Tranzeo's EnRoute500 Performance Analysis and Prediction

Introduction

Tranzeo has developed the EnRoute500 product family to provide an optimum balance between price and performance for wireless broadband mesh networks for public and private use. This white paper describes the aggregate network throughput and latency that an EnRoute500-based mesh network provides to WiFi clients that connect to it. Tranzeo offers a multi-layer solution to maximize the usable bandwidth available to the mesh in a given environment. Reductions in available bandwidth, due to noise or interference at the PHY layer, inefficiencies of the MAC layer such as perceived channel usage, or perceived congestion at the transport layer, all contribute to the degradation in performance and ability of a mesh network to scale. The EnRoute500 incorporates a tiered approach to ensure optimal use of available bandwidth resources.

The dual-radio EnRoute500 serves as a WiFi access point with a dedicated 802.11b/g radio, an intra-network repeater and router with a dedicated 802.11a mesh enabled radio, and as a gateway via Ethernet to an Internet point of presence (POP) as depicted in Figure 1. Tranzeo's patented technology separating the client access (in the crowded 2.4GHz ISM band) from the mesh backhaul channel allows explicit control of the backhaul and enables an extremely reliable high throughput solution. The EnRoute500 product family provides the lowest total cost of ownership (TCO) by leveraging omni-directional radio links providing a highly scalable network deployment, a self-forming network topology that reduces installation costs by simplifying the installation procedures, and a self-healing network ensuring the best possible and most reliable connectivity. An advanced quality of service (QoS) mechanism is employed to provide priority to VoIP and other high priority traffic, and fairly distribute network bandwidth to connected clients.

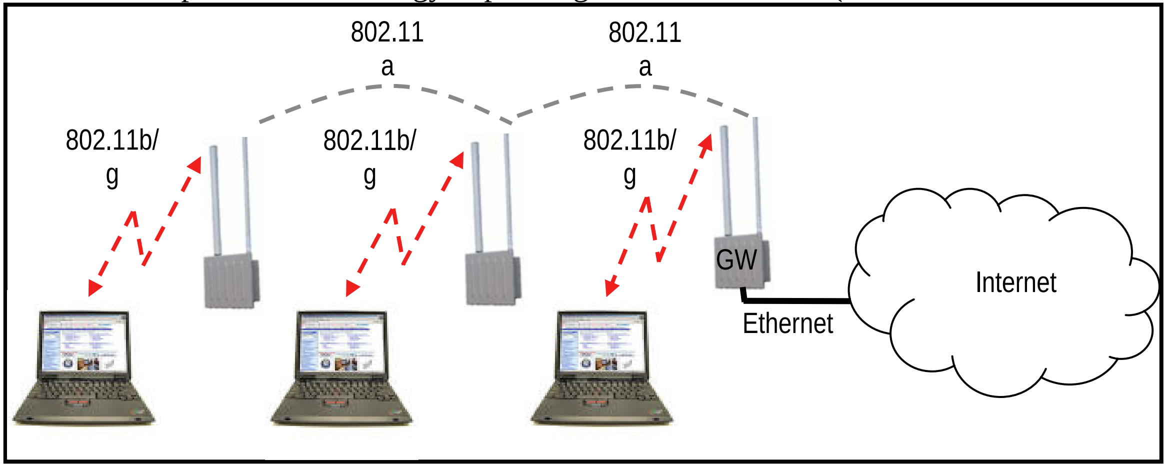

Figure 1. Three EnRoute500s serving as WiFi access points, mesh backhaul repeaters, and in the case of the EnRoute labeled GW to the Internet POP.

A Typical Mesh Network Topology

Figure 2 depicts a 5 X 5 array of EnRoute500 nodes deployed in an urban environment. The mesh network consists of node-to-node routes that are autonomously managed by the nodes' distributed routing algorithm using neighbor link quality. The routing algorithm also determines the best possible route from each node back to the POP gateway (node 33 in Figure 2).

