COURSE INFORMATION PACKAGE

Performance Improvement in Network Environments

GENERAL IDENTIFICATION DATA

Course Title: Performance Improvements in Network Environments
Type: Optional
Degree: Master "Advanced Sciences of Modern Telecommunications"
Cycle: Postgraduate studies
Department: Departamento de Informática
Responsible Professors: Juan Manuel Orduña, José Manuel Claver

Schedule: 1st quarter.
Credits: 5 ECTS credits
Classes: 3 hours each week (Tue. & Thur, 17:00 – 18:30), 10 weeks.
Level required: Graduate
Instructors: Dr. Juan M. Orduña, juan.orduna@uv.es , +34 9635(44489)

Dr. José M. Claver, jclaver@uv.es, +34 9635(43738),

GENERAL OVERVIEW

This course intends to be an introduction to the study of high performance network environments. The main purpose of this course is to provide the students with an in-depth view of the the interconnection networks in current high-performance computer architectures and applications. In order to achieve this purpose, this course covers three particular aspects: the design of current high performance networks, the performance improvements that a correct network design can provide to network applications like Distributed Virtual Environments, and the performance improvements that a communication-aware mapping technique can provide to both off-chip and on-chip distributed computer architectures.

COMPETENCES TO BE ACQUIRED

General

The general competencies that students must acquire are:

- Be able to propose improvements in network management and distributed network scheduling.
- Be able to obtain information from scientific papers and technical reports.
- Capability of writing and presenting a scientific paper to a qualified audience.
- Be able to collaboratively work with other students.
- Capability of criticism and synthesis of the information extracted from different bibliography sources.

Specific

- Capability of characterizing the behavior of distributed systems and applications.
- Efficiently evaluating the performance of distributed systems.
- Identify the most important components and mechanisms involved in current high speed networks
- Understand the problems that appear in the design of high speed and quality of service networks.
- Identify the most important issues involved in the design of Networks-on-Chip.
- Capability of designing strategies for improving the performance of distributed systems and applications.

CONTENTS

Part 1: Design of current high performance networks

Unit 1. Network Environments

1. Network Architecture and organization
   1.1. Networks Components
   1.2. Types of Networks
2. Design issues
3. Performance Metrics

Unit 2. Transmission technologies for LAN and SAN environments

1. Communication technology characterization
   1.1. Transport
   1.2. Protocols
   1.3. Media Access
   1.4. Memory management
2. Performance measurement
   2.1. Latency and peak bandwidth
   2.2. The half power point
   2.3. Ideal latency
3. Factors affecting performance
   3.1. Memory
   3.2. Links
   3.3. Data access
   3.4. Copies

3.5. Flow management
3.6. Network processors
3.7. Packet sizes
3.8. Current nodes (Multicore)

4. Comparing different technologies.

Unit3. Design of high speed network

1. Needs for high speed interconnects
2. Technological constrains
3. Mechanisms
   3.1. Flow control
   3.2. Congestion control
4. Routers architecture
   4.1. Evolution. Generations
   4.2. Need of high speed routers. The Moore's law.
   4.3. Queuing: OQ, IQ, VOQ, CIOQ, Buffered Crossbar.
   4.4. Memory needs and organization
   4.5. Routing in IP Networks
   4.6. Scheduling algorithms
      4.6.1.Taxonomy
      4.6.2.IQ: LQ, OC, W etc
      4.6.3.VOQ: WFA, RR, SLIP, etc

Unit 4. Quality of Service (QoS)

1. Introduction
   1.1. Definitions and Parameters
   1.2. Types of QoS: Classification
2. QoS in Internet
   2.1. IntServ
   2.2. DiffServ
   2.3. MPLS
3. Router architectures
   3.1. Memory needs
   3.2. Scheduling algorithms

Part 2 : Performance improvements in computerarchitectures and applications

Unit 5. Distributed Virtual Environments (DVEs)

4. Introduction
   4.1. Architectures for DVEs
   4.2. Current challenges ofDVEs
5. Networked-Server DVEs
   5.1. Network improvements.
      5.1.1. Partitioning techniques.
      5.1.2. Traffic characterization.
      5.1.3. Synchronization techniques.
   5.2. Heuristics for network improvements.
6. Peer-to-Peer DVEs
   6.1. DVE characterization
   6.2. Network improvements
7. Crowd simulations
   7.1. Computer architectures for crowd simulations

