RIPE NCC
RIPE NETWORK COORDINATION CENTRE

# Distance/RTT Analysis for Ukraine

RIPE NCC Days in Kyiv | 2019-09-26 | Alex Semenyaka

# Round Trip Time (RTT)

$$RTT = d_1 + d_2 + d_3$$

$$d_1 = \sum_i (propagatation\ delay)_i$$

Channels

$$d_2 = \sum_i (equipment\ delay)_i$$

Hardware

$$d_3 = \sum_i (queue\ delay)_i$$

Load

$$min(d_3) = 0 \quad (!)$$

# How RTT Depends on Distance

$$RTT = min(\{RTT_i\}) = d_1 + d_2$$

$$d_1 = \sum_i \frac{distance_i}{c_{fibre_i}}$$

Length of the fiber sector

Velocity of speed in this fibre

The speed of light in different fibres differs insignificantly (±10% max), so:

$$d_1 \approx \frac{total\ distance}{c_{fibre}}$$

# How RTT Depends on Distance

- Les Cottrell, Warren Matthews and Connie Logg from Stanford University, 2000-2002 гг.:

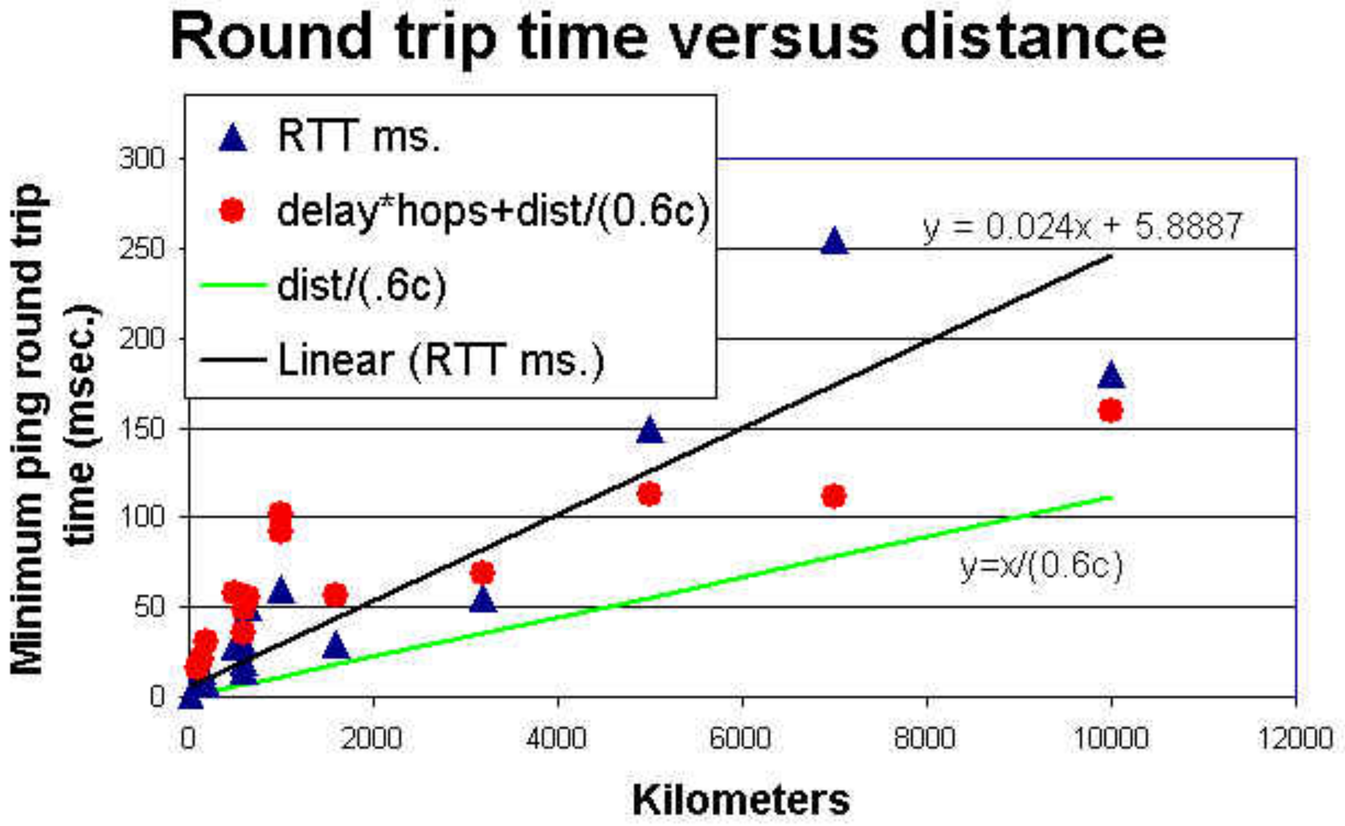

# How RTT Depends on Distance

- For 2000 (old slow equipment, no ASICs etc):

$$d_2 = 5.9\ ms$$

- In 2019, we can expect much less value
- And if so, the minimal RTT value between two nodes is approximately proportional to the distance between them:

$$RTT \approx \frac{total\ distance}{c_{fibre}}$$

# RTT Between Cities

- The speed of light in fibre is 60-70% of the speed of light in a vacuum, i.e.

$$c_{fibre} \approx 200 \frac{km}{ms}$$

- In other words:

