

# Idilio Drago – Curriculum Vitæ

---

<b>Address</b>	Piazza F. Carrara 5, 10132 - Torino - IT	<b>Phone</b>	+39 345 447 6229
<b>Date of Birth</b>	30 <sup>th</sup> March 1980	<b>Site</b>	<a href="https://sites.google.com/site/idiliod/">https://sites.google.com/site/idiliod/</a>
<b>Nationality</b>	Italian / Brazilian	<b>Email</b>	<a href="mailto:idilio.drago@polito.it">idilio.drago@polito.it</a>

## Background

- 2009 – 2013**      **Ph.D. in Computer Science (obtained on 13/12/2013)**  
**University of Twente, the Netherlands**  
*Understanding and Monitoring Cloud Services*  
Advisor: Prof.dr.ir. Aiko Pras and Prof.dr.ir. Boudewijn Haverkort  
Topics: Cloud storage, Internet measurements
- My thesis focused on cloud storage and network measurements. The contributions include methods to monitor the performance of cloud systems as well as practical analyses of cloud storage protocols. I studied popular storage services as a case study, showing how they impact the network and their performance bottlenecks.
- 2005 – 2007**      **M.Sc. in Computer Science**  
**UFES – Federal University of Espírito Santo, Brazil**  
*An Experimental Evaluation of Classification Algorithms for Time Series Databases*  
Advisor: Prof. D.Sc. Flávio Miguel Varejão  
Topics: Pattern recognition, time series classification
- I worked on data mining for fraud detection, aiming at deriving a methodology to assist on the selection of customers of the power delivery system for inspection. The system was based on classical classification algorithms, such as the nearest neighbor classifier and decision trees, and on the histories of electricity consumption.
- 1998 – 2004**      **B.Sc. in Computer Engineering**  
**UFES – Federal University of Espírito Santo, Brazil**  
*Feature Selection for Fraud Identification in the Electric Power Delivery System*  
Advisor: Prof. Ph.D. Thomas Walter Rauber  
Topics: Data mining, feature extraction and selection
- I obtained a Computer Engineering degree from UFES. My final paper was on “Feature Selection for Fraud Identification in the Electric Power Delivery System”.

## Prizes

- **IRTF Applied Networking Research Prize 2013**  
For the paper: “Inside Dropbox: Understanding Personal Cloud Storage Services”.  
Published at the 12th ACM SIGCOMM Internet Measurement Conference (IMC 2012).
- **Best Paper Award of PAM 2013**  
For the paper: “Measurement Artifacts in NetFlow Data”.

## Languages

Portuguese (native), English, Italian.

## Employment History – Academia

- Oct 2016 –** **Politecnico di Torino, Italy**  
Assistant Professor with a temporary contract (RTDa)
- Apr 2015 – Oct 2016** **Politecnico di Torino, Italy**  
Postdoctoral Researcher (Assegnista di Ricerca)
- I have focused on two main research lines, in the context of both EU-Projects and partnerships with private companies (e.g., Cisco Inc.):
- (i) Building a big data system to archive and analyze network measurements. This task has included the participation in the definition and engineering of the recently created BigData@Polito laboratory;
  - (ii) Employing big data techniques to analyze large collections of measurements and to evaluate the workings of Internet services – in particular, those hosted in cloud systems. My research focused on understanding how such services operate; modeling their workload; and uncovering how users actually perceive their performance.
- Jan 2014 – Apr 2015** **Politecnico di Torino, Italy**  
Postdoctoral Researcher (Borsa per Attività di Ricerca)
- I participated in research and development projects, in collaboration with Narus Inc., aiming at (i) building efficient systems for monitoring high-speed Internet links, leveraging methods typically employed in cloud storage, such as data de-duplication, for optimizing resources usage; (ii) evaluating the resulting traces to uncover performance impairments in networked applications.
- Mar 2009 – Mar 2013** **University of Twente, the Netherlands**  
Ph.D. Candidate
- During this time I developed my PhD thesis on cloud storage and internet measurements. I was enrolled as a PhD candidate in the “Design and Analysis of Communication Systems” group, which is part of both the Computer Science and the Electrical Engineering departments of the UT.
- Jan – Dec 2008** **UFES / FCAA – Federal University of Espírito Santo, Brazil**  
Research Assistant
- I worked as an assistant (masters’ level) in a research project in collaboration with Petrobras, the largest oil extraction company in Brazil. The project aimed at creating an intelligent system to anticipate failures in the machinery used to extract oil from the deep sea. My task was to develop an initial prototype for the system, based on machine learning and clustering algorithms.
- Feb – Jul 2006** **UFES – Federal University of Espírito Santo, Brazil**  
Teaching Assistant
- This assistantship was a requirement to obtain my M.Sc. Degree in Computer Science, and included lecturing for 6-month a course on data structures to first-year undergraduate students.

