

Correlating Wireless Link Cost Metrics to Capacity

Marianna Carrera^{◇‡}, Henrik Lundgren[◇], Theodoros Salonidis[◇], Christophe Diot[◇]

{firstname.lastname@thomson.net}

[◇]Thomson, [‡]UPMC Paris Universit s, Univ Paris 06

Abstract—In wireless mesh routing, it is common to use periodic estimation of link quality to identify high throughput paths. While extensive work has been devoted to the design and static throughput evaluation of path metrics, evaluating the accuracy of the underlying link cost metrics under dynamic conditions has not received due attention. We introduce an experimental and analytical methodology that quantifies the ability of link cost metrics to accurately estimate link capacity. We use this methodology on a wireless mesh testbed to evaluate network layer and cross-layer estimation approaches of ETT, a popular wireless link cost metric. Our results show that the network layer approach exhibits low correlation with link capacity across time and across multiple links. On the other hand, cross layer information obtained by our methodology can significantly improve accuracy and can aid in identifying the dominant factors that lead to link cost metric inaccuracies.

I. INTRODUCTION

Mesh networks are 802.11-based multi-hop wireless backbones that provide low-cost Internet access in urban areas. Crucial to the performance of mesh networks is the routing protocol operation and significant research effort has been devoted to metrics that identify high quality paths. First, it was shown that link cost metrics that estimate link quality by periodic probing can compose path metrics that yield significant throughput gains over minimum hop count [1]. Further improvements have been achieved by link cost and path metrics that incorporate multi-channel diversity [2], interference [3], [4], or congestion [5]. Path metrics are typically evaluated in terms of their static properties by running separate back-to-back experiments under identical traffic patterns and comparing UDP or TCP throughput [6], [7].

Despite extensive work on path metrics, the ability of the underlying link cost metrics to accurately estimate wireless link quality under dynamic conditions has not yet been quantified. Link cost metrics are inherently inaccurate because, instead of a direct measure like throughput or delay, they only estimate an indirect measure of link quality using a limited number of probe packets. The ability of a link cost metric to estimate link quality depends both on its definition and on its measurement method. Quantifying this ability is crucial to the performance of the mesh routing protocol, especially under dynamic operation. If a link cost metric does not track the variation of link quality, the routing algorithm cannot react properly. On the other hand, if a link cost metric varies while the link quality does not, the routing algorithm will take unnecessary or erroneous actions. Previous studies of link cost metric dynamics [3], [8] have observed sensitivity to network traffic but do not connect metric variations to variations of a

direct measure of link quality.

We introduce a methodology that quantifies the ability of link cost metrics to accurately estimate a reference measure of wireless link quality under dynamic conditions. Our methodology consists of a reference link quality measure definition, an experimental component that performs measurements and an analysis component that processes these measurements. We use as reference link quality measure the link capacity, defined as the maximum UDP goodput of each link in absence of contention from other links in the network. The experimental component includes simultaneous and cross-layer measurements of link capacity and all link cost metrics under evaluation. Simultaneous measurements enable fair and precise comparison of the estimation accuracy of all link cost metrics with respect to link capacity under identical channel conditions. Cross-layer measurements enable deeper analysis of the behavior of 802.11 wireless links. The analysis component of our methodology uses linear regression to correlate the simultaneous link cost metric values and link capacity values measured by the experimental component. The squared correlation coefficient is used to quantify and summarize the ability of each link cost metric to track link capacity variations both across time and across several links in the mesh network.

Using a 16-node mesh network testbed, we apply our methodology to evaluate the link quality estimation accuracy of ETT [2], a link cost metric used by several path metrics for mesh networks. We evaluate a network layer approach and a cross layer approach for ETT estimation. The network layer approach estimates the ETT components using network layer broadcast probes and packet pairs as originally proposed in [2]. The cross-layer approach estimates the ETT components using some of the cross layer information collected by our methodology. We evaluate both approaches under a wide range of wireless conditions, including fixed bit-rate and data rate adaptation scenarios.

Our investigation reveals that the popular network layer approach exhibits poor performance in tracking link capacity variations over time and in ranking links with different capacities. Hence, a routing protocol using this approach will make erroneous or unnecessary routing decisions with respect to link capacity. On the other hand, the cross layer approach is more accurate in all scenarios.

Finally, we use the cross layer information obtained by our methodology to explain the origins of the network layer approach poor performance. Among various potential sources of inaccuracy, the packet pair estimation technique and ACK delivery ratio approximation by broadcast probes are domi-

nant. A surprising finding is that the packet pair technique does not estimate the link bit-rate as proposed in [2]. Rather, it causes ETT to react to stable or varying link capacity in a non-correlative manner.

Our methodology provides a deeper understanding of the behavior of link cost metrics over 802.11 wireless links and enables fair and precise comparison of multiple metrics with respect to a reference link quality measure. Our results suggest that design of future wireless link cost metrics should exploit cross layer information and unicast probes to accurately estimate link capacity.

The rest of the paper is organized as follows. Section II introduces our methodology. Section III presents the two estimation approaches of ETT and describes our testbed and measurement system implementation. Section IV presents the results that correlate metric variations to link capacity variations across time and across multiple links. Section V identifies the sources of inaccuracies of the network layer approach. Section VI surveys related work. Section VII concludes.

