

TEACHING IN THE CLOUD MICROELECTRONICS UBIQUITOUS LAB (MULAB)

Original

TEACHING IN THE CLOUD MICROELECTRONICS UBIQUITOUS LAB (MULAB) / RUO ROCH, Massimo; Graziano, Mariagrazia. - STAMPA. - 1:(2012), pp. 131-135. (9th European Workshop on Microelectronics Education (EWME). Grenoble 9-11 May 2012).

Availability:

This version is available at: 11583/2500990 since:

Publisher:

EDA Publishing

Published

DOI:

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Teaching in the Cloud Microelectronics Ubiquitous LAB (MULAB)

M. Ruo Roch, M. Graziano
DET - Department of Electronics and Telecommunications
Politecnico di Torino
c.so Duca degli Abruzzi 24 10129 - Torino, Italy
Email: massimo.ruoroch@polito.it, mariagrazia.graziano@polito.it
Phone: +39-0110904147, +39-0110905172

Abstract—CAD laboratory students activity is mandatory for microelectronics teaching. This, applied in the deep-submicron era, creates new challenges to couple software management simplicity to user friendliness inside lab sessions, which requires the use of complex tools and concepts. In this paper, a new approach to microelectronics CAD deployment is presented, based on virtualization capabilities of new servers hardware and software technology. A test case, realized at Politecnico di Torino, degree of Electronic Engineering, is presented, with real world results on resource consumption and user satisfaction.

Keywords—cloud, cad, power, wireless, teaching, laboratory

I. INTRODUCTION

Microelectronics teaching is inherently related to the use of CAD software. This usage can be seen from different points of view, i.e., the software tool can either be the target of the course, or just a mean to improve learning of microelectronics specific topics. Different teaching tasks can so be identified.

The main one is the teaching of a design methodology. Theoretical lessons on methods to derive a hardware architecture from design specifications must be complemented by design exercises developed in the labs or as homeworks by students. Design entry requires anyway the use of CAD tools.

Another important task is Hardware Description Language learning. Even if not strictly as a programming language, HDL teaching needs extensive lab exercises, to let students understand details of hardware description through this class of languages.

Last, modern design flow in electronics design must be considered. CAD tools are of great help to learn currently used design flows, starting from higher levels of abstraction (specifications, hardware/software co-design), down to lower ones (gate and transistor level descriptions).

II. STATE OF THE ART

Laboratories involving CAD software are usually performed with two possible strategies. The most diffused one is the standard approach: several PCs are organized in a lab,




Figure 1. Standard PC based lab architecture

each one configured in such a way to allow direct software execution (Figure 1). Another possibility is the approach of *thin computing*. Each workspace is based on a so-called *thinclient*, i.e. a low-cost, reduced computing power PC. The clients are connected to a central server, on which CAD tools are executed. Only graphics rendering and user input is performed on the clients (Figure 2).

Due to hardware/software requirements of modern CAD tools, the first solution has the disadvantage to require big investments for the PCs used by the students. Moreover, the deployment of software configuration to each workspace requires, even with automatic techniques, a big effort in terms of man-power dedicated to system maintenance. Both free and commercial systems exist, suitable for this solution [1] [2] [3]. If the number of PCs is reasonably high, the problem of device failures due to aging or misuse must also be taken in account, as it can be important from an economical point of view. A significant advantage of this system architecture is the fact that students work in an environment which is familiar to their standard habits (their personal computer, was it laptop or desktop one).

On the other hand, the second approach permits to reduce costs on the client side, allowing to invest for a very robust server, which usually guarantees both better peak




Figure 2. Thinclient based lab architecture

performances, and a longer operative life. Maintenance cost, too, is reduced, because server appliances usually have greater reliability, and lower warranty services costs. The management of the software is easier, too, as just one installation is required, and deployment to clients is simple, based mainly on standard remote display protocols, like X for Unix-like appliances, or RDP for Microsoft Windows ones. Last, power consumption of the overall system is reduced, as thin client are typically machines without hard disks, or fans, and with reduced CPU requirements.

