that our new proposed distributed system architecture has indeed an improvement in the performance/power consumption ratio. As seen in Figure 3 and 4, the execution time is much lower for most selectivity rates.

The high execution time at high tuple numbers is caused by a bottleneck of the processor capacity in tuple processing. For that, we are improving the scan in our ongoing work as described in section 3. On contrast, the reference server shows a constant processing time over varying selectivity, as it is working on a ramdisk drive and furthermore it doesn't send the results over a network connection in this test.

2.2 Scalability Analysis

In the next experiment, we evaluated the scalability of the system, as this is a key point of the design of the approach. Figure 5 is showing the execution time and the energy consumption over an increasing number of active worker nodes at a constant data size. As assumed is the execution time proportional falling with the number of worker nodes while the energy consumption is constant.

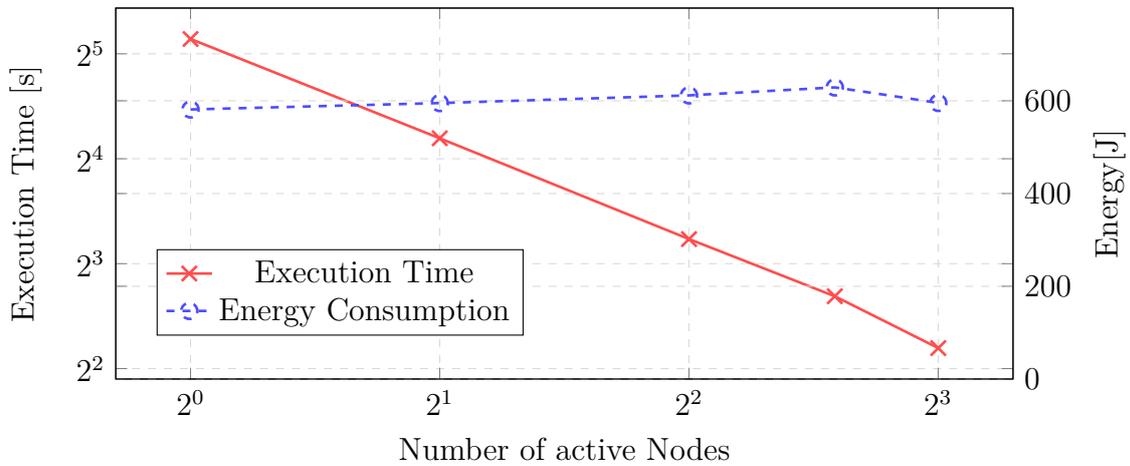


Figure 5: Aggregation over all values in one column with different number of worker nodes on 80 million lines

3 Ongoing Work

Our experiments have shown that our distributed storage approach with storage near pre-filtering can indeed increase the overall scan speed and results in lower network traffic in case of external storage. Moreover, this topology removes RAM and CPU constraints of a single machine solution and allows scaling up with the data sizes.

The current ongoing work is about more efficient inserts and updates as unordered transactions by multiple clients on the partitioned data with a more efficient balancing of parallel accesses on redundant replicated data and doing asynchronous read and write operations to the flash storage disks, what can improve data reading speed while significantly lowering the CPU-activity.

Moreover, we are currently evaluating on GPU-accelerated full table scans. The use of the available GPUs on the worker nodes can relieve the system's CPUs and make it even more efficient.

Measurement of CP violation in $B^0 \rightarrow J/\psi K_S^0$ decay with $J/\psi \rightarrow e^+ e^-$ and $K_S^0 \rightarrow \pi^+ \pi^-$

Gerwin Meier
Lehrstuhl für Experimentelle Physik 5
Technische Universität Dortmund
gerwin.meier@tu-dortmund.de

The European Organization for Nuclear Research (CERN) is located near Geneva at the border between Switzerland and France. There, the largest machine of the world, the Large Hadron Collider (LHC) is operated. At the particle detectors of the LHC, the aim is to study and constrain the Standard Model (SM), the most accurate model to date for particle physics, and search for physics beyond the SM, so called "New Physics". The LHCb experiment is one of the four main experiments at the LHC and study high precision measurements with b- and c-mesons. That includes charge-parity (CP) violating decays, one of three necessary conditions to explain the observed matter-antimatter asymmetry [4] in the universe, and rare decays.