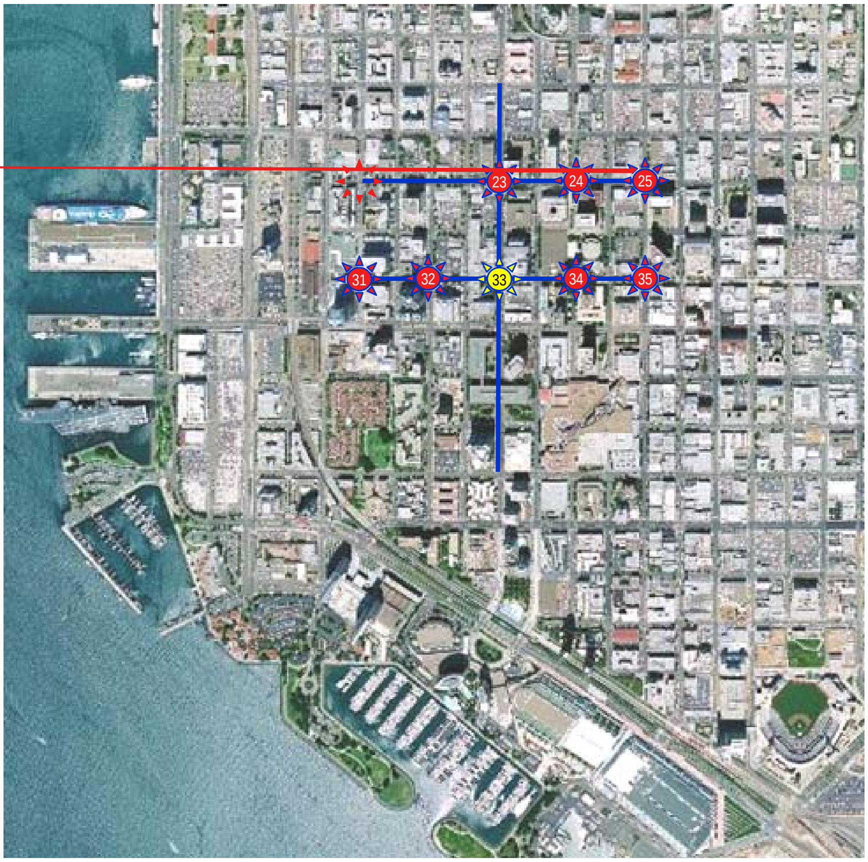

Every deployment is likely to have a different network topology based on the spacing of the nodes, the number and location of the Internet POPs, and the topography created by buildings, trees, and hills. The topology shown in Figure 2 is representative of a downtown urban environment and will be used to provide general performance guidelines for most deployment scenarios. Since each link in a deployment is a function of its environment, it would be impossible to test and evaluate every topology. However, we can examine basic topologies to create a model that can be used to predict network performance in a wide range of deployment configurations. In particular, there are two network topologies that can be evaluated that provide guidance, through extrapolation, on the overall performance of specific network deployments. The first useful network topology is a string of nodes with a single client communicating across the string to the POP gateway. The second network topology is a star topology where multiple clients communicate with the POP gateway one link away. The combination of these two topologies can provide a realistic prediction of performance across any network topology.

The distance between nodes and the terrain will have an effect on radio link quality, radio interference between nodes, and transmission overlaps causing collisions, for both the 802.11b/g (2.4GHz) access point links and the 802.11a (5.8GHz) mesh links. All of these factors will reduce available bandwidth in the network. These problems are examined in more detail in a separate white paper titled Describing Multi-Hop 802.11 Mesh Throughput [1]. For this analysis, both network topologies were tested in configurations where at a minimum, all neighbors one and two hops away could communicate or interfere with each other. Other 2.4GHz and 5.8GHz networks were observed to operate in the vicinity of the test deployment, which also contributed to the overall interference seen by the EnRoute500 network. The radio link quality, radio interference by in-network devices and extraneous radios, and signal degradation from obstacles, such as trees and vehicles can be highly variable between deployments, causing different degrees of connectivity degradation. There are also positive contributors to the wireless network such as buildings creating multiple reflections of the signal (multipath) which can be leveraged by the OFDM algorithm of 802.11g and 802.11a to improve the signal quality. Tranzeo's testing was done in several environments to take into account these negative and positive contributions to the wireless network. The following results represent a typical scenario to provide guidelines for network deployment.

