

On the virtualization overhead in consolidated servers for different hypervisors

Prof. Dr. Carlos Juiz and MSc. Belen Bermejo
{cjuiz, belen.bermejo}@uib.es

Computer Science Department, University of the Balearic Islands. Crta. Valldemossa km 7.5, Palma, E07122, Spain.

Introduction

- The virtualization technology has an inherent performance overhead (extra workload to manage virtualization).
- The server consolidation allows to manage datacenters in a more flexible way.
- The virtualization overhead depends on the hypervisor implementation (type-I, type-II and container-based) and the number of consolidated virtual machines.
- There are two sources of overhead [3]:
 - Due to hypervisor (OV_v)
 - Due to the co-allocation of virtual machines (OV_c)
- It is crucial to minimize the performance degradation due to virtualization overhead.

Objective

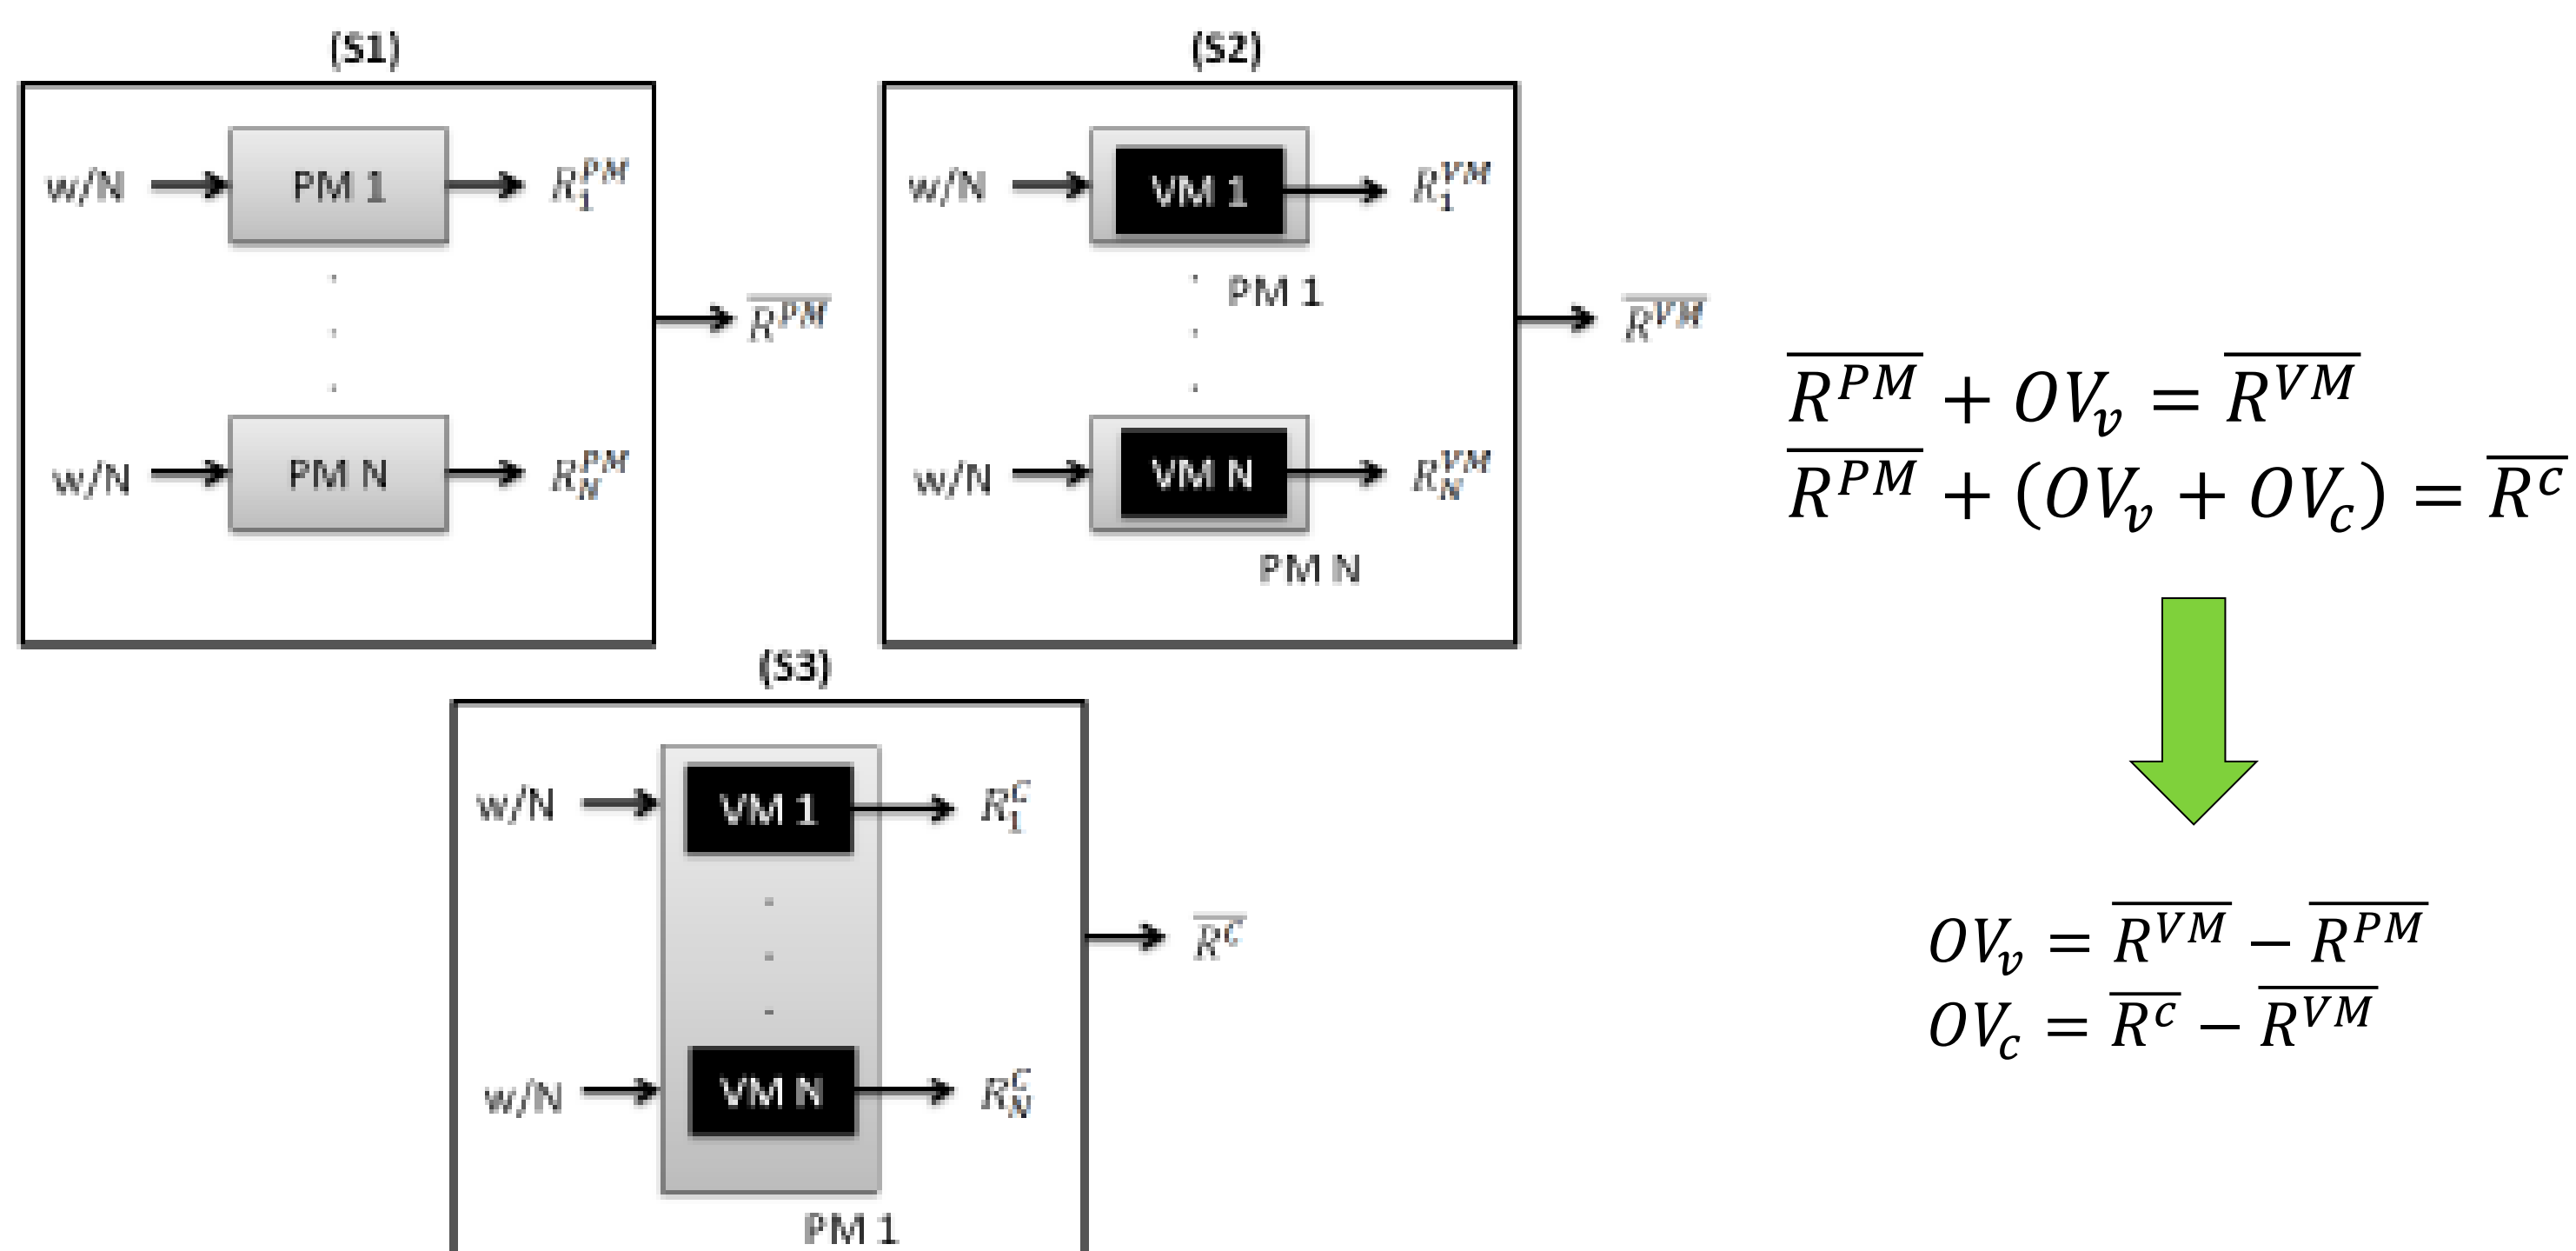
To determine how the virtualization overheads are affected by the hypervisor implementation.

