System Performance Evaluation 2

Irene Tallini

Internet of Things A.Y. 2020/2021

Prof: Chiara Petrioli

A Bit More Theory

Recap: confidence intervals

- Is the mean estimate good enough?

- A good way of answering is
  - choosing a **significance level** α,
  - finding the interval (a,b) (**confidence interval** for α) such that:

$$P(a \leq \mu \leq b) = 1 - \alpha$$

- A and b are computed with the Central Limit Theorem

Recap: confidence intervals

- Thus, the confidence intervals for μ can be obtained assuming the mean distribution is gaussian:

$$a = \bar{x} - Z_{(1-\alpha)/2} \cdot \frac{s}{\sqrt{n}}$$
$$b = \bar{x} + Z_{(1-\alpha)/2} \cdot \frac{s}{\sqrt{n}}$$

- Where $Z_{(1-\alpha)/2}$ is the (1-α)/2-quantile of half N(0,1). It is derived from tables.
- Remember s is the standard deviation.

How Long Should Our Simulation Be?

- Or, how large should n be for the mean to be in a smaller enough interval with a larger enough confidence?
- Suppose we want:

$$P(\bar{x} - r\bar{x} \leq \bar{x} \leq \bar{x} + r\bar{x}) = 1 - \alpha$$

- This means:

$$\bar{x} - r\bar{x} \ = \ a \ = \bar{x} - \ Z_{(1-\alpha)/2} \cdot \frac{s}{\sqrt{n}}$$

- And thus:

$$n = \left(\frac{Z_{(1-\alpha)/2} \cdot s}{r\bar{x}} \right)^2$$

Ns-3 Simulation Code Example

Code Structure

- Application (new module): All nodes send to a sink.
- Routing (new module): Flooding.
- MAC and Phy: Wi-fi ad-hoc (IEEE 802.11ah).
- We want to compute the packet latency of this system.