In proton-proton collisions, b quarks are produced in pairs mainly in a small cone around the beam line, and therefore, the LHCb detector measures in the forward region and is a one-arm forward spectrometer, as is shown in Fig. 1.

Closest to the beam is the Vertex Locator (VELO), which is used to identify the proton-proton collision point and the decay vertices of the produced particles with a small life time. The tracker turicensis (TT), the tracking stations (T1-T3) and the muon chambers (M1-M5) are used to reconstruct the tracks of the particles and determine the momentum in combination with the LHCb magnet. The particle type is determined with the help of the ring-imaging Cherenkov detectors (RICH1-RICH2) and the muon chambers. Finally the energy is measured with the calorimeter system, including the scintillating

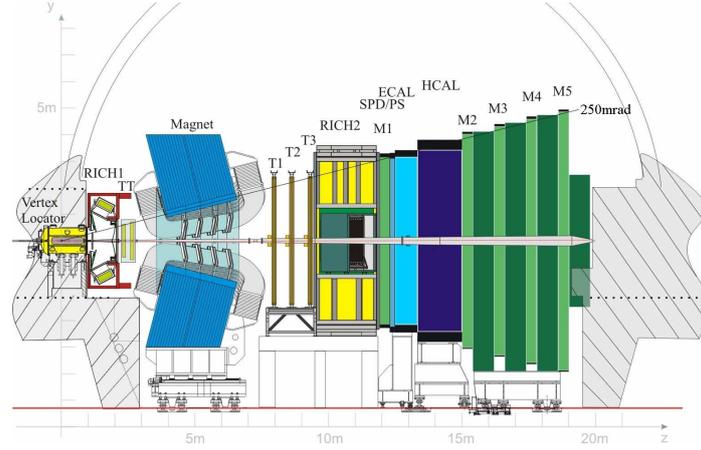


Figure 1: The LHCb detector with the various subdetectors for the identification of particles and reconstruction of their tracks.

pad detector (SPD), preshower (PS), electromagnetic and hadronic calorimeter (ECAL, HCAL).

There are several challenges for the reconstruction of all particles. In proton-proton collisions, many particles are produced and therefore there are thousands of hits in the detector, which makes it very challenging to reconstruct the hundreds of tracks, because the number of possible combinations of the hits to form a track increase quadratically with the number of hits. Another challenge is the high frequency with one proton-proton bunch collision every 25 ns and therefore the very short time to decide whether to save or reject the event, a so called trigger decision. The amount of data that needs to be recorded is very large and a huge effort is made in finding ways to efficiently and safely store huge amount of data.

A measurement of time-dependent CP violation in the golden mode $B^0 \rightarrow J/\psi K_S^0$ is performed, to show the necessity of handling the data in an efficient way. The J/ψ is reconstructed with two electrons and the K_S^0 with two oppositely charged pions. The analysis used the data collected with the LHCb detector in the second data taking period from 2015 to 2018, which doubles the number of signal candidates from the first period in 2011 and 2012 and where the center-of-mass energy was increased to $\sqrt{s} = 13$ TeV. The analysis from the first period is already published in [1].

Neutral mesons like B^0 can oscillate into their antiparticle and the analysed final state $J/\psi K_S^0$ is common for both mesons. Therefore a decay-time-dependent CP asymmetry in the interference between the amplitudes of the direct decay and the decay after B^0 - \bar{B}^0 mixing can be measured:

$$\mathcal{A}(t) \equiv \frac{\Gamma(\bar{B}^0(t) \rightarrow J/\psi K_S^0) - \Gamma(B^0(t) \rightarrow J/\psi K_S^0)}{\Gamma(\bar{B}^0(t) \rightarrow J/\psi K_S^0) + \Gamma(B^0(t) \rightarrow J/\psi K_S^0)} = \frac{S \sin(\Delta m t) - C \cos(\Delta m t)}{\cosh(\frac{\Delta\Gamma t}{2}) + A_{\Delta\Gamma} \sinh(\frac{\Delta\Gamma t}{2})}. \quad (1)$$

Here, B^0 and \bar{B}^0 indicate the flavour of the B meson at production, while t indicates the decay time. The parameters Δm and $\Delta\Gamma$ are the mass and the decay width differences between the heavy and light mass eigenstates of the $B^0-\bar{B}^0$ system, and S , C , and $A_{\Delta\Gamma}$ are CP observables. As $\Delta\Gamma$ is negligible for the $B^0-\bar{B}^0$ system [2], the time-dependent asymmetry simplifies to $\mathcal{A}(t) = S \sin(\Delta m t) - C \cos(\Delta m t)$.

