Masaryk University
Faculty of Informatics
Department of Computer Systems and Communications

Habilitation Thesis

Performance and Security in Scalable Distributed Systems: from Cellular Mobility to Cloud-Native Architectures

March 2026

Mgr. Lukáš Kencl, PhD

Declaration

I hereby declare that I have completed this thesis independently and that I have listed all the literature and publications used. Works presented here have been elaborated, with one exception (DNA-inspired information concealing) primarily with my PhD and Master students, at times together with other supporting partners.

If not mentioned otherwise, equal contribution may be assumed. To the best of my knowledge, my work does not infringe upon anyone's copyright nor violate any proprietary rights and that any ideas, techniques, quotations, or any other material from the work of other people included in my thesis, published or otherwise, are fully acknowledged in accordance with the standard referencing practices.

During the elaboration of this thesis, I have occasionally used generative AI tools, namely Google Gemini, for drafting article summaries, and polishing the text. I critically evaluated and modified all generated content. No AI tool has been used in generating any of the research results.

Acknowledgements

I would like to express my sincere gratitude for help, guidance and encouragement to doc. RNDr. Barbora Kozlíková, Ph.D. and prof. Ing. Tomáš Vojnar, Ph.D., for guiding me through the submission process and giving me an opportunity for habilitation.

I would also like to thank the Department of Telecommunication Engineering at FEE CTU, and its present and past heads, doc. Ing. Jiří Vodrážka, Ph.D. and prof. Ing. Boris Šimák, CSc., for fostering a highly collaborative environment, assembling a competent team and generously providing me with means and excellent environment for research and student supervision, as well as for always providing strong encouragement towards habilitation.

Furthermore, I wish to express my deep gratitude to my former supervisors and advisors, namely especially Prof. Jean-Yves Le Boudec, Prof. Refik Molva, Patrick Droz, Prof. Derek McAuley, Roberto Dorigo, Andrea Corda and Pierre Hallerstedt, for mentoring me and introducing me to the fields of distributed computing, network processing, cybersecurity, IoT and scalable clouds.

I would also like to thank my past students, Michal Ficek, Katerina Cizkova, Jan Stanek, Jiri Danihelka, Ondrej Tomanek, Pavol Mulinka, Jakub Klemsa, Tomas Pop, Jan Rudinsky, Zdenek Bumbalek, Zdenek Kouba, Vojtech Uhlir, Jakub Dolezal, Andrea Bergamini, Christian Schwarzer, Ramaswamy Ramaswamy, Ashish Gupta, Bozidar Radunovic, and many others, for strong technical contributions and great fun and collaboration throughout.

Sincere thanks also go to my inspirational colleagues, collaborators and mentors, especially prof. RNDr. Martin Loebl, CSc., prof. Dr. Ing. Michal Pěchouček, MSc., Prof. Milan Vojnovic, Prof. Sonja Buchegger, Prof. Gerhard Fettweis, Prof. Tanzeem Choudhury, Maja Vukovic, Prof. Milan Bjelica, Yoonho Choi, Weiguang Shi, Gianluca Milano, James Scott, Dina Papagiannaki, Tim Griffin, Sean Rooney, prof. Ing. Zdeněk Bečvář, Ph.D., Tomas Vanek, Pavel Strnad, Robert Bestak, Zbynek Kocur, Martin Rehak, Robert Haas, Andreas Kind, and many others.

I am also very grateful to the many industrial sponsors of my research, be it IBM, Intel, Microsoft, Cisco, Vodafone, Answers.com, iptel.org, CESNET, or, especially, Electrolux, for enabling this great innovation journey.

Supreme thanks go to the honorable Habilitation Committee and Thesis Reviewers, for dedicating time, energy and deep expertise to review my work.

Last, but certainly not least, I am very grateful to my family and, especially, to my wife, Jitka, for the eternal encouragement, without which this thesis would never have seen the light of day.

Preface

This habilitation thesis is conceived as a brief discourse covering a part of the author's career and his achievements. It is not intended as an overarching study on distributed infrastructure. Instead, it presents a peek into some of the most challenging aspects of the architecture of cellular networks and cloud-computing infrastructure of today. The thesis presents a body of research dedicated to the dual challenges of performance optimization and security assurance within large-scale distributed systems. Spanning a 15-year trajectory of technological evolution, the collected works transition from the analysis of physical mobility in cellular infrastructures to the engineering of dependable and observable cloud-native environments. The methodology of this work is strongly based on measured data from real, large, mobile-networking and cloud-computing infrastructures. Comprehensive measurements have been a necessary prerequisite for understanding these complex phenomena. Designing, implementing and operating the relevant measurement platforms is a complex problem on its own and original techniques how to address it represent a significant portion of this thesis.

It is recommended that the reader consult the main text together with the papers and manuscripts attached in Appendix A. For reader's convenience, all references directly enclosed in this thesis are explicitly stated as [Auth*]. List of all the publications enclosed can be found immediately after the Bibliography.

The habilitation thesis deals with some of the most challenging aspects of the architecture of the large-scale distributed infrastructures of today: measuring their performance, interpreting the measured data to build realistic mathematical models, predicting behavior using such models, and protecting the infrastructure and data. The content is arranged along three thematic pillars.

The first thematic pillar addresses Spatio-Temporal Network Intelligence. By leveraging Gaussian Mixture Models and advanced signaling analysis, this research refined coarse network metadata into high-fidelity mobility models. These contributions provided foundational insights into user-cell association and tracking, while simultaneously pioneering graph-based obfuscation techniques to address the emerging necessity for subscriber privacy in public telecommunications data.

The second pillar shifts the focus to Planetary-Scale Cloud Observability. As distributed computing migrated to globalized infrastructures, this research established empirical frameworks for auditing multidimensional cloud latency. Through the development of the CLAudit platform, we provided the first transparent benchmarks of global cloud service providers, revealing the intricate performance trade-offs inherent in multi-region deployments and Big Data frameworks such as Hadoop MapReduce.

The final pillar culminates in the design of Dependable and Secure Distributed Storage. Addressing the fundamental conflict between data deduplication efficiency and cryptographic privacy, this work introduced thresholded and popularity-based encryption schemes. These contributions, published in leading venues such as IEEE TDSC, ensure data integrity and confidentiality in the presence of semi-trusted cloud providers. The narrative concludes with a recent integration of Unsupervised Machine Learning, utilizing hierarchical clustering to provide interpretable anomaly detection in the face of increasingly automated network

threats.

Collectively, these works demonstrate a consistent commitment to bridging the gap between theoretical network modeling and practical system resilience. By synthesizing empirical measurement with cryptographic innovation, this thesis establishes a robust framework for managing performance and trust in the next generation of global digital infrastructures.

The original aspects of the thesis are discussed in Chapter 1. The (eighteen) papers, relevant to the development of the thesis, can be found attached in Appendix A.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

This chapter presents the overall motivation for works presented in this thesis, and summarizes the overall contribution.

1.1 Overview

This habilitation thesis presents a comprehensive body of research addressing the critical challenges of observability, modeling, and security in modern distributed networks. As digital infrastructures have evolved from localized cellular systems to planetary-scale cloud environments, the need for robust, evidence-based frameworks for performance and trust has become paramount. This work documents a 15-year scientific trajectory, moving from the Human Layer (Mobility) to the Infrastructure Layer (Observability) and culminating in the Trust Layer (Security and ML-driven Defense).

Pillar I: Spatio-Temporal Intelligence & Network Mobility

The first phase of this research focuses on the extraction of high-fidelity intelligence from massive, noisy datasets generated by cellular infrastructures.

Scientific Contribution: We pioneered the use of Gaussian Mixture Models (GMM) and probabilistic refinement to transform sparse Signaling/Call Data Records (CDRs) into precise spatio-temporal mobility maps [Auth2, Auth3, Auth4].

Impact: This work addressed the "Sparsity Paradox" in mobile networking—how to model human movement without high-frequency GPS tracking. Furthermore, it established methodologies for user-cell association and energy efficiency [Auth5], while pioneering graph-based obfuscation [Auth 6] to protect subscriber privacy at the topological level.

Pillar II: Planetary-Scale Cloud Observability & Performance

As services migrated to the cloud, the research focus expanded to the empirical behavior of globalized, "black box" distributed environments.

Scientific Contribution: I led the development of the CLAudit platform [Auth7], a first-of-a-kind system for auditing planetary-scale cloud latency.

Impact: This research provided a multidimensional evaluation [Auth8] of global Cloud Service Providers (CSPs), revealing how regional placement and internal datacenter topology directly influence the performance of Big Data frameworks such as Hadoop MapReduce [Auth10]. These findings replaced anecdotal evidence with empirical benchmarks [Auth9], enabling more dependable cloud-native deployments.

Pillar III: Dependable Security and ML-Driven Defense

The final pillar synthesizes theoretical insights with practical systems engineering to improve resilience in high-scale environments.

This work on protocol and data resilience bridges the protection of communication backbones—specifically SIP server resilience against DDoS and traffic anomalies [Auth11, Auth12, Auth13]—with protection of data content. Addressing the conflict between storage efficiency and privacy, we introduced enhanced thresholded deduplication schemes [Auth15, Auth16] that ensure data confidentiality in semi-trusted cloud environments.

Following-on the biological metaphors of DNA-inspired information concealing [Auth14], the research culminates in the development of unsupervised machine learning frameworks. By evolving from Hi-Clust [Auth17] to the HUMAN framework [Auth18], we provided a methodology for interpretable anomaly detection, utilizing hierarchical clustering to identify complex network threats, without the need for labeled training data.

1.2 Mobile Network and Cloud Computing Infrastructure: Scalable Distributed Systems

1.2.1 Definition

The objects investigated in this thesis are large-scale distributed infrastructures, such as cellular networks or cloud computing infrastructures, coupled with terminals, possibly mobile, and related applications.

This is an environment today of planetary scale, consisting of myriads of sub-domains and sub-systems; it is highly complex and constantly evolving. It is extremely difficult to obtain accurate and useful information about its status, and to correctly interpret such information. Furthermore, deploying novel services and optimizing the infrastructure requires accurate level of understanding of the processes happening within, to enable abstract modeling of the complex behavior of such system.

1.2.2 Methodology: measuring and modeling

Understanding processes within large infrastructures of the mobile cloud is a prerequisite to any performance or service optimization. Therefore, it is vital to be able to define and measure some key characteristics. We present three of such use cases in this thesis: tracking user-location within mobile networks, monitoring latency in remote clouds, and analyzing server traffic at a large communication server. The difficulty tackled rests in two aspects: (1) designing infrastructure that can capture data of interest at appropriate scale, without disruption to the running service, and (2) interpretation of such measurements to a useful end.

The first aspect is difficult, as pure logging of all vital events is not scalable and would harm performance, so specialized tools need to be built to facilitate measurements and allow aggregation or selection of representative data subset, good for interpretation. As shown in the study of the spatial extension of the reality mining dataset [Auth3] , very creative methods must be sometimes employed to complement incomplete measurements with missing attributes.

Perhaps an even more difficult challenge to address is interpretation of the measurements, which is discussed in the following part of the work. It is very hard to make correct conclusions about behavior from data characterized in many dimensions and influenced by many diverse factors. Deriving mathematical models is difficult and leads to highly complex structures. However, such models may prove useful in many ways, for example for analytically analyzing future resource planning or architecture optimizations. The challenge we address is to build models that prove useful for some further exploitation, such as prediction, resource optimization or troubleshooting. The case studies we present are motivated by potential for significant optimization of mobile networks, cloud services and communication services, should such data be gathered and utilized.

