CURRICULUM VITAE ABREVIADO (CVA)

Part A. PERSONAL INFORMATION

First name	Pedro		
Family name	Reviriego		
Gender (*)	Male	Birth date (dd/mm/yyyy)	31/10/1970
Social Security, Passport, ID number	09778438-B		
e-mail	pedro.reviriego@upm.es	URL Web	Pedro Reviriego - Google Académico
Open Researcher and Contributor ID (ORCID) (*)		0000-0003-2540-5234	

A.1. Current position

Position	Catedrático de Universidad (Professor)		
Initial date	25/04/2025		
Institution	Universidad Politécnica de Madrid		
Department/Center	Information Processing and Telecommunications Center		
Country	Spain	Teleph. number	626068722
Key words	Artificial Intelligence, Large Language Models, Computing, Networks		

A.2. Previous positions (research activity interuptions, indicate total months)

Period	Position/Institution/Country/Interruption cause
2022-2025	Associate Professor, Universidad Politécnica de Madrid, Spain
2018-2022	Visiting Professor, Universidad Carlos III de Madrid, Spain
2007-2018	Associate Professor, Nebrija University, Spain
2003-2007	DMTS, LSI/Broadcom, Spain
2003-2003	Visiting Professor, Universidad Carlos III de Madrid, Spain
2000-2003	Principal Engineer, Massana, Spain
1997-2000	Senior Engineer, Teldat, Spain

A.3. Education

PhD, Licensed, Graduate	University/Country	Year
Ing. Telecomunicación	Universidad Politécnica de Madrid, Spain	1994
Dr. Ing. Telecomunicación	Universidad Politécnica de Madrid, Spain	1997

Part B. CV SUMMARY

After completing my PhD in 1997, I joined the R&D department of Teldat a Spanish company that designs and manufactures routers. Then in 2000, I moved on to an Irish a start-up that was developing Ethernet transceiver ASICs and that later was acquired by an US company (now part of Broadcom). The transceivers that we developed were licensed to Intel and had a large commercial success. My experience in industry has given me a different perspective from that of purely academic researchers that has been very useful. I returned to academia in 2007.

Since then, I have made significant contributions in different areas and expanded my research topics in new directions that I describe in the following.

My first key contribution was on Energy Efficiency in Ethernet on which I was a pioneer contributing to the IEEE 802.3 standards and writing seminal papers on the field with more than 470 citations in Google Scholar. This work was recognized with two Google Research Awards funding my research on this area. My leadership enabled me to establish fruitful collaborations with Kenneth Christensen of University of South Florida and with Bruce Nordman and Mike Bennett of the Lawrence Berkeley National Laboratory, all of them pioneers in the field of energy efficiency in networking. The second topic on which I have become a leader in the last ten years is the design and implementation of Error Correction Codes (ECCs) to protect memories. In this area, I have pioneered the use of one step majority logic decodable codes for memory protection with a large number of publications. In this area, again my capacity to propose new ideas led to long term collaborations with Chris Bleakley

and Mark Flanagan of University College Dublin, with Dr. Costas Argyrides of AMD or with Marco Ottavi of University of Rome “Tor Vergata” among others. More recently, my focus on the fault tolerance area is on fault-tolerant machine learning in close collaboration with Fabrizio Lombardi from Northeastern University and Zhen Gao from Tianjin University. In this area we have proposed fault tolerant implementations for many machine learning algorithms including complex neural networks.

In parallel with those topics related to circuit level fault tolerance, my work on networking has evolved to focus on data structures for packet processing first and then more broadly on approximate data structures also known as data sketches that are widely used in large scale Big Data processing systems. Today switching ASICs handle more than 50 Tb/s performing advanced processing on each packet. This requires very efficient data structures. In this topic my research has both an academic and an industrial side. On the academic side, I have made significant contributions for example related to the implementation of cuckoo hashing or to support exact match with one memory access. I have also made contributions on the area of Bloom filters and proposed the Adaptive Cuckoo Filter. In this area, I have also been able to establish fruitful collaborations with Michael Mitzenmacher, of Harvard University, Ori Rottenstreich of Technion and more recently with Daniel Ting of Meta, Otmar Ertl of Dynatrace or Rasmus Pagh, the inventor of cuckoo hashing. On the industrial side, I work closely with the switch architecture group of Mellanox (now part of NVIDIA) on developing algorithms and data structures for their switches. This has produced more than twelve US/EU patents or patents applications so far. Many of the algorithms that I have developed are currently implemented in the Spectrum switching ASICs product line and have enabled an unprecedented flexibility and capacity of the packet processing tables and algorithms in which design I have participated are also implemented in the ConnectX SmartNIC product line.