The first step in the analysis is a signal and background selection. One background component comes from wrongly reconstructed particles, e.g. the J/ψ is reconstructed from two random electrons and not the two electrons, the J/ψ was decaying into. Another type of background are partially reconstructed backgrounds, physical decays where one or more particles are not reconstructed, like $B^0 \rightarrow J/\psi K^*(\rightarrow K_S^0(\rightarrow \pi^+\pi^-)\pi^0)$ where the neutral pion is lost. However, these backgrounds have different kinematical or topological properties. Therefore a cut based method on these variables is used at the beginning. Afterwards some special backgrounds are reduced like $B^0 \rightarrow J/\psi K^*(892)^0(\rightarrow K^\pm\pi^\mp)$, where the kaon is misidentified as a pion. These background are reduced with a decay time requirement for the K_S^0 to be more than 0.5 ps due to the fact, that the K^* is decaying almost immediately. The last step in the selection is a multivariate approach. A supervised learner is used, where a signal simulation is used for the signal proxy and wrongly reconstructed data as a background proxy. The data is used to train a boosted decision tree (BDT) [3], which uses a couple of decision trees to decide the signal likeliness for each candidate. After the BDT is applied to the full dataset and determined the signal likeliness, a cut on this prediction is performed in a way, that as many signal candidates as possible survives and as many background candidates as possible are rejected.

After this huge effort to reduce the background there are still some background candidates remaining. Therefore the sFit method [5] is used to unfold the signal shape. That is done with a fit to the invariant mass of the B^0 meson. The signal and each background component is modelled with a function and then a fit to the data is performed. Afterwards the signal function parameters described the signal component very well, which can be seen in Fig. 2, and one obtains about 50000 signal candidates. The signal parameters are then used to calculate signal weights, which can be used to extract the signal out of all data in all variables. This holds for variables uncorrelated to the mass of the B^0 meson like the decay time, which is needed to calculate the CP -parameters.

The next important step for the measurement is the determination of the flavour of the initial meson, *i.e.* knowing whether it is a B^0 meson or its antiparticle \bar{B}^0 . Therefore algorithms are used, which predict the flavour of the meson. The training of these algorithms is performed with decays, where the charge of the final state particles differ between the flavour which makes it easier to determine the initial flavour, so-called calibration channels. However, this means that the calibration channels need to be adjusted to match for the signal decay, *e.g.* in the kinematic variables. With the training of a BDT it is possible to take all the correlations of these variables into account and use the flavour tagging algorithms for the signal channel, although they are trained on the

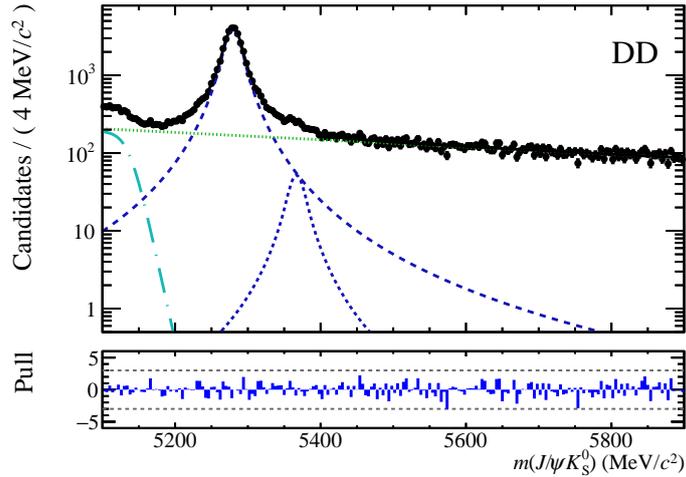


Figure 2: Massfit on selected data with the B^0 signal component in blue and long-dashed, the B_S^0 background in blue dashed, the wrong reconstructed background in green and dashed and the partial background in cyan and long-dashed-dotted. The DD indicate, that the K_S^0 decay after the vertex locator.

calibration channel. With all these input it is possible to achieve a sensitivity which is around the world average for the CP asymmetry parameter S .

