

# An Early Comparison of Commercial and Open-Source Cloud Platforms for Scientific Environments

Ivan Voras, Marin Orlić, Branko Mihaljević

University of Zagreb Faculty of Electrical Engineering and Computing, Zagreb, Croatia  
{ivan.voras, marin.orlic, branko.mihaljevic}@fer.hr

**Abstract.** Cloud computing promises efficient use of hardware resources through virtualization and elastic computing facilities. Various cloud computing solutions have emerged on the market from open-source communities and commercial vendors. In this paper we discuss criteria for feature comparison of private cloud platforms and compare several open-source and commercial products. We test performance of hypervisors used in these clouds with a set of benchmark suites containing tests for various aspects of the system. We discuss the results in the context of what is commonly described as a scientific workload. The described feature and performance differences can help make wiser platform choices.

**Keywords.** Cloud-computing, open-source, commercial/closed, hypervisors

## 1 Introduction

Cloud computing is a combination of technologies for resource management and provisioning with the goals of achieving cost control, elasticity, ease of use and mass deployment. In the wake of the cloud trend, numerous open-source and commercial cloud products have emerged. In this work, we focus on private-cloud infrastructure-as-a-service (IaaS) products, and present a comparison of features and performance of open-source and commercial solutions. We look at private clouds as primary platforms in scientific environments, with full local control over the platform.

Open-source solutions are often the first choice in scientific environments for their lower initial investments and support for scientific libraries and tools, as well as openness to customizations. Aggressive pricing and licensing options from vendors of commercial products leads us to expect a variety of cloud environments managing mixed virtualizations in the future. Previous performance comparisons of cloud services for scientific computing focused on open-source and modified open-source solutions [1]. In this paper, we compare performance of open-source and commercial hypervisors, the key elements and enablers of cloud platforms.

The rest of this paper is organized as follows. Sections 2 and 3 describe select open-source and commercial cloud platforms, respectively. Section 4 describes the evaluation criteria used to compare product features and presents the comparison

results. Section 5 compares the performance of open-source and commercial hypervisors, and Section 6 concludes the paper and discusses the future work in this area.

## 2 Open-Source Cloud Platforms

Open-source cloud platforms employ open-source hypervisors (KVM and Xen), but some of them also support commercial/closed hypervisors with exposed interfaces (VMware). Cloud platforms combine various tools of the underlying OS and virtualization layer with their own components in a more or less seamless cloud interface.

**Eucalyptus Community Cloud.** Eucalyptus is one of the most adopted cloud computing architectures, with open-source (Community Cloud) and commercial versions (Enterprise Edition). ECC is a flexible and highly modular system with components exposed in the form of web services interoperable with the Amazon AWS API which allows seamless integration with existing Amazon public cloud services. ECC currently supports Xen and KVM virtualizations and can be deployed on all major Linux distributions (Ubuntu, CentOS, Debian, RHEL, openSUSE, SLES etc.).

**OpenNebula.** OpenNebula is a management toolkit for private and public clouds, which orchestrates existing systems and services, relying on Linux and external products for virtualization, network, storage or security technologies. Infrastructure abstraction and modular approach support standardization and interoperability with most common virtualizations (Xen, VMware, and KVM), interfaces (Amazon AWS, VMware vCloud, and OGF OCCI), and APIs. Its features include secure management of virtual images, machines, networks and storage, authentication, multi-tenancy, quota management, and cloud-bursting with Amazon EC2.

## 3 Commercial Cloud Platforms

We have chosen the following commercial platforms for this discussion: VMware's (ESXi hypervisor), Microsoft's (Hyper-V hypervisor), and Citrix's (XenServer hypervisor, a modified version of the open-source product Xen). These hypervisors are a basis for other commercial offerings, especially those coming from OEMs<sup>1</sup>, and benchmark results also apply for those products.

**VMware vCloud.** VMware's offering is the oldest and the most feature-complete, relying on the popularity and proven track record of its virtualization products (ESXi, vSphere). On top of the virtualization platform, a comprehensive management platform was built and subsequently extended for use in cloud-like environments with the

---

<sup>1</sup> HP CloudSystem Matrix, <http://h18004.www1.hp.com/products/blades/components/matrix/main.html>

vCloud Director product. Though VMware's offer is truly complete, it is also somewhat fragmented, with services such as security management (vShield), billing (Chargeback Manager), high availability (Server Heartbeat), and others delivered and charged separately, which may significantly increase initially projected costs.

