

Investigation of Resource Reallocation Capabilities of KVM and OpenStack

Bachelor Thesis Presentation
26.08.2015

Andreas Gruhler <andreas.gruhler@uzh.ch>
Supervisor: Patrick Gwydion Poullie, Dr. Thomas Bocek

University of Zurich
Department of Informatics (IFI)
Communication Systems Group, Prof. Dr. Burkhard Stiller

Agenda

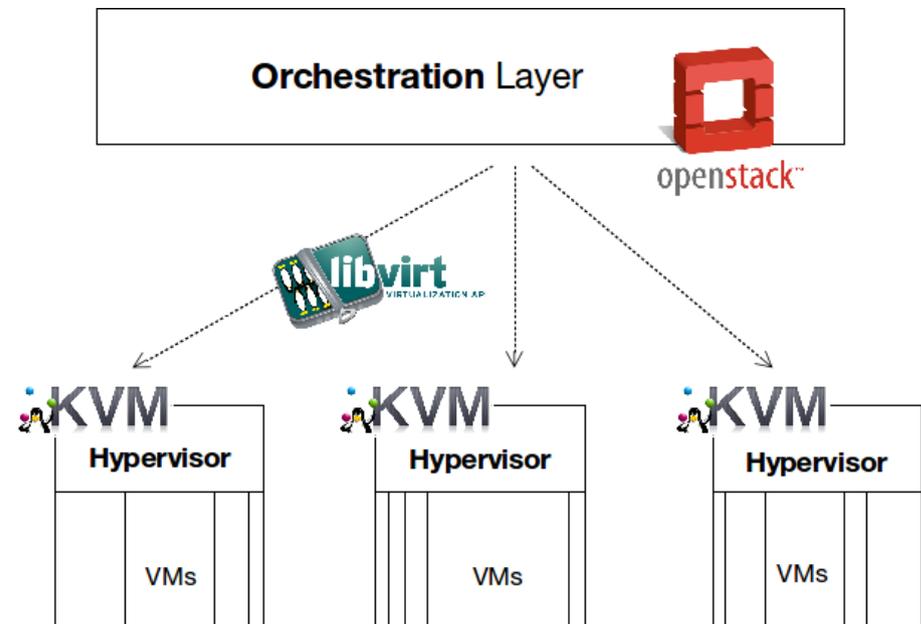
1. **Research Questions and Introduction**
2. libvirt Hard and Soft Limits
3. VM Resizing and Resource Reallocation
4. OpenStack Extension Design and Implementation
5. OpenStack Extension Evaluation
6. Conclusion
7. Future Work and Outlook

1. Research Questions

- How are physical resources (PR) reallocated in cloud computing?
- Which resources can be reallocated in KVM/OpenStack?
- Which methods are used in KVM/OpenStack to reallocate resources?

1. Introduction to Cloud Layers

- **OpenStack:**
scheduling and deployment of VMs
- **libvirt API:**
uniform access to resource limits
- Kernel based virtual machine (**KVM**):
provisions the PRs to the VMs
- Each VM has a hardware template referred to as its flavor.



Agenda

1. Research Questions and Introduction
- 2. libvirt Hard and Soft Limits**
3. VM Resizing and Resource Reallocation
4. OpenStack Extension Design and Implementation
5. OpenStack Extension Evaluation
6. Conclusion
7. Future Work and Outlook

2. libvirt Hard and Soft Limits

- Hard limit:
 - always active
 - upper boundary of the soft limit
- Soft limit:
 - only enforced if resource is scarce
 - thesis focus
- The libvirt API provides soft limits for
 - CPU
 - RAM
 - and disk IO
- tc's htb qdiscs allow to prioritize network bandwidth

Agenda

1. Research Questions and Introduction
2. libvirt Hard and Soft Limits
- 3. VM Resizing and Resource Reallocation**
4. OpenStack Extension Design and Implementation
5. OpenStack Extension Evaluation
6. Conclusion
7. Future Work and Outlook

3. VM Resizing and PR Reallocation

- VM **resizing**
 - describes a flavor change
 - **hard** limits change
 - rescheduling needed
 - implemented in OpenStack
- Resource **reallocation**
 - can be flavor independent
 - e.g. a change in VM resource **priorities (soft limits)**
 - no rescheduling needed