References

- [1] R. Aaij et al. Measurement of CP violation in $B^0 \rightarrow J/\psi K_S^0$ and $B^0 \rightarrow \psi(2S)K_S^0$ decays. *JHEP*, 11:170, 2017.
- [2] Y. Amhis et al. Averages of b -hadron, c -hadron, and τ -lepton properties as of summer 2016. *Eur. Phys. J.*, C77:895, 2017. updated results and plots available at <https://hflav.web.cern.ch>.
- [3] Tianqi Chen and Carlos Guestrin. XGBoost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, pages 785–794, New York, NY, USA, 2016. ACM.
- [4] A. D. Sakharov. Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe. *Pisma Zh. Eksp. Teor. Fiz.*, 5:32–35, 1967. [JETP Lett.5,24(1967); Sov. Phys. Usp.34,no.5,392(1991); Usp. Fiz. Nauk161,no.5,61(1991)].
- [5] Yuehong Xie. sFit: A method for background subtraction in maximum likelihood fit. 2009.

Studies of bottom meson decays to two open charm mesons

Antje Mödden

Lehrstuhl für Experimentelle Physik 5

Technische Universität Dortmund

antje.moedden@tu-dortmund.de

The LHCb experiment is one of the four big experiments located at the Large Hadron Collider near Geneva, Switzerland [1]. Its main focus is the search for New Physics effects in CP -violating and rare decays of beauty and charm hadrons. In order to maximize the sensitivity with respect to these specialized targets the LHCb detector is built as a forward single arm spectrometer (see Figure 1).

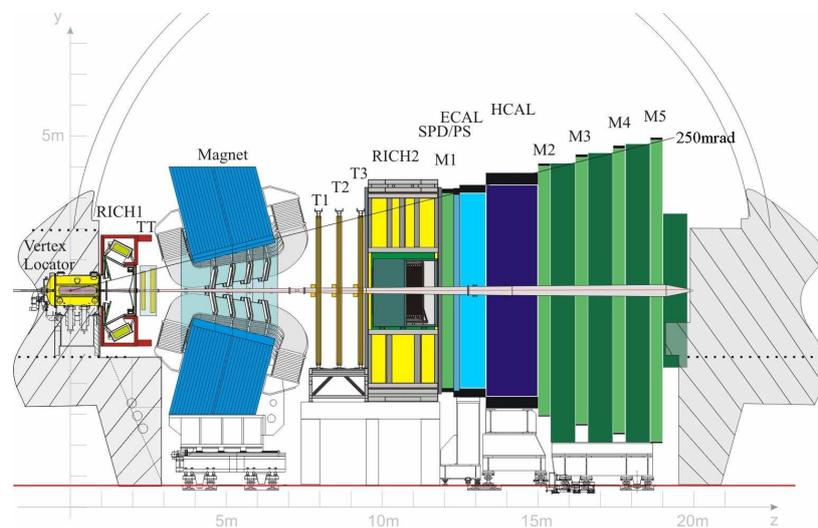


Figure 1: Scheme of the LHCb detector, illustrating the various subdetectors for identification and reconstruction of particles and their tracks. [1]

In the collisions of protons inside the vertex locator (VELO), a large amount of particles is produced by many different physical processes. These particles decay into new particles, which fly through the detector and leave traces in the various subcomponents by interacting with the detector material. These traces are hits in the tracking systems (VELO, TT, T1-T3, M1-M6), clusters in the calorimeters (ECAL, HCAL) and Cherenkov radiation in the Ring Imaging Cherenkov Detectors (RICH1, RICH2). To conclude on the presence of particles the information of these subdetectors needs to be reconstructed, e.g. by fits of trajectories to ensembles of hits and pattern recognition algorithms looking for clusters of energy deposition. Finally particle candidates need to be combined to heavier particles in order to perform physics measurements on the same.

