# Jose Maria Arnau

## CONTACT INFORMATION


| | |
|---|---|
| Dr. Jose Maria Arnau | Phone: +34 93 4054039 |
| Department of Computer Architecture | E-mail: jarnau@ac.upc.edu |
| UPC BarcelonaTech | www: http://jarnau.site.ac.upc.edu/ |
| Jordi Girona 1-3, C6-115, Campus Nord UPC | 08034 Barcelona, Spain |


## EDUCATION

- **UPC BarcelonaTech, March 2011 - April 2015**
  Ph.D. in Computer Architecture
  Grade: Excellent Cum Laude
  Thesis: “Energy-Efficient Mobile GPU Systems”
  Advisers: Prof. Joan-Manuel Parcerisa, Dr. Polychronis Xekalakis
- **UPC BarcelonaTech, October 2010 - September 2011**
  MSc. Degree on Computer Architecture, Networks and Systems
  Grade: 9.75/10 (Award for achieving the highest overall mark)
  Master’s Thesis: “High Performance, Ultra-Low Power Streaming Systems”
  Adviser: Prof. Joan-Manuel Parcerisa
  Co-Advisers: Dr. Polychronis Xekalakis, Prof. Antonio Gonzalez
- **Universitat Jaume I, September 2003 - June 2008**
  BSc. Degree on Computer Engineering
  Grade: 9.64/10 (Award for achieving the highest overall mark)
  Final Year Project: “Development of a Game Engine”
  Adviser: Prof. Miguel Chover Selles

## PROFESSIONAL EXPERIENCE

- **Sep 2017 - Present: Distinguished Researcher, UPC BarcelonaTech**
  I hold a distinguished researcher position at UPC in the ARCO (Architectures and Compilers) research group. My research investigates energy-efficient hardware architectures for cognitive computing. I am currently advising four PhD students.
- **May 2015 - Sep 2017: Postdoctoral Researcher, UPC BarcelonaTech**
  I held a postdoctoral research scientist position in the ARCO research group. My work investigated energy-efficient architectures for cognitive processors, exploring different approaches for low-power and high-performance deep learning acceleration such as GPU, FPGA and ASIC.
- **August 2010 - April 2015: Research Assistant, UPC BarcelonaTech**
  I held a research scientist position in the ARCO group, supported by an FI-Research grant from the Catalan Government. I investigated techniques that improve the energy-efficiency of mobile GPUs. My Ph.D. analyzed the suitability of decoupled access/execute architectures for graphics hardware to tolerate memory latency. My research also proposed memory bandwidth saving techniques to avoid redundant texture fetches and memoization schemes to remove redundant computations.
- **August 2014 - December 2014: Architecture Intern, NVIDIA Corporation**
  I worked as an intern in the Applied Architecture team at NVIDIA. My job consisted on performing GPU power and performance analysis. I also worked on OpenGL compute shader optimization.

- **July 2013 - September 2013: Visiting Researcher, University of Edinburgh**
  I performed a co-design space exploration of compiler optimizations and hardware parameters in GPGPU environments. I used an LLVM-based OpenCL compiler to evaluate the power-performance trade-offs of multiple compiler optimizations, and I explored the hardware design space by using GPGPUSim and GPUWattch.
- **September 2008 - July 2010: Software Developer, Institute of Ceramic Technology**
  I developed several applications for scientific data visualization using OpenGL for rendering 3D tomographies. I also implemented multiple image processing algorithms for data visualization.
- **January 2008 - June 2008: Research Assistant, Universitat Jaume I**
  I developed a game engine for teaching undergraduates with the aim of facilitating the learning of computer graphics. I designed the engine to be simple and clear, but also powerful enough to support visually compelling 3D scenes.