As an independent vendor not tied to specific hardware platforms or operating systems, VMware has created the most platform-agnostic product among the evaluated solutions, which has been adopted by hardware vendors and which fully supports the widest number of guest operating systems among the evaluated solutions.

**Microsoft Private Cloud.** Microsoft's presence as a server virtualization and cloud computing vendor is very recent, beginning in 2008 with the introduction of its Hyper-V hypervisor product. The features and abilities of Hyper-V are noticeably sub-par when compared to others we have evaluated, lacking important features such as live migration, support for guest OS booting from SCSI drives, memory over-commit, support for more than 4 virtual CPUs, and official support for popular non-Windows operating systems. On the other hand, Hyper-V is tightly integrated with other Microsoft's products and delivers an environment which is easy to use and manage, especially when virtualizing Windows workloads. The "Private cloud" solution integrates several products which automate resource management (most notably, the System Center Virtual Machine Manager).

**Citrix CloudStack.** Citrix's cloud solution is centered on its XenServer product, and supported by the infrastructure and the community of the CloudStack open source project, whose parent company (Cloud.com) it recently acquired. The relationship between Citrix and CloudStack seems to be somewhat complex, with Citrix supporting and endorsing it for its cloud computing efforts, as well as selling "Enterprise" and "Service Provider" editions. Of all the evaluated solutions, Citrix's is perhaps the most on the border between being a truly open and a commercial/closed solution.

Along with VMware's products, XenServer has the proven track record and presence in the business sector, but its openness and excellent Linux support has also made it attractive in the scientific and academic environments.

## 4 Evaluation Criteria

Evaluating complex products such as cloud platforms requires a common set of criteria for comparison [2], [3]. The primary goal for our criteria was to establish a comparison baseline for open-source IaaS products, which we later extended to include commercial/closed products. Some criteria, especially high-level ones, can be applied to other cloud models (PaaS and SaaS). The criteria are divided over six main groups: storage, virtualization, network, management, security and vendor support. The criteria admittedly favor platforms adapted to Linux workloads and hosts, and functionalities over usability and integration, which can be corrected using weights.

We have graded individual criteria in the range of 0 to 3, where 0 designates no support and 3 designates full support. In cases where support for a feature relies on the underlying operating system drivers or tools, we assigned the grade of 2 if the operating system support is mature. We describe each group in short, as their full description is provided in our earlier work and is out of scope of this paper.

**Storage criteria.** Management and functionalities of virtual machine storage are critical to achieve flexible and scalable implementations. Storage-related criteria focus on main technologies which implement cloud storage: direct-attached storage (DAS), storage area networks (SAN), and network-attached storage (NAS), as well as backup features and technologies. Main technology groups are further divided into specific technologies, e.g. file system and replication support.

**Virtualization Criteria.** Virtualization technologies are the enablers of cloud computing. Compared to commercial products, open-source platforms support a wider range of virtualization technologies, often three or more open-source and commercial hypervisors. Virtualization criteria focus features such as VM type and technology, quotas, migration and cloning, resource prioritization, hot configuration and provisioning. Platforms supporting a wide range of hypervisors fare better than platforms that focus on a single hypervisor, regardless of the level of supported functionalities, which we tried to amend by grading the level of support.

**Network Criteria.** Network connectivity binds the cloud components together, and connects it with its administrators and end-users. Network support is covered with criteria describing support for VLAN services such as tagging (IEEE 802.1q), network management and isolation, firewall support via network filtering, Ethernet quality of service (IEEE 802.1p), and integration capabilities such as IPv6 and virtual private networks for access and management.

**Management Criteria.** Comprehensive and usable management is directly related to the way a virtual machine cloud is handled. Criteria describe management facilities of a platform and focus on host-guest OS integration, individual and mass management of both hosts and guests, exposed and consumed APIs, data collection for billing and reporting, automatic recovery, high availability and alerting features.

**Security Criteria.** Cloud installations can pose a significant security risk. We describe security-related capabilities of a platform such as data encryption, directory services integration, authorization levels and auditing events for specific resources (e.g. VM server or storage servers). Security report types available to cloud administrators, third-party product evaluations and certifications of compliance with security standards, and secure management access are also looked into.