Methodology

Two overhead sources:

- OV_v : overhead due to hypervisor.
- OV_c : overhead due to consolidation management

To determine OV_v and $OV_c \rightarrow$ Comparison between scenarios

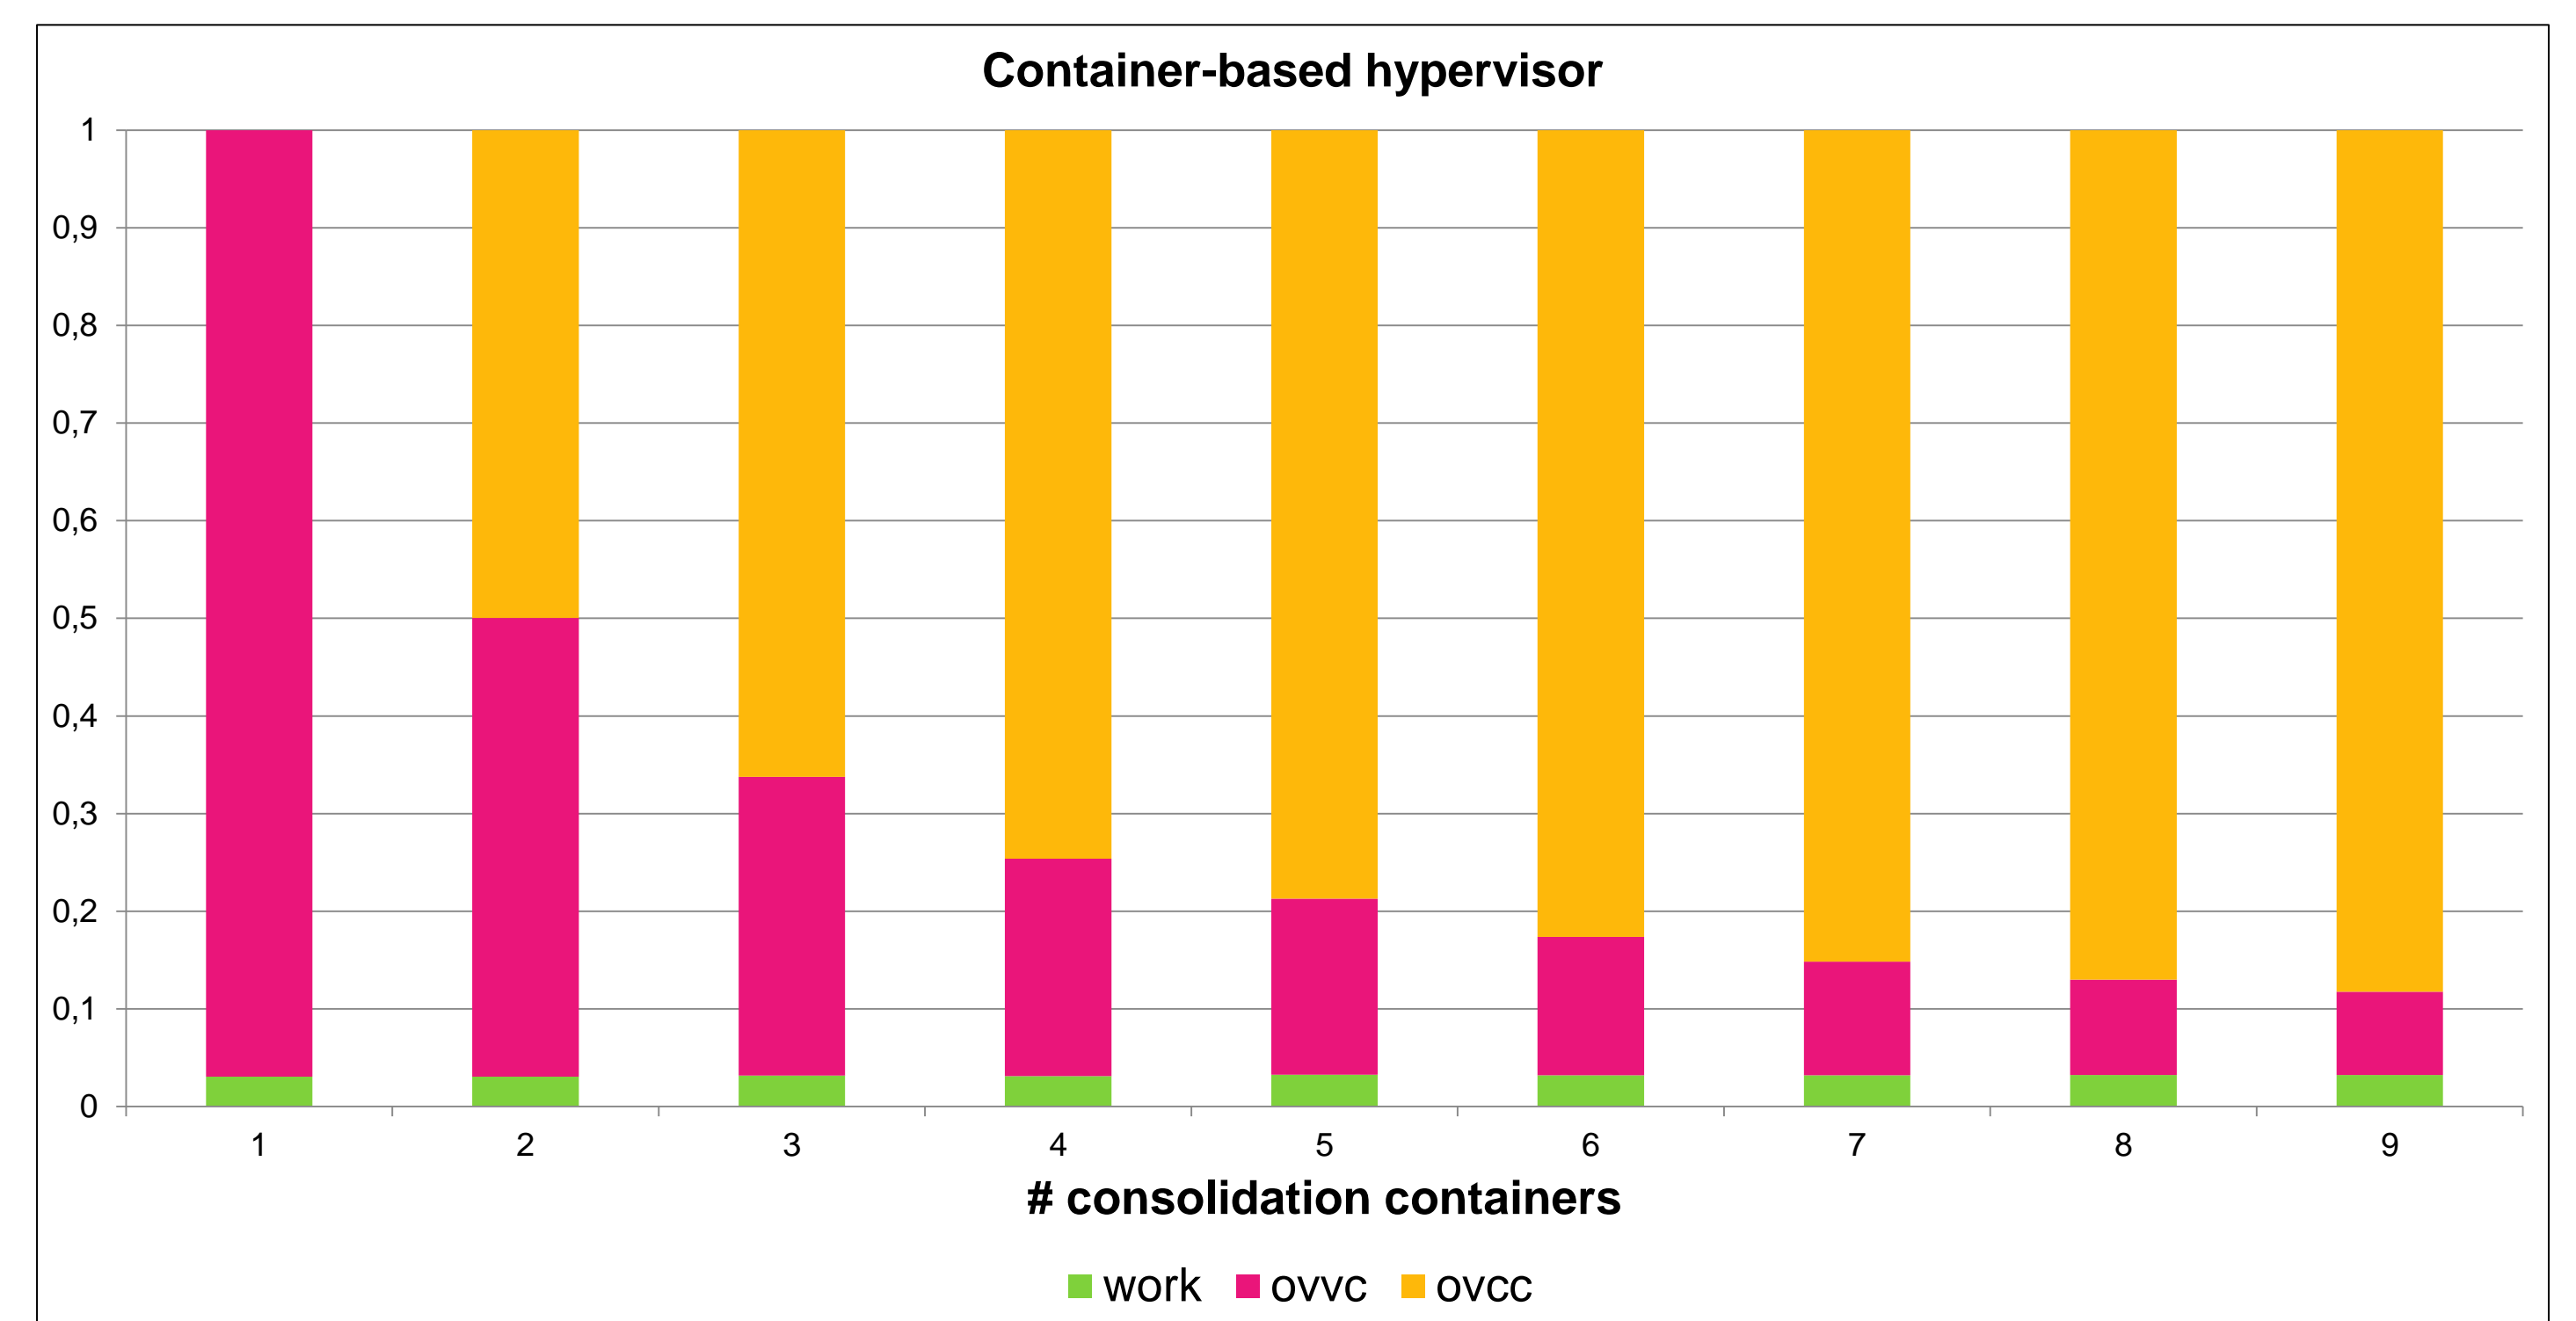