7.2. Heuristics for network improvements

Unit 6. Communication-aware task mapping techniques
1. Performance improvements in Cluster computing.
1.1 Message-passing paradigm. State-of-the-art
1.1.1 Effects of task scheduling on system performance
1.1.2 Task mapping techniques
1.2 Communication-aware task mapping technique
1.2.1 Model of communication cost
1.2.2 Model of communication requirements
1.2.3 Heuristic techniques for task mapping
2. Performance improvements on Networks-on-Chip (NoCs)
2.1 Netwoks-on-Chip: introduction
2.1.1 Architectural features
2.1.2 Performance parameters
2.1.3 State-of-the-art: applications, topologies, routing algorithms, power saving.
2.2 Task mapping techniques for NoCs
2.2.1 Applications. Task mapping. Topological mapping.
2.2.2 Communication-aware task mapping

WORKLOAD DISTRIBUTION

Attendance to lecture classes: 9.5 hours
Attendance to seminars/discussion classes: 16 hours
Attendance to guided training (in labs): 4.5
Attendance to tutorships: 10 hours
Homeworks: 3 hours/week x 10 weeks = 30 hours
Study-Preparation of seminars/discussion classes: 20 hours
Study-Preparation of labs and final course report: 20 hours
Study of preparation for exams: 1 hours/week x 10 weeks = 10 hours
Exams: 5 hours

GRADES

Assessment of each student is based on an exam of fundamental concepts of the course and assignments of homework to be delivered every lecture session all course long. Also, a final report of a study carried out over a subject of the course that will be submitted and presented at the end of the course. All these elements are mandatory.

Final examination includes questions of all subjects of the course and will be similar to the questions included in homework assignments. The relative weight of this exam will 30% of the evaluation.

Assignments are evaluated by the instructor and discussed by students. Participation in assignments discussion is part of the assessment of the assignments,. The relative weight of this assesment will be 40% of the evaluation.

Final report includes the study and analysis of a research paper carried out by a team of two students. Each team will write a final report of 4 pages and a little presentation of 10 minutes. Both works should include the main ideas shown in the paper. The relative weight of the final report will be 30% of the evaluation.

In a more detailed way:

Homework	40%
Final exam	30%
Report	30%

COURSE READINGS AND TENTATIVE SCHEDULING

The course will be based on research papers and selected chapters from some books.

Lectures and readings

Session 1: Presentation. Unit 1: Network environments. IP routers.

- James Aweya: "IP router architectures: an overview". *Int. J. Commun. Syst.*, 14, pp. 447 – 475. 2001.

Session 2: Unit 2: Transmitting technologies. Characterization and performance measurement.

- LL. Dickman, "Beyond Hero Numbers: Factors Affecting Interconnect Performance". *PathScale Inc.* June 2005.

Session 3: Unit 2: Transmitting technologies. Factors affecting interconnect performance. Comparing different LAN and SAN technologies

- B.M Bode *et al.*, "Cluster Interconnect Overview",2004.

Session 4: Unit 3: Design of high speed networks: Needs and mechanisms

- Flow control: Chapter 1-4 from the PhD thesis of Ferninand Gransamer, entitled "Scalable Flow Control Interconnection Networks"
- Congestion control: papers of James Aweya (TCP) and Pedro García (RECN)

Session 5: Unit 3: Design of high speed networks: Architecture of Routers. Evolution and new needs.

- James Aweya: "IP router architectures: an overview". *Int. J. Commun. Syst.*, 14, pp. 447 – 475. 2001.
- C. Mikenberg, "On Packet Switch Design". PhD Thesis. University of Eindhoven. Chapter 2 (first part). 2001.
- P. Gupta, "Scheduling in input queued switches: A survey", 1997.

Session 6: Unit 3: Design of high speed networks. Routers: Architecture and scheduling algorithms.

- C. Mikenberg, "On Packet Switch Design". PhD Thesis. University of Eindhoven. Chapters 2 y 3. 2001.
- P. Gupta, "Scheduling in input queued switches: A survey", 1997.

- N. Mckewon and T. Anderson, "A Quantitative comparison of scheduling algorithms for input-queued switches", 1996.

Session 7: Unit 4: Quality of Service.

- Z. Wang. Internet QoS. Architectures and Mechanisms for Quality of Service. *Morgan Kaufmann Publishers*. 2001. Cap. 1: The Big Picture.
- Nortel, "Introduction to Quality of Service". *White Paper*. 2003.
- CISCO, "Cap. 49: Quality of Service Networking". *Internetworking Technologies Handbook*. 2004.