The endeavour to find the particles of interest is hindered for two main reasons. Firstly, hundreds of particles are produced inside the angular acceptance which allows for a large number of combinations to be made in the reconstruction as well as the combination of particles. Secondly, the interaction rate of 50ns / 25ns together with the limitation on the bandwidth that can be written to disk enforces a fast reconstruction that leads to the selection of interesting events and the rejection of physically uninteresting ones. Both these points together set the frame for investigating these tasks in the context of resource limitation: The reconstruction and combination tasks can be parallelized and therefore performed faster.

To show the necessity of handling the data in an efficient way, two physics analyses are being performed.

1 Observation of the decay $B_s^0 \rightarrow D^{*\pm} D^\mp$

In this analysis the $B_s^0 \rightarrow D^{*\pm} D^\mp$ decay is observed and its branching ratio relative to the $B^0 \rightarrow D^{*\pm} D^\mp$ decay is measured. The decay $B_s^0 \rightarrow D^{*\pm} D^\mp$ has not been previously observed, but in a measurement of CP violation in $B^0 \rightarrow D^{*\pm} D^\mp$ [2] possible $B_s^0 \rightarrow D^{*\pm} D^\mp$ candidates have been seen.

To calculate the relative branching ratio, signal candidates for both decay channels have to be selected and the selection efficiencies have to be determined. An external parameter is the ratio of the b hadronisation fractions to B^0 and B_s^0 mesons, whose precision results in the dominant source of systematic uncertainty on the measurement. The selection efficiencies are calculated with simulated data samples and the signal candidates are determined by a fit to the mass distribution of the $D^{*\pm} D^\mp$ system.

A challenging part of this analysis is the identification and reduction of background candidates, which dominate the recorded data set. Due to wrong reconstruction or misidentification, these so-called background decays can mistakenly be present in the data set and need to be removed in the selection.

The first step of the selection is a centrally organized preselection for all analyses performed by the LHCb collaboration. This takes several weeks. Afterwards, the individual tuple of the size of several hundred GB is produced. The selection then proceeds with an offline cut-based preselection, which mainly consists of kinematical and geometrical requirements. In this context, rectangular means that certain measured quantities of the event have to lie in ranges. These ranges are well-defined and do not vary depending on other quantities in the event. This is not the case in multivariate analyses that are performed afterwards. At the end of the selection, background contributions are still present in the data set, but their amount has been reduced enormously and they can be modelled in the mass fit (see Figure 2).

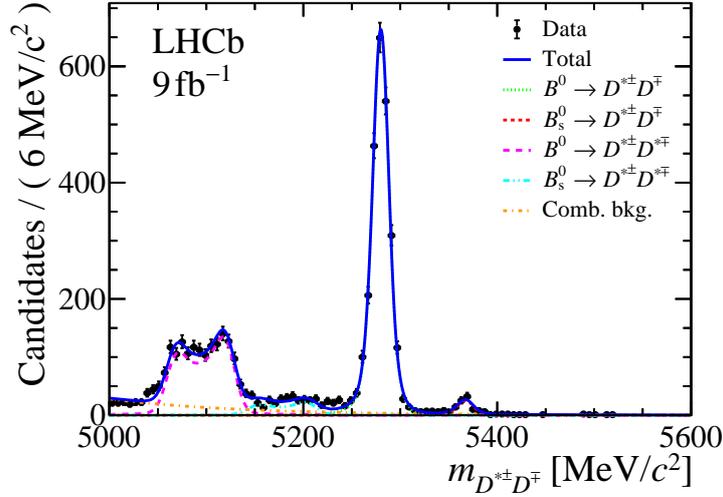


Figure 2: The $D^{*\pm} D^\mp$ mass distribution. The total fit projection is shown as the blue solid line. The green dotted and red dashed line correspond to the B^0 and B_s^0 signal contributions, respectively. The other lines describe different background contributions.

