CONSECUTIVE POINT ARCHITECTURE FOR BROADBAND WIRELESS ACCESS NETWORKS

Dan Gulliford, John Carter, David Oltman, Peter Chow

Triton Network Systems, Inc.
8529 SouthPark Circle, Orlando, FL 32819 USA
dgulliford@triton-network.com

Abstract - For wireless broadband systems to succeed as a major communication solution in the 21st century, they must be the primary communication solution in the bandwidth range of 10-100 Mb/s. This bandwidth range is not effectively served by copper or fiber optic networks. Key technical characteristics of a primary communication network are that it be reliable, maintainable, readily accessible, easy to deploy, cost effective, etc. This paper describes a reliable, maintainable, spectrally efficient wireless network architecture called the "consecutive point" architecture. A consecutive point network consists of a series of consecutive SONET or Ethernet radio hops typically arranged to form a ring. This paper describes the consecutive point wireless network architecture and its advantages which include self-healing, dense deployment, spectral efficiency, single POP manageability, in-service topology changes, and in-service software upgrades. Data from field trials and analysis are provided to support the consecutive point design.

I. CONSECUTIVE POINT ARCHITECTURE OVERVIEW

General

The consecutive point architecture is a strategy for deploying fixed broadband wireless networks using narrow beam, point-to-point radios, such as the Invisible Fiber™ unit (IFU) broadband radios from Triton Network Systems, Inc. The consecutive point network architecture is inspired by the ring-oriented designs used in many SONET networks. SONET ring designs, such as the Unidirectional Path Switched Ring (UPSR) or the Bidirectional Line Switched Ring (BLSR), automatically switch to protection channels in the event of an isolated network failure. The consecutive point architecture is also a ring-like network design in which traffic is automatically re-directed to an alternate route in the event of a single radio or a radio link failure. Consecutive point networks have the following advantages:

- self-healing
- dense deployment
- incremental deployment
- cost effectiveness
- scalability
- spectral efficiency
- high subscriber capacity
- manageability
- in-service upgradability
- ability to overcome line-of-sight obstructions

A consecutive point network consists of a series of customer access sites (typically office buildings) interconnected in a ring-like network by a series of broadband radio links. One of the access sites, designated as the "gateway site", serves as a link to a wide area network (e.g. the Internet) or a higher level backbone network. All non-local traffic passes through the gateway site. Figure 1 illustrates a small consecutive point ring consisting of a single gateway site and three customer access sites.

A consecutive point network differs from other common radio network designs, such as point-to-point and point-to-multipoint designs. It is unlike a point-to-point network in that it consists of a whole series of sites interconnected by radio links, rather than a single radio link interconnecting two sites. The interconnected radios act in concert to achieve high reliability. It is unlike a point-to-multipoint network in that all sites are served by the full available bandwidth, rather than some fraction of the bandwidth.

The IFU family of radios consists of two basic product groups. These are the Fast Ethernet IFU group and the SONET OC-3 IFU group. IFUs from either product group can be arranged in a consecutive point network.

The key features and advantages of the consecutive point network architecture are outlined below.

II. KEY FEATURES OF CONSECUTIVE POINT ARCHITECTURE

Self-Healing

A major feature of the consecutive point network design is its ability to reroute customer traffic around a single failed radio or radio link. In the example of Figure 1, if the radio link between customer building #1 and the gateway building were to fail, all traffic that would normally flow between those buildings is re-directed in the other direction around the ring.

In a Fast Ethernet IFU network, the third party routers and switches supporting the network perform the rerouting. In a network consisting of SONET IFUs, the SONET Add/Drop Multiplexers at each site perform the rerouting. When failures are artificially created in lab versions of Fast Ethernet IFU networks, the rerouting consistently occurs in less than 15 seconds.

In the Fast Ethernet version of a consecutive point network, no available capacity is reserved for protection purposes. All available routing paths carry customer data at all times. This contrasts with traditional SONET ring networks in which as much as half of the available capacity is reserved for protection. The Fast Ethernet version of the consecutive point network enables a more efficient use of the network resources by utilizing all available capacity.