Session 8: Unit 4: Quality of Service. Routers: Architecture and scheduling algorithms

- J. Duato, S. Yalamanchili, B. Caminero, D. Love, F. Quiles, "MMR: A High-Performance Multimedia Router - Architecture and Design Trade-Offs". *HPCA'05*. 2005.
- A. Martínez, F. J. Alfaro, J. L. Sánchez, J. Duato, "Providing Full QoS Support in Clusters Using Only Two VCs at the Switches". *Technical Report UCLM*. 2006.
- R. Martínez, F. Alfaro, J.L. Sánchez, "Implementing the Advanced Switching Minimum Bandwidth Egress Link Scheduler". *Technical Report UCLM*, 2006.

Session 9: Unit 5 Distributed Virtual Environments. Introduction.

- J. Duato, S. Yalamanchili, L. Ni: "Interconnection Network: An engineering approach". 2007.
- S. Singhal and M. Zyda, Networked Virtual Environments. ACM Press, 1999.
- J. C. Lui and M. Chan, "An efficient partitioning algorithm for distributed virtual environment systems," IEEE Trans. Parallel and Distributed Systems, vol. 13, 2002.

Session 10: Unit 5 Networked-server DVEs.Network improvements.

- P. Morillo, J. M. Orduña, M. Fernández, and J. Duato. On the characterization of avatars in distributed virtual worlds. In EUROGRAPHICS' 2003 Workshops, pages 215–220. The Eurographics Association, 2003.
- P. Morillo, J. M. Orduña, M. Fernández, and J. Duato. On the characterization of distributed virtual environment systems. In Euro-Par' 2003 - Lecture Notes in Computer Science 2790, pages 1190–1198. ACM, Springer-Verlag, 2003.

Session 11: Unit 5 Networked-server DVEs. Heuristics for network improvements.

- P. Morillo and M. Fernandez, "A GRASP-Based Algorithm for Solving DVE Partitioning Problem," Proc. 2003 Int'l Parallel and Distributed Processing Symp. (IPDPS 2003), Apr. 2003.
- P. Morillo, M Fernandez, and J.M. Orduña, "A Comparison Study of Modern Heuristics for Solving the Partitioning Problem in Distributed Virtual Environment Systems," Proc. Int'l Conf. Computational Science and its Applications (ICCSA 2003), pp. 458- 467, May 2003.

Session 12: Unit 5 Networked-server DVEs.Partitioning techniques..

- P. Morillo, J. M. Orduña, M. Fernández, and J. Duato. Improving the performance of distributed virtual environment systems. IEEE Transactions on Parallel and Distributed Systems, 16(7):637–649, 2005.
- N. Beatrice, S. Antonio, L. Rynson, and L. Frederick. A multiserver architecture for distributed virtual walkthrough. In Proceedings of ACM VRST'02, pages 163–170, 2002.

Session 13: Unit 5 Networked-server DVEs. Quality of Service

- Z. Choukair, D. Retailleau, and M. Hellstrom, "Environment for per forming collaborative distributed virtual environments with qos," in Proceedings of the International Conference on Parallel and Distributed Systems (ICPADS'00). IEEE Computer Society, 2000, pp. 111–118.
- Y. W. Bernier, "Latency compensating methods in client/server in-game protocol design and optimization," in Proc. of 15th Games Developers Conference, 2001.
- C. Faisstnauer, D. Schmalstieg, and W. Purgathofer, "Priority scheduling for networked virtual environments," IEEE Computer Graphics and Applications, vol. 20, no. 6, pp. 66–75, 2000.
- P. Morillo, J. M. Orduña, M. Fernández, and J. Duato, "A method for providing qos in distributed virtual environments," in Proceedings of 13th Euromicro Conference on Parallel, Distributed and Network-based Processing (PDP'05), pp. 152-159. IEEE Computer Society, 2005.

Session 14: Unit 5 Networked-server DVEs. Traffic characterization. Synchronization.

- P. Morillo, J.M. Orduña, and J. Duato. A scalable synchronization technique for distributed virtual environments based on networked-server architectures. In Proceedings of the 35th IEEE International Conference on Parallel Processing (ICPP'06) Workshops, pages 74–81. IEEE Computer Society Press, 2006.
- Eric Cronin, Burton Filstrup, Anthony R. Kurc, and Sugih Jamin. An efficient syn chronization mechanism for mirrored game architectures. Kluwer Multimedia Tools and Applications, 23(1), 2004.
- L. Gautier and C. Diot. Design and evaluation of mimaze, a multi-player game on the internet. In Proceedings of IEEE Multimedia Systems Conference, page 233, 1998.