Both preceding solutions have anyway the following limitations:

- Platform uniformity requirements. Sometimes, CAD tools have different software requirements, in term of Operating System type and version. To solve this problem the first approach requires client re-installation, possibly starting from a so-called *golden image*, but the number of choices must be limited, to reduce maintenance complexity. The second solution would require the re-installation of the server, typically impractical, so that often a small number of different servers is used, as an example, one for the Linux environment, and another one for MS Windows.
- Restriction of the lab to a single location. Both PCs and thinclients must be configured in a way suitable to perform in acceptable way, specifically form the point of view of graphics performance and available network bandwidth. Moreover, cost considerations lead to the use of clients which are not portable devices. It means the lab must be attended in a properly equipped laboratory room, in which a fixed number of workspaces is installed. This lead to underutilization of the resources, for courses with a low number of students, or to the

impossibility to have all students attending the lab at the same time for crowded courses.

Current information technology advances seem to be able to support new laboratory architectures and learning models. An example of hardware lab virtualization for electronics teaching is given in [4]. The distribution of software services to different client devices, suitable to delocalize user operations, is an active field of research, too, as described in [5].

III. MICROELECTRONICS UBIQUITOUS LAB

In this work, a new solution, developed by the authors, and used at Politecnico di Torino in the Electronics Engineering degree, will be presented. The main idea is to apply the concept of virtualization and cloud computing to laboratory teaching. In fact, it is an extension of the thincomputing approach, applying the following two modifications:

- The main physical server is replaced by a set of virtual servers, running on a cloud of physical high-end machines. Each virtual server is dedicated to a specific course, i.e. it satisfies exactly the system requirements of used CAD tools, both from the point of view of OS, and hardware resources.
- Software applications are installed in a centralized location, based on a virtual file server, shared among all virtual servers. The same apply for user data.
- The thinclients are not physical devices. Instead, students laptop are used, just deploying a client application, suitable to connect to the main server. Anyway, conventional thinclients are available for students without a PC, even if in a greatly reduced number.
- The communication infrastructure between clients and servers is based on the campus wireless LAN. Of course, wired LAN or WAN access can be used, too, to exploit maximum flexibility.

The overall system architecture is shown in figure 3. This approach leads to a significant advantage from the point of view of optimization of resource usage. In fact, many virtual servers can run concurrently, allowing parallel execution of different laboratories. Moreover, computing resources not used for teaching can be allocated to research purposes, just carefully selecting virtual server priority in memory, I/O bandwidth and CPU time sharing.

Scalability is greatly improved, too. The proposed architecture can scale efficiently, both increasing memory, CPUs, and storage of a single physical server, or just replicating it. The increase in obtainable performance has no cost from the point of view of software maintenance, as no re-installation is required.

Also user experience get an improvement. The fact to use his own PCs to attend the lab has a positive impact from the point of view of student comfort. Great attention must be used, anyway, to allow easy data interchange between the server and personal laptops.

