Musings on Analysis of Measurements of a Real-Time Workload

Analysis of real-time performance is often focused on the worst case "bad" value, such as maximum latency or maximum interrupt disabled time. This data is critical for a hard real-time application, but may be overly strict for a soft real-time application. This talk examines some statistics that may be interesting and useful for examining soft real-time applications.

Frank Rowand, Sony Corporation of America

April 6, 2009

DISCLAIMER

This is not a statistics class (which would require a bit more than 45 minutes).

Apologies to any true statisticians in the audience! Please try to not get too annoyed with the over-simplifications presented.

At the end of this talk you should know enough to be dangerous. I hope you become sufficiently intrigued to delve more deeply into the concepts, and thus less dangerous.

Overview

- Some statistics (theory)

- Applying the theory to some data

- Some ad-hoc data analysis,
  not so much using the theory

Statistics

This will be fun. Really...

Describing Data

How do you characterize data?

Describing Data

How do you characterize data?

Some typical goals of characterization:

- Understand how multiple data sets differ (when I changed the system, did the performance get better or worse?)

- Understand if my real-time application will complete its work before the deadline

Describing Data

- Maximum

Describing Data

- Maximum

Everyone knows what this means...

Describing Data

- Maximum

Everyone knows what this means...

There is obvious utility for characterizing real-time systems.

The real-time task must meet a deadline:

Is maximum latency + work < deadline?

Describing Data

- Maximum
- Minimum
- Mean

$$mean = \frac{1}{n} \sum_{i=1}^{n} x_i$$

Describing Data

- Maximum
- Minimum
- Mean

Can be useful to understand how behavior (performance) responds to changes in the system.

Low overhead to collect this data,
thus often available.

Describing Data

- Maximum
- Minimum
- Mean
- Median
- [Mode]

50% of data values < Median < 50% of data values

Describing Data

- Maximum
- Minimum
- Mean
- Median
- [Mode]

Median is in the same family as Mean, but may provide different insights.

Mode is usually not useful for my work, but is mentioned for completeness.

Describing Data

- Maximum
- Minimum
- Mean
- Median
- [Mode]
- Visualization (look at the raw data or graph)

Describing Data

- Maximum
- Minimum
- Mean
- Median
- [Mode]
- Visualization (look at the raw data or graph)

Pattern recognition can be very powerful.

Describing Data

- Maximum
- Minimum
- Mean
- Median
- [Mode]
- Visualization (look at the raw data or graph)
- Standard Deviation

Describing Data

- Maximum
- Minimum
- Mean
- Median
- [Mode]
- Visualization (look at the raw data or graph)
- Standard Deviation

A possible measure of how much the data varies.

standard deviation – What is it?

First, what is a frequency distribution?

standard deviation – What is it?

First, what is a frequency distribution?

Start by creating some random data
for an example.

Next slides are the data.

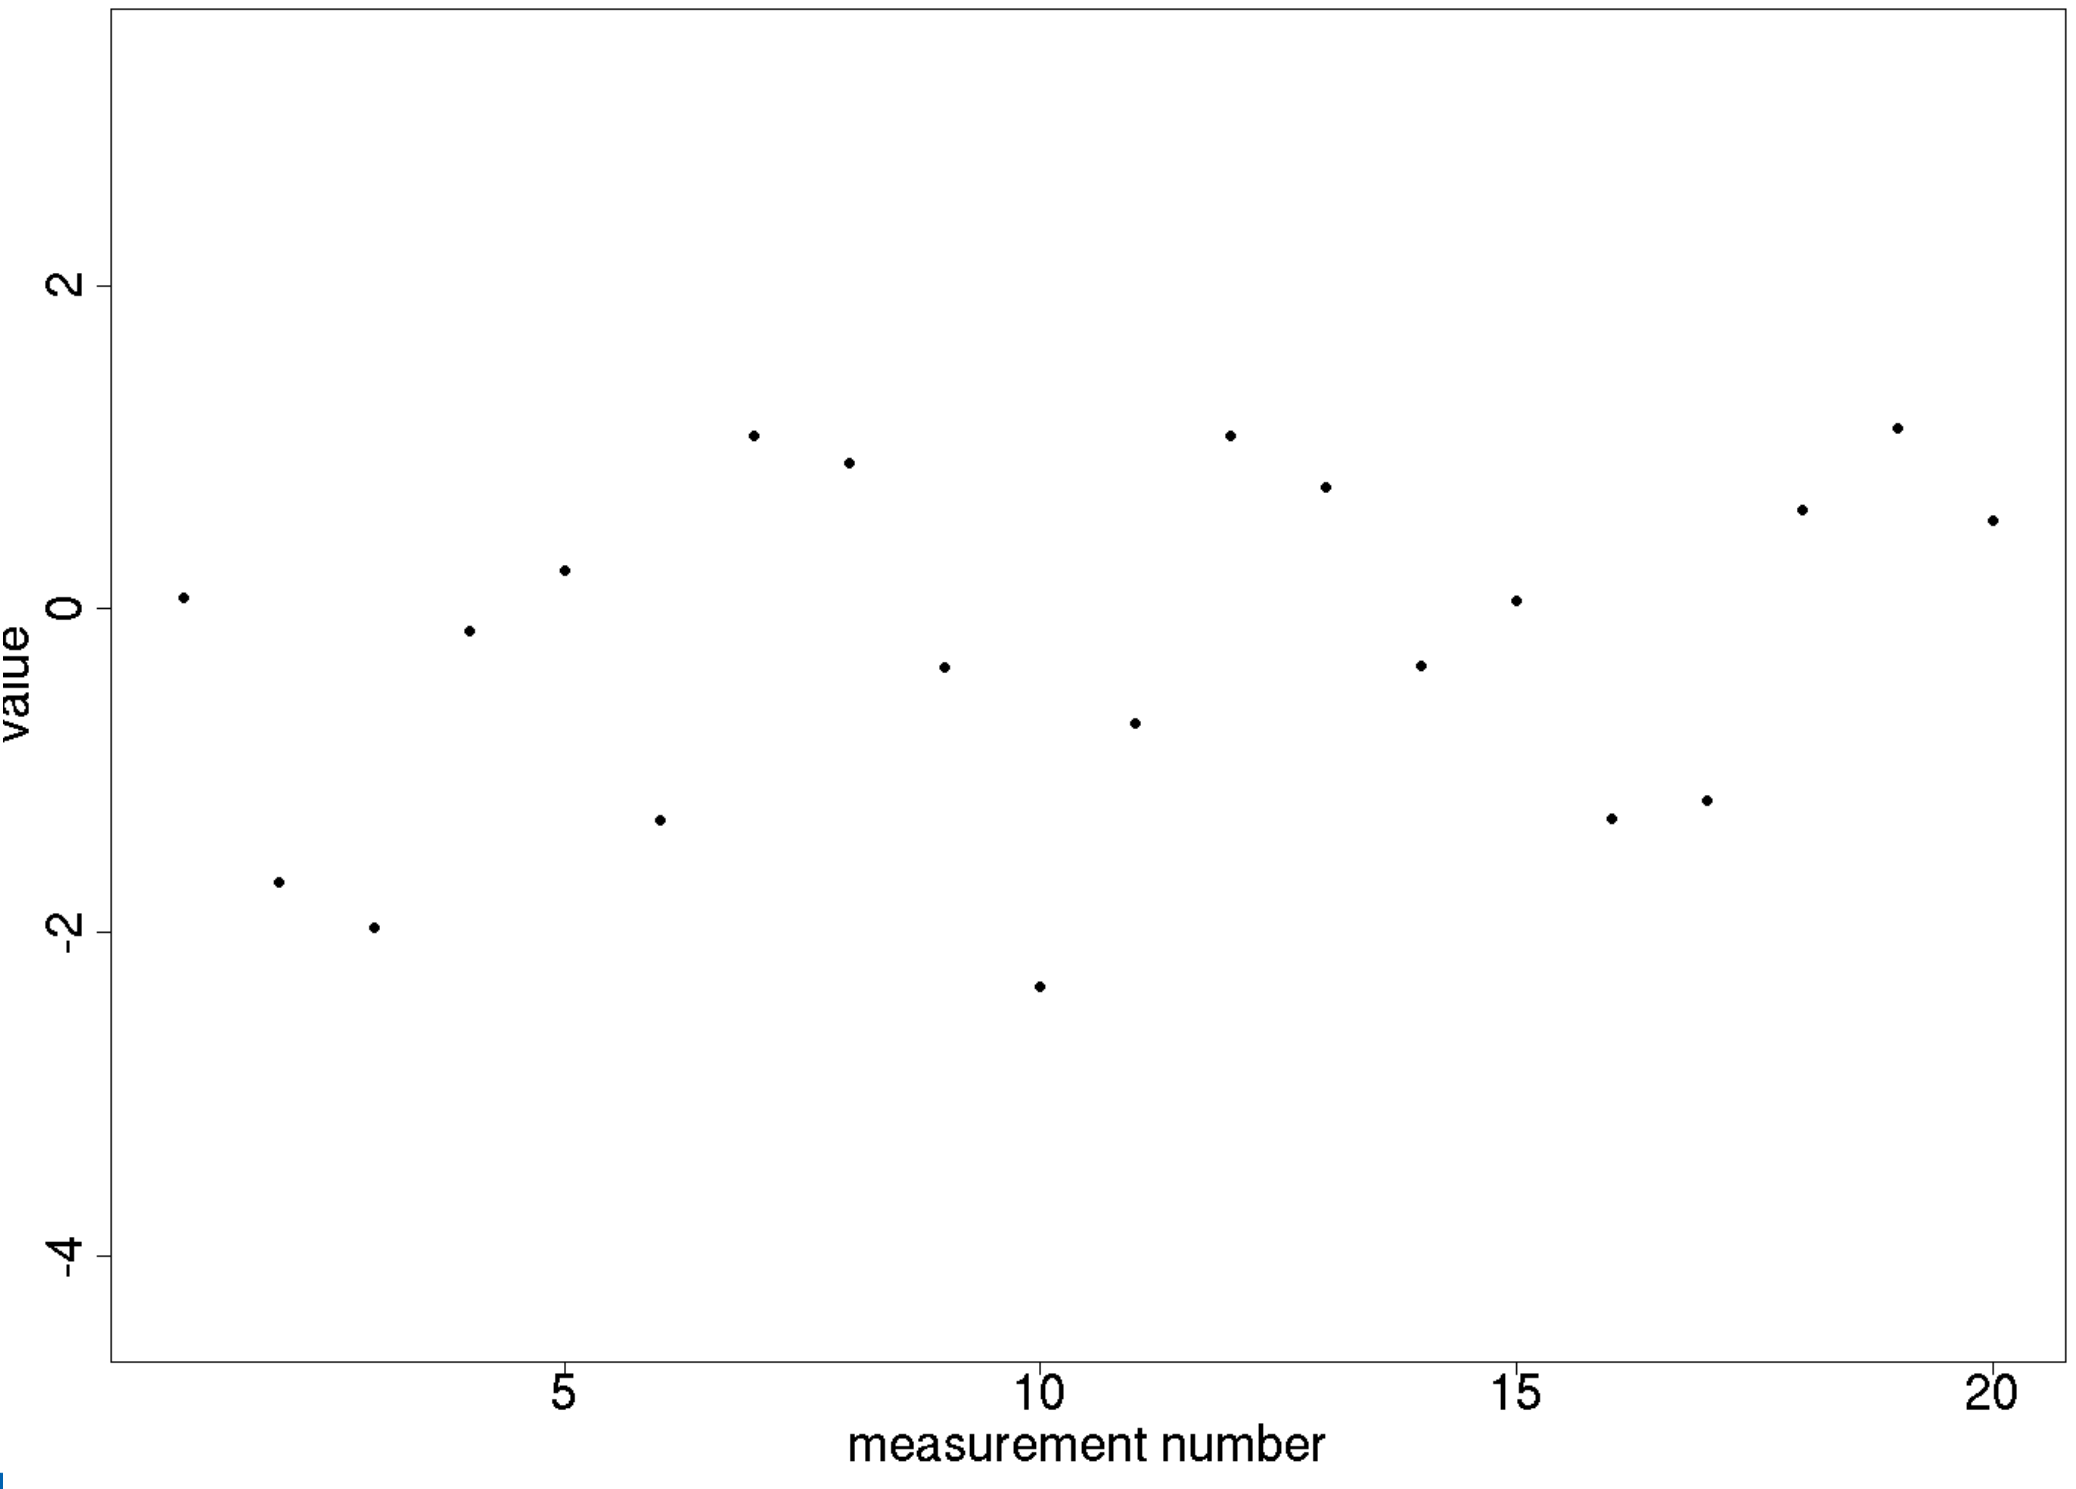
value
2
0
-2
-4
5
10
15
20
measurement number

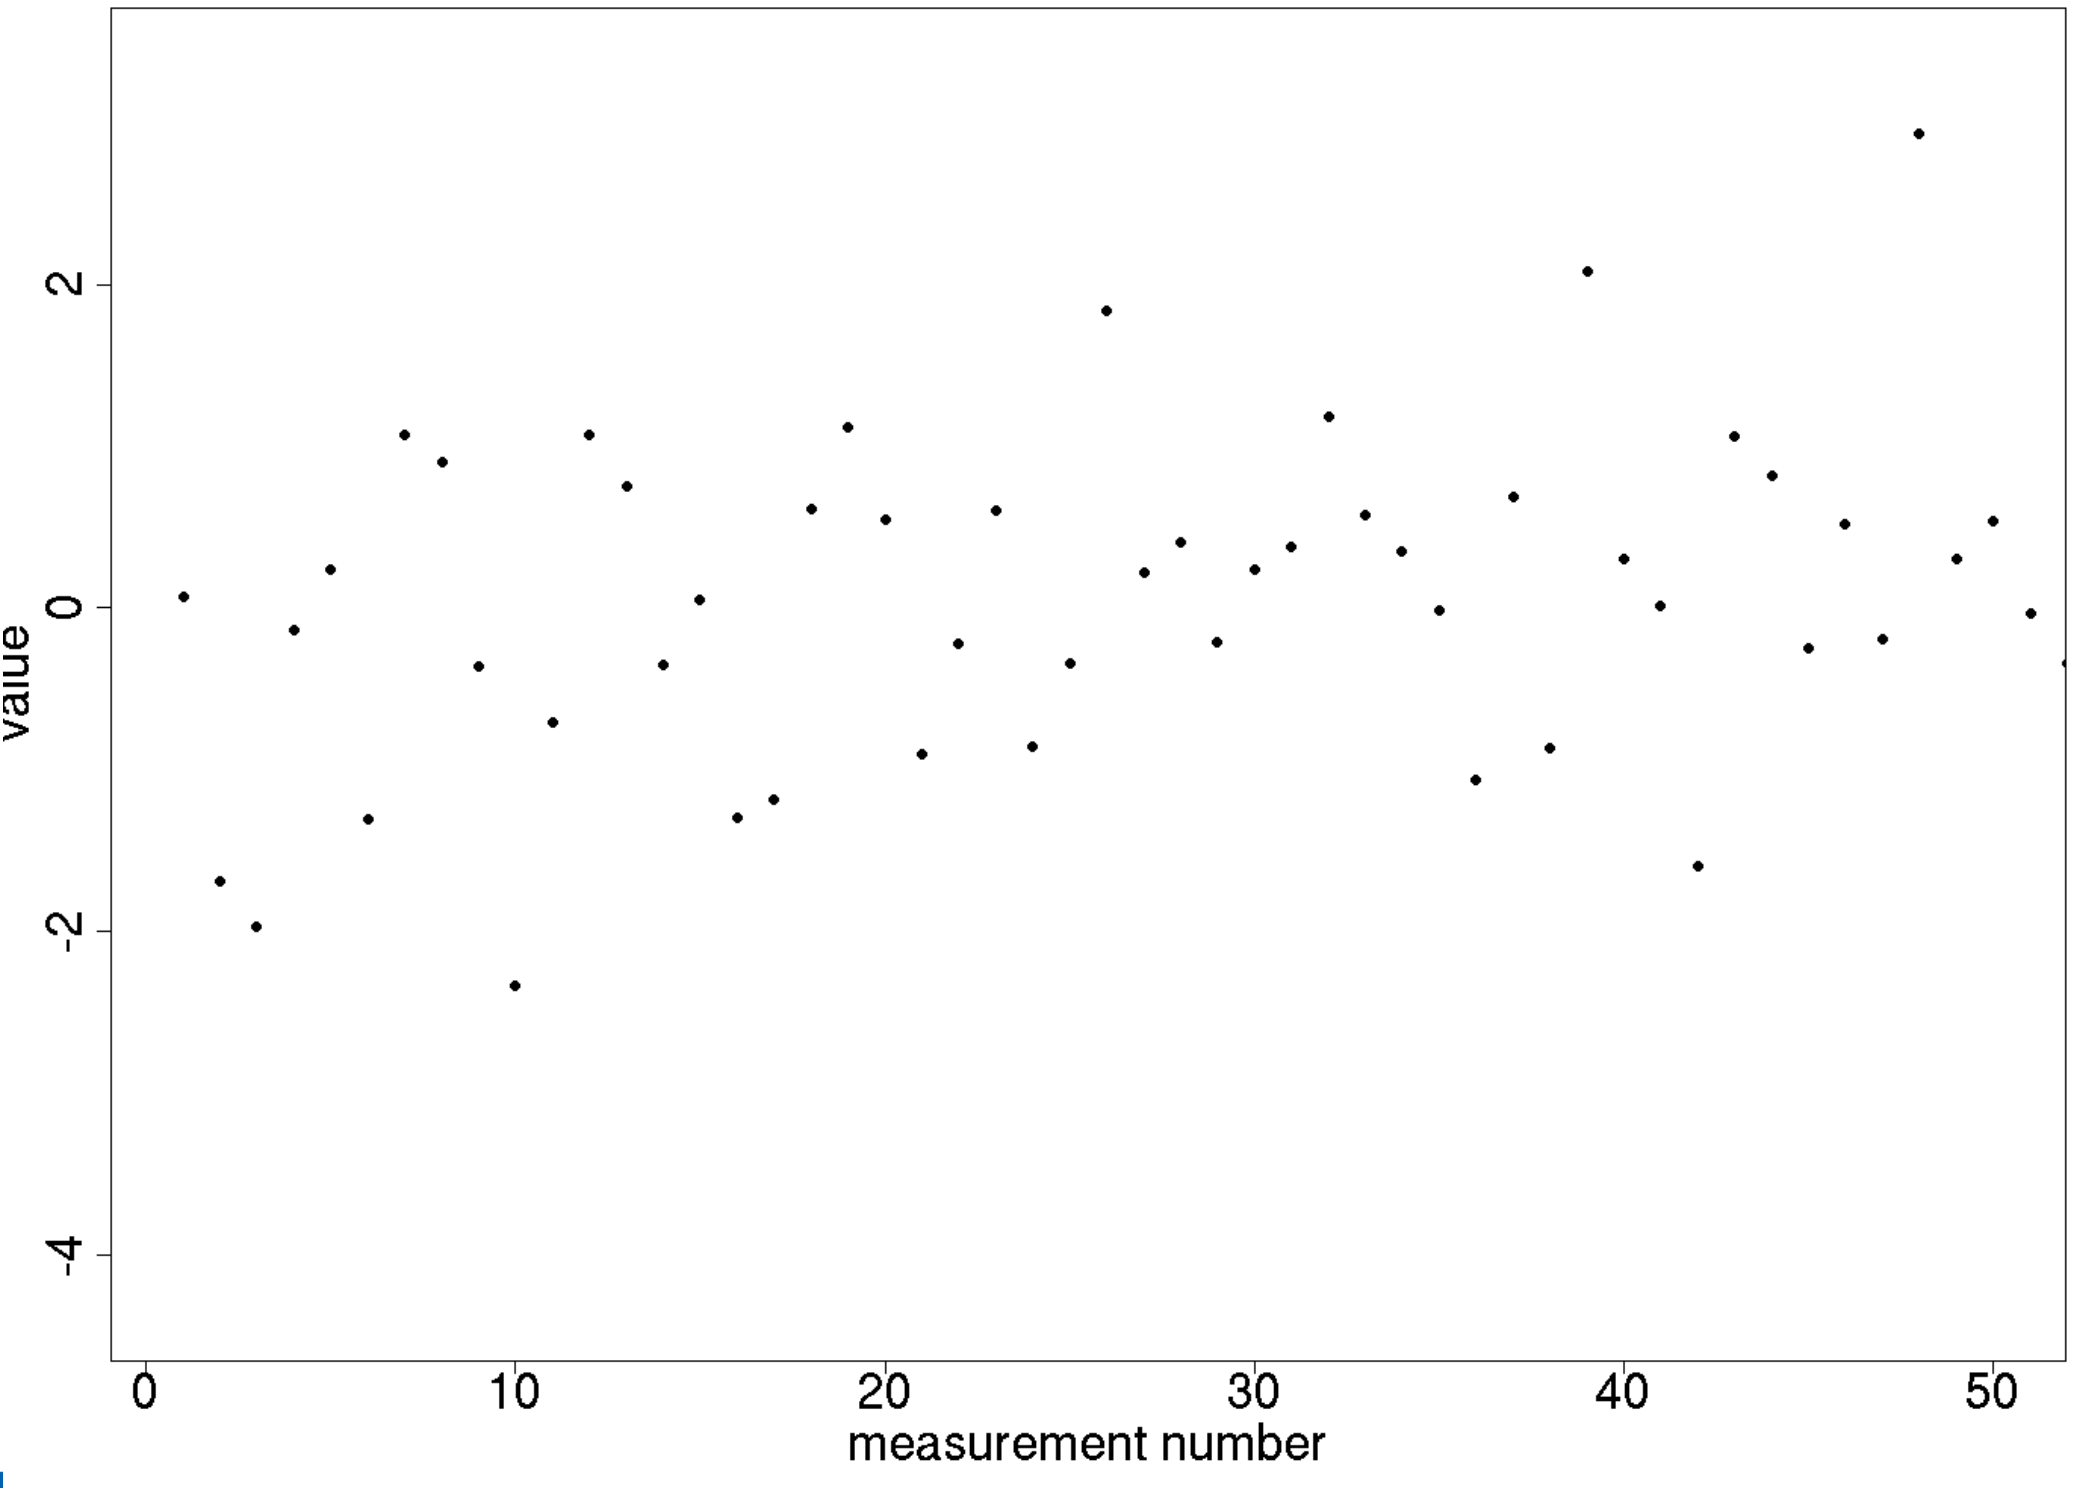
value
2
0
-2
-4
0
10
20
30
40
50
measurement number

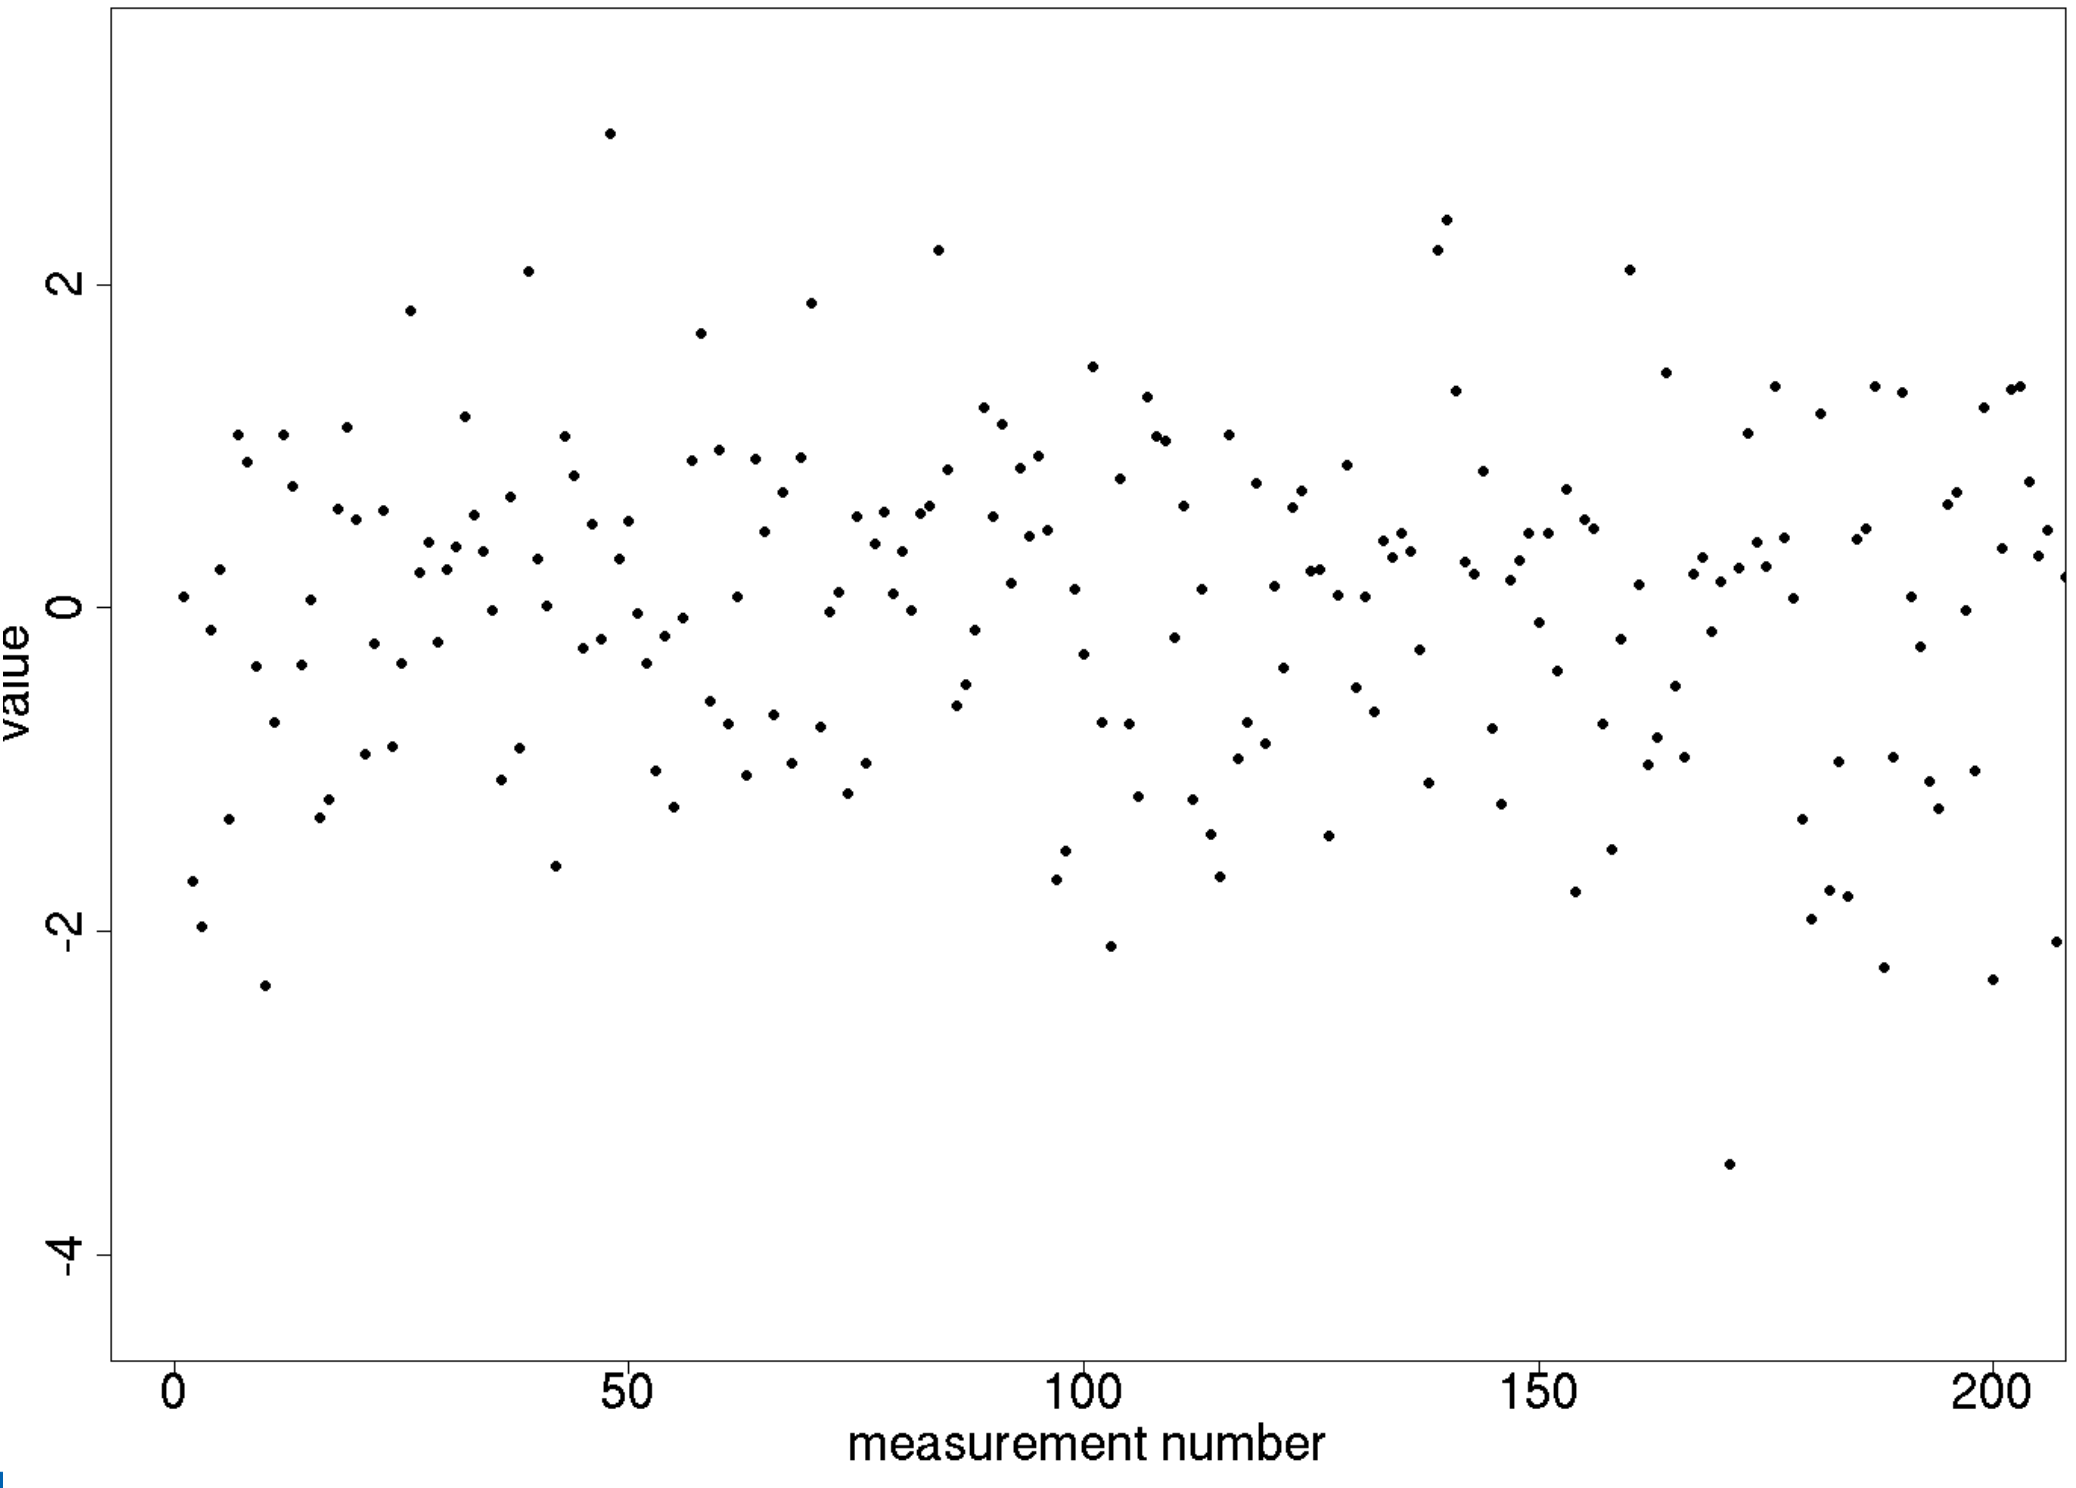
value
2
0
-2
-4
0
50
100
150
200
measurement number

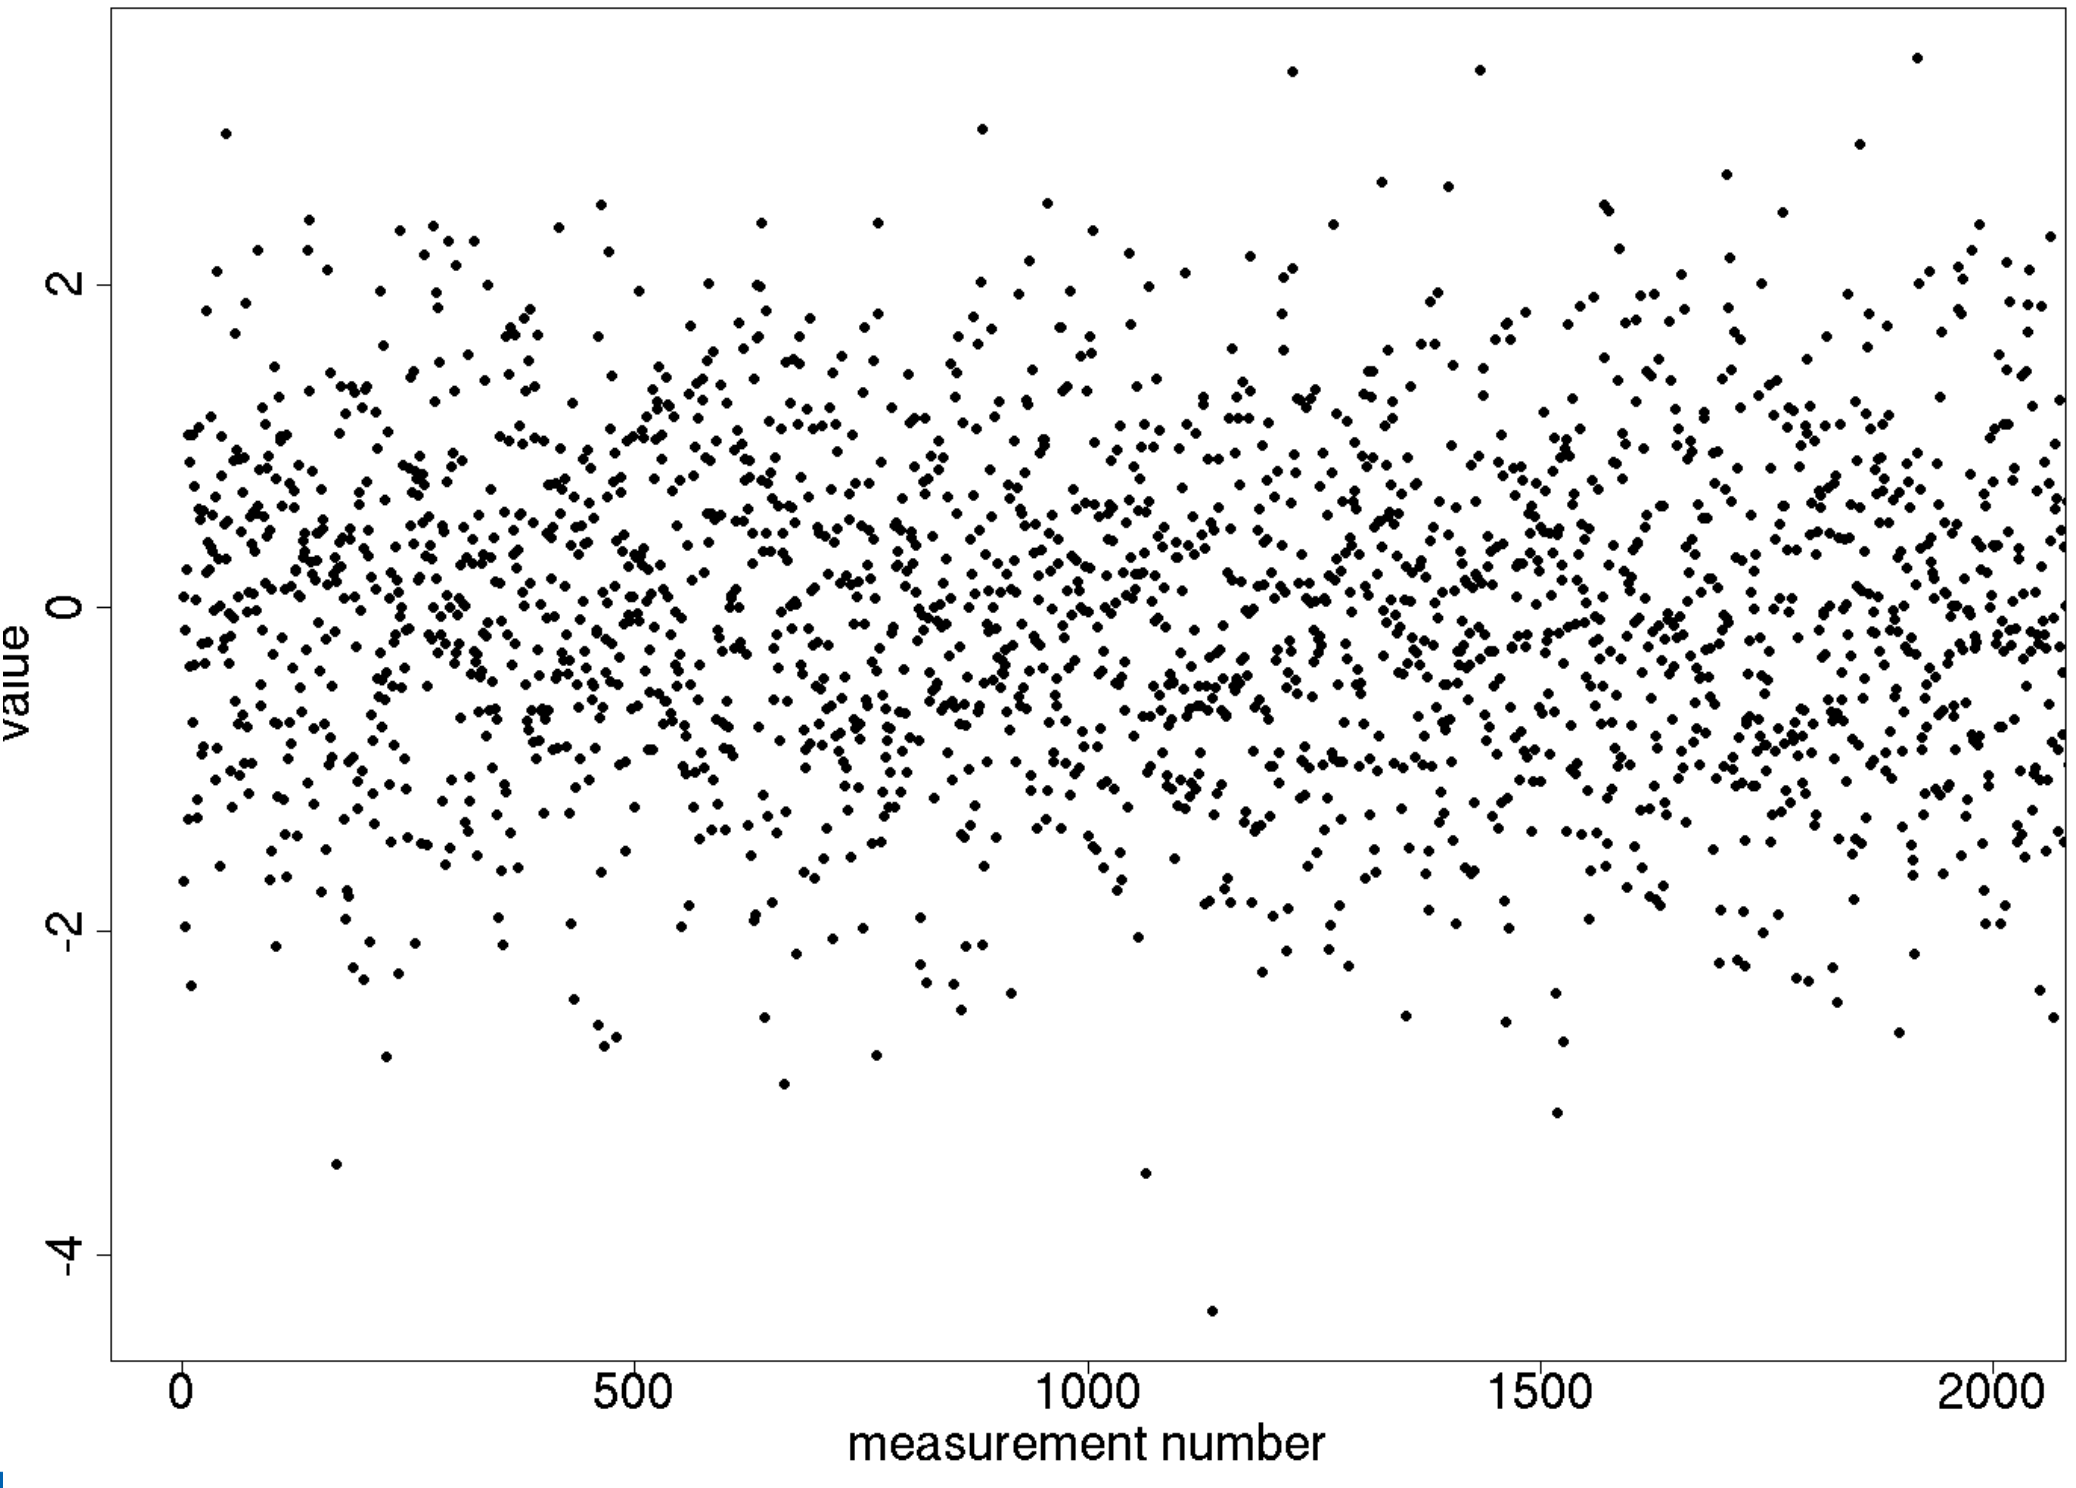
value
2
0
-2
-4
0
500
1000
1500
2000
measurement number

