Ollivier's Ricci Curvature of real low-power wireless network testbed

Pedro Henrique Gomes,Chi Wang, Bhaskar Krishnamachari, Edmond Jonckheere

July 31, 2017

This technical note reports on our experimental investigation of the Olivier-Ricci curvature characteristics [1, 2] of a low-power wireless network testbed, the USC Tutornet testbed consisting of about 100 nodes, illustrated below [3].

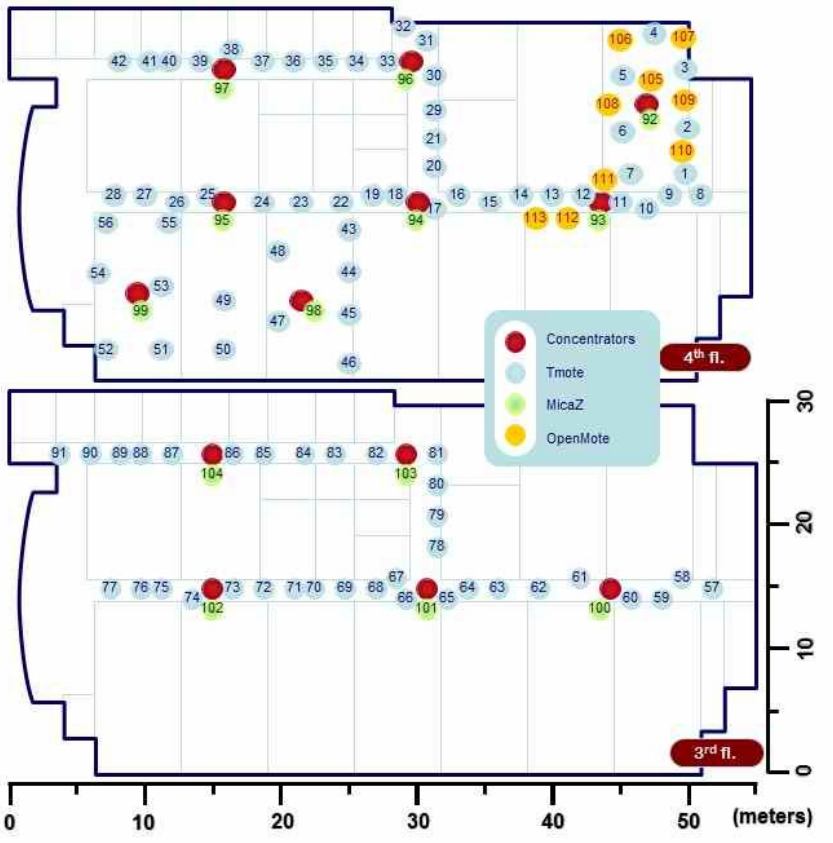

In the experiments we considered the packet delivery ratio between 40 nodes at 16 different channels. The traces are gathered every 15 minutes, with a total of 96 snapshots of the network, resulting in 16 NxN PDR matrices.

We found a high correlation between quality of links and quality of channels and the OR curvature.

Results

Below we present and discuss a few graphs obtained empirically, that show:
(i) Average of curvature and Average PDR for different channels
(ii) Curvature variance versus Variance of PDR

(iii) Average curvature versus Average size of schedule (supposing a scheduled TDMA network, the size size is inversely proportional to the throughput)