**Each 100 kilometres add 1 ms to RTT**

- Therefore, in the case of perfectly straight fibre the ratio

$$S = \frac{RTT_{real}}{RTT_{ideal}} = 100 \frac{RTT, ms}{distance, km}$$

is equal to 1

# How RIPE Atlas Can Help

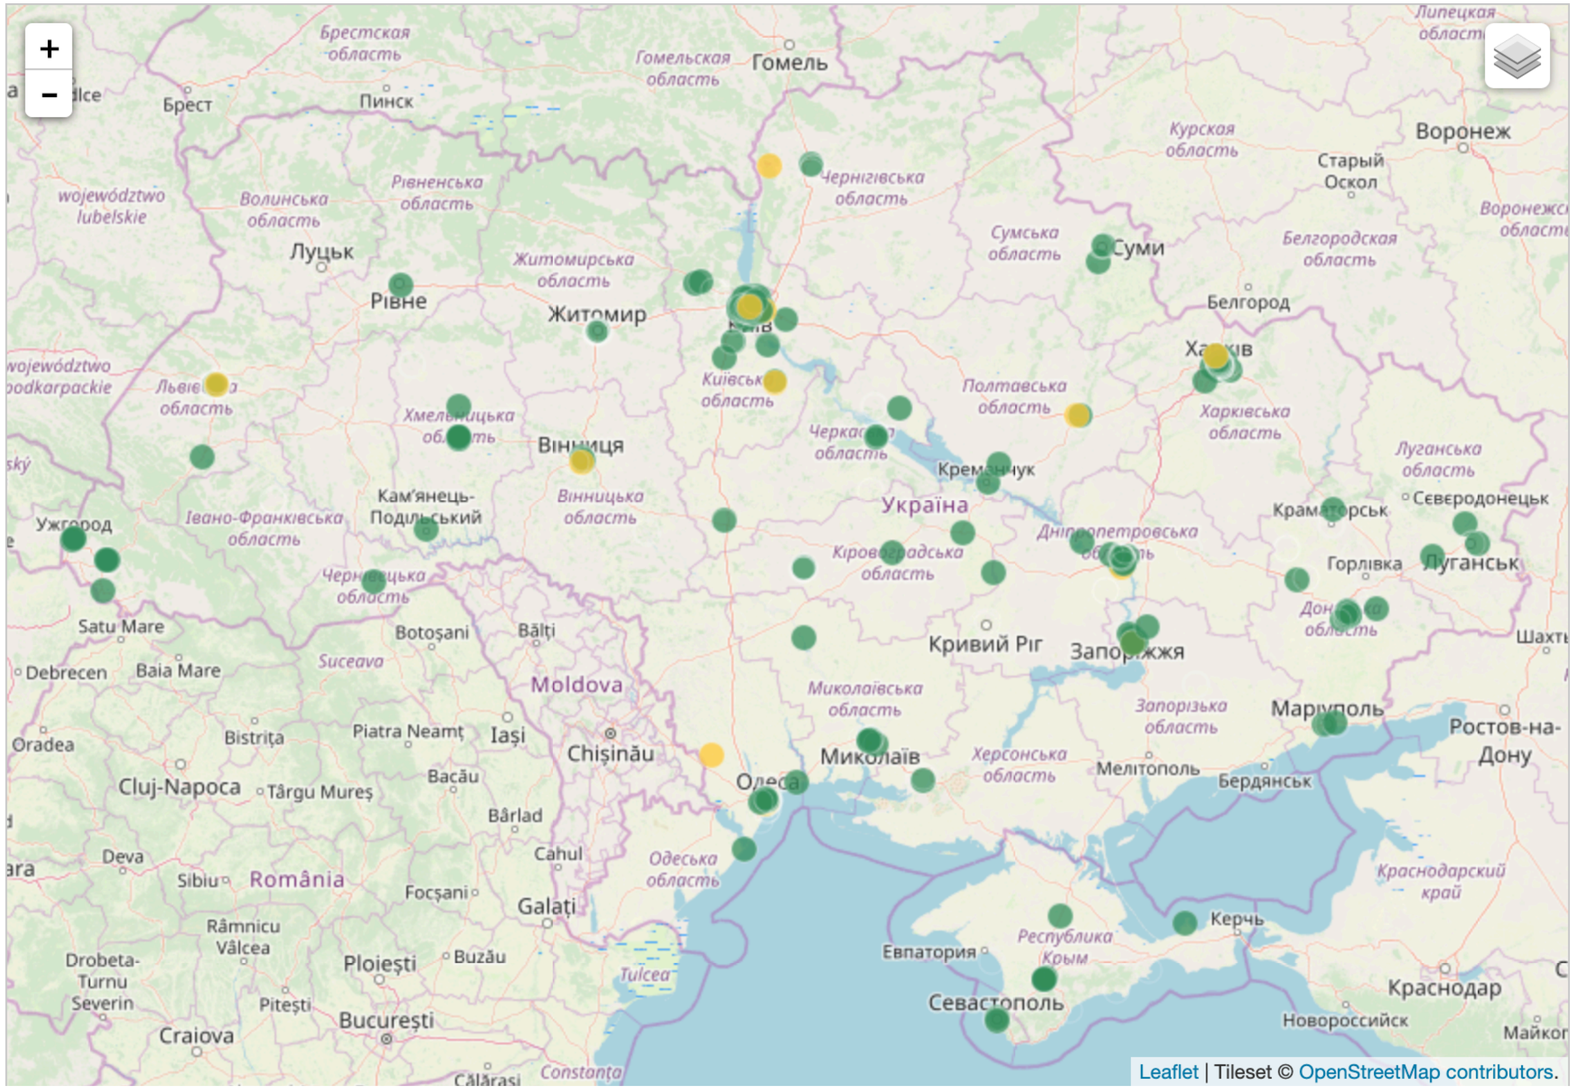


**https://atlas.ripe.net/results/maps/network-coverage/?filter=b%27Ukraine%27+(ua)**

# How RIPE Atlas Can Help

- There are approximately 200 working probes in Ukraine
- They are distributed all over the country
- We can start measuring the RTT between samples and build a matrix of delays between cities
- By comparing the delay and distance, we can conclude if the cabling is geographically optimal

# Prerequisites

- **No “special features” of the RIPE NCC employee were used to obtain results**
  - (…and I am not sure that such a feature even exists)
- **The measurements were performed with a personal account on my own accumulated credits**
  - (…yes, it took a lot of them)
- **Nevertheless, nothing would have happened without interaction with the RIPE Atlas team**
  - (…they answer the questions from the community exactly the same as they do mine!)

# Setup

- **3,000 lines on Python3**
- **Active usage of the library ripe.atlas.cousteau (from github)**
- **Five new classes were added to this library (local repository yet)**
- **Found and fixed by the team:**
  - Two mistakes in the documentation
  - Four errors in the system code
- **70 million measurements were done in one month**

# Disclaimer

**The study is indicative in nature, it can not be perceived as “the ultimate truth”:**

- We don’t have samples for every Autonomous System in Ukraine, nor for each IPv4 prefix, so we can’t “see” everything
- Some phenomena may have been temporary; today the picture may have changed
- Despite checking the results, it is still possible that peculiar behaviour of RIPE Atlas samples and target nodes may have had an influence

# The Points Selection

- **The obvious criteria is the size of the cities**
  - Top-10 excluding Crimea and occupied districts of Donetsk and Luhansk regions
- **Problem: in some cities we have just a few probes:**
  - Mariupol: one probe
  - Kherson: one probe
  - Poltava: one probe
  - Kryvyi Rih: no probes at all (thus this town was excluded from the research)
- **Two cities were aggregated into the agglomeration:**
  - Dnipro-Zaporizhia - 122 км