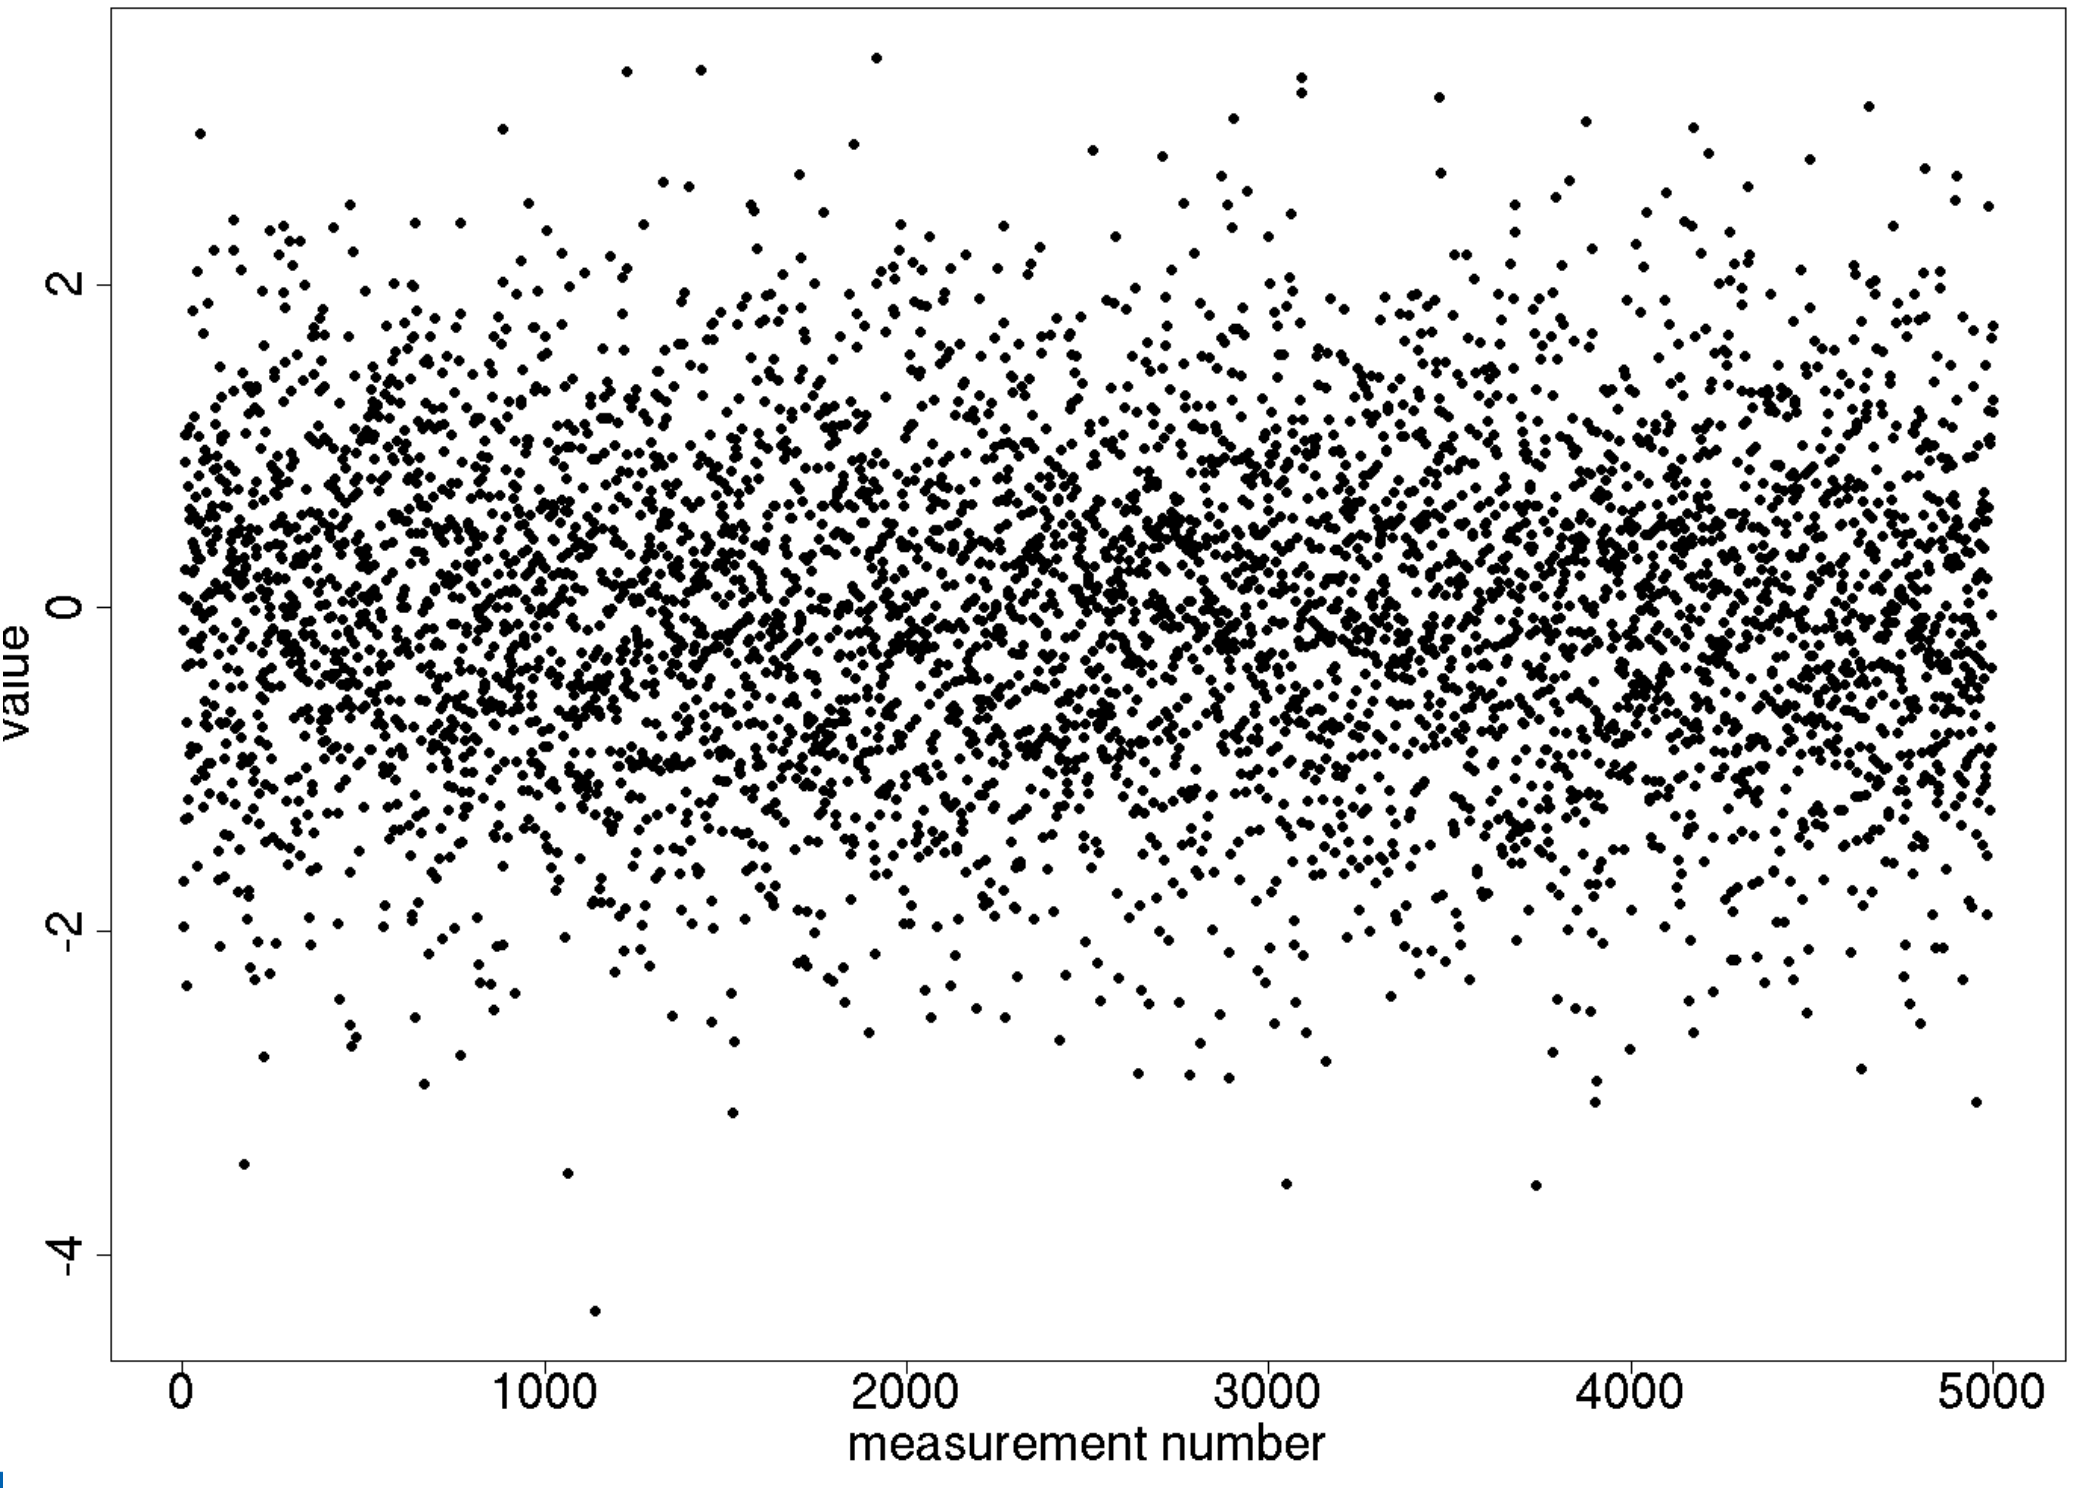
value
2
0
-2
-4
0
1000
2000
3000
4000
5000
measurement number

Value for each event

Large overhead to collect this data, thus rarely available.

Value for each event

Large overhead to collect this data, thus rarely available.

One compromise is to collect enough data to create a histogram.

The next slides are the same data as before, presented in histogram format.

