

Modem Link-Property Advertisements

William D. Ivancic
NASA Glenn Research Center
Cleveland, OH 44135
United States
+1-216-433-3494
william.d.ivancic@nasa.gov

Lloyd Wood
University of Surrey alumni
L.Wood@society.surrey.ac.uk

Dan Shell
University of Akron
Akron, OH 44325
United States
+1-216-970-8260
dshell@uakron.edu

Daniel Floreani
Cisco Systems
Adelaide, South Australia 5000
Australia
+61-8-8124-2206
danielf@cisco.com

Rajiv Asati
Cisco Systems
Research Triangle Park, NC 27709
United States
+1-919-392-8558
rajiva@cisco.com

Abstract— Devices and associated applications are increasingly connected to a variety of smart modems that have dynamically varying link characteristics. For example, incoming and outgoing link rates can be varied over time with adaptive coding and modulation to suit the channel characteristics. In addition, those links may go up and down due to a variety of factors. The link rate and conditions offered by the modem to connected devices and associated applications therefore vary. In order to autonomously condition traffic and allow the network to get the most out of the modem's link capacity, the upstream devices and applications need to be informed of the modem's link conditions. This becomes a form of cognitive networking. This paper will discuss the problem and present a simple protocol that can be used to provide upstream devices and applications with downstream link conditions. While the protocol in this document is described in the context of modem radio-frequency link properties, it can also be broadly applied to other scenarios such as cryptographic devices.

TABLE OF CONTENTS

1. INTRODUCTION	1
2. SMART MODEMS	1
3. TERMINOLOGY	1
4. SCENARIOS	2
5. PROTOCOL CONCEPTS	3
6. ALTERNATIVE APPROACHES	4
7. SUMMARY	5
REFERENCES	5
BIOGRAPHIES	6

1. INTRODUCTION

Cognitive networking, self-configuring systems, and autonomous systems are of great interest to reduce the cost of managing and configuring complex, heterogeneous systems and to provide resilient applications and services. “Cognitive networks enable the vision of pervasive computing, seamless mobility, ad hoc networks and dynamic spectrum allocation”[1]

An autonomous system perceives current conditions, and then plans, decides, and acts on those conditions based on

predefined rules and algorithms. A cognitive system adds the following capability: A cognitive system learns from the consequences of its actions and uses this knowledge to improve the future decisions. In either case, the key is the system needs to be aware of its environment. Thus, techniques and technologies must be developed that expose system parameters and provide methods for adjusting those parameters. The protocol described provides a method to advertise the modem characteristics, which can be used by intelligent systems and applications to adapt to their environment.

2. SMART MODEMS

A number of potential solutions for reacting to changes in links from smart modems have been proposed for router/modem interactions:

- RFC-5578, Point-to-Point Protocol (PPP) over Ethernet (PPPoE) Extensions for Credit Flow and Link Metrics [2]
- Dynamic Link Exchange Protocol (DLEP) [3]
- Link Properties Advertisement from Modem to Router [4]

These solutions do not necessarily address devices and associated applications that may be multiple network hops away from the modem.

This paper discusses the problem, and presents a simple protocol that can be used to provide upstream applications and network-attached devices, such as routers, with downstream link conditions. Much of the ideas presented here build on previous work documented in reference 4, “Link Properties Advertisement (LPA) from Modem to Router.”

3. TERMINOLOGY

Attached device - a host (computer) running some type of application or a router.

Uplink - Inbound data on the RF link of the modem.

Downlink - Outbound data on the RF link of the modem.

Upstream - The direction from the modem to the device or application; i.e. away from the RF link.

Downstream - The direction from the device or application to the modem; i.e. toward the RF link.

LPA - Link Property Advertisement

4. SCENARIOS




Figure 1 Space-based Sensor Satellite

Figure 1 illustrates the configuration of a space-based sensor onboard a satellite. The host is co-located with the radio and connected via a serial link. Since all clocking can be derived from the modem and how it drives the serial link, a rate-based application on the host can be told to run at the outgoing line rate without overwhelming the modem. Furthermore, congestion control is not of concern if the space/ground link is a dedicated, private link.