The mass fit results in 2509 $B^0 \rightarrow D^{*\pm} D^\mp$ and 95 $B_s^0 \rightarrow D^{*\pm} D^\mp$ candidates. The $B_s^0 \rightarrow D^{*\pm} D^\mp$ decay is observed with a high significance calculated using Wilk's theorem [4].

The ratio of branching fractions is determined to be

$$\frac{\mathcal{B}(B_s^0 \rightarrow D^{*\pm} D^\mp)}{\mathcal{B}(B^0 \rightarrow D^{*\pm} D^\mp)} = 0.137 \pm 0.017 \pm 0.002 \pm 0.006,$$

where the first uncertainty is statistical, the second systematic and the third is due to the uncertainty of the hadronisation fraction ratio f_s/f_d . The $B_s^0 \rightarrow D^{*\pm} D^\mp$ branching fraction is determined to be

$$\mathcal{B}(B_s^0 \rightarrow D^{*\pm} D^\mp) = (8.41 \pm 1.02 \pm 0.12 \pm 0.39 \pm 0.79) \times 10^{-5},$$

where the fourth uncertainty is due to the $B^0 \rightarrow D^{*\pm} D^\mp$ branching fraction [5].

The analysis resulted in a paper, which will be published on arXiv and submitted to Journal of High Energy Physics before the end of 2020.

2 Time-dependent analyses of $B_s^0 \rightarrow D_s^+ D_s^-$ and $B^0 \rightarrow D^+ D^-$ decays

Decay-time dependent analyses of the $B_s^0 \rightarrow D_s^+ D_s^-$ and $B^0 \rightarrow D^+ D^-$ decays allow the measurement of the weak mixing phases, ϕ_s and β , respectively. The analyses of both decays involve a very similar candidate selection and subsequent extraction of decay-time-dependent parameters, allowing both analyses to be combined.

The selection, where signal and background decays are separated, uses similar methods as are used to observe the $B_s^0 \rightarrow D^{*\pm} D^\mp$ decay. Good control over the different backgrounds is achieved, but small background contributions are still present in the data set after the selection. A statistical background subtraction is performed with the sPlot method [3].

The next steps include the calibration of Flavour-Tagging performance parameters, the parametrisation of the decay-time resolution and a decay-time dependent efficiency and the extraction of decay-time dependent parameters to calculate the mixing phases. Work on these three parts is ongoing and know-how from previous analyses is exploited. To finalise the analysis, the size of systematic uncertainties due to different sources have to be determined.

References

- [1] A. A. Alves, Jr. et al. The LHCb detector at the LHC. *JINST*, 3:S08005, 2008.
- [2] Roel Aaij et al. Measurement of CP violation in $B^0 \rightarrow D^{*\pm} D^\mp$ decays. 2019.
- [3] Muriel Pivk and Francois R. Le Diberder. sPlot: A statistical tool to unfold data distributions. *Nucl.Instrum.Meth.*, A555:356–369, 2005.
- [4] S. S. Wilks. The large-sample distribution of the likelihood ratio for testing composite hypotheses. *Ann. Math.Stat.*, 9:60–62, 1938.
- [5] P. A. Zyla et al. Review of particle physics. *to be published in Prog. Theor. Exp. Phys.*, 6:083C01, 2020.

GPU accelerators in the future

LHCb-Triggersystem

Holger Stevens
Lehrstuhl für Experimentelle Physik 5
Technische Universität Dortmund
holger.stevens@tu-dortmund.de

The LHCb experiment entered an upgrade phase in Dezember 2018, which will last 2 years. Not only the detector but also the computing farm will be changed. The planning for this started 10 years ago. A certain amount of computing power was planned to be bought by a fixed budget, but meanwhile "Moore's Law died" so the plan cannot be converted. There are different solution approaches in the LHCb collaboration to reach the necessary computing power, one is the usage of GPUs.

Computing tasks can only benefit from the multiprocessor architecture of a GPU, if they are massively parallelizable. Different algorithms are running in the online farm to reconstruct the events. Most of them are independent and parallelizable. On the one hand it is the combination of hits to reconstruct the tracks on the other hand the processing of RAW information.

This report will describe the decoding of the so-called RAW banks of the SciFi tracker. These contain all information of the channels which detect a signal. The known channel IDs need to be translated to physical positions in the detector.