## PUBLICATIONS

### Peer-Reviewed Conferences

- "Boosting LSTM Performance Through Dynamic Precision Selection". Franyell Silfa, Jose-Maria Arnau, Antonio Gonzalez. In Proceedings of the 27th IEEE International Conference on High Performance Computing, Data, and Analytics (**HiPC**), December 2020.
- "Demystifying Power and Performance Bottlenecks in Autonomous Driving Systems". Pedro Exenberger, Jose-Maria Arnau, Antonio Gonzalez. In Proceedings of the IEEE International Symposium on Workload Characterization (**IISWC**), October 2020.
- "Neuron-Level Fuzzy Memoization in RNNs". Franyell Silfa, Gem Dot, Jose-Maria Arnau, Antonio Gonzalez. In Proceedings of the 52nd IEEE/ACM International Symposium on Microarchitecture (**MICRO**), October 2019.
- "SCU: A GPU Stream Compaction Unit for Graph Processing". Albert Segura, Jose-Maria Arnau, Antonio Gonzalez. In Proceedings of the 46th International Symposium on Computer Architecture (**ISCA**), June 2019.
- "E-PUR: An Energy-Efficient Processing Unit for Recurrent Neural Networks". Franyell Silfa, Gem Dot, Jose-Maria Arnau, Antonio Gonzalez. In Proceedings of the 27th International Conference on Parallel Architectures and Compilation Techniques (**PACT**), November 2018.
- "The Dark Side of DNN Pruning". Reza Yazdani Aminabadi, Marc Riera, Jose-Maria Arnau, Antonio Gonzalez. In Proceedings of the 45th IEEE/ACM International Symposium on Computer Architecture (**ISCA**), June 2018.
- "Computation Reuse in DNNs by Exploiting Input Similarity". Marc Riera, Jose-Maria Arnau, Antonio Gonzalez. In Proceedings of the 45th IEEE/ACM International Symposium on Computer Architecture (**ISCA**), June 2018.
- "A Novel Register Renaming Technique for Out-of-Order Processors". Hamid Tabani, Jose-Maria Arnau, Jordi Tubella, Antonio Gonzalez. In Proceedings of the 24th IEEE International Symposium on High-Performance Computer Architecture (**HPCA**), February 2018.
- "UNFOLD: A Memory-Efficient Speech Recognizer Using On-The-Fly WFST Composition". Reza Yazdani Aminabadi, Jose-Maria Arnau, Antonio Gonzalez. In Proceedings of the IEEE/ACM International Symposium on Microarchitecture (**MICRO**), October 2017.
- "An Ultra Low-Power Hardware Accelerator for Acoustic Scoring in Speech Recognition". Hamid Tabani, Jose-Maria Arnau, Jordi Tubella, Antonio Gonzalez. In Proceedings of the 26th International Conference on Parallel Architectures and Compilation Techniques (**PACT**), Sep. 2017.

- "An Ultra Low-Power Hardware Accelerator for Automatic Speech Recognition". Reza Yazdani Aminabadi, Albert Segura, Jose-Maria Arnau, Antonio Gonzalez. In Proceedings of the IEEE/ACM International Symposium on Microarchitecture (**MICRO**), October 2016.
- "Eliminating Redundant Fragment Shader Executions on a Mobile GPU via Hardware Memoization". Jose-Maria Arnau, Joan-Manuel Parcerisa and Polychronis Xekalakis. In Proceedings of the 41st IEEE/ACM International Symposium on Computer Architecture (**ISCA**), June 2014.
- "Parallel Frame Rendering: Trading Responsiveness for Energy on a Mobile GPU" . Jose-Maria Arnau, Joan-Manuel Parcerisa and Polychronis Xekalakis. In Proceedings of the 22nd IEEE/ACM International Conference on Parallel Architectures and Compilation Techniques (**PACT**), September 2013.
- "TEAPOT: A Toolset for Evaluating Performance, Power and Image Quality on Mobile Graphics Systems" . Jose-Maria Arnau, Joan-Manuel Parcerisa and Polychronis Xekalakis. In Proceedings of the 27th ACM International Conference on Supercomputing (**ICS**), June 2013.
- "Boosting Mobile GPU Performance with a Decoupled Access/Execute Fragment Processor". Jose-Maria Arnau, Joan-Manuel Parcerisa and Polychronis Xekalakis. In Proceedings of the 39th IEEE/ACM International Symposium on Computer Architecture (**ISCA**), June 2012.

### Peer-Reviewed Journals

- "Design and Evaluation of an Ultra Low-Power Human-Quality Speech Recognition System". Dennis Pinto, Jose-Maria Arnau, Antonio Gonzalez. Accepted for publication in ACM Transactions on Architecture and Code Optimization (**TACO**).
- "LAWS: Locality-AWare Scheme for Automatic Speech Recognition". Reza Yazdani, Jose-Maria Arnau and Antonio Gonzalez. **IEEE Transactions on Computers**, vol. 69, no. 8, pp. 1197-1208, 1 Aug. 2020, doi: 10.1109/TC.2020.2991002.
- "A Low-Power, High-Performance Speech Recognition Accelerator". Reza Yazdani, Jose-Maria Arnau and Antonio Gonzalez. **IEEE Transactions on Computers**. doi: 10.1109/TC.2019.2937075
- "CGPA: Coarse-Grained Pruning of Activations for Energy-Efficient RNN Inference". Marc Riera, Jose-Maria Arnau and Antonio Gonzalez. **IEEE Micro**. doi: 10.1109/MM.2019.2929742
- "Performance Analysis and Optimization of Automatic Speech Recognition". Hamid Tabani, Jose-Maria Arnau, Jordi Tubella and Antonio Gonzalez. IEEE Transactions on Multi-Scale Computing Systems (**TMSCS**), vol. 4, no. 4, pp. 847-860, 2018.
- "Low-Power Automatic Speech Recognition Through a Mobile GPU and a Viterbi Accelerator". Reza Yazdani, Albert Segura, Jose-Maria Arnau and Antonio Gonzalez. **IEEE Micro**, vol. 37, no. 1, pp. 22-29, 2017.
- "Study of the Pressing Operation of Large-sized Tiles Using X-ray Absorption". J.L. Amoros, G. Mallol, D. Llorens, J. Boix, J.M. Arnau, C.Feliu, J.A. Cerisuelo and J.J. Gargallo. Journal of the Spanish Ceramic and Glass Society, vol. 49, no. 4, pp. 279-288, 2010.
- "Apparent Density Measurement of the Ceramic Tiles in a Quick, Harmless and Non-destructive Way". G. Mallol, D. Llorens, J. Boix, M. Aguilella, L. Foucard and J.M. Arnau. Journal of the Spanish Ceramic and Glass Society, vol. 49, no. 6, pp. 393-398, 2010.