Agenda

1. Research Questions and Introduction
2. libvirt Hard and Soft Limits
3. VM Resizing and Resource Reallocation
- 4. OpenStack Extension Design and Implementation**
5. OpenStack Extension Evaluation
6. Conclusion
7. Future Work and Outlook

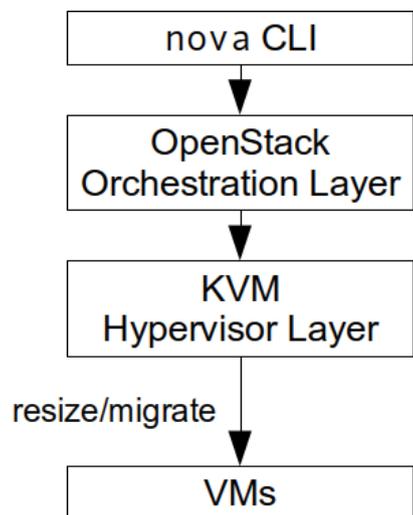
4. OpenStack Extension: Conceptual Formulation

- Design, implementation, evaluation and documentation of an OpenStack extension to control the PR allocation to VMs of individual PMs.
- Extend the OpenStack nova API to allow changes of libvirt's
 - CPU
 - RAM
 - disk IOand tc's network bandwidth soft limits at run-time.
- Extend the nova python client to make use of the new nova API methods.

4. OpenStack Extension: High Level Design

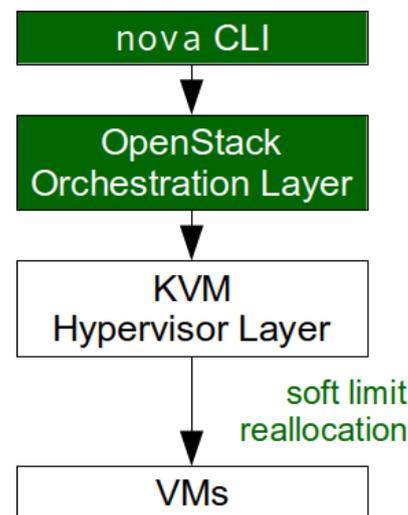
Now:

- flavor bound, boot-time resource limits (soft and hard)
- no network bandwidth soft limit



With extension:

- instance bound, run-time soft limit manipulation for all four resources



4. OpenStack Extension: Implementation

- Different extension architectures in OpenStack API version 2 and version 2.1
- Incomplete documentation and fuzzy entry points
- Extension is based on the stable *kilo* OpenStack release and the legacy OpenStack API version 2
- Devstack development environment (1 host)