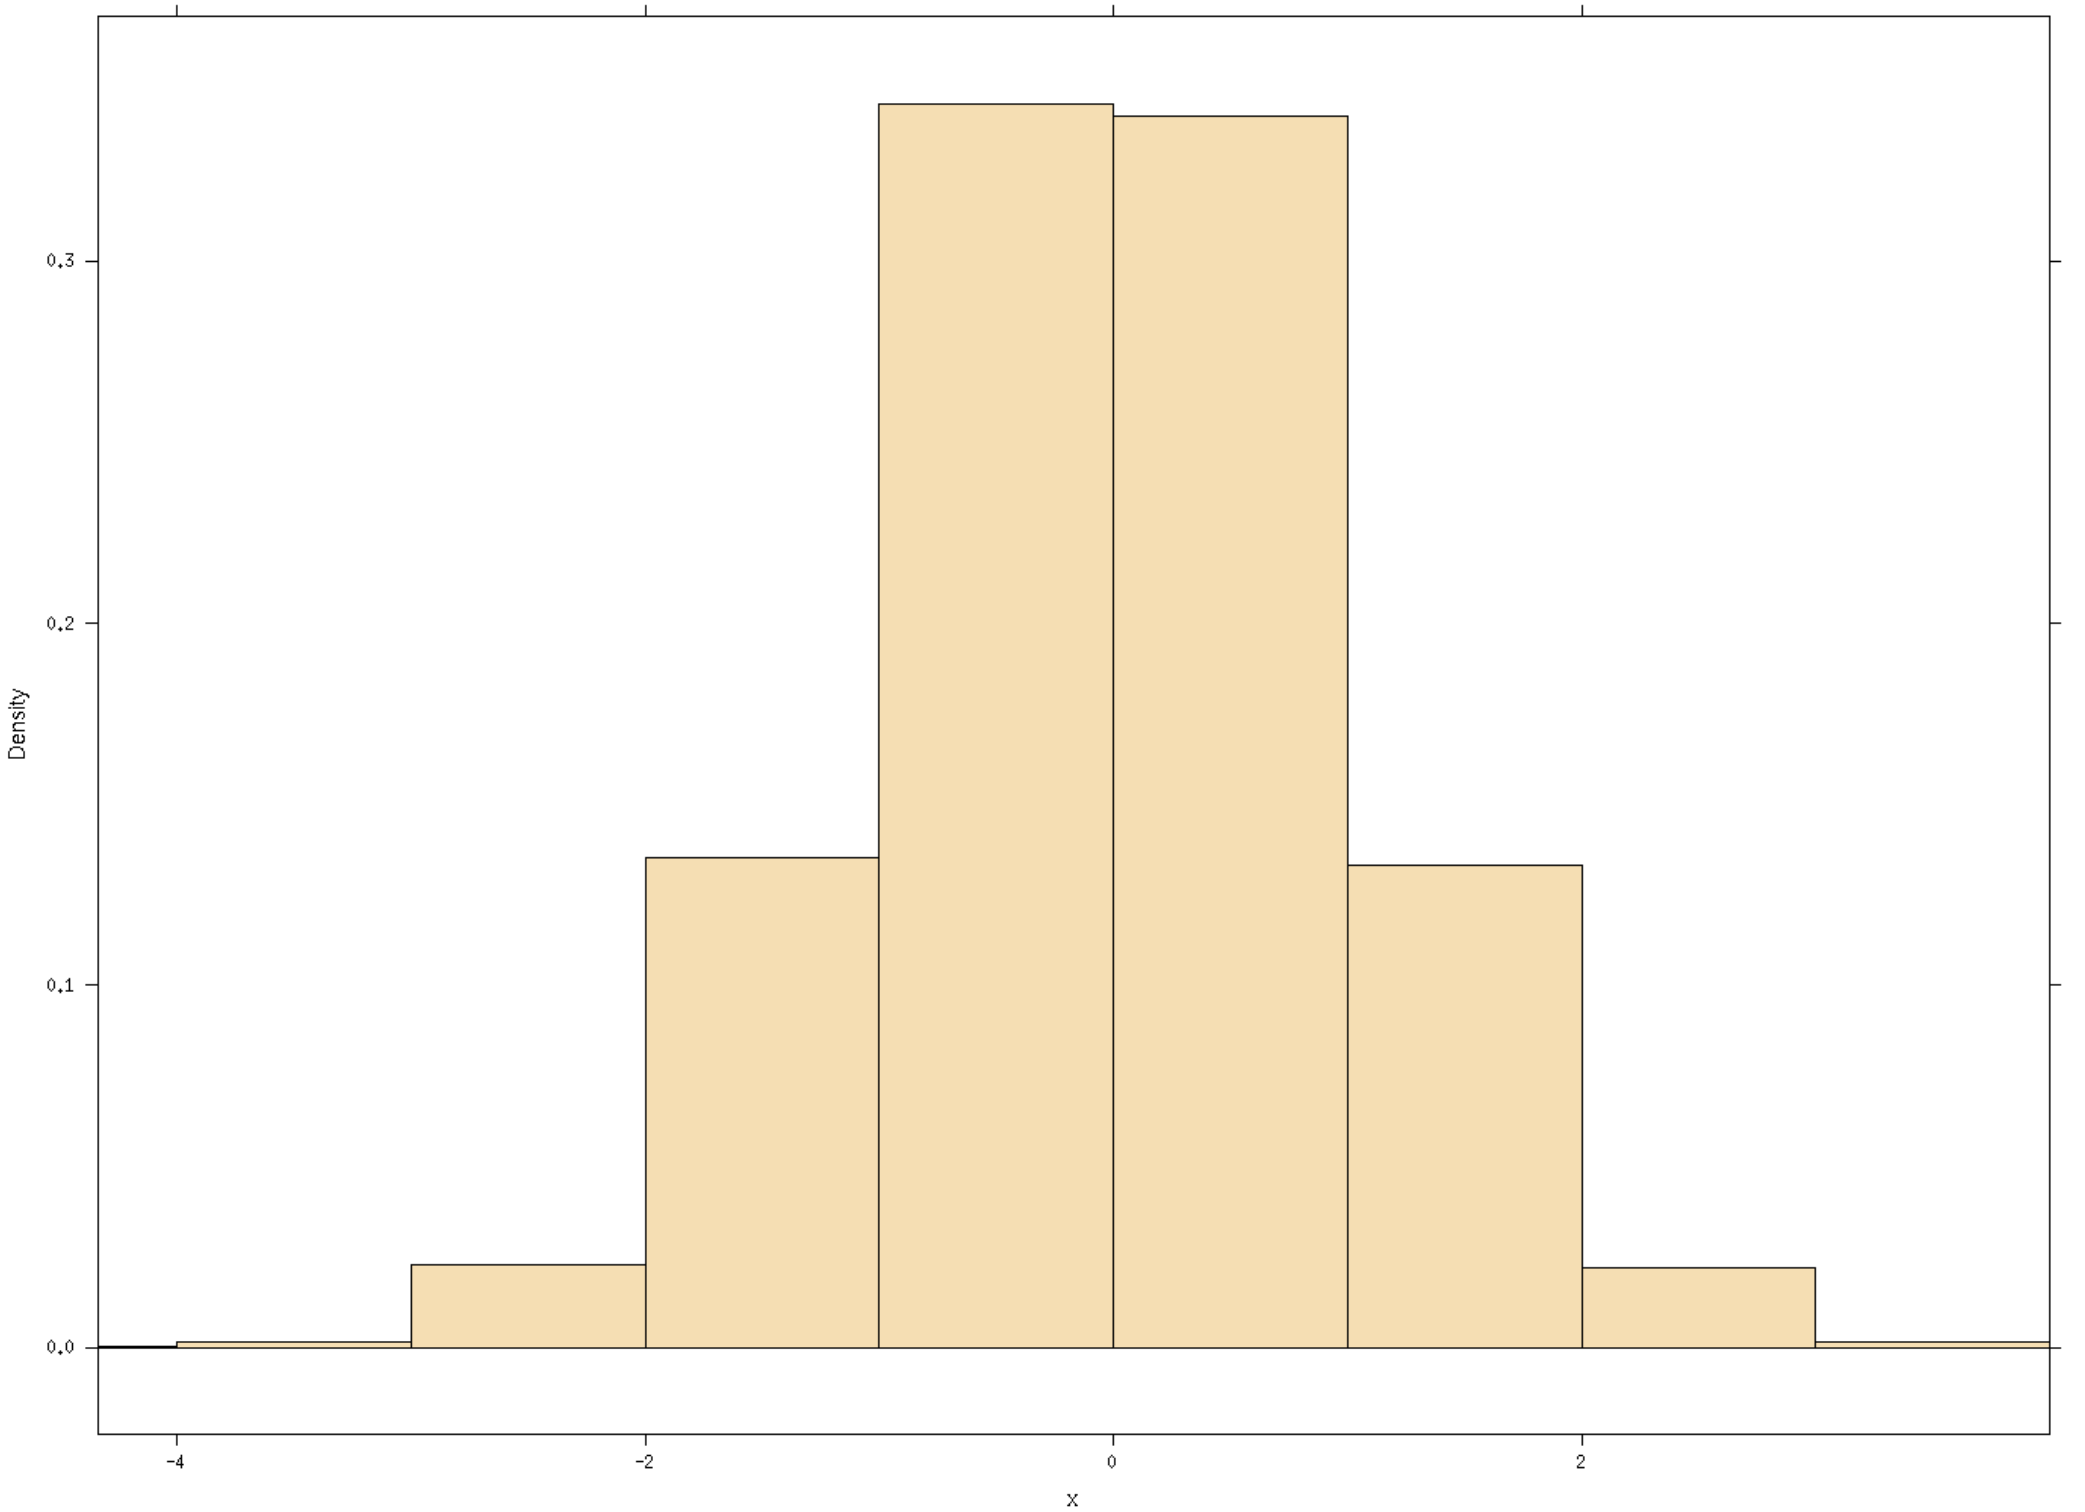
0.3
0.2
0.1
0.0
Density
-4
-2
0
2
x

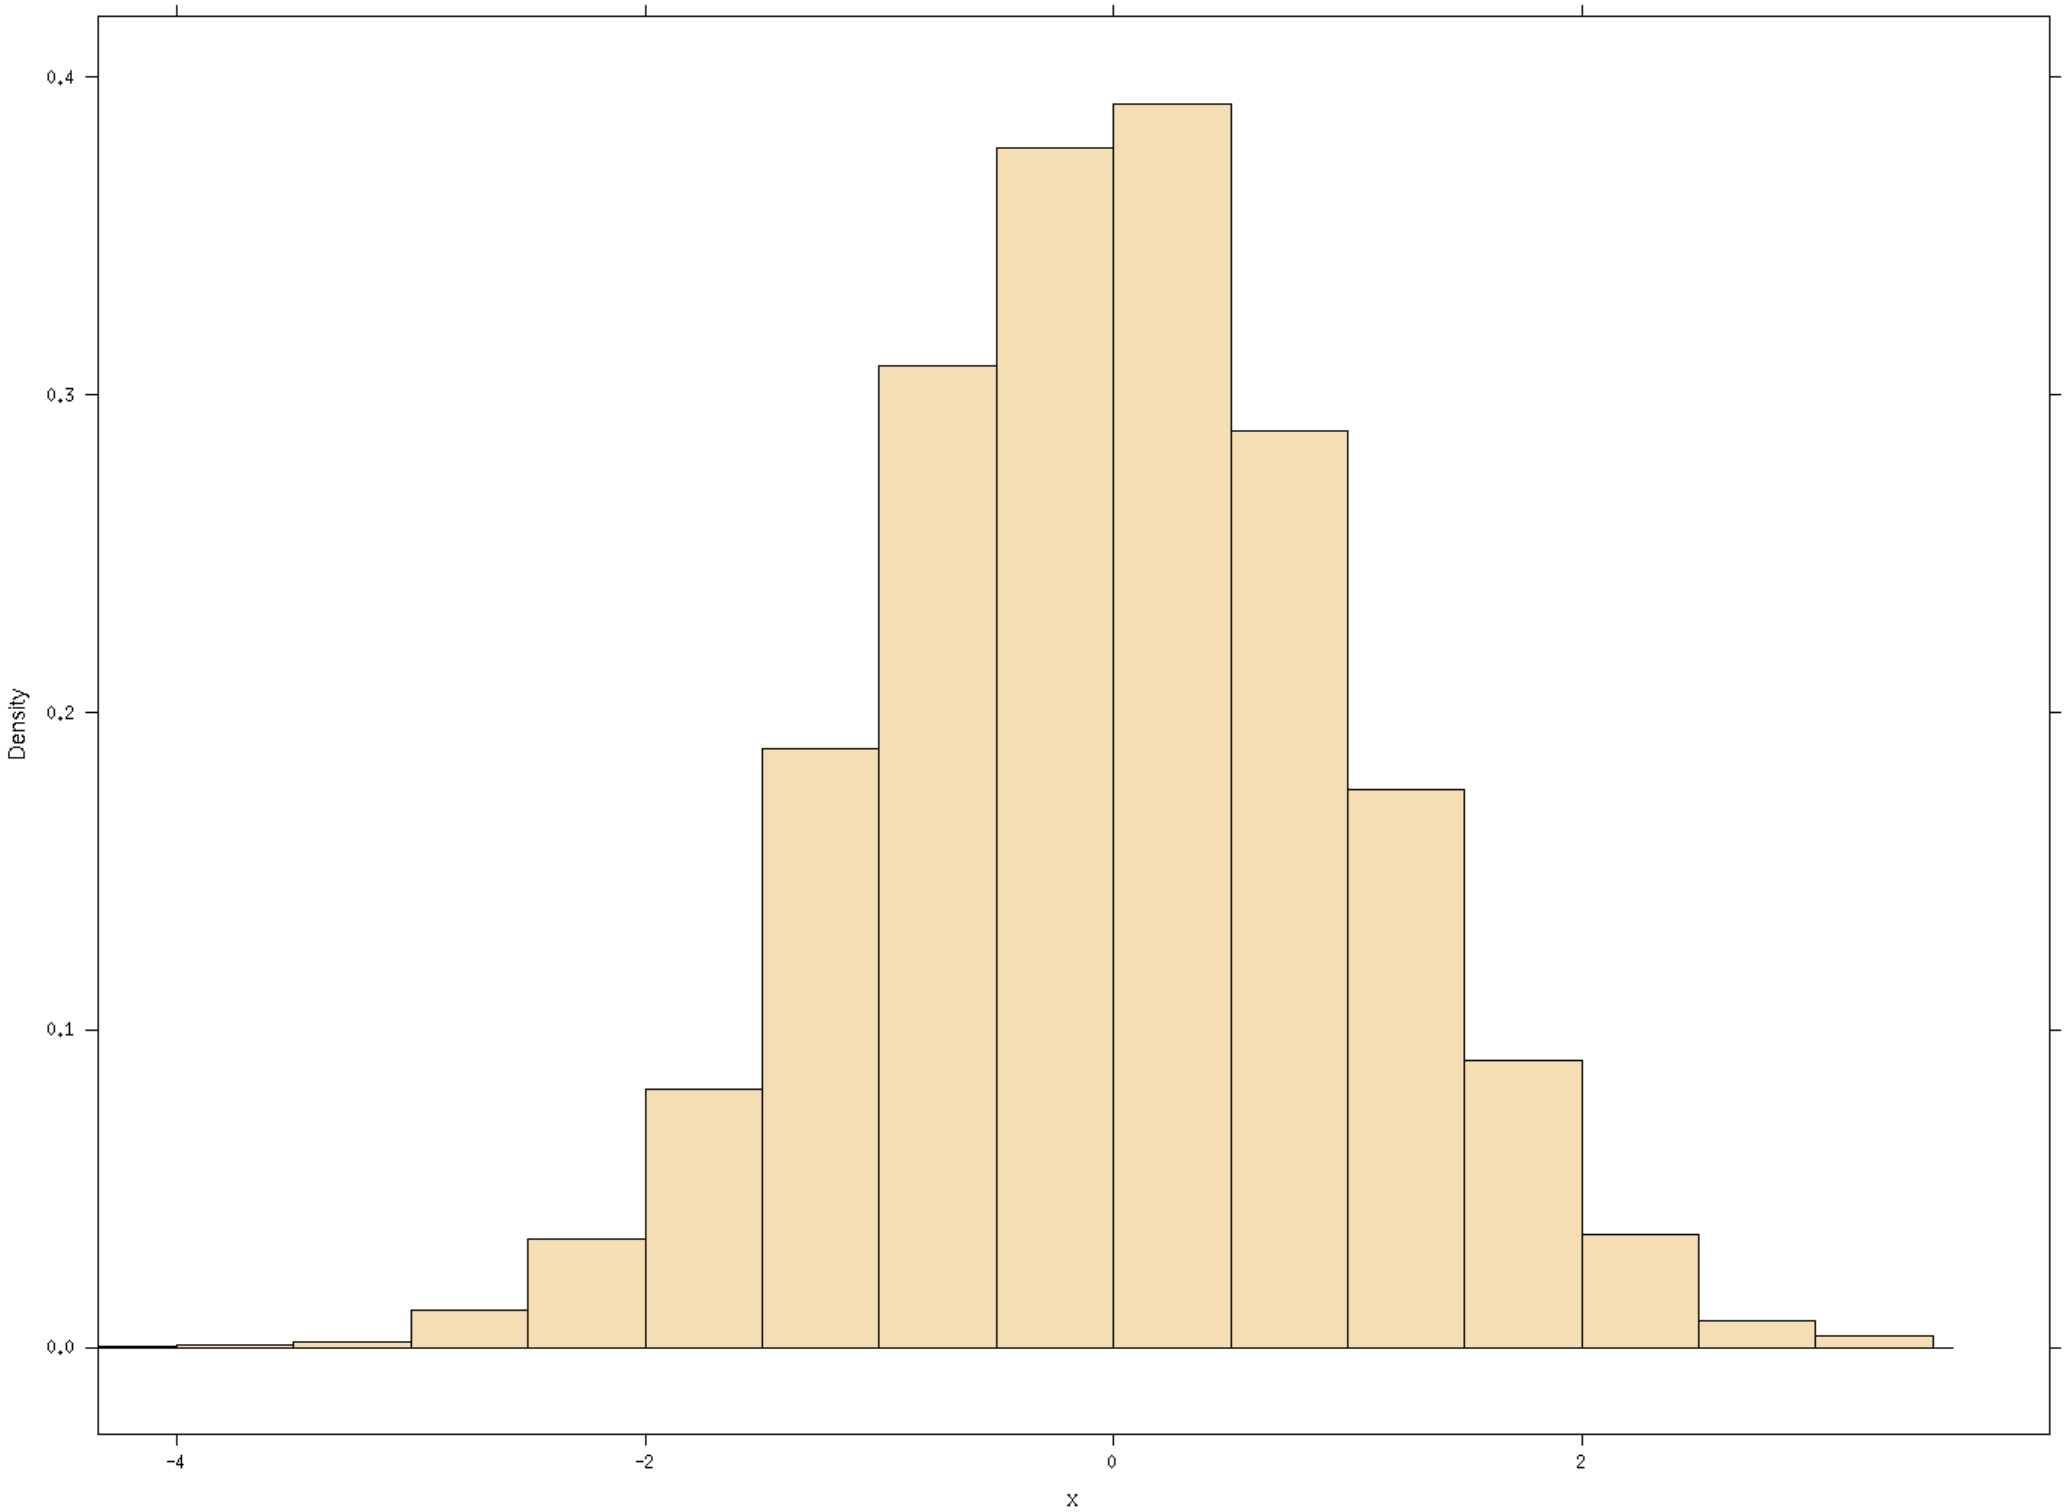
0.4
0.3
0.2
0.1
0.0
Density
-4
-2
0
2
x

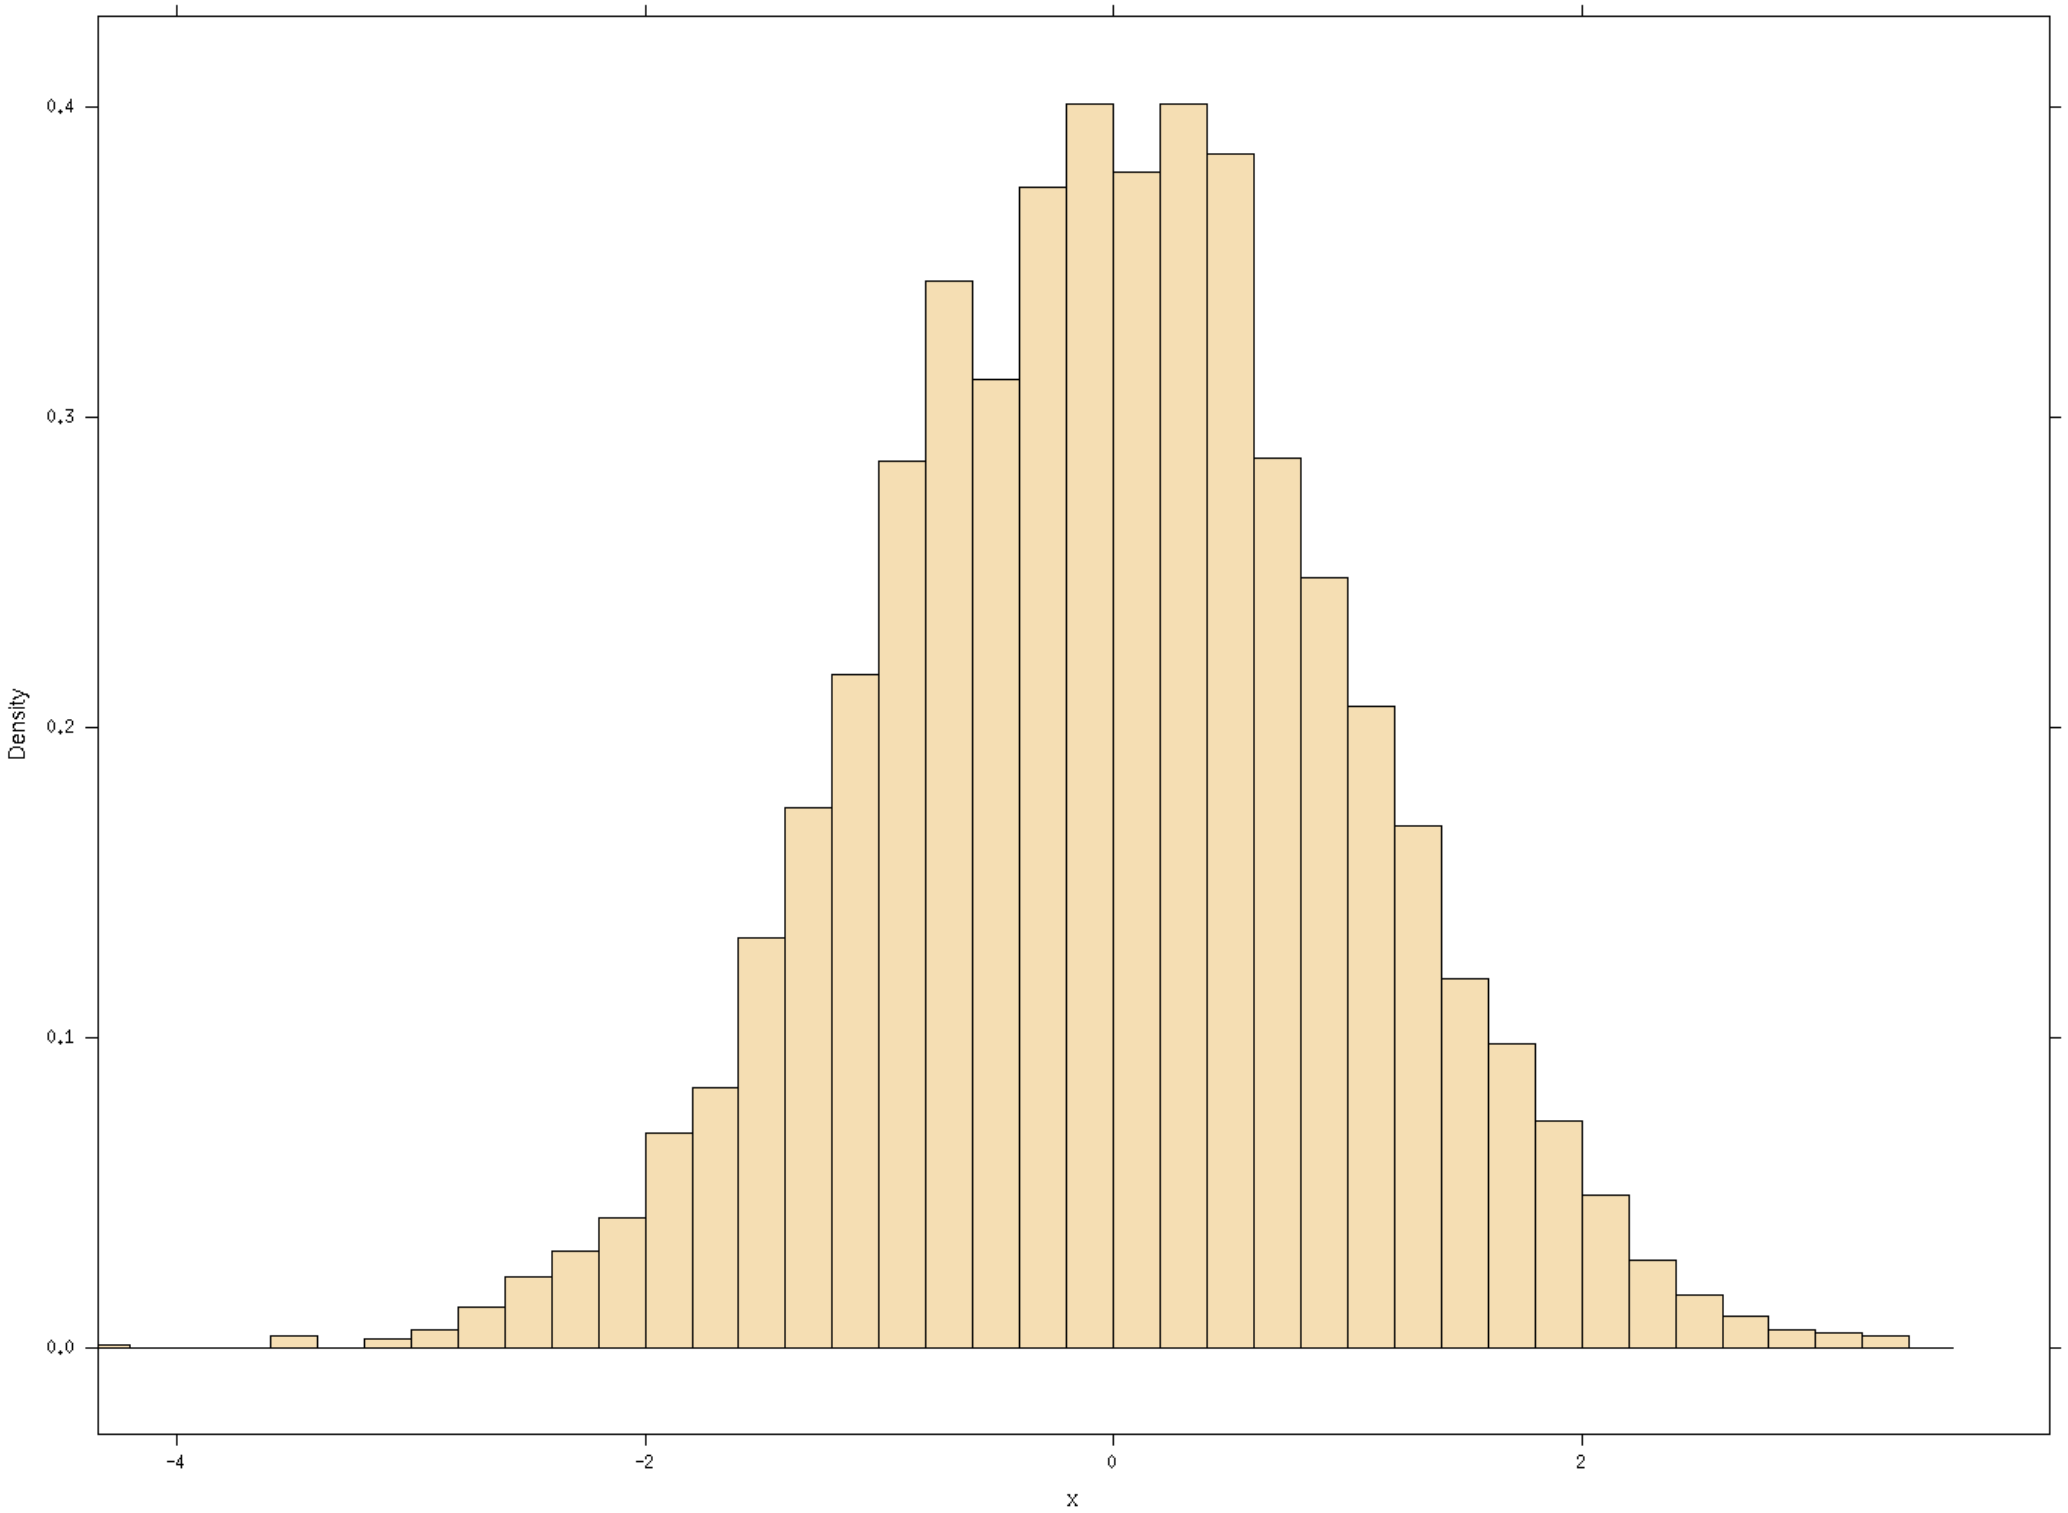
0.4
0.3
0.2
0.1
0.0
Density
-4
-2
0
2
x

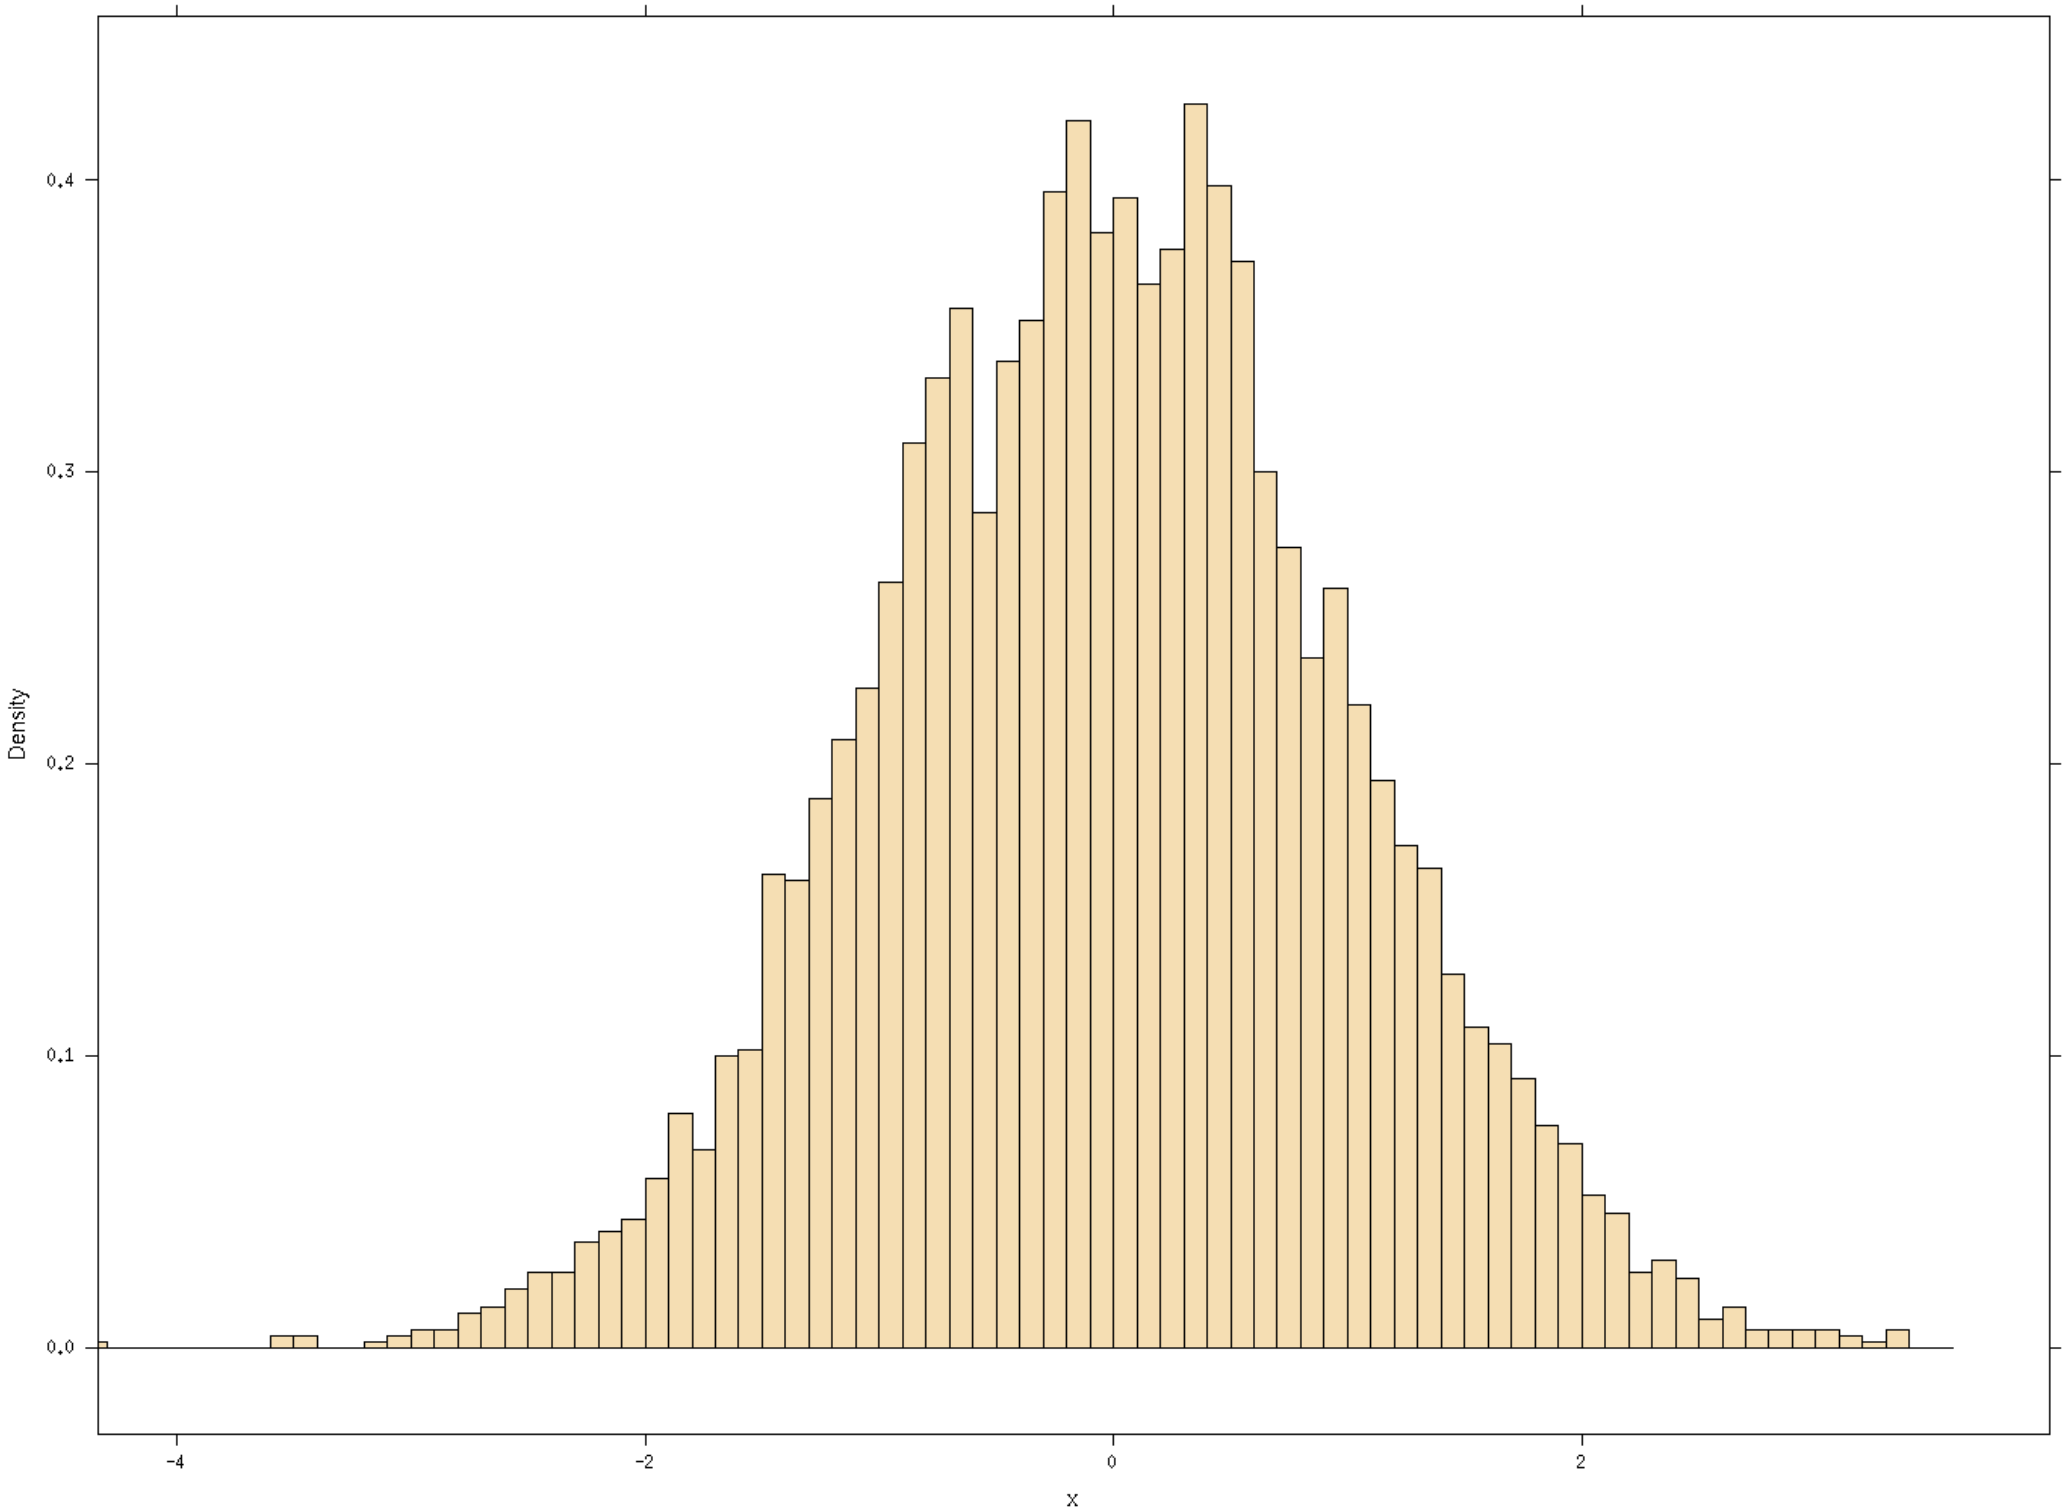
Density
0.0
0.1
0.2
0.3
0.4
-4
-2
0
2
x

standard deviation – What is it?

If the data was a perfect “normal distribution”, it would be described by the red line on the next slide.

The starting point for many statistics analysis is “assume a normal distribution”.

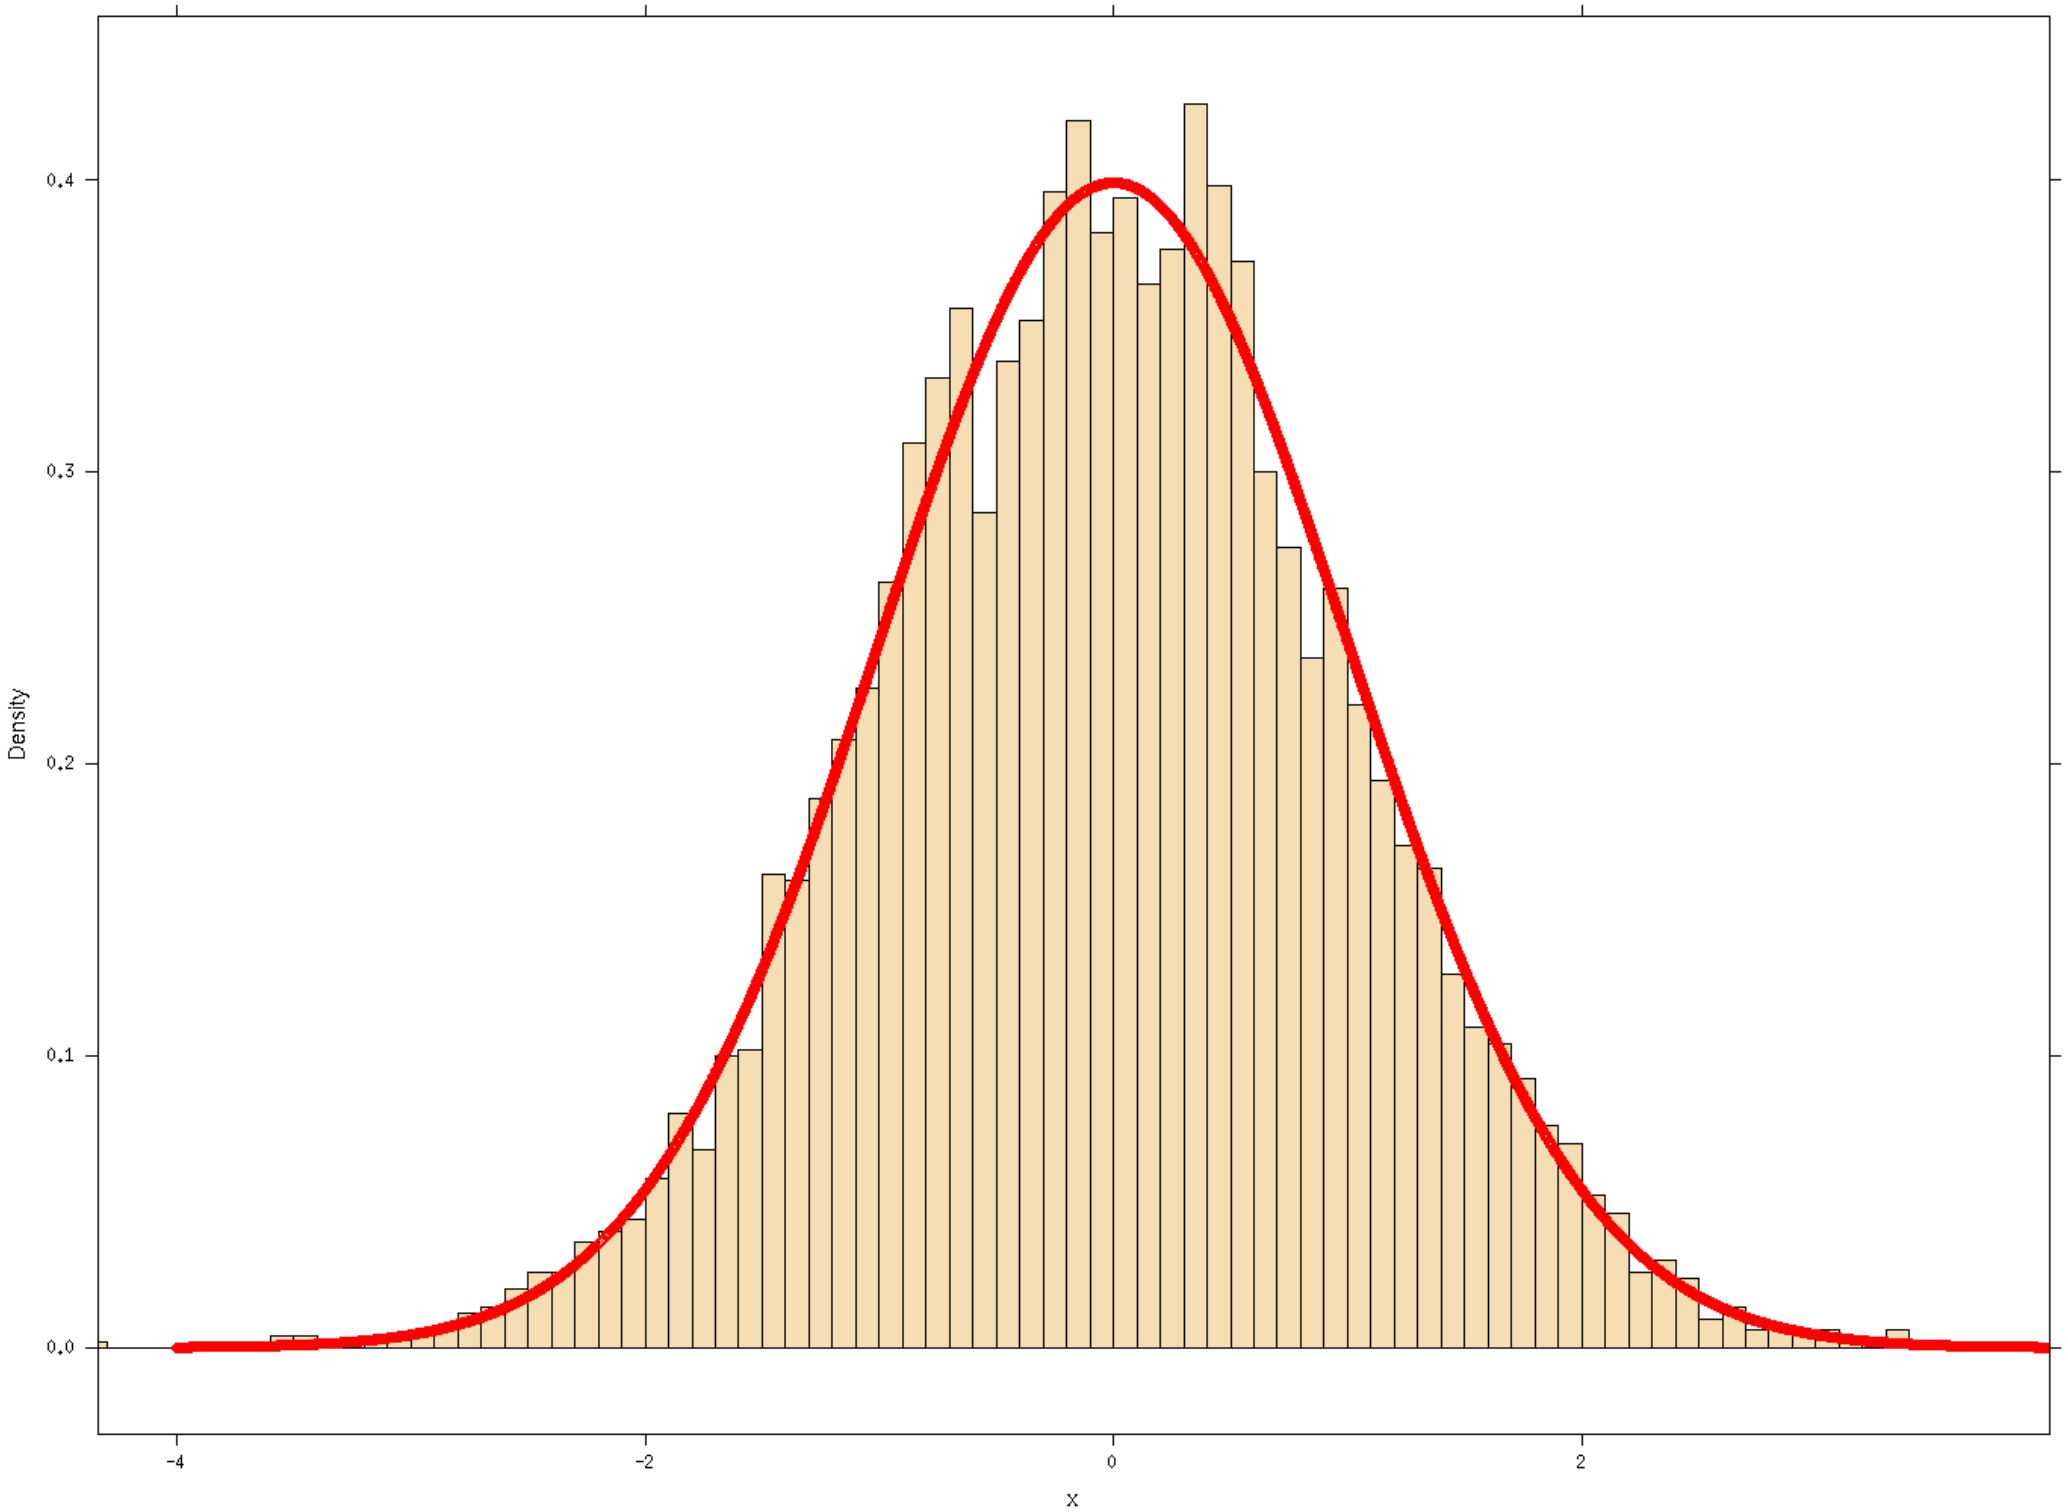
0.4
0.3
0.2
0.1
0.0
Density
-4
-2
0
2
x

standard deviation - What is it?

Intuitive description

standard deviation - What is it?

Intuitive description

> A measure of the width of a frequency distribution.

A smaller standard deviation means the distribution of the data points is narrower.

A smaller standard deviation means the distribution of the data points is closer to the mean.

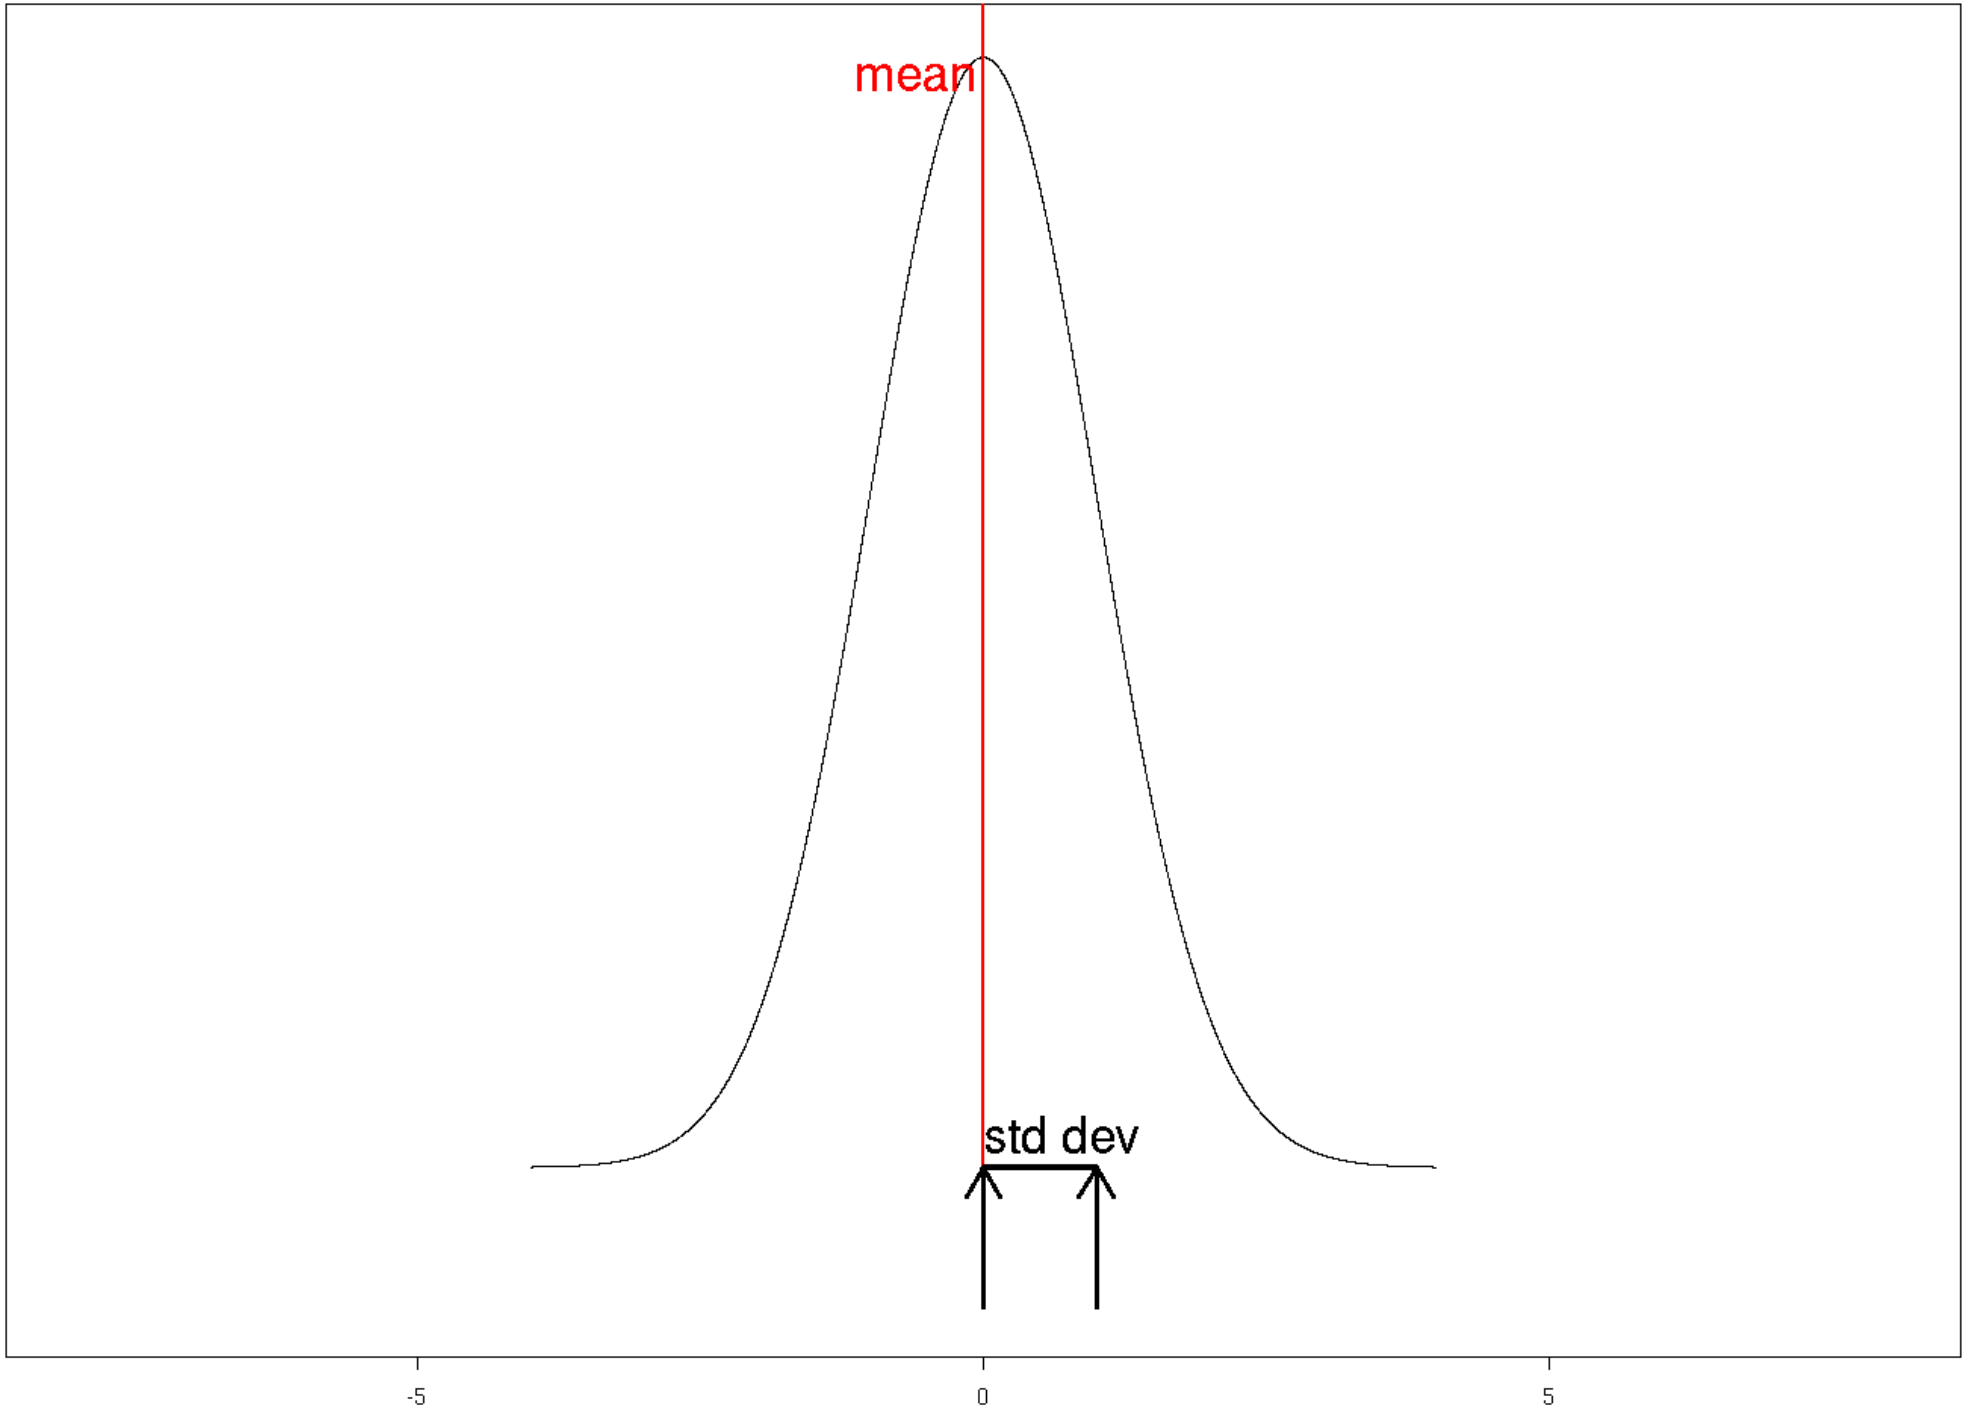
mean
std dev
-5
0
5

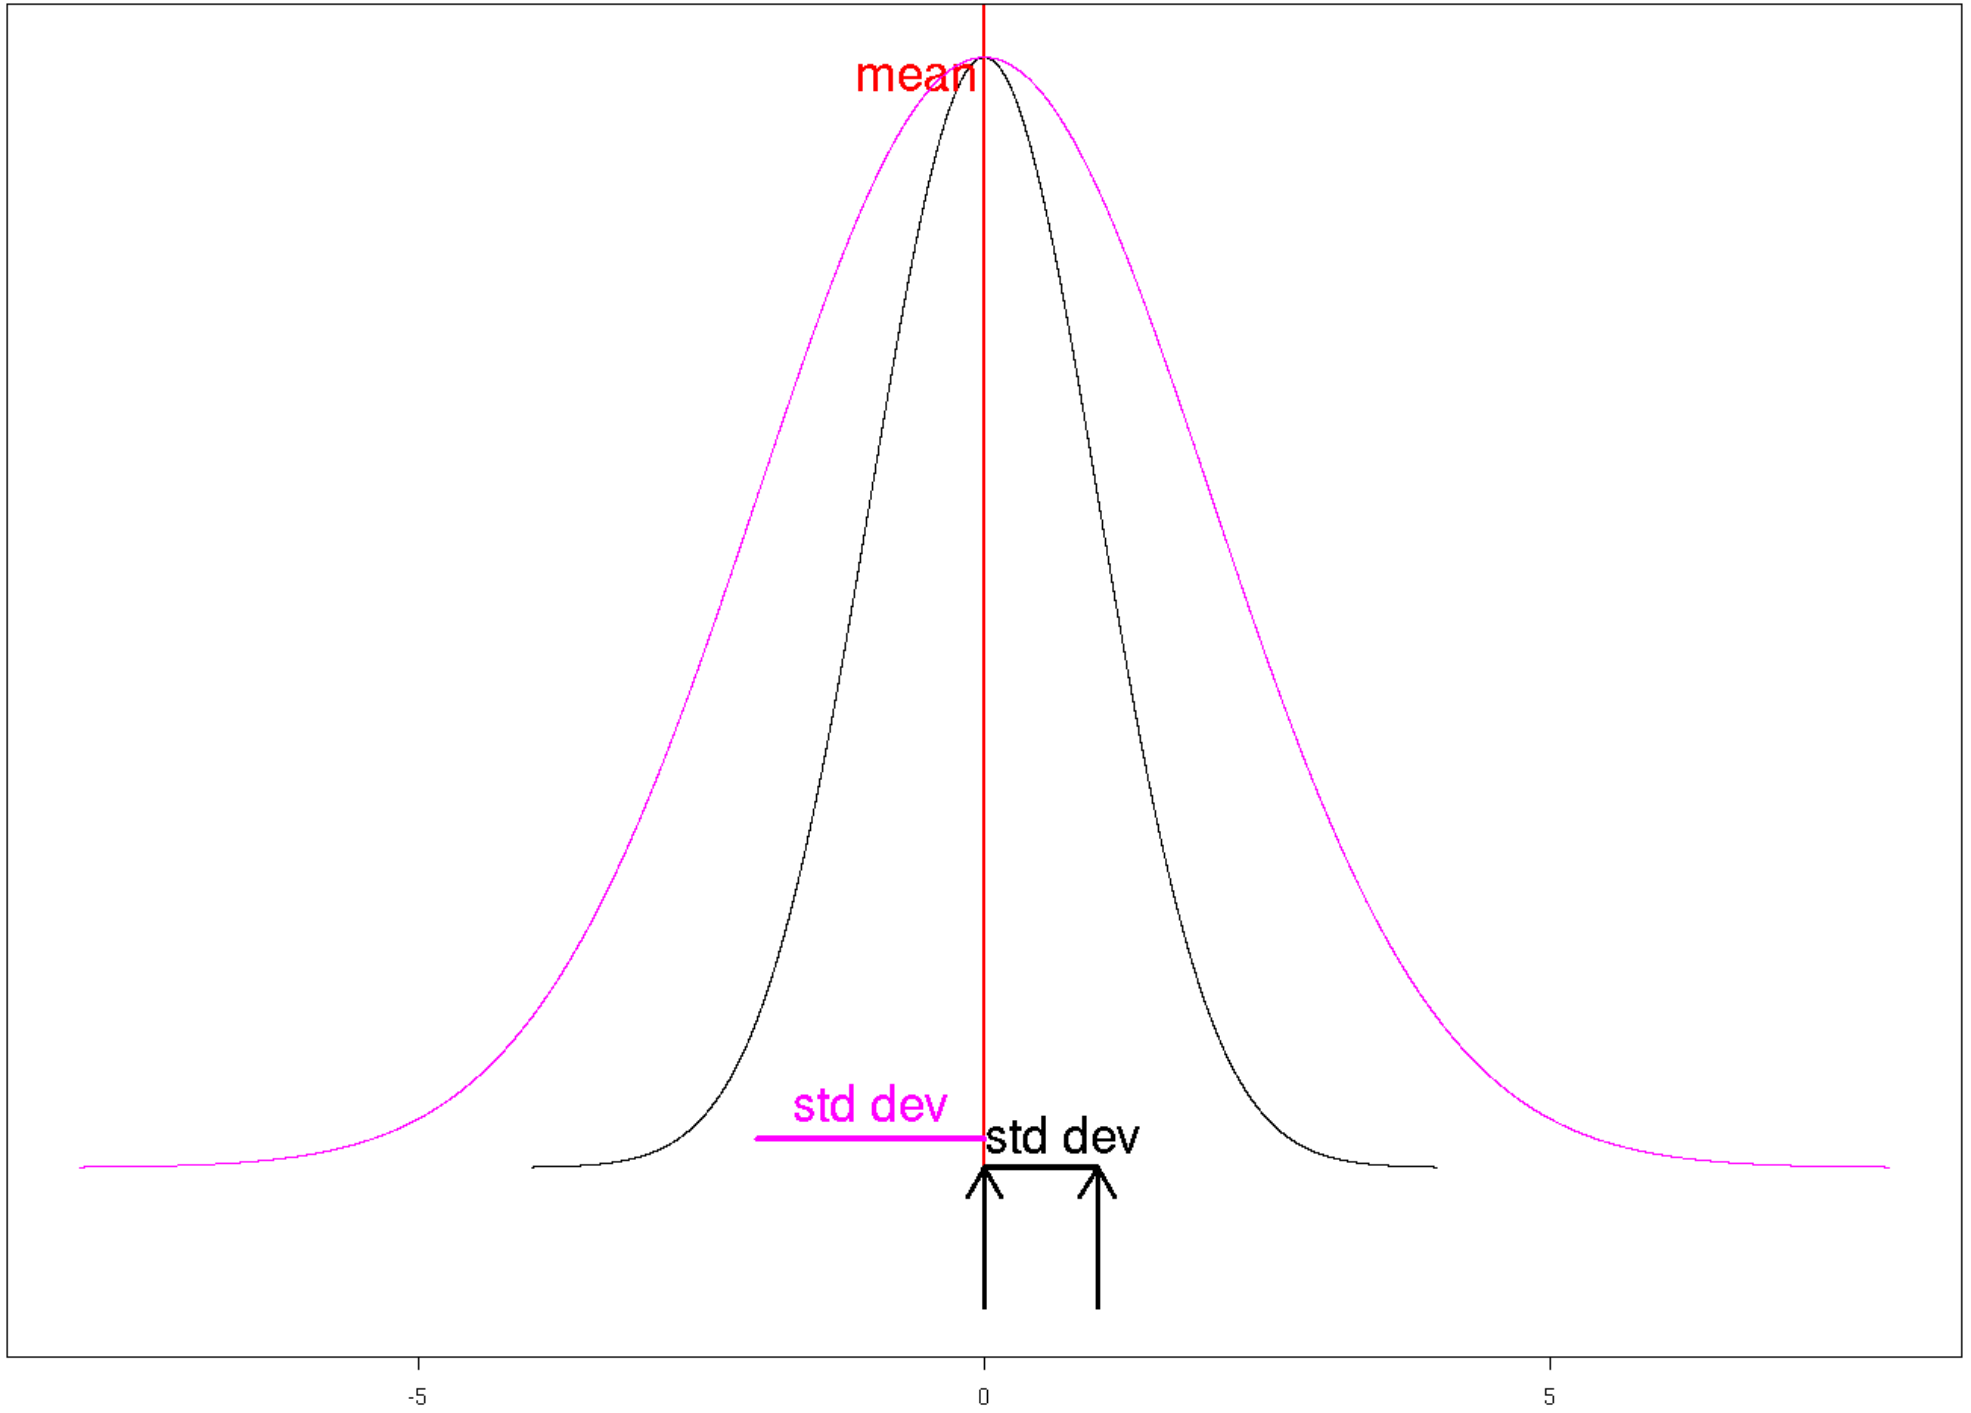
mean
std dev
std dev
-5
0
5

standard deviation - What is it?

Mathematical description of population standard deviation

$$\sigma = \sqrt{\frac{1}{n-1}\sum_{i=1}^{n}\left(x_i - \bar{X}\right)^2}$$

n *is number of samples*

x_i *are the individual samples*

$\bar{X}$ *is the population mean*

Standard Deviation

Is it a meaningful metric for real-time data?

Standard Deviation

Is it a meaningful metric for real-time data?

Yes, if data is a normal distribution.

Less so, as a tail develops.