II. A NEW EVALUATION METHODOLOGY FOR WIRELESS LINK COST METRICS

This section describes our experimental and analysis methodology, which quantifies the ability of link cost metrics to accurately estimate a reference measure of wireless link quality. We present the main components of our methodology in the following.

A. Reference Link Quality Measure

Link cost metrics typically measure characteristics of probe packets (e.g., losses and delays) and combine them to *estimate* the link's quality. This estimation can be inaccurate: the measured characteristics of probes approximate only to some extent the characteristics of data packets, the combination of such characteristics only captures some of the factors which affect link quality. In order to quantify the accuracy of this estimation, we need to define a reference link quality measure. We define link quality as the link capacity: the maximum goodput achievable on the link. Unlike the capacity of wired links, the capacity of wireless links varies over time, due to a variable transmission bit-rate, an unreliable physical medium and a variable number of MAC-layer retransmissions. This definition of link quality can be directly measured as the goodput of a backlogged UDP flow in absence of interference due to contention with other links in the network. As we show in Section IV and V, even in absence of contention interference, link cost metrics can be inaccurate. Once the accuracy in absence of contention interference has been understood, other definitions of link quality may be used to study the accuracy of link cost metrics under contention.

According to our definition, each experiment consists in a backlogged UDP flow between two wireless neighbors. During each experiment, the UDP flow is the only traffic in the network, except for the probe packets used to estimate the link cost metrics of interest. The goodput of the UDP flow is measured as received bytes per time unit. We perform long

lasting experiments, on the order of hours, to study the time evolution of all the quantities of interest.

B. Simultaneous Measurements

All the quantities of interest (i.e., the reference measure of link quality and all the link cost metrics under evaluation) are measured simultaneously. This aspect of our experimental methodology yields a set of quantities which have experienced identical link conditions. Since they all estimate an identical link quality it is fair to compare them.

In order to perform simultaneous measurements, our system continuously sends different types of probes to estimate the link cost metrics under evaluation together with the traffic needed to measure the link quality (in this study, the UDP backlogged flow). We ensure that the probes are periodically sent also on the link loaded by the backlogged UDP flow. We denote the estimations of all the link metrics and the link capacity referring to the same time interval $(t - \tau, t]$, as $(metricA, metricB, \dots, capacity)_t$. Each of these sets is a sample of the relation between each link cost metric and the link capacity. Moreover, by running long lasting experiments we collect time-series of such sets of simultaneous measurements, and are thus able to compare the evolution of each quantity over time.

C. Quantifying the Accuracy of Link Cost Metrics Estimations

Our analysis methodology quantifies the accuracy of a link cost metric estimation by computing its correlation with the reference link quality measure, i.e., as defined in Section II-A, the link's capacity. We evaluate two complementary aspects: to what extent the variations of the link capacity are captured and reflected by each link metric, and, vice versa, to what extent the variations of a link metric are explained by variations in the link capacity. Our analysis methodology returns a coefficient, ranging between 0 and 1, which quantifies the accuracy of the estimation of link quality done by a link cost metric, according to both these aspects.

The analysis uses data set of simultaneous estimations, $(metricA, metricB, \dots, capacity)_t$, obtained as detailed in Section II-B. For simplicity, we describe the analysis considering a single link cost metric, referring to a set of $(metric, capacity)_t$ measurements. The same analysis applies to each link cost metric under evaluation. Recall that each $(metric, capacity)_t$ pair is a sample of the relation between the link metric and the link capacity. We expect this relation to be inversely proportional: the higher is the cost of a link, the lower is its capacity, and vice versa. We express this inversely proportional relation in a general form as:

$$metric = \frac{\alpha}{(capacity)^\beta}. \quad (1)$$

We consider the same relation in log-scale:

$$M = a - bC, \quad (2)$$

where a and b are functions of α, β in Eq. (1), and M and C are the logarithms of the link metric and the link capacity,

respectively. Then, we compute the *squared correlation coefficient* (R^2) of the samples of M and C . This squared correlation coefficient ranges between 0 and 1, and is defined as:

$$R^2 = \frac{[\text{cov}(M, C)]^2}{\text{var}(M) \cdot \text{var}(C)}, \quad (3)$$

where $\text{cov}(M, C)$, $\text{var}(M)$ and $\text{var}(C)$ denote covariance, and sample variance of M and C , respectively.

Since the relation of M and C is linear, R^2 measures the amount of variation of the link metric which is explained by the variation in the link capacity, and vice versa. Likewise, the complement of the squared correlation coefficient ($1 - R^2$) represents both the amount of metric variation that is not related to capacity variation, and the amount of capacity variation that is not captured by the metric. If a link cost metric varies while the capacity does not, it may induce the routing algorithm to avoid a good link or select a bad link. On the other hand, if a link cost metric does not capture the variation of the link capacity, the routing algorithm will not be able to react to improvements or degradations of a link's condition.

The interpretation of the results of this methodology depends on the characteristics of the input data set. For example, given a data set consisting of subsequent measurements from a single link, we evaluate the ability of a metric to track the capacity variation over time; while, given a data set consisting of measurements from several different links, we evaluate the consistency of the relation metric-capacity across links.