Figure 2 Modem/Router Directly Connected

Figure 2 illustrates a wireless modem directly connected to a router. Because Ethernet interfaces are plentiful and cheap, the modem router is connected via high-speed Ethernet. The router needs to know the actual link rate offered by the modem in order to properly set the quality-of-service (QoS) and traffic shaping for that link; simply sending 10/100 Mbps traffic to the modem and having the modem drop most of that traffic, because its outgoing link is slower, is not acceptable. An outgoing traffic flow using the transmission control protocol (TCP) will autonomously adapt to rate changes by way of the TCP congestion control algorithms, but will do so with inefficient link utilization as TCP's control loop hits the limit of the outgoing link. Rate-based protocols may not autonomously adapt – particularly if there is no effective way to implement congestion control such as on a unidirectional link.

It is possible to configure QoS and shaping on the router's Ethernet interface to match the modem's link rate, effectively extending the modem's link an extra hop to the router. However, such a static configuration is only useful if the modem's link rate is also static. Many modern modems are able to vary their communications according to the channel characteristics they experience, or due to negotiation with the modem at the other end of the link. Examples include the Adaptive Coding and Modulation

(ACM) and Variable Coding and Modulation (VCM) used in Digital Video Broadcast for Satellite (DVB-S2), and ADSL2's (Asymmetric Digital Subscriber Line) Seamless Rate Adaptation (SRA). (In general, “varying coding” has one end varying its link speed, while “adaptive coding” has both ends communicating and negotiating changes in link speed.) Link characteristics may also vary due to layer-2 handovers, e.g. IEEE 802.21 media-independent handoffs. The router needs to learn about changes in modem link characteristics in order to vary its current QoS and shaping configurations to match the current characteristics and get the most from the modem's link. The aim is to make the link between router and modem behave as an extension of the modem's own link.




Figure 3 Generic Single-Hop Topology

Consider the topology in Figure 3. Here, an application is running on a host that is behind the router. The application needs to know the downlink status of the modem and the outgoing link rate in order to properly shape the data. For example, the application may be a video camera capable of setting the frame rate relative to the available bandwidth. By having the modem advertise its link properties, the application can autonomously adapt to the offered rate. Note, in this instance, the host and application are one hop removed from the router. Thus, advertisements of modem properties have to be able to pass beyond the link-local network attachment of the modem.




Figure 4 Generic Multi-hop Topology

Figure 4 depicts a multi-hop system with one router attached to multiple radios. This configuration shows the need for modem identification, as there may be more than one downstream link that needs to be considered by the router and applications. In addition, the applications that need to shape their traffic are multiple hops from the modems. Knowing the current data rates and whether or not either link is available can enable the applications to make intelligent decisions regarding traffic shaping and when to send. For example: one such application may be to implement reactive fragmentation in Disconnection/Delay Tolerant Networking (DTN) as soon as the link is down[5].

One can replace the modem in Figures 1 through 4 with a cryptographic device and have the same basic problem. For




Figure 5 Crypto/Router Directly Connected

example, Figure 5 is the same basic scenario as Figure 3. Figure 5 illustrates a cryptographic device directly connected to a router. Here, both interfaces may be high-rate Ethernet interfaces with a cryptographic device in between. The rate on each interface may not match. For example, the red-side interface may operate at 100 Mbps with the black-side interface is operating at 10 Mbps. The cryptographic devices normally operate at line rate, thus we have a rate-mismatch problem. Also, during rekeying the offered rate to the normal traffic may be reduced for a period of time or the cryptographic tunnels may drop, resulting in performance hits or momentary loss of communications.