We present the cases of energy consumption optimization, and of troubleshooting of server communication or of a cloud-service behavior, based on detecting and analyzing anomalies.

1.2.3 Protecting: privacy and security

Collecting and studying data from networking and cloud-computing infrastructure presents another challenge - protecting this data from abuse, or simply just from revealing such information about users that is considered private, while maintaining useful properties of such data intact. The work on augmenting the reality mining dataset clearly shows that obfuscating some parts of the dataset is non-trivial and might in the future be broken by utilizing other related data sources. Various methods have been proposed to achieve the goal of sufficient privacy, but there is no one-size fits-all solution available. We present a few alternative solutions to this challenge, protecting server data, mobility data or data stored in the cloud, using various mathematical techniques: graph distortion, query field replacement, or an entirely novel cryptosystem which supports deduplication.

Another vexing cybersecurity problem is protecting the infrastructure on its own, as it otherwise is prone to attacks. Increasingly, attacks of the distributed denial-of-service type have been targeting important infrastructure. Innovative means of defending the perimeter of the infrastructure are necessary, such as the one present on defending a SIP server. It

might not even be that easy for a large infrastructure to discover that it is under attack! Analyzing monitored latency data in real time might be very helpful to discover and correctly identify such situation, as presented in [Auth8, Auth18].

1.3 Organization and scope of the habilitation thesis

The thesis is divided into three chapters, covering various aspects of large infrastructures: mobility tracking and modeling, performance auditing and privacy and security protection, as have been addressed by the author and co-authors. The contributions of the authoring teams are briefly presented, together with short summaries of related research work in each area. This thesis is meant to be a short integrating text, referring mainly to the appended papers.

1.4 Summary of contributions

The collective 18 publications represent a significant contribution to state-of-the-art in understanding scalable distributed infrastructures. By bridging the gap between theoretical modeling and empirical auditing, this research provided insights into building future networking and security solutions that are not only high-performing but also mathematically and autonomously secure.

Research results have been published at respected venues, including the IEEE INFOCOM conference as well as journals like IEEE Transactions on Dependable and Secure Computing and Computer Networks.

Furthermore, this work synthesizes a wide interdisciplinary breadth: a combination of Telecommunications (SIP/LTE), Cloud Computing and Bio-inspired methods for Data Protection and Machine Learning. It represents an evolution from measuring and modeling network signals to designing an intelligent, self-protecting distributed architecture.

Chapter 2

Mobility Tracking and Modeling

2.1 Introduction

Proliferation of mobile telephony at the end of the first decade of the 21st century transformed cellular networks into the world's largest sensors of human activity. However, the data generated by these networks—primarily Call Data Records (CDRs) and signaling logs—were never intended for high-fidelity positioning. They were spatially coarse (tower-level) and temporally sparse (event-driven). The challenge at the time was to move beyond descriptive statistics toward predictive, high-resolution models of mobility that could inform infrastructure planning while respecting the nascent requirements of data privacy.

This chapter documents a progression from early empirical data expansion to sophisticated probabilistic modeling. In [Auth3, Auth4], we addressed the initial limitations of available research datasets by providing spatial extensions to the Reality Mining corpus and improving user-cell association logic. The core contribution of this period is centered on the application of Gaussian Mixture Models (GMM) [Auth2], which allowed for the spatio-temporal refinement of sparse CDRs into continuous trajectories.

Recognizing that increased visibility creates privacy risks, the research then pivoted to protection. In [Auth1, Auth6], we provided a deep analysis of active tracking and proposed a novel graph-based obfuscation methodology. Unlike traditional k-anonymity, which often fails in high-dimensional mobility data, our approach secured the underlying network topology itself. Finally, [Auth5] connected these behavioral models back to the physical layer, demonstrating how understanding user flux can enable significant energy savings in cellular deployments without sacrificing Quality of Service (QoS).

2.2 Background and Related Work

Historically, mobility modeling has evolved around statistical methods and privacy-protective techniques.

Early work by Gonzalez et al. [1] (2008) and Song et al. [2] (2010) established the inherent predictability of human movement. However, these studies relied on "Reality Mining" datasets (like the one expanded in [Auth3]) which were often spatially coarse.

Research by Caceres et al. [3] (2007) attempted to use network signaling for traffic estimation, but lacked the probabilistic refinement needed to handle the "ping-pong" effect of cell-tower switching.

Current trends utilize Graph Neural Networks (GNNs) and Differential Privacy (DP) to synthesize trajectories. Fan et al. [4] (2021) introduced TrajGANs for masking, while Yang et al. [5] (2025) have moved toward modeling 6G ultra-dense networks.

Most modern AI-driven mobility models act as "black boxes" that ignore the physical constraints of cellular topology. Our work [Auth1, Auth2, Auth3] bridged this by using Gaussian Mixture Models to refine sparse signaling data, while [Auth6] provided a unique graph-based alternative to traditional k-anonymity [6].

2.3 Location tracking

The ability to track user's location is a pre-requisite to understanding patterns of his or her movement within the geographic space as well as within the network itself. There are a number of ways mobility or location tracking may be achieved, in this work we have focused on tracking from the perspective of the network operator. Such approach is difficult, as the client device may remain passive - not communicating - and in such case its exact location is opaque to the network. Therefore, some form of active probing may be necessary in order to trigger and share updates of passive device location within the network. The advantage of such approach is its non-disruptive nature and complete transparency to the user. In such situation, a clear option is to utilize tools and methods already available in the network. For this purpose, we have developed an active-probing tool, SS7Tracker, that allows to track user movement across a 2G mobile network without any cooperation necessary from the user. This tool is lightweight and presents no disruption or additional deployment cost to the actual network itself. It utilizes standard signaling messages as present in the network and operates within such constraints so as not to disrupt actual network operation.

In "Active tracking in mobile networks: An in-depth view" [Auth1] we have not only presented the architecture of the tracker (see Fig.2.1), but also analysis of the tracked data, its application, and recommended limits on the number of tracked users and tracking frequency.

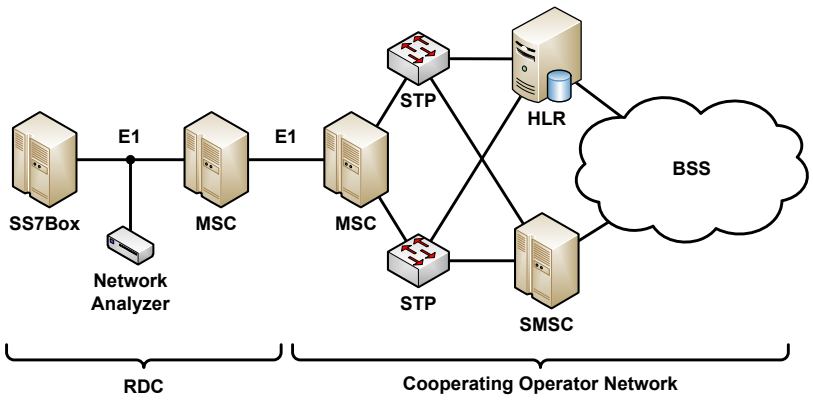

Figure 2.1: Interconnection of SS7Box into GSM

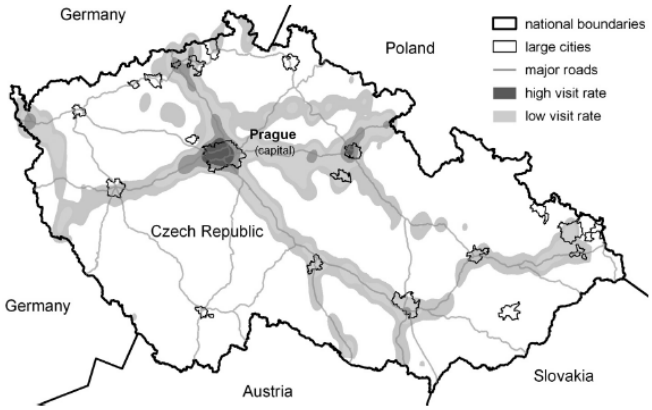

Figure 2.2: Network-based active tracking use example: places in the Czech Republic visited by a sample group of tourists—spatial interpolation of the numbers of tourists that visited each discrete cell.

This tool was implemented and deployed by the publishing team and subsequently used by a commercial partner for optimization of its roaming service in the Czech Republic.

Utilization results example, mapping of frequently visited places by tourists, is presented in Fig. 2.2.

2.4 Modeling and prediction

Mobility or location tracking tools produce wealth of mobility data. However, such passive, ex-post knowledge may not necessarily be always useful. For practical application, simulation or prediction, one might prefer to obtain an analytical model, which would allow to model mobility using a parametric mathematical model, or predict user location within the next time interval. In paper [Auth2], we derive such model, called Inter-Call Mobility Model (ICM), in the form of a Gaussian mixture.

The ICM model is derived as follows: Given the origin position $(x_A, y_A, t_A) = (0, 0, 0)$ and the destination position at $(x_B, y_B, t_B) = (1, 0, 1)$, the probability $\Phi(x, y, t)$ of finding a user at coordinates $(x, y) \in R^2$ at a time $t \in R, 0 \leq t \leq 1$, is defined as a Gaussian mixture

$$p(z = (x, y, t)^T; \theta) = \sum_{k=1}^{K} \pi_k \mathcal{N}(z; \mu_k, \Sigma_k) \tag{2.1}$$

of $K = 10$ components with parameters θ empirically derived from the data of Reality Mining Dataset 2.3. A graphical representation of such distribution can be observed in figure 2.4. A good fit between real data and the model is clearly visually observable, even if shown also statistically in the paper. Note the two extremes at the beginning and end of the call, showing users who have not yet commenced or already completed their movement across the network.

The Gaussian mixture model of probability distribution of user location across a mobile network over time is an extremely useful tool, allowing the use of analytical formulas when studying user mobility or utilizing user mobility for modeling other phenomena.

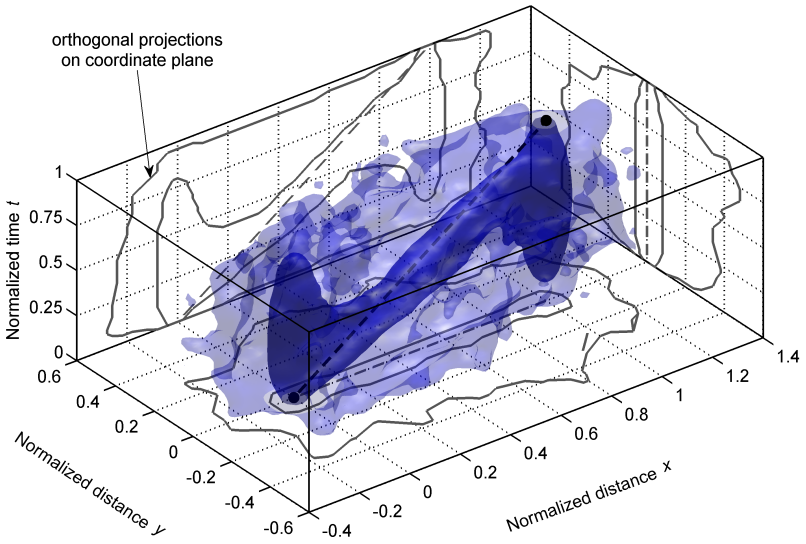

Figure 2.3: Kernel density estimation of the spatio-temporal probability distribution of users' position between calls (places A and B in normalized common reference frame), computed from the Reality Mining Dataset. Isosurfaces enclose 0.5- and 0.9-quantiles (dark and pale blue, respectively).