More recently, since 2019, I have been working on artificial intelligence focusing on fault tolerant implementations and on the evaluation of large generative AI models such as LLMs or text to image generators and on the evaluation of Generative AI expanding the scope of my research. Currently my research is completely focused on LLMs and generative AI with one of the main topics being the intersection of LLMs and psycholinguistics working in close collaboration with Marc Brysbaert from Ghent University.

I have also been the deputy director for research of the school of engineering and the director of the Doctoral program at Nebrija University for more than four years. I have eight graduated PhD students that are currently working in industry (AMD-Xilinx, ASML, TTTech and Indra) and in academia (Saint Louis University, Universidad Complutense, Universidad Rey Juan Carlos and Universidad Politécnica de Madrid). I am also associate editor of IEEE Transactions on Emerging Topics in Computing. I have more than 300 papers on journals indexed on the JCR, most of them in Transactions of the IEEE and more than 6500 citations on Google Scholar.

Part C. RELEVANT MERITS

C.1. Selected contributions

1. *Contribution to the Development of Ethernet Transceivers*

I participated in the first European team (an Irish-Spanish startup) that developed a 1 Gb/s Ethernet transceiver, leading the design of interference cancellation algorithms. The startup was acquired by Agere Systems (formerly the microelectronics division of ATT) for 35 million USD. These transceivers were among the most advanced mixed-signal devices of their time, and the design was commercialized by Intel and used in tens of millions of devices. During this period, I was the leading contributor of the company to the IEEE 802.3an standard for 10 Gb/s transceivers making contributions to the transmitter based equalization schemes (IEEE 802.3an-C1 or IEEE 802.3an-C2).

2. *Contribution to Energy Efficiency in Ethernet*

I contributed to the development of the Energy Efficient Ethernet standard (IEEE 802.3az) in collaboration with researchers from the Lawrence Berkeley National Laboratory and the University of South Florida. This initiative was a pioneer in network energy efficiency and is currently implemented in hundreds of millions of devices. We presented the first analytical (Performance

Evaluation of Energy Efficient Ethernet) and empirical (An Initial Evaluation of Energy Efficient Ethernet) evaluations of the standard showing the dependency of the energy savings on the traffic shaping and the need to implement packet coalescing to maximize energy savings. In collaboration with researchers from Lawrence Berkeley National Laboratory and the University of South Florida we published the reference paper on this topic (IEEE 802.3az: the Road to Energy Efficient Ethernet) which has close to 500 citations in Google Scholar. I also made contributions to the IEEE 802.3 standards (Reach and Energy Efficiency in NG-BASE-T). This work was recognized with two Google Research Awards, where I was as Principal Investigator.

3. *Contribution to the Development of Ethernet Equipment for Space Environments*

I participated in the H2020 SEPHY project (SEPHY) to develop a European Ethernet transceiver focusing on the transceiver architecture and algorithms. The transceiver was successfully implemented and tested leading also to a few publications like The Space Ethernet Physical Layer Transceiver (SEPHY) Project: a Step Towards Reliable Ethernet in Space. I am currently involved in the development of an Ethernet switch (Space-L3S) for space led by an US company, Ecliptic Enterprises which will be used in leading space missions.

4. *Contribution to Advanced Packet Switching*

I have developed advanced algorithms for high-speed switches in collaboration with NVIDIA/Mellanox resulting in more than 12 patents. In particular I have designed innovative solutions for example for:

- Advanced data structures for switching (US10496680B2 - High-performance bloom filter array - Google Patents and US9892057B2 - Single double cuckoo hash).
- Longest Prefix Match (LPM) for IP forwarding (US20170366459A1 - Jump on a Match Optimization for Longest Prefix Match using a Binary Search Tree).
- Algorithmic Ternary Content Addressable Memories (TCAMs) for access control lists and other packet classification rules (US20170052731A1 - Efficient lookup of TCAM-like rules in RAM).
- Caching schemes for both LPM and TCAMs (EP3258657A1 - IP route caching with two search stages on prefix length and US20190036821A1 - Efficient caching of TCAM rules in RAM).
- Machine learning optimizations for packet processing (US20240146652A1 - Optimizing hash table selection for packet processing driven by machine learning).