## Employment History – Industry

Apr – Nov 2004

**Xerox, Software Development Center, Brazil**  
Software Developer

Software Developer for *Xerox Inc.*

Dec 1998 – Mar 2004

**ESCELSA – Energy Company of Espírito Santo, Brazil**  
Web Developer and System Administrator

I worked as a System Administrator in what was the first Internet Service Provider (ISP) in my state in Brazil. Thus, I have also a large practical experience on telecommunications from my time in the industry.

## Publications

Profile on Google Scholar (556 citations):

<https://scholar.google.it/citations?user=0K72LOkAAAAJ>

Profile on Scopus (170 citations):

<https://www.scopus.com/authid/detail.uri?authorId=25926966100>

- Gonçalves, G., Drago, I., Vieira, A.B., Silva, A.P.C., and Almeida, J.M. (2016) *The Impact of Content Sharing on Cloud Storage Bandwidth Consumption*. IEEE Internet Computing Magazine. 20 (4). pp 26-35. DOI 10.1109/MIC.2016.89.
- Gonçalves, G., Drago, I., Vieira, A.B., Silva, A.P.C., Almeida, J.M., and Mellia, M. (2016) *Workload Models and Performance Evaluation of Cloud Storage Services*. Computer Networks. 109 (2). pp 183-199. DOI 10.1016/j.comnet.2016.03.024.
- Trevisan, M., Drago, I., Mellia, M. (2016) *Impact of Access Line Capacity on Adaptive Video Streaming Quality – A Passive Perspective*. In: Proceedings of the Workshop on QoE-based Analysis and Management of Data Communication Networks (Internet-QoE). Florianopolis, Brazil. pp. 7-12. DOI 10.1145/2940136.2940139.
- Trevisan, M., Drago, I., Mellia, M., H.H. Song, Baldi, M. (2016) *WHAT: A Big Data Approach for Accounting of Modern Web Services*. In: IEEE Workshop on Big Data and Machine Learning in Telecom (BMLIT). Washington, DC, USA.
- Manzoor, J., Drago, I., and Sadre, R. (2016) *The Curious Case of Parallel Connections in HTTP/2*. To Appear: Proceedings of the International Conference on Network and Service Management (CNSM). Montreal, Canada.
- Gonçalves, G., Drago, I., Vieira, A.B., Silva, A.P.C., and Almeida, J.M. (2016) *Analysing Costs and Benefits of Content Sharing in Cloud Storage*. In: Proceedings of the Workshop on Fostering Latin-American Research in Data Communication Networks (LANCOMM). Florianopolis, Brazil. pp. 43-45. DOI 10.1145/2940116.2940128.
- Trevisan, M., Drago, I., Mellia, M., and Munafò, M.M. (2016) *Towards Web Service Classification using Addresses and DNS*. In: Proceedings of the 7th International Workshop on Traffic Analysis and Characterization (TRAC). Paphos, Cyprus. pp. 38-43. DOI 10.1109/IWCMC.2016.7577030.
- Vassio, L., Drago, I., and Mellia, M. (2016) *Detecting User Actions from HTTP Traces: Toward an Automatic Approach*. In: Proceedings of the 7th International Workshop on Traffic Analysis and Characterization (TRAC). Paphos, Cyprus. pp. 50-55. DOI 10.1109/IWCMC.2016.7577032.
- Bocchi, E., Drago, I., and Mellia, M. (2015) *Personal Cloud Storage Benchmarks and Comparison*. IEEE Transactions on Cloud Computing. 99 (PP). pp 1–14. DOI 10.1109/TCC.2015.2427191.
- Bocchi, E., Drago, I., and Mellia, M. (2015) *Personal Cloud Storage: Usage, Performance and Impact of Terminals*. In: Proceedings of the 4th IEEE International Conference on Cloud Networking (CLOUDNET). Ontario, Canada. pp. 106–111. DOI 10.1109/CloudNet.2015.7335291.