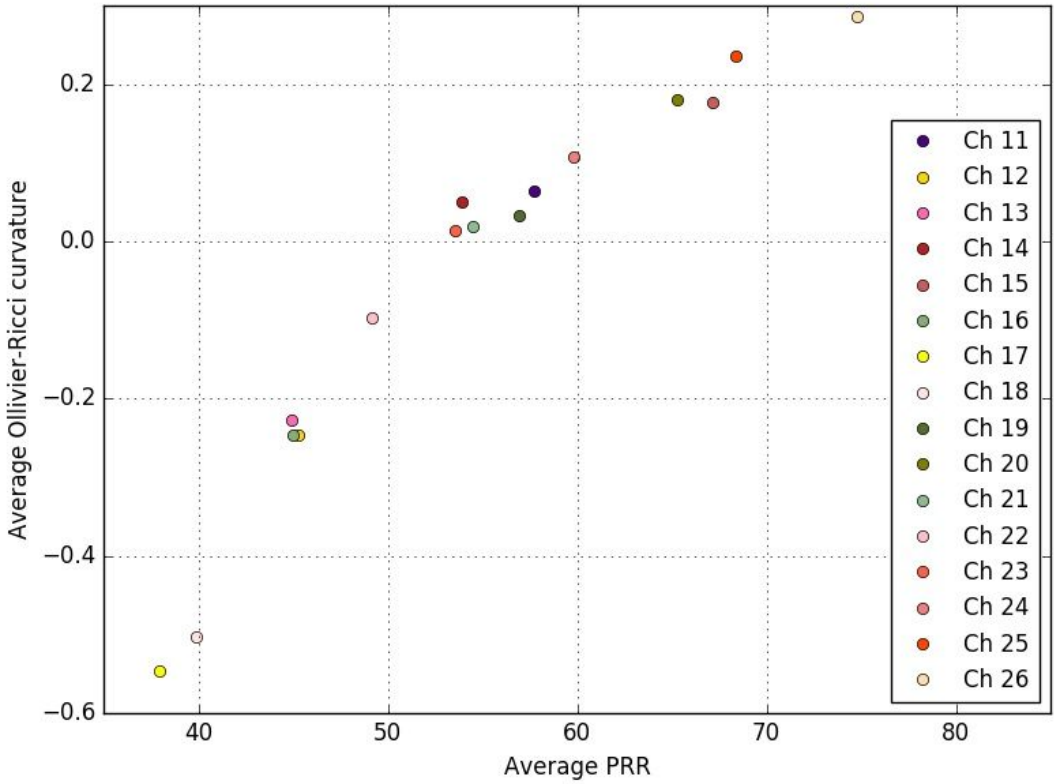

Figure 1: Average OR curvature vs. PRR

Figure 1 shows the scatter-plot of average OR curvature versus average packet reception rate (PRR) for all 16 channels of IEEE 802.15.4 measured on the testbed. Some channels of this radio, such as channels 25, 26 are relatively immune from interference from IEEE 802.11 (WiFi) networks active in the building where Tutornet is deployed. On these channels, the average PRR is quite high. The figure shows that channels with high PRR (more than 50% on average) tend to have a positive OR curvature, while those with low average PRR (less than 50%) tend to have a negative OR curvature.