1 Introduction

The LHCb experiment is one of the four big experiments located at the Large Hadron Collider (LHC) near Geneva, Switzerland. Its main focus is the search for rare decays and effects of CP -violation in decays of beauty and charm hadrons [1]. Due to some physical constraints in the production of b and c quarks through proton proton collisions the

LHCb detector is designed as a single-arm forward spectrometer. Over the past years the understanding of the detector and its systematical effects has reached an almost perfect level. At the moment, the most limiting factor for analyses is the statistical uncertainty. The only way to improve this is to massively increase the dataset. For this reason, as mentioned before an upgrade of the experiment is foreseen [6]. Most of the

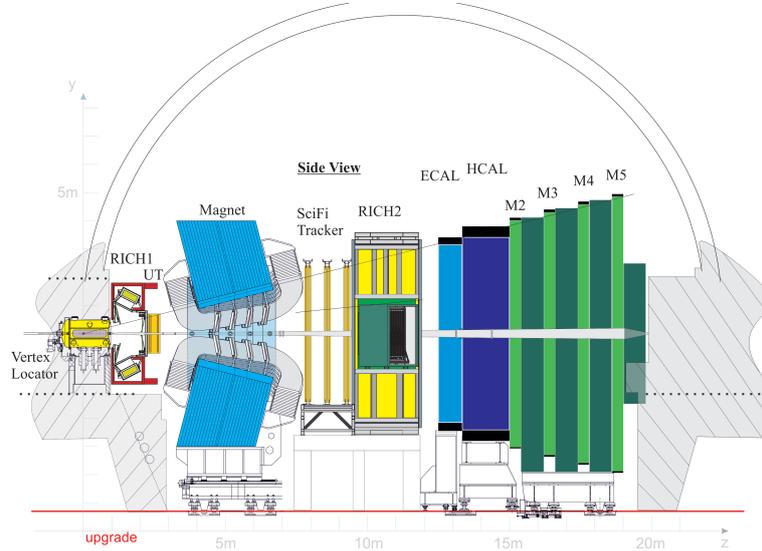


Figure 1: The LHCb upgrade detector with the various subdetectors for the identification of particles and reconstruction of their tracks [3].

existing detector will be replaced, but the general structure of the detector will remain the same. In the Vertex Locator (Velo) the position of the primary interaction is detected. The Upstream Tracker (UT) and the SciFi Tracker also belong to the tracking system. Other components like the Ring Imaging Cherenkov Detectors (RICH), the Electronic Calorimeter (ECAL), the Hadronic Calorimeter (HCAL) and the Muon Chambers (M2-M5) are used for the particle identification.

A major change is the new triggerless readout of the full detector with 40 MHz. This leads to a data rate of 40 Tb/s which is the input of the high level trigger (HLT) [7].

2 Details of the SciFi Tracker

Before the actual decoding part is described, some technical details of the SciFi Tracker are explained to understand the data format.

In total the Tracker consists of 3 stations with 4 layers each. These layers are built out of so-called modules, which again are composed of 8 fibre mats. Each mat is connected to 4 SiPM arrays. These are connected to the so-called FrontEnd electronics (FE). An FPGA based algorithm is looking for clusters of hits. A real signal has a different

width and height than a background hit, caused for example by thermal noise. Only the information from found clusters, so-called zero suppressed data, is sent from the FE to the BE (BackEnd electronics). The BE has to pack the data from several FE boards and send it to the computing farm where the actual decoding will take place. It is important to mention that the stations have a different width, T1 and T2 consist of 10 modules, and T3 is larger with its 12 modules. This difference leads to a branching of the decoding code, which causes a significant longer runtime.

3 Status and outlook

The aim of one working group of the LHCb collaboration is to run the first Trigger stage (HLT1) on GPU [5]. The project, Allen, can be found at <https://gitlab.cern.ch/lhcb/Allen>. Multiple streams can send data to the GPU. Then a visitor service will load the necessary Algorithms. One of these is the SciFi hit decoding. Every data-package has a global header which contains information about the associated proton-proton collision, like a timestamp. This information is followed by detector component-specific informations, for the SciFi this are the station number, MatID etc. with these information a position in the detector is unambiguously defined. The actual numbering scheme is constant for a specific time and can be a loud as a constant in the code [8]. During the development of the decoder, the data-format was changed to reduce the needed bandwidth. The amount of Hits is no longer in the header [4]. To allocate the correct size of memory for the result vector of decoded Hits, a pre count is necessary.