In other traditional Internet Protocol (IP) cryptographic devices ("cryptos") it is hard to sense the real rates offered through "the system". It is feasible for such devices to work this out on their black side - and the red side can simply advertise the "offered rate". The black side can obtain knowledge of its downstream link state via modem advertisements, router advertisements, or probes, and pass this on to the red side via approved methods. The red side can then advertise its rate via the LPA protocol.

The red side can then advertise its rate via the protocol - but the crypto can offer this rate information on the red side (perhaps by doing its own probes on behalf of the connected devices).

5. PROTOCOL CONCEPTS

The simplest approach to this problem is to have the modem advertise its link conditions on its other local interfaces using user datagram protocol (UDP) packets [6] sent to link-local multicast addresses. This approach requires no explicit configuration setup to provide information to connected devices. UDP is well understood and widely available, making deployment relatively easy for all types of modems, routers and other connected devices.

One area of debate, yet to be resolved, is the best way to advertise link properties to upstream systems. Use of the "all routers" multicast address for devices attached to the local link is straightforward and greatly simplifies implementation. However, for upstream devices not directly connected to the local link either or both use of IPv4 organizational-scoped multicast and IPv6 site-local multicast or unicast advertisements are possible.

Use of scoped multicast simplifies implementation within the modem as there is no need for configuration or knowledge of what is connected upstream - particular unicast addresses or device identities. However, the

upstream network must support and be configured to propagate scoped multicast.

Use of IPv4 organizational-scoped multicast and IPv6 site-local multicast could handle both devices that are directly connected to the modem as well as host and applications that are multiple hops away [7][8]. These may be used with no explicit configuration in the modem, but would require approval of a new organizational/site-local multicast type by the Internet Assigned Number Authority (IANA). However, there are potential deployment issues with scoped multicast. Administratively scoped multicast can have some unusual deployments that may result in unforeseen global traffic. For example, in some implementations, site-local is the global corporation. This could result in *e.g.* modem LPAs suddenly flooding a planet-wide multi-protocol-label-switched (MPLS) net.

Conversely, unicast packets may be sent to known hosts. Use of unicast messages for link property advertisements requires no scoped multicast support but does require explicit configuration within the modem; however, this is simple and straightforward as the end systems where such configurations apply are not expected to be large or complex and likely to consist of only a handful of hosts/applications.

Protocol Signaling

Link property advertisements are sent periodically or as notifications of link-layer events when they happen. This falls into the scope of Detecting Network Attachment (DNA) processes as specified in RFC4957 [9].

The modem sends UDP updates about changing link and interface conditions to connected devices using link-local multicast packets, and to upstream hosts and applications using unicast packets (*i.e.* a link rate changes due to a coding change, or the link and its interface go up or down). As well as sending on event-triggered updates, the modem should send periodic advertisements about link conditions, in case new devices have been connected, *e.g.* on a broadcast Ethernet LAN. These updates are sent over both IPv4 and IPv6.

The specifics of the protocol message format are provided in the draft specification, draft-ivancic-moboopts-modemlpa [10].

Currently, only one Block Type has been considered, the Rate Block [Table 1]. A few of the key components of this block type are:

- The maximum data rate is defined by a 32-bit word corresponding to a limit of 4 Gbps.
- This format enables one to describe link characteristics simultaneously for both link directions, or separately for the incoming and outgoing links.

- Link Up / Link Down information is conveyed via the interface flag “U”.

Other possible blocks, not yet defined here, could express measured error rate, current forward error-coding rate, latency (propagation delay, serialization delay), link MTU size, indicate link-level security mechanisms in use, or provide authentication.