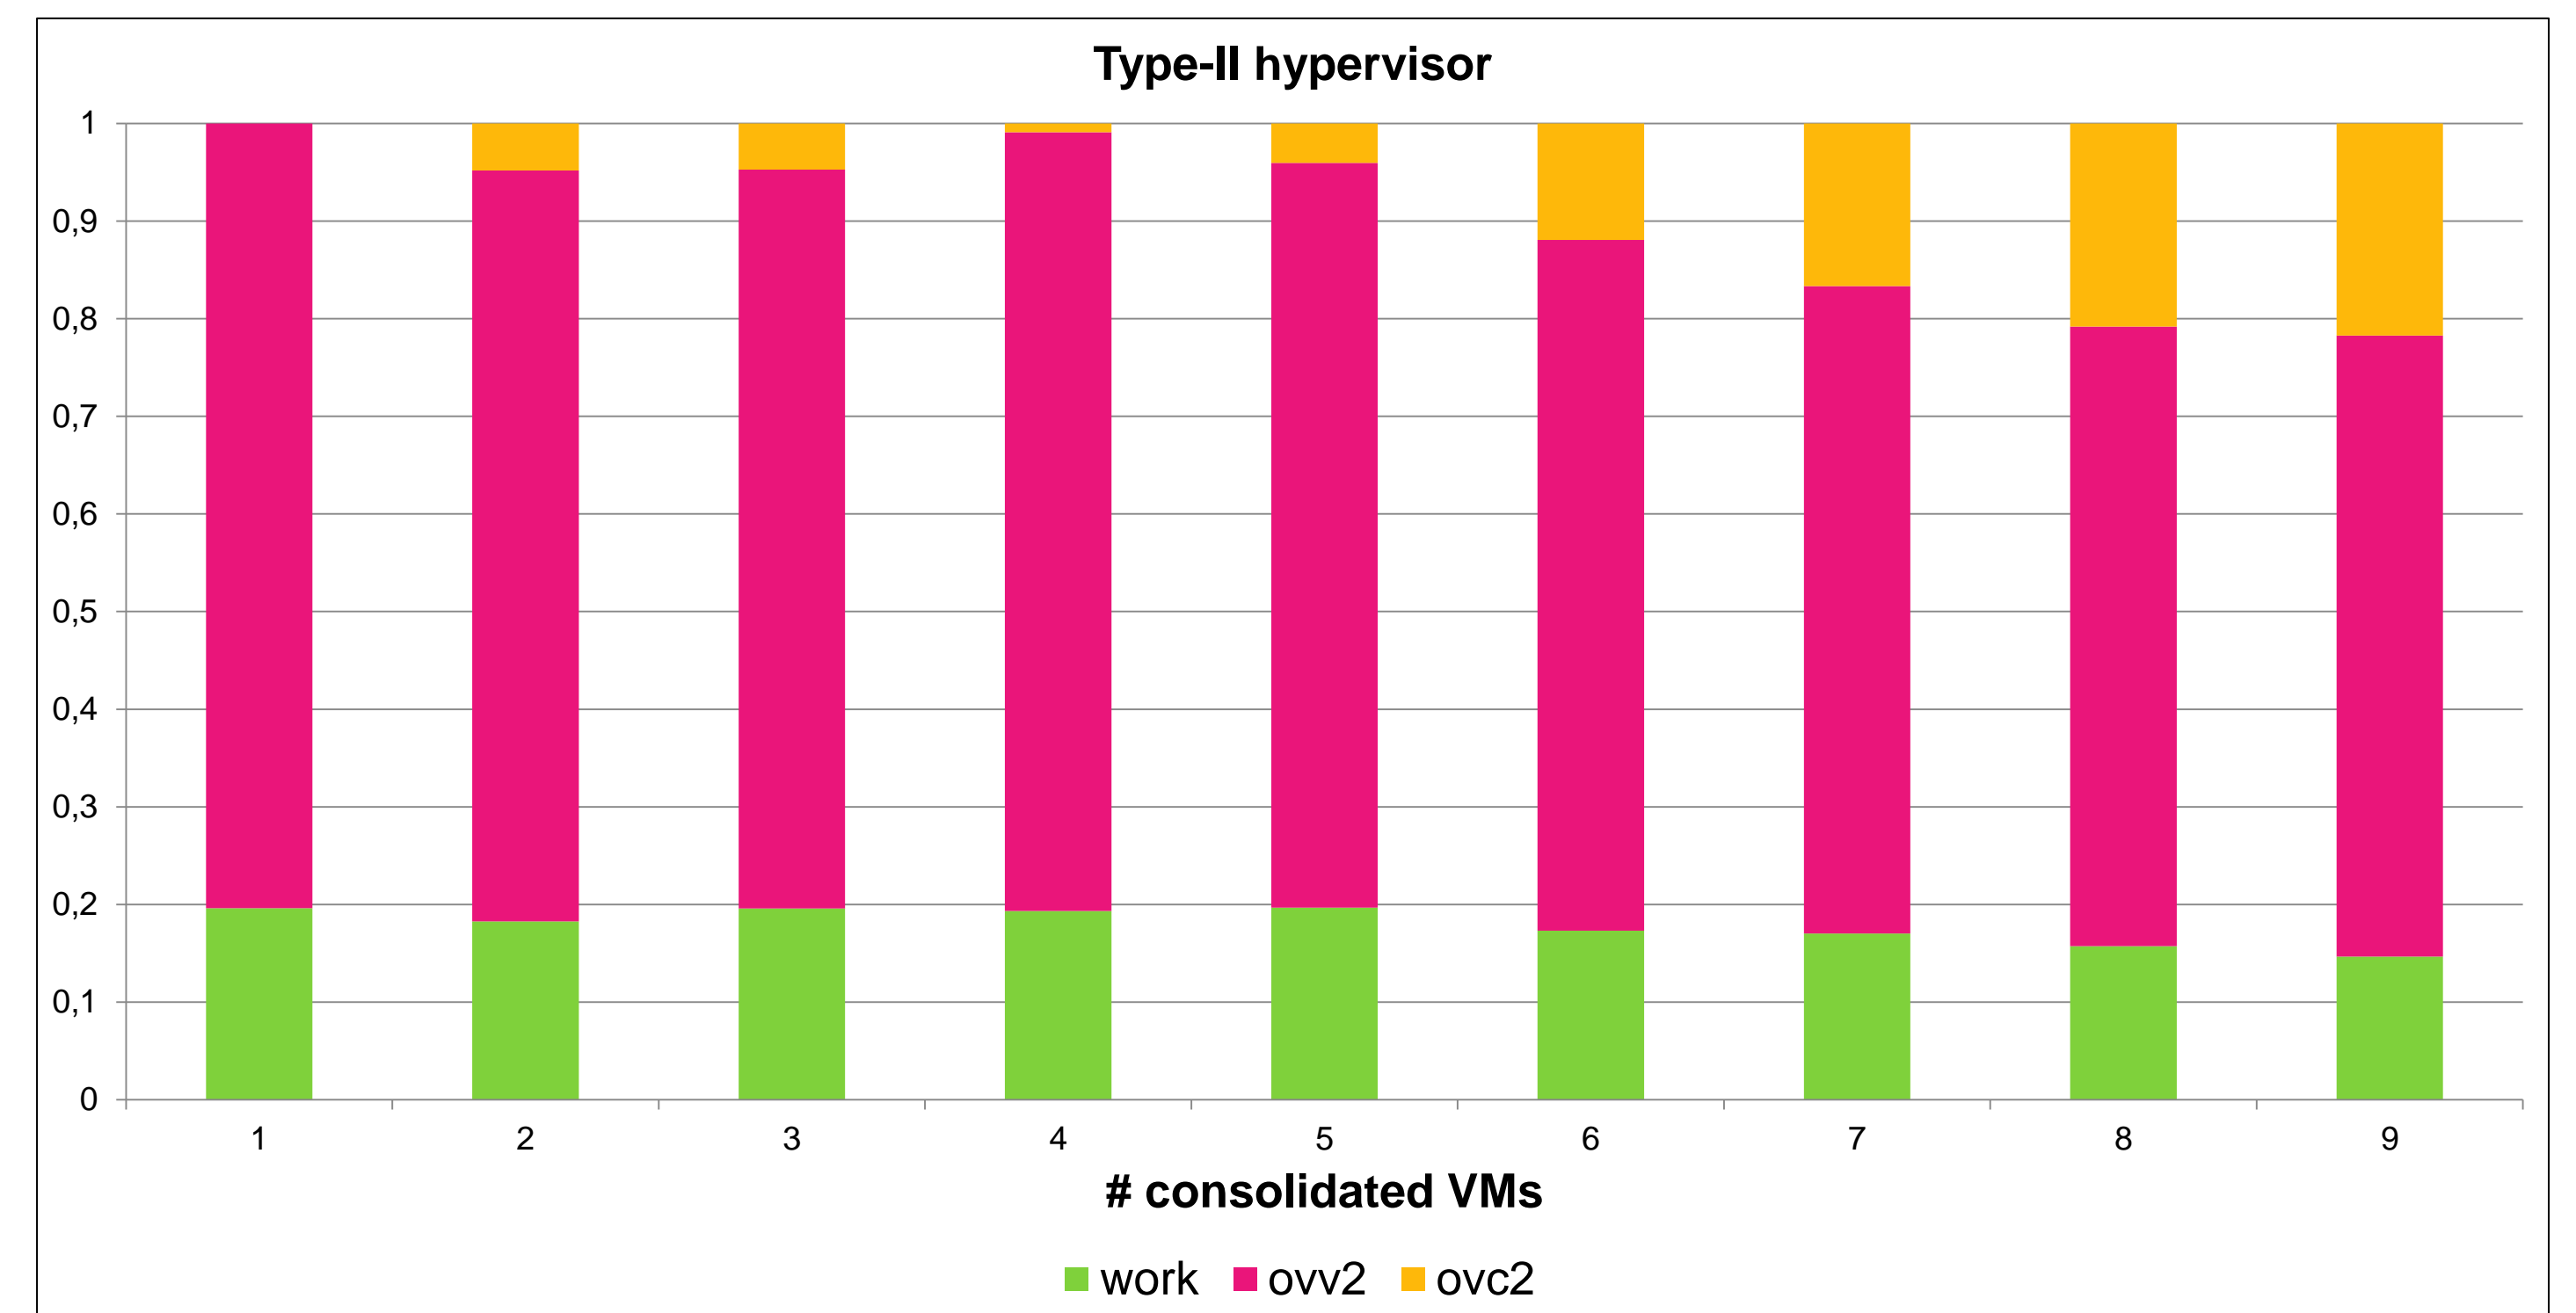