Dense Deployment

Unlike some point-to-multipoint network designs where subscriber radios must be placed within narrow angular sections, the consecutive point architecture allows for very flexible and potentially dense deployment of IFUs at customer sites. There is no geographical limit on exactly where radios can be placed as long as the hop distances are kept within the allowed limits and the radios do not interfere with each other.

The potential for interference is minimized by the use of a high-gain antenna with a narrow beam width of 10 degrees. A narrow beam width allows all radio links to be placed within close proximity to each other. Another IFU feature that limits interference is the use of Adaptive Transmit Power Control™ (AdTPC). AdTPC automatically adjusts the transmit power of each radio to keep it at the minimum output level needed to maintain a 99.999 percent link availability, regardless of rain events.

Rapid Deployment

A consecutive point wireless network is relatively quick to deploy, especially compared to cable-based networks, which require digging permits to be obtained and cable to be buried in the ground. In urban environments, burying cable can be a very expensive, time-consuming proposition.

To deploy a wireless consecutive point network, roof rights and access to the telecom room must be obtained for each building. Figure 2 illustrates a typical configuration for a single customer building. Two IFUs are mounted on each rooftop and fiber optic cable is pulled through each building's risers leading from the IFUs to the telecom room and from there to each customer's wiring room. A power supply, battery backup, switch, and other equipment must be installed in the telecom room. Finally, each radio link is aligned and commissioned. This process is much simpler than burying cable in the ground.

Incremental Deployment

A consecutive point network can be built out incrementally such that service to existing customers is not interrupted, even as new buildings and customers are added to the network. Adding a new customer in a building already serving customers is very straightforward. A new cable is run from the switch in the telecom closet to the customer's switch or router. Basic configuration tasks are performed and the new customer is brought on-line. Adding a new building to an existing ring requires a new radio link to be spliced in. The primary service path must be temporarily broken and traffic rerouted to alternate paths for a short period while the new link is spliced in.

If needed, a new ring can be created at any time to add a new set of buildings to a service area. A metropolitan area backbone network can interconnect a series of rings scattered throughout a metropolitan area, or all rings can share a common gateway building. In the former case a fiber-based backbone network interconnects a set of subtending wireless networks. In the latter case a set of wireless networks converge at a single gateway building where a kind of collapsed backbone is implemented in a high capacity switch.

Cost Effectiveness

In addition to allowing for rapid deployment, the advantage of not having to lay cable also contributes to the cost effectiveness of a consecutive point wireless network. Obviously, installing cable in the ground is expensive. The cost savings of a wireless network are somewhat offset by the expense of spectrum ownership, but ultimately installing a wireless network is less expensive than installing a cable network.

Scalability

Consecutive point ring networks are scalable to whatever size is needed. If an existing ring is not yet fully populated, a new customer building can be spliced into it. If an existing ring becomes too large to add another building, a new ring can be created to accommodate the additional building and any others in the area.

Spectral Efficiency

The IFU design enables an efficient use of the available RF spectrum. Because of the IFU's narrow beam width and its use of AdTPC, a single transmit and receive channel pair can be reused by every radio link in a network.

High Subscriber Capacity

A consecutive point network delivers the full network capacity to every customer building on the network.

For a Fast Ethernet IFU network, shared access to the full 100 Mbps full-duplex network is delivered to each building. Individual customers can be given either full 100 Mbps access or some lesser increment, such as 10 Mbps.

For a SONET OC-3 IFU network, each radio link acts as an "invisible fiber" to carry the 155 Mbps OC-3 signal in both directions between the add/drop multiplexers located in each building. An IFU-based OC-3 network design is exactly analogous to a typical fiber-optic deployment. The only difference is that the bidirectional OC-3 signal is carried through the air.

Manageability

A consecutive point network is managed through a separate channel of RF bandwidth reserved for that purpose. This network management channel is called the "radio overhead" channel. It does not steal any capacity from the either the 100 Mbps Fast Ethernet channel or the 155 Mbps SONET OC-3 channel. Any IFU in a network can be reached through the radio overhead channel using Ethernet and TCP/IP protocols. Each IFU includes two 10 Mbps Ethernet access ports, either of which can be used to manage that IFU and every other IFU in the network. It is possible to manage an entire IFU ring through a single 10 Mbps Ethernet connection.