## AWARDS AND HONORS

- **Intel Doctoral Student Honor Programme: $**35000 awarded by Intel Corporation for performing outstanding research in the area of Computer Architecture (2012).
- **FI Research Grant:** funding from the Catalan Government for a three year Ph.D. (2011).
- **Best Student Graduating in Master's Degree on Computer Architecture, Networks and Systems:**

awarded by the Barcelona School of Informatics at the UPC BarcelonaTech (2011).

- **Second Best Student Graduating in Computer Engineering in Spain:** awarded by the Spanish Government (2010).
- **Best Student Graduating in Computer Engineering in the Valencian Community:** awarded by the Valencian Government (2009).
- **Best Student Graduating in Computer Engineering:** awarded by the Universitat Jaume I (2008).
- **Best Student Graduating in Computer Engineering:** awarded by the School of Technology and Experimental Sciences of Castellon (2008).
- **Research Collaboration Grant from the Spanish Ministry of Education:** funding for a six month research collaboration in the Computer Graphics Group at the Universitat Jaume I of Castellon (2007).
- **HiPEAC Paper Award** for the publication of “SCU: A GPU Stream Compaction Unit for Graph Processing” in the International Symposium on Computer Architecture (2019).
- **HiPEAC Paper Award** for the publication of “The Dark Side of DNN Pruning” in the International Symposium on Computer Architecture (2018).
- **HiPEAC Paper Award** for the publication of “Computation Reuse in DNNs by Exploiting Input Similarity” in the International Symposium on Computer Architecture (2018).
- **HiPEAC Paper Award** for the publication of “A Novel Register Renaming Technique for Out-of-Order Processors” in the International Symposium on High-Performance Computer Architecture (2018).
- **HiPEAC Paper Award** for the publication of "UNFOLD: A Memory-Efficient Speech Recognizer Using On-The-Fly WFST Composition" in the International Symposium on Microarchitecture (2017).
- **HiPEAC Paper Award** for the publication of "An Ultra Low-Power Hardware Accelerator for Automatic Speech Recognition" in the International Symposium on Microarchitecture (2016).

## TEACHING

**Past lectures:**

- Sep 2017 – Jan 2020: *Computer Organization*, BSc. Degree on Computer Engineering, UPC
- Sep 2019 – Jan 2020: *Processor Design*, Master MIRI, UPC
- Feb 2020 – Jul 2020: *Nanoelectronic Circuit Design*, Master MIRI, UPC

**Current students under my supervision:**

- 4 PhD students: Daniel Pinto, Raul Taranco, Pedro Exenberger, Mojtaba Abaie
- 1 Master student: Elisabet Valle

**Graduated students under my supervision (first job after PhD)**:

- 5 PhD students: Hamid Tabani (BSC), Reza Yazdani (Microsoft), Marc Riera (UPC), Albert Segura (Apple), Franyell Silfa (UASD)
- 2 Master students: Berta Delgado (now at HP), Azmath Syeda (now at Qualcomm)
- 4 Bachelor students: Ferran Olivera (now at HP), Alfonso Castaño, Oscar Mañas, David Ramal

## LANGUAGES

- **Spanish:** Native speaker
- **Catalan:** Native speaker
- **English:** Fluent

**OTHER**

- Organization of conferences:

  - Publications Chair of the IEEE/ACM International Symposium on High Performance Computer Architecture (HPCA), Barcelona (Spain), February 2016.

  - Shadow program committee member of the 23rd ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS).

- Participation in research projects:

  - Researcher staff. "CoCoUnit: An Energy-Efficient Processing Unit for Cognitive Computing". European Research Council. Sep 2019 - Present.

  - Researcher staff. “Arquitecturas de Sistemas de Computación Inteligentes, Ubicuos y Energéticamente Eficientes”. Spanish Ministry of Economy. TIN2016-75344-R. Jan 2017 - Present.

  - Researcher staff. “Microarquitecturas y Compiladores para Futuros Procesadores III”. Spanish Ministry of Economy. TIN2013-44375. Jan 2014 - Dec 2016.

  - Researcher staff. “Microarquitecturas y Compiladores para Futuros Procesadores II”. Spanish Ministry of Science and Technology, TIN2010-18368, Jan 2011 - Dec 2013.