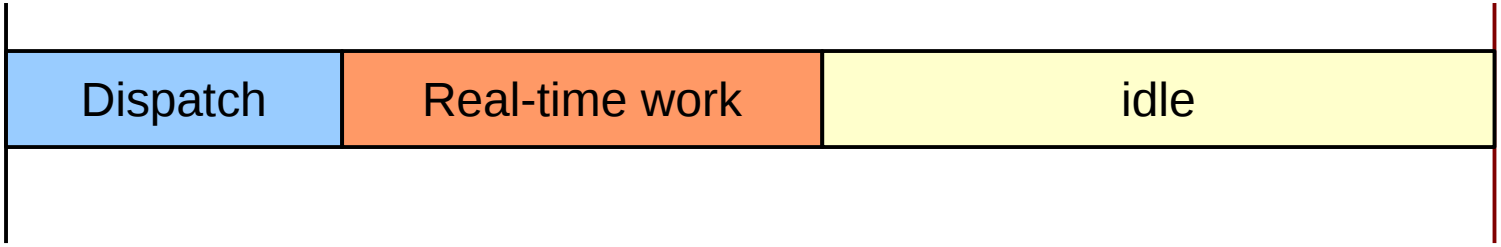

good

event

deadline

Recalling one of the reasons we might be analyzing data...

The real-time task must meet a deadline:

Is maximum latency + work < deadline?

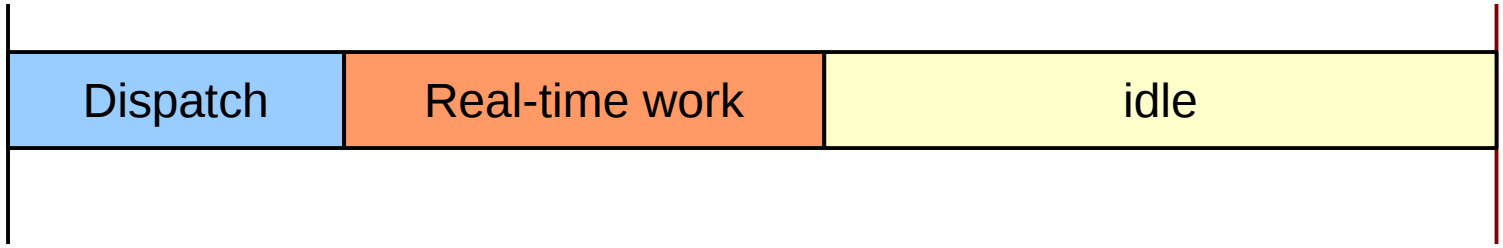

good

event

deadline

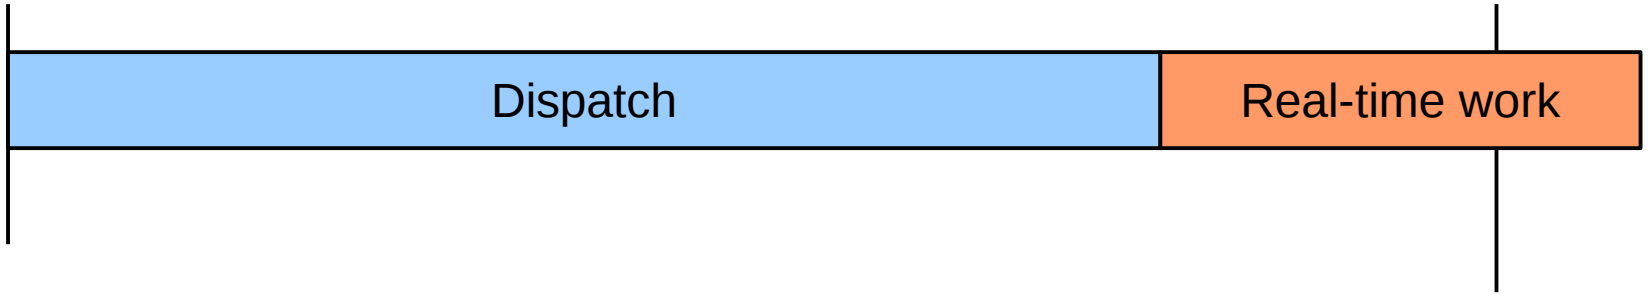

bad

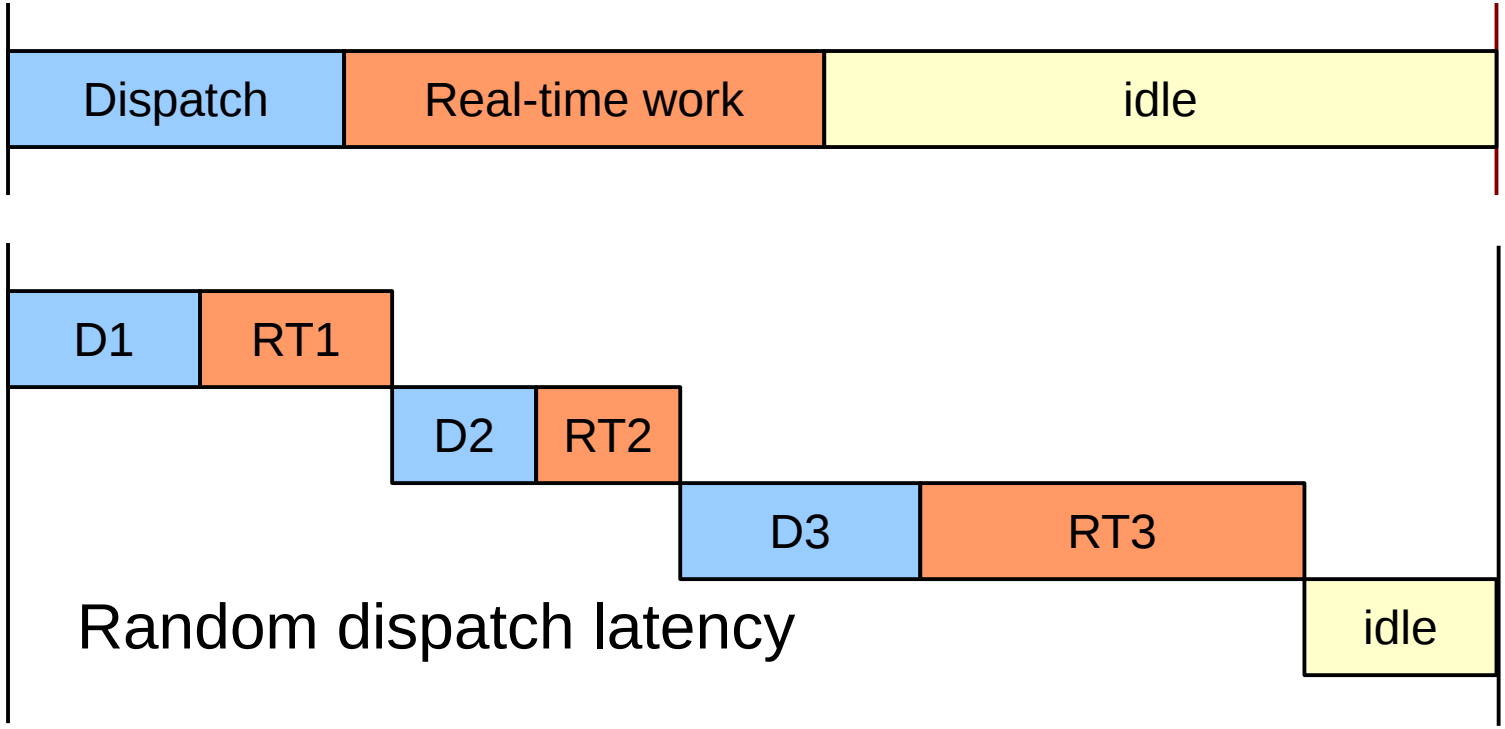

good

event

deadline

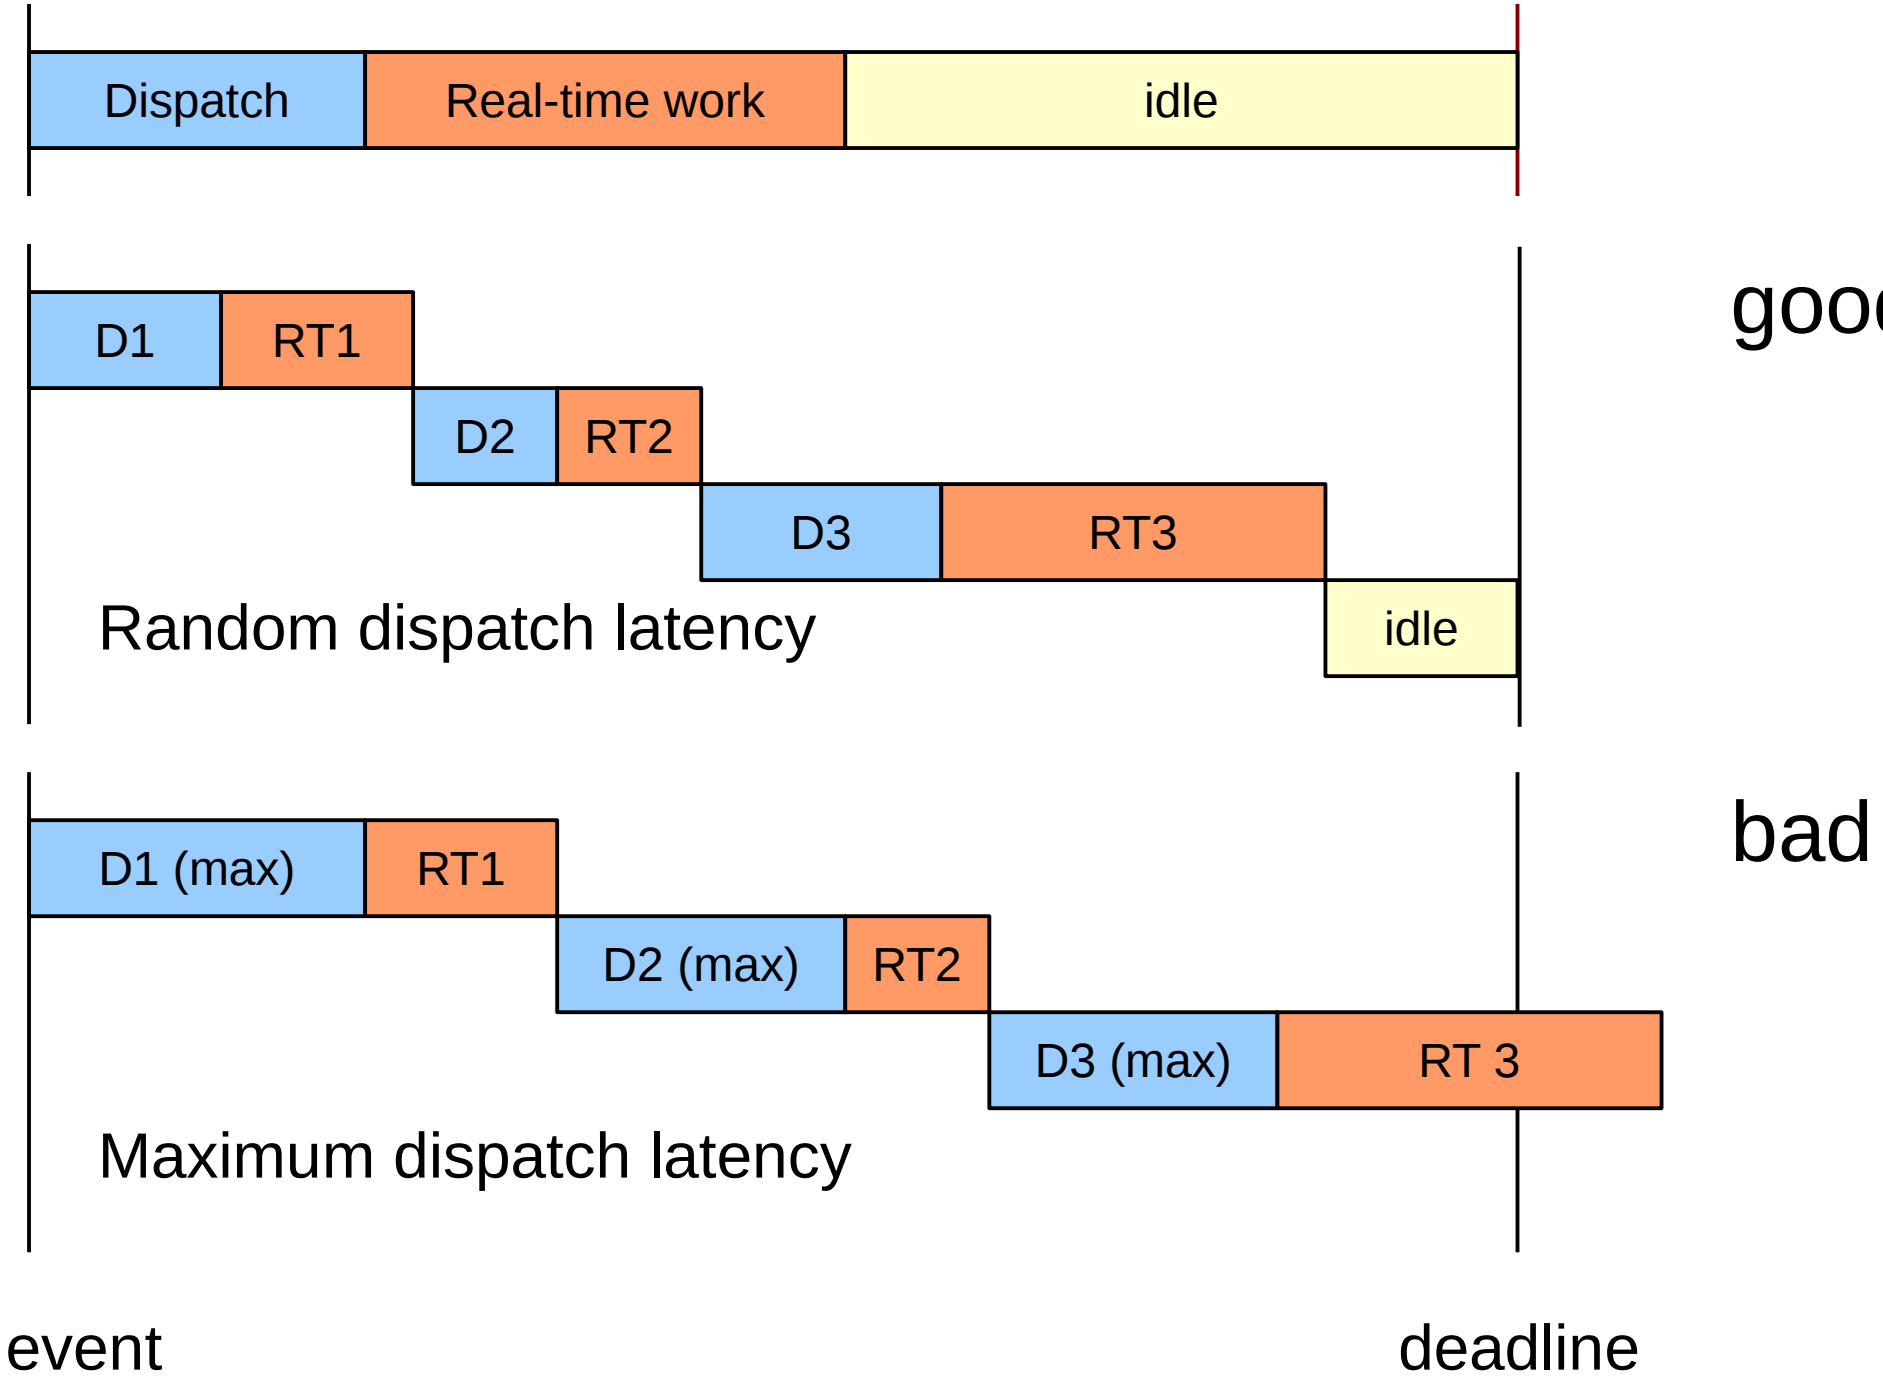
Dispatch
Real-time work
idle
D1
RT1
D2
RT2
D3
RT3
idle
Random dispatch latency
good
D1 (max)
RT1
D2 (max)
RT2
D3 (max)
RT 3
Maximum dispatch latency
bad
event
deadline

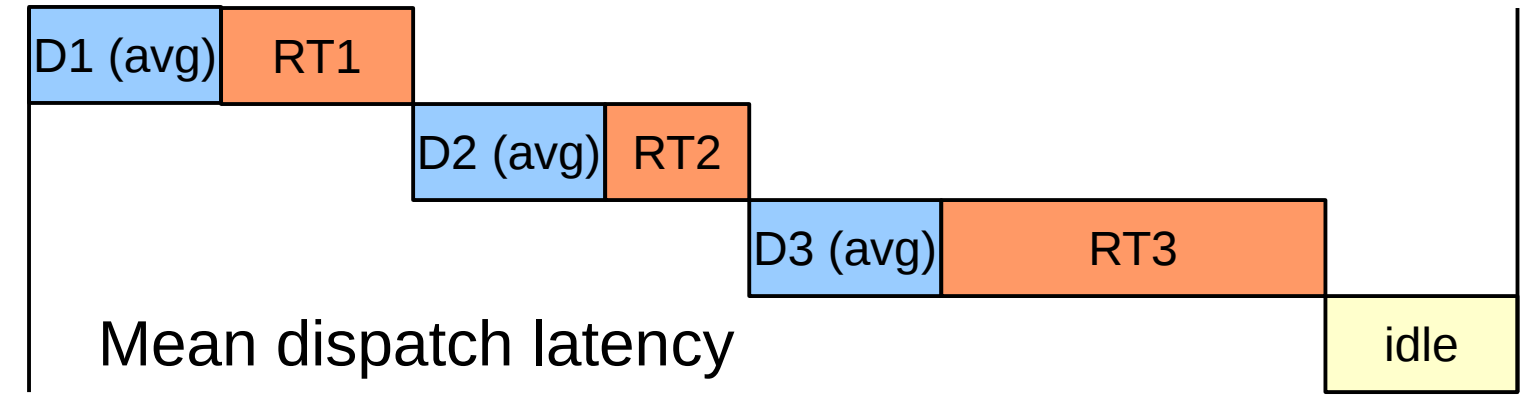

good

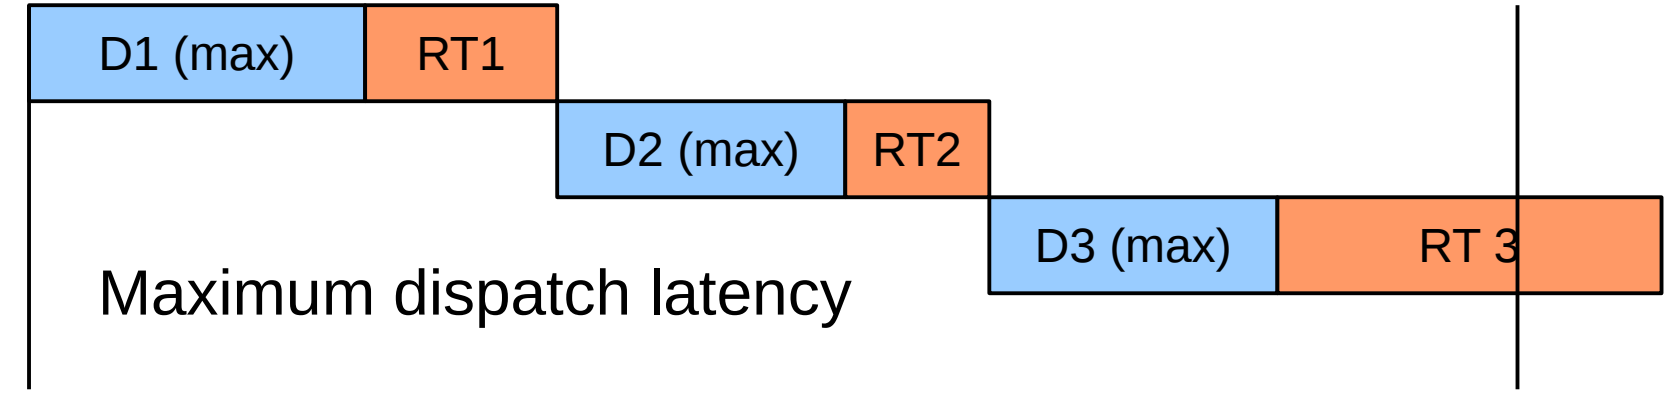

bad

event

deadline

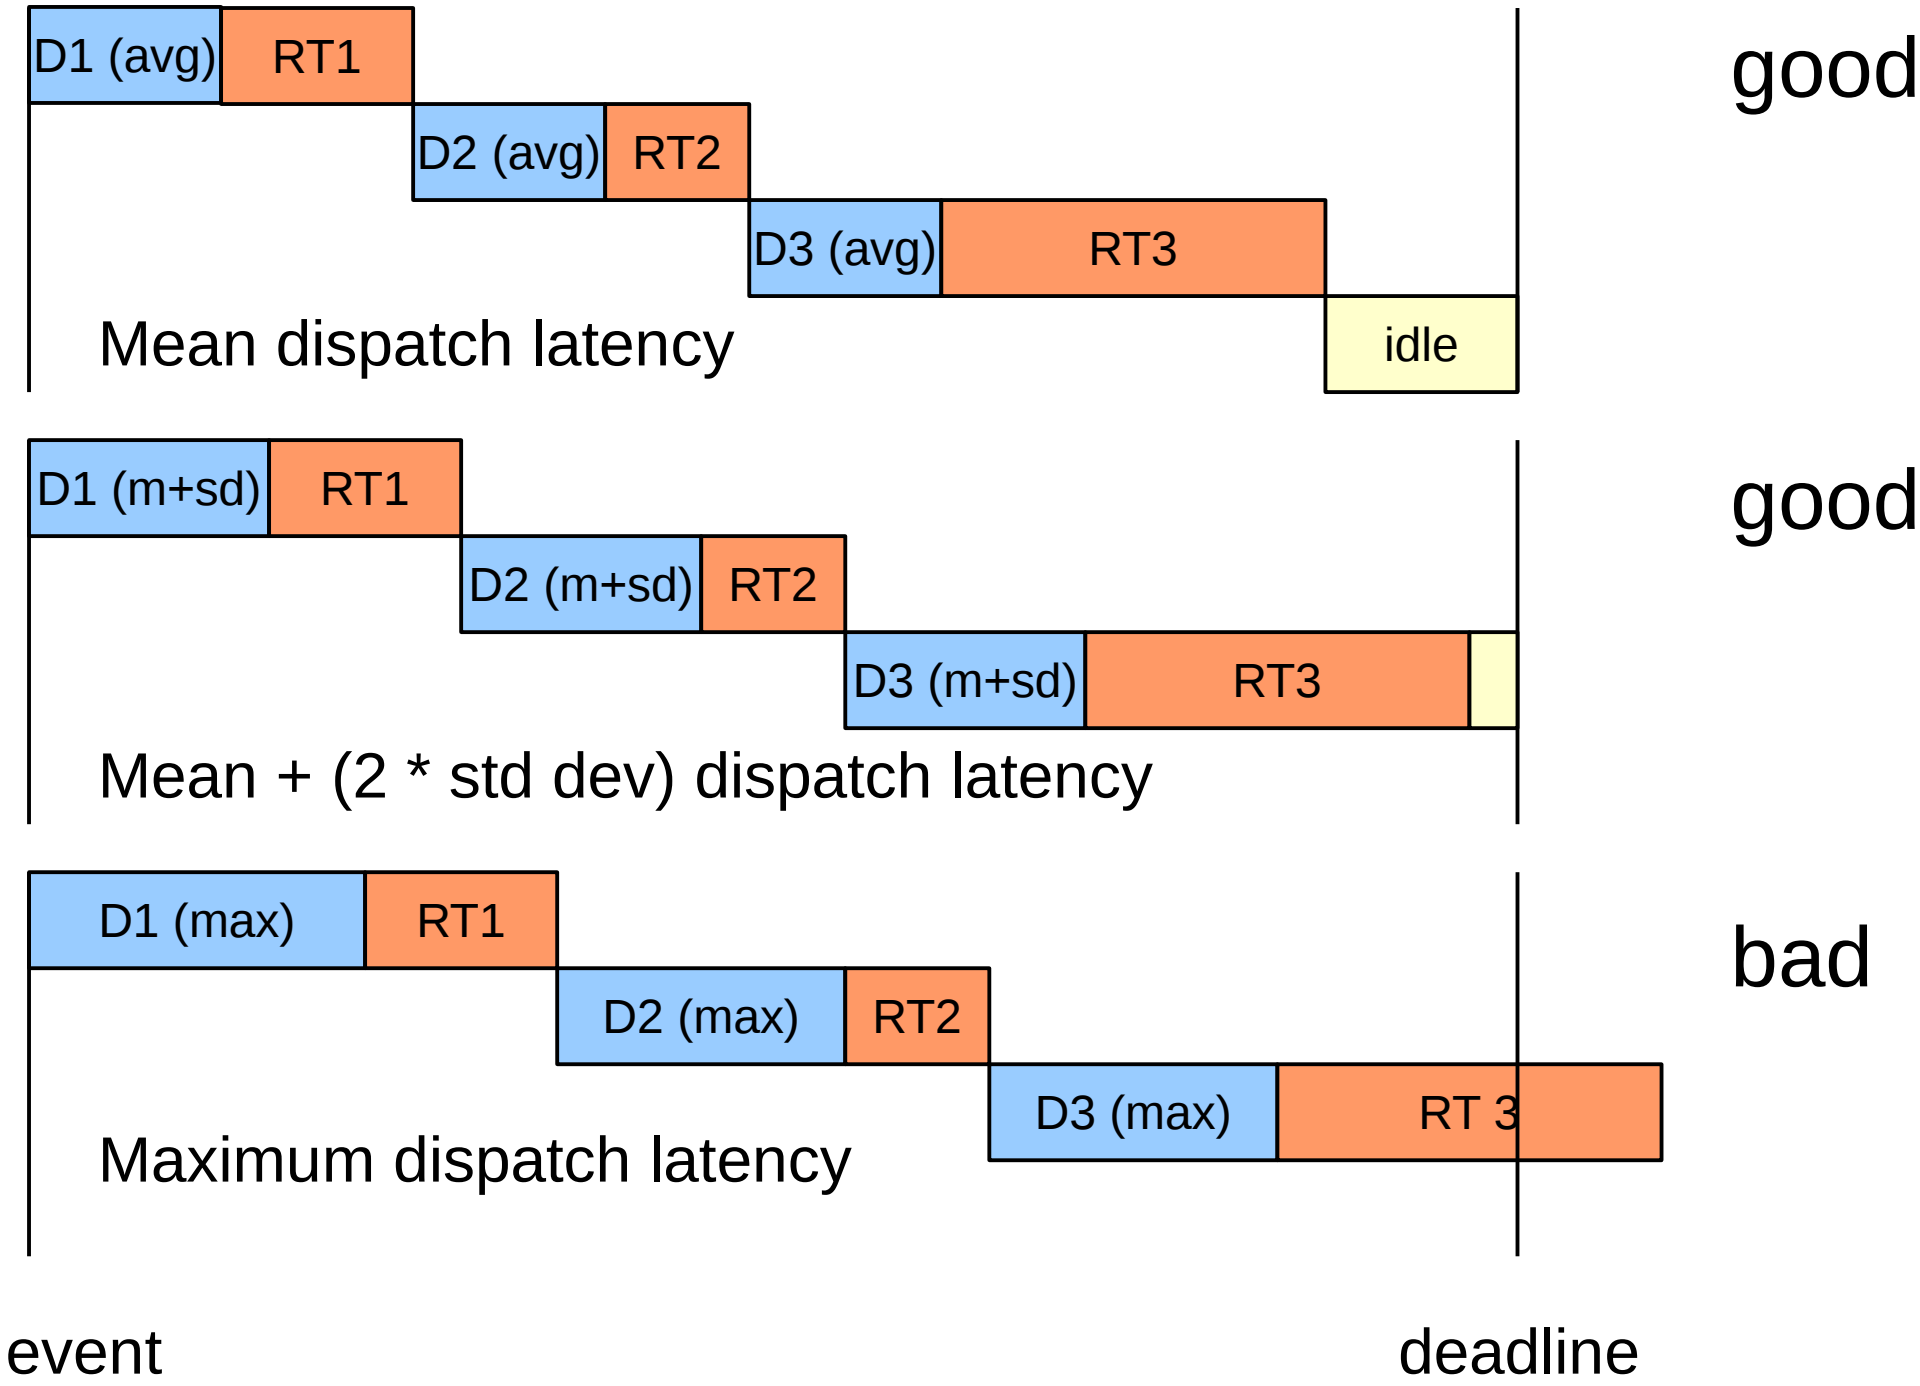
D1 (avg)
RT1
D2 (avg)
RT2
D3 (avg)
RT3
idle
Mean dispatch latency
good
D1 (m+sd)
RT1
D2 (m+sd)
RT2
D3 (m+sd)
RT3
Mean + (2 * std dev) dispatch latency
good
D1 (max)
RT1
D2 (max)
RT2
D3 (max)
RT 3
Maximum dispatch latency
bad
event
deadline

Is “Mean + (2 * std dev)” special?

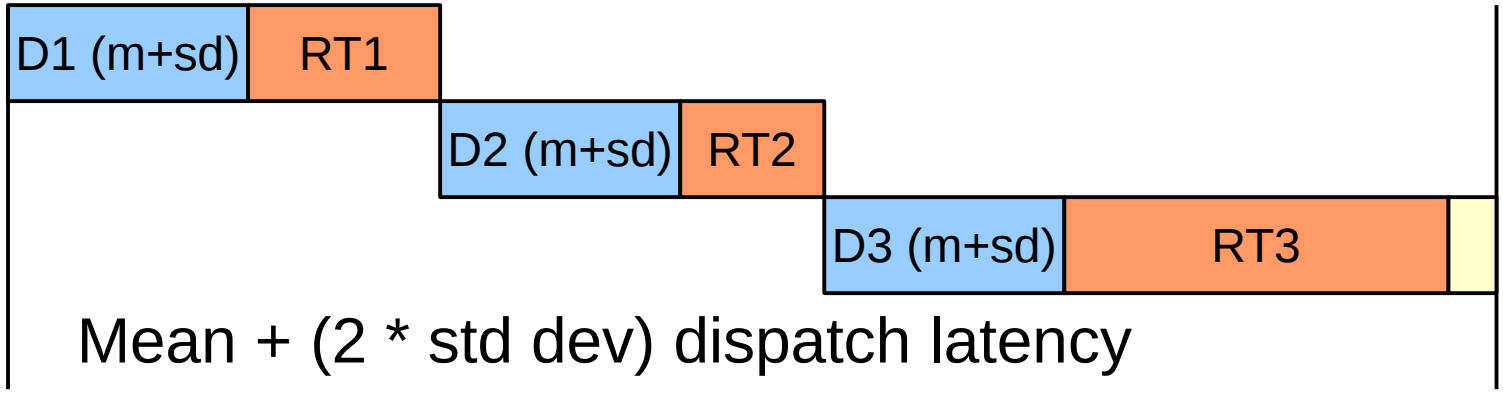

good

Is “Mean + (2 * std dev)” special?

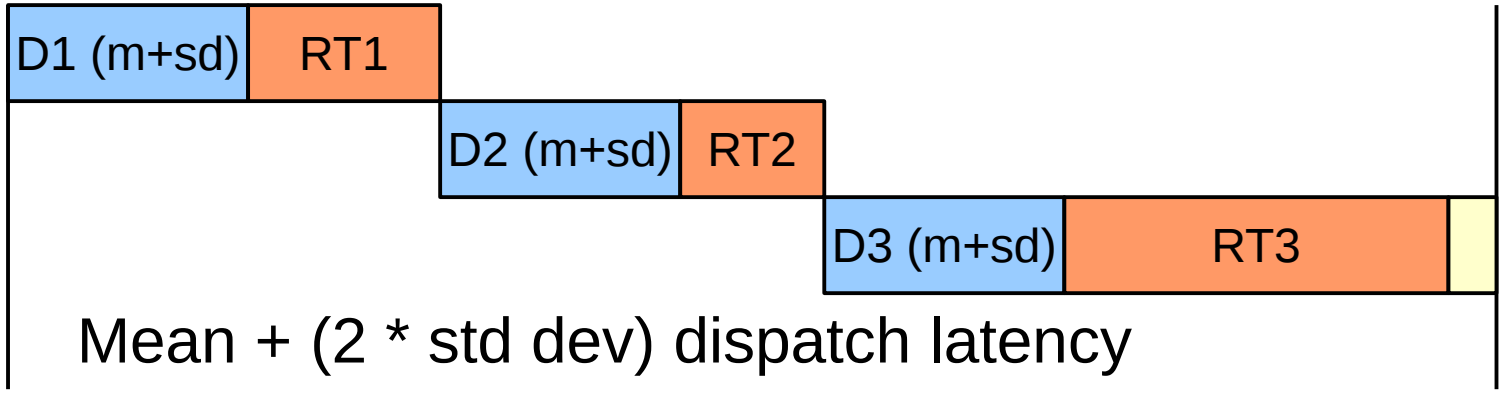

good

No – it is just an example that maximum may not always be the appropriate value

Is “Mean + (2 * std dev)” special?

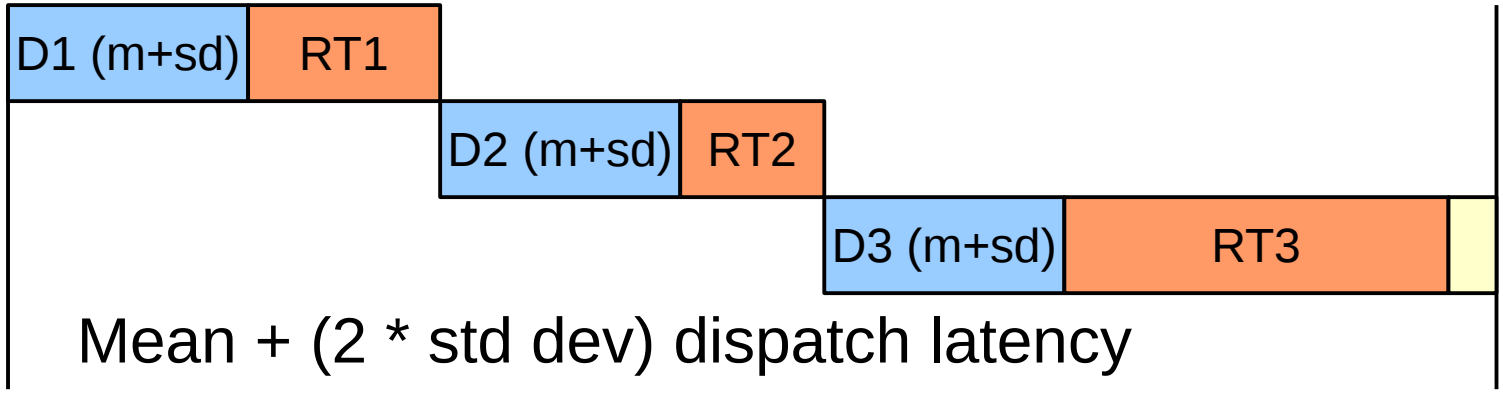

good

No – it is just an example that maximum may not always be the appropriate value

I do not have a magic formula to determine which value of latency is appropriate for any given system

Is “Mean + (2 * std dev)” special?

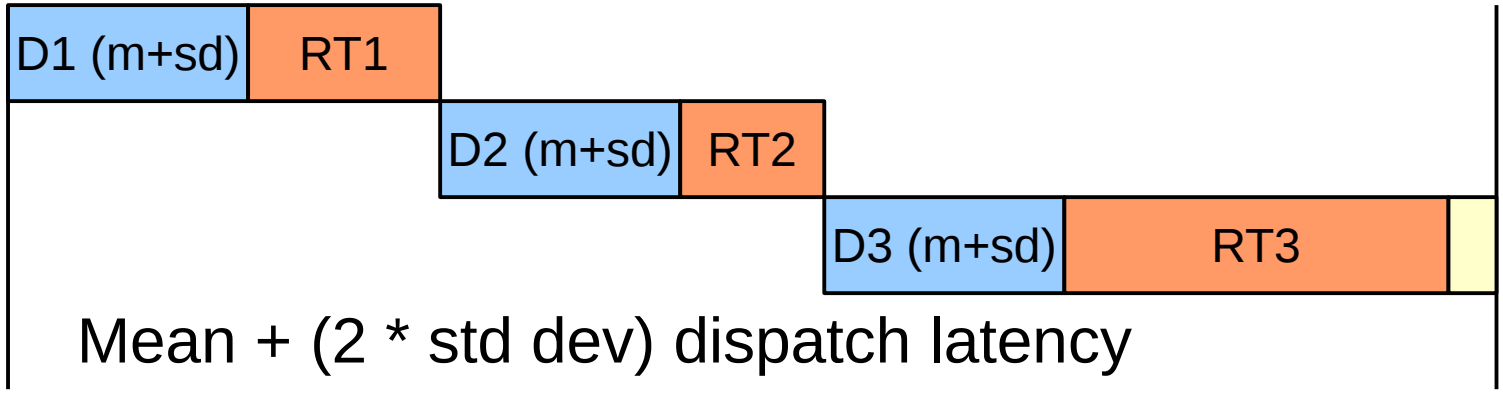

good

Even though “Mean + (2 * std dev)” is not special, it does have a useful meaning.