The radio overhead format and usage is exactly the same for both the Fast Ethernet and the SONET OC-3 versions of the IFU. In both cases Ethernet and TCP/IP protocols are used to send and receive SNMP messages. Both IFU product versions use 10BASE-T and 10BASE-FL interfaces for network management connections.

Like the customer data channel, the radio overhead network is fully self-healing. In a consecutive point ring network, like the one illustrated by Figure 1, if a radio link fails or an IFU fails, network management traffic is re-routed along an alternate path to avoid the failure.

In-Service Upgradability

Once a consecutive point network is installed and commissioned, the need occasionally arises to upgrade the software of all the IFUs in the network. Fortunately, it is possible to do this without interrupting service to customers for more than a few seconds.

The software update process for a single IFU involves first downloading new software to the IFU's non-volatile memory through the OAM&P channel. The next step is to reboot the IFU. While the IFU is rebooting and the radio link is down, affected customer traffic is re-routed along alternate paths as previously described. Customers see only a brief service interruption lasting no more than 15 seconds for Fast Ethernet or less for OC-3. To limit the effect on customers, software upgrades can be done during off-peak hours.

Overcoming Line-of-Sight Obstructions

A consecutive point network lends itself well to situations where there are obstructions to the line-of-sight between the gateway building and the outlying customer buildings. The network can be geographically arranged so that the radio links angle around line-of-sight obstructions such as tall buildings. Linking the radios from building to building also makes it possible to reach customer buildings that are much further away from the gateway building than the maximum single hop distance.

III. PERFORMANCE DATA

During field trial of a Fast Ethernet consecutive point network, a variety of measurements were made of the network's packet handling performance. This section summarizes the results of these measurements for throughput, packet loss, and latency.

The trial network that these measurements were taken from consists of four buildings and four radio links. This trial network is very similar to the example illustration of Figure 1. Measurements were taken on the entire IFU network such that all measured Ethernet traffic flowed all the way around the ring, through four radio links and eight IFUs. For example, the latency measurements indica

te the amount of time it took for packets to traverse four radio links and eight IFUs. No Ethernet traffic flowed through any third party switches or networking equipment other than IFUs.

All testing was done using a SmartBits™ SM-200™ from Netcom Systems, Inc. All tests were performed using bidirectional traffic, flowing both ways around the ring simultaneously. The listed results are all averages of measurements taken in each direction.

Throughput

This test measured the highest Ethernet traffic rate for various packet sizes at which no traffic loss or errors occurred during the sampling interval. Measurements were taken of seven packet sizes ranging from 64 to 1518 bytes. The results are expressed as a percentage of the maximum throughput for the corresponding packet size. For example, a result of 90% for 64-byte packets would mean that the maximum throughput at which no loss or errors occurred was 90% of the maximum possible throughput for 64-byte packets.

Packet Loss

This test measured the ratio of traffic lost vs. traffic offered for a variety of packet sizes and offered throughput rates. The measured packet sizes were 64, 512, 1014, and 1518 bytes. The offered throughputs were 55, 75, and 95 percent of the maximum. For each packet size, measurements were taken with traffic loads of each of the three throughput rates.

The results of this testing showed no packet loss at all for any packet size or throughput.

Latency

This test measured the average interval of time between the moment the first bit of each frame left the transmitting port of the SmartBits™ tester, to the moment the first bit of each frame was received on the receiving port of the tester. Traffic was offered at a 99% rate in both directions during the test. Latency was measured for seven frame sizes ranging from 64 bytes to 1518 bytes. The results are expressed as an interval in microseconds.

IV. CONCLUSION

Consecutive point networks have several advantages over other common network designs. When radios are deployed in a ring configuration, alternate routing paths protect customer traffic from radio failures or radio link failures. Alternate routing paths also make it possible to replace or upgrade radios without shutting down the entire network. In combination with narrow beamwidth and AdTPC technologies, consecutive point technology allows denser networks to be deployed. The same technologies also enable radio frequencies to be reused in close proximity without interference. Consecutive point networks are flexible with regard to radio placement, such that obstructions can be avoided and large areas covered.