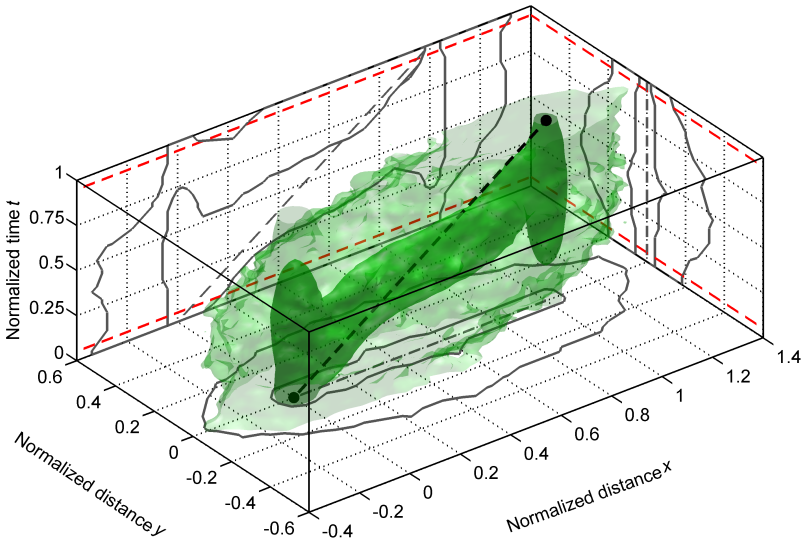

Figure 2.4: Inter-Call Mobility model. Isosurfaces enclose 0.5- and 0.9-quantiles (dark and pale green).

User location, even if within the network, may be considered private information, and may not thus be readily available to others. In reality, it is very hard to come by experimental data that would contain location traces, as network operators fear violating regulations and user privacy. Even in data collected as part of networking research, some attributes of the data are typically obfuscated, so as to prevent individual tracing or user localization. However, in "Spatial extension of the reality mining dataset" [Auth3] we have shown that sophisticated data mining techniques and additional knowledge of the organizing principles of mobile networks may help to subsequently determine the location information, even if originally obfuscated. Adding location information to the Reality Mining Dataset [7], which is widely used as reference, was highly appreciated by the research community and utilized in multiple subsequent works [8, 9]. At the same time, though, it has foreshadowed issues which are becoming much more relevant nowadays: with the growing availability of various datasets and the power of data mining tools, obfuscation tools that may be strong enough for a particular dataset on its own may not suffice in the case where cross-domain data mining is applied. The need for strong data obfuscation mechanisms is thus becoming ever more pressing. We present work on privacy protection in captured data in Section 2.6.

A question valuable to be answered for a mobile service or infrastructure is how to determine, with a reasonable success rate, the position, or location of association to network, of a user, within the next time interval. Such prediction is complicated, as large number of users remain idle or static, and thus even a primitive prediction may be relatively successful, but methods with greater accuracy must capture even user mobility patterns well. We present a technique based on past-data analysis using the Market-Basket Analysis method [Auth4], which allows to predict user location in the next interval with great accuracy. For the learning phase, we have used data provided by the measurements described in Section 2.3. We also demonstrate that even greater accuracy can be achieved when considering aggregates of users, not only individuals. Applications of the mobility modeling and prediction methods are plentiful. We present an example in the following Section 2.5.

2.5 Utilization: saving energy

In subsequent research work, we have focused on practical exploitation of mobility modeling and location tracking.

We present the example of studying how to exploit mobility knowledge for potential energy savings within a cellular network. Cellular networks are typically designed for peak performance and workload, thus wasting a lot of energy at off-peak times, where a large portion of their capacity remains idle. In Ënergy savings for cellular network with evaluation of impact on data traffic performance¨ [Auth5] we have presented a proposal how large energy savings could be achieved by powering off underutilized parts of the network at times of lower load. Users attached to the powered-off base stations seamlessly migrate to the next nearest ones (see Fig. 2.5 for illustration of the adjustments).

Owing to the realistic mobility data, we can demonstrate that such adjustments in the cellular network are feasible without hardly any performance decrease to the users, yet at substantial energy savings. We have shown that finding an optimal solution to the user-cell re-assignment is an NP-hard problem, and yet derived a heuristic optimization technique for

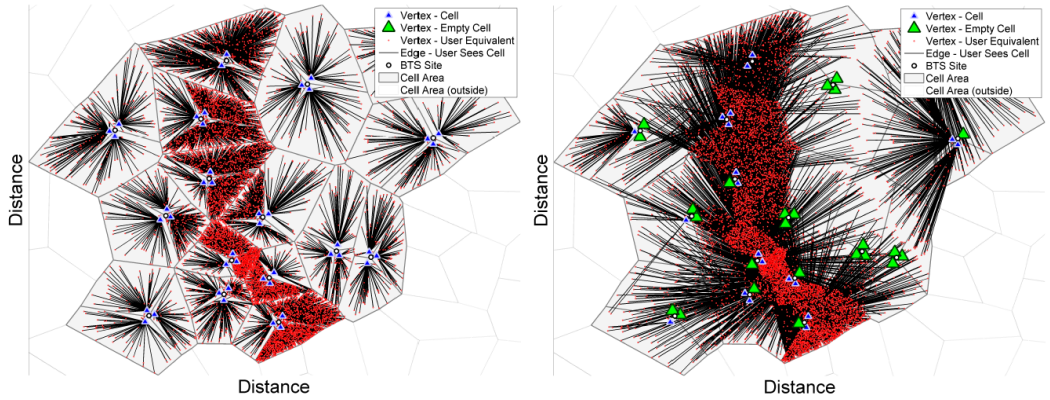

Figure 2.5: Visualization of energy-efficient optimization results. Cell radius R = 10 km, maximum number of active users allowed to associate to a single cell 240. Formerly under-utilized cells start to serve user equivalents from neighbouring cells. Left: Original user-cell association. Right: Optimized user-cell association

the selection of the network nodes targeted for powering down, and evaluate both voice and data communication performance impact on the users 2.6.

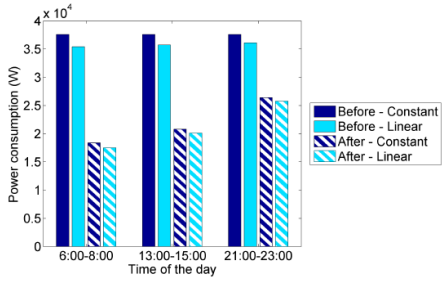

(a) Energy savings summary—the impact of user allocation optimization and cell switch-off on power consumption. On the x–axis is the time of day, for each time there are results for constant and linear energy profile, both for the situations before and after the user allocation optimization. On the y–axis is the total network power consumption in watts.

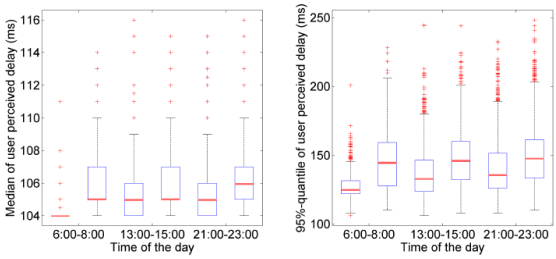

(b) User perceived delay — the impact of user allocation optimization. On the x–axis is the time of the day, for each time there are two boxplots, the one on the left for the situation before user allocation optimization and the one on the right for the situation after. On the y–axis is the delay in milliseconds. On a boxplot, the central thick line is the median, the box edges are the 25th and 75th percentiles, the whiskers mark data points not considered outliers, outliers plotted as dots. Left: Median delay of all requests of a user equivalent. Right: 95% quantile delay of all requests of a user equivalent.

Figure 2.6: Impact of energy-saving heuristics on energy consumption and user-perceived delay.

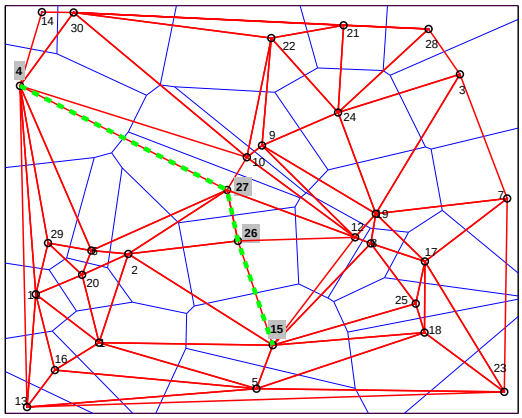

Figure 2.7: Cellular Network Topology Graph. Black circles represent Base Stations, blue lines define Voronoi cells, red lines stand for Delaunay triangulation, green dashed line shows a sample path S=4,27,26,15.

2.6 Obfuscating mobility patterns

Another type of data considered as private are data about user's movement within a cellular network. Since the sequence of visited (and associated) cells is recorded, it is feasible to track someones trajectory or daily habits in such network.

In paper "Mobility data anonymization by obfuscating the cellular network topology graph" [Auth6], we present a technique to transform the graph of the cellular network in such manner that it is still feasible to carry out aggregate studies of user mobility dynamics within the network, but individual trajectories are transformed in such manner that they are highly non-trivial to reconstruct.

The Cellular Network Topology Graph (CTG) 2.7 is transformed into an Obfuscated Cellular Network Topology Graph (OCTG) 2.8 while preserving some much desired properties such as node neighbourhood relationship and length and shape of users' trajectories are preserved, allowing to carry out longitudinal mobility studies without actually invading user privacy.

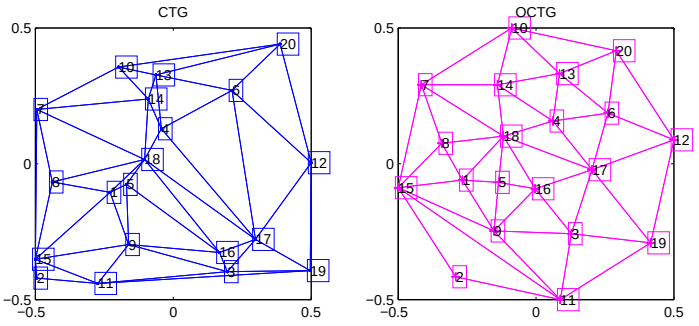

Figure 2.8: Cellular Network Topology Graph undergoing feature-preserving obfuscation.

Chapter 3

Planetary-Scale Cloud Observability

3.1 Introduction

As distributed computing shifted from local infrastructure to Public Cloud, an "observability gap" emerged. Cloud Service Providers (CSPs) operated as "black boxes," offering Service Level Agreements (SLAs) that were difficult for customers to independently verify. There was a critical lack of tools capable of performing longitudinal, multi-region auditing of the global Internet performance.

This chapter outlines development and deployment of the CLAudit platform [Auth7], our answer to the need for planetary-scale observability. CLAudit enabled us to move from anecdotal evidence to a multidimensional evaluation of cloud latency [Auth8]. We demonstrated that cloud performance is not merely a function of bandwidth but is deeply influenced by geographic placement and the provider's internal routing policies.

The research further investigated the intersection of network topology and Big Data. In [Auth10], we provided one of the first empirical evaluations of how datacenter network architectures (e.g., Fat-Tree vs. traditional hierarchies) directly impact the performance of Hadoop MapReduce clusters. By providing a framework for latency-based benchmarking [Auth9], this body of work empowered organizations to make evidence-based decisions during cloud migration, establishing a rigorous methodology for independent infrastructure auditing that remains a standard in the era of Edge and Multi-Cloud deployments.