Standard Deviation of a Normal Distribution

z	percent of samples within mean +/- (z * std dev)
1.96	95%
2.326	98%
2.576	99%
3.291	99.9%
4.0	99.993666%
5.0	99.99994267%
6.0	99.9999998027%
7.0	99.9999999997440%

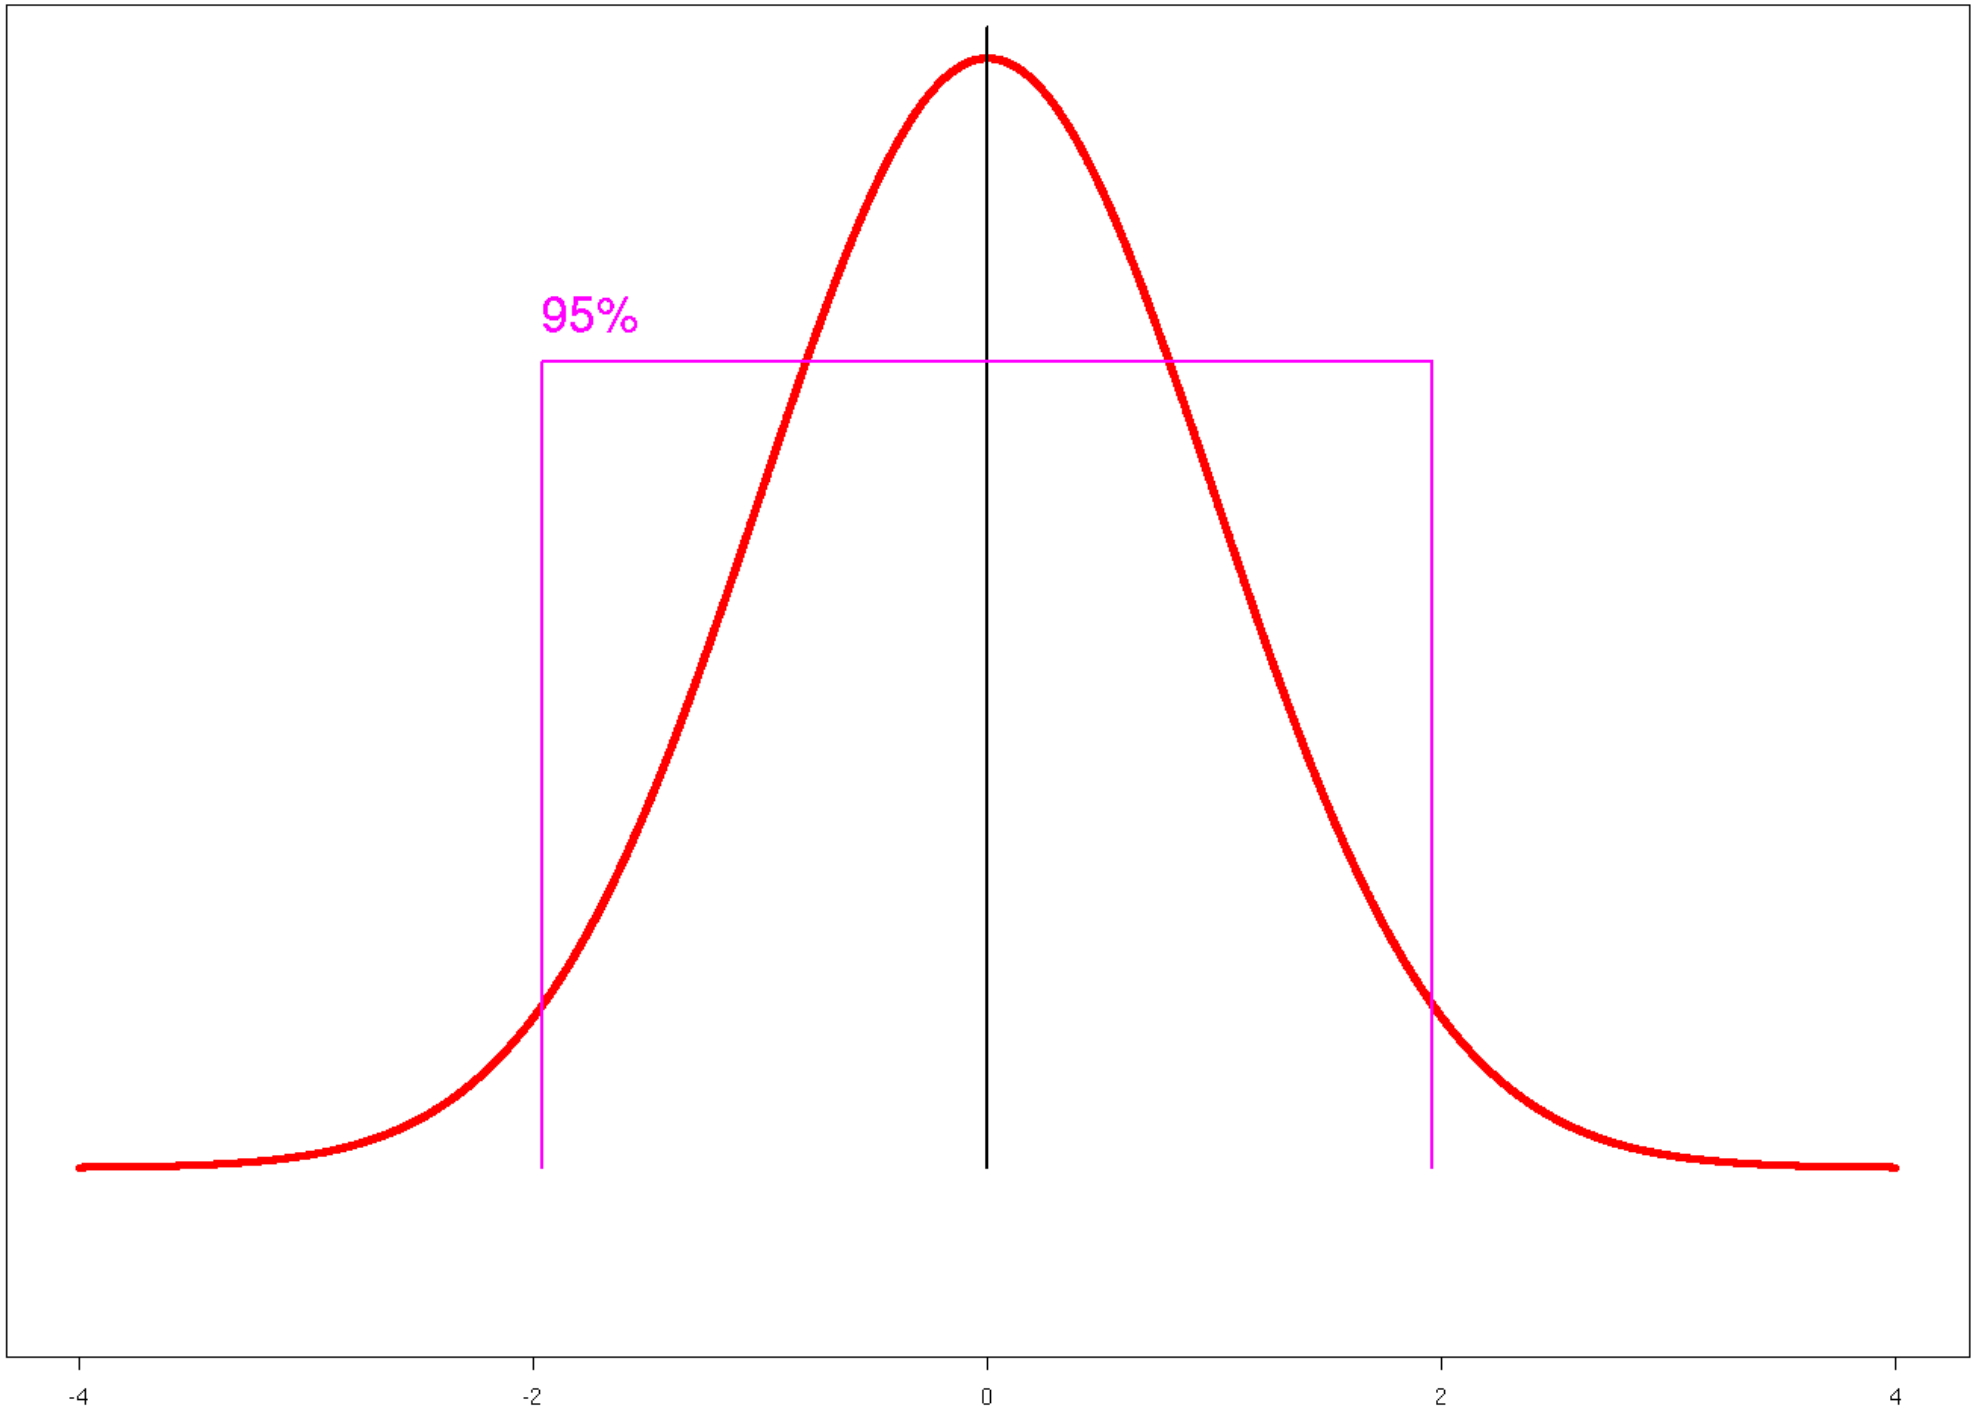
95%
-4
-2
0
2
4

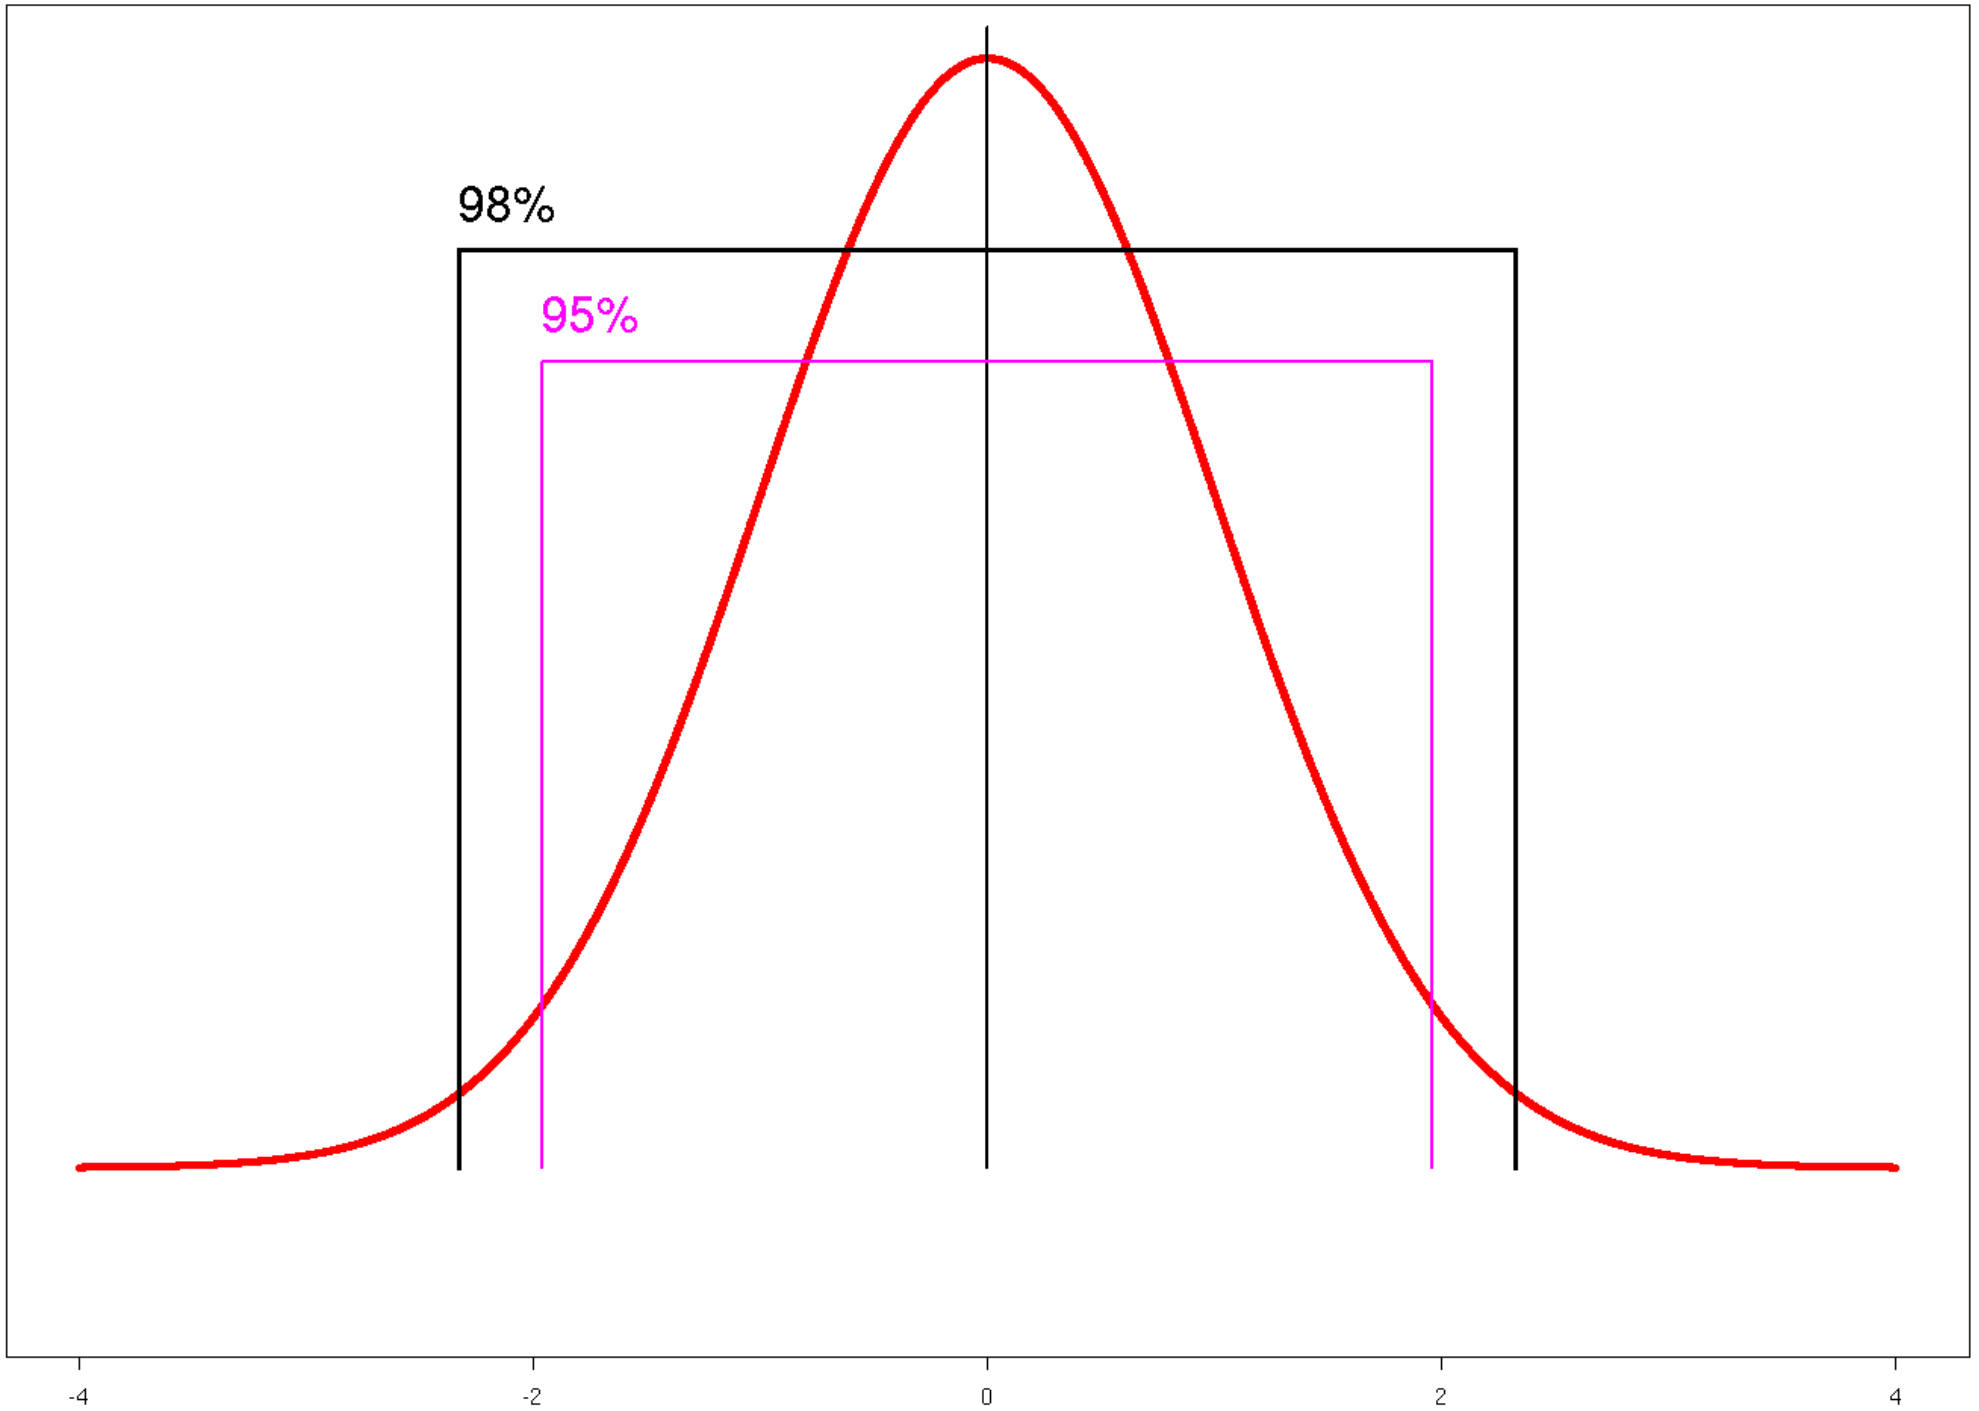
98%
95%
-4
-2
0
2
4

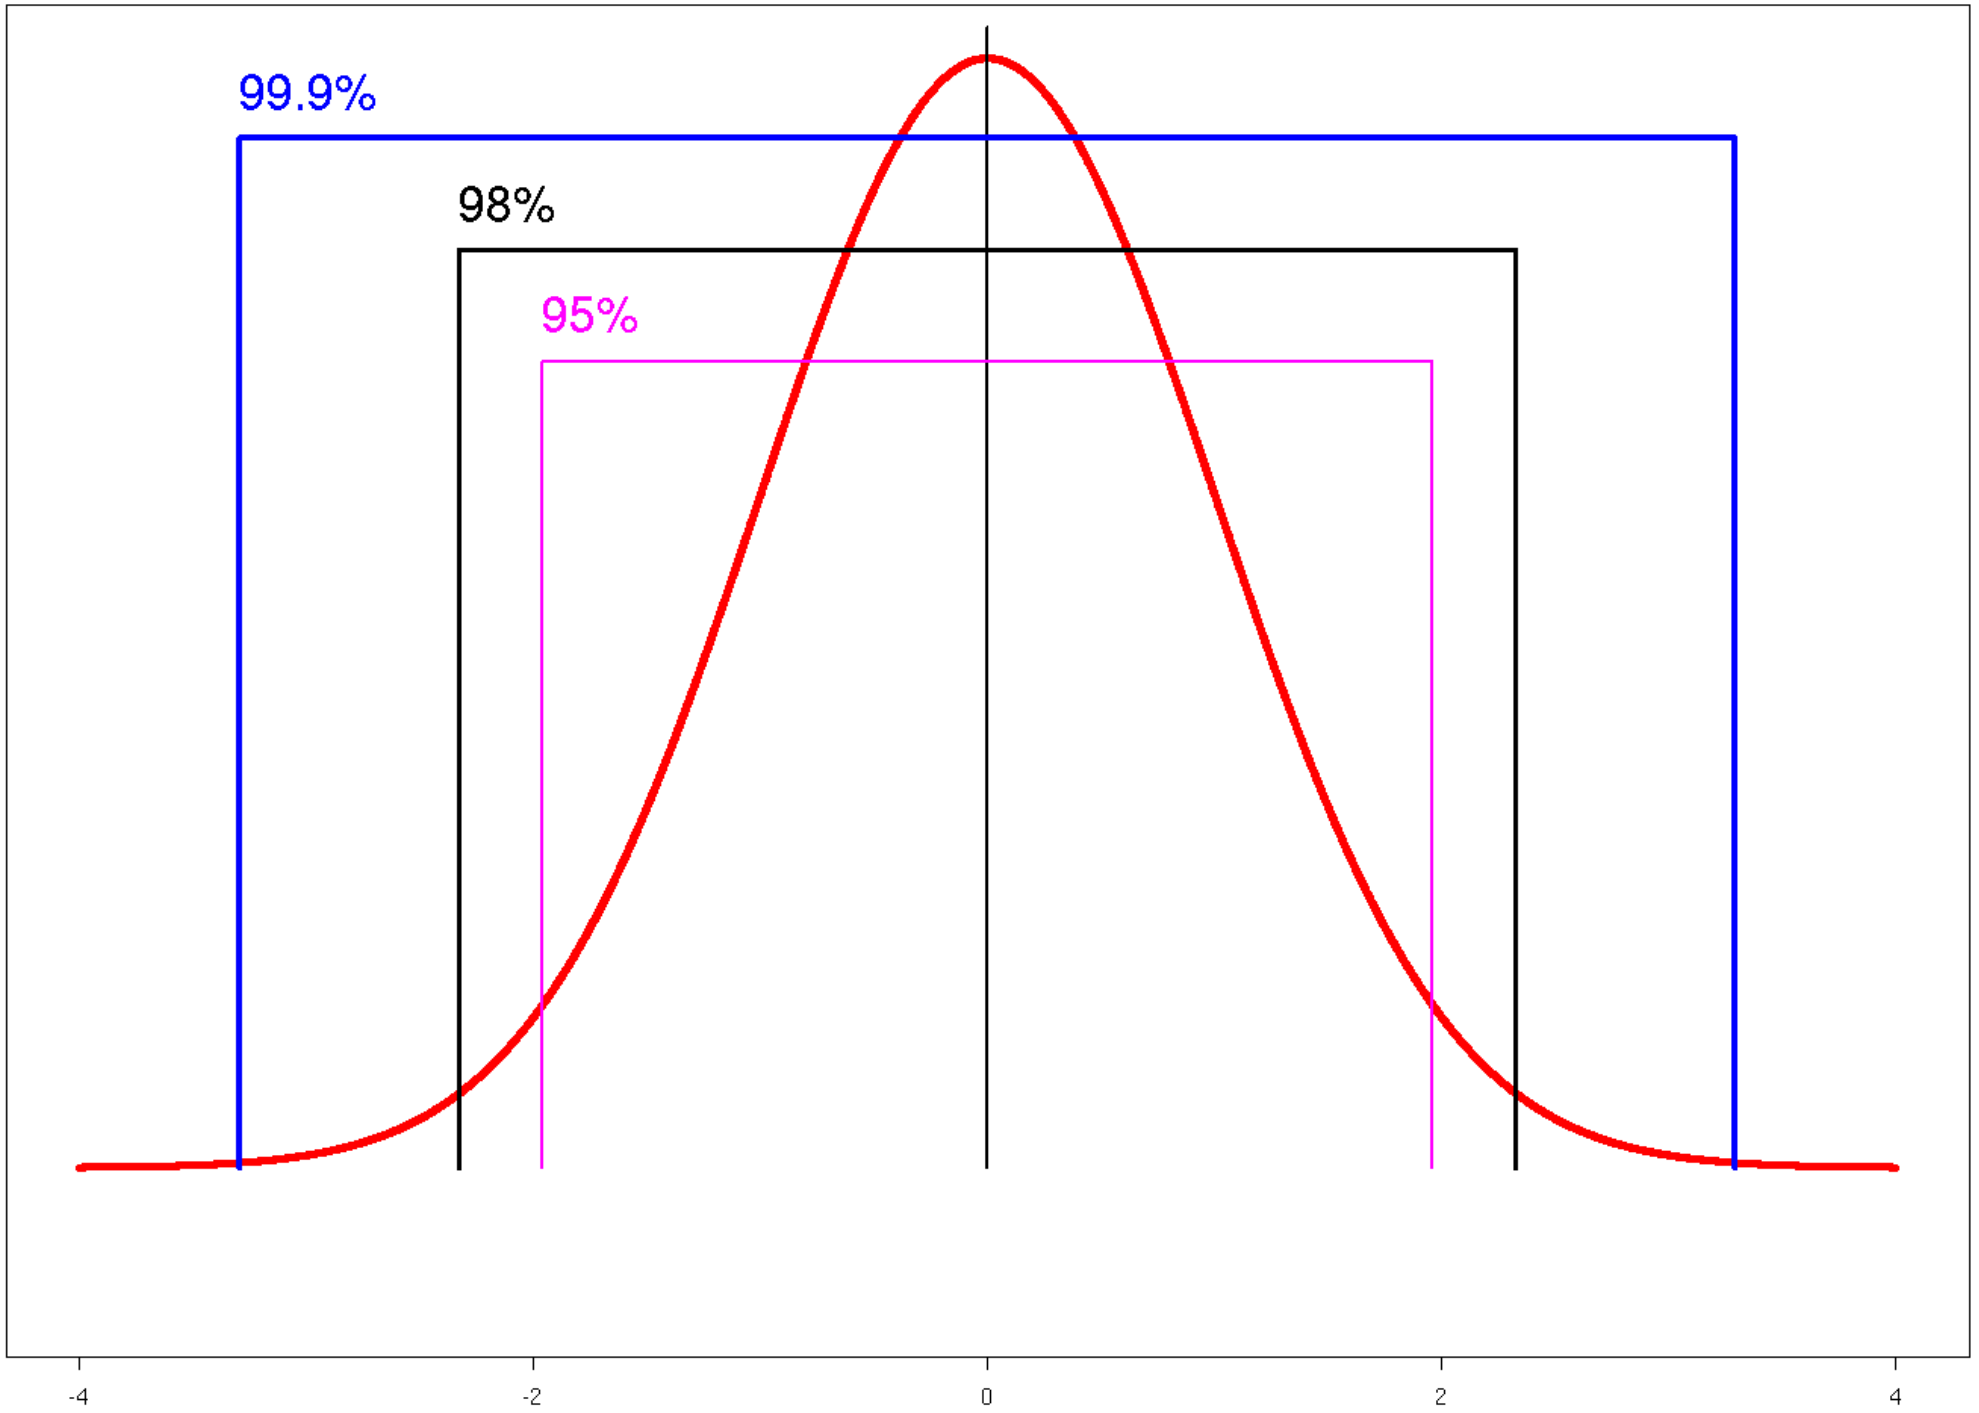
99.9%
98%
95%
-4
-2
0
2
4

Test for “normal distribution”

Test for “normal distribution”

- Are more values outside of the range than predicted on the previous slide?

Test for “normal distribution”

- Are more values outside of the range than predicted on the previous slide?

  + Easy test, even for the statistically challenged
  + A real statistician would be appalled if I were to suggest this as a real answer.

Test for “normal distribution”

- Are more values outside of the range than predicted on the previous slide?

- Graph the data. Does it look normal?

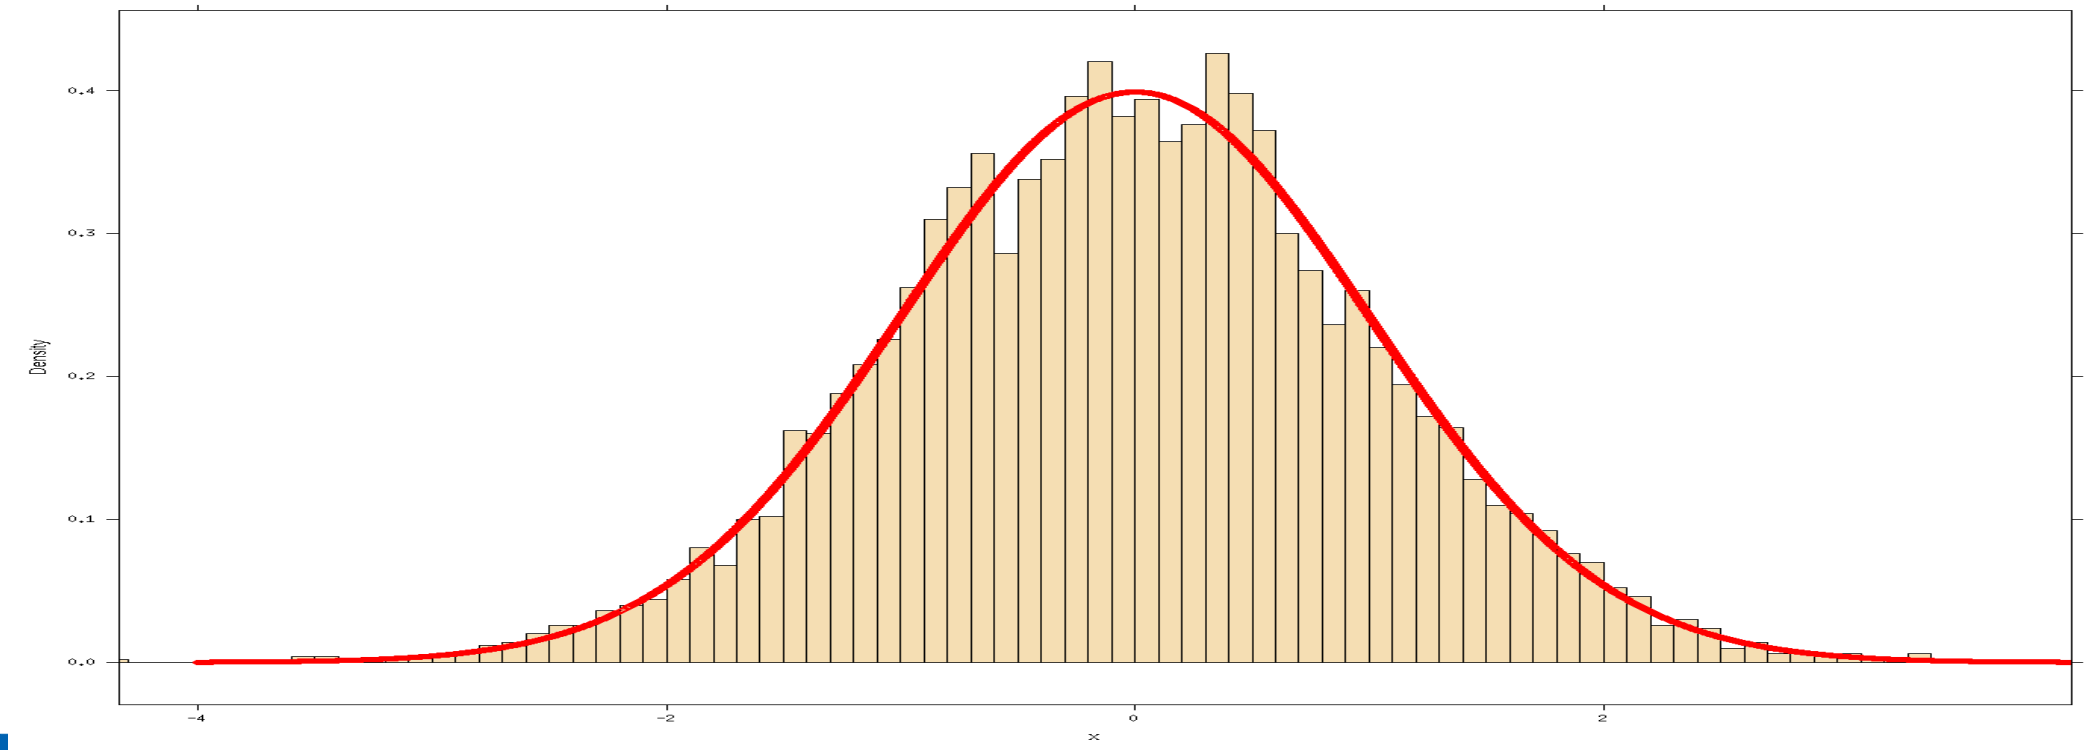

Test for “normal distribution”

- Are more values outside of the range than predicted on the previous slide?

- Graph the data. Does it look normal?

    + Once again, a real statistician would be appalled.

Test for “normal distribution”

- Are more values outside of the range then predicted on the previous slide?

- Graph the data. Does it look normal?

- Use a real statistical test.

Test for “normal distribution”

- Are more values outside of the range than predicted on the previous slide?

- Graph the data. Does it look normal?

- Use a real statistical test.

  + skewness
    http://en.wikipedia.org/wiki/Skewness

  + kurtosis
    http://en.wikipedia.org/wiki/Kurtosis

Test for “normal distribution”

My data is a normal distribution, is there any reason I should still be cautious?

Test for “normal distribution”

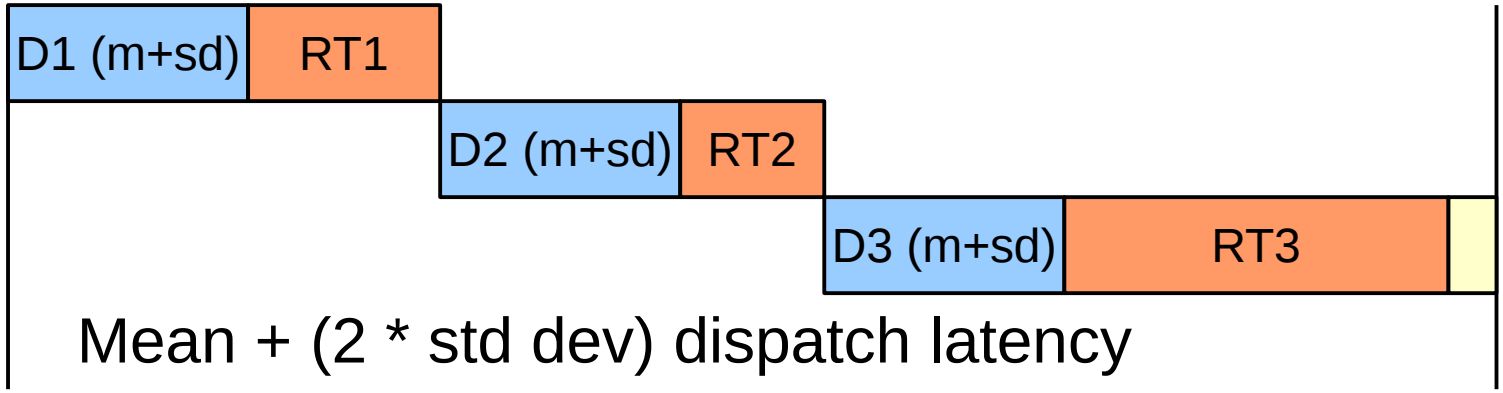

good

My data is a normal distribution, is there any reason I should still be cautious?

Including increasing portions of the tail of the distribution

The uncertainty and risk increase.