**Community and Vendor Support Criteria.** Product support is often divided between a user community and a vendor. Community-related criteria deal with freely available support channels of the user or developer communities gathered around the product, and their quality. Vendor-related criteria concern with direct vendor support channels and the possibility of a SLA contract. We also cover methods of customer relations such as public issue tracking, proactive updates and CRM-like approaches, completeness of open-source or free product versions and related documentation, track record, future viability, and the possibility of third-party auditing.

#### 4.1 Feature Comparison

We have compared the features of Eucalyptus Community Cloud 2.0.3 (ECC), OpenNebula 3.2 (OpenNebula), VMware vCloud 5 (vCloud), Microsoft Private Cloud 2012 (MSPC), and Citrix CloudStack 2.2 (CloudStack). Table 1 presents a sum of grades by criteria group, normalized to the maximum possible score per group. The full evaluation table contains 97 grades per product, 485 in total.

**Table 1.** Feature comparison with evaluation criteria

Criteria group	ECC	OpenNebula	vCloud	MSPC	CloudStack
Storage	22 (49%)	24 (53%)	19 (42%)	21 (47%)	26 (58%)
Virtualization	18 (27%)	30 (45%)	54 (82%)	21 (32%)	51 (77%)
Management	33 (48%)	44 (64%)	58 (84%)	37 (54%)	45 (65%)
Network	11 (41%)	6 (22%)	14 (52%)	10 (37%)	14 (52%)
Security	13 (39%)	16 (48%)	26 (79%)	14 (42%)	14 (42%)
Vendor Support	30 (59%)	37 (73%)	39 (76%)	23 (45%)	37 (73%)
<b>Total score</b>	<b>127 (44%)</b>	<b>157 (54%)</b>	<b>210 (72%)</b>	<b>126 (43%)</b>	<b>187 (64%)</b>

As expected, vCloud is the most feature-complete solution, with CloudStack and OpenNebula following suit. CloudStack has a capable hypervisor, while Microsoft has the edge with its Windows ecosystem hardware support, management and integration, as well as vendor support, but most of the advanced features it offers apply only for Windows workloads. Open-source products exhibit a lower level of component integration, and excel in some areas while lacking in other.

## 5 Hypervisor Performance Comparison of Commercial/Closed and Open-Source Cloud Solutions

While virtualization is not a necessary part of cloud computing solutions, it is almost universally present for the convenience it brings to the implementation of elasticity and automation [4]. Virtualization is commonly achieved by emulating complete virtual machines under the control of hypervisor software (as opposed to more light-weight virtualization environments [5]), making such software essential for stability, security and performance of the whole cloud infrastructure. The scientific com-

puting environment is somewhat special as it commonly uses private or semi-private infrastructure (i.e. where both the users and the infrastructure are known and controlled by a single entity), which makes both security and software stability less risky. System performance of virtualized environments is of importance to many potential users, which drove us to investigate this aspect of cloud computing. We have set our goal to evaluate common computing workloads and provide a comprehensive answer on the influence of modern hypervisor software on the overall system performance.

Our evaluation was centered on four of the most common virtualization products: VMware ESXi 5.0, Citrix XenServer 6.0, Microsoft Hyper-V 2008 R2 SP1, and KVM from Linux 2.6.32. The baseline was an un-virtualized system on an identical hardware platform (2x Xeon 5460 CPUs, 16 GB of memory, 2x 15 kRPM SAS disk drives). We have used the CentOS 6.0 operating system for the virtual machines, with the same virtual machine disk image used on all hypervisors. The latter was necessary to reduce the number of variables and outside influences on the measured system, but has lead us to conclude that, sadly, there is no standard for migrating virtual machines between virtualization environments adopted by all the major vendors and projects. We hoped that the Open Virtualization Format [6] ratified as an ANSI standard would have been widely adopted, but its support is notably absent from Microsoft's products.

## 5.1 Notes on the Hypervisors

Of the evaluated hypervisors (summarized in Table 2), ESXi, XenServer and Hyper-V can be described as “Type 1” hypervisor [7], running on “bare-metal” without a supporting operating system, while KVM is decidedly “Type 2”, making full use of regular Linux kernel features and subsystems in its operation. The most important assumption which can be derived from this distinction is that the Type 1 hypervisors are more light-weight and introduce less latency. One of our goals was to verify the validity of this assumption for the chosen hypervisors and the guest operating system.