# The Points Selection

- **In total, including agglomerations, there are 10 logical “areas”:**
  - Kyiv, Kharkiv, Odessa, Dnipro, Lviv, Mykolayv, Mariupol, Vinnytsya, Kherson, Poltava
- **Known problem:**
  - Remembering the low number of probes in some cities, we have to enrich our sampling

# How to Enrich Sampling

- **RTT(IP1, IP2) = RTT(IP2, IP1)**
  - Meaning we need to find nodes in each region that can be “pinged” from our RIPE Atlas probes
- **The primary criteria:**
  - The host responds to ICMP (nmap scanning was used)
  - Geographical “binding” in the RIPE Database
  - Direct and reverse name resolution match
  - Additional “bonus”:
    - The IP address belongs to the local website
    - Geo-referencing inside the domain name
- **At this stage, 30 nodes were selected**
- **In total, 127 probes and 153 (127+26) targets**

# Verification of Results

- **At each stage, results cross-verified, questionable nodes and probes excluded**
- **Examples:**
  - The probe, which RTT to itself exceeds two ms
  - RIPE Atlas probe marked as from Odessa which has high RTT to anywhere (probably, it is actually located somewhere else)
- **Some other oddities:**
  - Probes in MOAS (Multiple-Origin Autonomous System)
  - Rare results like 77.47.130.162 -> 89.162.136.106:
    - Attempt 1: 8.69552 ms
    - Attempt 2: ‘Network unreachable’ ⇠ i.e. ICMP unreach was received; who sent it?
    - Attempt 3: 8.342565 ms
- **In total, eight IP addresses were excluded**

# Processing of Results

- **Between each pair of IP addresses we take a minimum RTT**
- **How do I aggregate data between different areas or Autonomous Systems?**
  - The minimum RTT value (MinRTT)
  - The median RRT (MedRTT) - which is the arithmetic mean of the central elements of the sorted set:

0.2, 0.4, 0.5, 1, **1, 1,** 1, 4, 10, 25

Median: **1**

Arithmetic mean: **5.4**

- **Let's calculate S for minRTT and medRTT and build “heat maps”**

# Local RTT Inside AS

- **Degenerate case: RTT inside the city and within each Autonomous System there**
- **Important as a stage intermediate inspection**
- **32 results**
- **As expected:**
  - minRTT almost never exceeds three ms
  - MedRTT almost never exceeds five ms
- **Not expected (and probably should be fixed):**
  - Kyiv, AS15895 (Kyivstar PJSC): MinRTT=5.5ms, MedRTT=16.2ms
  - Kyiv, AS25229 (Kyivski Telekomunikatsiyni Merezhi LLC, aka Volia): MinRTT=7ms, MedRTT=13ms
  - Kharkiv, AS25229: MinRTT=3.5ms, MedRTT=3.5ms

# RTT Inside Cities

- **We can process data between all possible IP addresses**
  - This approach seems to correlate with users' experiences since they are randomly distributed among operators
  - These columns of the table below are labeled "between IP"
- **We can first group the addresses by ASN, and thus not take into account RTT within a single Autonomous System**
  - This approach reflects the operator's view on connectivity
  - These columns of the table below are labeled "between ASes"

# RTT Inside Cities

| Area | Min RTT between IPs | Median RTT between IPs | Min RTT between ASes | Median RTT between ASes |
|---|---|---|---|---|
| Дніпро | 0,24 | 17,92 | 1,24 | 17,41 |
| Київ | 0,17 | 14,94 | 0,38 | 1,38 |
| Львів | 0,50 | 31,94 | 0,50 | 16,13 |
| Миколаїв | 0,22 | 23,02 | 20,22 | 24,15 |
| Одеса | 0,43 | 10,75 | 0,43 | 14,59 |
| Харьків | 0,62 | 7,57 | 0,65 | 11,63 |

# RTT Inside Cities: Conclusions

- **Overall, the situation doesn’t look good:**
  - However, all but one MinRTT are less that two ms
    - Confirmation of the hypothesis that $d_2 \approx 0$
  - MedRTT between ASes never exceeds 25 ms
- **Nevertheless, there is a huge space to improve the situation**
  - Mykolayv
  - RTT=15ms inside the city is still too much

# MinRTT Between Cities: matrix

| Area/Area | Вінниця | Дніпро | Київ | Львів | Маріупол | Миколаї | Одес | Полтав | Харьків | Херсон |
|---|---|---|---|---|---|---|---|---|---|---|
| **Вінниця** | X | 2,21 | 1,53 | 3,11 | 2,72 | 3,72 | 3,12 | 2,01 | 1,62 | 3,16 |
| **Дніпро** | 2,21 | X | 1,73 | 1,62 | 7,82 | 2,81 | 1,34 | 7,24 | 2,08 | 2,08 |
| **Київ** | 1,53 | 1,73 | X | 1,52 | 2,14 | 2,37 | 1,61 | 1,58 | 1,53 | 1,82 |
| **Львів** | 3,11 | 1,62 | 1,52 | X | 2,20 | 2,51 | 1,56 | 1,53 | 1,50 | 2,07 |
| **Маріуполь** | 2,72 | 7,82 | 2,14 | 2,20 | X | 5,43 | 3,89 | 5,94 | 2,68 | 5,47 |
| **Миколаїв** | 3,72 | 2,81 | 2,37 | 2,51 | 5,43 | X | 1,93 | 4,01 | 3,22 | 2,47 |
| **Одеса** | 3,12 | 1,34 | 1,61 | 1,56 | 3,89 | 1,93 | X | 2,69 | 0,11 | 2,44 |
| **Полтава** | 2,01 | 7,24 | 1,58 | 1,53 | 5,94 | 4,01 | 2,69 | X | 6,63 | 3,95 |
| **Харьків** | 1,62 | 2,08 | 1,53 | 1,50 | 2,68 | 3,22 | 0,11 | 6,63 | X | 2,31 |
| **Херсон** | 3,16 | 2,08 | 1,82 | 2,07 | 5,47 | 2,47 | 2,44 | 3,95 | 2,31 | X |

**…Well, doesn’t look so easy to comprehend, right?**

# Heatmaps

- Heatmapping is the natural and easy way to visualise matrices
- To build heat maps, we need thresholds
- Let’s analyse typical cases of the geo-topology and corresponding value of S