D. Collecting Cross Layer Information

The last aspect of our methodology consists in collecting cross layer information, throughout the whole duration of all the experiments. The performance of 802.11 links are affected by several mechanisms, such as the retransmissions of unicast DATA frames and the multirate support. Collecting cross layer information gives an important insight on the link's status and behavior, which can be used to explain the inaccuracy of the studied link cost metrics.

We modify the wireless card driver to record, for each unicast network layer probe, the exact number of MAC-layer retransmissions, the list of the bit-rates used for all the retransmissions, and the outcome of the final transmission, which determines if the network layer packet is successfully delivered. We collect this information periodically and continuously over time.

III. EXPERIMENTAL STUDY

In this section we show how our methodology can be used to evaluate the accuracy of one or more link cost metrics. We study the accuracy of two estimation approaches of ETT, a widely used link cost metric and a building block of many routing metrics. We detail how the two approaches measure the components of ETT; and we describe the testbed and system implementation which allows to apply our methodology to this experimental study.

A. Two ETT Estimation Approaches

Expected Transmission Time, ETT, defines the cost of a wireless link as the expected time the MAC-layer takes to successfully deliver a unicast packet. ETT is defined as $ETX \cdot \frac{S}{B}$, where ETX is the *Expected Transmission Count* [1], S is the packet size and B is the link's bit-rate. Both ETX and B may vary over time and need to be continuously evaluated. The technique used to measure these components of ETT may introduce errors and therefore affect the accuracy of the estimation of link quality done by ETT.

The Network Layer Approach. Draves et al. [2] propose to infer ETX and B by measurements of network layer quantities. We call this approach the *network layer approach*. To compute ETX , each node periodically sends broadcast packets, and counts the number of packets received from each neighbor. Each node computes the delivery ratio of its incoming links as the ratio of the number of received packets over the number of expected packets. To account for the bi-directionality of a unicast transmission (DATA + ACK) the nodes exchange the computed ratios in the payload of each probe packet. ETX is then computed as $\frac{1}{d_f \cdot d_r}$, where d_f is the delivery ratio of the forward direction of the links (approximating the DATA delivery ratio) and d_r is the delivery ratio of the reverse direction of the link (approximating the ACK delivery ratio). The link's transmission bit-rate (B) is computed using the packet pair technique: each node sends two unicast packets back-to-back on each of its outgoing links and records the packets' dispersion. The link's bit-rate is the ratio between the size of the second packet and the minimum dispersion among the last 10 samples.

The Cross Layer Approach. The cross layer information collected using our experimental methodology can be used to estimate the components of ETT. Recall that we collect from the driver the exact number of retransmissions along with all the bit-rates used to transmit network layer probes. Instead of approximating the expected number of retransmissions (ETX) by the delivery ratios of the broadcast probes, we average the number of retransmissions performed for the unicast probes. Similarly, instead of inferring the link's bit-rate (B) by the packet pair technique, we average the bit-rates used for all the retransmissions of unicast probes.

B. Testbed and System Implementation

We implement, on a 802.11g based testbed, the two approaches to estimate ETT and the mechanisms to collect and analyze the measurements according to our methodology.

The testbed consists of 16 nodes deployed on Thomson premises, and spans two buildings and an outdoor parking garage (Figure 1). Each node is a mini-ITX form factor PC (1.7GHz Intel CPU, 512MB RAM, 80GB HDD), equipped with an Atheros 802.11a/b/g mini-PCI card (AR5212) connected to an external 5 dBi antenna. Each node runs the Linux OS (kernel 2.6.18.5) and the Atheros card is driven by the MadWifi driver (version 0.9.4) with the SampleRate rate adaptation algorithm. SampleRate implements the Multi-Rate

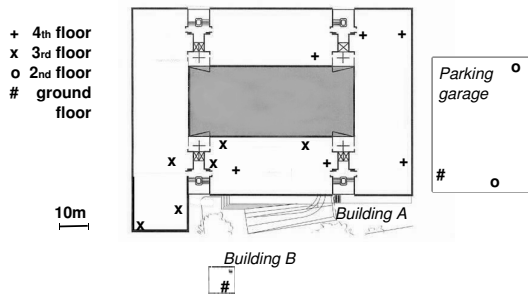


Fig. 1. Location of the nodes in our testbed.

Retries (MRR) option, i.e., retransmissions of the same MAC frame can use different bit-rates. RTS/CTS is turned off.

We implement the probing system using Click (version 1.5.0) [9]. We send three types of probes: (i) broadcast packets to estimate the number of retransmissions (ETX) according to the network layer approach; (ii) a pair of unicast packets to estimate the link transmission bit-rate (B) according to network layer approach; and (iii) unicast packets to collect the cross layer information as part of our methodology. In this experimental study, the cross layer information is also used to estimate ETX and B according to the cross layer approach.

The size of the first unicast packet of a packet pair is 137 bytes, the size of the second packet is 1137 bytes; each pair is sent back to back every minute, as in [2]. The other two types of probes are used to estimate the number of retransmissions, and their sizes may therefore impact the results. We select these probes to be of medium size to not favor neither small nor large packets. We set the broadcast packets to be 400 bytes and the unicast packets to be 523 bytes. We modify the driver to record the cross layer information only for packets 523 bytes long. These probes, type (i) and (iii), are sent every 10 seconds.