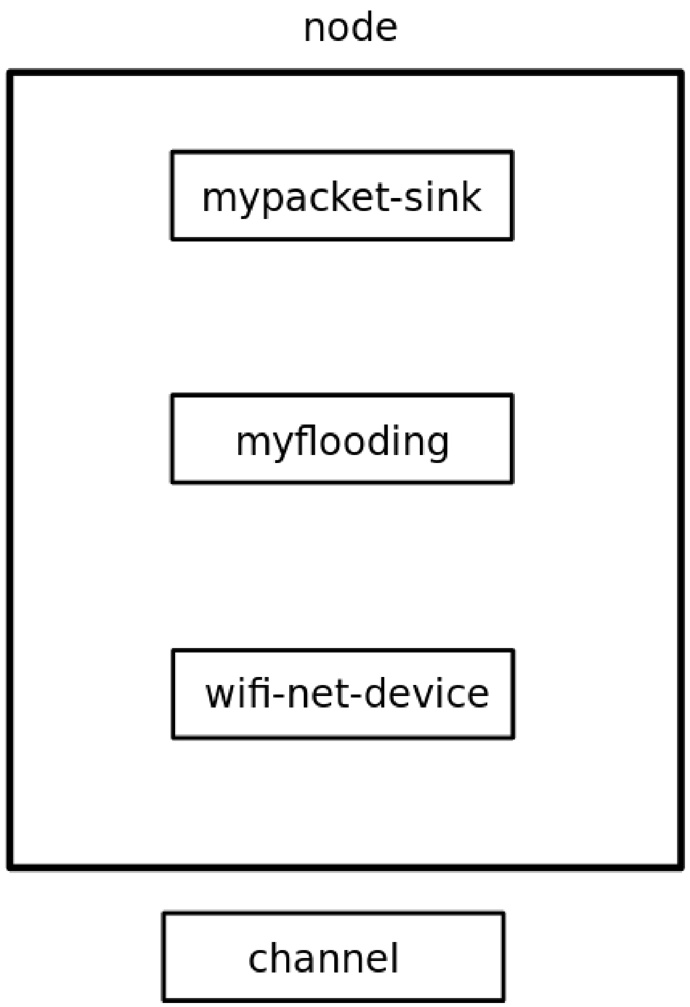

Flooding: Brief Recap

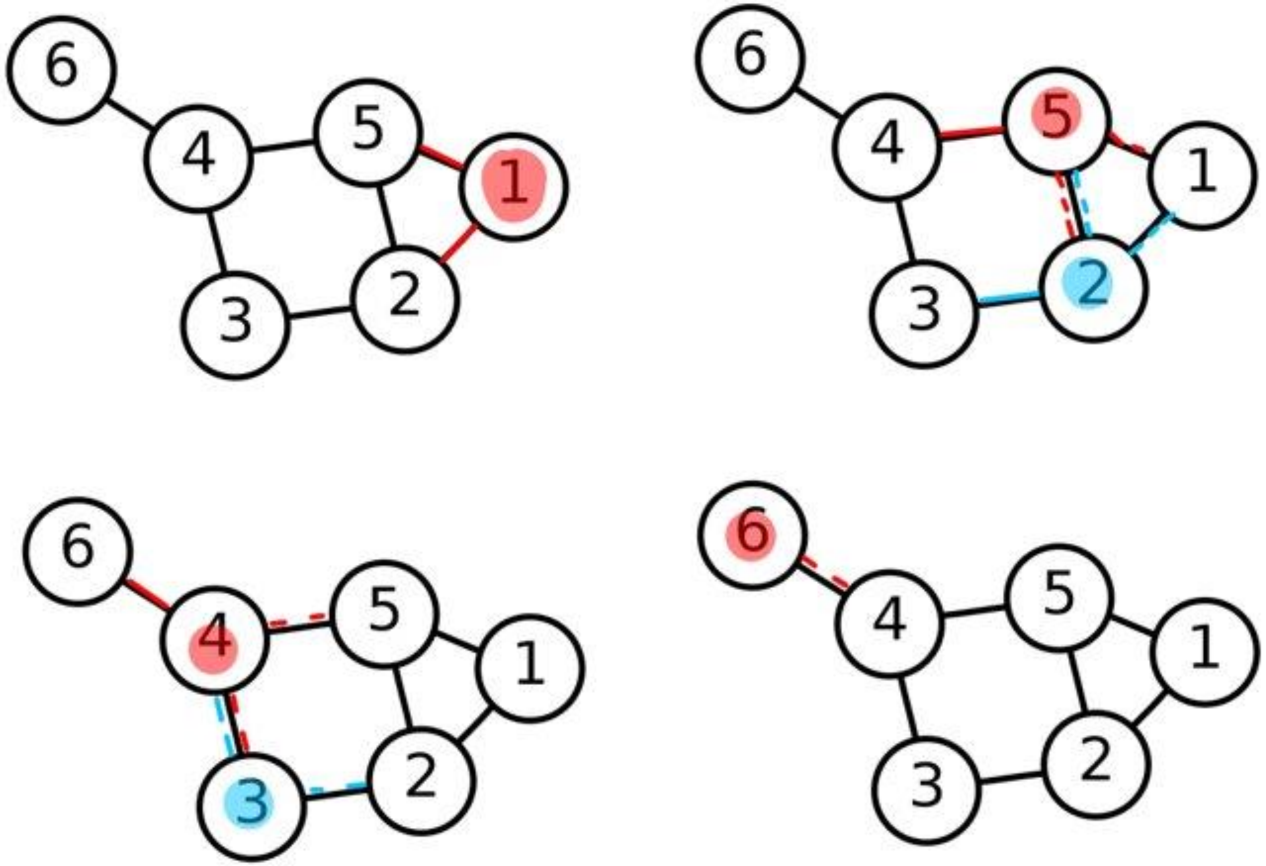

Create New Module

- `cd <your-ns3-folder>`
- Create new module in contrib names mymodule:

```
./utils/create-module.py
    contrib/mymodule
```