# What is “S”? Example: direct fibre

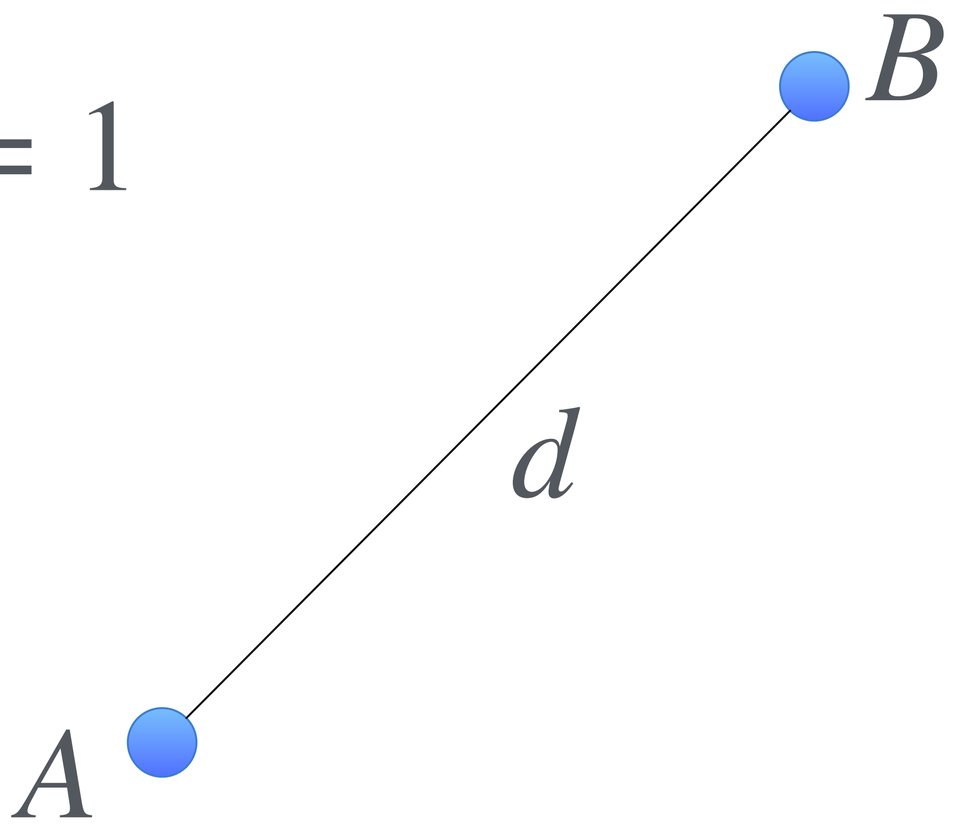

# Now, add some angles…

$$S = \sqrt{2} \simeq 1.41$$

# …or circle

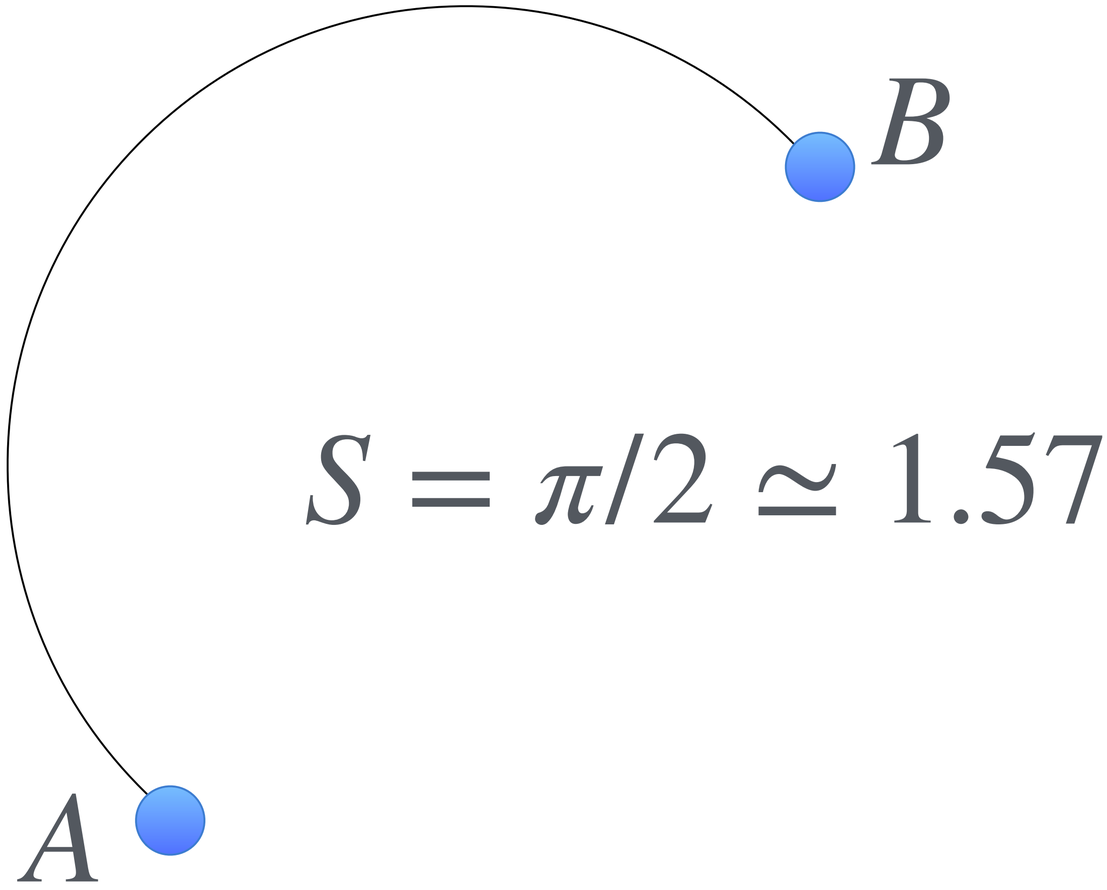


$$S = \pi/2 \simeq 1.57$$

# …time for the triangle

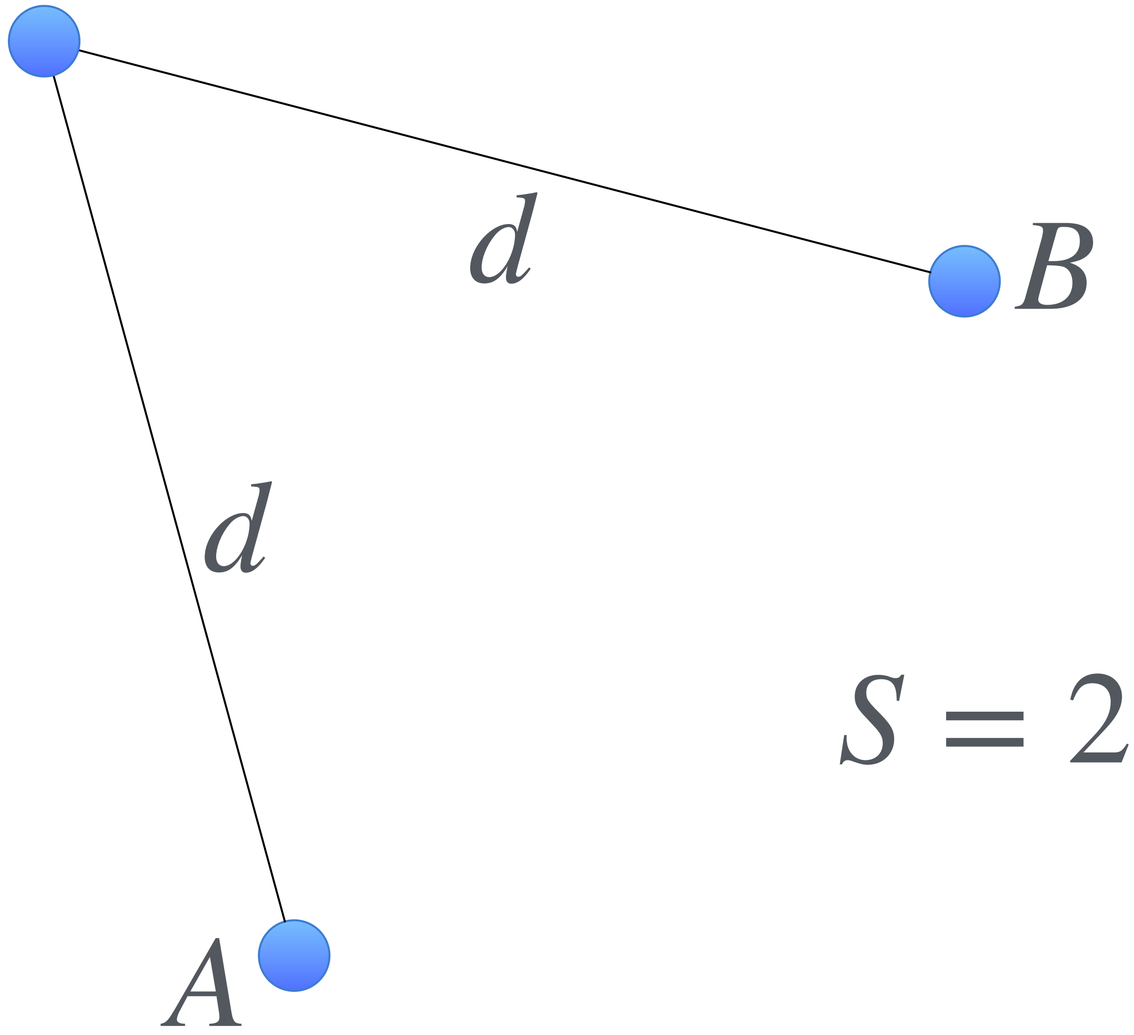


$$S = 2$$

# …and square

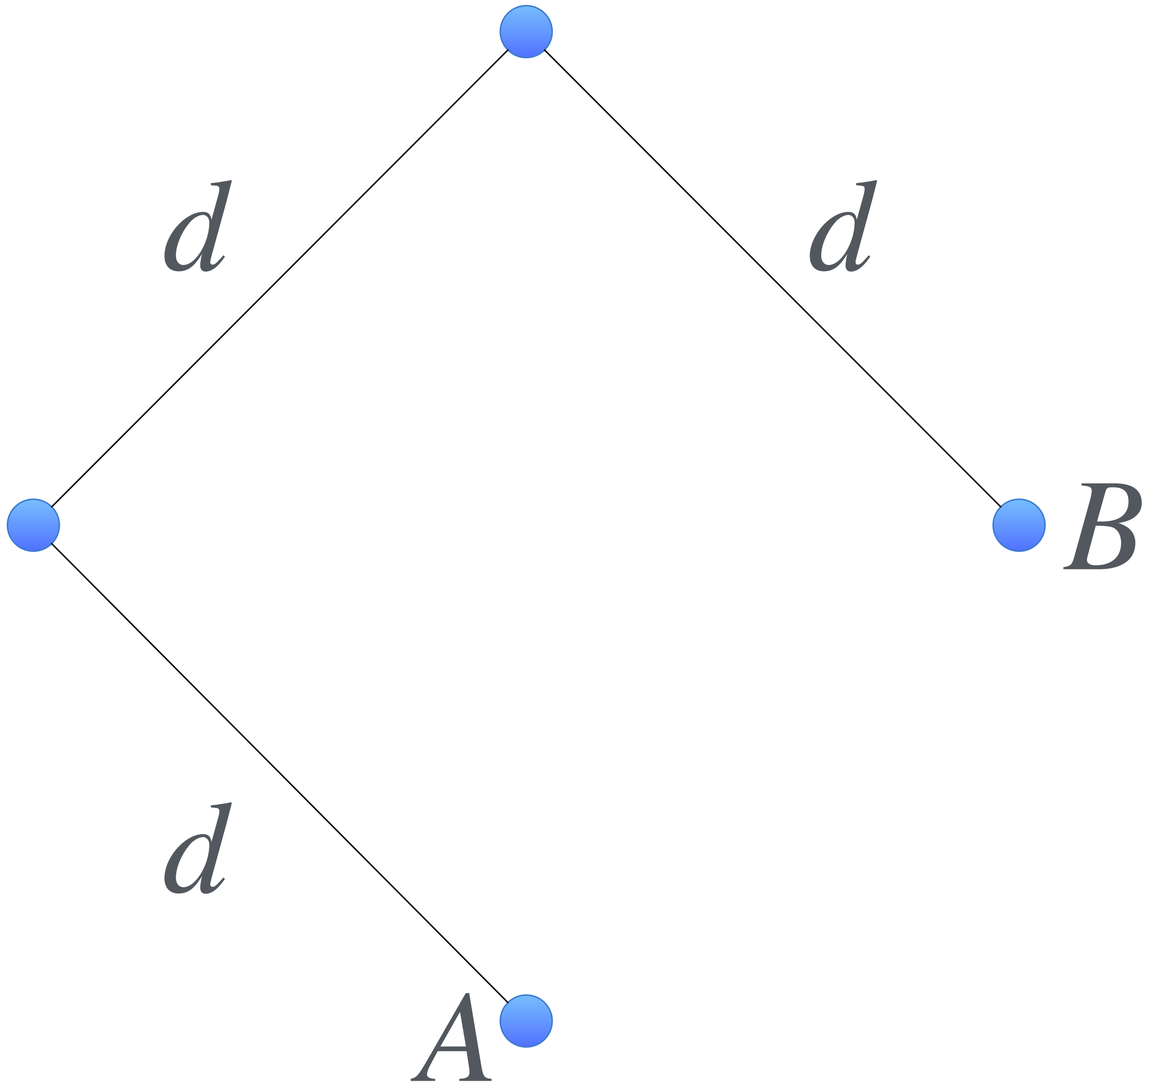


$$S = 3$$

# …bent one

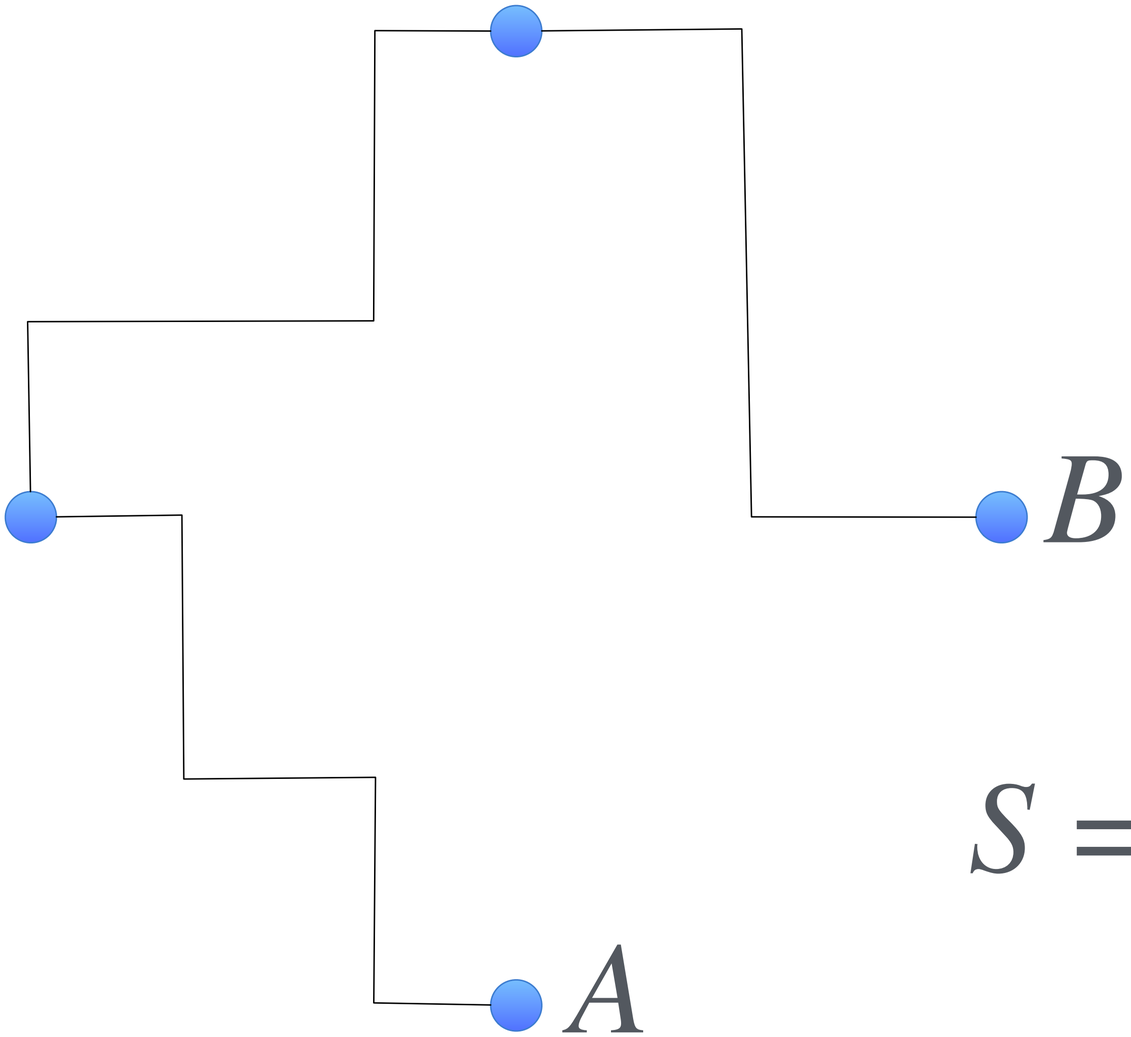


$$S = 3 \times \sqrt{2} \simeq 4.24$$

# …triangle again, the long one

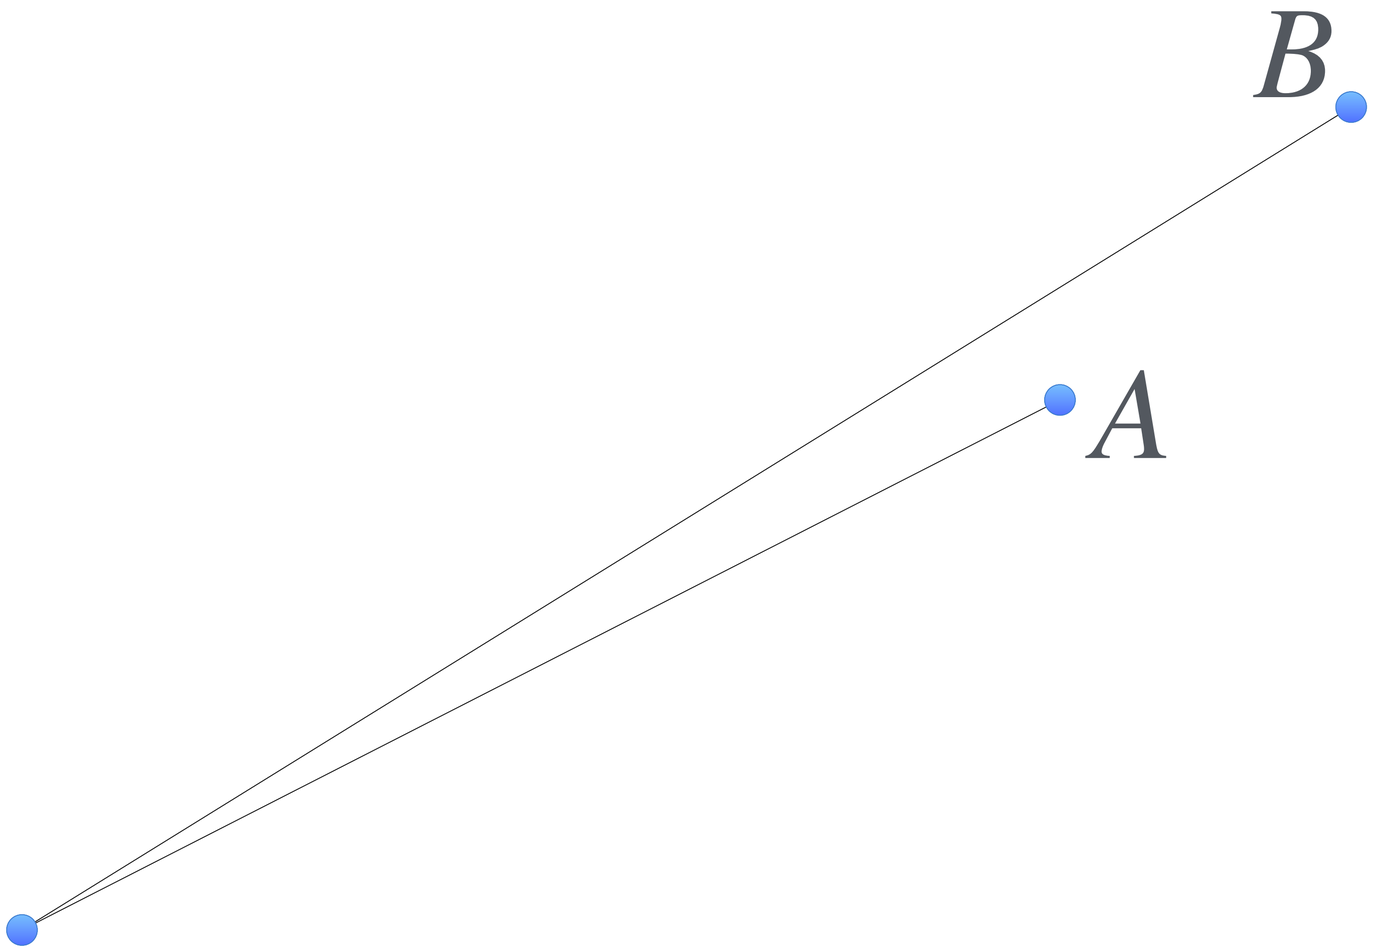


$$S \simeq 6.6$$

# …now let’s bend it

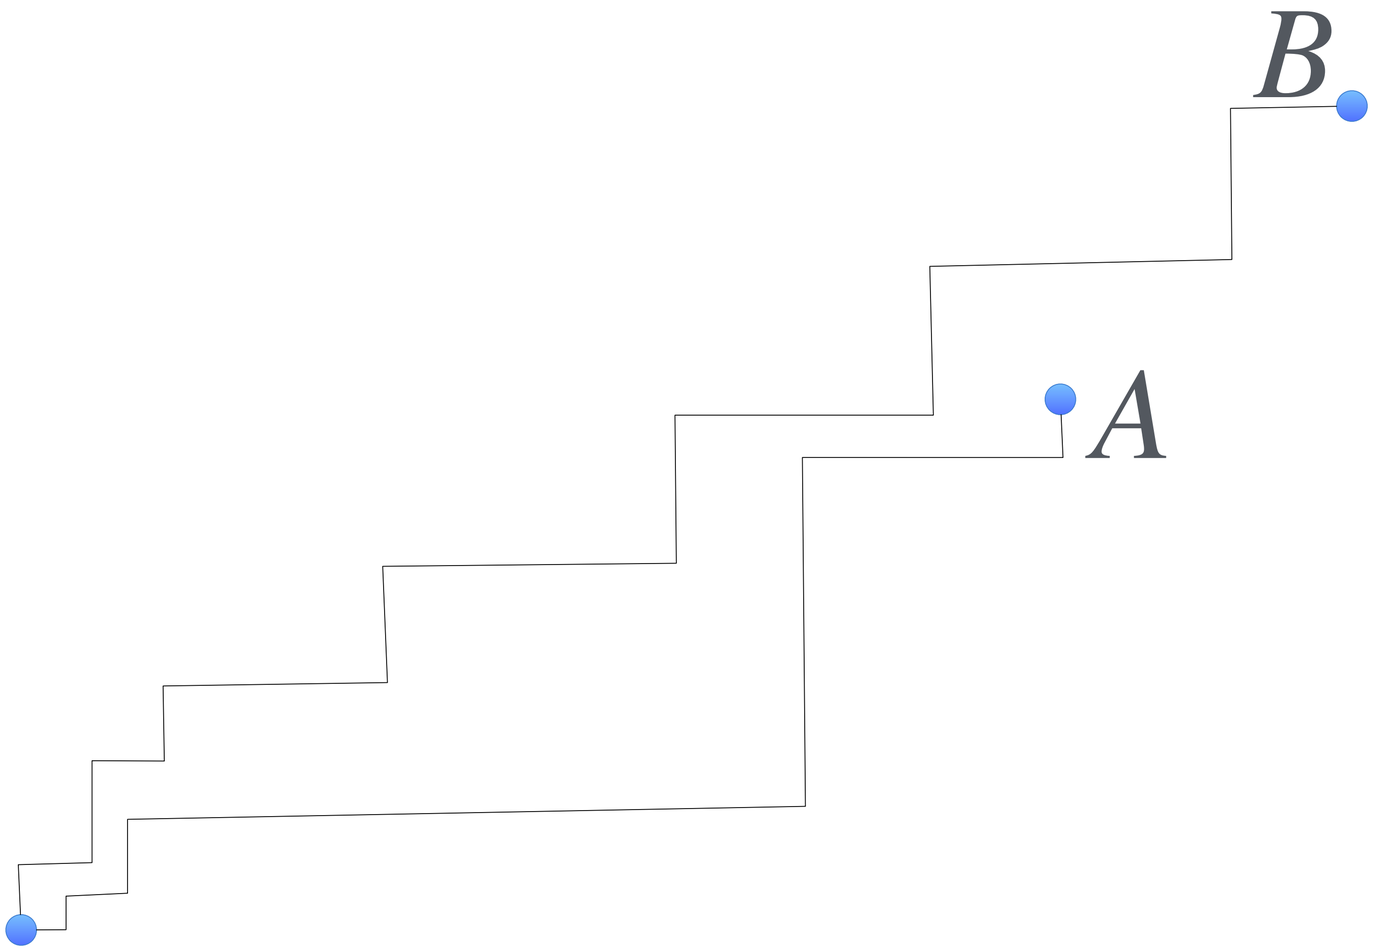


$$S \simeq 9.3$$

# Time for heatmaps! MinRTT

| | |
|---|---|
| S⩽2 | |
| 2<S⩽3 | |
| 3<S⩽5 | |
| S>5 | |