Successfully a GPU version of the SciFi decoder was implemented. Also, the other parts of the HLT1 sequence are ported to a GPU based version. The sequence can now be used to compare the performance to the baseline solution, which are standard x64 CPU server. The general idea and the status are presented in an article in *Computing and Software for Big Science* [2]. Different GPUs were tested to optimize the throughput of data, respectively Events. Figure 2 shows these performance studies. As the decision of using GPUs in LHCb or not wasn't taken yet, there is also no need to choose a specific GPU. It is possible to wait for the next generation ,which might have an even better cost/performance ratio. This can be expected as the performance of the sequence scales almost linear with the computing power (FLOPS) of the different GPUs.

References

- [1] Letter of Intent for the LHCb Upgrade. Technical Report CERN-LHCC-2011-001. LHCC-I-018, CERN, Geneva, Mar 2011.

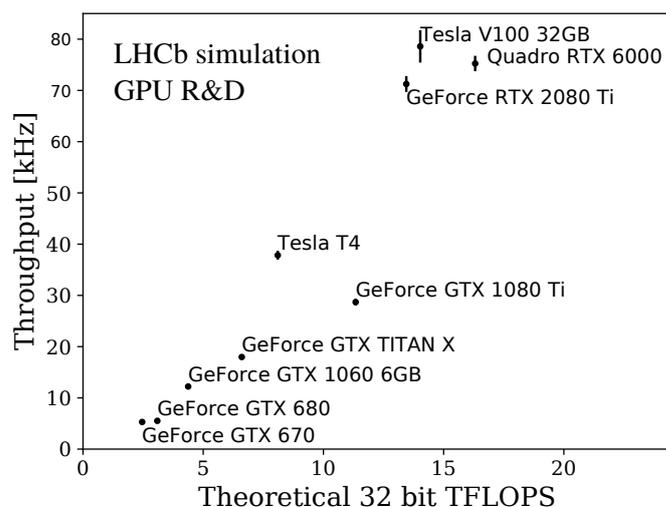


Figure 2: Allen throughput on various GPUs with respect to their reported peak 32-bit FLOPS performance.

- [2] Roel Aaij et al. Allen: A high level trigger on GPUs for LHCb. 2019.
- [3] LHCb Collaboration. LHCb Tracker Upgrade Technical Design Report. Technical Report CERN-LHCC-2014-001. LHCb-TDR-015, Feb 2014.
- [4] Sevda Esen, Jeroen Van Tilburg, Luigi Del Buono, Pierre Billoir, and Louis Henry. Clustering and rawbank decoding for the SciFi detector. Technical Report LHCb-INT-2018-024. CERN-LHCb-INT-2018-024, CERN, Geneva, Jul 2018.
- [5] Stefano Gallorini, Donatella Lucchesi, Alessio Gianelle, Silvia Amerio, and Marco Corvo. First experiences with a parallel architecture testbed in the LHCb trigger system. Technical Report LHCb-PUB-2017-015. CERN-LHCb-PUB-2017-015. 3, CERN, Geneva, Apr 2017.
- [6] Christian Joram. LHCb Scintillating Fibre Tracker Engineering Design Review Report: Fibres, Mats and Modules. Technical Report LHCb-PUB-2015-008. CERN-LHCb-PUB-2015-008, CERN, Geneva, Mar 2015.
- [7] CERN (Meyrin) LHCb Collaboration. Computing Model of the Upgrade LHCb experiment. Technical Report CERN-LHCC-2018-014. LHCb-TDR-018, CERN, Geneva, May 2018.
- [8] Jeroen Van Tilburg. SciFi readout numbering scheme. Technical Report LHCb-INT-2016-044. CERN-LHCb-INT-2016-044, CERN, Geneva, Nov 2016.