Multi-hop Performance: Client Access Through A String Of Nodes

The first network topology considered in measuring multi-hop performance is a linear string of six nodes forming a five-hop mesh network to the POP Gateway 33 as illustrated in Figure 3.

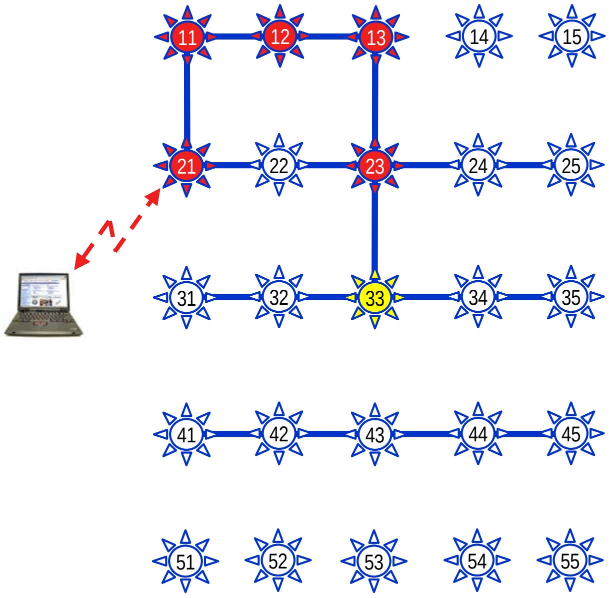

Figure 3. Client access through a string of nodes.

Multi-hop Bandwidth Test

Node 21 connects to the POP Gateway (node 33) using nodes 11, 12, 13, and 23 as repeaters. An 802.11g client connects to node 21 and the TCP/IP throughput for a flow from the client over the five hops to the Gateway (node 33) is measured. The 802.11g client is then moved to connect through node 11 to measure the throughput over the four hops to the Gateway. The 802.11g client is subsequently connected to each node in the string of repeaters to the Gateway to measure the effective throughput over every hop in the string. Figure 4 shows that multi-megabit throughput is achieved with an EnRoute500 network, even if as many as 5 hops are required to relay the data from the mesh node to a POP. Results from similar tests using Nortel and Tropos mesh solutions were reported in a presentation by the Network Research Center of Tsinghua University [2].