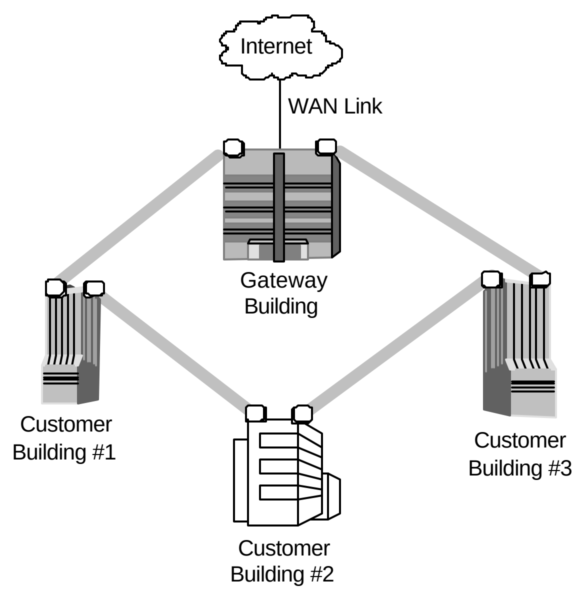

Figure 1. Consecutive Point Network Example

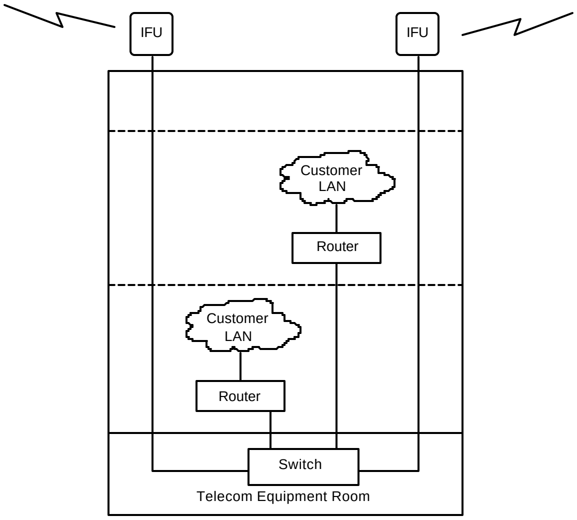

Figure 2. Typical Building Configuration

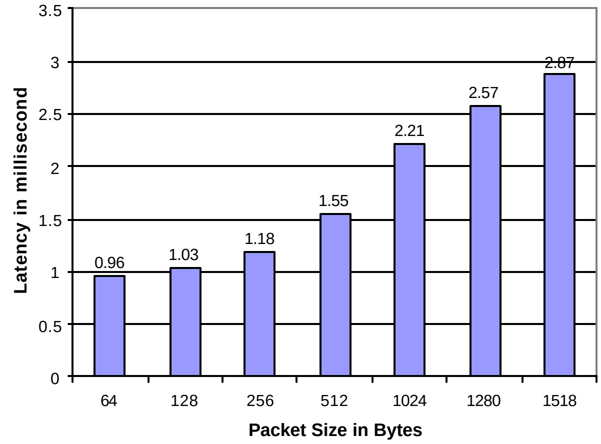

Figure 3. Trial Network Latency

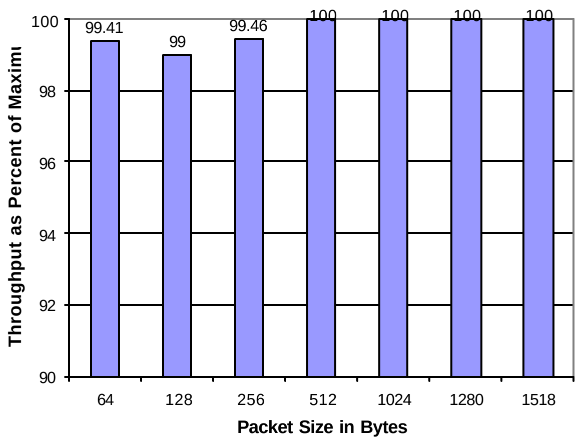

Figure 4. Trial Network Throughput