Another goal was to check the impact of paravirtualization support for certain hypervisors, most notably Xen (used in XenServer) and KVM. Xen's paravirtualization capabilities are much more extensive than KVM's, which led us to expect significant performance improvements. Neither VMware nor KVM provide paravirtualization at the level of virtual memory management or other operating system components which would influence our benchmark.

**Table 2.** Hypervisors tested in this work

Hypervisor	License	Type	Cloud product
VMware ESXi 5.0	Commercial	Type 1	vCloud
Citrix XenServer 6.0	Commercial/GPL	Type 1	CloudStack
Microsoft Hyper-V 2008 R2 SP1	Commercial	Type 1	Microsoft Private Cloud
CentOS 6.0 KVM (Linux 2.6.32)	GPL	Type 2	OpenNebula, Eucalyptus...

According to industry surveys [8], [9], the three commercial virtualization solutions dominate the market, with a cumulative share upwards of 80% in product deployments. Of these, VMware's products are more popular by a significant margin.

Scientific and academic environments are traditionally very strongly tied to open source solutions for reasons of cost-efficiency and historically better support for specific workloads. We have observed ever stronger industry lobbying emphasizing “enterprise-grade” features and more refined management tools, which could mean a wider adoption of commercial hypervisors in scientific computing environments.

## 5.2 Simulating Scientific Workloads

There is of course no single definition of a “scientific workload”, but there are regular attempts to characterize typical workloads of certain types of scientific environments [10-13]. We decided on an approach of describing typical performance by using a number of benchmark suites, with the intent that the overall results show the relative differences between the tested systems, with respect to CPU, memory, and OS interaction performance. Table 3 summarizes the selected benchmark suites.

**Table 3.** Benchmark suites used in this work

<b>Benchmark</b>	<b>Platform/environment</b>	<b>Threading</b>	<b>Result interpretation</b>
DaCapo 9.12	Java (OpenJDK 6 b17)	Mixed	Time (lower is better)
Pysysbench 0.9	Python/C (Python 2.6)	Multi-process	Throughput (higher is better)
Bullet Cache 1.0.1	C, C++ (gcc 4.4.6)	Multi-threaded	Transactions per second (higher is better)
SciMark2/C	C (gcc 4.4.6)	Single-threaded	Throughput (higher is better)

The DaCapo benchmark suite [11] evaluates real-world performance by running a set of open source applications of a varied behavior. From this set of benchmarks we have used the following: avrora (AVR microcontroller simulation, single-threaded), batik (Scalable Vector Graphics rendering, single-threaded), eclipse (Eclipse development environment, multi-threaded), h2 (in-memory database, multi-threaded), jython (Python interpreter in Java, multi-threaded), luindex (Lucene full-text searching library, single-threaded), lusearch (Lucene, multi-threaded), pmd (source code analyzer, multi-threaded), sunflow (ray-tracing graphics, multi-threaded), and xalan (XML transformer, multi-threaded). The DaCapo results are reported as average execution time for each benchmark, in milliseconds.

The Pysysbench benchmark measures basic throughput of common low-level CPU-intensive tasks implemented in C libraries called from Python: Hash-SHA256, Hash-SHA512, Zlib-Compress, Zlib-Decompress, Socket-Syscalls (a loop executing the socket() and close() system calls), Socket-OnePipe (a loop transferring 1 byte of data over a pair of sockets). All tasks are executed in parallel processes and the results indicate throughput (amount of data processed or the number of operations per second), as a composite dimensionless score.

The Bullet Cache benchmark runs the multi-threaded Bullet Cache<sup>2</sup> [12] memory database with 5 threads on the server and on the client side, with a varied number of client connections, and reports the number of transactions per second, in thousands.

The SciMark2/C benchmark is the C implementation of the SciMark2 benchmark, implementing a set of scientific numerical algorithms: FFT, SOR, Monte Carlo integration, sparse matrix multiplication, and dense LU matrix factorization. The benchmark was used in the “large” configuration and the reported results present an approximate MFLOPS performance.