Table 1 Rate-Block Type

1 1 1 1 1 1 1 1 1 1 2 2 2 2 2 2 2 2 2 2 3 3																														
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1									
UDP source port										UDP destination port										B										
UDP length										UDP checksum										T										
Block Type ID (Rate Type 1)										Length										L										
no. of links										link rate size										modem flags (15 bits unused)										Y
unique modem interface identifier															O															
MTU										interface flags										P										
current link rate (varies) - 32 bits is rate size of 1															E															
minimum supported rate - 32 bits is rate size of 1															K															
maximum supported rate - 32 bits is rate size of 1															S															
IPv4 address of modem's local link interface, if 4 flag set															U															
IPv6 address of modem's local link interface, if 6 flag set															B															
															F															
															4															
															6															
															U															
															I															
															S															

Uses of these notifications

A device attached to the local link must be able to receive link property advertisements via UDP/IP packets sent to the "all routers" multicast address. Information from this message is used to construct or update the QoS and to shape policies applied on the interface for traffic directed through the modem's link. Other network-attached devices may receive advertisements via IPv4 organizational-scoped multicast and IPv6 site-local multicast or unicast advertisements.

An attached router may use this information to update its routing table so that the routing information associated with a route through the modem is either deleted or added or updated with a new metric. Changes in the routing table information are then propagated further.

The modem may damp changes to Link Instance information, to prevent advertising transient changes, in line with RFC4907 [11]. A router may also damp responses to frequent changes in Link Instance information received, so that related QoS policies and routing information are not updated until a certain time period has elapsed. This hysteresis would be useful in the case of a rapidly varying

link rate, where the router would stick to the minimum rate seen to ensure stable operation.

A router may also use this information as input to *e.g.* Call Admission Control (CAC) functionality, to reserve capacity for calls. This can prevent use by registered applications such as telephony call setup *etc.* in the event of less-than-desired available link capacity through the modem's interface.

6. ALTERNATIVE APPROACHES

Other approaches to this problem have been proposed. None of these approaches addresses the need to extend the modem advertisement beyond the locally attached device.

The simplest approach to this problem is to have a fixed serial interface between modem and router, with the modem altering the serial rate clock to match the speed of its link, and the router measuring the rate of the received clock [Figure 1]. However, fast serial interfaces are unfashionable, and faster Ethernet now dominates the world.

We considered using Internet Control Message Protocol (ICMP) [12] [13] to provide this rate information, using the framework for carrying extended information in ICMP messages [14]. However, this extended information can only be carried in Destination Unreachable and Time Exceeded ICMP messages (for both IPv4 and IPv6) and Parameter Problem (for IPv4 only) ICMP messages. These messages are responses to specific problems, and should not be overloaded for general advertisements. The appropriate ICMP message type would be the obsolete Information Request message (type 15), though requests are also sent unsolicited for this use. Other factors in preferring UDP to ICMP is that sockets programming for UDP is easier than for ICMP, easing implementation, and that UDP is less likely to be blackholed in networks.

We have previously outlined an approach leveraging the Access Node Control Protocol (ANCP), used in the head-ends of Digital Subscriber Line (DSL) networks, within Digital Subscriber Line Access Multiplexers (DSLAMs) [15]. However, ANCP is not lightweight and depends on the General Switch Management Protocol (GSMP), which depends on TCP. The ANCP workgroup has been focused on the DSL headend, and has yet to extend beyond that environment, while the LPA protocol is more generally useful at the tail-end in customer networks.

Another approach aimed at improving communication between modems and routers is outlined in RFC5578. That approach assumes a PPPoE infrastructure. PPPoE is not always architecturally suitable to network needs, and requiring PPPoE infrastructure and introducing authentication dependencies for what was just a simple local modem-router (or other attached device) problem is, in our opinion, complexity overkill. That approach may be suitable as an upgrade to existing environments that are already using PPPoE. Adoption of the metrics described in

RFC5578, but with communication separate from PPPoE, could be very useful for a wider market, and would provide more information than the link rate information outlined in this draft.

Ethernet pause frames have also been suggested as one way of slowing the Ethernet link to match the modem's link a hop further along [16]. This approach has the disadvantage of being tied to a particular layer-2 technology, while fine-tuning the pause frames to match the modem's offered link rate has the potential to introduce complex control loops and problems as the paused and now-bursty Ethernet rate only approximates the modem link rate and interacts with QoS and shaping delay mechanisms on the router.