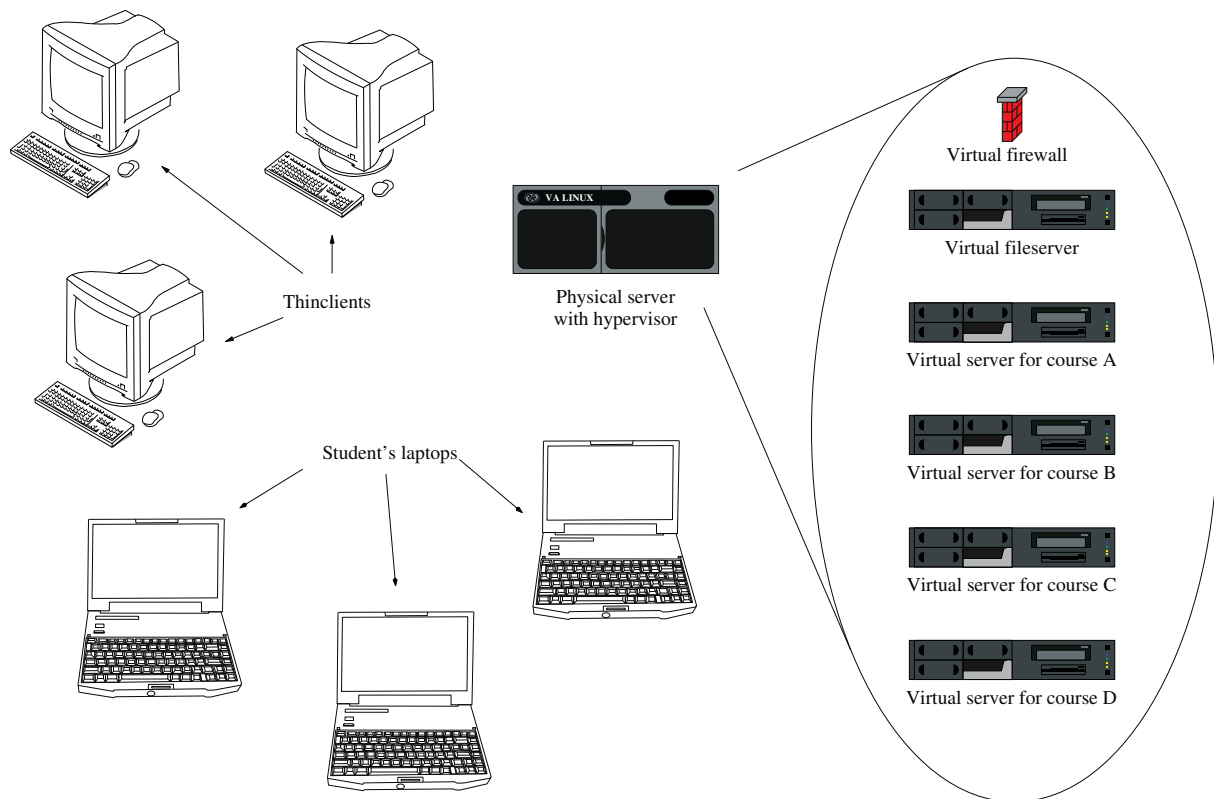


Figure 3. Virtual server based lab architecture

Last but not least, better integration between theory lessons and lab experiences can be achieved. Both can now be held in the same place, typically a course classroom, and time interleaved, too. Besides, student access to the virtual server is possible also from the entire campus network, or, if needed, directly from home, and at any time.

IV. ON-FIELD EXPERIMENTATION

To validate the proposed system architecture, an on-field test has been performed, applying this approach to the Microelectronics Systems course of the Electronics Engineering degree at Politecnico di Torino. The experiment has been performed using a scaled-down architecture, just sufficient to support one course lab at a time, but enough to have interesting data both from a technical point of view (resource usages), and from a user perspective.

This minimum system is based on an enterprise server (IBM x3850-X5) with 300 GB of RAID1 hard disk space, 16 MB of RAM, and 16 Xeon X7xxx cores. The choice of this specific system is due to the high degree of expandability of the machine, which can reach 4 TB of hard disk, 2 TB of RAM, and 64 cores. Moreover, the moderate cost of the base configuration server, joined to the presence of very high I/O bandwidth (2x10Gb/s, and 2x1Gb/s network links) drove our choice.

From the point of view of client software, the choice fell on NX, by NoMachine [6] [7] [8]. This client is a remote desktop protocol application, which has several features letting it ideal for our application:

- Low minimum bandwidth requirements. The network bandwidth used by this protocol can be tailored to the available media, and several optimizations are present, allowing to have a reasonable remote desktop even on 64 kb/s lines. In our experiment, we use a 54 Gb/s WLAN, with two or three access point per room. Besides increasing available shared bandwidth, this configuration guarantees communication fault tolerance, too.
- Availability of the client for every possible platform. It is freely available for MS Windows, Linux and MacOSX, allowing each student to access the virtual server.
- Server side application is available at no-cost. The NX protocol is supported by a free daemon, which easily run on Unix-like systems. item Communication protocol encryption. This feature is very important, as access to CAD tools can be granted to students even if they are not inside the campus, so that access and data protection is mandatory.

Students had access to the lab from the classroom, but the

virtual server was turned on 24/7, so that they can access it at anytime, when they are in the campus, or even at home. To fulfill security requirements, a firewall, running on a virtual appliance, has been introduced, to block accesses of the server to the WAN. The only possible outbound connection is to the campus webserver.




Figure 4. Virtual server CPU activity




Figure 5. Virtual server disk activity




Figure 6. Virtual server network activity

Technical feedback has been acquired through the use of appropriate monitoring tools, supplied from the chosen hypervisor (VMware ESXi 4 [9], [10]). Figures 4, 5, 6 and 7 show a sample of CPU, disk, network and memory activities during a typical laboratory session. Data are relative to a




Figure 7. Virtual server memory activity

single virtual server with 30 groups of 3 students each accessing at the same time, and executing Mentor Graphics Modelsim HDL simulator. Time range spans across a 1 hour interval. Similar data have been recorded for other CAD tools, ranging from analog simulation to standard cell place and route. The virtual server is configured with 12 GB of RAM and 4 cores.