All benchmarks were installed and ran from a single virtual machine image, converted as necessary for various hypervisors. The hypervisors themselves were installed on identical hardware systems equipped with two Xeon 5460 CPUs and 16 GB of memory. The virtual machines were configured with 8 virtual CPUs and 8 GB of memory, except in the case of Hyper-V where, due to its limitations, the virtual machine was configured with 4 virtual CPUs.

### 5.3 Benchmark Results

The chosen benchmark suites produce detailed output for each individual benchmark, which can be used to interpret specific aspects of the tested systems (discussed in following sections). The benchmark results are the averages of five consecutive runs, and the overall results are presented summarized per benchmark and per hypervisor in Table 4 and Figure 1.

**Table 4.** Hypervisor benchmark results (compared to non-virtualized baseline)

	<b>ESXi</b>	<b>XenServer</b>	<b>Hyper-V</b>	<b>KVM</b>	<b>Non-virt.</b>
DaCapo	5530 (+25%)	6010 (+37%)	6240 (+42%)	6020 (+37%)	4390
Pysysbench	1013 (-5%)	1043 (-2%)	532 (-50%)	1020 (-4%)	1068
BulletCache	217 (-46%)	218 (-46%)	165 (-59%)	184 (-54%)	401
SciMark2/C	522 (-2%)	520 (-3%)	507 (-5%)	511 (-4%)	533

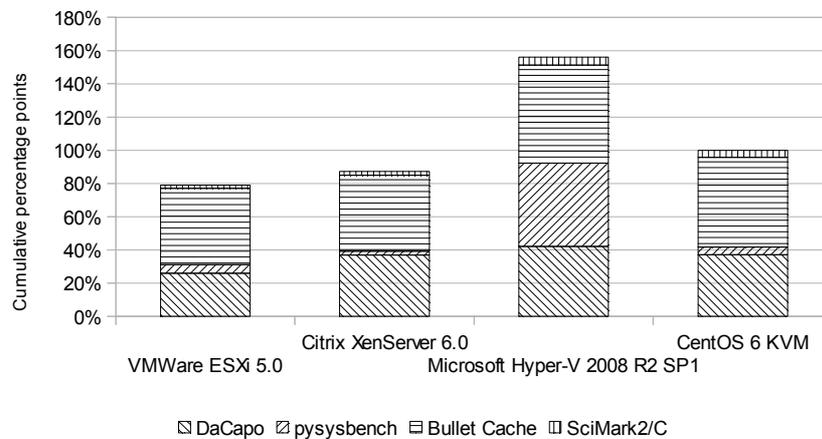
The results show that the performance of current generation of hypervisors is reasonable and there are, at first sight, no significant outstanding problems. The one exception is Microsoft’s Hyper-V whose results are skewed in aggressively multi-threaded benchmarks (Pysysmark and Bullet Cache) because it does not support more than 4 virtual CPUs. We draw the following conclusions from the results:

1. Performance differences among the latest versions of the hypervisors, tested under the specified operating system(s) and hardware, are mostly within 10%. The performance difference to the non-virtualized case, however, can be larger than 50% (in the DaCapo and Bullet Cache benchmarks);
2. Single-threaded, purely computational workloads (SciMark2/C) exhibit the smallest performance impact, between 2% (ESXi and XenServer) and 5% (Hyper-V);
3. The largest differences are in extensively multithreaded workloads (Bullet Cache);

<sup>2</sup> Bullet Cache / MultiDomain Cache Daemon, <http://mdcached.sourceforge.net/>

4. From the detailed results we can also conclude that simple system calls do not exhibit a significant performance impact. Combined with the previous point it is likely that execution context switching comes with huge penalties in virtualized environments. However, more testing is required on newer equipment (CPUs implementing Nested Page Tables) to verify that the impact is still significant.

The KVM hypervisor exhibited good overall performance, but it was never at the top. The smallest virtualization impact (4%) was measured for SciMark2 and Pysysbench but this result is still far from the best (2%, measured for the XenServer).



**Fig. 1.** Relative performance impact of virtualized environments, compared to non-virtual environments, by cumulative percentage points

The absolute performance impact of KVM compared to the non-virtualized case is large (figure 1). The relative impact when compared to the leading hypervisor (ESXi) is much smaller, 15% in the worst case (Bullet Cache) but only 0.7% in the best case (Pysysbench). The difference for the single-threaded compute-intensive workload (SciMark2) is 2% and nearly 9% for the mixed application workload (DaCapo). Whether this is a problem in practice depends on a specific deployment, but we are confident that the performance is suitable for most uses. XenServer and Xen hypervisors share the same code base and we expect them to have identical performance.