To obtain reasonable understanding of processes and service-performance within such complex infrastructure as contemporary clouds, one must be able to measure various performance parameters and track the data and control flows within. The exact purpose may be varying, but in principle the goal is to determine normal behavior and separate it from anomalous behavior. At the same time, one might attempt to optimize the normal behavior, and troubleshoot the causes of the anomalous one. One way or another, large-scale studies are necessary, and adequate large-scale monitoring and appropriate trace analysis are research problems on their own. With this perspective in mind, we present several case studies of monitoring and analyzing performance of particular components of the networked cloud infrastructure - a large communication server in Section 4.3, a particular cloud service, and large-scale measurements of cloud infrastructure.

In Section 3.3 we first describe the various infrastructure components monitored, in Section 3.4 we analyze some of the anomalous phenomena uncovered and present methodology used for troubleshooting. Then, Section 3.5 presents a study how the described cloud infrastructure monitoring is employed to benchmark cloud operators. Finally, in Section 3.5 we present an study how the described cloud infrastructure monitoring is employed to benchmark cloud operators.

3.2 Background and Related Work

Cloud measurement has evolved from Availability Monitoring to Deep Introspection of fragmented infrastructures.

Classic Benchmarking: Armbrust et al. [10] (2010) identified performance unpredictability as a top obstacle for cloud adoption. Early tools like CloudCmp (Li et al. [11], 2010) attempted to compare providers but focused on static compute metrics.

When focussing on process performance, the logic of Hadoop MapReduce (Dean and Ghemawat [12], 2008) was well-understood, but the impact of the underlying physical network topology—theoretical work by Al-Fares et al. [13] (2008)—lacked empirical validation in public, multi-tenant cloud.

Simulations of datacenter application performance on different network topologies has been presented for example in [14]. This paper evaluates in detail throughput and delay in DCell, FatTree and *3-Tier hierarchical* topology (a tree topology consisting of 3 layers of switches and hosts at the leaves) given a typical traffic in a datacenter (not just MapReduce). The simulated traffic patterns were based on logs from several existing datacenters.

Performance of MapReduce specifically with respect to different topologies has been the focus of [15]. Topologies simulated in this paper were star (hosts are connected to a central "hub" node), double rack (essentially two star topologies with a connection between the central nodes), tree and DCell. The paper also presented a simulator called *MRPerf* and evaluated its performance during those simulations.

Research focused on optimizing effectivity of private clouds is also conducted at Department of Cybernetics of Czech Technical University. In the paper [16], the authors present a setup consisting of *Eucaliptus* private cloud system and queue engine *Cloud Gunther*. This setup integrates control of VM provisioning within the job scheduler, thus allowing for maximal utilization of the cloud's resources.

Current focus has shifted to observability based on eBPF (Extended Berkeley Packet Filter) [17] and Serverless (FaaS, Function-as-a-Service) [18] performance. Shahrad et al. [19] (2020) and Zhang et al. [20] (2024) have highlighted that "cold-start" latencies are the new primary bottlenecks in global infrastructures.

While modern observability focuses on internal microservice tracing, our CLAudit platform [Auth7] and multidimensional framework [Auth8] established the methodology for independent, external auditing of the "Black Box" cloud.

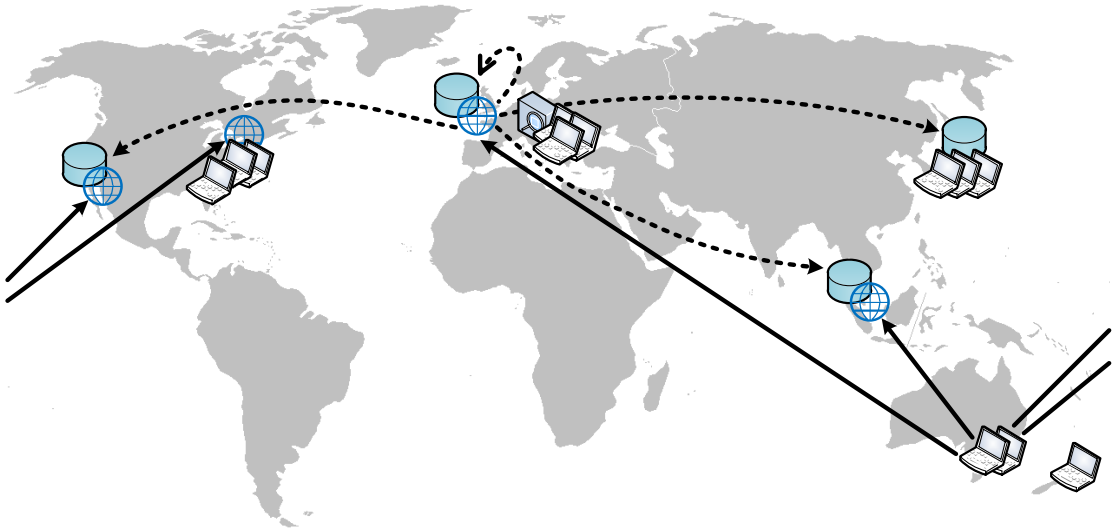

Figure 3.1: **CLAudit deployment.** Vantage Points (VPs) are represented by laptops, frontend servers by globe icons and backend servers by cylinder icons. Many of the monitor functions are executed on the central server, represented by lens icon. Continuous and dashed curves depict subset of VP–to–frontend and frontend–to–backend measurements, respectively.

3.3 Measurements: latency and traffic

For a comprehensive study on cloud-service performance, we have designed, implemented and operated an original global cloud-service evaluation platform, CLAudit [Auth7]. CLAudit focuses on continuous, periodic probing of specific datacenter services, measuring latency (the round-trip time) of the relevant task 3.1. The measurements are executed at various layers of the TCP/IP stack 3.2, to multiple datacenters across the globe and from multiple probing locations, also distributed globally.

Different services are being measured, a simple web response or a more elaborate database query, requiring a further redirect to a different datacenter. This variability of measurements

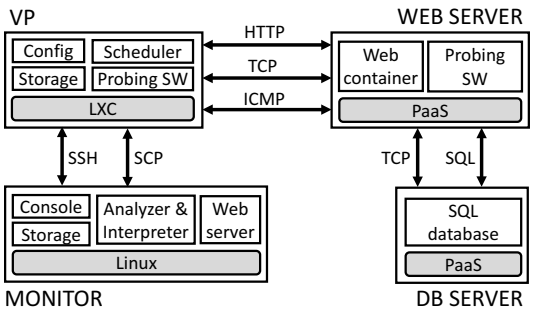

Figure 3.2: **Component view of CLAudit.** Four major components along with execution environments that host functions and modules for inter–component protocol interactions.

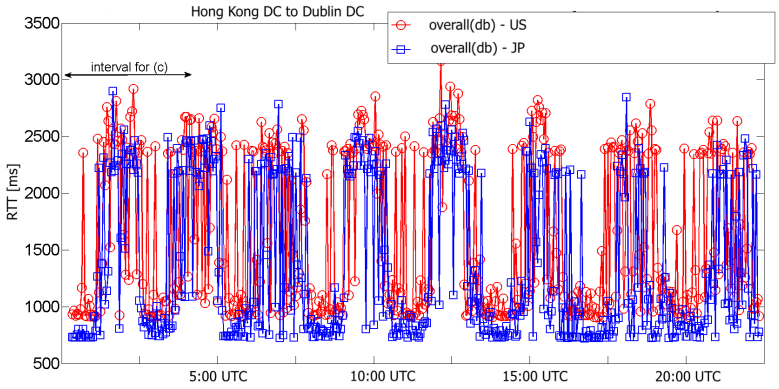

(a) Periodic latency (US and Japan VP retrieving dynamic webpage from Hong Kong web server.)

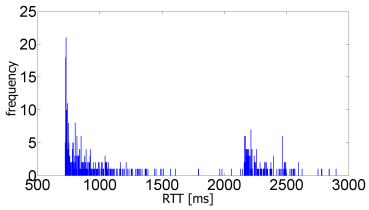

(b) RTT histogram (JP)

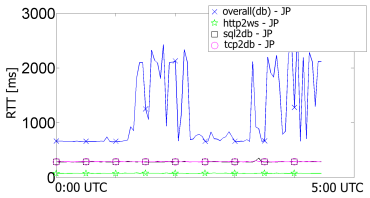

(c) Protocol breakdown

Figure 3.3: **Periodic latency anomaly of multitier geo–distributed application**. Alternating low stable and high volatile periods of user–perceived latency of multitier web application.

composes a multi-dimensional view of a cloud service offered globally by one operator. This can be employed for real-time decisions as well as longer-term evaluations of quality of a cloud service. CLAudit was in continuous operating mode for a couple of years, gathering measurements from the prominent cloud operator infrastructure, Microsoft Azure and Amazon AWS, publishing its measurements and sharing those with the research community.

3.4 Interpretation: troubleshooting

Wealth of data was amassed by multidimensional cloud latency monitoring using the CLAudit [Auth7] platform. This vast trace was thoroughly analyzed in a follow-up journal publication "Multidimensional Cloud Latency Monitoring and Evaluation" [Auth8], which also focused on comparison between two distinct years of measurements: 2013 and 2016. It contains some innovative methodology, such as comparison to speed-of-light minima, and introducing the metric of latency per kilometer. It has uncovered some very interesting phenomena, as summarized in the paper. One observation is that cloud services are quite unstable. Latency of various cloud services exhibits a multi-modal distribution 3.3, which, for latency-sensitive applications, could result in sub-standard performance. Furthermore, network paths across the globe are not equally reliable, and their performance fluctuates, thus for each location, at different times, a different datacenter may be optimal. A clear

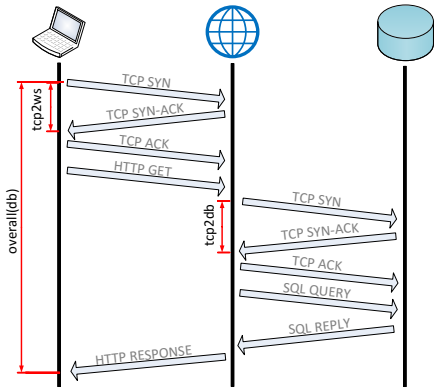

Figure 3.4: Claudit dynamic probe into remote sql database.

improvement in quality and reliability of the service between 2013 and 2016 is also visible.

This follow-up work [Auth8] has also developed an original methodology how to find the root cause of anomalies found in the multidimensional latency traces. As there are inherent dependencies among the layers of the TCP/IP stack (i.e. if the network layer is not functioning properly, neither can the application layer) and measurements are made from multiple locations to multiple locations, aggregate anomalies at particular location can help localize a problem, whether it occurs at the probe location, in the network, or within the datacenter (see probe example in 3.4. This innovative methodology is summarized in a combination of two decision trees, an Interpretation Tree 3.5 and Impact Tree, which help determine the root cause for a particular anomaly. A summary of applying this methodology to the entire trace is then provided, again demonstrating improvements between 2013 and 2016.

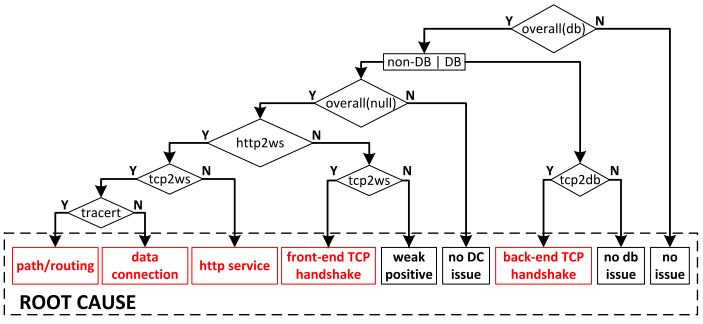

Figure 3.5: Claudit event interpretation tree for anomaly detection.