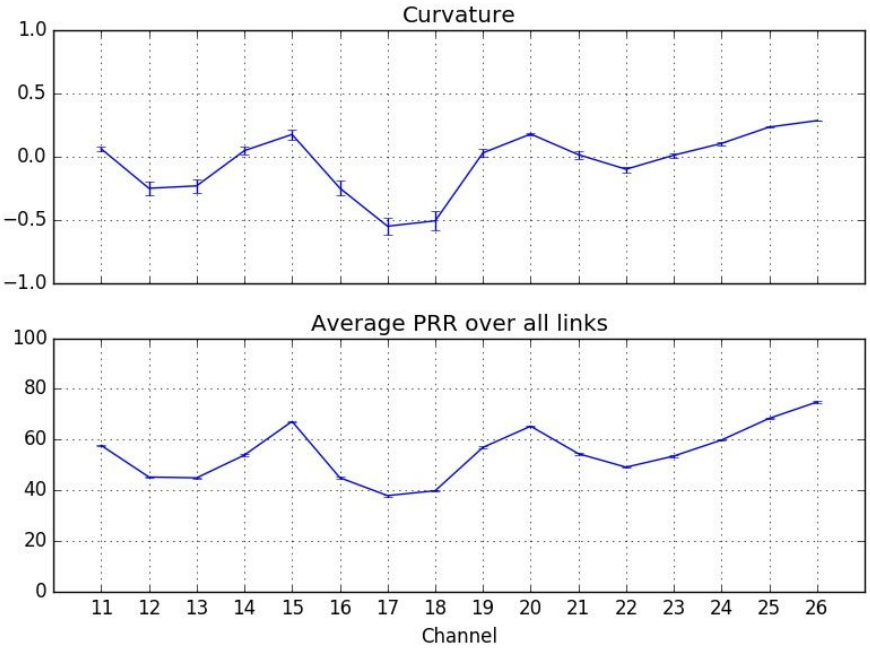

Figure 2: Average OR curvature and PDR across channels

Figure 2 shows more starkly the same finding. We can see that channels that have the most overlap with WiFi channels: 12-13, 16-18 show low PRR, while other channels such as 15, 25, 26 have a high PRR. The channels with high PRR are all associated with positive average OR curvature, while the channels with low PRR all exhibit negative OR curvature on average.

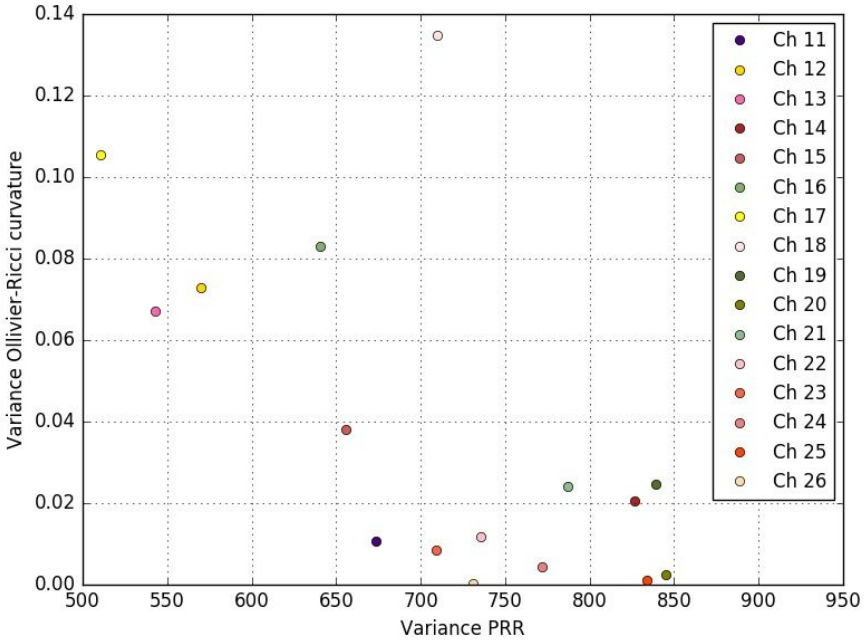

Figure 3: Variance of OR curvature versus variance of PRR

Figure 3 shows the variance of OR curvature (over the different time samples) plotted against the variance of PRR for the sixteen channels. It shows that channels with low variance in PRR tend to be associated with higher variance in OR curvature, while channels with higher variance in PRR tend to have very little variance in OR curvature.

