# Operating System Noise: Linux vs. Microkernel

Stefan Wächtler[1], Michal Sojka[1,2]

[1]Technische Universität Dresden
Operating-Systems Group

[2]Czech Technical University in Prague
Faculty of Electrical Engineering

14th Real-Time Linux Workshop
October 18, Chapel Hill, NC, USA

# Introduction

- Computers running Linux are often used as a tool to measure time in various experiments.
- Benchmarking etc.

How does the operating system influence the results of benchmarks?

# Introduction

- Computers running Linux are often used as a tool to measure time in various experiments.
- Benchmarking etc.

How does the operating system influence the results of benchmarks?

## Expectation

- Microkernels should give results with less variance.
- Less interference from the OS.

# This work

- Statistical comparison of benchmark results.
- The same set benchmarks run on different OSes:
  - Linux, Debian kernel
  - Linux, minimal kernel
  - NOVA microkernel

# NOVA microkernel

- Microkernel-based hypervisor.
- Extremely small (9 kLoC) and fast.
- Almost nothing runs behind application's back.

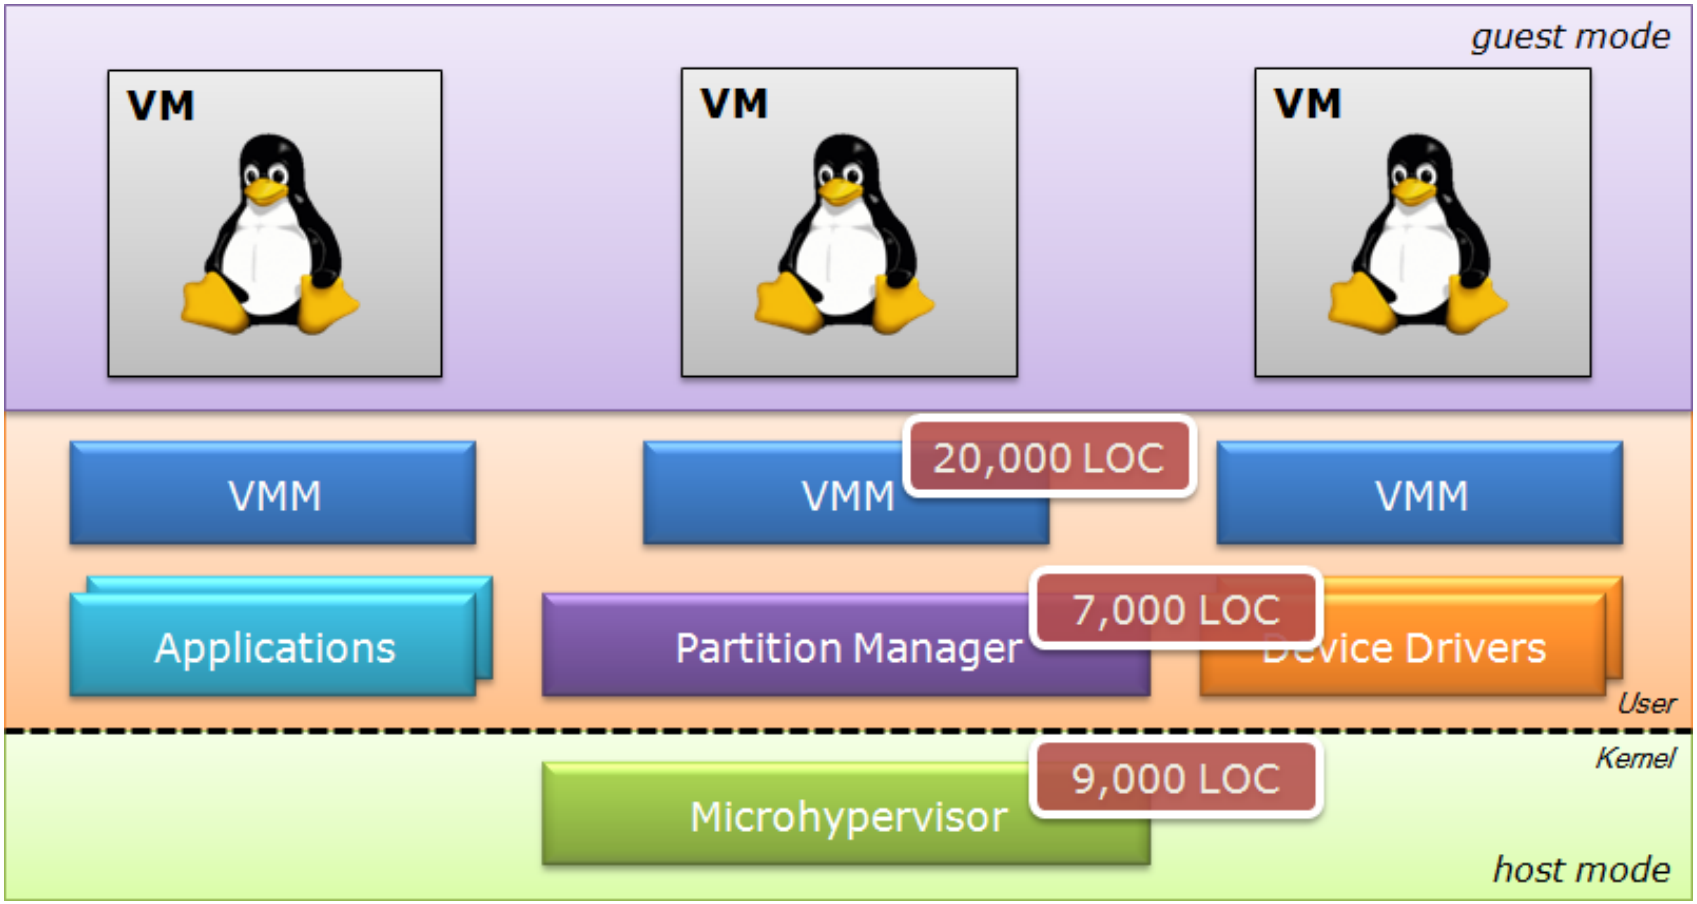