The Dynamic Host Configuration Protocol (DHCP) is intended for address configuration of hosts, so is not considered suitable as a way of piggybacking modem link property advertisements.

7. SUMMARY

We have discussed the problem of taking advantage of the capabilities offered by smart modems with varying link speeds, and have presented a simple protocol that can be used to provide upstream devices and applications with downstream link conditions. The protocol in this document is described in the context of modem RF link properties, but can also be broadly applied to other scenarios such as cryptographic devices. The ability to sense and react to such information is critical for new and developing technologies such as cognitive radios and cognitive networks, and deserves further attention.

REFERENCES

- [1] Q. H. Mahmoud et al., "Cognitive Networks – Towards Self-Aware Networks," Wiley & Sons, Ltd., 2007.
- [2] B. Berry, S. Ratliff, T. Kaiser, M. Adams: "PPP Over Ethernet (PPPoE) Extensions for Credit Flow and Link Metrics", RFC 5578, February 2010, <http://tools.ietf.org/html/rfc5578>
- [3] S. Ratliff, B. Berry, G. Harrison, S. Jury, D. Satterwhite: "Dynamic Link Exchange Protocol (DLEP)," draft-sratliff-dlep-01 (work in progress as an internet-draft, May 2011 <http://tools.ietf.org/html/draft-ietf-manet-dlep>
- [4] L. Wood, R. Asati, D. Floreani: "Link properties Advertisement from Modem to Router," draft-wood-dna-link-properties-advertisement-01, (expired internet draft, <http://tools.ietf.org/html/draft-wood-dna-link-properties-advertisement>
- [5] K. Scott, S. Burleigh: "Bundle Protocol Specification," RFC5050, November 2007, <http://tools.ietf.org/html/rfc5050>
- [6] J. Postel "User Datagram Protocol," RFC 768, August 1980, <http://tools.ietf.org/html/rfc768>
- [7] D. Meyer: "Administratively Scoped IP Multicast," RFC 2365, July 1998, <http://tools.ietf.org/html/rfc2365>
- [8] R. Hinden, S. Deering IP: "Version 6 Addressing Architecture," RFC 2373, July 1998, <http://tools.ietf.org/html/rfc2373>
- [9] S. Krishnan, N. Montavont, E. Njedjou, S. Veerapalli, A. Yegin: "Link-Layer Event Notifications for Detecting Network Attachments," RFC 4957, August 2007 <http://tools.ietf.org/html/rfc4957>
- [10] W. Ivancic, L. Wood, R. Asati, D. Floreani, D. Shell: "Modem Link Properties Advertisement," draft-ivancic-mobopts-modemlpa, qork in progress as an internet draft, October, 2011.
- [11] Internet Architecture Board (B. Aboba, Ed.): "Architectural Implications of Link Indications," RFC 4907, June 2007, <http://tools.ietf.org/html/rfc4907>
- [12] J. Postel, "Internet Control Message Protocol," RFC 792, September 1981, <http://tools.ietf.org/html/rfc792>
- [13] A. Conta, S. Deering, M. Gupta: "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification," RFC 4443, March 2006, <http://tools.ietf.org/html/rfc4443>
- [14] R. Bonica D. Gan C. Pignataro, "Extended ICMP to Support Multi-Part Messages," RFC 4884, April 2007, <http://tools.ietf.org/html/rfc4884>
- [15] Floreani, D. and L. Wood, "Extension of ANCP for Satellite and Terrestrial Wireless Modem Networks", draft-floreani-ancp-wireless-00, expired internet-draft June 2007, <http://tools.ietf.org/html/draft-floreani-ancp-wireless>
- [16] Ge, A. and G. Chiruvolu, "Diffserv Compatible Extended Pause (DiffPause) for Fair Congestion Control in Metro-Ethernet", IEEE International Conference on Communications, vol. 2 pp. 1248-1252, June 2004.