We instruct the driver to use a separate queue for probes. This eliminates the risk of probe packets being dropped by the kernel in case of queue overflow due to the backlogged UDP traffic. The probes do not have priority over the UDP traffic. We record the broadcast and packet pair probes received by each node using tcpdump. We generate the backlogged UDP flow using iperf¹, which sends 1470 Bytes UDP packets and records the number of received bytes each second. We load only one link at a time with the backlogged flow, and we perform experiments during night time and weekends to minimize any potential interference from external networks. We exclude from our analysis the set of simultaneous measurements where the capacity is lower than 100 Kbps, since links with such low capacity are not interesting for transferring data.

Although the evaluation of the link capacity and each link cost metric could be performed on-line, the current implementation evaluates them off-line, using the information gathered throughout each experiment. The two approaches to measure ETT may have different overhead. We estimate this overhead as the average number of MAC frames injected in the network by each approach. We then sample the traces of the cross layer information in order to compare measurements of ETT

obtained with the same overhead. We read these traces with a sliding window algorithm to mimic an on-line evaluation. Each metric is computed based on the probe packets sent or received in the time interval $(t - \tau, t]$, where τ is the *estimation period* and t is the time-stamp of the last received probe. Each time a new packet is received (i.e., read from a trace), the window slides. Likewise, for the capacity estimation we consider the bytes received during the time interval $(t - \tau, t]$ for the same *estimation period* τ . In this study, we select τ to be 10 minutes. We believe that this is a good trade-off between the number of probes considered to compute one metric value, and the age of such probes. Finally, we consider the pairs of $(metricA, metricB, capacity)_t$ from the same time interval $(t - \tau, t]$ and compute the squared correlation coefficient as described in Section II-C.

IV. RESULTS OF THE CORRELATION ANALYSIS

In this section, we study the accuracy of the ETT estimation with respect to two important properties of a link cost metric: i) the ability to track the link capacity variation over time, and ii) the consistency across several links of the inverse proportional relation of each metric against the link's capacity.

A. Metrics and Capacity Variations Over Time

We consider six links in our testbed and we perform, link by link, the simultaneous measurements of the capacity and the link cost metrics, using in turn all available bit-rates. We organize the data into sets, each consisting of measurements from a single link using one fixed bit-rate. In each dataset the measurements last up to two hours, during which the capacity may vary considerably or remain stable. In this part of our study, we are interested in the data sets where the capacity varies, in order to study the ability of the link cost metrics to track its variation.

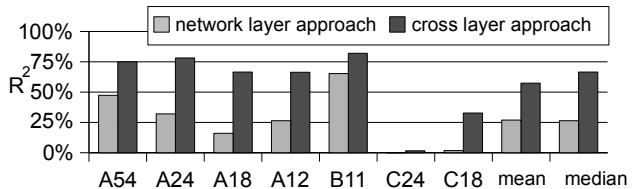


Fig. 2. Ability of each link cost metric to track the capacity variation over time.

Figure 2 compares the squared correlation coefficients (R^2) of ETT measured by the two approaches, for the seven data sets that demonstrate the highest variability in link capacity. The seven data sets are labeled by a letter denoting the link, followed by the bit-rate in Mbps. As we can observe in the graph, in all the datasets, ETT measured by the network layer approach is less accurate than ETT measured by the cross layer approach. Across these data sets, the average R^2 of ETT computed by the network layer approach is 27%, while the average R^2 for ETT computed using cross layer information is 57%. Moreover, in four datasets over seven, the network layer approach is less accurate or as accurate as its average accuracy, in fact, its median R^2 (26%) is close to its average (27%). On the contrary, the cross layer approach, in five datasets out of

¹<http://dast.nlanr.net/projects/Iperf/>

seven, is more accurate than its average accuracy, in fact, its median R^2 (67%) is higher than its average (57%).

The results in Figure 2 underline a severe inability of the network layer approach to track the link capacity variations over time: either the metric is varying with no correlation with the link capacity, or it does not vary while the capacity does. In Section V we investigate two reasons for the inaccuracies of the network layer approach in the dataset “A24” in Figure 2.

B. Metrics and Capacity Variation Across Links

We now study the consistency across links of the inversely proportional relation between a link cost metric against the link capacity, with auto rate disabled and enabled. We perform one hour long experiments, link by link, collecting the simultaneous measurements according to our methodology. In contrast with the study in the previous section, we now consider data sets consisting of measurements from several links. On each link the capacity may vary considerably or remain stable. By considering measurements taken on several links, we obtain a variety of samples of the relation between metric and capacity, spanning a wide range of link capacity values and metric values. Ideally, all the metric-capacity pairs should describe the same inversely proportional relation with the link capacity, regardless of the link where the measurements have been taken. The squared correlation coefficient of such data sets expresses the consistency of this relation across links. A metric with low consistency across links would be incapable of ranking different links according to their capacity.

Auto Rate Disabled. We first perform our experiment across multiple links at four fixed bit-rates: 1, 11, 36 and 54 Mbps. The data gathered at the two highest bit-rates corresponds to fewer links than the data gathered at the lowest bit-rates (several links in the testbed cannot carry any traffic at such high bit-rates).