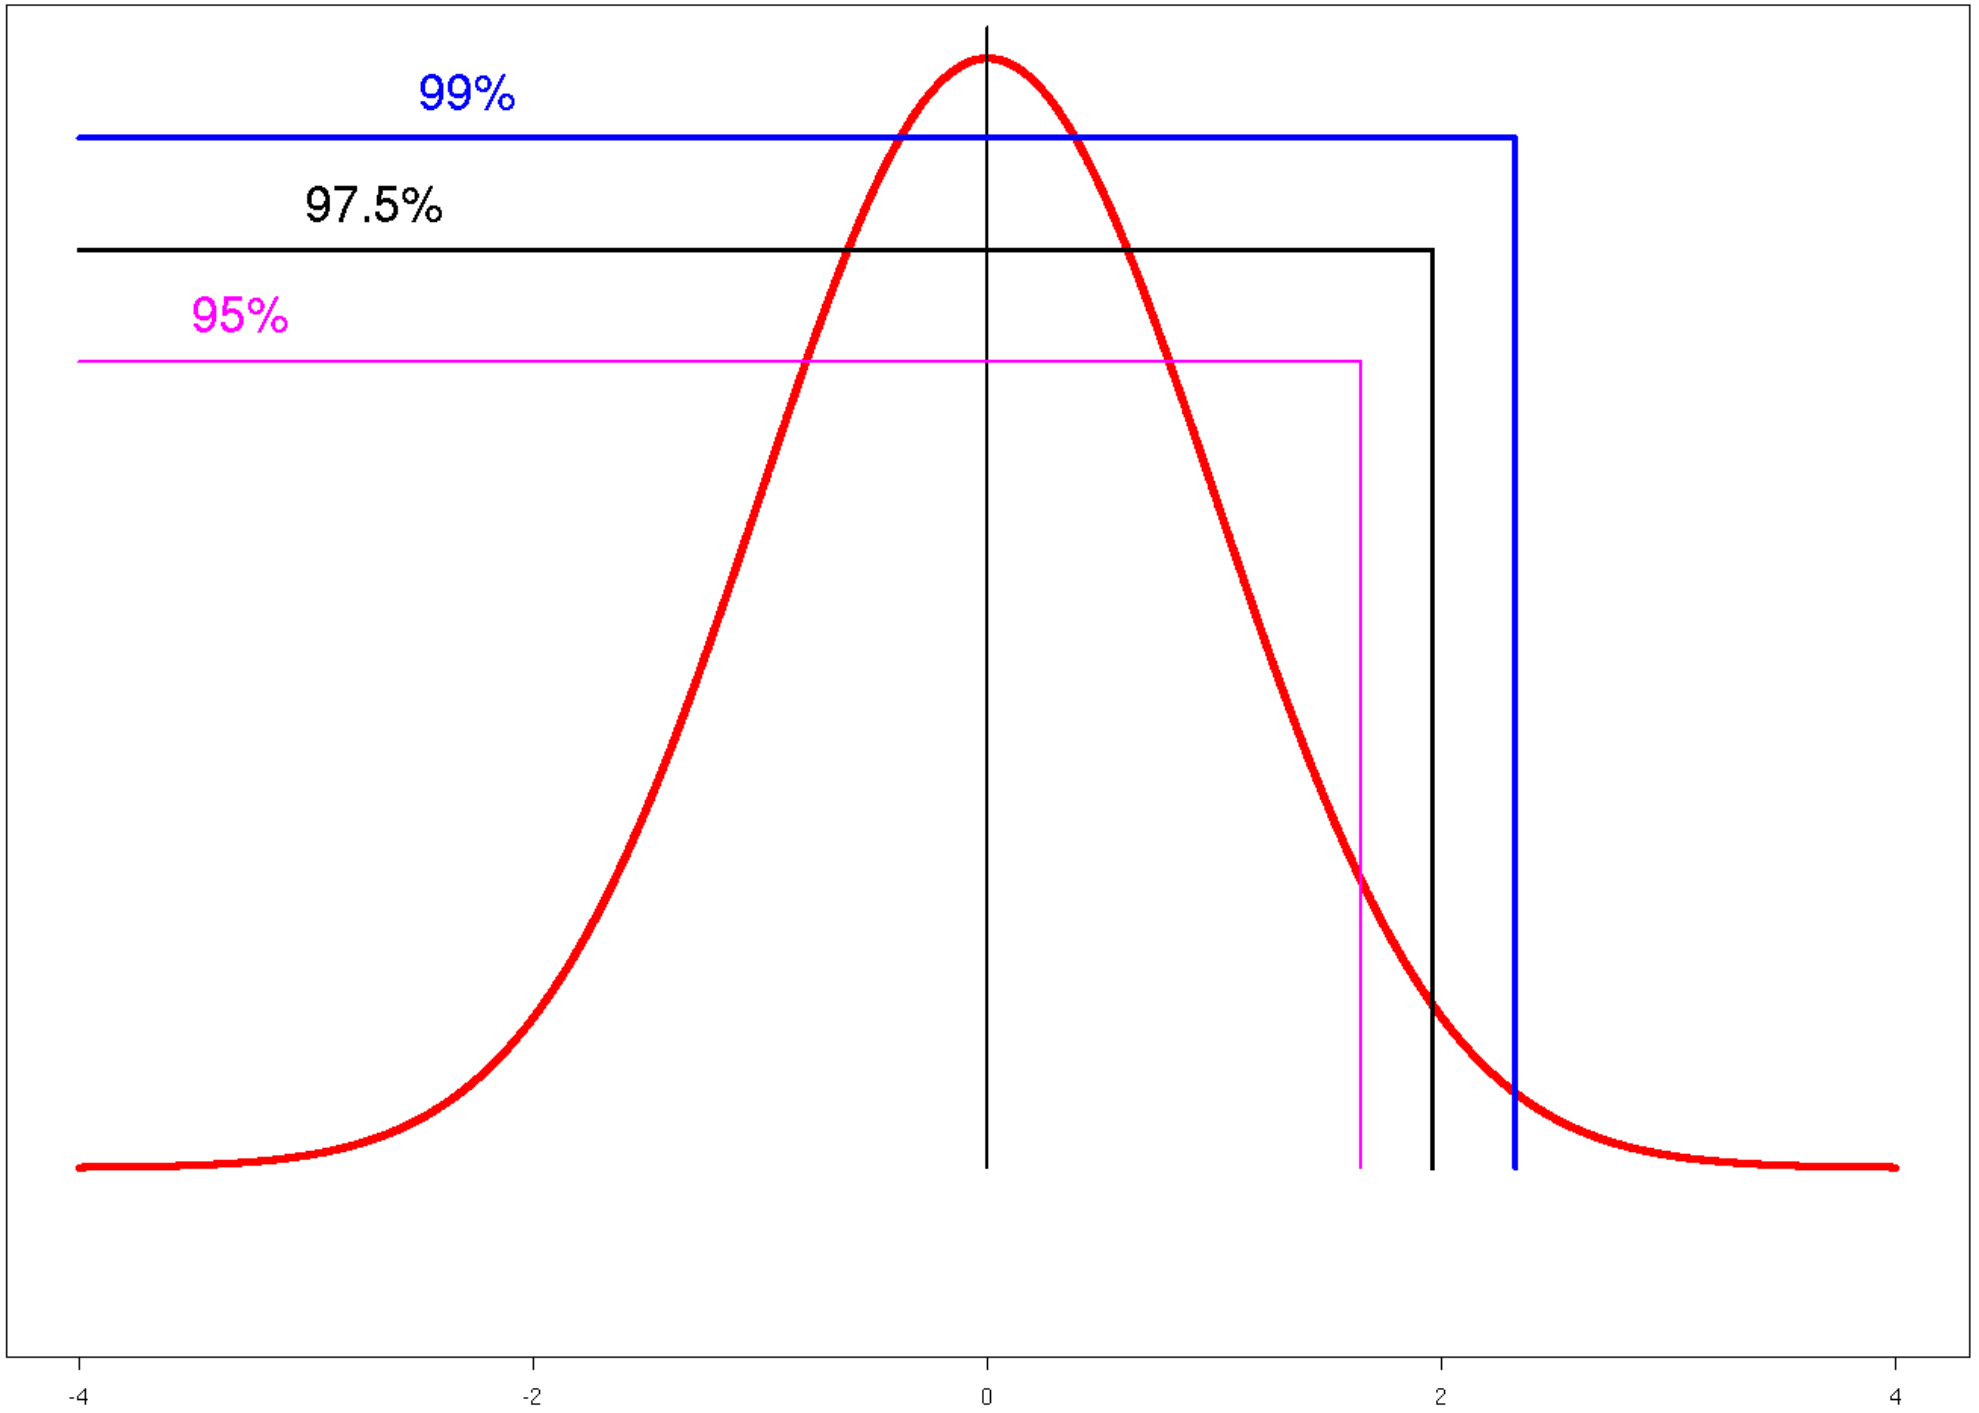
99%
97.5%
95%
-4
-2
0
2
4

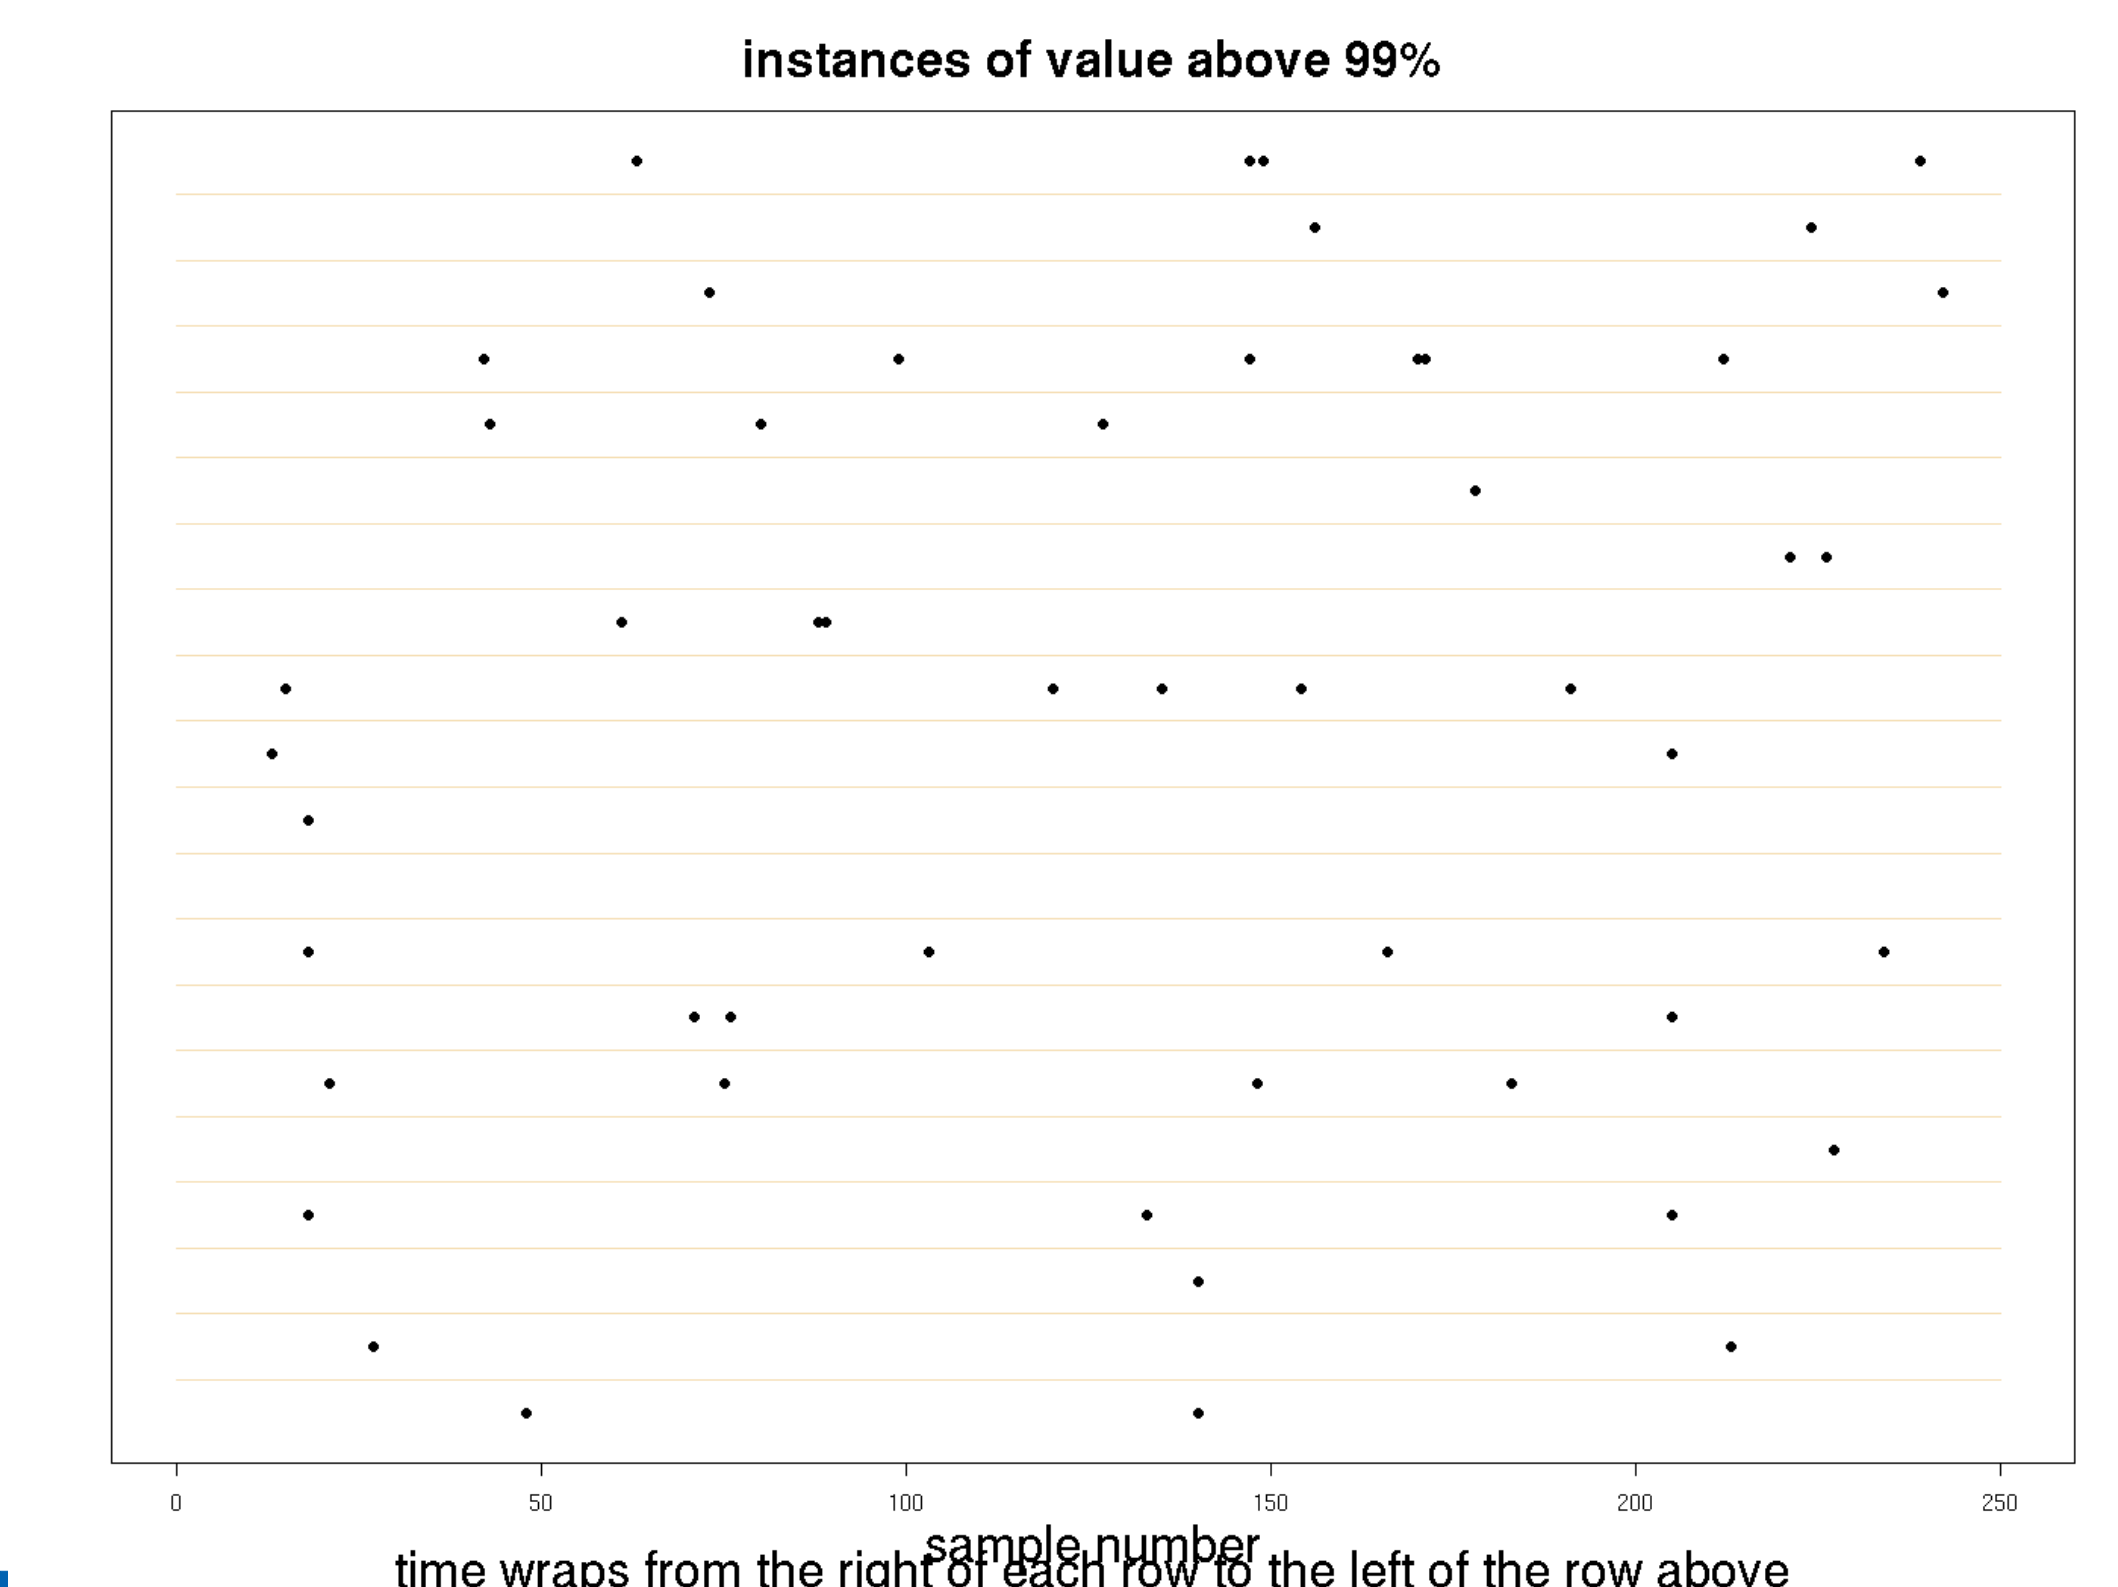
instances of value above 99%
0
50
100
150
200
250
sample number
time wraps from the right of each row to the left of the row above

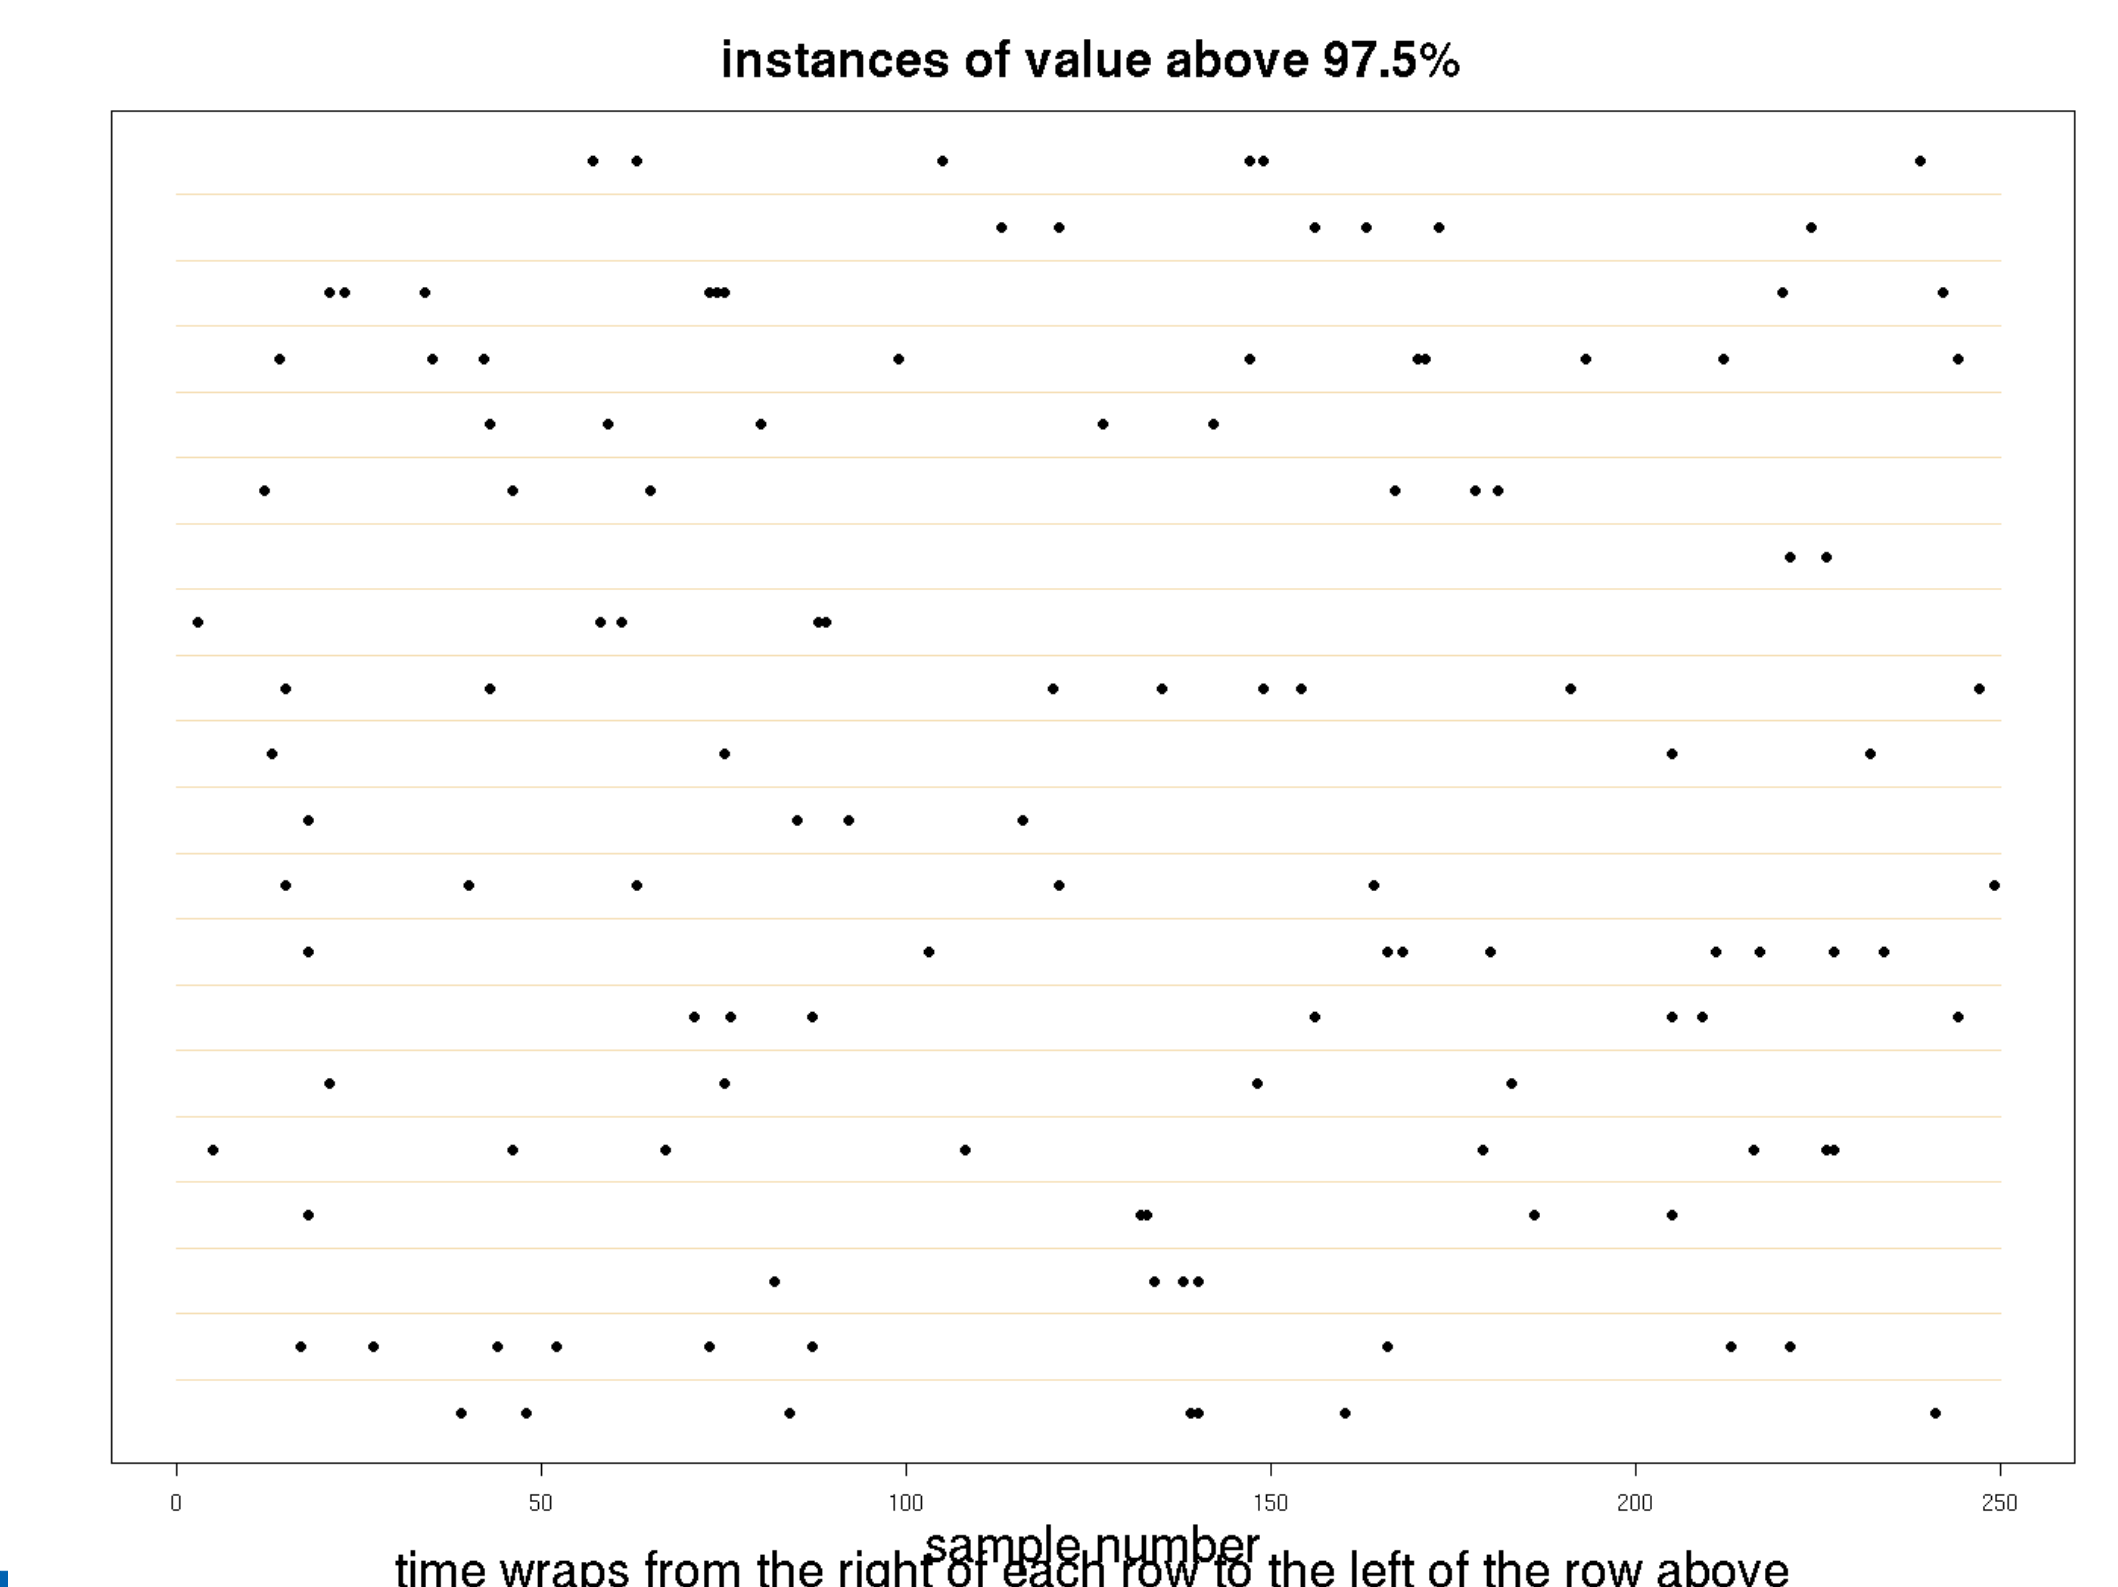
instances of value above 97.5%
0
50
100
150
200
250
sample number
time wraps from the right of each row to the left of the row above


instances of value above 95%
0
50
100
150
200
250
sample number
time wraps from the right of each row to the left of the row above

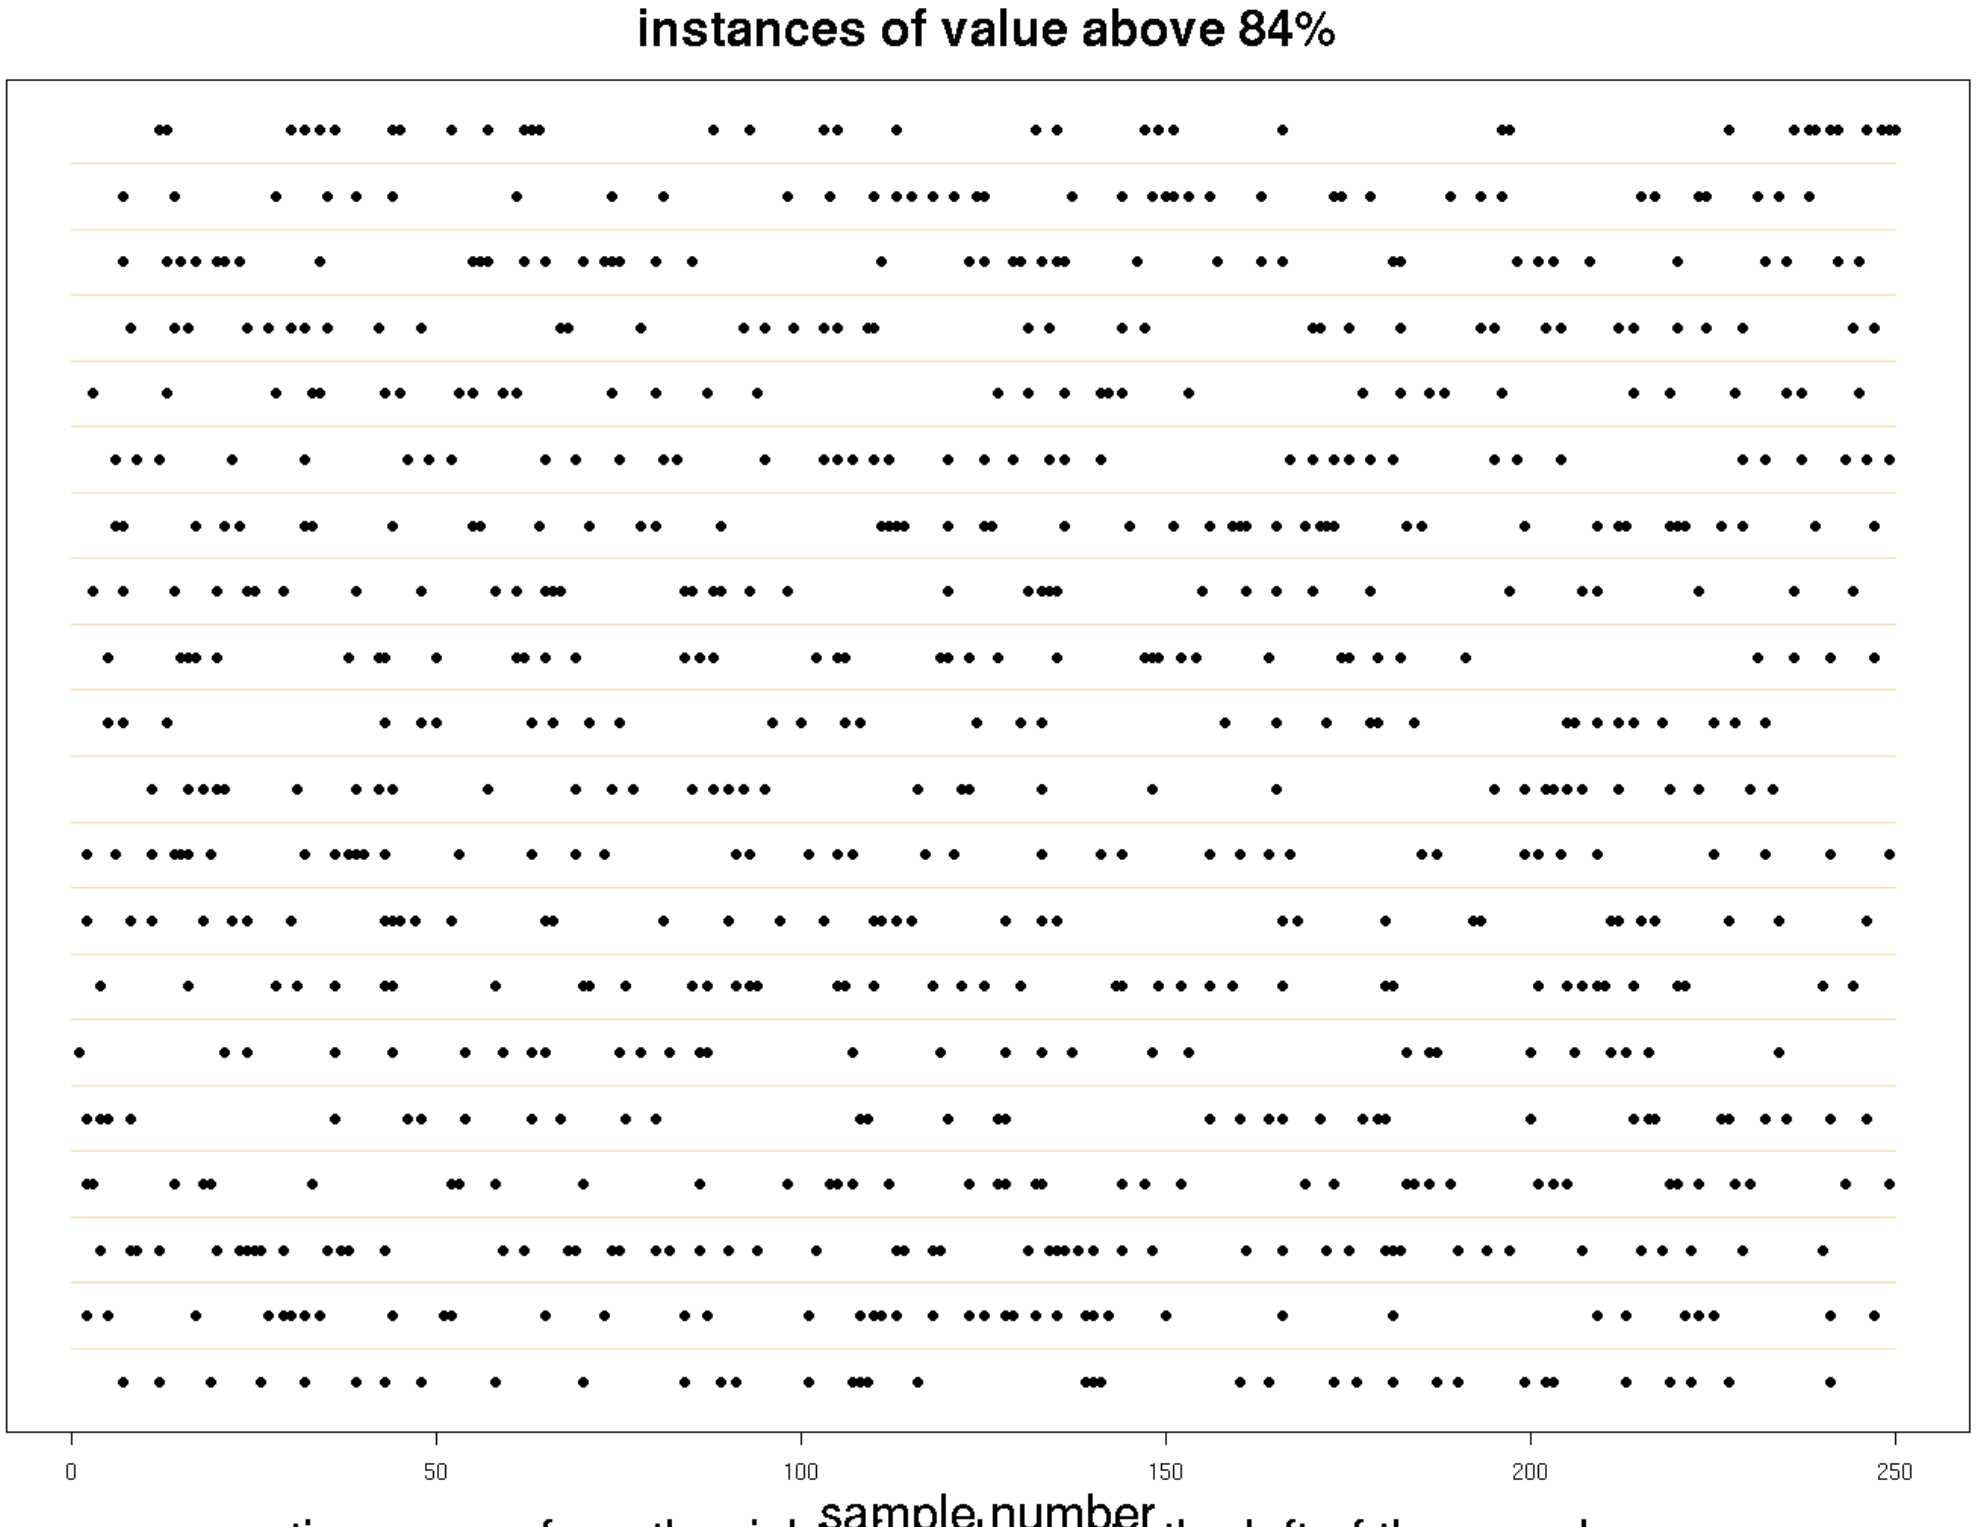
instances of value above 84%
0
50
100
150
200
250
sample number
time wraps from the right of each row to the left of the row above

Next slide is a perfect data set for soft real-time.

And an opportunity for tuning and debugging (even for hard real-time).

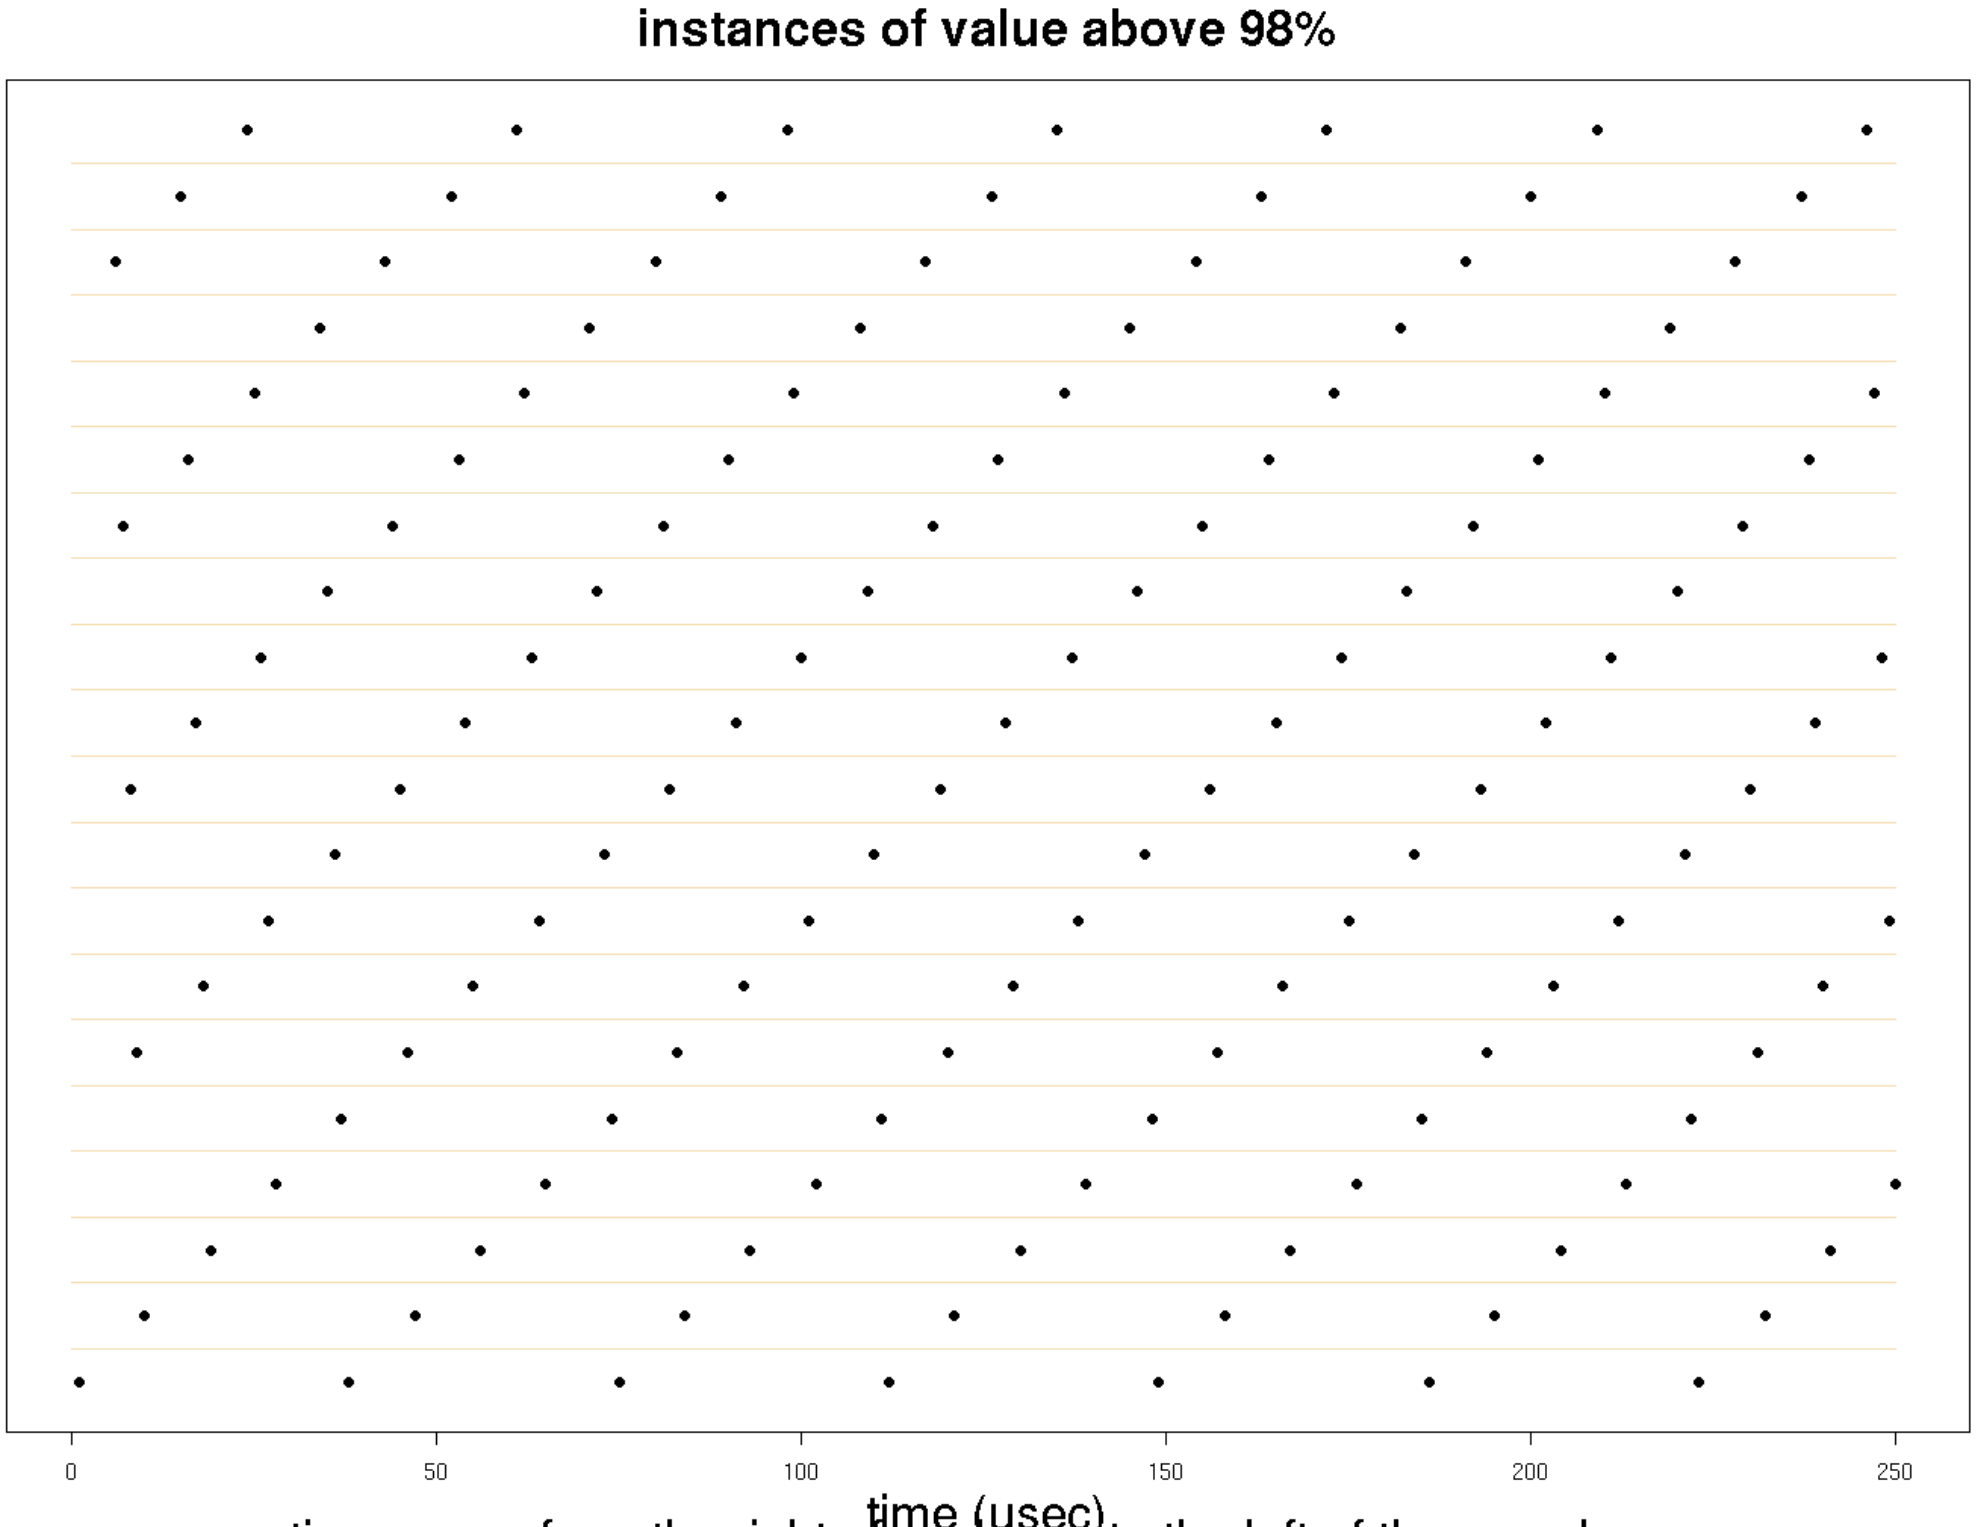
instances of value above 98%
0
50
100
150
200
250
time (usec)
time wraps from the right of each row to the left of the row above

Next slide is a pathalogical case.