Many of these schemes are implemented in NVIDIA products and used in data centers. The research was funded by NVIDIA/Mellanox with more than ten research contracts and produced also a few papers, for example: Algorithmic TCAMs: Implementing Packet Classification Algorithms in Hardware or Flexible Packet Matching with Single Double Cuckoo Hash

5. *Contribution to Fault-Tolerant Memory Design*

I have proposed innovative error correction schemes for memories both by optimizing error correction codes widely used to protect memories and by proposing the use of other codes not previously used for memory protection. Some examples of works with a relevant number of citations:

- Optimized the decoding speed of classical SEC codes used for memory applications (A Method to Construct Low Delay Single Error Correction Codes for Protecting Data Bits Only).
- Introduced the use of difference set codes for memory protection and optimized the decoding for memory applications (Efficient Majority Logic Fault Detection With Difference-Set Codes for Memory Applications).
- Designed new codes to correct the typical multibit errors in memories (MCU Tolerance in SRAMs Through Low-Redundancy Triple Adjacent Error Correction).
- Modeled the impact of multibit errors on memory reliability (Reliability Analysis of Memories Suffering Multiple Bit Upsets) .

- Analyzed the placement of the error correcting blocks on the memory (Protection of Memories Suffering MCUs Through the Selection of the Optimal Interleaving Distance).

Many of these works were done in collaboration with international researchers from institutions such as AMD, Cisco and University College Dublin for example.

6. *Contribution to Energy-Efficient Computing Systems*

Designed blocks and algorithms for stochastic computing and conditional computing schemes to reduce energy consumption, in collaboration with researchers from various institutions. A few examples of relevant works:

- An invited review paper on building blocks for stochastic computing (From Multipliers to Integrators: A Survey of Stochastic Computing Primitives).
- Optimized dividers which are one of the key blocks in stochastic computing (Stochastic Dividers for Low Latency Neural Networks or Delta Sigma Modulator-Based Dividers for Accurate and Low Latency Stochastic Computing Systems).
- Optimized implementations of Neural Networks with stochastic computing (Energy-Efficient Stochastic Computing (SC) Neural Networks for Internet of Things Devices With Layer-Wise Adjustable Sequence Length (ASL)).
- Conditional computing to reduce energy in IoT implementations of machine learning algorithms (Adaptive Resolution Inference (ARI): Energy-Efficient Machine Learning for Internet of Things).

All the publications are in collaboration with leading experts in the field like Jie Han from the University of Alberta or Fabrizio Lombardi form Northeastern University.

7. *Contribution to the Design of Probabilistic Data Structures*

I have optimized existing probabilistic data structures (aka as data sketches) and proposed new algorithms with applications in networking and computing. Some examples of relevant works are:

- Data structures for traffic monitoring in networks (Lightweight Acquisition and Ranging of Flows in the Data Plane | 2024 ACM SIGMETRICS).
- Data structures for databases (InfiniFilter: Expanding Filters to Infinity and Beyond | 2023 ACM SIGMOD).
- Among the first to propose adaptive approximate membership check filters (Adaptive Cuckoo Filters | ACM Journal of Experimental Algorithmics).
- Security and privacy of data sketches (Security of HyperLogLog (HLL) Cardinality Estimation: Vulnerabilities and Protection or On the Privacy of Counting Bloom Filters).

Most of the works are in collaboration with international researchers, for example Michael Mitzenmacher (Harvard) and Rasmus Pagh (University of Copenhagen) and also from industry Daniel Ting (Meta) or Otmar Ertl (Dynatrace).