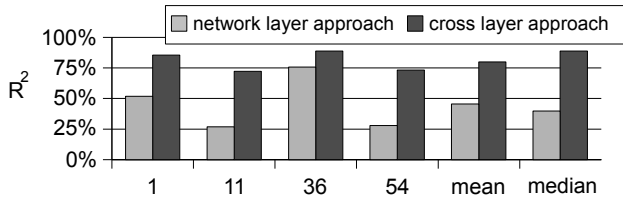


Fig. 3. Consistency across links of the relation between link metric and link capacity.

Figure 3 shows the squared correlation coefficients (R^2) of ETT computed according to the two approaches, for each used bit-rate (x-axis). We observe that the relation of the link capacity and ETT estimated with the network layer approach has little consistency, with an average R^2 equal to 45%. In contrast, the cross layer approach shows better consistency across links, with an average R^2 of 79%, and a maximum of 88%. We also note that the maximum R^2 of the network layer approach (76%) is comparable with the minimum R^2 of the cross layer approach (72%).

We now investigate in more detail the data set where the network layer approach obtains its best squared correlation coefficient (data set “36” in Figure 3). Figure 4 shows the

scatter plot of ETT versus capacity, and the fitting curve approximating the inversely proportional relation (least squared fitting), for the two approaches. We assign the same color to measurements points belonging to the same link².

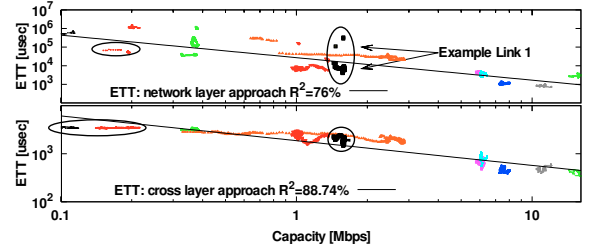


Fig. 4. Scatter plots and regression lines, in logarithmic scale, of ETT as a function of capacity, across multiple links using a fixed bit-rate equal to 36 Mbps.

We observe in Figure 4 that on several links, such as the one labeled “Example Link 1”, the capacity is fairly stable (little variation on the x-axis), while the network layer estimation of ETT (y-axis) spans a wide range of values above and below the fitting line. Such a spread of metric values on a stable link has two effects: first, the link appears to have highly varying quality while the capacity is stable; second, the spread of the ETT values are such that they are comparable with ETT values for other links with significantly different capacity. The consequence of the first effect is that a routing algorithm which uses these estimations of link quality may keep using a poor link, believing that the link is better than what it is in reality, or it may trigger unnecessary route changes, believing that the link is worse than reality. This may lead to network performance degradation. The consequence of the second effect is that these metric estimations cannot be used to correctly rank links according to their capacity. In Section V we investigate the underlying reasons for the observed metric spread, and show examples where this is due to uncorrelated variation of the bit-rate estimated by the packet pair technique.

Figure 4 shows another source of inaccuracy: in the circled regions on the left of the plots, the capacity varies considerably, while the link cost is constant. This is visible in both approaches. It is due to the definition of ETT itself: the very low capacity of these links is due not only to the very high transmission time, but also to losses at network layer. Network layer losses are not considered by the definition of ETT. This problem leads to the inability of ETT to rank poor links.

Auto Rate Enabled. We perform the same study with auto rate enabled. We fit the data set consisting of measurements from all the considered links. As above, the squared correlation coefficient measures the consistency of the relation existing between the link metric and the link capacity. A low squared correlation coefficient corresponds to a metric which is unable to rank different links according to their capacity.

Figure 5 shows the scatter plot of ETT versus capacity, the fitting curve and the corresponding squared correlation coefficients (R^2) for the two approaches. Points of the same

²Figures 4 and 5 read better in color, however, in case of b&w printing, we indicate with arrows and labels the most interesting phenomena.

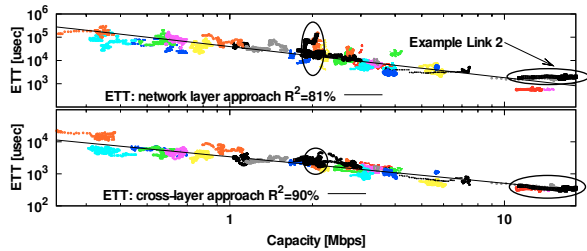


Fig. 5. Scatter plots and fitting lines, in logarithmic scale, of ETT as a function of capacity, across multiple links using auto rate.

color in Figure 5 corresponds to measurements from the same link. The network layer approach shows an accuracy of 81%, while the cross layer approach obtains 90%. The network layer approach shows better correlation with the link capacity compared with the auto rate disabled scenario, where its maximum R^2 is 76%. However, also in this scenario, there are links (in the center of the plot) where ETT varies considerably while the capacity is stable. For the same links, ETT estimated using the cross layer approach takes values very close to the fitting line. Moreover, the network layer approach shows a clear inability to represent very high capacity, as indicated by the label “Example Link 2”. In Section V we explain this inaccuracy by the inadequateness of the packet pair technique in the presence of auto rate.

V. INVESTIGATION OF METRIC INACCURACIES

In this section we present and explain the reasons for the inaccuracies noticed in the previous section. We analyze the time evolution of the components of ETT, estimated by each approach. We use the cross layer information to understand several inaccuracies of the network layer approach.