S1: physical machines. Black box composed by N servers which executes a whole workload in a balanced manner.
S2: virtual machines. Black box composed by N servers, each one hosting a single virtual instance, which executes the N workload fraction.
S3: consolidated physical machine. Black box composed by a single server, which has N consolidated virtual instances. Each VM executes the N workload fraction.

SUT: rx600s5-1 (HPI)

Workload: Sysbench CPU under CPU saturation conditions

Results



Conclusions

- We identified two sources of overhead: from the hypervisor (OV_v) and the consolidation management (OV_c).
- We determined OV_v and OV_c through scenarios comparison (S1, S2 and S3).
- The portion of useful work time is higher when consolidating with type-I hypervisor.
 - It spends less time managing the virtual instances.
- OV_v value is lesser when we consolidate thought type-I hypervisor (it remains constant).
 - This is due to the proximity with physical hardware.
- OV_c value is lesser when we consolidate thought type-II.
 - Type-II virtual machines are considered as a system process.
- The CPU-intensive workload is not suitable for containers consolidation.
- Type-I hypervisor allows consolidating more virtual instances than type-II and container-based hypervisors.
- OV_v and OV_c depend on hypervisor type, the number of consolidated instances and the workload nature.
- For CPU-intensive workload in the HPI server, the best option to consolidate is type-I hypervisor.

References

- B. Bermejo, C. Juiz, and C. Guerrero. "Virtualization and consolidation: a systematic review of the past 10 years of research on energy and performance". In: The Journal of Supercomputing (2018). issn: 1573-0484. doi: 10.1007/s 11227-018-2613-1.
- B. Bermejo, C. Juiz, and C Guerrero. "On the Linearity of Performance and Energy at Virtual Machine Consolidation: the CiS2 Index for CPU Workload in Server Saturation". In: Proceedings of the IEEE 20th International Conference on High Performance Computing and Communications. Exeter, UK, 2018, pages 928–933.
- B. Bermejo, C. Juiz, and N. Thomas, "On the virtualization overhead and energy consumption in consolidated servers." in UK- Performance Engineering Workshop (UKPEW), 2018.