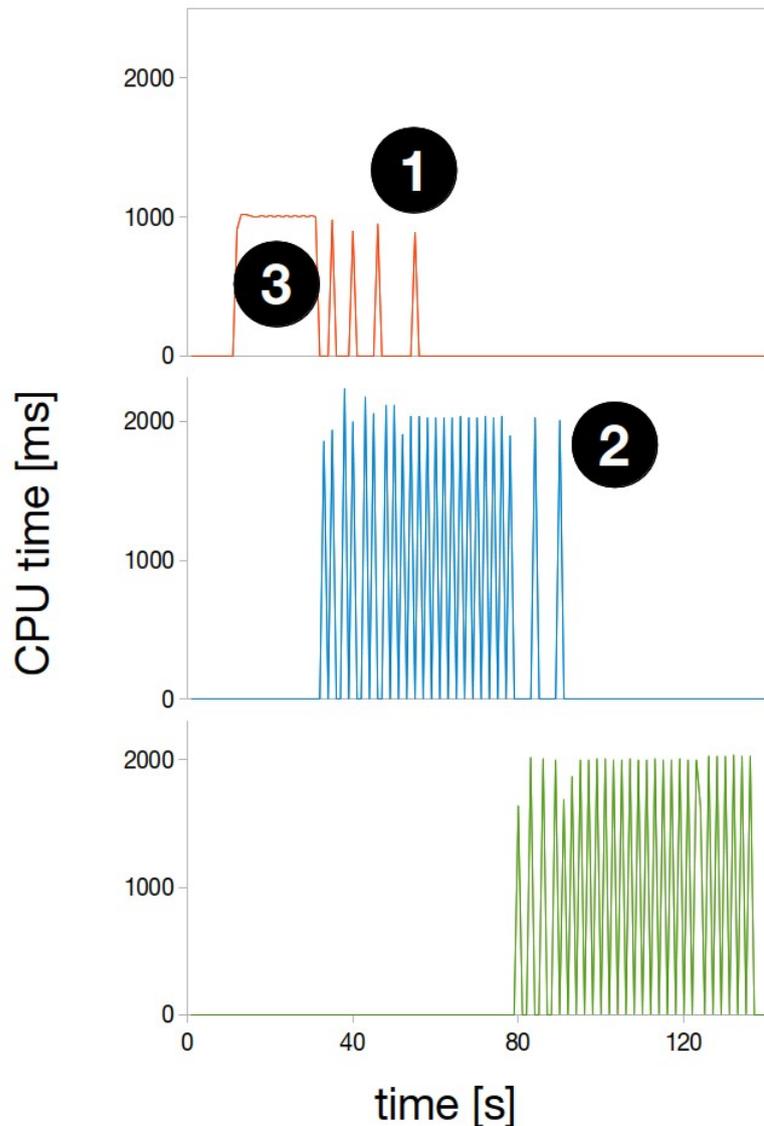
Agenda

1. Research Questions and Introduction
2. libvirt Hard and Soft Limits
3. VM Resizing and Resource Reallocation
4. OpenStack Extension Design and Implementation
- 5. OpenStack Extension Evaluation**
6. Conclusion
7. Future Work and Outlook

5. Evaluation

- The data series in **blue**, **green** and **red** on the following slides represent different VM workloads.