- Gonçalves, G., Drago, I., Vieira, A.B., Silva, A.P.C., and Almeida, J.M. (2015) *Analyzing the Impact of Dropbox Content Sharing on an Academic Network* (in Portuguese). In: Proceedings of the XXXIII Brazilian Symposium on Computer Networks and Distributed Systems (SBRC). Vitoria-ES, Brazil. pp. 100–109. DOI 10.1109/SBRC.2015.22.
- Hofstede, R.J., Celeda, P., Trammell, B., Drago, I., Sadre, R., Sperotto, A., and Pras, A. (2014) *Flow Monitoring Explained: From Packet Capture to Data Analysis with NetFlow and IPFIX*. IEEE Communications Surveys & Tutorials. 16 (4). pp. 2037–2064. DOI 10.1109/COMST.2014.2321898.
- Gonçalves, G., Drago, I., Silva, A.P.C., Vieira, A.B., and Almeida, J.M. (2014) *Modeling the Dropbox Client Behavior*. In: Proceedings of the IEEE International Conference on Communications (ICC). Sydney, Australia. pp. 1332–1337. DOI 10.1109/ICC.2014.6883506.
- Gonçalves, G., Drago, I., Silva, A.P.C., Vieira, A.B., and Almeida, J.M. (2014) *Characterizing and Modeling the Dropbox Workload* (in Portuguese). In: Proceedings of the XXXII Brazilian Symposium on Computer Networks and Distributed Systems (SBRC). Florianópolis-SC, Brazil. pp. 266–274. DOI 10.1109/SBRC.2014.32.
- Drago, I. (2013) *Understanding and Monitoring Cloud Services*. PhD Thesis, University of Twente. CTIT Ph.D. thesis Series No. 13–279. ISBN 978-90-365-3577-9. DOI 10.3990/1.9789036535779.
- Drago, I., Bocchi, E., Mellia, M., Slatman, H., and Pras, A. (2013) *Benchmarking Personal Cloud Storage*. In: Proceedings of the 13th ACM SIGCOMM Internet Measurement Conference (IMC). Barcelona, Spain. pp. 205–212. DOI 10.1145/2504730.2504762.
- Hofstede, R.J., Drago, I., Sperotto, A., Sadre, R., and Pras, A. (2013) *Measurement Artifacts in NetFlow Data*. In: Proceedings of the 14th International Conference on Passive and Active Measurement (PAM). Hong Kong, China. pp. 1–10. DOI 10.1007/978-3-642-36516-4\_1.
- Drago, I., Hofstede, R.J., Sadre, R., Sperotto, A., and Pras, A. (2015) *Measuring Cloud Service Health Using NetFlow/IPFIX: the WikiLeaks Case*. Journal of Network and Systems Management. 23 (1). pp. 58–88. DOI 10.1007/s10922-013-9278-0.
- Drago, I., Schmidt, R.O., Hofstede, R.J., Sperotto, A., Karimzadeh, M., Haverkort, B.R., and Pras, A. (2013) *Networking for the Cloud: Challenges and Trends*. PIK – Praxis der Informationsverarbeitung und Kommunikation. 36(4). pp. 207–214. DOI 10.1515/pik-2013-0035.
- Drago, I., Mellia, M., Munafò, M.M., Sperotto, A., Sadre, R., and Pras, A. (2012) *Inside Dropbox: Understanding Personal Cloud Storage Services*. In: Proceedings of the 12th ACM SIGCOMM Internet Measurement Conference (IMC). Boston, USA. pp. 481–494. DOI 10.1145/2398776.2398827.
- Hofstede, R.J., Drago, I., Sperotto, A., and Pras, A. (2011) *Flow Monitoring Experiences at the Ethernet-Layer*. In: Proceedings of the 17th Workshop on Energy-Aware Communications (EUNICE). Dresden, Germany. pp. 134–145. DOI 10.1007/978-3-642-23541-2\_15.
- Hofstede, R.J., Drago, I., Moreira Moura, G.C., and Pras, A. (2011) *Carrier Ethernet OAM: An Overview and Comparison to IP OAM*. In: Proceedings of the 5th International Conference on Autonomous Infrastructure, Management and Security (AIMS). Nancy, France. pp. 112–123. DOI 10.1007/978-3-642-21484-4\_14.
- Drago, I., and Pras, A. (2010) *Scalable Service Performance Monitoring*. In: Proceedings of the 4th International Conference on Autonomous Infrastructure, Management and Security (AIMS). Zurich, Switzerland. pp. 175–178. DOI 10.1007/978-3-642-13986-4\_26.
- Mendel, E., Mariano, L.Z., Drago, I., Loureiro, S., Rauber, T., and Varejão, F.M. (2008) *Automatic Bearing Fault Pattern Recognition using Vibration Signal Analysis*. In: Proceedings of IEEE International Symposium on Industrial Electronics (ISIE). Cambridge, UK. pp. 955–960. DOI 10.1109/ISIE.2008.4677026.
- Fabris, F., Drago, I., and Varejão, F.M. (2008) *A Multi-Measure Nearest Neighbor Algorithm for Time Series Classification*. In: Proceedings of 11th Ibero-American Conference on AI (IBERAMIA). Lisbon, Portugal. pp. 153–162. DOI 10.1007/978-3-540-88309-8\_16.