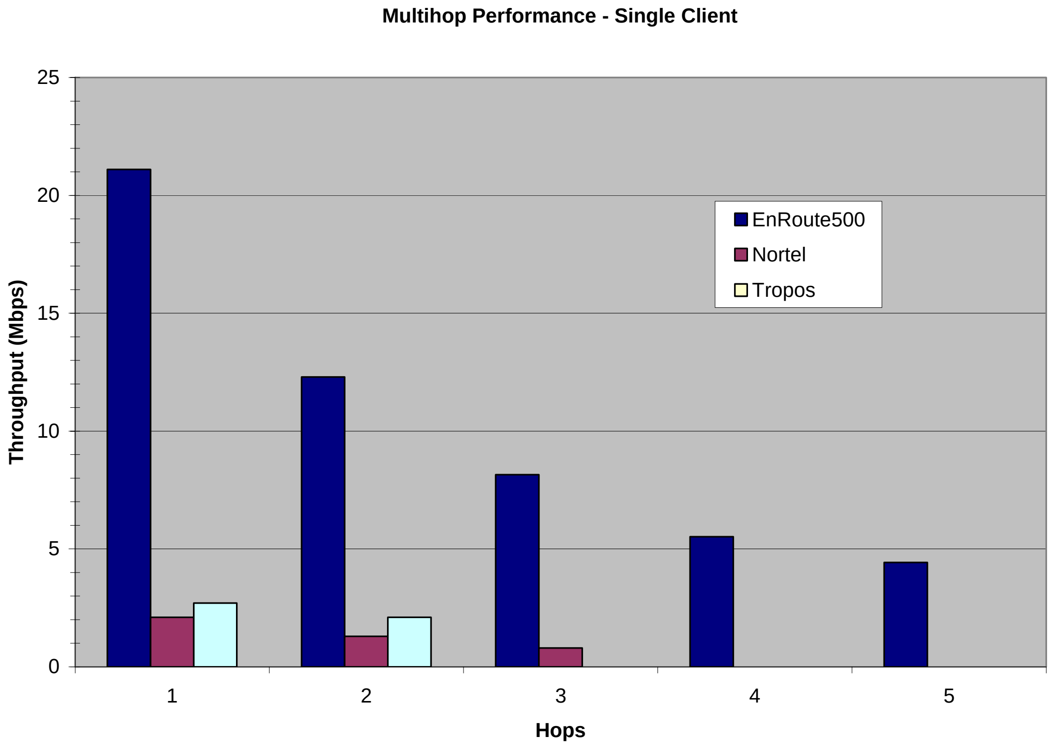

Figure 4. Multi-hop performance for a single client scenario.

The EnRoute500 delivers approximately ten times more throughput than these other solutions over the network and can support more hops in the network. The superior performance of the EnRoute500 clearly supports more bandwidth per access point compared to competitive solutions for the same network topology which translates into enabling more subscribers on the network. In addition, the ability to support more hops reduces the number of POPs needed, which reduces the cost of the overall network and provides greater flexibility as to where POPs must be located.

Multi-hop Latency Test

Latency was measured from remote nodes to the Gateway using the same network topology that was used in the bandwidth tests. Typical round trip latency when the network was lightly loaded averaged 5ms over five hops, or 1ms per hop. When the network was heavily loaded, the round trip network latency average increased to 35ms over five hops. All round trip latency measurements made when the network was heavily loaded exhibited less than 15ms of delay per hop. This low latency allows an EnRoute500 network to effectively carry VoIP traffic, which typically requires latency less than 150ms [3].

Quality of Service

While the single client use case is commonly referred to, it does not adequately reflect a real use case since the requirement of the network is not to serve a single client, but to support multiple concurrent users. In a multi-hop environment, clients connected closest to the Gateway will typically consume more bandwidth and starve clients further away from the Gateway. The EnRoute500 incorporates an advanced Quality of Service (QoS) framework that controls and distributes the available bandwidth to all nodes in the network. Tranzeo leverages the 802.11e QoS standard (WMM) and expands the concept to all data sources within a mesh network to provide the following capabilities:

- Multiple priority levels for VoIP, video, high priority data and low priority data
- Bandwidth control per subscriber in support of Service Level Agreements (SLAs)
- Peak balancing to prevent peak traffic from saturating the RF channel and overwhelming the 802.11 CSMA/CA MAC
- Load balancing to manage the use of multiple flows to ensure a fair distribution of the available bandwidth across the entire mesh cloud

To provide a more realistic use case, Figure 5 shows measured test results of the same string of six nodes supporting five hops, now with a client laptop associated with each mesh node.