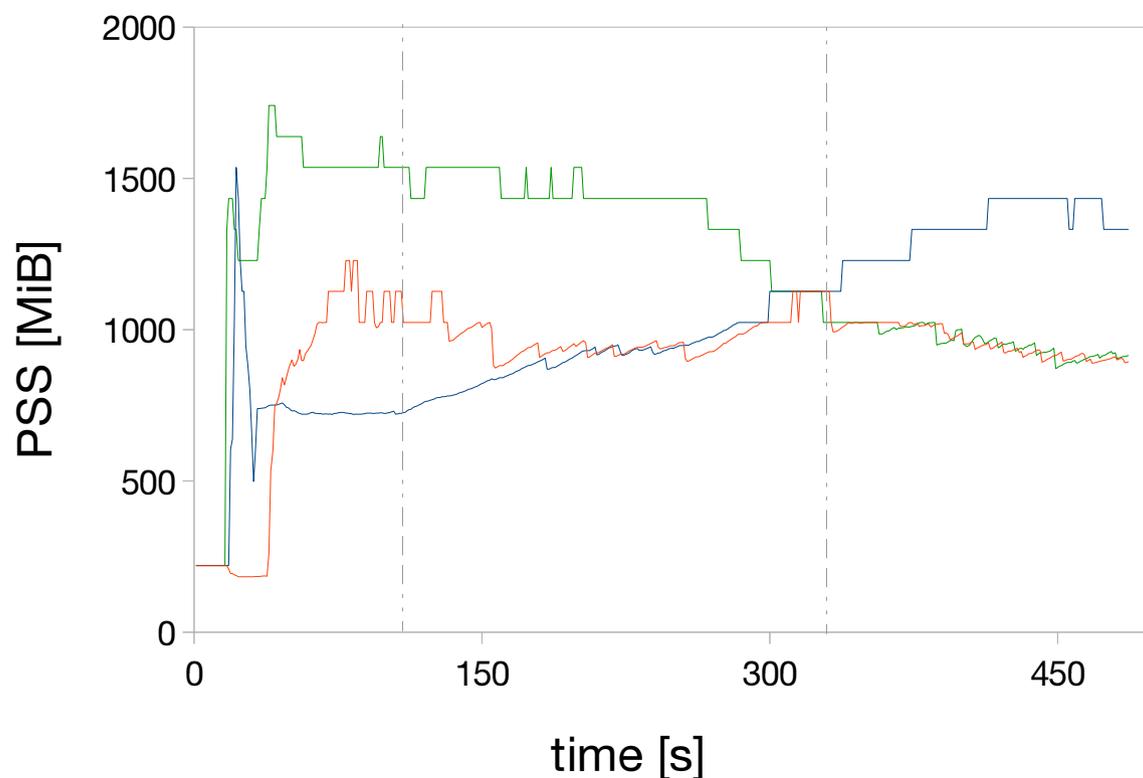
5. Evaluation: CPU Shares



Priorities: 100 500 1'000

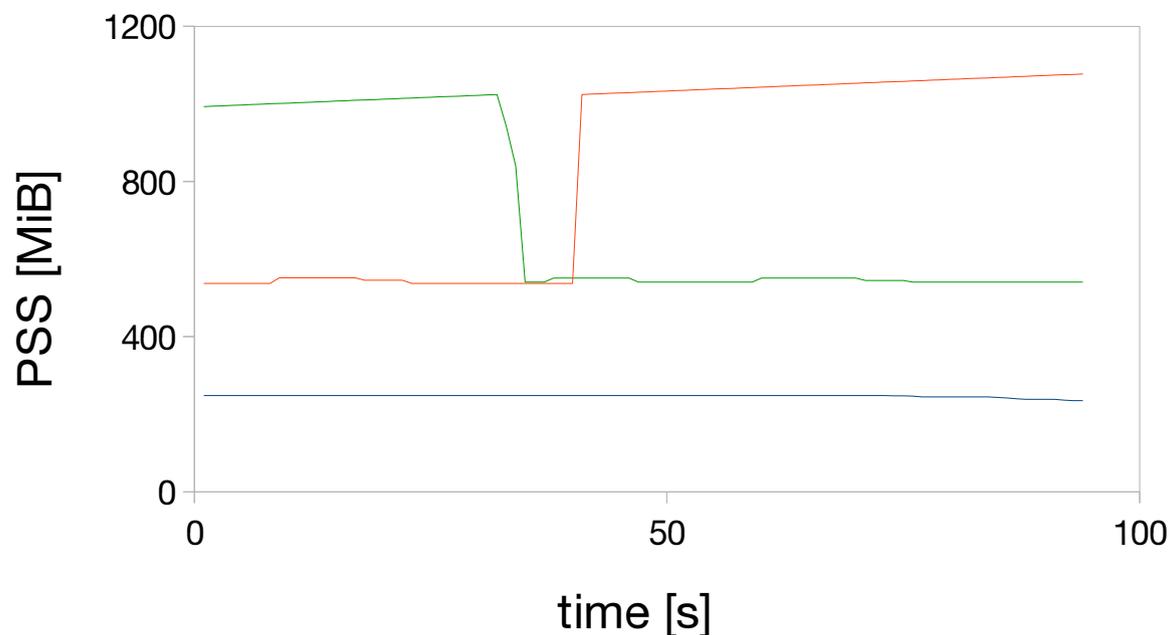
- 1 2 priorities are applied
- 3 constant usage

5. Evaluation: Memory Soft Limit



- best-effort feature
- balancing reclaimed memory takes time (slow adaptation)

