Ph.D. in Information Technology:

Ashouri, Pagano and Stamelakos Final Dissertations

DEIB Conference Room

December 20th, 2016

10.00 am

First Ph.D. presentation and discussion:

Amir H. ASHOURI – XXVIII Cycle

"Compiler Autotuning Using Machine Learning Techniques"

Advisor: Prof. **Cristina Silvano**

Abstract:

Diversity of today's architectures have forced programmers to spend additional effort to port and tune their application code across different platforms.

Compilers within this process need additional tuning which is a hard task itself. Recent compilers offer a vast number of multilayered optimizations, capable of targeting different code segments of an application.

Choosing among these optimizations can significantly impact the performance of the code being optimized. The selection of the right set of compiler optimizations for a particular code segment is a very hard problem, but finding the best ordering of these optimizations adds further complexity. In fact, finding the best ordering is a long standing problem in compilation research called the phase-ordering problem. The traditional approach of constructing compiler heuristics to solve this problem simply can not cope with the enormous complexity of choosing the right ordering of optimizations for every code segment in an application.

In the PhD thesis, we provide break-through approaches to tackle and mitigate the well-known problems of compiler optimization using machine learning techniques. We hope that it will be useful for a wide range of readers, including computer architects, compiler developers, researchers and technical professionals.

Second Ph.D. presentation and discussion:

Roberto PAGANO – XXIX Cycle

"Context-Driven Recommender Systems"

Advisor: Prof. **Paolo Cremonesi**

Abstract:

Recommender systems always relied on personalization and context has only been used to improve personalization performance. However algorithms that decouple the context exploitation from personalization were not fully explored.

This thesis presents Context-Driven Recommender Systems, a new family of recommendation algorithms that personalize on user's intent and situation, exploring their applicability in five application domains. This work shows that leveraging contextual information in these domains helps in addressing the challenges that they pose to traditional recommendation approaches and improves the performance. This new paradigm improves serendipity and can "pop" the filter bubble, can provide recommendation to unregistered or unlogged users, does not suffer from cold start problem and respects the privacy of users.

Third Ph.D. presentation and discussion:

Ioannis STAMELAKOS – XXVIII Cycle

"Near-Threshold Computing with Performance Guarantees for Manycore Architectures"

Advisor: Prof. **Cristina Silvano**

Abstract:

Near-Threshold voltage Computing (NTC) has emerged as a promising approach to overcome the manycore power-wall, at the expense of reduced performance values and higher sensitivity to process variation. Given that several application domains operate over specific performance constraints, the performance sustainability is considered a major issue for the wide adoption of NTC.

In this thesis, we investigated how performance guarantees can be ensured when moving towards NT manycores through defferent variability-aware voltage and frequency allocation schemes. The research was extended to NT power delivery architectures, showing that when the workload characteristics of the application are analyzed and considered at runtime, big power savings can be obtained even when using existing, cost-effective power delivery techniques.

Finally, NTC runtime optimizations were considered and a lightweight runtime algorithm for balancing throughput under process and workload variability at NTC was introduced, providing significant energy savings.