# OS Noise in High Performance Computing

- Timer interrupts etc.
- Scalability problem

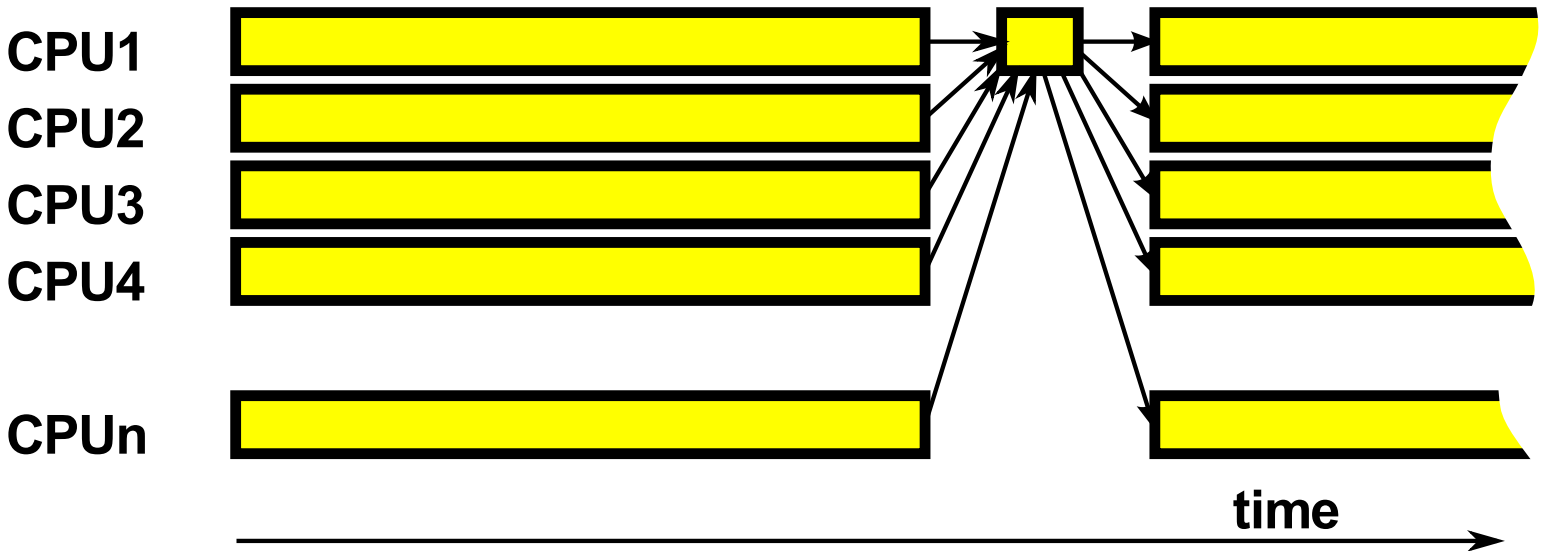

# OS Noise in High Performance Computing

- Timer interrupts etc.
- Scalability problem

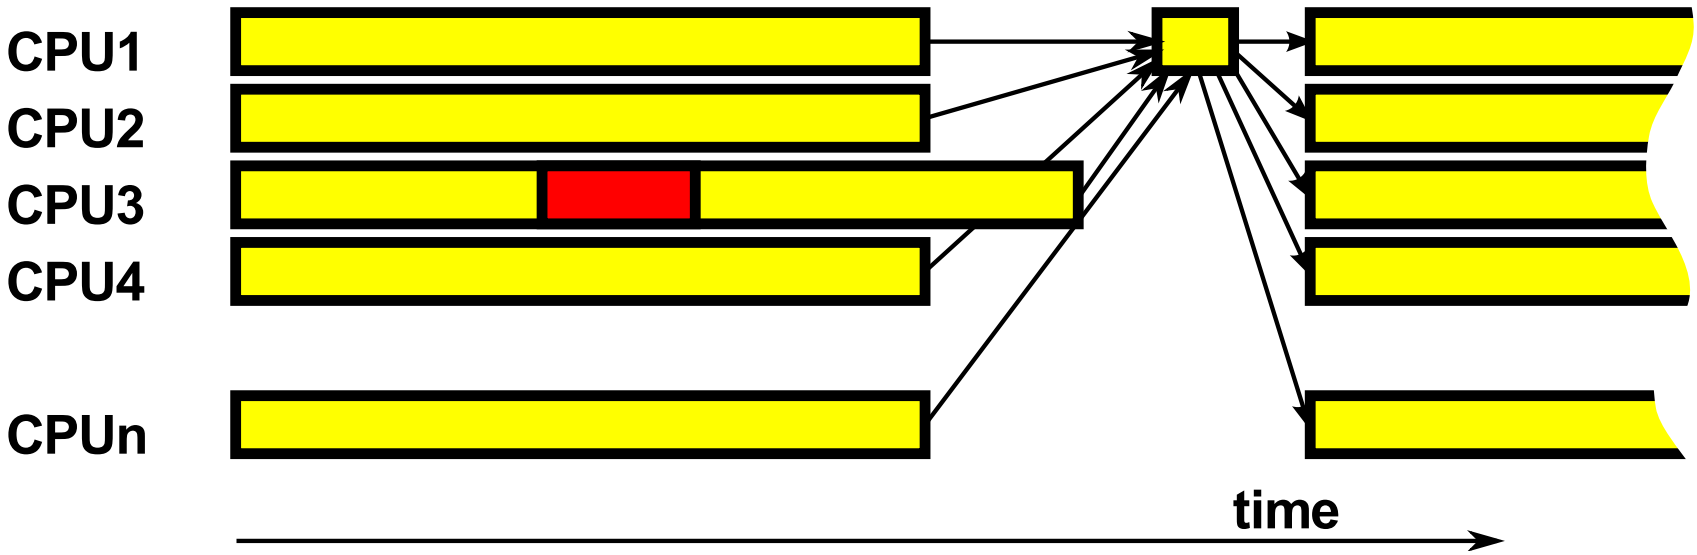

## Benchmarks

- Measure time of certain operations
- In total 46 experiments:
- cache/memory bandwidth (8 experiments),
- cache-related preemption delay (32 experiments),
- cache-related migration delay (6 experiments)

# Benchmarks

## Benchmarks

- Measure time of certain operations
- In total 46 experiments:
- cache/memory bandwidth (8 experiments),
- cache-related preemption delay (32 experiments),
- cache-related migration delay (6 experiments)

## Method

- Treat benchmarks as black boxes
- Look at the variance of results

# Benchmarks

## Benchmarks

- Measure time of certain operations
- In total 46 experiments:
- cache/memory bandwidth (8 experiments),
- cache-related preemption delay (32 experiments),
- cache-related migration delay (6 experiments)

## Method

- Treat benchmarks as black boxes
- Look at the variance of results

## Testbed

- Dell Precision T7500, 2$\times$ Intel Xeon X5650 (6 cores each), NUMA architecture