## 6 Conclusion and Future Work

We have given a brief overview of several open-source and commercial/closed private-cloud platforms. Present popularity of cloud computing paradigm and headroom for further development guarantee continuous growth and improvement of cloud platforms. Using a set of criteria developed for evaluating open-source platforms, we have evaluated the features of leading commercial/closed solutions, and confirmed that open-source and commercial solutions have comparable features.

We also compared performance of hypervisors running at the core of evaluated cloud products. In an attempt to simulate scientific workloads, we used a set of benchmark suites testing specific system aspects – CPU, memory, and operating system interaction performance. Benchmark results led us to conclude that modern hypervisors performance is close to non-virtualized case with single-threaded workloads. Large performance differences were observed with multi-threaded workloads, both between the hypervisors and compared to non-virtualized hardware. The tests were made on an older generation of hardware, and we plan to continue the tests once we have state-of-the-art equipment available.

## References

1. Iosup, A., Ostermann, S., Yigitbasi, N. et.al.: Performance Analysis of Cloud Computing Services for Many-Tasks Scientific Computing. *IEEE Transactions on Parallel and Distributed Systems*. 22, 931-945 (2011).
2. Voras, I., Mihaljević, B., Orlić, M. et.al.: Evaluating Open-Source Cloud Computing Solutions. *Proc. of the 34th International Convention MIPRO 2011*. 209-214, Zagreb (2011).
3. Voras, I., Mihaljević, B., Orlić, M.: Criteria for Evaluation of Open Source Cloud Computing Solutions. *Proc. of the ITI 2011 33rd International Conference on Information Technology Interfaces (ITI)*. pp. 137-142., Dubrovnik, Croatia (2011).
4. Armbrust, B., Griffith, R., Joseph, A.D. et.al.: A view of cloud computing. *Communications of the ACM*. 53, 50-58 (2010).
5. Kamp, P.-H., Watson, R.N.M.: Jails: Confining the Omnipotent Root. *Proc. of the 2nd International System Administration and Networking Conference SANE 2000* (2000).
6. VMware Inc.: The Open Virtual Machine Format Whitepaper for OVF Specification, version 0.9, [http://www.vmware.com/pdf/ovf\\_whitepaper\\_specification.pdf](http://www.vmware.com/pdf/ovf_whitepaper_specification.pdf).
7. Goldberg, R.P.: *Architectural Principles for Virtual Computer Systems*, doctoral dissertation, Harvard University (1972).
8. Bigelow, S.J.: State of virtualization and cloud computing: 2011, <http://searchdatacenter.techtarget.com/feature/State-of-virtualization-and-cloud-computing-2011>.
9. Citrix Systems, Inc.: Cloud.com 2011 Cloud Computing Outlook Survey Results, <http://www.cloud.com/cloud-computing-outlook/survey.pdf>.
10. Bradley, D., Cybenko II, G., Larson, J. et.al.: Supercomputer workload decomposition and analysis. *Proc. of the 5th international conference on Supercomputing*. pp. 458-467. ACM (1991).
11. Cirne, W., Berman, F.: A comprehensive model of the super-computer workload. *Proc. of the Fourth Annual IEEE International Workshop on Workload Characterization WWC4 Cat No01EX538*. 140-148 (2001).
12. Numrich, R.: A note on scaling the Linpack benchmark. *Journal of Parallel and Distributed Computing*. 67, 491-498 (2007).
13. Gray, L., Kumar, A., Li, H.: Workload Characterization of the SPECpower\_ssj2008 Benchmark. In: Kounev, S., Gorton, I., and Sachs, K. (eds.) *Performance Evaluation Metrics Models and Benchmarks*. pp. 262-282. Springer Berlin Heidelberg (2008).
14. Blackburn, S.M., Guyer, S.Z., Hirzel, M. et.al.: The DaCapo benchmarks. In *Object Oriented Programming Systems Languages and Applications OOPSLA*. 41, 169 (2006).
15. Voras, I.: Cache server for distributed applications adapted to multicore systems, doctoral dissertation, Faculty of Electrical Engineering and Computing, University of Zagreb (2011).