But an opportunity for tuning and debugging.

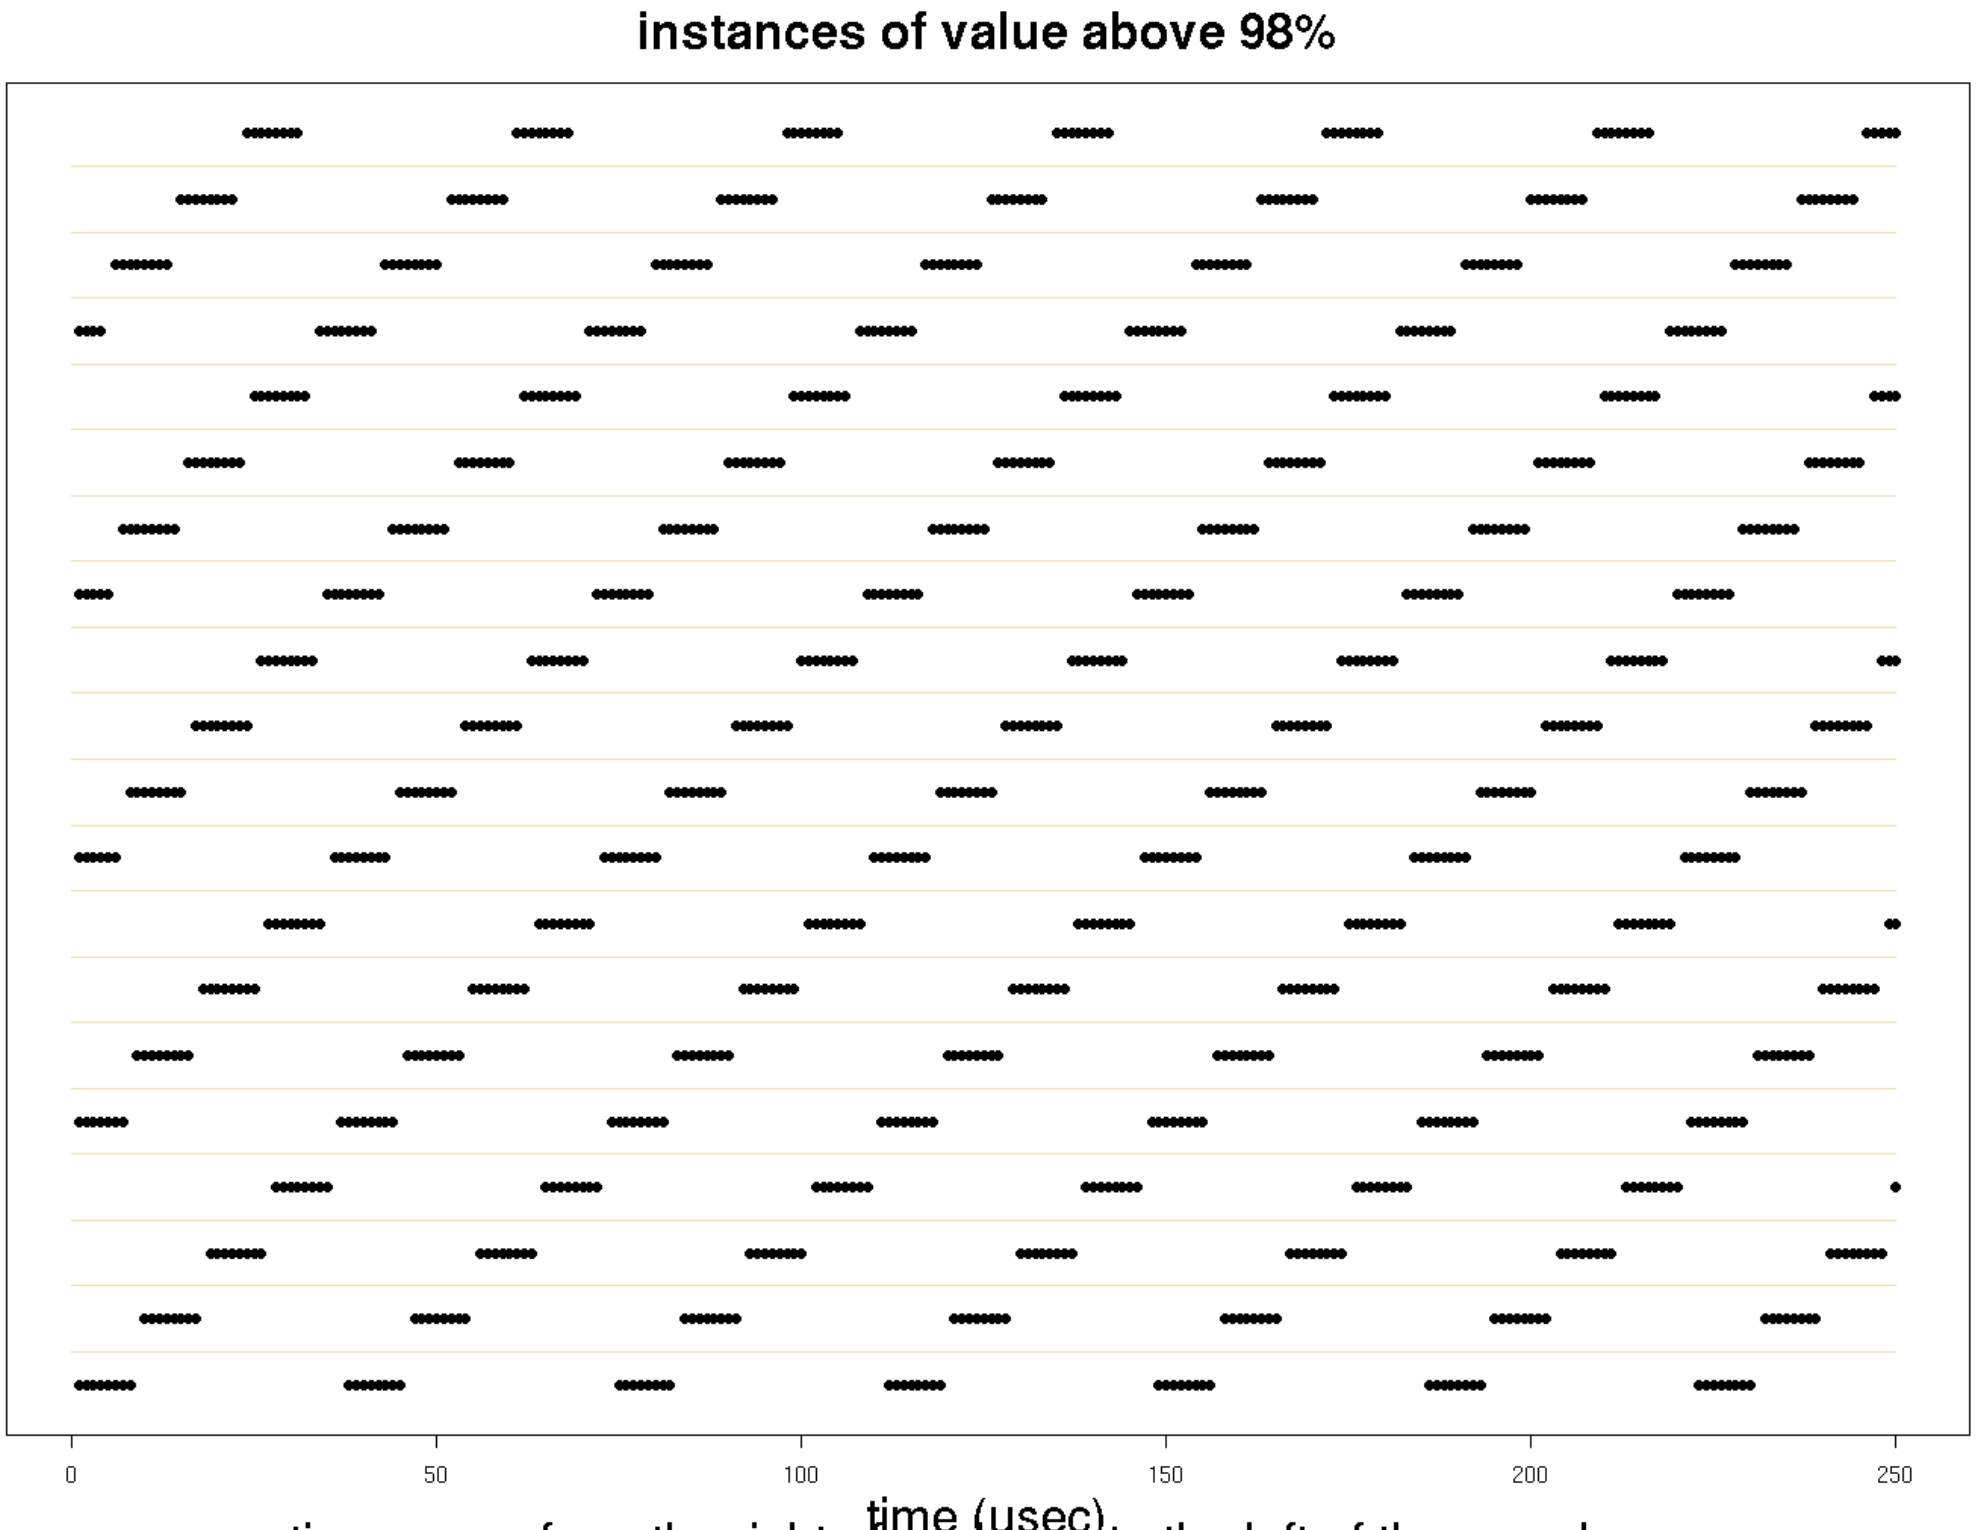
instances of value above 98%
0
50
100
150
200
250
time (usec)
time wraps from the right of each row to the left of the row above

Next slide is a less clear cut situation for soft real-time.

But an opportunity for tuning and debugging.

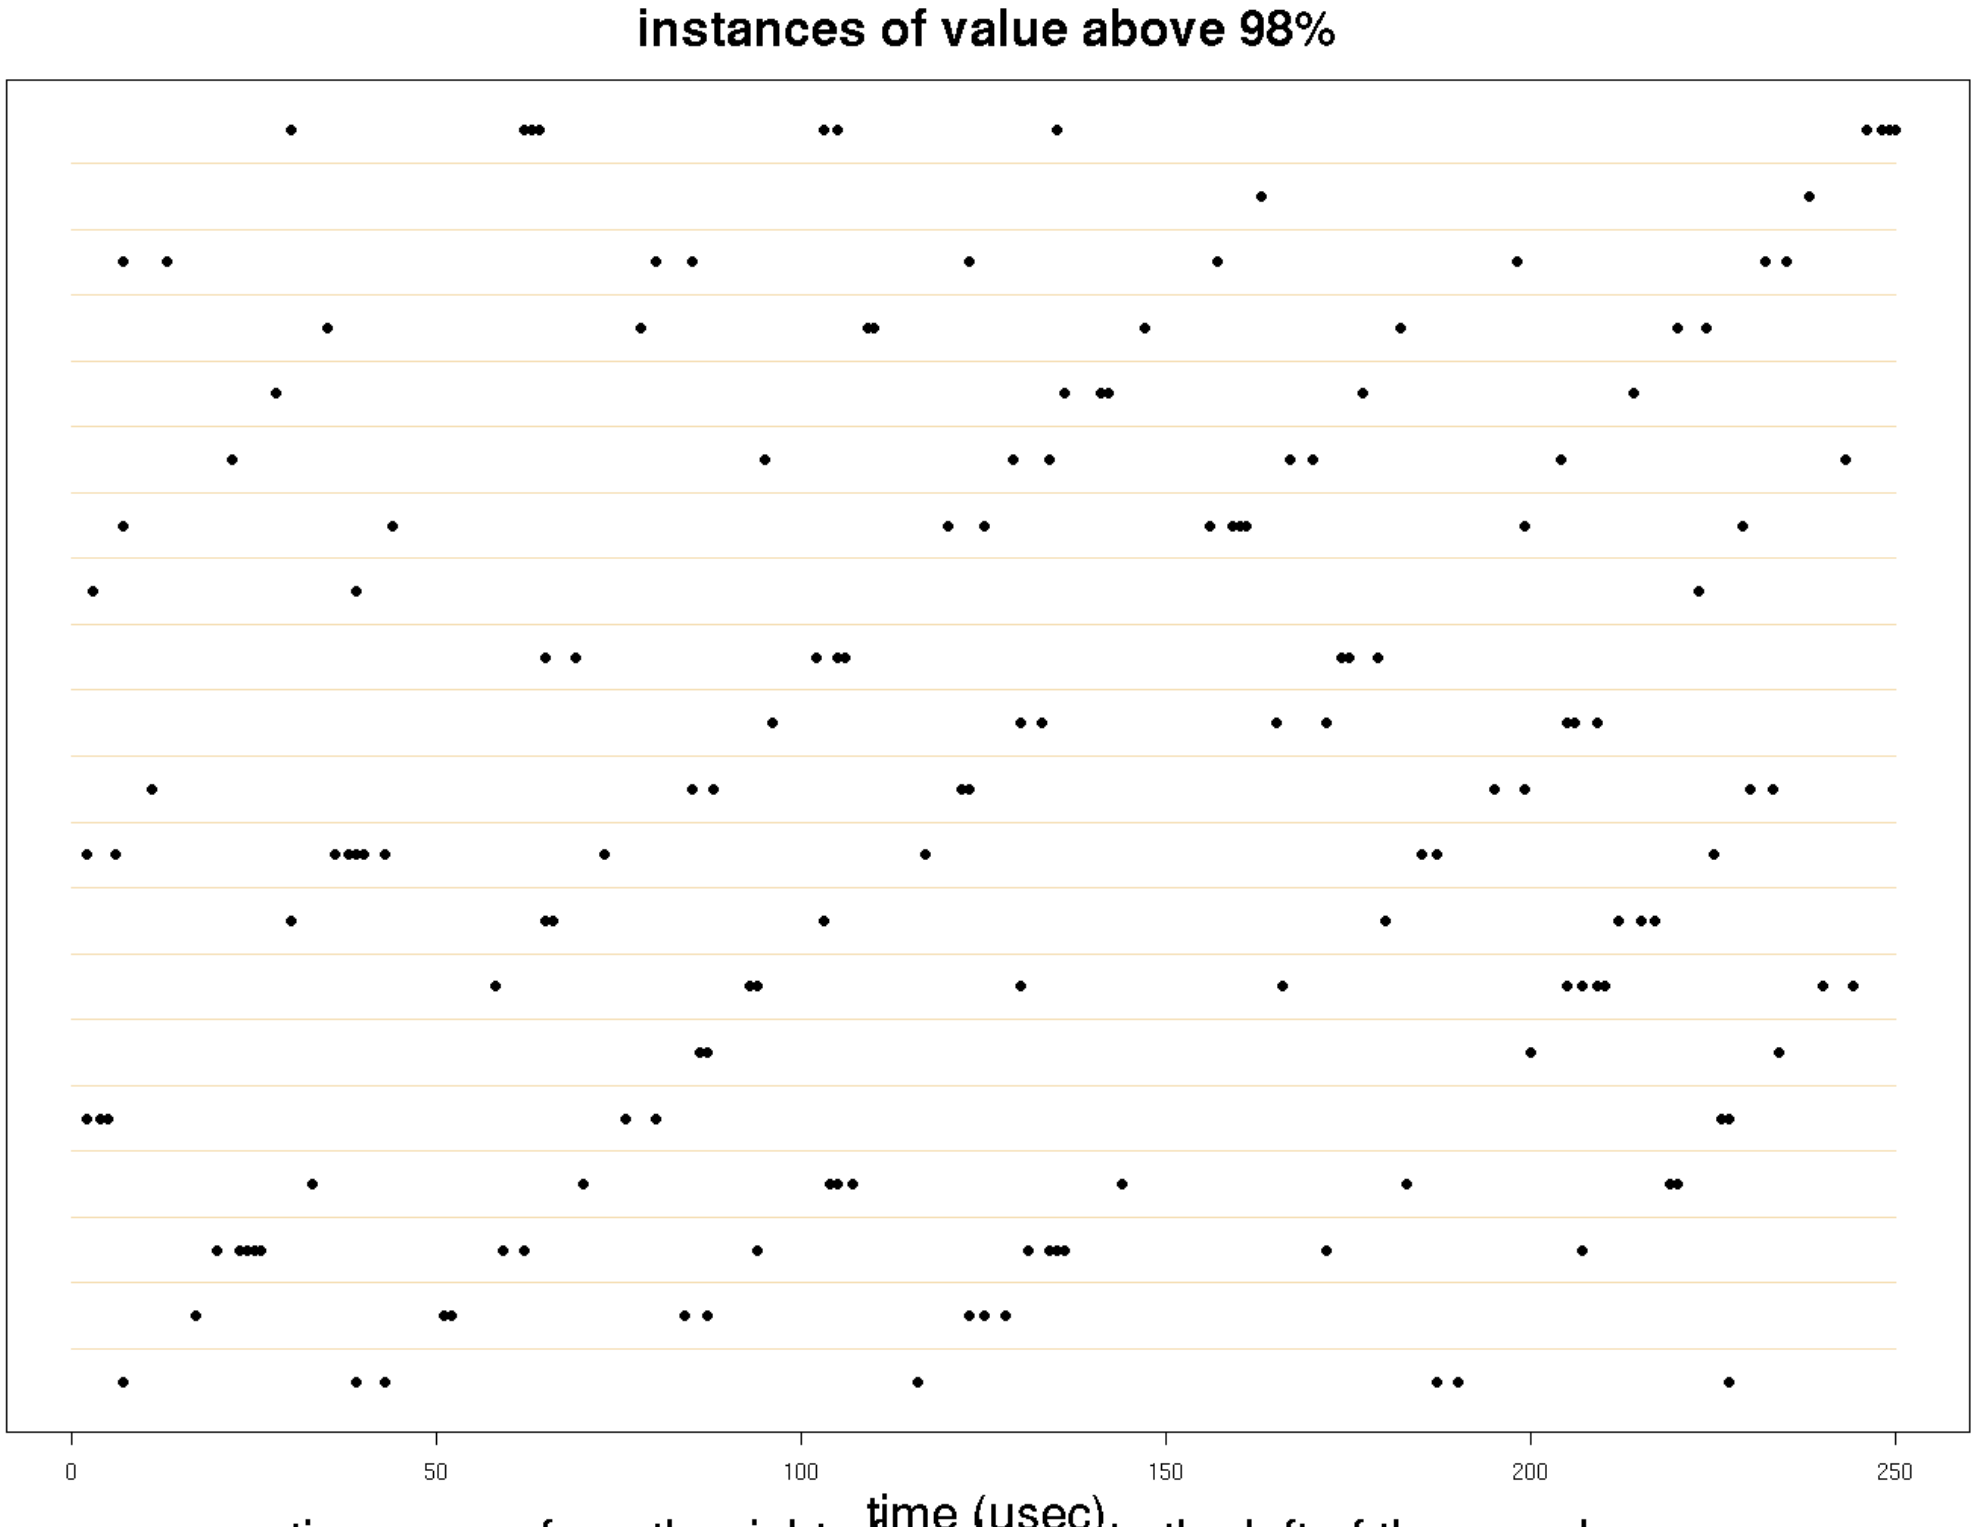
instances of value above 98%
0
50
100
150
200
250
time (usec)
time wraps from the right of each row to the left of the row above

Test for “normal distribution”

- Graph the data. Does it look normal?

Some real-life examples.

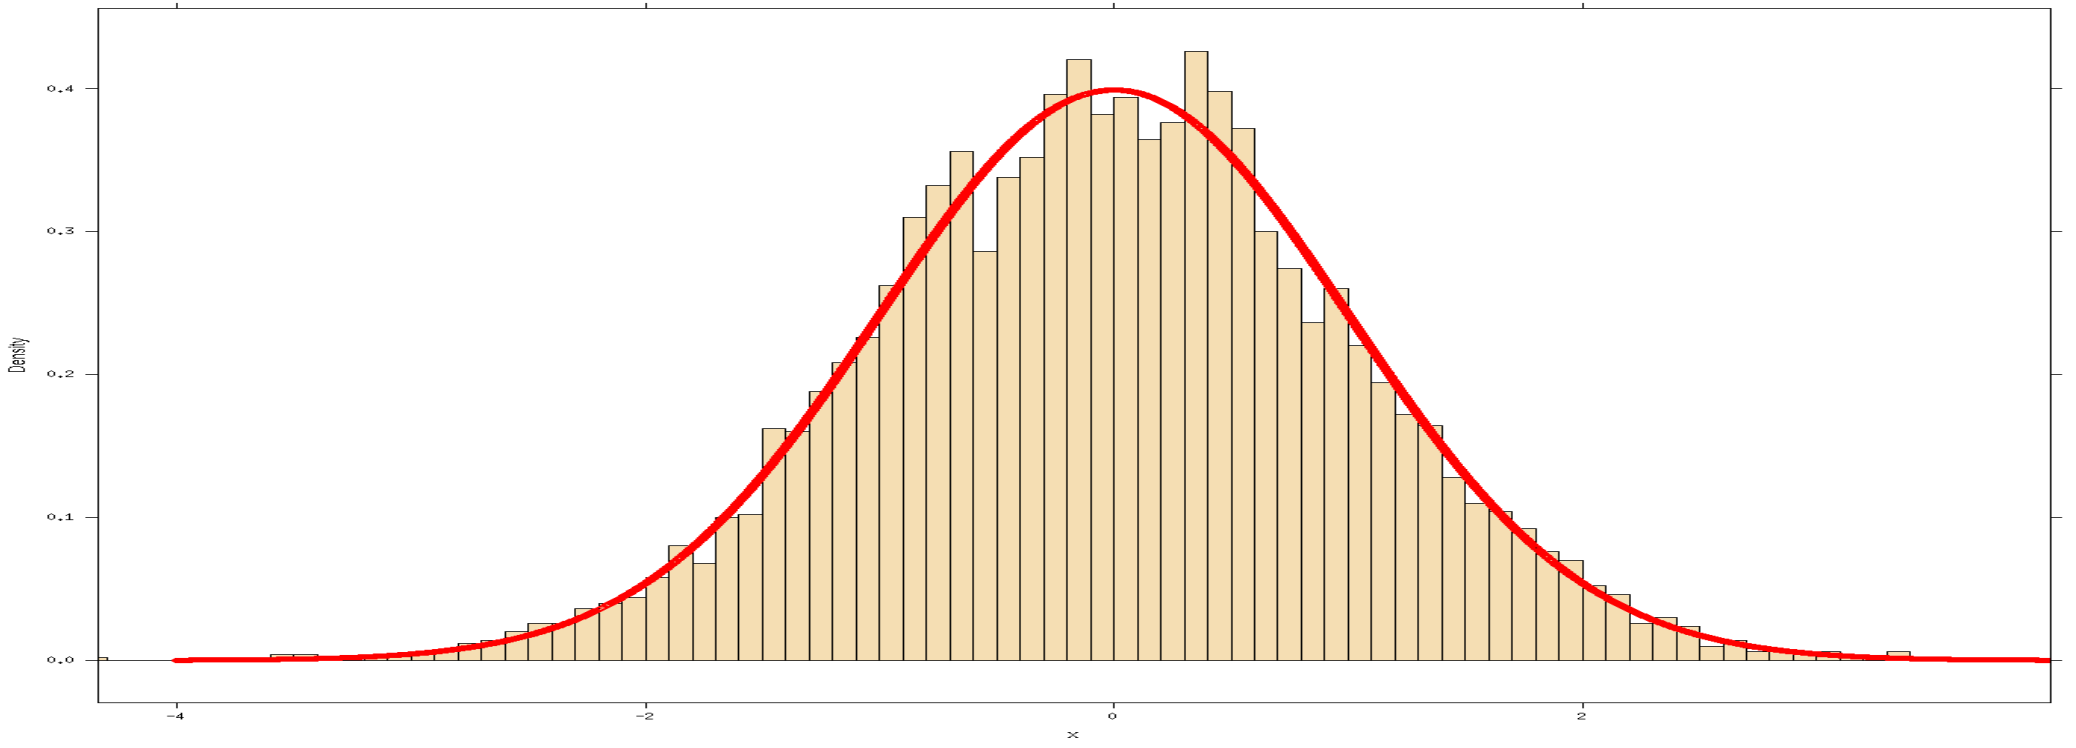

2.6.23 Task Wake Up Time

A toy benchmark result

The data is not normal distribution

2.6.23 Task Wake Up Time

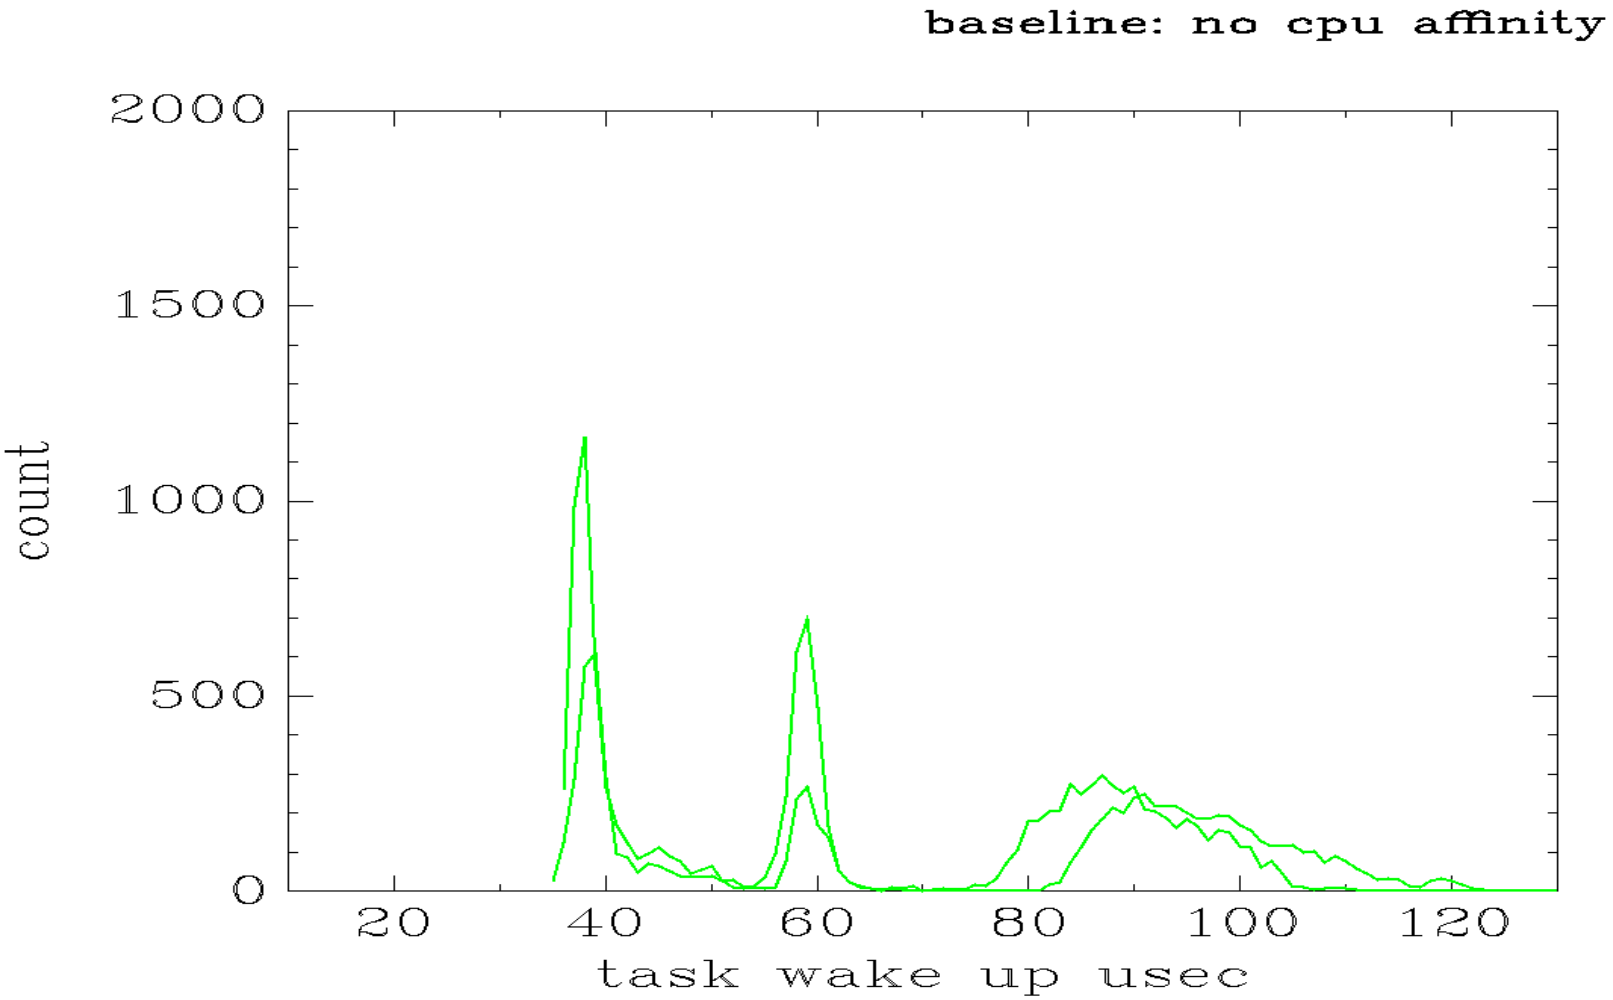

Try to measure the peaks separately

Note that the benchmark was changed

- CPU affinity was set for each process

- Separate data collection for each affinity combination

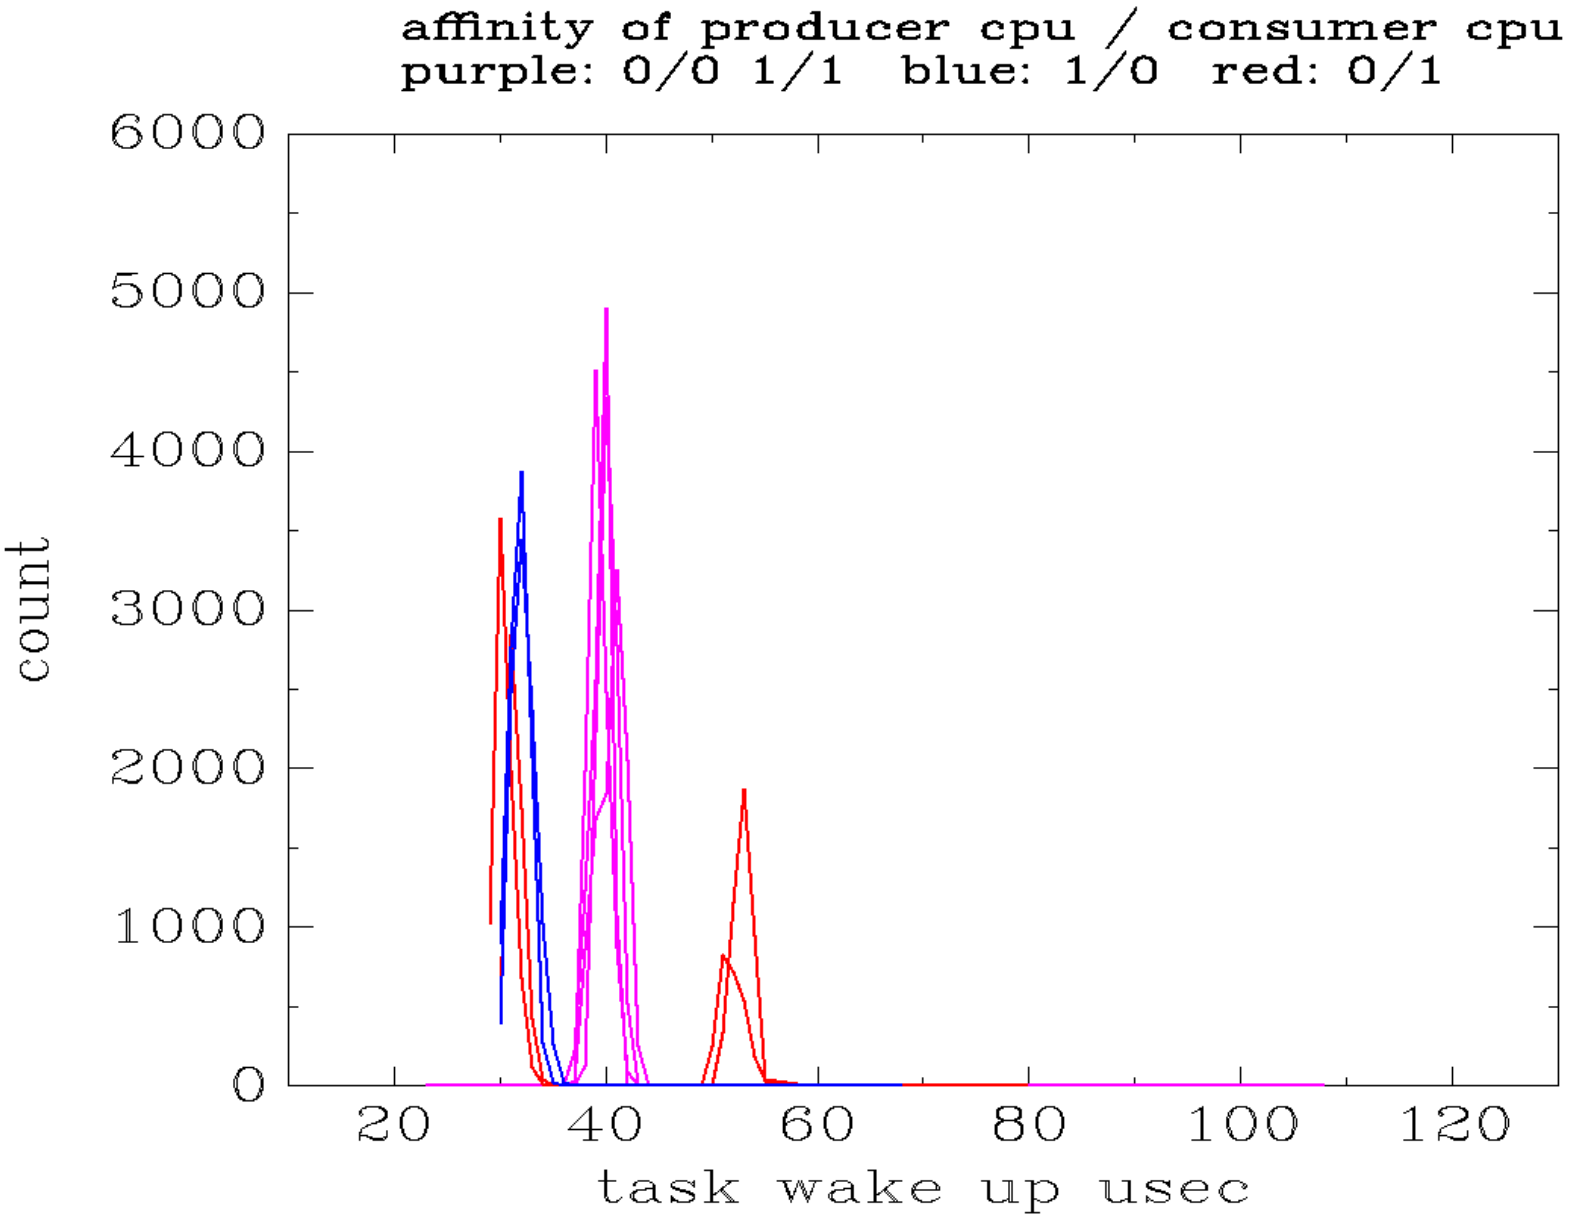
affinity of producer cpu / consumer cpu
purple: 0/0 1/1 blue: 1/0 red: 0/1
6000
5000
4000
3000
2000
1000
0
count
20
40
60
80
100
120
task wake up usec

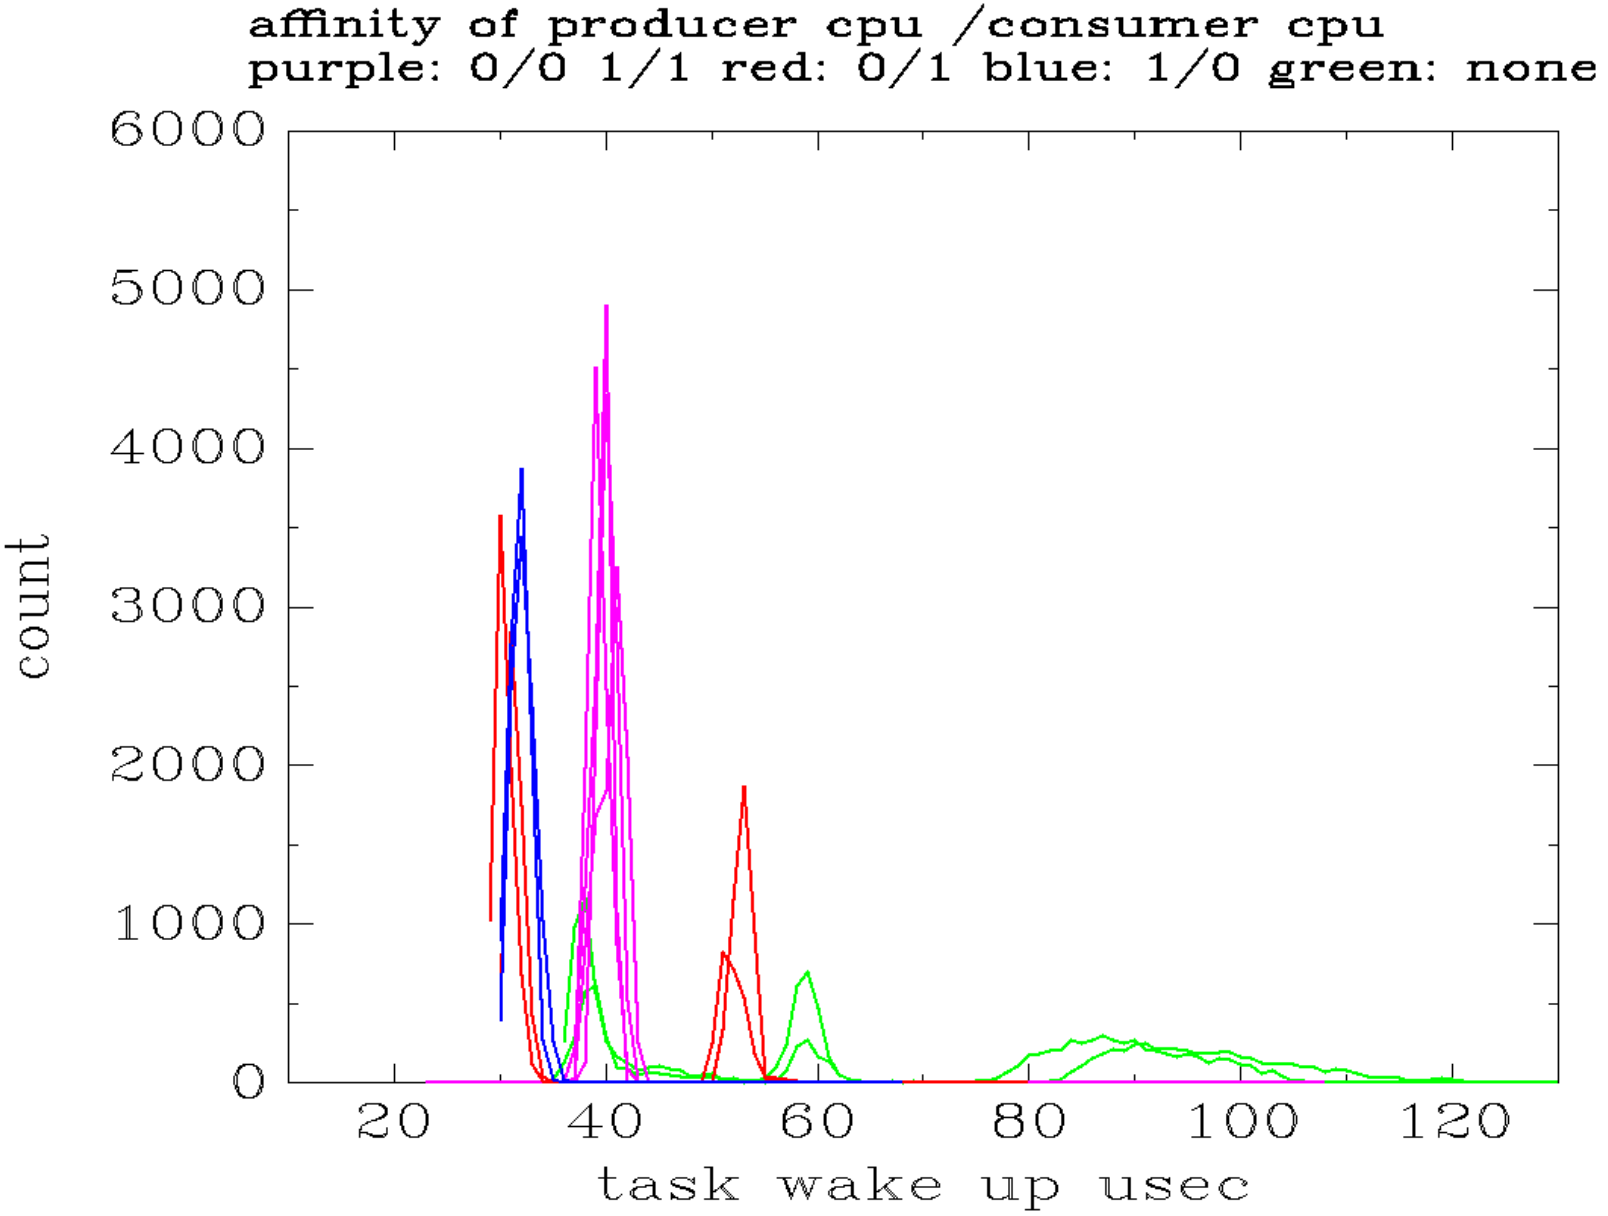
affinity of producer cpu /consumer cpu
purple: 0/0 1/1 red: 0/1 blue: 1/0 green: none
6000
5000
4000
3000
2000
1000
0
count
20
40
60
80
100
120
task wake up usec