# First Benchmark – Memory Access Times

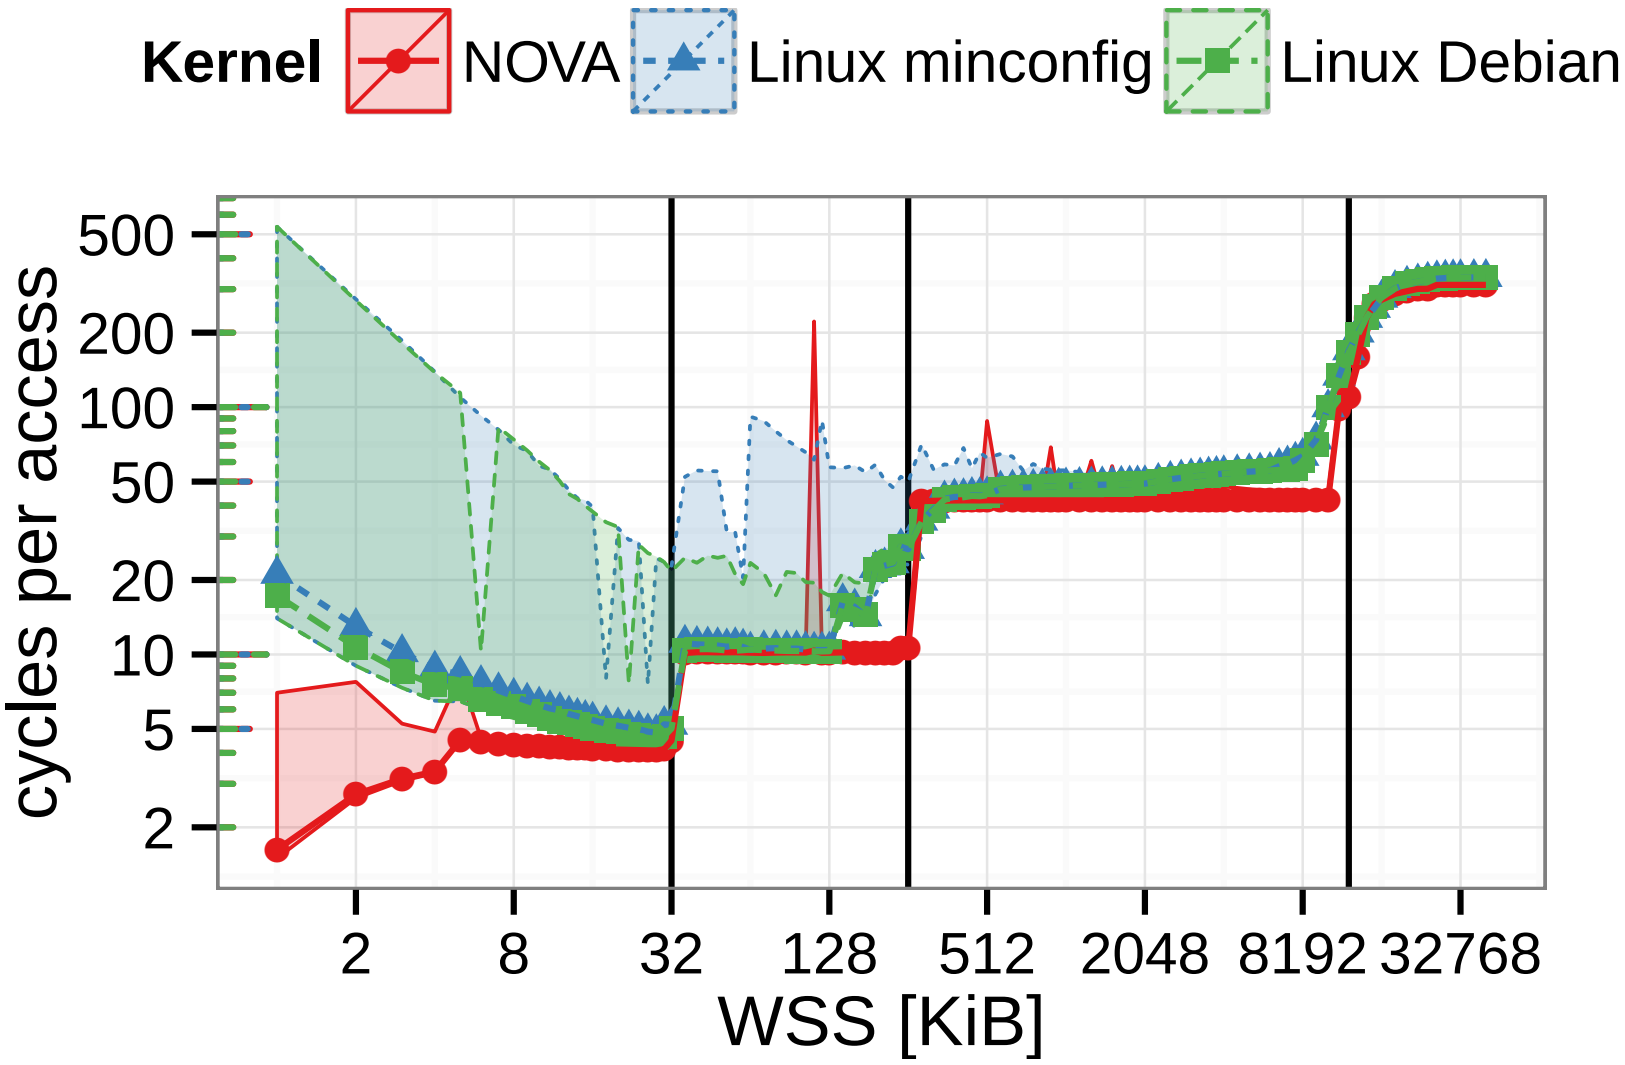

# Observations

1. Steps on the NOVA curve are much sharper.
2. The differences between the mean values for the two Linux kernels is very small.
3. The differences at the left end of the graph is caused by inaccuracy of our measurement method.
4. There are big differences in the observed noise (min/max values).
5. Peaks on the NOVA curve: perhaps caused SMIs.