- Configure examples building:

```
./waf configure --enable-
    examples
```

- Build:

```
./waf
```

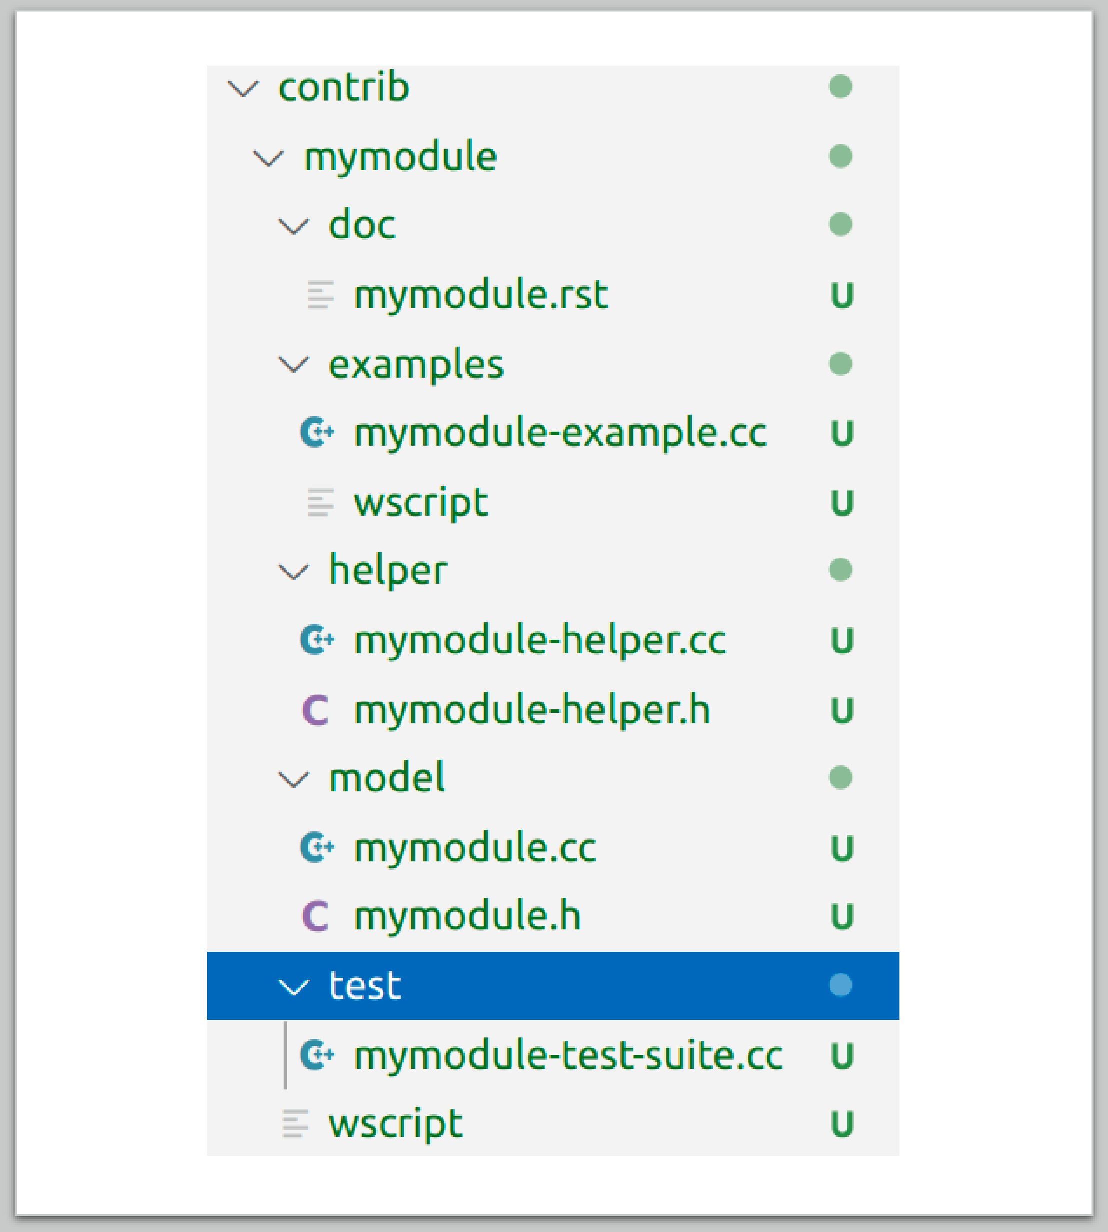

To code!

Latency Analysis

Note the unity of measure is microseconds (10^{-6} seconds)

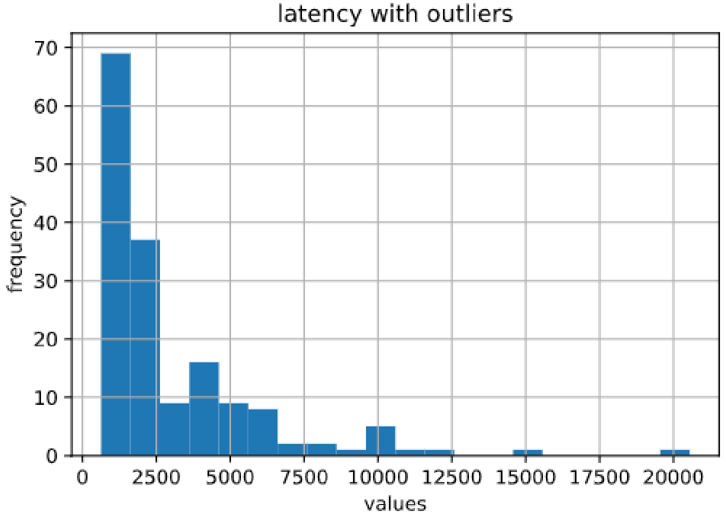

	latency
count	160.000000
mean	2725.993750
std	2517.251997
min	627.000000
25%	642.000000
50%	1926.000000
75%	4160.500000
max	11606.000000

