Resilient Wireless Networking

M. Scott Corson
Flarion Technologies

1. Introduction

Redundancy underpins resilient communications. Robust network architectures necessarily avoid single points of failure and rely upon distributed or decentralized control mechanisms. We will argue here that Internet Protocol (IP)-based networks possess the best-suited architecture for designing distributed networks and, consequently, that IP-based Cellular Networks enable the best solutions for cost-effective, survivable, wireless access.

2. IP Design Principles

The fundamental goal behind the development of the Internet was to enable efficient *multiplexing* of heterogeneous, interconnected communication networks [1]. Packet switching was chosen as the underlying multiplexing technology. Subsequently the resultant IP architecture possesses the following properties: packet-switched forwarding, asynchronous operation and unidirectional traffic flows. In contrast, traditional telecommunications networks generally possess the following properties: circuit-switched forwarding, synchronous operation and bi-directional traffic flows. From an architectural perspective, it is apparent that packet-switching is the more general of the two approaches. Essentially any service that can be realized within a circuit-switched network can also be realized in a packet-switched network, but the converse does not hold. For example, one can realize connection-oriented transport over a connectionless, packet-switched fabric, but one cannot effectively realize connectionless transport over a connection-oriented, circuit-switched fabric. Packet switching is the more general concept in practice.

A secondary goal behind the Internet's development was "survivability in the face of failure". A key methodology employed in achieving this characteristic is network design in keeping with the 'end-to-end argument'—a principle suggesting that functions placed at low levels of a system may be redundant or of little value when compared with the cost of providing them at that low level [2]. In the Internet this principle manifests itself in the placement of protocol complexity/state at the *edge* of the network (end hosts, servers) whenever possible, leaving the network core functioning as a collection of relatively stateless, datagram forwarding entities (routers). This is the so-called 'dumb network, smart terminal' scenario.

This approach is key to the design of robust distributed systems as it results in a state distribution model termed 'fate-sharing'—"The fate-sharing model suggests that it is acceptable to lose the state information associated with an entity if, at the same time, the entity itself is lost" [1]. This is a powerful concept and it ties in closely with the end-to-end argument and survivable system design. To maintain synchronization of a communication session between two end hosts, a network architect can choose to distribute session synchronization state regarding the session along all network nodes on the session's path (a state 'replication' model), or concentrate it only in the communicating end hosts. Should any of the intermediate nodes fail, the former 'replication-based' state model requires complex, distributed session state management algorithms to correctly react to the failure and preserve session integrity. In contrast, state management in the end-to-end model employing fate-sharing is greatly simplified and concerned only with the health of the end hosts. Their existence and that of their communication session are inherently bound together. In practice such systems are far simpler to engineer and deploy.

3. Resilient Wireless Access

How does all this relate to wireless access? Traditional wireless Wide-Area Networks (WAN) (e.g. TDMA, CDMA) are circuit-switched, and are essentially mobility-enabled replicas of the voice-centric PSTN. Their Radio Access Networks (RAN) contain numerous Base Stations (BS) connected to a smaller number of Base Station/Radio Network Controllers (BSC/RNC) which are further connected into a centralized Mobile Switching Center (MSC)—the cellular equivalent of a 5ESS switch. Each of these RAN elements is stateful with respect to a given voice call. Furthermore, these RANs are hierarchical in nature, with an MSC being a single point of failure for a very large number of cells. These are classic embodiments of the 'smart network, dumb terminal' architecture. Providing these networks the ability to survive the loss of any of these session-stateful network elements requires the complex engineering associated with the replication state management model—a daunting task given the inherent complexity of a circuit-switched, mobile RAN which requires intelligence in the circuit data path.

In contrast, the deployment of an IP-based RAN—a RAN architecture where session state is principally held in and managed by a mobile end host—greatly facilitates the practical deployment of survivable wireless networks. Firstly, the concentration of control/session state in the end host is consistent with the fate-sharing model, such that only the loss of the host itself need destroy its active sessions. Secondly, the provision of a stateless core greatly facilitates the deployment of redundant network communication paths, as 'routed' networks based on connectionless packet forwarding are inherently robust with respect to router/link failures and require no particular topology to function, in which case non-hierarchical, 'mesh' topologies of many forms are possible. Base stations may be provisioned with multiple backhaul links, and multiple redundant, virtual overlay networks may be accessible in parallel over these backhaul links, providing base stations and mobiles access to redundant communication services. Thirdly, this characteristic of topology flexibility enables rapid deployment of IP RANs in areas of need, where the backhaul link to/from base stations can be provided via wireless/satellite if necessary.

This brings us to the nature of base stations in IP RANs. In traditional networks, base stations are essentially Radio Frequency (RF) 'physical layer' processors, taking traffic to/from a BSC/RNC and converting it from/to RF in communication with the mobile. Virtually no 'higher layer' intelligence exists in these devices. In contrast, an IP base station is necessarily a layer 3 device as it must handle the routing of datagram traffic, and the delivery of said traffic to/from the mobiles/network in an IP Quality-of-Service (QoS)-aware manner. Moreover, service layer intelligence as it relates to controlling network access and IP QoS over the air must also reside here. Thus, it is the nature of an IP RAN that the base station is second only to the mobile as a loci of control—the mobile orchestrates mobility and session management, whilst the base station controls access to the network and its services. This latter characteristic of distributing access and service control across a large set of intelligent IP base stations greatly facilitates the provision of survivable network access. Each base station can manage its interaction with backend systems, selecting from sets of non-failed servers over redundant communication paths. It can also coordinate communications between end hosts in its coverage area and the areas of nearby base stations in the absence of connectivity to the larger network. Thus as long as the base station itself exists, mobiles in its coverage can communicate, and communication can be extended to mobiles under reachable surviving base stations.

With the preceding in mind, the feasibility of a resilient IP RAN depends on the existence of wireless access technology designed in accordance with the requirements of packet switching. The air interface itself must be packet-switched for efficient operation, and support reliable, low latency delivery of data traffic such that the protocols of the Internet require no modification for correct operation. The technology must also enable a RAN architecture that is packet-switched—connectionless, asynchronous with unidirectional traffic flows—such that the base stations can directly connect to standard IP routers using any means of backhaul. Given such a technology, it is hopefully apparent that IP-based networks possess the best architecture for designing distributed systems and, consequently, that IP-based Cellular Networks enable the best solutions for cost-effective, survivable, wireless access.

References

[1] D. Clark, "The Design Philosophy of the DARPA Internet Protocols", in Proc. SIGCOMM 1988.

[2] J. Saltzer, D. Reed, and D. Clark, "End-to-End Arguments in System Design", in ACM Trans. on Computer Systems, 1984.