# Second Benchmark – Memory Access Times

## Modified cache lines

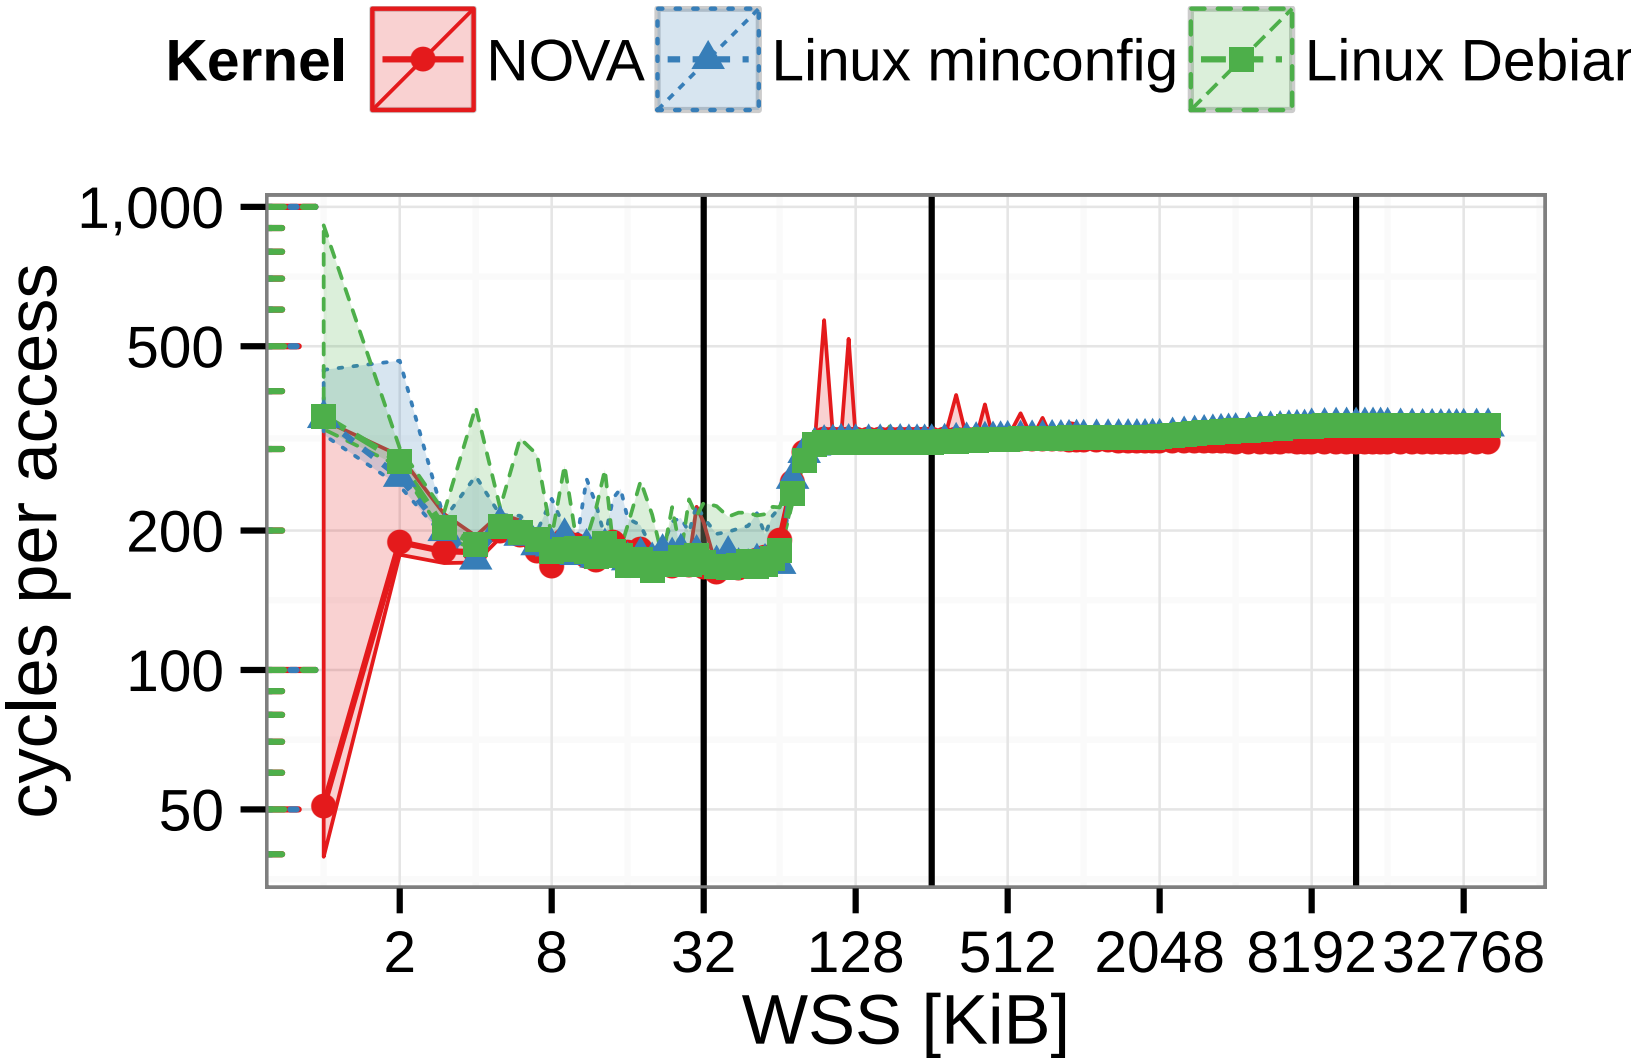

# How to compare “the noise”?

Several options:

- Standard deviation $\sigma$
- Variance $\sigma^2$
- Coefficient of variation $c_v = \sigma/\mu$
- Variance to mean ratio (index of dispersion) $D = \sigma^2/\mu$

# How to compare “the noise”?

Several options:

- Standard deviation $\sigma$
- Variance $\sigma^2$
- Coefficient of variation $c_v = \sigma/\mu$
- Variance to mean ratio (index of dispersion) $D = \sigma^2/\mu$