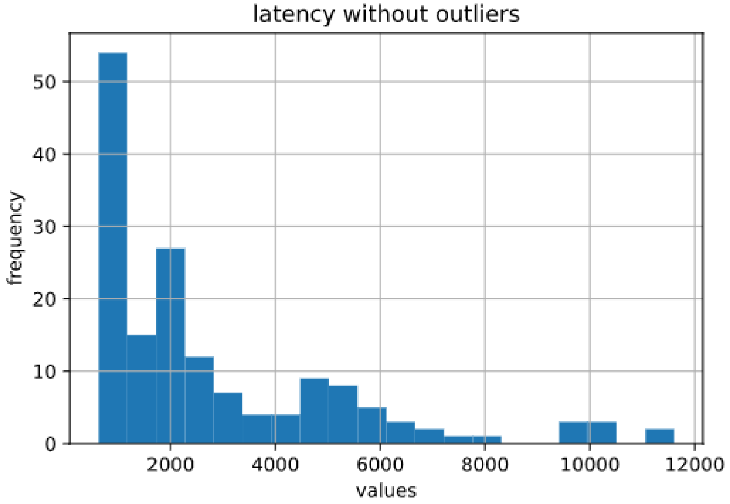

Distribution

Median and Quantiles Representation

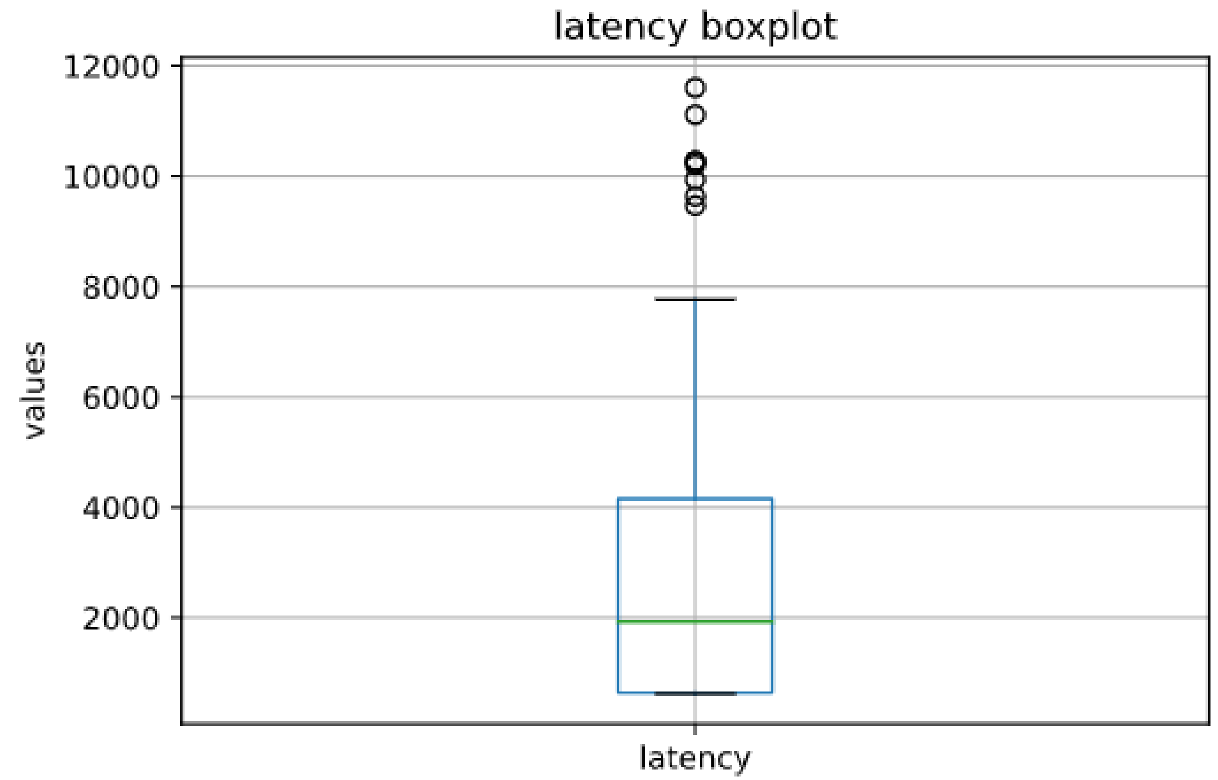

Median representation: is the number of samples big enough?

- We choose r = 0.15 and 1-α=0.9.
- This means $n = \left(\frac{Z_{0.45} \cdot s}{0.1\bar{x}}\right)^2 = \left(\frac{1.65 \cdot s}{0.15\bar{x}}\right)^2 \simeq 121\frac{s^2}{\bar{x}^2}$
- Our test has:

```
mean = 2725.993750
standard deviation = 2517.251997
n = 160
squared std/mean = 1.172725
```

- 121*1.17 = 141.57 < 160. So, our n is big enough!

Our Estimate For the Mean Latency

- We can say that, with 90% confidence: $\bar{x} - 0.15\bar{x} \leq \bar{x} \leq \bar{x} + 0.15\bar{x}$
- I.e., $2317 \leq \bar{x} \leq 3134$
- Moreover, our standard deviation is 2517 microseconds.