3.5 Benchmarking

The world of cloud computing today is not standardized (in comparison to e.g. the Internet itself), but rather highly diverse and fragmented. Multiple Cloud Operators have built competing global infrastructures and platforms, providing analogical, yet not interoperable, services. A cloud service user is thus forced to either choose one operator only, or implement the missing interoperability part by themselves. Selecting a Cloud Operator, either long-term, or at run-time, may be a tricky task, as not much comparison data is available.

In paper "Latency–based Benchmarking of Cloud Service Providers" [Auth9], we have presented an innovative methodology how to compare multiple cloud operators based on measurements. First, a comparable set of datacenter locations, cloud services and probing locations is established. Upon measurements taken, preprocessing, normalization, and selection of particular metrics of choice, together with their weights (i.e. setting their importance) is carried out. As a result, a single multi-dimensional vector in the unit space is obtained, reflecting the distance from optimal performance (point 0). The size of the vector thus determines the actual aggregate performance of a particular cloud operator. In the study, we have compared two major cloud operators, Microsoft Azure and Amazon AWS. The identity of each remained anonymous though - the interest of the article is in the comparison methodology. Nevertheless, in the sample evaluation presented in paper [Auth9], there is a clear observable performance difference among the two.

3.6 MapReduce performance evaluation

Paper [Auth10] focuses on the effect of network topology on performance of different workload types of the MapReduce distributed computing model. To study this problem, several simulation scenarios have been proposed. These simulations cover cases of different MapReduce job dimensions (numbers of Mappers and Reducers), intermediary data distributions, and network topologies. The performance of MapReduce in individual scenarios is evaluated based on job finish times.

Performance of MapReduce comprises of several factors - among others task scheduling, skew mitigation, struggler detection, topology optimization.

In each scenario, a single MapReduce job is launched in the simulated cluster. Nodes in the cluster are configured with only one task slot and the scheduler only assigns one task to each machine (as opposed to one machine running several tasks sequentially). These constraints were introduce in order to increase performance of the simulator, but more importantly to reduce effect of other aspects of MapReduce performance than the impact of network topology. With that in mind, the simulated jobs were configured with relatively small amount of computational workload compared to the configured volume of data.

The simulator chosen to run the simulations is the extension of CloudSim [21] called CloudSimEx MapReduce [22] chosen for its flexibility, level of detail, as well as ease of implementation and integration of new modules.

A network topology simulation module, called *CloudDynTop*, has also been implemented as an extension of CloudSimEx MapReduce.

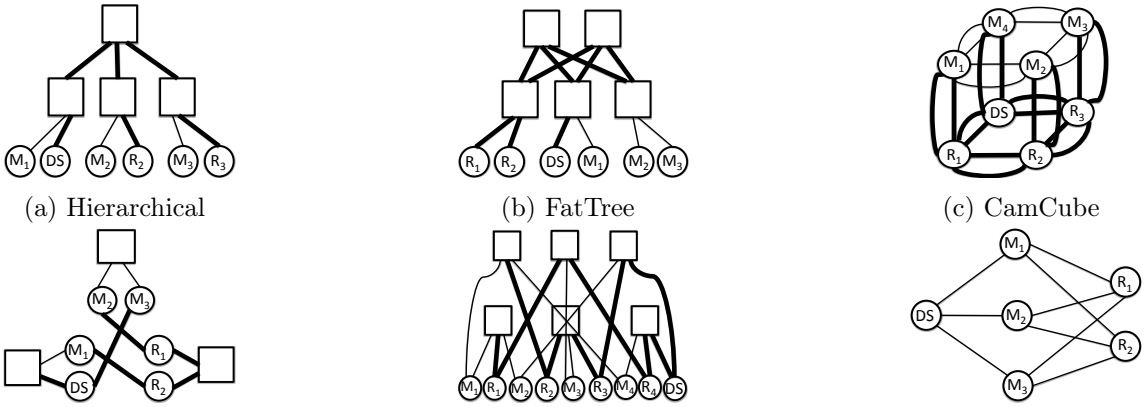

(a) Hierarchical (b) FatTree (c) CamCube

(d) DCell (e) BCube (f) MapReduce

Figure 3.6: *Illustration of topologies considered in the simulations.* Circles denote hosts, squares denote switches and lines denote links of two capacities (distinguished by line weight).

We thoroughly evaluate the influence which the popular Data Center network topologies have on Hadoop MapReduce performance. Considered are two standard, three experimental and a single hypothetical topology (illustrated in Fig. 3.6):

- *Hierarchical*
- *FatTree*
- *DCell*
- *BCube*
- *CamCube*
- *MapReduce*

Hierarchical and FatTree topologies are variations on a standard n-tier Tree. The *FatTree* topology has been described in paper [23]. It is a tree topology, where processing nodes are located at the leaves and internal nodes of the tree are switches. Another Tree-like topology is the recursive BCube [24], defined by parameters N (number of hosts in the level-0 BCube) and K (index of the highest recursion level). $N = 3, K = 1$ BCube is depicted in Fig. 3.6. The CamCube [25] is based on 3D Torus topology, where servers are directly interconnected. DCell [26] topology combines switches and server routing. The hypothetically optimal, "MapReduce", topology is modeled after the MapReduce data flow, providing a dedicated connection between Data Source and Mappers as well as between Mappers and Reducers such that no link is traversed by more than a single flow.

The main output of this study are finish times of the simulated workloads on individual topologies. These results are presented in Figure 3.7

For all distributions the best performance was achieved by the MapReduce topology, where each flow has its own dedicated link. This topology represents optimal performance given the limited capacity of links in the network.

The CamCube topology performed only slightly worse than the MapReduce topology, especially in case of uniformly distributed intermediary data. The BCube topology has performed very well, although not as well as CamCube. An interesting result is that increasing throughput of link connecting the Reducer nodes has had a great impact on the MapReduce job's performance.

Performance on hierarchical topology has been the worst for all the simulated workloads, because of the bottlenecks between edge and aggregation and aggregation and core layer. Because of these bottlenecks the hierarchical topology has been clearly outperformed.

In summary, it has been convincingly demonstrated that datacenter topology has strong impact on performance of critical processes such as Hadoop MapReduce and that 3D toroidal or recursive topologies may perform close to optimum.

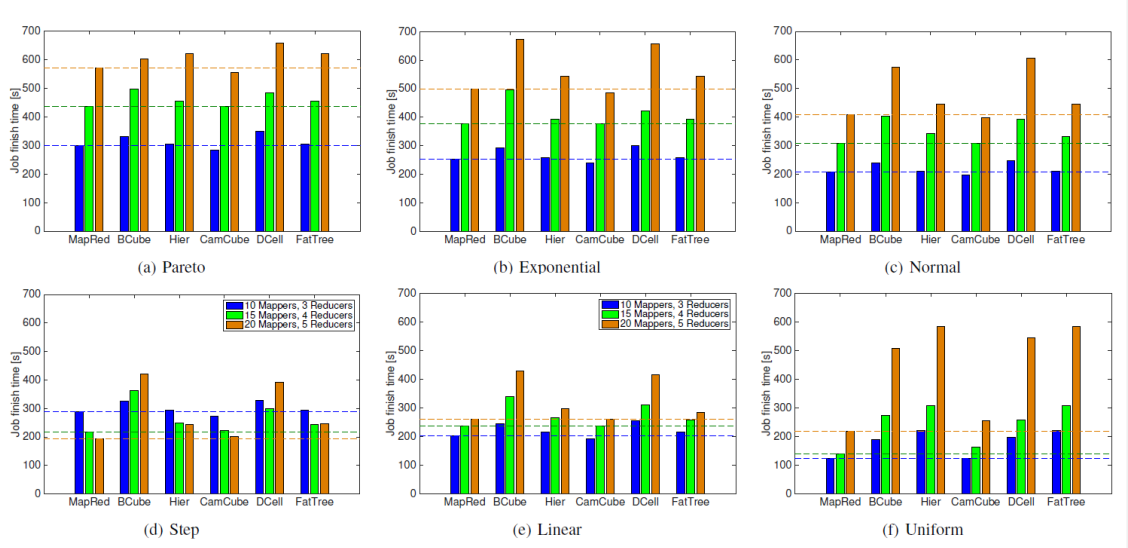

Figure 3.7: Job finish times. Results of simulation runs of MapReduce job, varying three job dimensions (14, 20, 26 hosts), six Reducer–workload divisions (Pareto, Exponential, Normal, Step, Linear, Uniform) and six datacenter network topologies (MapReduce, BCube, Hierarchical, CamCube, DCell, FatTree).

Chapter 4

Security in Scalable Distributed Systems

4.1 Introduction

This Chapter discusses results in the areas of security and privacy in relation to Cloud infrastructure and data. At the age of moving large amounts of services, often critical or highly personal, to distant infrastructure networked over the Internet, questions about robustness and cybersecurity of such infrastructure and of protection of confidential data stored within, arise. This area has attracted a lot of research interest. First, attackers may exploit various vulnerabilities in communication protocol design, implementation and deployment, to execute various types of denial-of-service attacks, which may disrupt infrastructure operation and impact its availability and robustness. Thus, various enhanced manners of protocol deployment and infrastructure protection are being sought, learning from practical, real-world scenarios. Second, another objective of the attacker may be abuse of the data stored within or informing about these infrastructures. Various data-encryption or data-transformation methods may thus be used, to protect such information. However, such protection may thus be in contrast to the actual use of the data. This often leads to interesting mathematical problems, as computing efficiency is often in contradiction to the level of protection of infrastructure or data, and a compromise needs to be sought. This means that solutions arise which for example allow to execute some partial operations, but prevent others, while preserving some level of confidentiality.

This research thus addresses a fundamental requirement of distributed systems: Dependability. This theme evolved through three stages: securing communication protocols (SIP), securing stored data (Deduplication), and finally, securing the network through intelligent, autonomous defense (Hierarchical Clustering).

In this Chapter we present several innovative methods of this type. Firstly, in Section 4.3 we describe a method to protect a large and widely used Internet infrastructure server against distributed denial-of-service attacks (DDoS). Further on, in Section 4.4, we present several methods of data transformation in the interest of its protection, while preserving some of its usability properties, in the interest of anomaly detection, mobility analysis or de-duplication.

The research in this chapter begins with the resilience of the Session Initiation Protocol (SIP) [27, 28], the backbone of modern VoIP. In [Auth11, Auth12, Auth13], we analyzed real-world SIP traffic anomalies and proposed "SIP Protector," a defense architecture designed to mitigate DDoS flood attacks. This work bridged the gap between traditional protocol engineering and behavioral traffic analysis.

Simultaneously, we addressed the "Deduplication-Encryption Paradox." In [Auth15, Auth16], we developed an enhanced secure thresholded deduplication scheme. This solved a major roadblock in cloud storage by allowing for efficient data reduction while maintaining high cryptographic privacy for sensitive information.

The logical climax of this thesis is the transition toward Intelligent Defense. Drawing from the theoretical foundations of DNA-inspired information concealing [Auth 16], we moved away from signature-based security toward unsupervised machine learning. The evolution from Hi-Clust [Auth 17] to the HUMAN framework [Auth 18] represents our current state-of-the-art: a system capable of interpretable anomaly detection. By using hierarchical clustering, we provide network administrators not just with an "alert," but with a structural visualization of the threat, bridging the gap between automated AI and human-led network forensics.