5. Evaluation: Memory Hard Limit

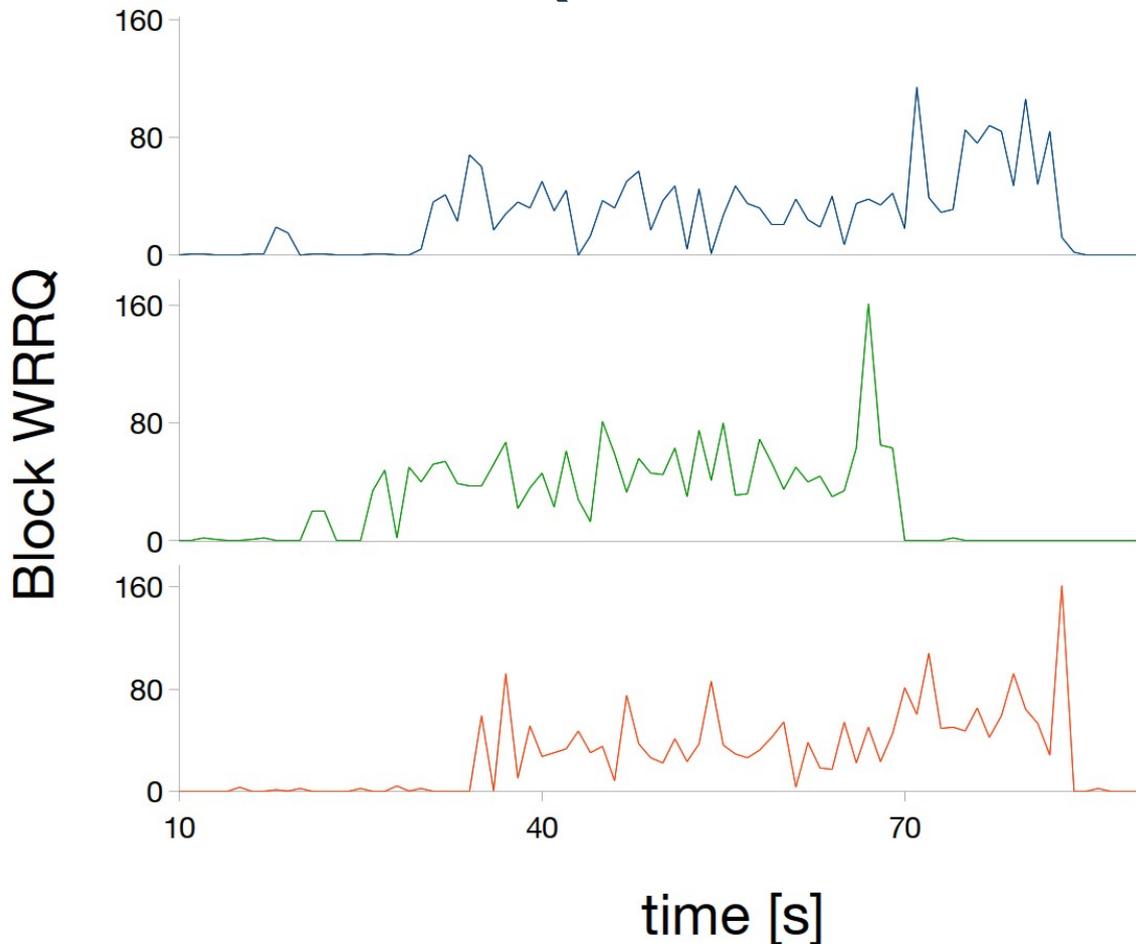


- always active
- upper boundary for the soft limit
- fast adaption

5. Evaluation: Disk IO Soft Limit

Disk IO priorities: 500 900 300

CFQ Scheduler

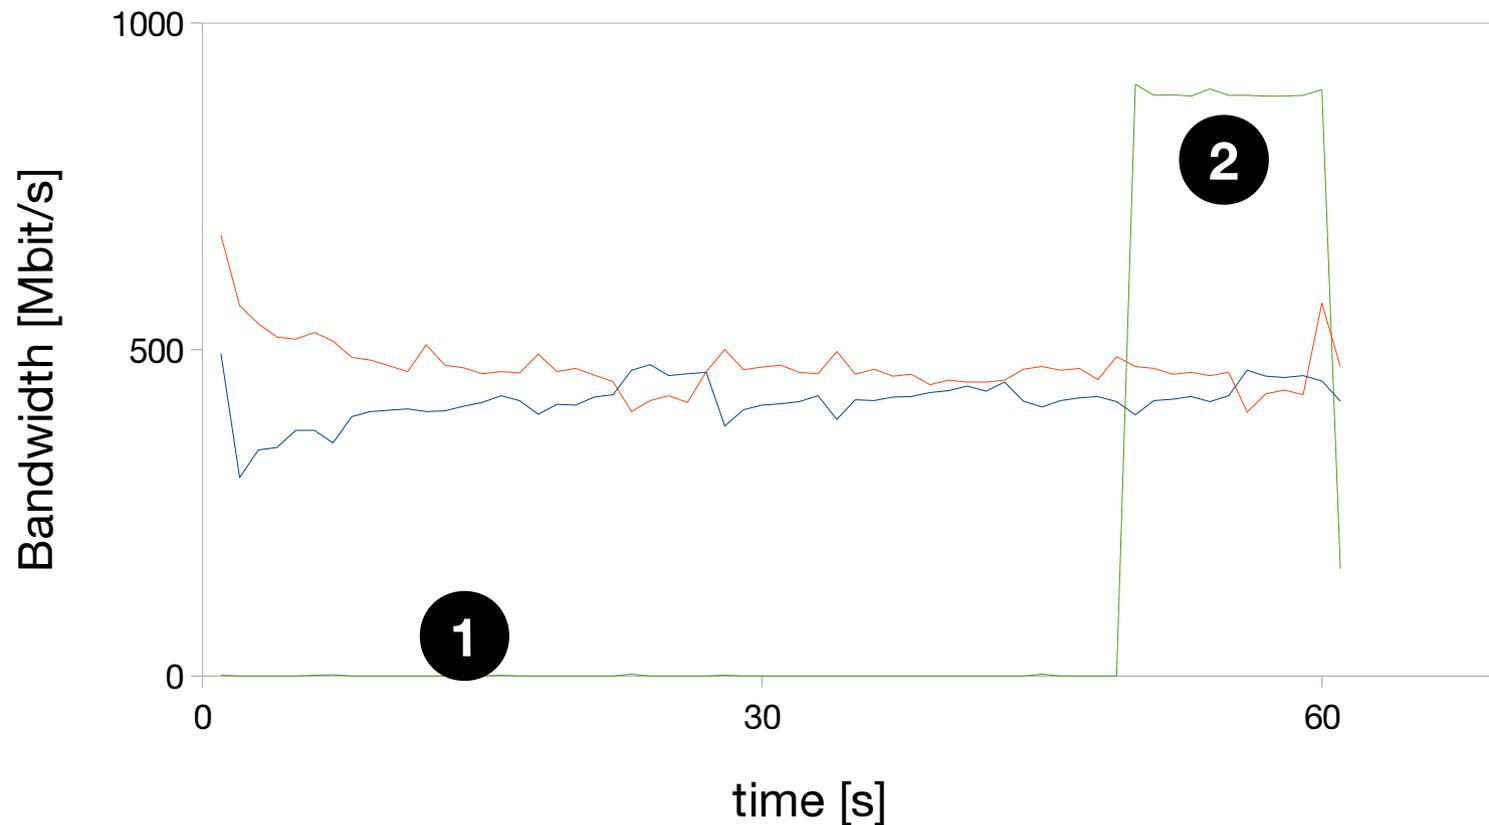


early start,
late finish

early start,
early finish

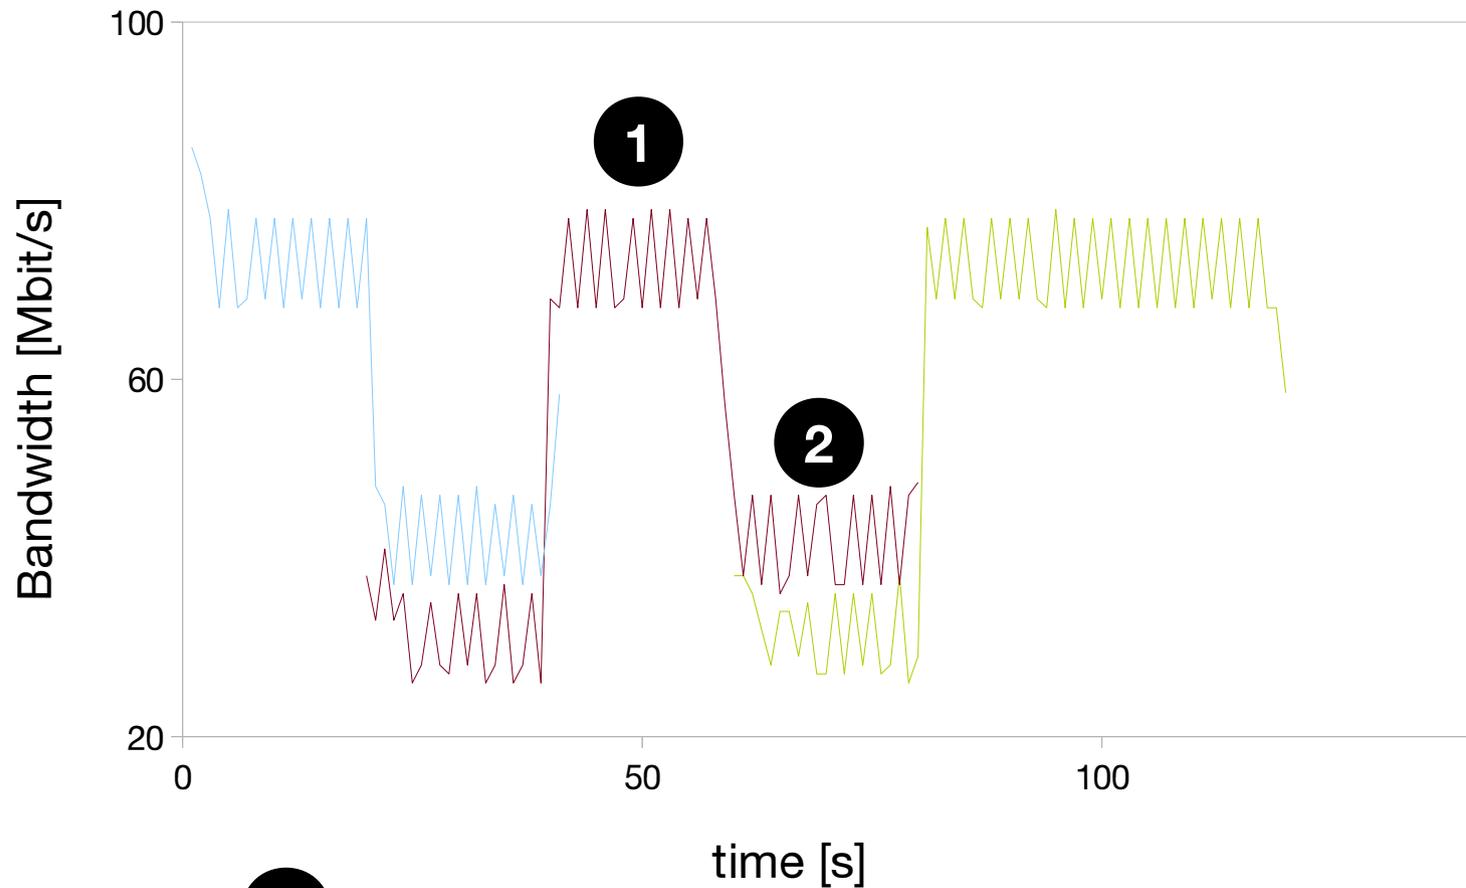
late start,
late finish

5. Evaluation: Network Bandwidth Soft Limit



- 1 starvation of low priorities
- 2 unreliable, unexpected behaviour

5. Evaluation: Network Bandwidth Soft Limit



Priorities:

39 26 15

1 borrowing

2 proportional borrowing

Agenda

1. Research Questions and Introduction
2. libvirt Hard and Soft Limits
3. VM Resizing and Resource Reallocation
4. OpenStack Extension Design and Implementation