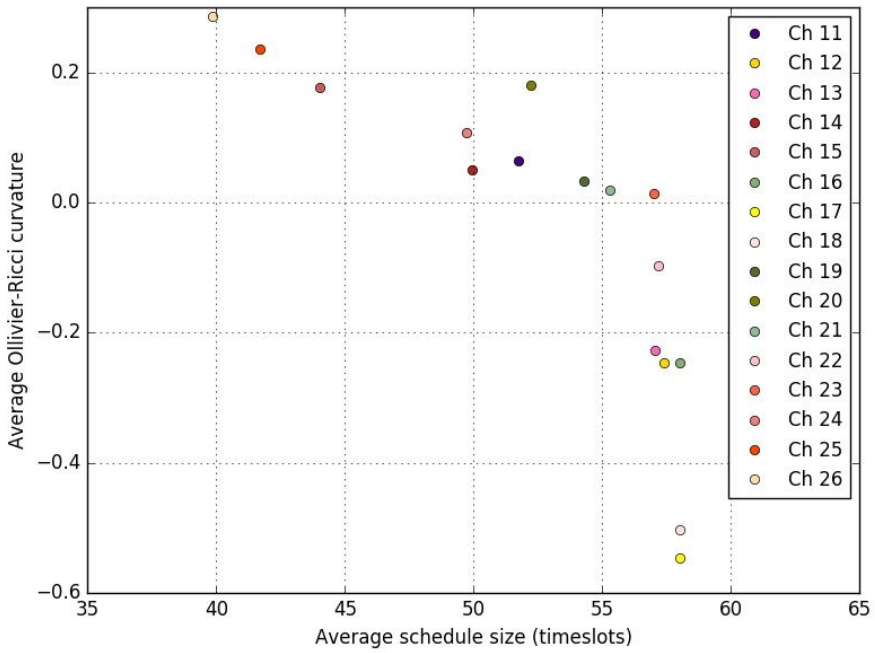

Figure 4: Average OR curvature versus multihop TDMA schedule size

We also ran simulations calculating the schedule size for a multi-hop time slotted scheme using a greedy algorithm for allocating non-interfering time slots to links in the network. For these simulations we considered as minimum PRR for a link to be 50%, which caused some instances at some channels to create disconnected trees. For those cases we considered a 'penalty' schedule of size 58. This value was obtained empirically as the maximum schedule size that was created if we run all experiments forcing the minimum PRR to be the maximum possible for each instance. We can see that for channels 17 and 18 most of the instances did not create a connected tree, as average PRR on these channels is below 50%. Conversely, we see that the good channels with positive average OR curvature show a small average schedule size. In such a Time Slotted network, the schedule size is effectively the inverse of the throughput, hence we are seeing that the channels with the highest OR curvature are able to provide the highest throughput.

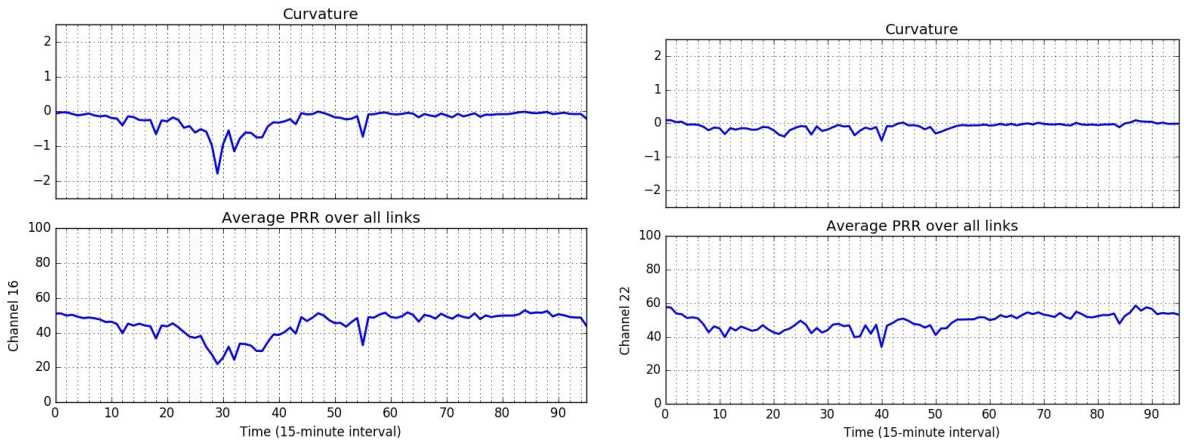

Figure 5 Average PRR over all links and Curvature as a function of time on channels 16, 22

Figure 5 shows for two particular channels (16, 22), how the average PRR over all links and the OR curvature behave over time. We find a high positive correlation with dips in PRR reflected in dips in the Curvature as well, showing that the positive correlation holds not only for their average values, but also over time.

Acknowledgement

This work was supported in part by NSF through grant number 1423624.

References

[1] J. Jost, S. Liu, “Ollivier's Ricci curvature, local clustering and curvature dimension inequalities on graphs,” Discrete Computational Geometry, 51 (2014), no. 2, 300-322.

[2] C. Wang, E.A. Jonckheere, R. Banirazi, “Wireless network capacity versus Ollivier-Ricci curvature under Heat Diffusion (HD) protocol,” ACC 2014.

[3] USC Tutornet Low Power Wireless IoT Testbed, URL: http://anrg.usc.edu/www/tutornet/