4.2 Background and Related Work

This pillar reflects the most significant evolution: from protecting protocols and data to engineering intelligent defense.

Protocol Security (SIP): Early VoIP security research by Salsano et al. [29] (2002) and Geneiatakis et al. [30] (2006) identified the vulnerability of SIP to INVITE floods and DDoS attacks. Traditional defenses relied on static rate-limiting or signature-based detection.

Storage and Deduplication: The "Deduplication Paradox" was defined by Douceur et al. [31] (2002). Later, Bellare et al. [32] (2013) formalized Message-Locked Encryption (MLE), but it remained vulnerable to brute-force attacks on low-entropy data. Current research (Li et al. [33], 2022) utilizes Confidential Computing (Intel SGX) to perform secure deduplication.

Evolution of AI Defense (Clustering): Unsupervised learning for anomaly detection was pioneered by Laskov et al. [34]. However, traditional clustering methods often suffered from high dimensionality. The current state of the art (Liu et al. [35]) emphasizes Self-Supervised Learning (SSL), yet these models often lack the "interpretability" required for forensic action.

The Gap: This research fills three critical voids. First, it addresses SIP resilience [Auth11, Auth12, Auth13] through behavioral traffic analysis. Second, it solves the deduplication conflict via thresholded security [Auth15, Auth16]. Finally, it moves from basic clustering (Hi-Clust [Auth17]) to interpretable unsupervised defense (HUMAN [Auth18]), bridging the gap between high-accuracy detection and human-readable network forensics.

4.3 Infrastructure Protection

As already indicated in Chapter 3, the Session Initiation Protocol (SIP) [27, 28] belongs among the key communication protocols of the Internet today. Therefore, its control servers,

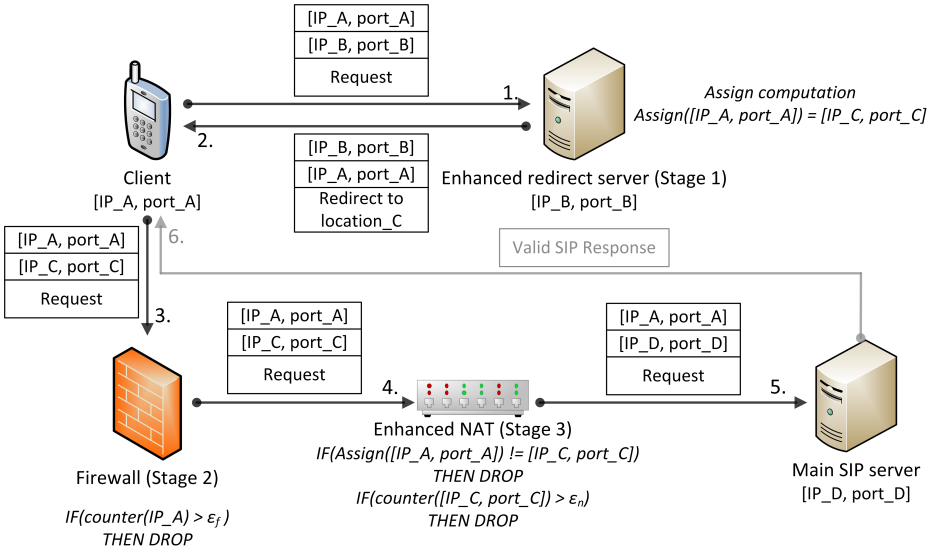

Figure 4.1: SIP Protector architecture overview.

typically deployed in a remote cloud, are an attractive target for denial-of-service attacks. In "SIP Protector: Defense architecture mitigating DDoS flood attacks against SIP servers" [Auth11] we have presented a simple but effective mechanism how to prevent major types of DDoS attacks, by inserting a redirect server in front of the actual SIP server. An opaque assignment mapping is computed for each requester, and sent back to each. Only subsequent requests conforming to the mapping are then allowed to pass through the Network Address Translation (NAT) element deployed in front of the actual SIP server (see architecture in Fig. 4.1).

In a practical experiment using real SIP traffic we then demonstrate the effectiveness of mitigating such attack (see Fig. 4.2.

The study of the SIP server trace [Auth12] has uncovered some very interesting general phenomena, such as that due to its specific properties, the SIP workload does not follow the Power-Law probability distribution, otherwise so prominently present in the networking world.

The first measurements work we present is a study of several days of operation of a large communications server [Auth12]. It is a server operated by an external commercial entity,

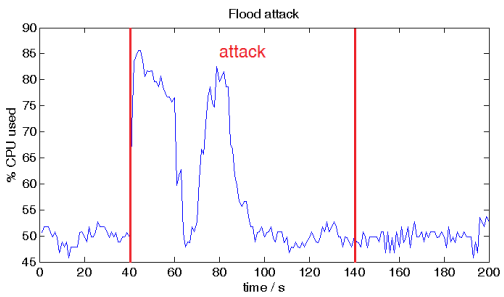

Figure 4.2: SIP Protector mitigating DoS attack on SIP server.

one of the technical leaders in this domain. The server acts a central point for the Session Initiation Protocol (SIP) [27, 28], one of the key components of contemporary cloud services. It is a lightweight and highly intuitive signaling protocol, enabling to set up real-time Voice-over-IP (VoIP) and video interaction. We capture and analyze a full trace of three days of operation of the said SIP server. The main technical challenge here is efficient way of working with the large dataset, for which we design and employ parallelism to be able to utilize multiple machines.

Another interesting phenomenon uncovered is the so-called *registration storm*, which occurs after a short outage. Due to the protocol design, a massive workload spike appears in such case, threatening yet another outage. Thorough analysis of the registration storms, and possible remedy thereof, was subject of our follow-up work [Auth13] .

4.4 Data Protection

Various types of data are stored or logged within contemporary networked systems, with many of such data being sensitive from one perspective or another. Such data typically needs to be protected by encryption, which though often prevents any meaningful operation over the data.

We present three examples of innovative techniques that address this area as a compromise - executing partial obfuscation or encryption to enable some particular type of data processing:

4.4.1 Obfuscating identity in SIP Logs

In the paper "Analyzing anomalies in anonymized SIP traffic" [Auth13] we presented a method to obfuscate private information in SIP server logs (see Fig. 4.3) in such manner that prevents identifying the originator yet enables to reconstruct complex protocol behavior. We demonstrate the method usability by successfully detecting anomalous behavior even over such obfuscated logs.

4.4.2 Information Concealing techniques in DNA

The paper explores the fascinating intersection of biological data storage and digital security. DNA is presented not just as a biological blueprint, but as a hyper-dense, highly redundant, and naturally secure medium for information. The survey posits that the mechanisms evolved by nature to protect genetic information—such as redundancy, error-correction, and non-coding regions ("Junk DNA")—offer a perfect template for digital steganography and cryptography.

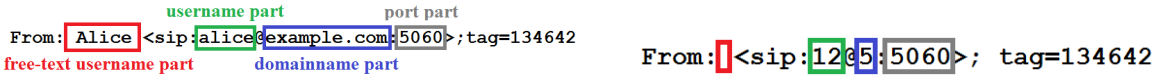

(a) SIP trace data prior to anonymization.

(b) SIP trace data after anonymization.

Figure 4.3: Sensitive data obfuscated in SIP trace.

The survey categorizes the field into three distinct domains:

DNA Steganography: Hiding a secret message within a DNA sequence. The paper reviews methods like the "Genomic Steganography" approach, where digital data is mapped to the four-letter alphabet of nitrogenous bases (A,C,G,T).

DNA Watermarking: Techniques for marking synthetic DNA sequences or digital data to ensure provenance and intellectual property protection, drawing parallels to how biological signatures persist across generations.

DNA Cryptography: Using the massive parallelism of DNA computing to solve NP-hard problems or using DNA hybridization as a physical "one-time pad."

While much of the early literature focused on physical DNA strands in a lab, Kencl and Loebl specifically highlighted the digital application of these metaphors. The paper identifies how biological principles can be abstracted into algorithms to protect digital assets against eavesdropping and tampering.

A central concept discussed in the survey is the mapping of digital bits to nucleotide bases. A common representation is: 00→A,01→C,10→G,11→T . This simple mapping allows a digital binary stream to be represented as a "synthetic" DNA strand, which can then

Example 1: Procedure $S(\omega, o, lb, ub)$
Input $\omega =$ **theEaimEofEthisEpaperEisEtoEpresentEanEinformationEconcealingEalgorithm**, parameters $o = 3, lb = 4, ub = 6$.

1. First, the input sequence is partitioned randomly into blocks of length 4, 5 or 6. The blocks are divided by '+' below:
   theEai+mEofEt+hisE+pape+rEisE+toEpr+esent+EanEi+nfor+matio+nEco+ncea+lingE+algor+ithm+
2. Next we add overlap (of length $o = k - 1 = 3$) in front of each block:
   thmtheEai+EaimEofEt+fEthisE+isEpape+aperEisE+isEtoEpr+Epresent+entEanEi+nEinfor+formatio+tionEco+Econcea+cealingE+ngEalgor+gorithm+
3. Next we add the dust behind each block (of length approximately 2), and we get the cards:
   thmtheEaip+EaimEofEtim+fEthisEcon+isEpapeEin+aperEisEa+isEtoEproEp+Epresentese+entEanEilgo+nEinforiE+formatiofo+tionEcoE+EconceaEci+cealingEpa+ngEalgorEp+gorithmap+
4. Finally the output is given by arranging the cards in a random order (here we use the order 14, 9, 10, 13, 5, 3, 12, 1, 6, 4, 7, 11, 8, 15, 2):
   ngEalgorEpnEinforiEformatiofocealingEpaaperEisEafE thisEconEconceaEcithmtheEaipisEtoEproEpisEpapeEinE presentesetionEcoEentEanEilgogorithmapEaimEofEtim

Figure 4.4: Example of inserting repeats into text string by the concealing procedure.

be hidden within a larger genomic sequence using steganographic algorithms. We conjecture that repeats in DNA may form such data protection, define concealing procedures of introducing repeats (see example in Fig. 4.4) and prove that information sequence concealed in this manner is mathematically difficult to reconstruct.

This technique allows to build mechanisms for preserving some sequence characteristics (e.g. all short sequences) while having difficulty reconstructing the entire sequence. Digital applications are numerous, such as virus detection or voice recognition in otherwise uninterpretable file.

4.4.3 Secure deduplication

Data deduplication is a well known technique to save space needed for storage of large filesets by storing copies of identical files only once. This is particularly relevant in the cloud scenario, where the operator might save huge resources by deduplicating across large filesets. At the same time though, it is desirable that data stored in the cloud are encrypted, at least when at rest, so as to prevent abuse of the data. Deuplication is unfortunately unusable when dealing with encrypted files or data, as any initial file similarity must be hidden and encryption must possess proper entropy, rendering file redundancy extremely unlikely. The method of convergent encryption (encryption which uses as seed the actual file to be encrypted) has been designed to solve this problem. It works with respect to enabling file deduplication, but unfortunately is prone to some basic guessing attacks.

In works [Auth15] and [Auth16], we present a method that builds on convergent encryption, but for some files, ads an extra, non-convergent, encryption layer. This extra layer applies to less popular files, i.e. files that only appear within the fileset a few times, up to a given threshold (see Figure 4.5).

The rationale being that once a file is shared among many cloud customers, its confidentiality may be less stringent, and convergent encryption would thus be sufficient to apply. This compromise also makes sense from the deduplication point of view, as it is the popular files (i.e. those of which many instances appear) who will bring the greatest savings when deduplication applies.