Technical reports:

- Drago, I., Ricciato, F., and Sadre, R. (2016) *Report from the 6th PhD School on Traffic Monitoring and Analysis (TMA)*. Computer Communication Review. 46(4). pp. 36-40. DOI 10.1145/3027947.3027954
- Drago, I., Sadre, R., and Pras, A. (2011) *Report of the Third Workshop on the Usage of NetFlow/IPFIX in Network Management*. Journal of Network and Systems Management. 19(4). pp. 529-535. DOI 10.1007/s10922-011-9204-2.
- Drago, I., Barbosa, R.R.R., Sadre, R., Pras, A., and Schonwalder, J. (2011) *Report of the Second Workshop on the Usage of NetFlow/IPFIX in Network Management*. Journal of Network and Systems Management. 19 (2). pp. 298-304. DOI 10.1007/s10922-010-9163-z.
- Pras, A., Sperotto, A., Moreira Moura, G.C., Drago, I., Barbosa, R.R.R., Sadre, R., Schmidt, R.O., and Hofstede, R.J. (2010) *Attacks by "Anonymous" WikiLeaks Proponents not Anonymous*. Technical Report TR-CTIT-10-41. CTIT. Enschede, the Netherlands.

## Participation in Research Projects

**Sep 2015 – Jul 2016**    **Clue – In-depth and Automated Metering of Cloud Usage**  
 Sponsor: **Cisco Inc., USA**  
 Reference: **Prof. Marco Mellia and Prof. Mario Baldi (Polito)**

CLUE aims at building a system that (i) automatically measures performance and usage of cloud services; (ii) transforms the obtained knowledge in policies on the network control and forward planes. CLUE was envisioned to be deployed at large enterprises that outsource critical services to cloud providers. Such enterprises need to know, for example, whether internal users are able to interact with the remote services, and whether non-accredited services are used.

My goal in CLUE was to research new methodologies to identify traffic of cloud services, even in situations where classical methods are not effective anymore (e.g., because of encryption).

**Apr 2015 – Nov 2015**    **mPlane – Building an Intelligent Measurement Plane for the Internet**  
 Sponsor: **European Commission - FP7**  
 Reference: **Prof. Marco Mellia (Polito)**

mPlane was a FP7 project focusing on Internet traffic monitoring and analysis. mPlane is a distributed and flexible infrastructure to perform measurements, composed of several layers to collect, store and analyze data.

I joined mPlane during my time as an "assegnista di ricerca" at Polito, participating in the last year of the project. I was involved on the evaluation of data collected via the mPlane infrastructure on real use cases, e.g., by applying big data and data mining algorithms and methodologies to extract knowledge from the data.