# Variance of the Second Benchmark

- Benchmark executed 1024 times

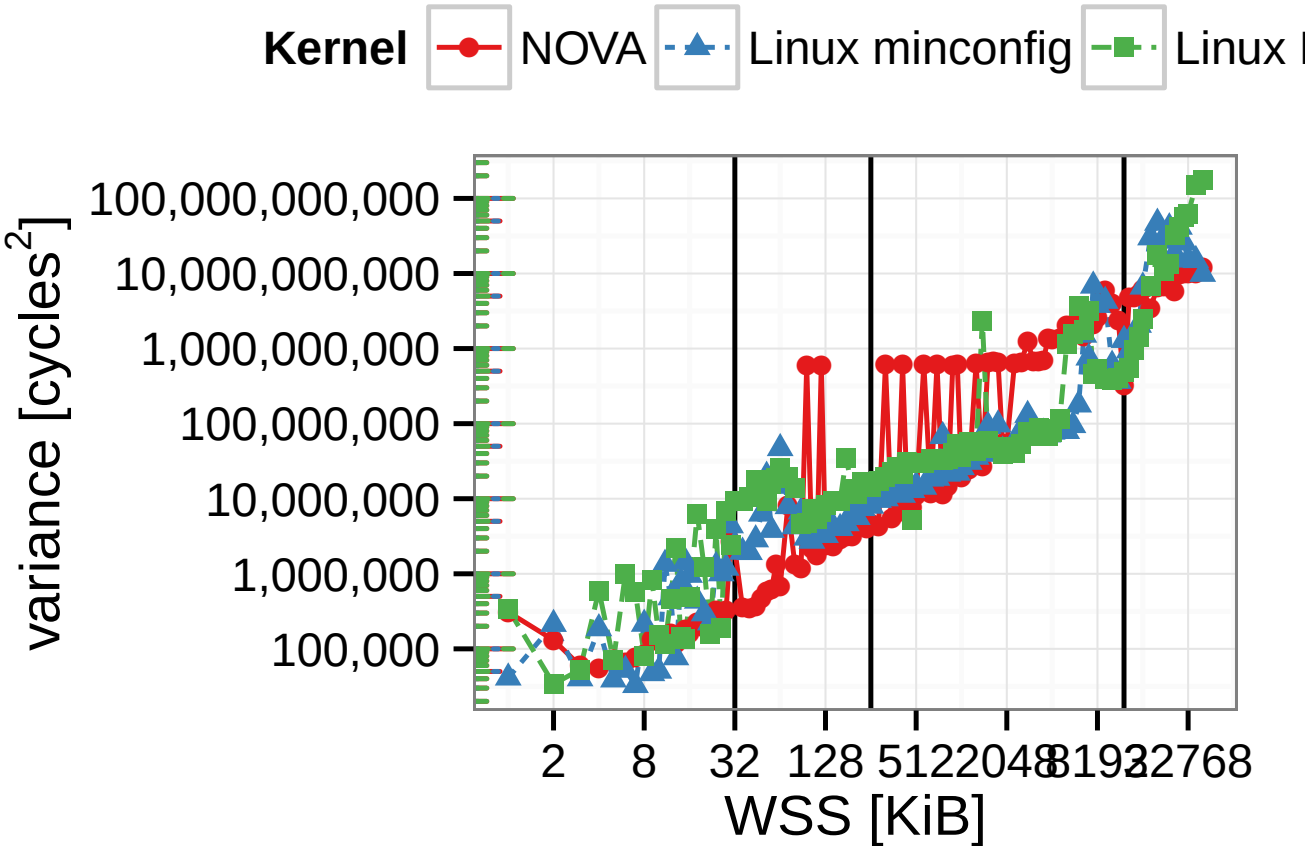

# Variance to mean ratio

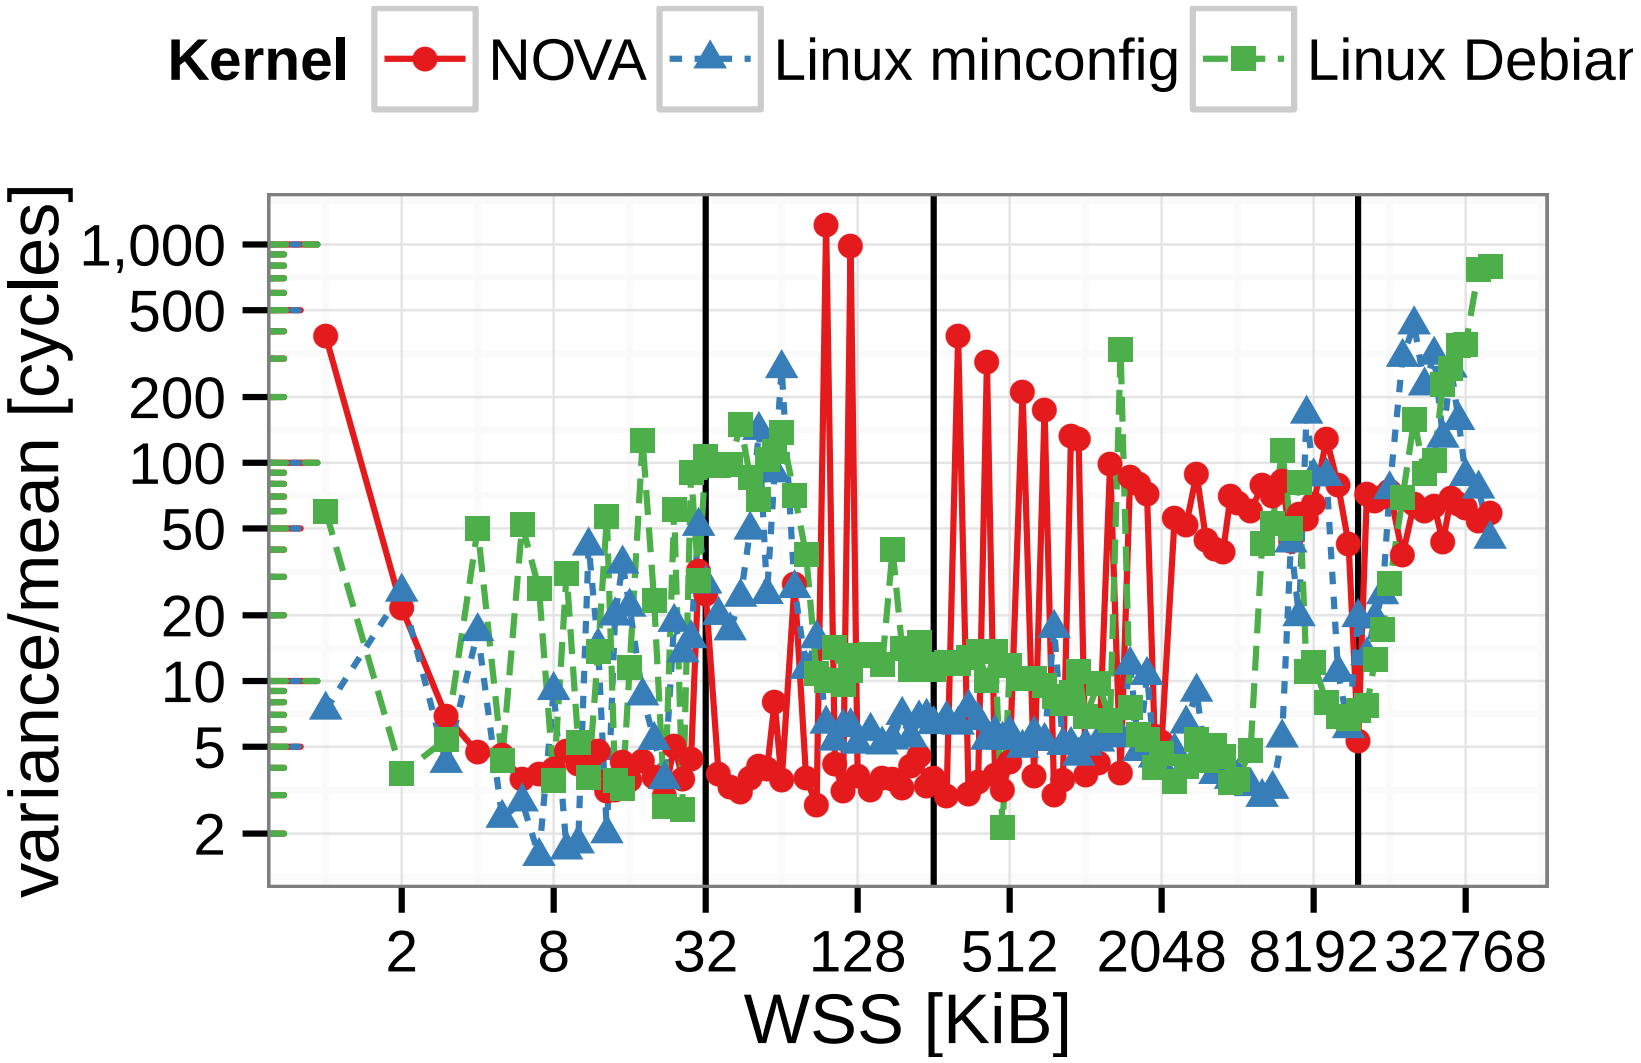

# First benchmark – Aggregated results

- Split to 4 WSS areas according to cache sizes
- Median calculation