The delivery ratio of broadcast packets does not approximate the ACK packets delivery ratio. We consider the dataset A24 in Figure 2, where ETT measured by the network layer approach has a squared correlation coefficient of 32%, while the squared correlation coefficient of the cross layer approach is 78%. Figure 6 shows, from top to bottom, the timeline evolution of: the link capacity, the network layer measure of ETT (we zoom the y-axis to visualize the important features), the number of retransmissions (ETX) according to the two approaches, and the approximation of the ACK delivery ratio by the network layer approach.

One of the reasons for the inaccuracy of the network layer estimation of ETT is the variation of ETT labeled “A” in the second plot of Figure 6. In this region, just before 15:00, the network layer estimation of ETT increases to values comparable to the ones it takes between 14:00 and 14:15. However, the first plot of Figure 6 shows that in region A the capacity is steadily increasing to values higher than those that it takes between 14:00 and 14:15. Therefore, the variation of ETT in region A is uncorrelated with the capacity variation. This inaccuracy is due to the approximation of the ACK delivery ratio by the broadcast probes delivery ratio: the third plot of Figure 6 shows that this variation of ETT is due to the estimation of ETX , which varies, while the average number of retransmissions computed on the cross layer information

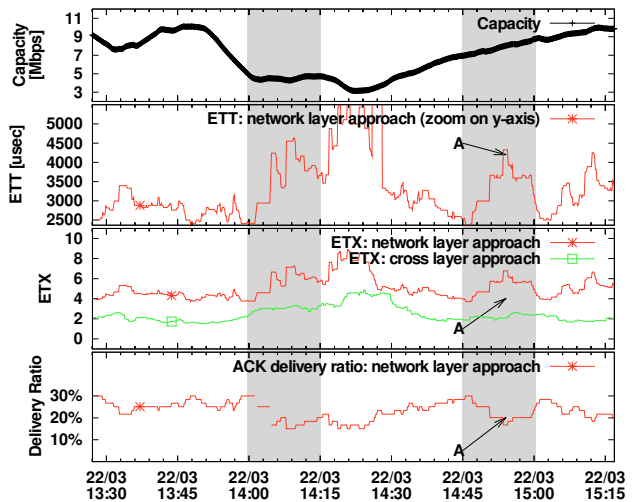


Fig. 6. Dataset A24 in Figure 2. Timeline evolutions of the link capacity, the network layer ETT and some of its components. The gray areas indicates the time intervals discussed in the text.

does not vary. In particular, as shown by the last plot, the variation is due to the approximation of the ACK delivery ratio, which decreases from roughly 30% to 15% and then increases again. We verified that the forward delivery ratio is instead fairly stable, around 90%. Our results confirm the inaccuracy of the estimation of the ACK delivery ratio, which was first pointed out in [1]. In addition, we show that the effects of this inaccuracy that can lead to very low correlation between the estimation of ETT and the link capacity.

The packet pair technique does not approximate the link’s bit-rate and reacts abruptly to capacity variations. Figure 7 considers again the dataset A24 of Figure 2. It shows, from top to bottom, the time evolution of the link capacity, of ETT, and of the estimation of the link’s bit-rate by the two approaches.

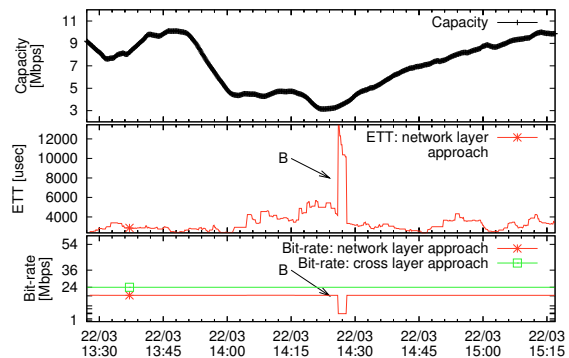


Fig. 7. Dataset A24 in Figure 2. Timeline evolutions of the link capacity, network layer ETT and the link bit-rate estimated by the two approaches.

The first and second plots in Figure 7 show another cause of the low correlation between the link capacity and ETT measured by the network layer approach: the spike labeled “B” does not correlate to a similar variation in the link capacity. The bottom plot in Figure 7 shows that the bit-rate measured by the packet-pair technique is responsible for this sharp variation of ETT. We observe that the bit-rate estimation varies considerably while it should not, because in this scenario the

bit-rate is fixed. The bit-rate estimation is sensitive to the link's capacity variation: it reacts to a capacity degradation with a certain delay and to capacity improvement with no delay, due to the selection of the minimum dispersion in the last estimation period. Recall that ETT estimates the bit-rate from the inverse of the minimum packet dispersion over the last 10 samples. When capacity decreases, the minimum dispersion increases. However, the bit-rate estimation needs 10 samples of high dispersion to settle to the new minimum. When the minimum is updated, ETT abruptly increases to a high cost. On the other hand, when the capacity increases the minimum dispersion decreases. The newly recorded dispersion is immediately used to recompute the bit-rate. This explains the sudden variations and results in ETT responding slowly to capacity decreases and rapidly to capacity increases.

The packet pair technique induces abrupt metric variations also on links with stable capacity. Figure 8 compares the timeline evolution of the capacity to ETT and its components, for the link labeled "Example Link 1" in Figure 4. This link exemplifies one type of inaccuracy of the network layer estimation of ETT.