Session 15: Unit 5 Peer-to-Peer DVEs. Characterization.

- J. Duato, S. Yalamanchili, L. Ni: "Interconnection Network: An engineering approach". 2007.
- J.M. Orduña P. Morillo and M. Fernánndez. Workload characterization in multiplayer online games. Lecture Notes on Computer Science, 3980:490–499, 2006.
- S. Rueda, P. Morillo, J. M. Orduña, and J. Duato, "On the characterization of peer-to-peer distributed virtual environments," in Proceedings of the IEEE Virtual Reality 2007 (IEEE-VR07), Charlotte, USA. IEEE Computer Society Press, 2007, pp. 107–114.

Session 16: Unit 5 Peer-to-Peer DVEs. Network improvements.

- Y. Kawahara, T. Aoyama, and H. Morikawa. A peer-to-peer message exchange scheme for large scale networked virtual environments. Telecommunication Systems, 25(3):353–370, 2004.
- P. Morillo, W. Moncho, J. M. Orduña, and J. Duato. Providing full awareness to distributed virtual environments based on peer-to-peer architectures. Lecture Notes on Computer Science, 4035:336–347, 2006.
- Shun-Yun Hu, Jui-Fa Chen, and Tsu-Han Chen. Von: a scalable peer-to-peer network for virtual environments. IEEE Network, 20(4):22–31, 2006.

Session 17: Unit 5 Crowd simulations. Network improvements.

- D. Diller, W. Ferguson, W. Leung, A. Benyo, and D. Foley. Behavior modelling in comercial games. In BRIMS '04: Proceedings of the 2004 Behavior Representation in Modelling and Simulation Conference, 2004.
- S. Dobbyn, J. Hamill, K. O'Conor, and C. O'Sullivan. Geopostors: a real-time geometry/impostor crowd rendering system. ACM Trans. Graph., 24(3):933–933, 2005.
- H. Tianfield, J. Tian, and X. Yao. On the architectures of complex multi-agent systems. In Proc. of the Workshop on "Knowledge Grid and Grid Intelligence", IEEE/WIC International Conference on Web Intelligence / Intelligent Agent Technology,, pages 195–206. IEEE Press, 2003.
- M. Lozano, P. Morillo, J. M. Orduña, V. Cavero, On the design of an efficient architecture for supporting large crowds of autonomous agents, in: Proceedings of IEEE 21th. International Conference on Advanced Information Networking and Applications (AINA'07), 2007.

Session 18: Unit 6 Communication-aware task mapping techniques. Performance improvements in Cluster computing.

- H. Topcuoglu, S. Hariri, M.Y. Wu, Task scheduling algorithms for heterogenous processors, in: Proceedings of 8th IEEE Heterogeneous Computing Workshop (HCW'99), pp. 3–14.
- V. Yarmolenko, J. Duato, D.K. Panda, P. Sadayappan, Characterization and enhancement of dynamic mapping heuristics for heterogeneous systems, in: Proceedings of the 2000 ICPP Workshop on Network-Based Computing, 2000, pp. 437–444.
- Orduña, J., Silla, F., Duato, J.: On the development of a communication-aware task mapping technique. Journal of Systems Architecture 50(4) (2004) 207–220

Session 19: Unit 6. Task mapping techniques for NoCs (Networks-on-Chip).

- Ascia, G., Catania, V., Palesi, M.: Mapping cores on network–on–chip. Interna tional Journal of Computational Intelligence Research 1(1–2) (2005) 109–126
- Lei, T., Kumar, S.: A two–step genetic algorithm for mapping task graphs to a netword on chip architecture. In: Proceedings of the Euromicro Symposium on Digital System Design (DSD'03). IEEE Computer Society Press. (2003)
- Hu, J., Marculescu, R.: Energy-and-perfomance-aware mapping for regular noc architectures. IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems 24(4) (2005) 551–562
- Murali, S., Micheli, G.D.: Bandwidth-constrained mapping of cores onto noc architectures. In: Proceedings of Design Automation, and Test in Europe. IEEE Computer Society Press. (2004) 896–901

Session 20: Future research lines.

Dic 20: Final exam.