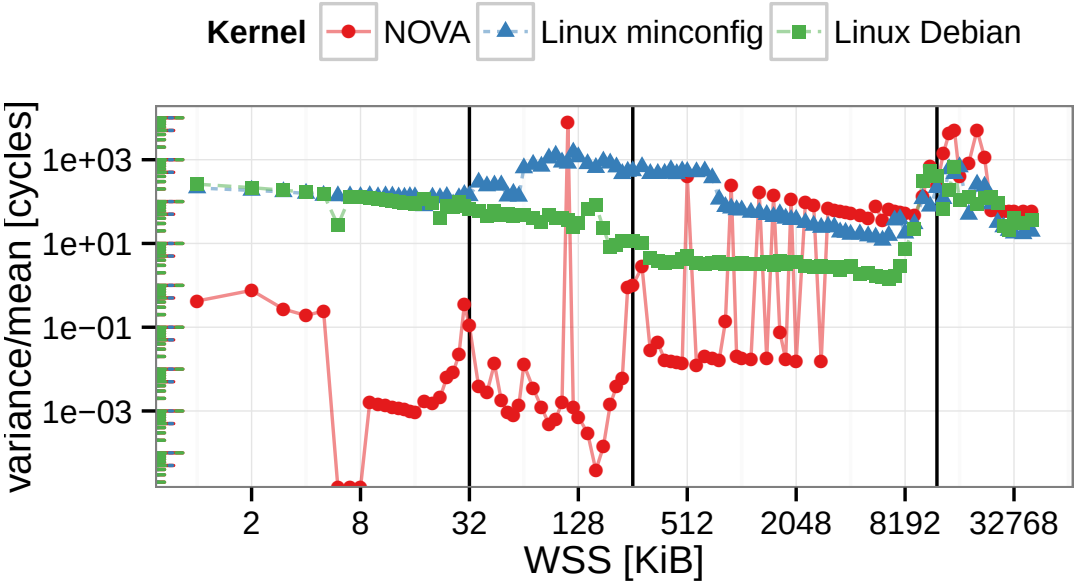


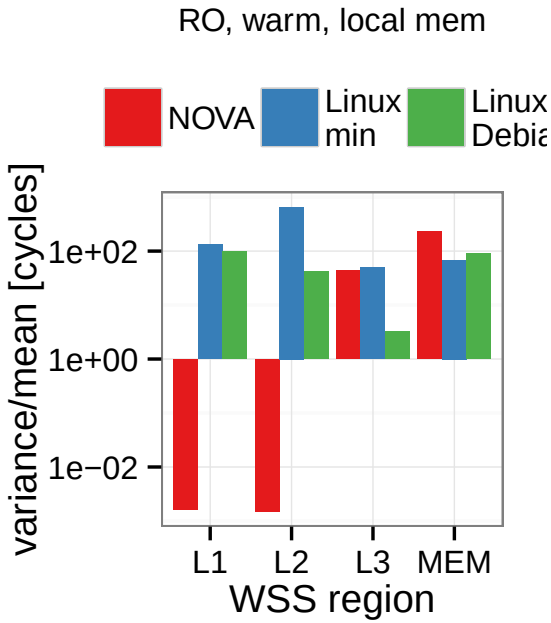

# Aggregated results of all 46 experiments

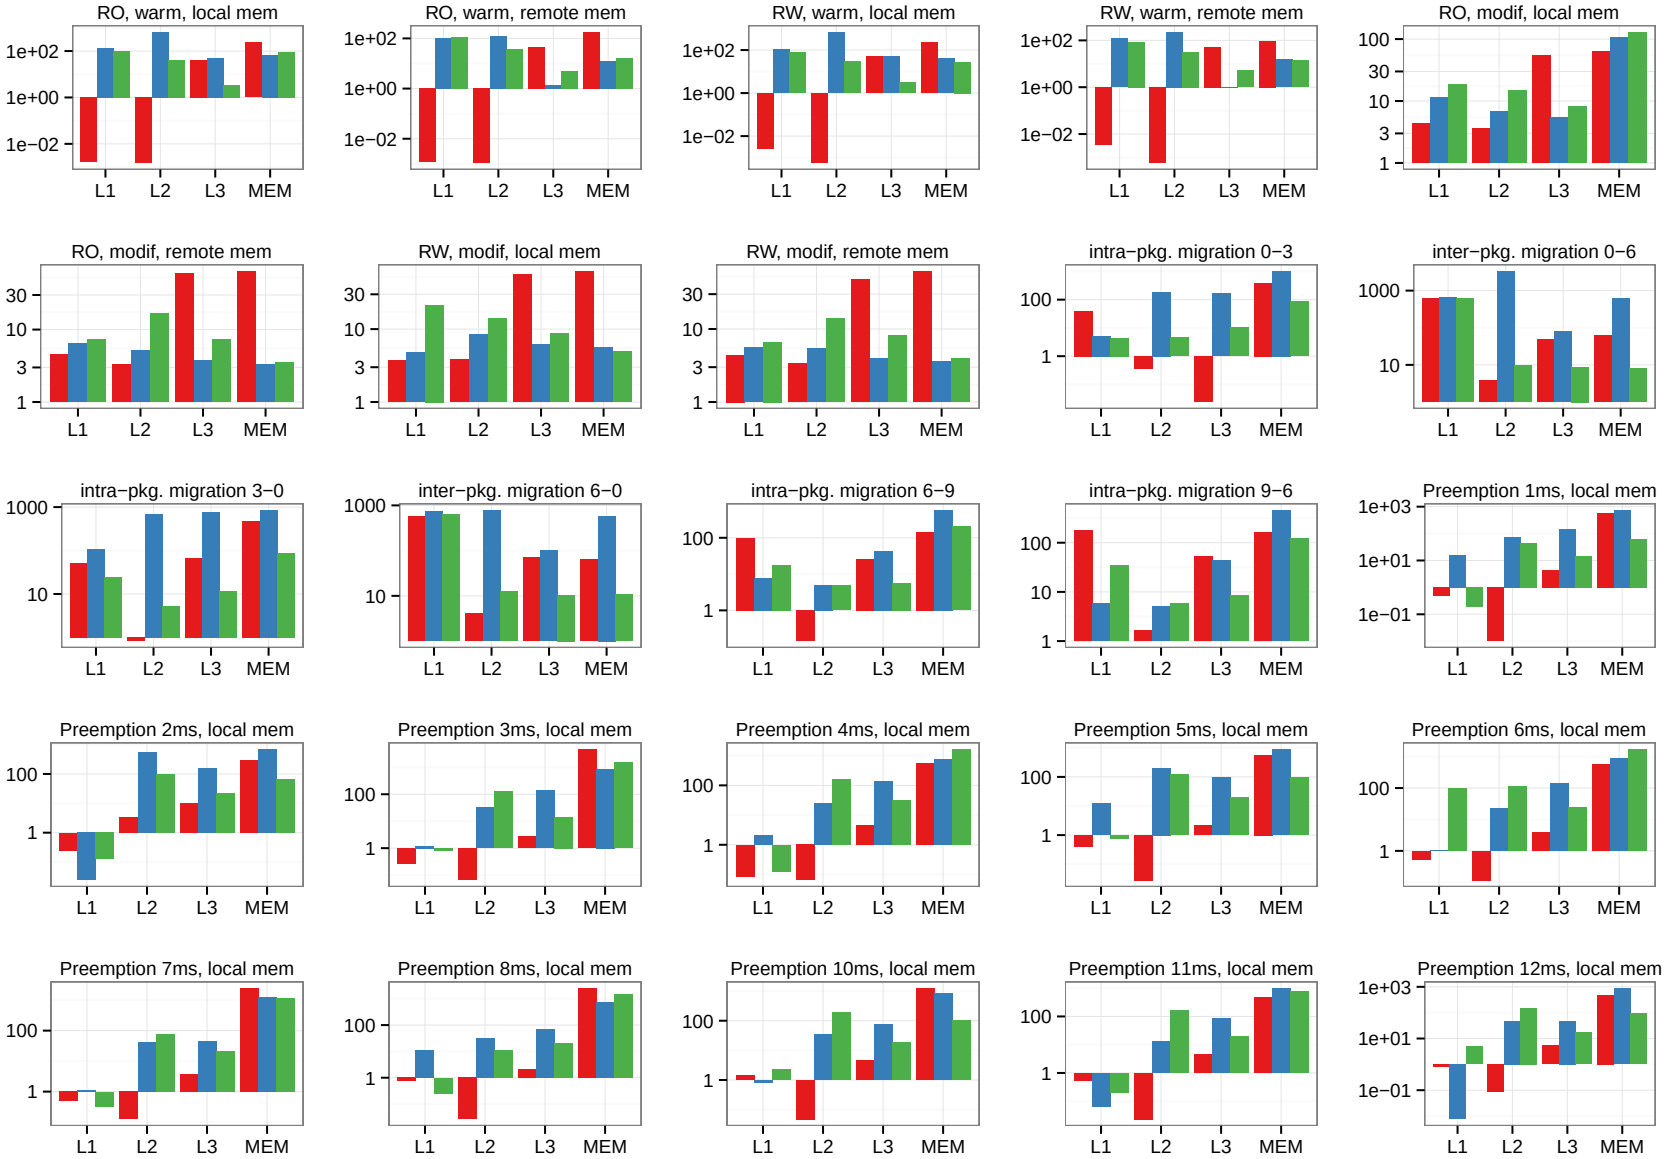

# Aggregated results of all 46 experiments

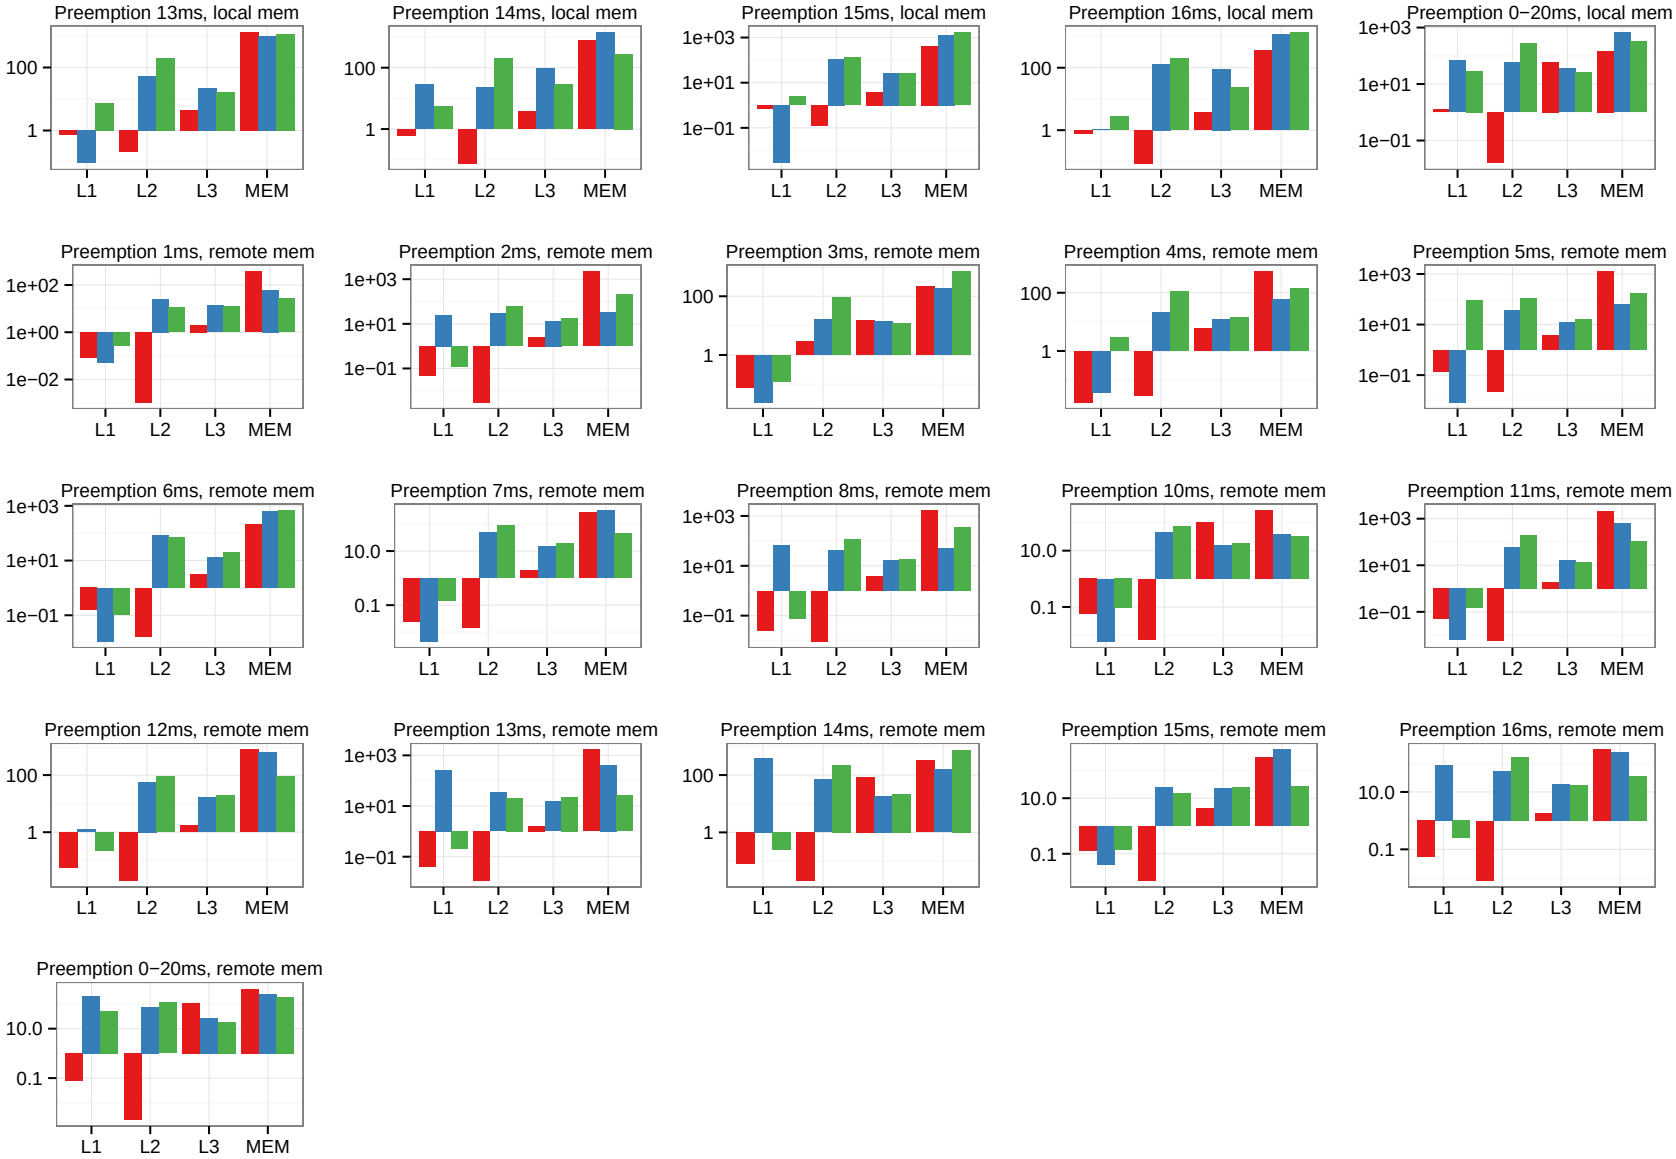

# Summarized results

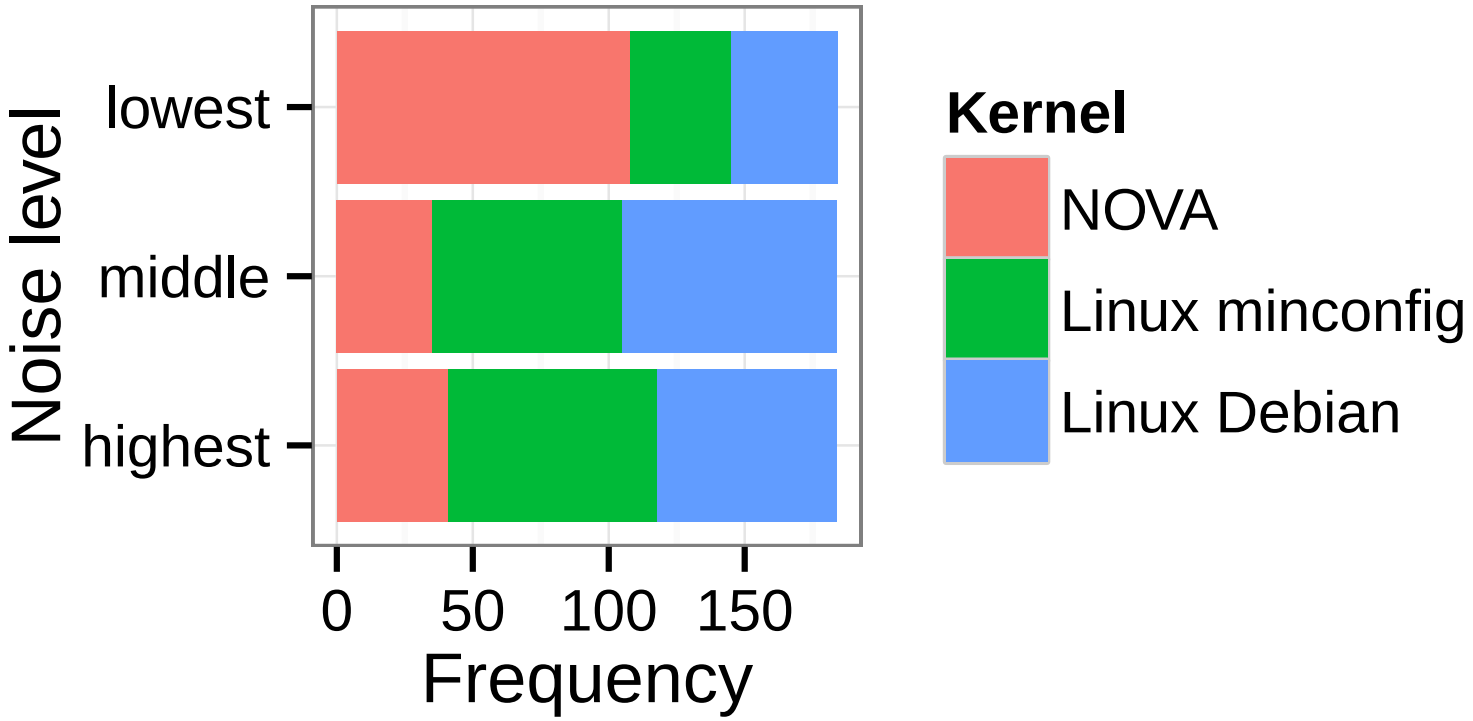


| OS | lowest | middle | highest |
|---|---|---|---|
| NOVA | 108 (59%) | 35 (19%) | 41 (22%) |
| Linux min. | 37 (20%) | 70 (38%) | 77 (42%) |
| Linux Debian | 39 (21%) | 79 (43%) | 66 (36%) |

# Conclusions

- In most cases NOVA microkernel produced lowest noise.
- In several other cases NOVA produced highest noise.
- System management interrupts(?) are really a big source of noise.
- Operating system is not the only dominant source of noise. Off-chip access increases the noise.
- Configuration of Linux kernel has significant impact on the observed noise.