5. OpenStack Extension Evaluation
- 6. Conclusion**
7. Future Work and Outlook

6. Conclusion

All four resources,

- CPU,
- disk IO,
- RAM and
- network bandwidth

can be reallocated between VMs competing for scarce resources with the help of libvirt's soft limits and tc's htb qdisc.

6. Conclusion

- How are PRs reallocated in cloud computing?
 - libvirt virtualization API
 - tc
- Which resources can be reallocated in KVM/OpenStack?
 - All four resources, CPU, disk IO, RAM and network bandwidth
- Which methods are used in KVM/OpenStack to reallocate resources?
 - Resizing (live migration)

Agenda

1. Research Questions and Introduction
2. libvirt Hard and Soft Limits
3. VM Resizing and Resource Reallocation
4. OpenStack Extension Design and Implementation
5. OpenStack Extension Evaluation
6. Conclusion
- 7. Future Work and Outlook**

7. Future Work

- Test the new API extension in a “real world” environment with multiple compute nodes
- Reset network priorities (delete tc qdiscs)
- Improve nova client usability
 - RAM limits: different units (chosen by the user)
 - Uniform priority scale for CPU, disk IO and network bandwidth
 - Priority based RAM reallocation (to be consistent with the other three resources)
- Rewrite extension for the OpenStack API version 2.1 (new microframework plugin architecture)

7. Outlook

- Flavor change at run-time will not be part of the next stable OpenStack release.
 - Maybe resource priorities soon will?
- Is there interest in resource reallocation with soft limits?

Thank You for Your Attention!

Andreas Gruhler
<andreas.gruhler@uzh.ch>