8. *Contribution to Fault-Tolerant Machine Learning Systems*

Conducted pioneering studies on the robustness of both classical and generative machine learning algorithms. Most of this work is done in collaboration with professors Fabrizio Lombardi (Northeastern University), Shanshan Liu (University of Electronic Science and Technology of China) and Zhen Gao (Tianjin University). Relevant examples of contributions in this area are:

- Proposed efficient concurrent error detection schemes for large machine learning systems (Concurrent Classifier Error Detection (CCED) in Large Scale Machine Learning Systems).
- Analyzed the impact of errors on FPGA implementations of LLMs and CNNs (Robustness against Faults in Configuration Memories of FPGA-based LLMs and Modeling the Effect of SEUs on the Configuration Memory of SRAM-FPGA-Based CNN Accelerators).
- Proposed fault tolerant schemes for machine learning classifiers (Error-Tolerant Computation for Voting Classifiers With Multiple Classes).
- Proposed fault tolerant schemes for neural networks (Selective Neuron Re-Computation (SNRC) for Error-Tolerant Neural Networks).

- Analyzed the impact of errors on LLMs (On the Dependability of Bidirectional Encoder Representations from Transformers (BERT) to Soft Errors).

9. Contribution to the Evaluation of Large Language Models (LLMs)

Developed new methodologies, tools and benchmarks for evaluating and applying LLMs to other disciplines. For example:

- By incorporating insights from cognitive science, in collaboration with Professor Marc Brysbaert (Ghent University), a leading figure in psycholinguistics (Using large language models to estimate features of multi-word expressions: Concreteness, valence, arousal | Behavior Research Methods or Moving beyond word frequency based on tally counting: AI-generated familiarity estimates of words and phrases are an interesting additional index of language knowledge | Behavior Research Methods).
- By considering the implications of LLM on linguistic features (Beware of Words: Evaluating the Lexical Diversity of Conversational LLMs using ChatGPT as Case Study | ACM Transactions on Intelligent Systems and Technology).
- By analyzing the knowledge of LLMs of the words in a language (Open Conversational LLMs do not know most Spanish words | Conde | Procesamiento del Lenguaje Natural).
- By exploring the limitations of LLMs and its implications (Can ChatGPT Learn to Count Letters?)

Again, most of the works are in collaboration with international researchers and I have been recently awarded a MIT Seed Fund to collaborate with Prof. Ted Gibson of the Department of Brain & Cognitive Sciences to study the interrelations of AI and cognitive science.

10. Contribution to IEEE publications

I have collaborated with the IEEE for more than two decades and I am currently Associate Editor for IEEE Transactions on Emerging Topics in Computing, a high-impact journal focused on current computing trends. I have also edited a special issue on fault-tolerant machine learning algorithms in the same journal and have been a member of the Nano Technology Computing initiative.

C.2. Research projects and Contracts

The information on the research projects on which I have been PI is summarized in the following:

1. One MIT-UPM Seed Fund project with Ted Gibson of the MIT on psycholinguistics and AI.
2. Four Spanish National Plan Projects (PID2022-136684OB-C22, PID2019-104207RB-I00, ESP2014-54505-C2-1-R and AYA2009-13300-C03) two of them including an FPI scholarship. The most recent one (PID2022-136684OB-C22) is a coordinated project with the UC3M team of subproject 2 in the current Psycholinguistics and Artificial Yntelligence (PLAY) proposal.
3. One Digital Transition project the Agencia Estatal de Investigación (TED2021-130118B-I00).
4. One Proof of Concept funded by the Agencia Estatal de Investigación (PDC2022-133888-I00) associated to National Project (PID2019-104207RB-I00).
5. One H2020 project (SEPHY, 2015-2018) with more than 340K in funding.
6. One HE-CJU project (SMARTY, 2023- 2027) with more than 470K in funding.
7. Two Google Research Awards.
8. One Comunidad de Madrid Research Project (2009-2010).

I have also participated as a researcher in several H2020/HE (for example PASSION, SMOOTH, METROHAUL, PIMCITY, B5G-OPEN) and Comunidad de Madrid (TUCAN6-CM, TAPIR-CM, EMPATIA-CM) projects.

Contracts: I have been the PI in more than ten research contracts with Mellanox/NVIDIA with more than 250K euros in funding in the period 2014-2022. Those contracts were all related to algorithms and optimizations of their switches and SmartNICs and thus directly related to the network topics. I have also been the PI in more than 8 research contracts with Arquimea/Ecliptic in the period 2019-2025. These contracts were all related to the use of Ethernet in space systems.