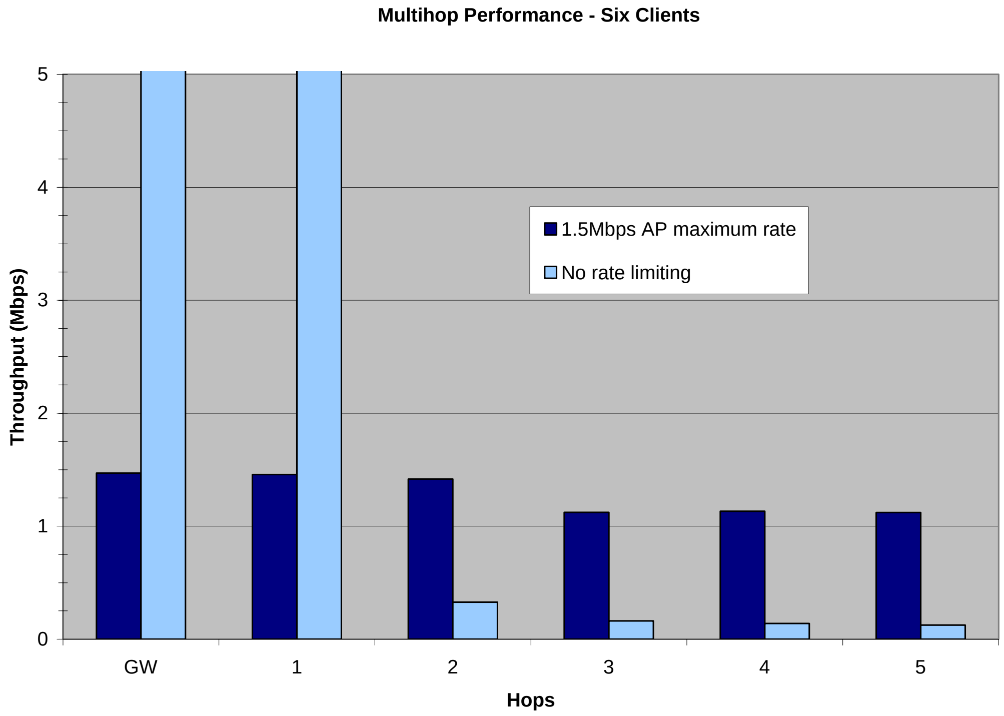

Figure 5. The benefit of QoS for multi-hop performance with multiple clients.

Tranzeo's QoS mechanism has limited each access point to 1.5Mbps of sustained throughput so as not to saturate the RF channel and eliminate unwanted bandwidth reduction by the 802.11 collision avoidance mechanism. In this test each client initiated a TCP flow for a period of one minute. The aggregate throughput available to these six clients was 7.7Mbps which is greater than the 4.8Mbps the single client furthest from the gateway achieved in the previous test shown in Figure 4. When Tranzeo's QoS is not rate limiting the access points, the clients nearest the Gateway benefit from higher bandwidth while virtually starving the clients further from the Gateway. This unfair distribution of bandwidth would be present for any WiFi mesh solution without some form of QoS [4]. Tranzeo's QoS guarantees a fair level of performance to every node in the network regardless of its location relative to the Gateway.

Single-Hop: Multiple Clients

The second network topology that is used to model network performance is a star topology where multiple clients connect to nodes one hop away from the Gateway 33 as shown in Figure 6.

The POP Gateway 33 has wireless mesh connections with nodes 23, 32, 34, and 43. 802.11g clients connect to each of the four nodes and the Gateway, and the maximum one-way TCP/IP throughput is measured for all five clients simultaneously accessing the network. Figure 7 shows the measured results of this test. Each of the four clients attached to node 23, 32, 34, and 43 enjoy 5-6Mbps of throughput with aggregate throughput of 21.5Mbps. Clients connected to the Gateway do not communicate over the mesh backhaul and therefore are only limited by the POP backhaul and the capacity of the Gateway. In real deployments, the POP backhaul has finite throughput often provided by a WiMax wireless link, T1, or DSL connection. The QoS mechanism should be employed to ensure clients connected to the Gateway do not absorb an unfair proportion of the available POP backhaul.

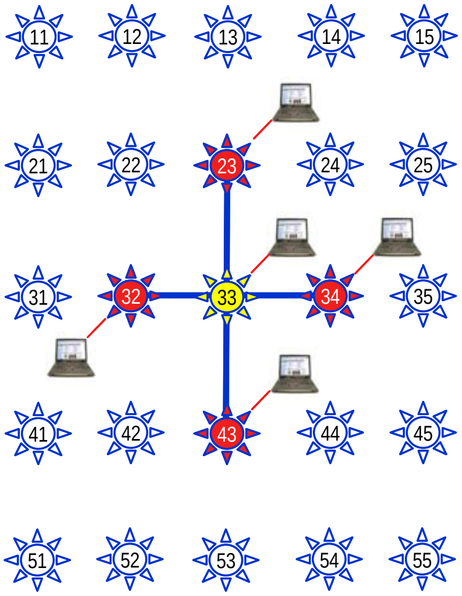

Figure 6. Single-hop scenario with multiple clients.