| Area/Area | Вінниця | Дніпро | Київ | Львів | Маріуполь | Миколаїв | Одеса | Полтава | Харьків | Херсон |
|---|---|---|---|---|---|---|---|---|---|---|
| Вінниця | X | | | | | | | | | |
| Дніпро | | X | | | | | | | | |
| Київ | | | X | | | | | | | |
| Львів | | | | X | | | | | | |
| Маріуполь | | | | | X | | | | | |
| Миколаїв | | | | | | X | | | | |
| Одеса | | | | | | | X | | | |
| Полтава | | | | | | | | X | | |
| Харьків | | | | | | | | | X | |
| Херсон | | | | | | | | | | X |

## Place

| |
|---|
| 6 |
| 5 |
| 1 |
| 2 |
| 10 |
| 9 |
| 4 |
| 7 |
| 3 |
| 8 |

# MinRTT: Observations

- The picture is not perfect in any sense, but still acceptable
- There is a certain correlation between “rightness” and the size of the city:

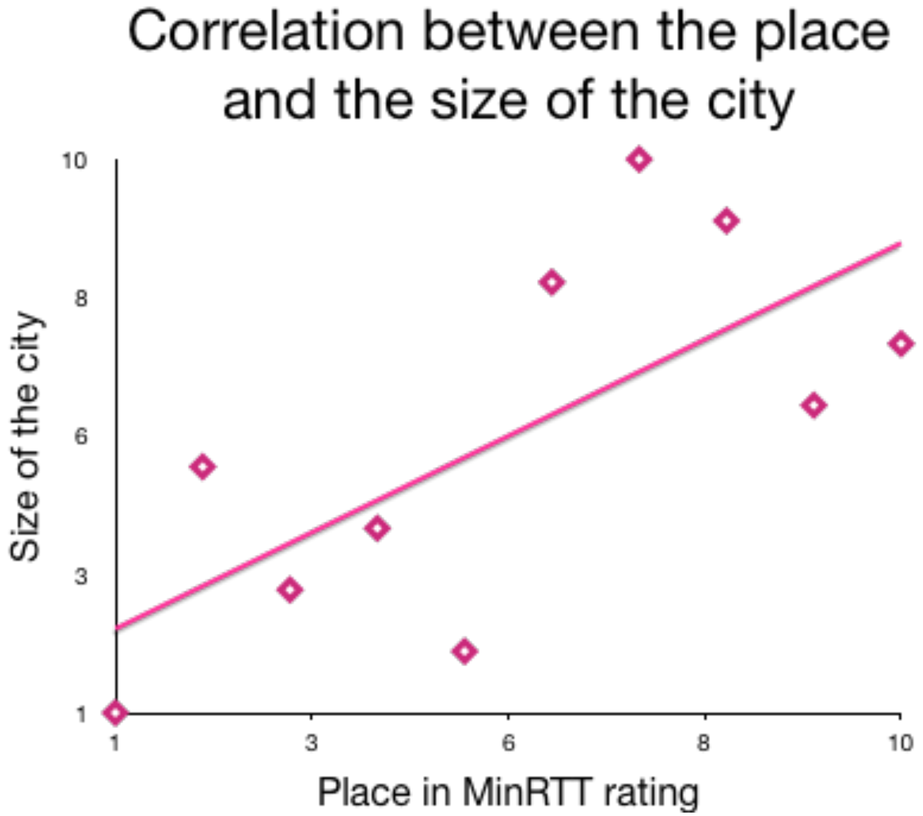

# Time for heatmaps! MedRTT

| | |
|---|---|
| S⩽2 | |
| 2<S⩽3 | |
| 3<S⩽5 | |
| S>5 | |

| Area/Area | Вінниця | Дніпро | Київ | Львів | Маріуполь | Миколаїв | Одеса | Полтава | Харьків | Херсон |
|---|---|---|---|---|---|---|---|---|---|---|
| Вінниця | X | | | | | | | | | |
| Дніпро | | X | | | | | | | | |
| Київ | | | X | | | | | | | |
| Львів | | | | X | | | | | | |
| Маріуполь | | | | | X | | | | | |
| Миколаїв | | | | | | X | | | | |
| Одеса | | | | | | | X | | | |
| Полтава | | | | | | | | X | | |
| Харьків | | | | | | | | | X | |
| Херсон | | | | | | | | | | X |

# MedRTT: Observations

- The picture is really bad
- There are a lot of “pathological” values medRTT that clearly need attention
- No any correlation: almost all values are horrible
- Further analysis makes no sense 😟

# Conclusions

- RIPE Atlas provides you the opportunity to explore RTT on a large scale
- Within cities, the situation is still acceptable
- Indicators of “straightness” of paths between cities on average are just bad

# Some speculations

- **Global underfunding of the industry leads to the approach to the fibreoptic cables like “at least, some exist”**
- **Mass extreme price optimisation leads to strange and suboptimal traffic paths**
  - There are no local resources for which access speed is critical
  - “Gamers must suffer”
  - “Traders must suffer bigger time”
- **Market consolidation is likely to significantly change this picture**

# What's Next

- **Conduct a more in-depth analysis of the results already obtained**
- **Make improvements in the code (more readability, even less manual work in the future)**
- **More automation on results verification**
- **Perform similar measurements for other countries and between them**
  - Next step: RTT measurements between European capitals
- **On verified data sets, make these measurements regular and the results public**
  - Need more RIPE Atlas probes in Ukraine!

# Acknowledgements


- To Emile Aben for the IXP Country Jedi tool, which inspired me to do this research
- To the RIPE Atlas team for the toolset that they developed
- To Viktor Naumov for advice and assistance during the work

# Questions