**Jan 2014 – Apr 2015**    **Partnership between Narus Inc. and Polito**  
 Sponsor: **Narus Inc., USA**  
 Reference: **Prof. Marco Mellia (Polito)**

The partnership consisted of several sub-projects with a common goal of defining and engineering a novel platform for cyber-security analysis.

I worked on this project while receiving a "borsa per attività di ricerca" from Polito. My tasks were to build a system for monitoring high-speed Internet links, as well as to evaluate the resulting data, aiming at the identification of performance impairments on networked applications.

**Dec 2011 – Feb 2012**    **TMA – Traffic Monitoring and Analysis**  
 Sponsor: **European Commission – COST Action IC0703**  
 Reference: **Prof.dr.ir. Aiko Pras (University of Twente)**



TMA was a COST action aiming at increasing the quality and the impact of European research in the field of traffic monitoring and analysis.

The TMA COST Action supported a number of Short-Term Scientific Missions (STSM). I was awarded one of those STSM while I was a PhD student at the University of Twente. I expended three months visiting the group of Prof. Marco Mellia in Polito, studying the characteristics of emerging cloud services. Results of this collaboration were documented in a paper published at the IMC 2012, and later recognized by the IRTF with an Applied Networking Research Prize.

**Mar 2009 – Oct 2012**    **SEQUAL – Service Optimization and Quality**  
Sponsor: **Ministry of Economic Affairs, the Netherlands**  
Reference: **Prof.dr.ir. Aiko Pras (University of Twente)**

The goal of this project was to design and develop a quantitative-model based framework to measure and control end-to-end quality of complex services running over the web and relying on heterogeneous third-party services.

My responsibility was to build methods for measuring the parameters needed for such framework, considering the complexity of modern Internet services. This project was the starting point for my PhD thesis, where a methodology to monitor cloud storage services is presented.

**2006 – 2008**            **ADD-RPD – An Intelligent System to Detect Failure Patterns in Motor-Pumps**  
Sponsor: **Petrobras S.A. and ANP – National Agency of Petroleum (Brazil)**  
Reference: **Prof. D.Sc. Flavio Miguel Varejao (UFES)**

ADD-RPD is a system for the early detection of failures in motor-pumps used to extract oil from the deep sea. Platforms extracting oil in such conditions are remote, and companies want to reduce engineers' travels for routine inspections, as well as downtime resulting from maintenance. ADD-RPD employs data mining to evaluate data collected from several sensors deployed in the platforms, anticipating preventive interventions in motor-pumps.

I worked on this project while I was an assistant at UFES, after finishing my master's. My tasks included understanding and pre-processing the input data, as well as building a prototype of ADD-RPD for validating the approach.

**2003 – 2005**            **MIP – Improving Fraud Identification**  
Sponsor: **ESCELSA and ANEEL – Electricity Regulatory Agency (Brazil)**  
Reference: **Prof. D.Sc. Flavio Miguel Varejao (UFES)**

Illegal installations in the power grid are a major problem in Brazil, and they are fought by manually inspecting customers' premises. MIP is a system to select which customers should be inspected by the power delivery companies. MIP applies machine learning on datasets of customers' records (e.g., histories of electricity consumption) to increase the odds that inspections are successful.

I was part of the team responsible for developing MIP, from selecting the most suitable algorithms for the classification task, to the deployment and validation of a prototype in the company sponsoring the project. This project was the basis for my work as master's student at UFES.

## Topics of Interest

- Computer networking
- Internet measurements
- Big data
- Quality of experience
- Network security

## Additional Activities

- I have participated in the **organization of international conferences** (e.g., *MANWEEK 2009* and *TMA 2016*), as well as acted as Technical Programme Committee Member (e.g., *IFIP Networking 2016*).
- I am a regular **reviewer of several journals and conferences**, such as the *IEEE Journal on Selected Areas in Communications*, the *IEEE Communications Magazine*, the *IEEE Transactions on Network and Service Management*, and the *IEEE Transactions on Storage*.
- I have **advised a number of master and bachelor students**, as well as performed several teaching assistantship during the time I was a PhD Student and a Postdoctoral Researcher.

Torino, 1st January 2017,

Idilio Drago