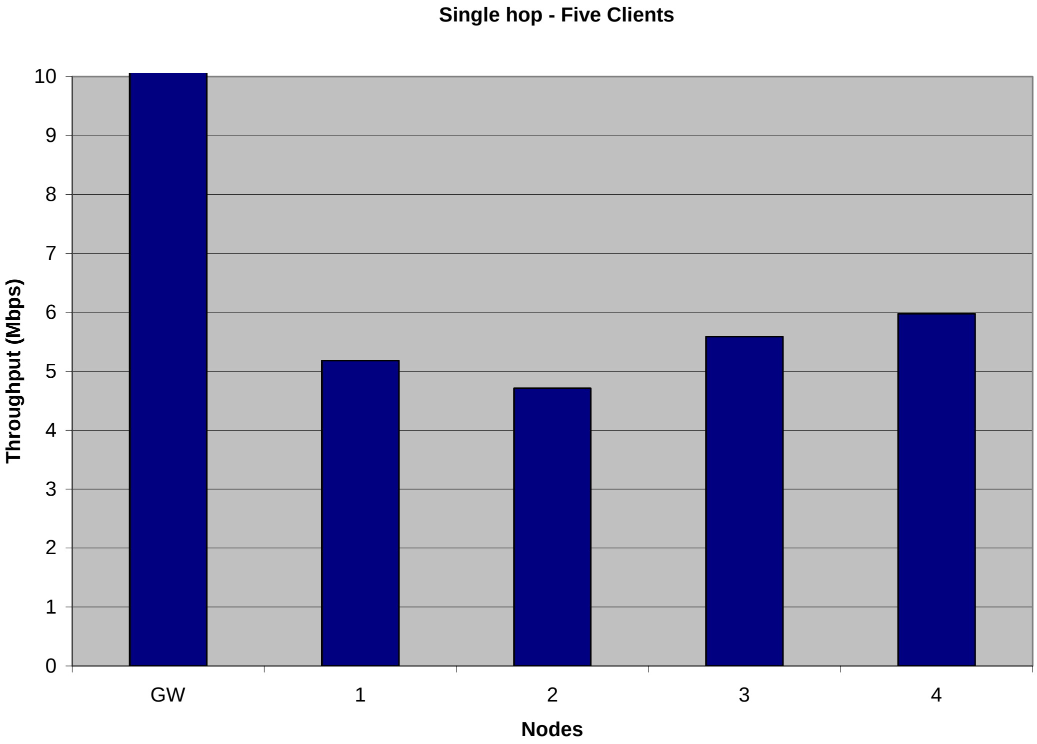

Figure 7. One-hop bandwidth.

From Measurement to Modeling

These two network topologies discussed in this white paper, the string of nodes and star topology, form the basic building blocks for any network deployment. By analyzing the network performance of each, a general model can be developed to predict network throughput. Based on the analysis of the referenced white paper Describing Multi-Hop 802.11 Mesh Throughput [1], network throughput can be estimated by dividing the maximum throughput by the sum of the hops for each client. While the maximum throughput for one-hop one-client scenario is measured to be 21Mbps as shown in Figure 4, this is partially limited by the QoS mechanism. The unbounded (without QoS) one hop throughput has been measured to be 24Mbps which we will use to predict the multi-hop performance of the network. For example, in the single-client five-hop test case in Figure 4 the measured throughput is 4.8Mbps. The model predicts the throughput available to the access point five hops from the gateway to be 4.8Mbps (24Mbps / 5 = 4.8Mbps). The six-client five-hop test case in Figure 5 is rate limited to 1.5Mbps. The model predicts the available throughput to each access point to be 1.6Mbps (24Mbps / {1+2+3+4+5} = 1.6Mbps). The 1.5Mbps ceiling was selected for the QoS rate limiting so as not to saturate the available throughput of 1.6Mbps per access point. For the five-client single-hop test case in Figure 7 the available throughput is 5Mbps. The model predicts 6Mbps (24Mbps / 4).

Based on this model, the available throughput can be predicted for the full 5 X 5 array of nodes illustrated in Figure 1. The total number of hops for each node is 60. The predicted available throughput for each access point is 400Kbps (24Mbps / 60). A 4 X 4 array of nodes can expect throughput of 750Kbps for each node (24Mbps / 32) and a 3 X 3 array of nodes can expect throughput of 2Mbps for each node (24Mbps / 12).

Conclusion

Tranzeo's EnRoute500 provides excellent performance through the combination of the dual radio architecture, a robust mesh routing algorithm, and an advanced QoS mechanism. The superior performance of the EnRoute500 clearly supports more bandwidth per access point compared to other solutions for the same network topology