UCL Department of Computer Science
CS M038/GZ06: Mobile and Cloud Computing
2010–2011, Term 2
Kyle Jamieson and Brad Karp

One-pager: Roofnet (Bicket et al., 2005)   **Due: Start of lecture, 14th January 2011**

Instructions: **in your own words,** *answer the following question as* **succinctly** *as possible (in 200–500 words, but shorter answers within this range are encouraged). Quoting figures or text from the assigned reading or from any other source is specifically prohibited.*

In Figure 10 in the paper, the authors compare the throughput achieved on multi-hop paths in Roofnet (*y*-axis) with that predicted by the ETT metric (*x*-axis). Clearly, the achieved throughput often falls short of the throughput predicted by the ETT metric.

Why does the ETT metric tend to overpredict the throughput of a path? Refer to the details of the definition of the ETT metric in your answer, and any relevant assumptions made by this definition that may not hold in deployment.