In depth analysis of recorded data shown that the virtual server is never used near saturation and that typical resource usage is around 30-40%, with peaks around 55%. Network requirements, too, are very limited, with cumulative peak requests around 500 KByte/s.




Figure 8. Results of a student survey, composed by 10 different questions. Left bars are obtained after 3 labs. Right ones are obtained at end of course. Score can vary between 1 and 5

Next step is the analysis of system responsiveness to user interface requests, particularly when students are using CAD tools requiring intensive graphic operations, like simulation viewer, or ASIC layout place and route user interfaces. User satisfaction has been measured through a questionnaire, distributed first after three labs, and again at the end of the

course (7 labs). Answered questions were:

- 1) Do you like the fact you are not in a dedicated laboratory-room?
- 2) Do you like the fact you are using your own laptop to access the laboratory sessions?
- 3) Do you like the fact that you are accessing the facility when you want and from everywhere?
- 4) What do you think in general about the service outside the lab sessions?
- 5) What do you think in general about the service during the lab sessions?
- 6) Do you think the service performance are adequate for the requested exercises?
- 7) Do you think the real time responsiveness of the user interface is sufficient out side the lab sessions?
- 8) Do you think the real time responsiveness of the user interface is sufficient during the lab sessions?
- 9) Do you think the performance during simulation and synthesis is sufficient out side the lab sessions?
- 10) Do you think the performance during simulation and synthesis is sufficient during the lab sessions?

As shown in figure 8, students gave a quite good feedback to the methodology used to access the server, and, more important, test results have a slight improvement toward the end of the course, probably related to the solution of some technical problems (software bugs, performance optimization issues) solved on the field.

A further indication of a positive learning experience for students, can be derived by an analysis of the quality of hardware designs produced both in the labs and in the homeworks. In fact, we think this could be due to a higher availability of CAD tools to students (compared with minor results of previous years), so that they can use their spare time with greater effectiveness.

V. CONCLUSIONS AND FUTURE WORKS

This first experiment has shown good results from the point of view of overall system cost (hardware/software setup and management), service level adequacy, and students satisfaction.

For current academic year (2011/2012), the experiment has been scaled up, upgrading physical server hardware, and increasing the number of courses using this new lab architecture. In particular, now we have 96 GB of RAM and 600 GB of hard disk, thus allowing us to have 10 virtual servers running at the same time.

Monitoring techniques similar to the ones described above are currently used to gather further resource utilization and student satisfaction data, up to the end of current semester. Weak points to be optimized, yet, are mainly related to the need to exclude misuse of a so powerful computing platform, without limiting too much user accessibility.

ACKNOWLEDGMENT

The authors want to acknowledge Danilo Demarchi, Luciano Lavagno, Francesco Gregoretti and Maurizio Martina for their patience during lab experimentation's. A great thank to Paolo Motto, and all the staff of LED, for their precious work. The authors would like to thank VMware Inc., too for the grant of Academic Program Licenses, used in this work.

REFERENCES

- [1] Chao-Tung Yang, Ping-I Chen, Ya-Ling Chen, "Performance Evaluation of SLIM and DRBL Diskless PC Clusters on Fedora Core 3" Parallel and Distributed Computing, Applications and Technologies, 2005. PDCAT 2005. Sixth International Conference on 05-08 Dec. 2005, pp 479-482
- [2] <http://clonezilla.sourceforge.net>
- [3] <http://it.norton.com/ghost>
- [4] Jianchu Yao, Limberis, L., Warren, S., "Work in progress A ubiquitous laboratory model to enhance learning in electronics courses offered by two universities with dissimilar curricula", Frontiers in Education Conference (FIE), 2010 IEEE , pp.F3C-1-T1A-2, 27-30 Oct. 2010.
- [5] Surendar Chandra, "Beacon: A Peer-to-Peer System to Teach Ubiquitous Computing", SIGCSE03, February 19-23, 2003, Reno, Nevada, USA.
- [6] <http://www.nomachine.com>
- [7] <http://www.nomachine.com/documents/getting-started.php>
- [8] <http://www.nomachine.com/news-read.php?idnews=342>
- [9] <http://www.vmware.com/products/vsphere/esxi-and-esx/overview.html>
- [10] Tan Wenhui. "Data center achieves server virtualization by using VMware", (J). Ship Electronic Engineering, 2008 (6) pp. 156-159.