For the above scheme, we build an elaborate system description and define a threshold-based multi-party encryption scheme, where each participant holds to a particular decryption share. This further encryption layer is removed once the number of instances of a particular

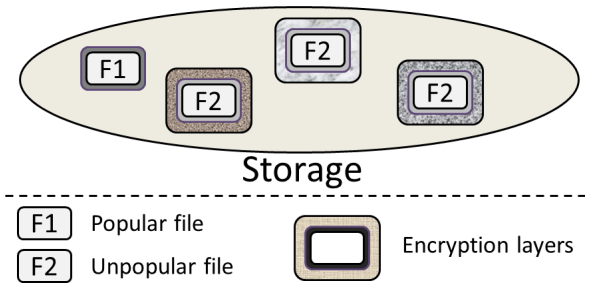

Figure 4.5: Secure deduplication: popular and unpopular files.

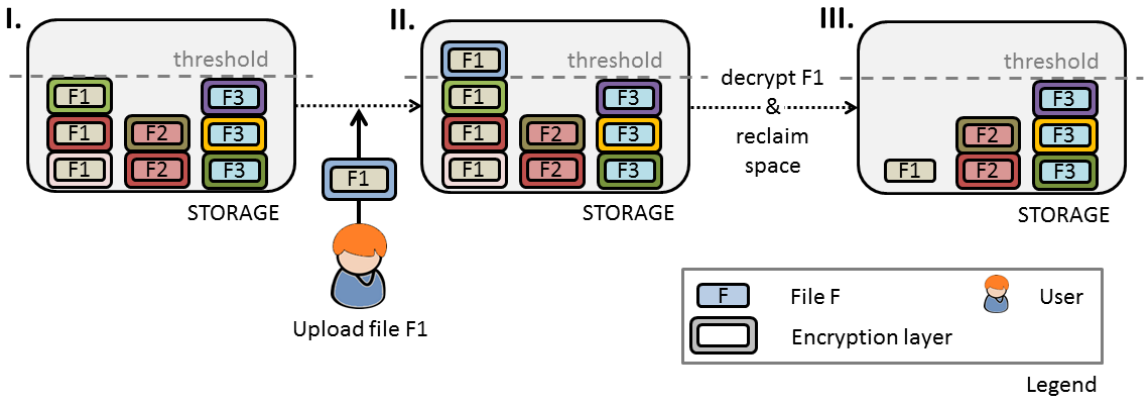

Figure 4.6: File encryption evolution over popularity threshold.

file reaches a certain threshold, which means sufficient number of decryption shares exists (see Figure 4.6).

In subsequent journal publication [Auth16], we enhance the overall security of the scheme by moving the overall handling of decryption shares outside of the cloud storage 4.7 and proving theoretical validity of the cryptosystem setup. Furthermore, we provide a thorough performance analysis of the scheme and comparison to other state-of-the-art proposals.

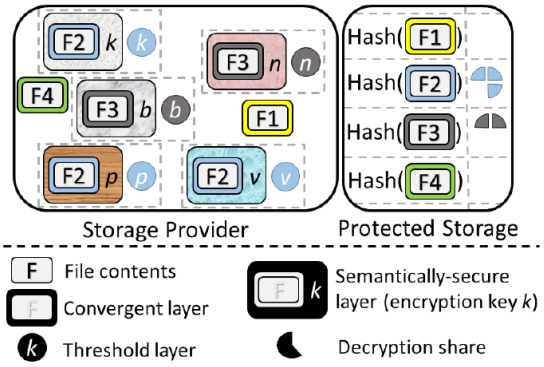

(a) The multi-layered cryptosystem for deduplication. Unpopular files (F2 and F3) are protected using two layers, semantically-secure and convergent, whereas popular files (F1 and F4) have the upper layer already removed and are deduplicated.

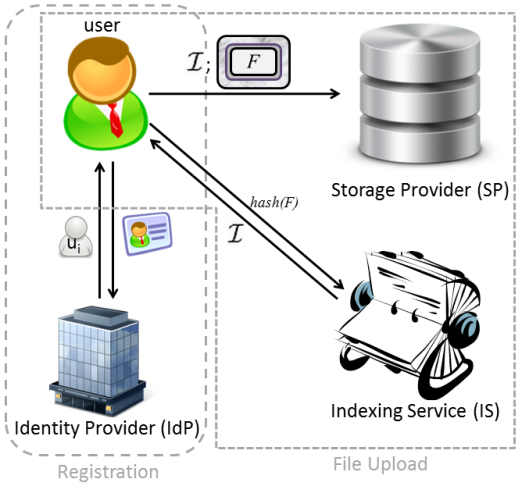

(b) System model for secure de-duplication. The schematic shows the main four entities and their interaction for registration and file upload process

Figure 4.7: Enhanced secure de-duplication.

4.5 Machine Learning for Intelligent Defense

The final phase of this research track addresses the increasing complexity of distributed data by moving away from hard-coded thresholds toward Autonomous, Interpretable AI. This trajectory is defined by two major works that utilize Hierarchical Agglomerative Clustering (HAC) to solve the "Black Box" problem in network management.

4.5.1 Hi-Clust: Unsupervised Latency Profiling [Auth17]

Published at IEEE CloudNet (2018), Hi-Clust was developed to handle the large, high-dimensional datasets produced by the CLAudit platform 3.3. Traditional monitoring often relies on the mean (μ) or standard deviation (σ) of latency, which fails to capture the "multi-modal" nature of cloud performance (e.g., where a single region exhibits three distinct latency behaviors simultaneously due to varying internal paths).

Hi-Clust applies a bottom-up hierarchical clustering approach. It does not require a predefined number of clusters (k), which is essential in cloud environments where the number of performance profiles is unknown.

The framework proved that hierarchical structures could effectively isolate "noisy neighbor" effects and network jitter from baseline performance. By utilizing the Cophenetic Correlation Coefficient to validate cluster consistency, Hi-Clust provided the first unsupervised "fingerprinting" of global Cloud Service Providers.

4.5.2 HUMAN: Interpretable Anomaly Detection [Auth18]

The HUMAN framework (Hierarchical clustering for Unsupervised iNterpretation of Anomalies) [Auth18], published at IEEE NoF (2020), represents the security-centric evolution of Hi-Clust. It addresses the "Alert Fatigue" in modern Security Operations Centers (SOCs).

While Deep Learning models can detect anomalies with high precision, they rarely explain why a flow was flagged. HUMAN solves this by leveraging the hierarchical tree. HUMAN processes raw network features—packet sizes, inter-arrival times, and protocol flags—and groups them into clusters. Anomalies are identified as "outlier clusters" or small, isolated branches. The key innovation is that the framework allows an administrator to "drill down" into an anomaly cluster to see the common features that define it (e.g., "All these flows use port 5060 and have a 0-byte payload"). This provides immediate forensic context.

Table 4.1: Comparative Analysis of the Hi-Clust and HUMAN Frameworks

Feature	**Hi-Clust [Auth 17]**	**HUMAN [Auth 18]**
Research Goal	Infrastructure Profiling	Threat Detection & Forensics
Primary Metric	Network Latency (RTT)	Flow-based Traffic Features
Clustering Logic	Distance-based (Euclidean)	Multi-feature Similarity
Learning Type	Unsupervised (Clustering)	Unsupervised (Interpretive AI)
Output	Performance Profiles	Interpretable Anomaly Clusters
Operational Value	SLA Verification	Zero-day Attack Mitigation

Chapter 5

Conclusion and Future Outlook

5.1 Synthesis of Research Contributions

This Thesis has documented a 15-year scientific journey dedicated to observability, performance and trust in distributed infrastructures. It contains selected original results and innovative works focusing on cellular networks, cloud computing and data protection. Approximately, it can be partitioned into sections on analyzing and modeling network mobility, on measuring and modeling cloud performance, and on protecting the cloud infrastructure and user data stored therein. By bridging the gaps between telecommunications, cloud computing, and machine learning, the body of work presented (papers [Auth1]– [Auth18]) offers a unified framework for managing the complexities of modern digital ecosystems.

Pillar I – Mobility: We explored the deterministic, idealized mobility models and then evolved toward a probabilistic reality. By utilizing Gaussian Mixture Models [Auth2], we proved that sparse carrier data could be refined into high-fidelity intelligence. At the same time, data abuse may be prevented by underlying network topology obfuscation process [Auth6].

Pillar II – Observability: We addressed the Cloud transparency gap by architecting the CLAudit platform [Auth7]. This research transitioned cloud auditing from provider-centric metrics to a global, longitudinal perspective, establishing that planetary-scale performance is fundamentally dictated by the interplay of physical topology and distributed framework logic [Auth10].

Pillar III – Security and Machine Learning: First, we established a paradigm of Bio-inspired Resilience [Auth14]. Then, beginning with protocol-level defense for SIP [Auth13] and moving to secure data reduction [Auth16], we culminated in the development of the HUMAN framework [Auth18]. This work showed that unsupervised hierarchical clustering may provide the missing link between high-accuracy anomaly detection and the interpretability required for human-led forensics.

In all the areas above, we have presented novel solutions on how to collect, interpret and utilize data for increased understanding of the observed phenomena. This data-driven methodology has allowed us to build innovative solutions on top of the data, be it to conserve energy in cellular networks, to compare and interpret cloud networking and computing performance, to protect SIP traffic or data stored in the cloud.

The emerging *central thesis* of this work then is that *Performance* and *Security* are *not competing interests*, but *two sides of the same, dependability, coin*. The mobility patterns derived in Pillar I inform the cloud placement strategies of Pillar II, which in turn define the traffic profiles analyzed by the ML defenses in Pillar III. This holistic view - treating the user, the network, and the data as a single, interdependent system - is the primary contribution of this Habilitation thesis.

5.2 Future Outlook

As we look toward 2030, the scientific challenges identified in this work will only intensify. The transition to 6G ultra-dense networks and the ubiquity of AI will require a new generation of observability and cyber-resilience tooling.

Building on the foundations of [Auth6] and [Auth16], my future research will investigate confidential computing for verifiable federated learning, ensuring that AI models can learn from sensitive data without truly seeing it.

Evolution starting with anomaly detection, as in HUMAN [Auth18], will lead towards autonomous response and self-healing networks. Adaptive methods were studied by the author already in works predating this thesis [36, 37] and thus we intend to continue in this direction in the new context. By integrating Large Language Models (LLMs) with hierarchical clustering, we can move toward a system that not only interprets an attack but autonomously generates and deploys the necessary mitigation scripts in real-time. The advent of autonomous AI Agents shall trigger needs for research in the domains of cyber resilience and trust mechanisms, which we intend to explore.

Another vital direction is sustainability. Following the energy-efficiency goals established in [Auth5], future performance studies must treat "Carbon-per-Packet" as a metric as critical as Latency or Throughput.

5.3 Final Statement

The 18 publications included in this thesis represent a sustained commitment to rigorous, independent research. By providing the tools to measure, model and secure complex scalable distributed systems, this work contributes to the foundation of a more resilient, transparent, and trustworthy internet.

Bibliography

[1] Marta C Gonzalez, Cesar A Hidalgo, and Albert-Laszlo Barabasi. Understanding individual human mobility patterns. *Nature*, 453(7196):779–782, 2008.

[2] Chaoming Song, Zehui Qu, Nicholas Blumm, and Albert-László Barabási. Limits of predictability in human mobility. *Science*, 327(5968):1018–1021, 2010.

[3] N Caceres, JP Wideberg, and FG Benitez. Deriving origin–destination data from a mobile phone network. *IET Intelligent Transport Systems*, 1(1):15–26, 2007.