BIOGRAPHIES



William Ivancic has over twenty-five years of experience in network and system engineering for communication applications, communication networking research, state-of-the-art digital, analog and RF hardware design and testing. He is currently a senior research engineer at NASA's Glenn Research

Center, where he directs the hybrid satellite/terrestrial networking, space-based Internet, and aeronautical Internet research. He has led research efforts to deploy commercial-off-the-shelf (COTS) technology into NASA missions, including the International Space Station and Shuttle. Mr. Ivancic has performed research on advanced routing research for space-based and aeronautic-based networks. Of particular interest is large scale, secure deployment of mobile networks, including Mobile-IP and mobile router technology.

Mr. Ivancic is also principle of Syzygy Engineering, a small consulting company specializing in communications systems and networking as well as advanced technology risk assessment. Mr. Ivancic is currently performing research and development on identity-based security and key and policy management and distribution for tactical networks - particularly mobile networks.



Lloyd Wood is a Chartered Engineer with experience in computing, networking and aerospace. As a space initiatives manager for Cisco Systems, Lloyd had responsibility for CLEO, the Cisco router in Low Earth Orbit.

Lloyd spent some years contributing to the Internet Engineering Task Force and modifying IOS, Cisco's router software...so he's gone on to fly his own code in space. Working with colleagues at NASA Glenn Research Center and Surrey Satellite Technology Ltd, Lloyd achieved the first tests from space of IPv6 and of the delay-tolerant networking bundle protocol for the "Interplanetary Internet." Lloyd gained his PhD from the Centre for Communication Systems Research at the University of Surrey, where he researched internetworking and satellite constellations, where he later returned as a research fellow.



Mr. Dan Shell, CEO of DSHELL Network Architects, LLC, has been focused on Wireless, Mobile Networking and Satellite. Dan was the Primary Investigator for the following projects:

- Communications for the Aerostat project
- Low-Cost MANET
- Identity-Key-Based Security

Dan is one of the lead developers of Mobile Networking and was the main driver behind the development of the Cisco Systems 3251 PC/104 Router. Dan has been the lead Engineer in the support of the Cisco/NASA Space Act Agreement for joint research into networking over high delay and high data rate networks. In addition to supporting NASA Glenn Dan Shell also supports various other Federal Agencies in deploying network solutions. Previously Dan was the WAN network technical specialist for BP America and then GE Lighting. Recently, he has also been working with NASA-Glenn in researching IP over satellite and internet nodes in space. Dan currently teaches networking classes at University of Akron, Ohio and Lorain County Community College, Ohio.



Rajiv Asati, a Distinguished Engineer who has been at Cisco Systems for over twelve years, is responsible for driving "Next Generation Network" architecture & strategies. He incubated the first IP NGN system (CESNA) for Cable SPs and has led its world-wide adoption since 2006. He is a Subject Matter

Expert, a key go-to technical resource, and well respected thought-leader across many key IP/NGN technologies such as MPLS, IP Routing, Multicast, L2/L3VPN, IPv6, QoS and Video. He has approximately forty pending or issued patents, Cisco CCIE certification, and has written 100+ engineering documents/design guides/white-papers. Rajiv is a frequent speaker at various international conferences/seminars and regularly contributes to CableLabs and the IETF (with over ten RFCs and internet-drafts).

Prior to joining Cisco, Rajiv worked with AT&T Labs and Nortel.



***Dr. Daniel Floreani** is currently a Network Architect with the Advanced Services Team within Cisco Systems, specializing in space, defense and mobility. Based in Adelaide, he supports business and technical development of the space and defense markets for Cisco Systems. His primary responsibilities are to develop network architectures, observe technical trends and develop strategies to enter new markets. Daniel was also a pre-sales engineer for five years at Cisco, focusing in voice, video and wireless technologies.*

Prior to joining Cisco, Daniel was a senior communications engineer on a large defense project and carried out research and development for the department of defense at Defense Science and Technology Organization (DSTO). Daniel has a PhD from the University of South Australia and both a science and an engineering degree from the University of Adelaide.