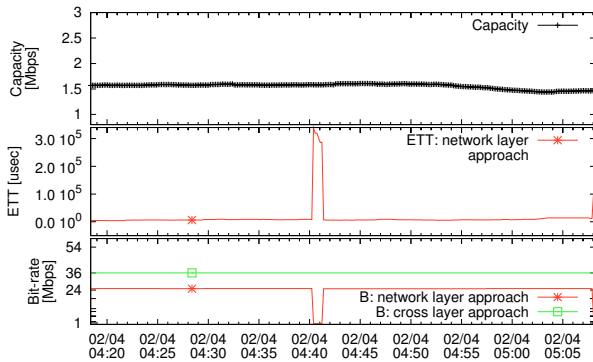


Fig. 8. Timelines of the ETT's components of the "Example Link 1" in Figure 4 (bit-rate fixed to 36 Mbps).

As we can observe in Figure 8, ETT presents a significant variation while the capacity is fairly stable. The variation of ETT spans several orders of magnitude, and it is concentrated in time, as a sharp peak around 4:40. The third plot shows that this variation is due to the estimation of the bit-rate by the packet pair technique. However, in contrast with the inaccuracy described above, here the capacity is not varying. This link has a low capacity, resulting from many MAC-layer retransmissions and some network layer losses. The cross layer information collected by our methodology registered an average number of retransmissions over the estimation period, ranging between 4.5 and 6.5 and a network layer delivery ratio between 80% and 95%. In such conditions, even the minimum dispersion of an estimation period includes multiple retransmissions of the second packet of the pair. Therefore, the computed bit-rate is very low (below 1 Mbps). This problem is not completely avoided by the use of the auto rate mechanism, which aims to maximize the throughput and avoids bit-rates which need a high number of retransmissions. In fact, in the circled region in the middle of the top plot of Figure 5 (auto

rate enabled), we can observe a similar vertical spread of ETT values estimated by the network layer approach, for constant capacities. We verified that this spread is due to the variation of the bit-rate estimation.

The packet pair technique cannot track the average bit-rate changes when auto rate is enabled. Figure 9 presents the timeline evolution of capacity, ETT and ETT's components for the link labeled "Example link 2" in Figure 5. The

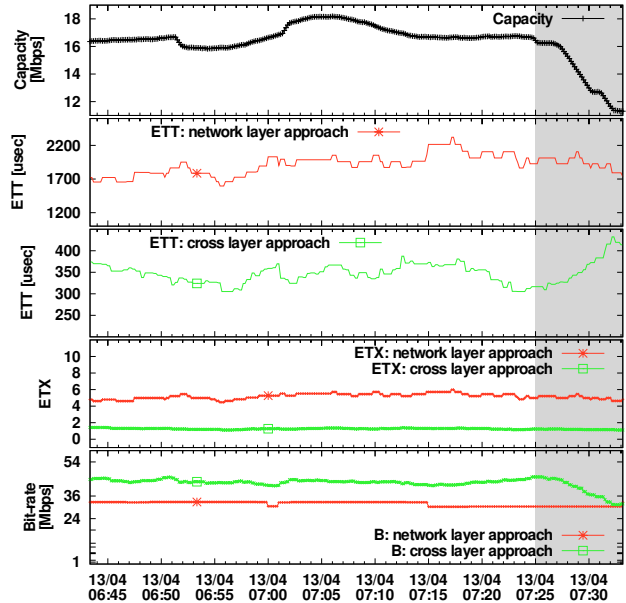


Fig. 9. Timelines of the ETT's components of the "Example link 2" in Figure 5 (auto rate enabled). The gray area indicates the time interval discussed in the text.

top plot shows how the capacity varies around 16 Mbps for roughly 40 minutes, and then, at 7:25, decreases till 11 Mbps. In contrast, we observe in the second plot that the network layer estimation of ETT varies without clearly tracking these capacity variation, in fact, in Figure 5 ETT appears as a stable horizontal line. On the contrary, the third plot shows that ETT estimated by the cross layer approach shows a trend inversely proportional to the capacity. The bottom plots of Figure 9 are the timelines of ETT's components, comparing in each plot the cross layer approach and the network layer approach. The number of retransmissions (ETX) remains stable according to both approaches. On the other hand, the used bit-rate is stable according to the packet pair technique, while it decreases according to the cross layer approach, which explains why the cross layer estimation of ETT tracks the capacity variation while the network layer estimation does not. On such high capacity links, where the retransmissions are few and a range of high bit-rates are used, the auto rate adaptation is responsible for the capacity variation. The cross layer approach considers all the bit-rates used for each transmission of the unicast probes and it computes the average bit-rate. On the contrary, the network layer approach selects the minimum dispersion and reports the highest bit-rate in each estimation period, which may be neither the most frequently used nor the average bit-rate.

VI. RELATED WORK

In this section we discuss related work which study the accuracy or the behavior of link cost metrics, focusing on the used methodology.

De Couto et al. [1], propose ETX as link cost and path metric. They study the accuracy of ETX (as a link metric) by comparing it to the number of retransmissions experienced by UDP packets. These two quantities are measured at different points in time (one minute apart), and for a very short duration (1 second and 10 seconds, respectively). Although the authors list some of the weaknesses of this estimation technique, they do not quantify the impact of these shortcomings. In contrast, we compare simultaneously estimated quantities, and show how different shortcomings manifest in the metrics' correlation to the link capacity.