Try to measure the peaks separately

Modify the instrumentation to collect separately the duration of

- five worst case interrupt handlers

- do_local_timer()

then subtract them out from all interrupts

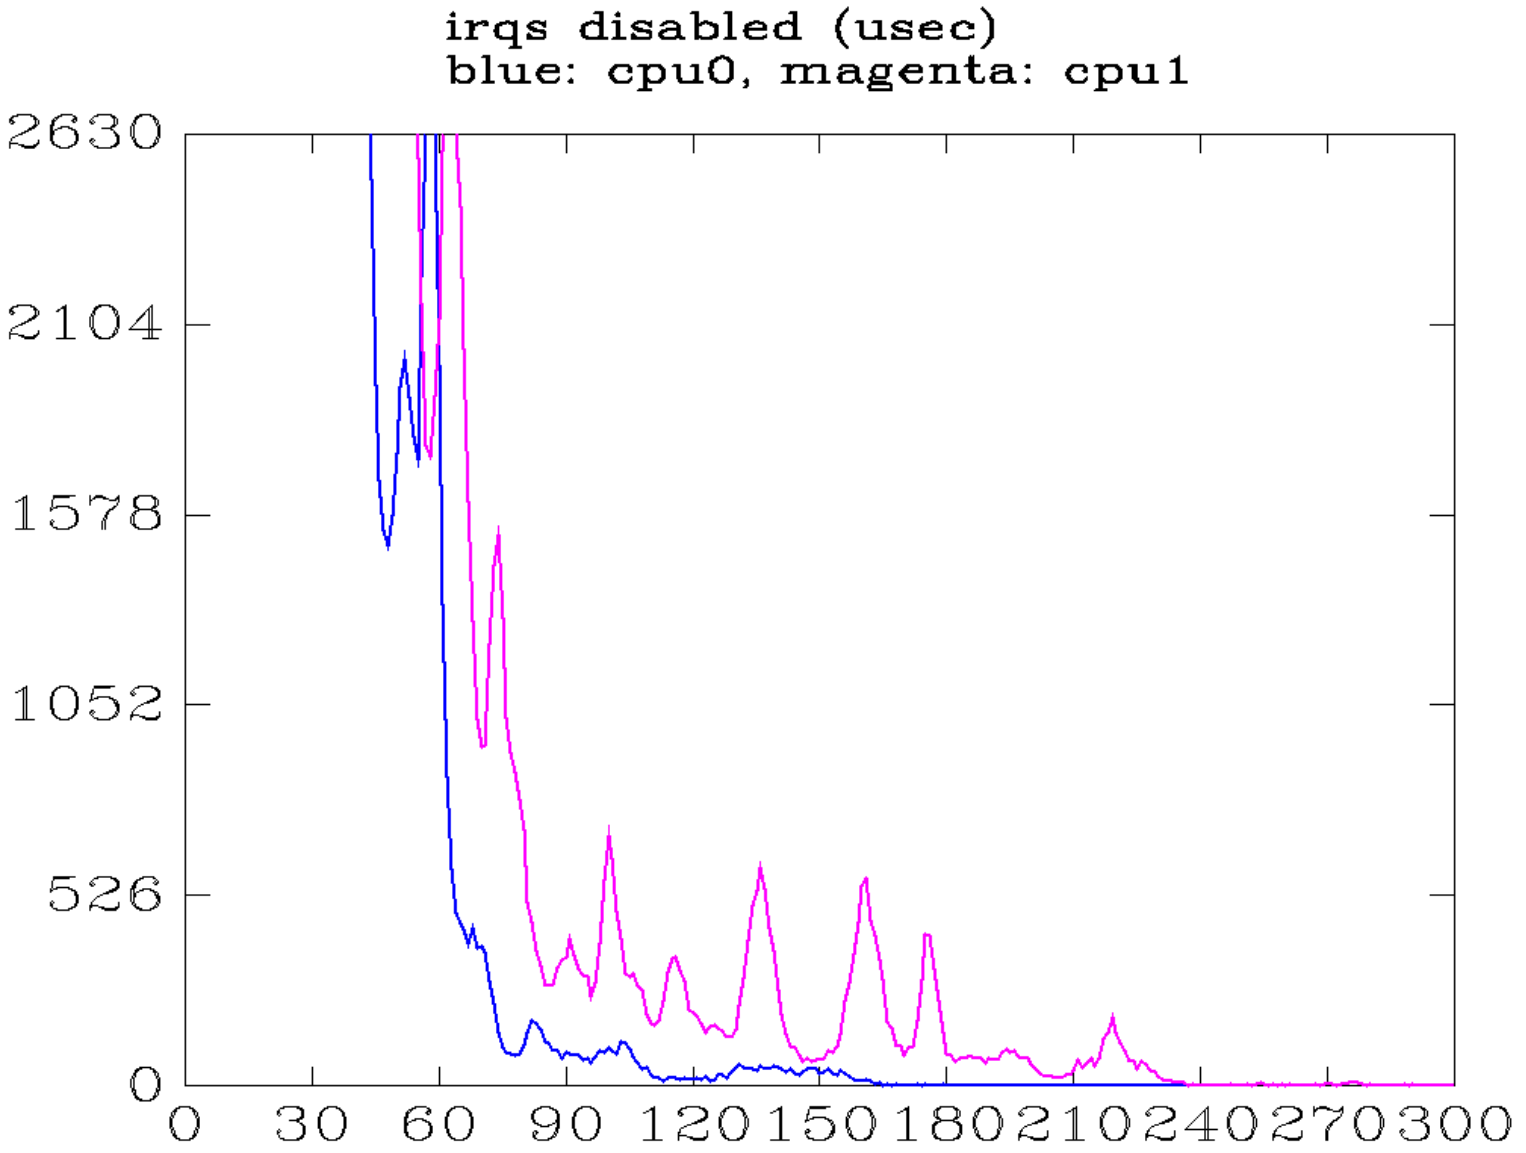
irqs disabled (usec)
blue: cpu0, magenta: cpu1
2630
2104
1578
1052
526
0
0 30 60 90 120 150 180 210 240 270 300

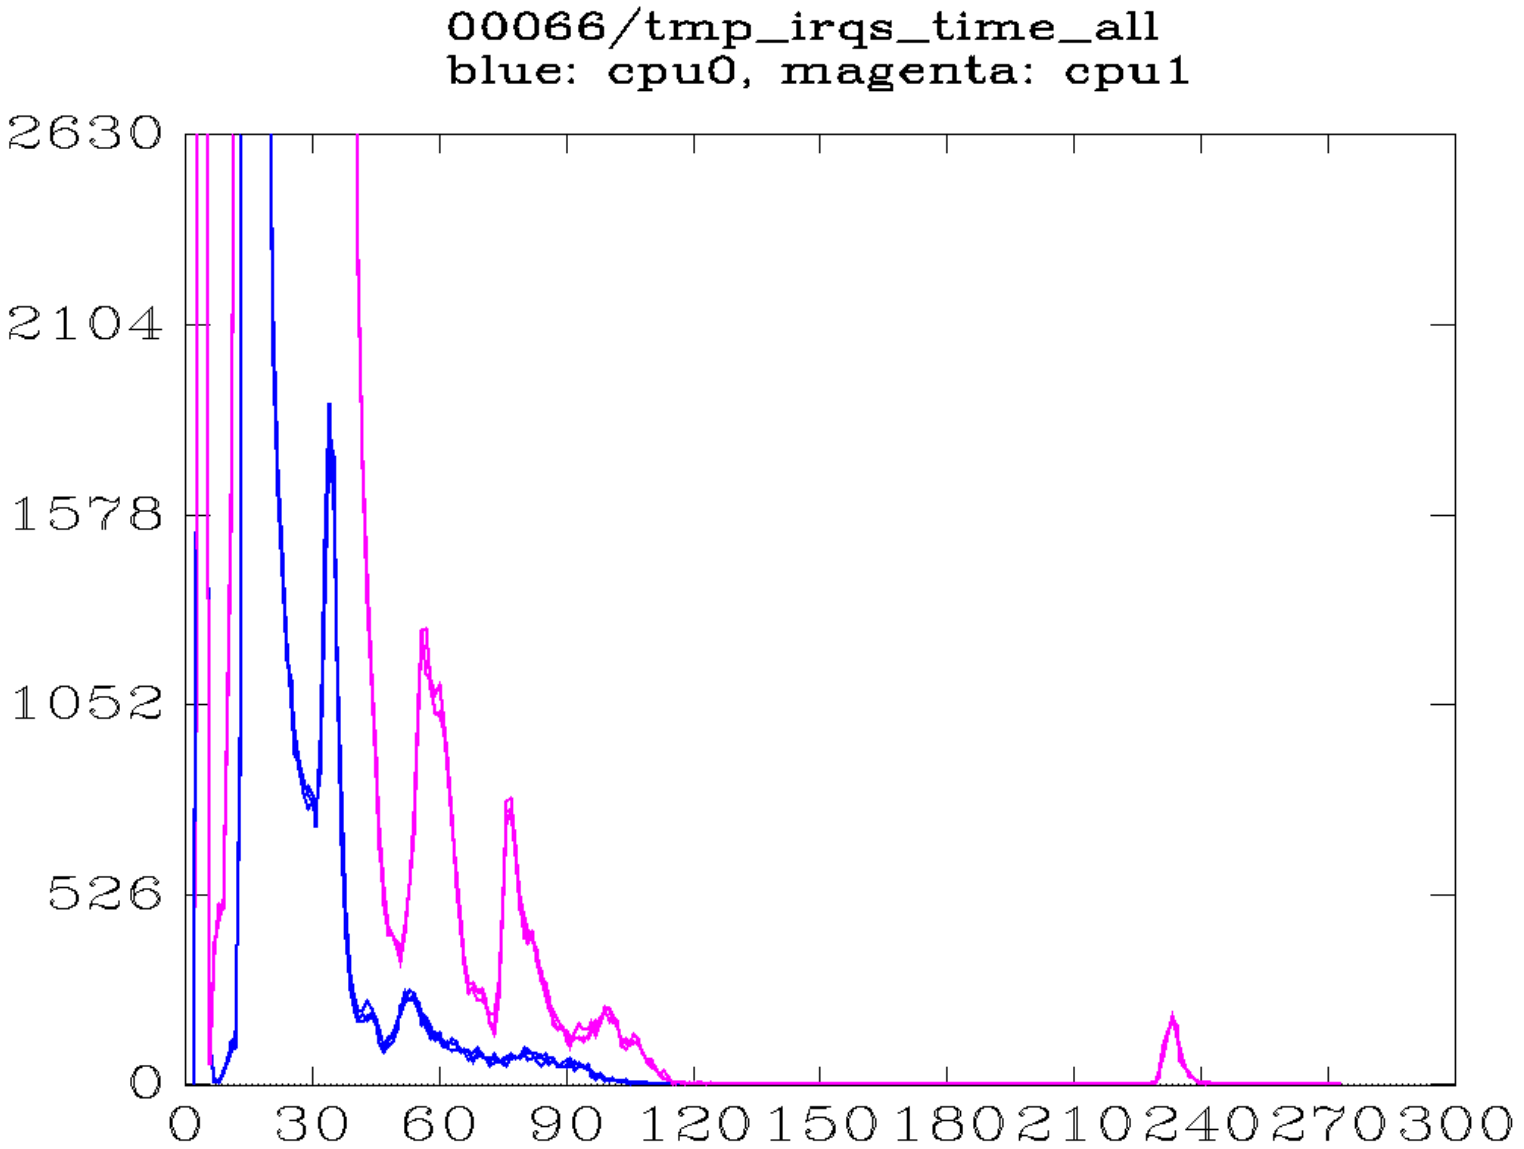
00066/tmp_irqs_time_all
blue: cpu0, magenta: cpu1
2630
2104
1578
1052
526
0
0
30
60
90
120
150
180
210
240
270
300

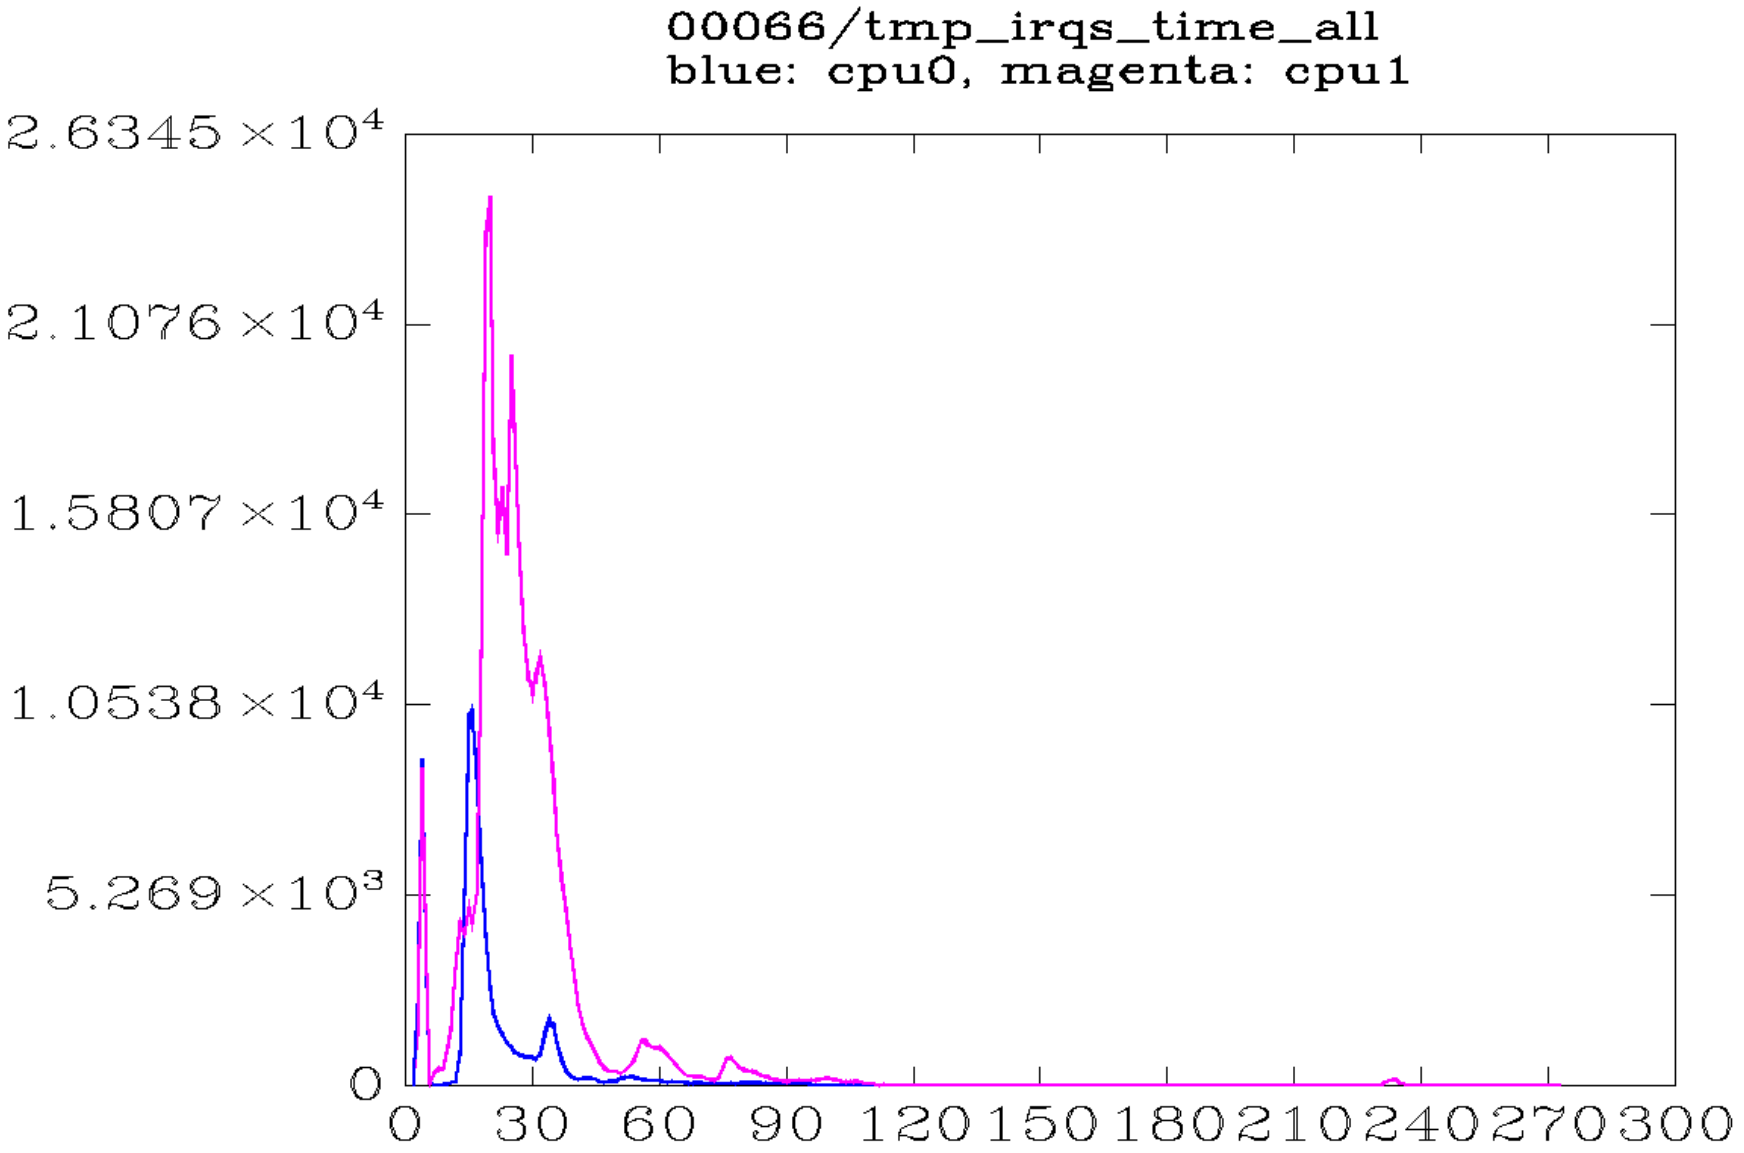
00066/tmp_irqs_time_all
blue: cpu0, magenta: cpu1
2.6345×10⁴
2.1076×10⁴
1.5807×10⁴
1.0538×10⁴
5.269×10³
0
0 30 60 90 120 150 180 210 240 270 300

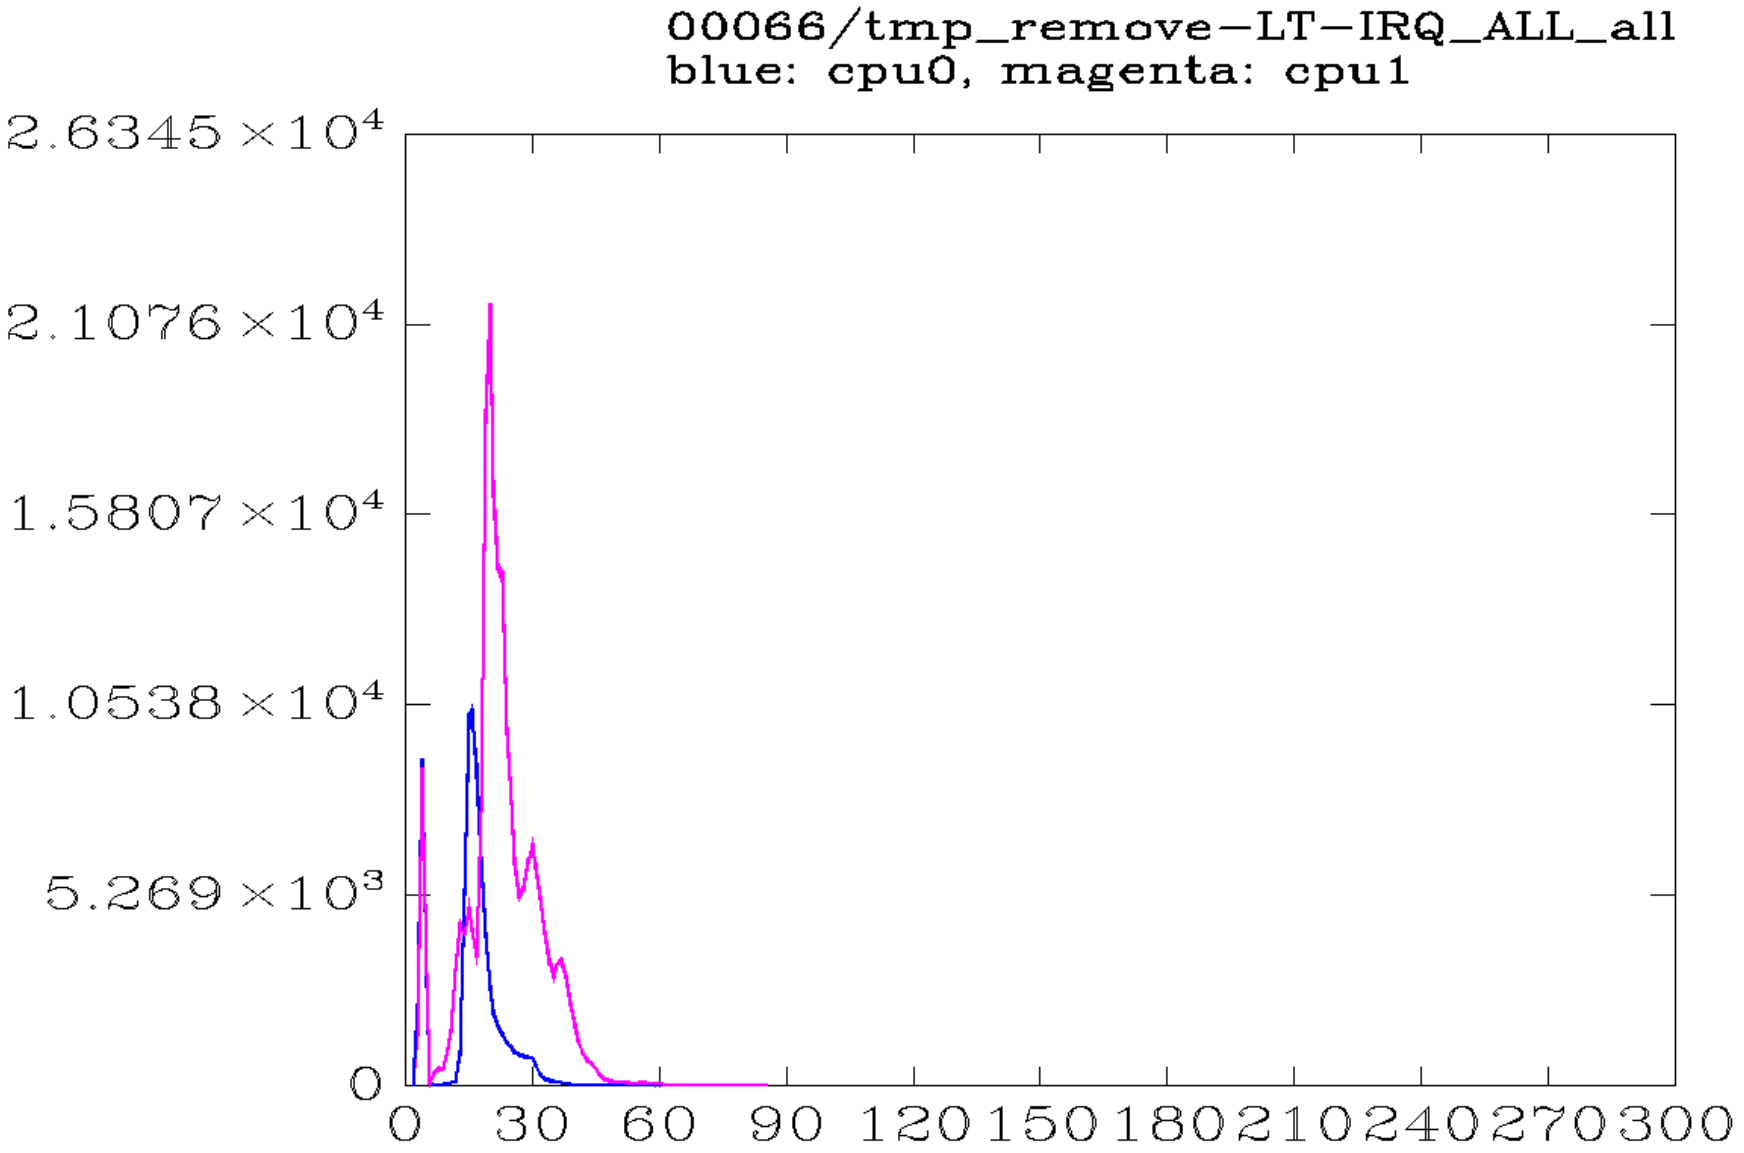
00066/tmp_remove-LT-IRQ_ALL_all
blue: cpu0, magenta: cpu1
2.6345×10⁴
2.1076×10⁴
1.5807×10⁴
1.0538×10⁴
5.269×10³
0
0 30 60 90 120 150 180 210 240 270 300

How did subtracting out some peaks help me?

- The five worst irq handlers are part of the real-time application, so not directly part of my focus when tuning the kernel. Thus some noise removed from my tuning.

- Removing do_local_timer() allowed me to determine whether there were other masked problems lurking and other causes of the long latency tail.

A closer look at one of the peaks

The do_local_timer() peak is also not a normal distribution.

Is there an applicable ad-hoc technique?

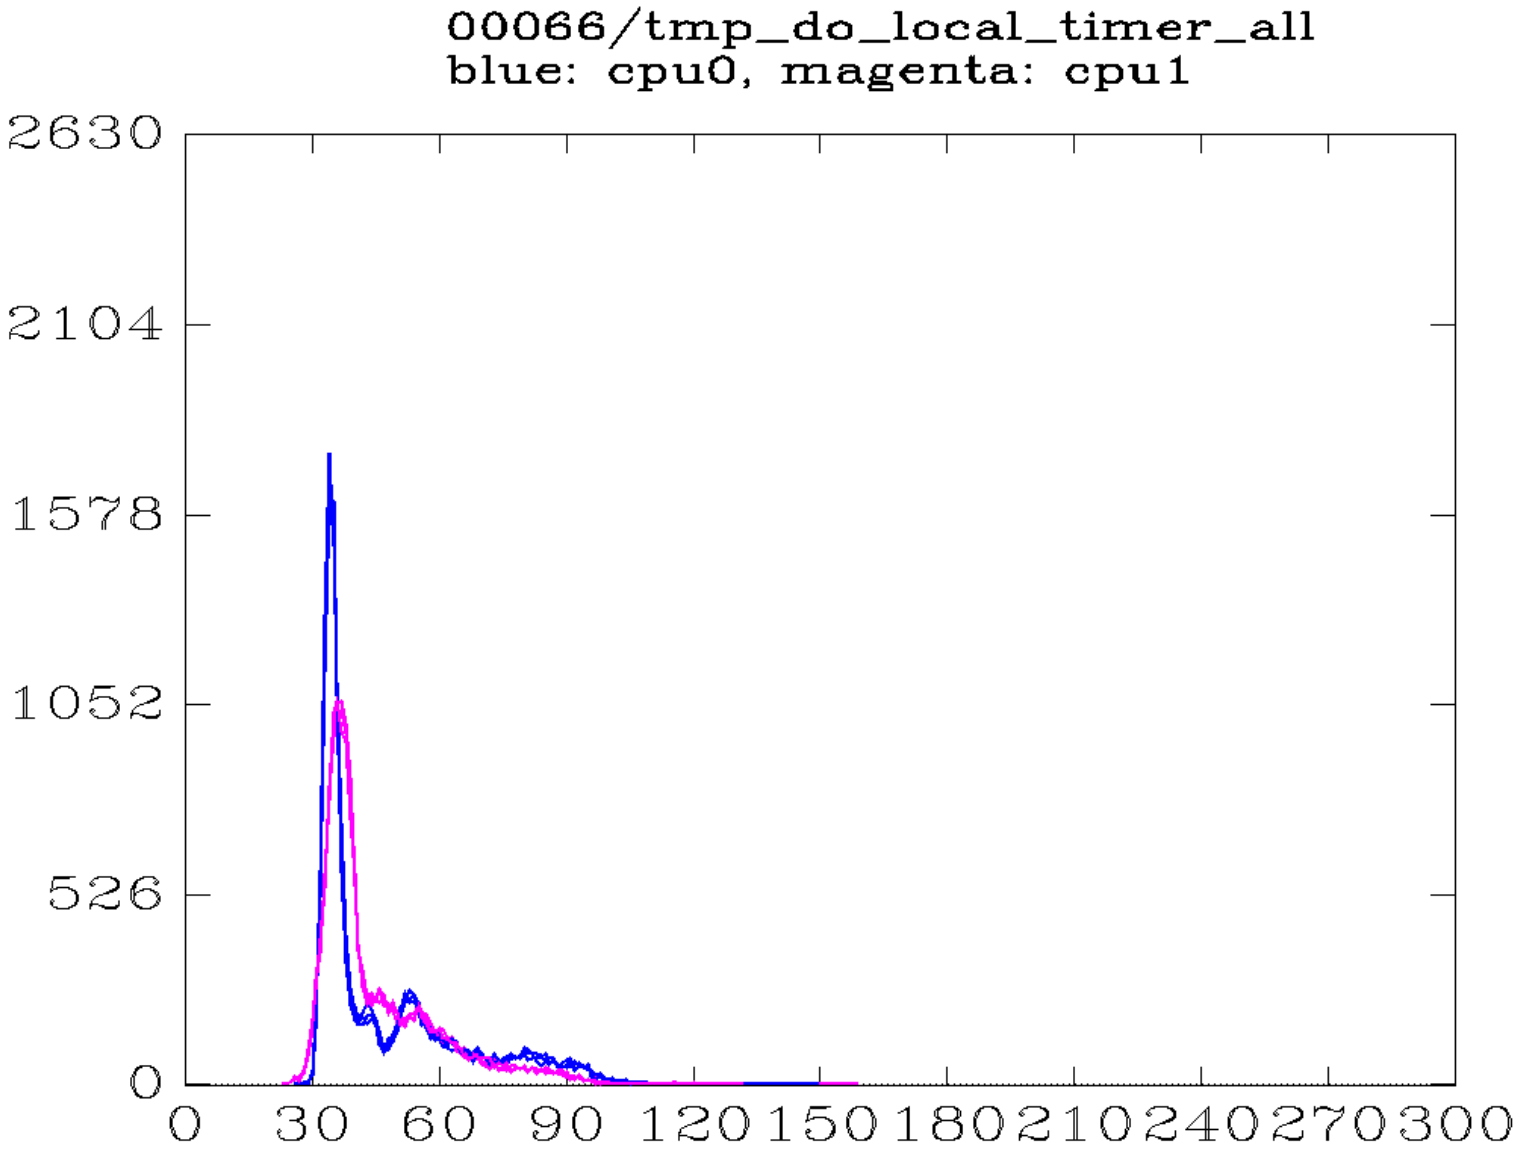
00066/tmp_do_local_timer_all
blue: cpu0, magenta: cpu1
2630
2104
1578
1052
526
0
0 30 60 90 120 150 180 210 240 270 300

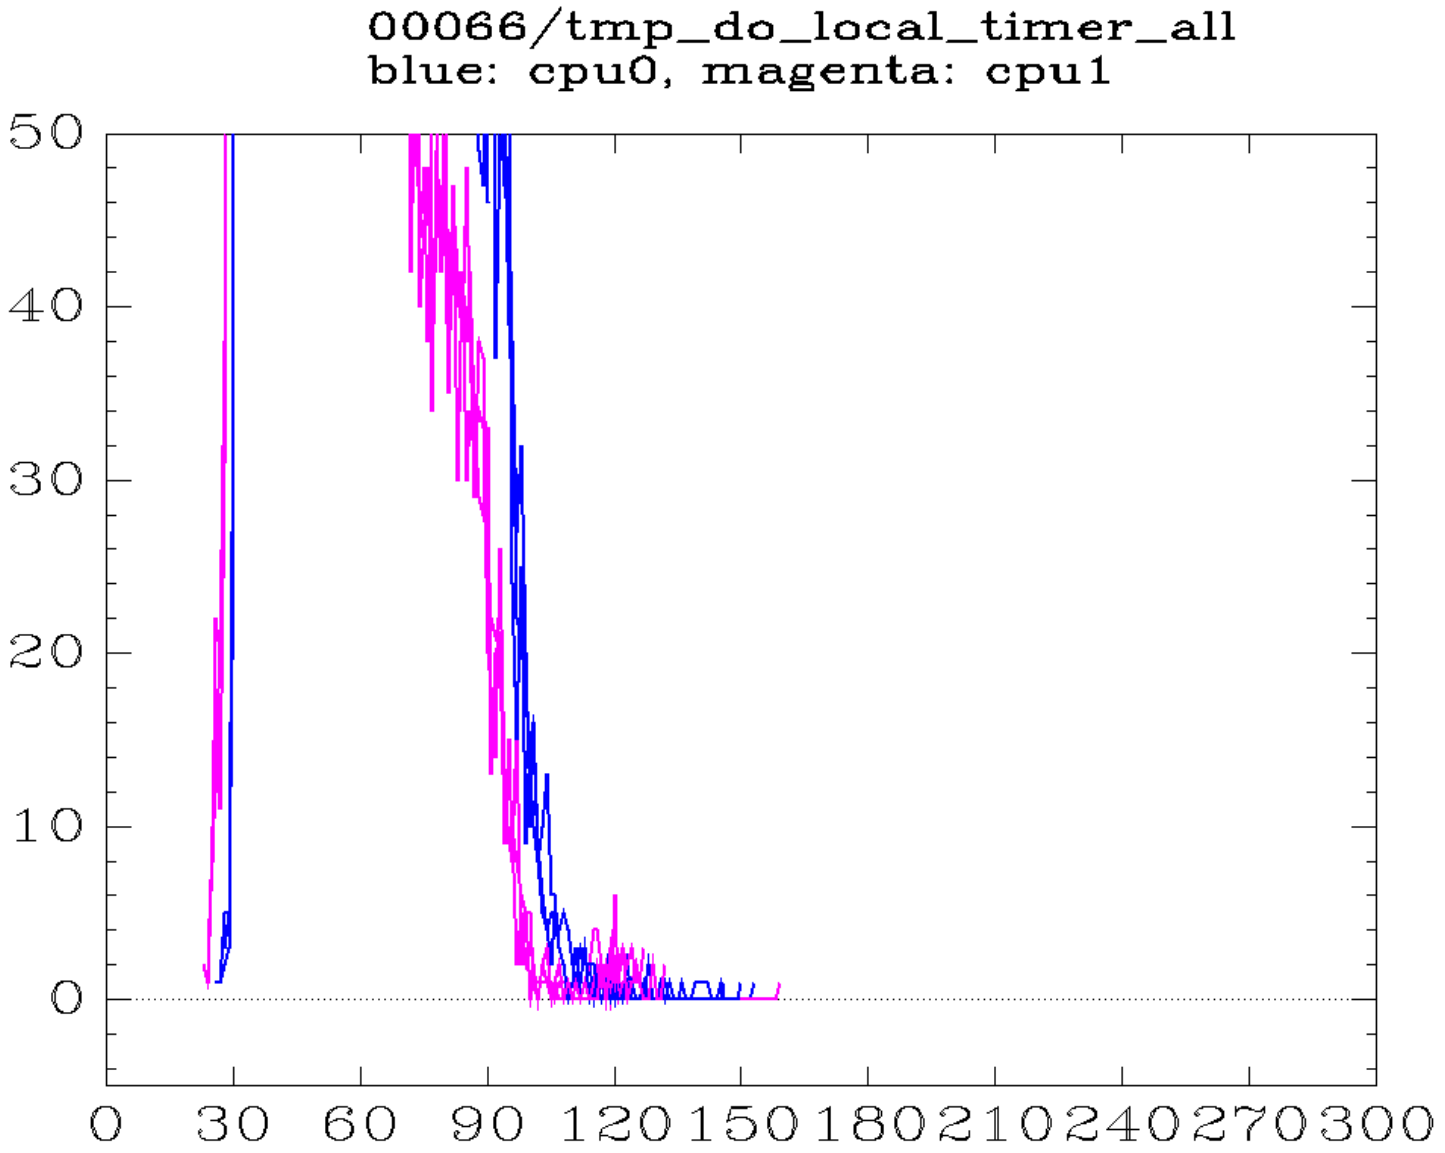
00066/tmp_do_local_timer_all
blue: cpu0, magenta: cpu1
50
40
30
20
10
0
0
30
60
90
120
150
180
210
240
270
300

Recap

- Some statistics (theory)

- Applied the theory to some artificial data

- Some ad-hoc data analysis of real data, not so much using the theory

Recap

- You know enough to be dangerous.

I mentioned that if the data distribution is not normal, then analysis that expects normal data will yield incorrect results.

Another example that will result in incorrect results is if the data from the measurement is not representative of system.

Recap

- You know enough to be dangerous.

There are more traps awaiting you. A deeper study of statistics is highly recommended.

QUESTIONS?