[4] Y. Fan et al. TrajGANs: Using generative adversarial networks for geo-textual trajectory masking. *IEEE Transactions on Intelligent Transportation Systems*, 2021.

[5] L. Yang et al. GNN-based mobility forecasting in 6G ultra-dense networks. *IEEE Communications Magazine*, 2025.

[6] L. Sweeney. k-anonymity: a model for protecting privacy. *International Journal on Uncertainty, Fuzziness and Knowledge-based Systems*, 10(5):557–570, 2002.

[7] A Pentland, N Eagle, and D Lazer. Inferring social network structure using mobile phone data. *Proceedings of the National Academy of Sciences (PNAS)*, 106(36):15274–15278, 2009.

[8] Suman Nath. Ace: exploiting correlation for energy-efficient and continuous context sensing. In *Proceedings of the 10th international conference on Mobile systems, applications, and services*, pages 29–42, 2012.

[9] Longxiang Gao, Ming Li, Alessio Bonti, Wanlei Zhou, and Shui Yu. Multidimensional routing protocol in human-associated delay-tolerant networks. *IEEE Transactions on Mobile Computing*, 12(11):2132–2144, 2013.

[10] Michael Armbrust, Ion Stoica, Matei Zaharia, et al. A view of cloud computing. *Communications of the ACM*, 53(4):50–58, 2010.

[11] Ang Li, Xiaowei Yang, Srikanth Kandula, and Ming Zhang. CloudCmp: comparing public cloud providers. In *Proceedings of the 10th ACM SIGCOMM Conference on Internet Measurement*, pages 1–14, 2010.

[12] Jeffrey Dean and Sanjay Ghemawat. MapReduce: simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008.

[13] Mohammad Al-Fares, Alexander Loukissas, and Amin Vahdat. A scalable, commodity data center network architecture. In *ACM SIGCOMM Computer Communication Review*, volume 38, pages 63–74, 2008.

[14] Kashif Bilal, Samee U Khan, Limin Zhang, Hongxiang Li, Khizar Hayat, Sajjad A Madani, Nasro Min-Allah, Lizhe Wang, Dan Chen, Majid Iqbal, et al. Quantitative comparisons of the state-of-the-art data center architectures. *Concurrency and Computation: Practice and Experience*, 25(12):1771–1783, 2013.

[15] Guanying Wang, Ali Raza Butt, Prashant Pandey, and Karan Gupta. A simulation approach to evaluating design decisions in mapreduce setups. In *Modeling, Analysis & Simulation of Computer and Telecommunication Systems, 2009. MASCOTS'09. IEEE International Symposium on*, pages 1–11. IEEE, 2009.

[16] Tomas Vondra and Jan Sedivy. Maximizing utilization in private IaaS clouds with heterogenous load. In *CLOUD COMPUTING 2012, The Third International Conference on Cloud Computing, GRIDs, and Virtualization*, pages 169–173, 2012.

[17] Sebastiano Miano, Matteo Bertrone, Fulvio Risso, Massimo Tumolo, and Mauricio Vásquez Bernal. Creating complex network services with ebpf: Experience and lessons learned. In *2018 IEEE 19th International Conference on High Performance Switching and Routing (HPSR)*, pages 1–8, 2018.

[18] Theo Lynn, Pierangelo Rosati, Arnaud Lejeune, and Vincent Emeakaroha. A preliminary review of enterprise serverless cloud computing (function-as-a-service) platforms. In *2017 IEEE International Conference on Cloud Computing Technology and Science (CloudCom)*, pages 162–169, 2017.

[19] Mohammad Shahrad et al. Serverless in the wild: Characterizing and optimizing the Azure functions production workload. In *USENIX Annual Technical Conference (ATC)*, pages 205–218, 2020.

[20] X. Zhang et al. Mitigating cold starts in FaaS: A survey of 2024 techniques. *ACM Computing Surveys*, 2024.

[21] Rodrigo N Calheiros, Rajiv Ranjan, Anton Beloglazov, César AF De Rose, and Rajkumar Buyya. Cloudsim: a toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms. *Software: Practice and Experience*, 41(1):23–50, 2011.

[22] Aa Bb. Cloudsimex home page. https://github.com/Cloudslab/CloudSimEx/, cited on May 5th, 2015.

[23] Charles E Leiserson. Fat-trees: universal networks for hardware-efficient supercomputing. *Computers, IEEE Transactions on*, 100(10):892–901, 1985.

[24] Chuanxiong Guo, Guohan Lu, Dan Li, Haitao Wu, Xuan Zhang, Yunfeng Shi, Chen Tian, Yongguang Zhang, and Songwu Lu. Bcube: a high performance, server-centric network architecture for modular data centers. *ACM SIGCOMM Computer Communication Review*, 39(4):63–74, 2009.

[25] Hussam Abu-Libdeh, Paolo Costa, Antony Rowstron, Greg O'Shea, and Austin Donnelly. Symbiotic routing in future data centers. *ACM SIGCOMM Computer Communication Review*, 41(4):51–62, 2011.

[26] Chuanxiong Guo, Haitao Wu, Kun Tan, Lei Shi, Yongguang Zhang, and Songwu Lu. Dcell: a scalable and fault-tolerant network structure for data centers. *ACM SIGCOMM Computer Communication Review*, 38(4):75–86, 2008.

[27] Handley et. al. SIP: Session initiation protocol (RFC 2543). http://www.ietf.org/rfc/rfc2543.txt.

[28] Rosenberg et. al. SIP: Session initiation protocol (RFC 3261). http://www.ietf.org/rfc/rfc3261.txt.

[29] Stefano Salsano, Luca Veltri, and David Papalilo. SIP security issues: The state of the art. *IEEE Communications Magazine*, 40(4):48–55, 2002.

[30] Dimitris Geneiatakis et al. A survey of SIP security. *IEEE Communications Surveys & Tutorials*, 8(2):68–81, 2006.

[31] John R Douceur et al. Reclaiming space from duplicate files in a serverless distributed file system. In *ICDCS*, pages 617–624, 2002.

[32] Mihir Bellare, Sriram Keelveedhi, and Thomas Ristenpart. Message-locked encryption and secure deduplication. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 296–312. Springer, 2013.

[33] J. Li et al. Secure and efficient deduplication with confidential computing. *IEEE Transactions on Parallel and Distributed Systems*, 33(9):2132–2145, 2022.

[34] Pavel Laskov, Patrick Düssel, Christin Schäfer, and Konrad Rieck. Learning intrusion detection: Supervised or unsupervised? In Fabio Roli and Sergio Vitulano, editors, *Image Analysis and Processing – ICIAP 2005*, pages 50–57, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.

[35] R. Liu et al. Self-supervised learning for network anomaly detection: A survey. *IEEE Communications Surveys & Tutorials*, 2023.

[36] Lukas Kencl and Jean-Yves Le Boudec. Adaptive load sharing for network processors. *IEEE/ACM Transactions on Networking*, 16(2):293–306, 2008.

[37] L. Kencl and C. Schwarzer. Traffic-adaptive packet filtering of denial of service attacks. In *2006 International Symposium on a World of Wireless, Mobile and Multimedia Networks(WoWMoM'06)*, pages 5 pp.–489, 2006.

Appendix A

Selected Author's Publications

[Auth1] Michal Ficek, Tomáš Pop, and Lukáš Kencl. Active tracking in mobile networks: An in-depth view. *Computer Networks*, 57(9):1936 – 1954, 2013.

[Auth2] M. Ficek and L. Kencl. Inter-call mobility model: A spatio-temporal refinement of call data records using a Gaussian mixture model. In *2012 Proceedings IEEE INFOCOM*, pages 469–477, 2012.

[Auth3] M. Ficek and L. Kencl. Spatial extension of the Reality Mining Dataset. In *The 7th IEEE International Conference on Mobile Ad-hoc and Sensor Systems (IEEE MASS 2010)*, pages 666–673, 2010.

[Auth4] Kateřina Dufková, Jean-Yves Le Boudec, Lukáš Kencl, and Milan Bjelica. Predicting user-cell association in cellular networks from tracked data. In *Mobile Entity Localization and Tracking in GPS-less Environnments*, pages 19–33. Springer, 2009.

[Auth5] K. Dufková, M. Bjelica, B. Moon, L. Kencl, and J. Le Boudec. Energy savings for cellular network with evaluation of impact on data traffic performance. In *2010 European Wireless Conference (EW)*, pages 916–923, 2010.

[Auth6] E. B. Martinez, M. Ficek, and L. Kencl. Mobility data anonymization by obfuscating the cellular network topology graph. In *2013 IEEE International Conference on Communications (ICC)*, pages 2032–2036, 2013.

[Auth7] Ondrej Tomanek and Lukas Kencl. Claudit: Planetary-Scale Cloud Latency Auditing Platform. In *Cloud Networking (CloudNet)*. IEEE, 2013.

[Auth8] Ondrej Tomanek, Pavol Mulinka, and Lukas Kencl. Multidimensional Cloud Latency Monitoring and Evaluation. *Computer Networks*, 107:104–120, 2016.

[Auth9] Vojtech Uhlir, Ondrej Tomanek, and Lukas Kencl. Latency–based Benchmarking of Cloud Service Providers. In *Proceedings of the 9th International Conference on Utility and Cloud Computing*, pages 263–268. ACM, 2016.

[Auth10] Zdenek Kouba, Ondrej Tomanek, and Lukas Kencl. Evaluation of datacenter network topology influence on Hadoop MapReduce performance. In *2016 5th IEEE International Conference on Cloud Networking (Cloudnet)*, pages 95–100, 2016.

[Auth11] Jan Stanek and Lukas Kencl. SIP protector: Defense architecture mitigating DDoS flood attacks against SIP servers. In *2012 IEEE International Conference on Communications (ICC)*, pages 6733–6738. IEEE, 2012.

[Auth12] Jan Stanek, Lukas Kencl, and Jiri Kuthan. Characteristics of real open SIP-server traffic. In *International Conference on Passive and Active Network Measurement*, pages 187–197. Springer, 2013.

[Auth13] Jan Stanek, Lukas Kencl, and Jiri Kuthan. Analyzing anomalies in anonymized SIP traffic. In *Networking Conference, 2014 IFIP*, pages 1–9. IEEE, 2014.

[Auth14] Lukas Kencl and Martin Loebl. DNA-inspired information concealing: A survey. *Computer Science Review*, 4(4):251–262, 2010.

[Auth15] Jan Stanek, Alessandro Sorniotti, Elli Androulaki, and Lukas Kencl. A secure data deduplication scheme for cloud storage. In Nicolas Christin and Reihaneh Safavi-Naini, editors, *Financial Cryptography and Data Security*, pages 99–118, Berlin, Heidelberg, 2014. Springer Berlin Heidelberg.

[Auth16] J. Stanek and L. Kencl. Enhanced securere thresholded data deduplication scheme for cloud storage. *IEEE Transactions on Dependable and Secure Computing*, PP(99):1–1, 2016.

[Auth17] Pavol Mulinka, Pedro Casas, and Lukas Kencl. Hi-Clust: Unsupervised analysis of cloud latency measurements through hierarchical clustering. In *2018 IEEE 7th International Conference on Cloud Networking (CloudNet)*, pages 1–7, 2018.

[Auth18] Pavol Mulinka, Pedro Casas, Kensuke Fukuda, and Lukas Kencl. HUMAN - hierarchical clustering for unsupervised anomaly detection interpretation. In *2020 11th International Conference on Network of the Future (NoF)*, pages 132–140, 2020.