Draves et al. [2] propose ETT, a link cost metric, and WCETT, a routing metric. They do not study the accuracy of ETT as link metric, but study the accuracy of the estimation of the link's bit-rate by the packet pair technique. However, their evaluation is limited to two nodes in close proximity, and this experimental setting suppresses the shortcomings of this technique since the number of retransmissions is likely to be small and stable. Our study covers a large set of links with different characteristics, and consequently provides a more realistic evaluation.

Das et al. [8] study three link cost metrics: ETX, ETT and the links bandwidth estimated by the packet pair technique. They investigate the link metrics' dynamic behavior and their sensitivity to traffic in the network. However, they do not quantify the accuracy of these link metrics and do not correlate to the variation of a direct measure of the link's quality. Subramanian et al. [3] propose iAWARE, a link and routing metric. The link version of iAWARE [3] is compared to ETT and IRU [4] in one controlled scenario to show the poor interference awareness of ETT and IRU.

Kim and Shin [10] propose a complex probing system (EAR), which relies on the standard MAC MIB interface to query the driver for MAC layer statistics. This work is complementary to our work and this system may be used to retrieve the cross layer information used by our methodology. They study the accuracy of the estimation of the frame delivery ratio by several approaches. However, the compared estimations are taken at different points in time. On the contrary, one fundamental aspect of our methodology is the comparison of simultaneous measurements.

Draves et al. [6] compare four path metrics, Hop Count, Per-hop Round Trip Time, Per-hop Packet Pair Delay, and ETX. Yang et al [4] propose the MIC path metric which relies on several link's characteristics, including ETT. Both these work present and compare path metrics, focusing on the performance and characteristics of the discovered routing paths. However, neither the behavior nor the accuracy of the underlying link cost metric has been evaluated.

VII. CONCLUSIONS

In this paper we proposed a novel combined experimental and analysis methodology for quantifying the correlation between wireless link metrics and link quality. Our methodology allows to quantify this correlation for two critical properties of wireless link metrics: (i) the metrics' capability of tracking link quality changes over time, and (ii) the metrics' capability of correctly ranking links according to their quality.

We demonstrated the feasibility and strength of our methodology by applying it to an experimental study of the widely used ETT link metric. In this study, we compared two different link quality estimation techniques of ETT (network layer vs. cross layer approach) and quantified their correlation to link capacity. Our results show that the widely adopted network layer approach lead to poor correlation: an average correlation of 27% in tracking link capacity variation over time; an average correlation of 45% and 81% in ranking different links according to their capacity (with auto-rate disabled and enabled, respectively). The corresponding results for the cross layer approach clearly show better correlation, with 57%, 79%, and 90% correlation, respectively. In addition, by using detailed cross layer information we explained why the use of broadcast and packet pair probes typically lead to poor metric-capacity correlation. Our results suggest that exploiting cross layer information lead to performance gains, since a more accurate link quality estimation translates into selection of more appropriate links.

Finally, we believe that our methodology takes an important step forward by not only studying link metrics' characteristics in isolation, but also taking a careful and systematic approach to relating the link metrics' behavior to a reference link quality measure.

ACKNOWLEDGMENT

The authors would like to thank Martin May for the helpful discussions during the writing of the paper.

REFERENCES

- [1] D. D. Couto, D. Aguayo, J. Bicket, and R. Morris, "A high-throughput path metric for multi-hop wireless routing," in *Proc. MobiCom*, 2003.
- [2] R. Draves, J. Padhye, and B. Zill, "Routing in multi-radio, multi-hop wireless mesh networks," in *Proc. MobiCom*, 2004.
- [3] A. Subramanian, M. Buddhikot, and S. Miller, "Interference Aware Routing in Multi-radio Wireless Mesh Networks," in *Proc. WiMesh*, 2006.
- [4] Y. Yang, J. Wang, and R. Kravets, "Interference-aware load balancing for multihop wireless networks," University of Illinois at Urbana-Champaign, Tech. Rep., 2005.
- [5] T. Salonidis, M. Garetto, A. Saha, and E. Knightly, "Identifying High-throughput Paths in 802.11 Mesh Networks: a Model-based Approach," in *Proc. ICNP*, 2007.
- [6] R. Draves, J. Padhye, and B. Zill, "Comparison of Routing Metrics for Static Multi-hop Wireless Networks," in *Proc. SIGCOMM*, 2004.
- [7] Y. Yang, J. Wang, and R. Kravets, "Designing Routing Metrics for Mesh Networks," *Proc. WiMesh*, 2005.
- [8] S. Das, H. Pucha, K. Papagianakki, and Y.-C. Hu, "Studying wireless routing link metric dynamics," in *Proc. IMC*, 2007.
- [9] E. Kohler, R. Morris, B. Chen, J. Jannotti, and M. F. Kaashoek, "The click modular router," *Transactions on Computer Systems*, 2000.
- [10] K.-H. Kim and K. Shin, "On accurate measurement of link quality in multi-hop wireless mesh networks," in *Proc. MobiCom*, 2006.