# Adaptive Link Rate (ALR) for Full-Duplex Ethernet

Pre-Proposal for an IEEE 802.3 Study Group

**January 9, 2006**

**Summary:** There is a substantial energy savings potential from implementation of a standard method for quickly changing link data rates in response to traffic levels for both 1 Gb/s and 10 Gb/s Ethernet. The appendix to this document contains a Frequently Asked Questions (FAQ) for ALR.

## Background

There is increasing interest in the energy use and power requirements of network interfaces and the effect that network connectivity has on connected devices (particularly on PCs) [2], [3], [6]. A variety of methods could be used to reduce the direct energy use of network interfaces and links with the most immediately promising being Adaptive Link Rate (ALR) for switched Ethernet links. ALR is a proposal that Ethernet links dynamically change their data rates in response to traffic levels. For high levels of traffic, a high link data rate is necessary and must be used. For low levels of traffic, a low data rate is sufficient and should be used. It is well known that most Ethernet links – especially PC to LAN switch links – have low traffic levels most of the time [4], [5]. PCs are often idle and left on, or their use is for non-network applications (e.g., editing or viewing of locally stored documents) [6]. Networked printers, and other networked devices, are also often idle the majority of time that they are powered-on.

Measurements of several NICs, PCs, and a typical LAN switch have shown savings from dropping the link rate on the order of 2 W and up for a 1 Gb/s NIC, and ten times that for a 10 Gb/s link [1]. While more measurements should be collected to refine these savings figures, the savings potential is clearly substantial (on the order of at least $100 million/year for the U.S). An estimate for 2000 [6] put U.S. commercial-sector network equipment consumption at about $500 million/year; including residential products and the stock increases since 2000 would raise this figure greatly. Ethernet NICs may now easily require about $400 million/year of electricity in the US alone.

Since all 1000BASE-T NICs have the ability to operate at 100 Mb/s or 10 Mb/s, the issue is not the capability to operate at different data rates. Rather, the issue is that the existing auto-negotiation mechanism is too slow to be acceptable for most applications. The key challenges are to define a standard method for quickly changing rates and create policies to change rates without adversely affecting network utility.

## Initial Work

Initial work has focused on characterizing network traffic on Ethernet links (over a broad range of usage environments) and building simulation models of ALR [1]. Network characterizations by many researchers (including [2], [4], and [5]) show that average link utilization is low. Usage of Ethernet links is largely bursty – that is, most of the time links are idle (or very close to idle). Thus, there is clearly "head room" for reducing the link data rate for much of the time with little or no impact to network operation. The amount of energy that can be saved is a function of the time to switch link date rates (this time is a sum of the handshaking and link re-synchronization time). Simulation results [1] show that even with link data rate switching time of 10s of milliseconds, substantial energy savings are possible.

## Assumptions

Some reasonable assumptions are:

- Either end of a link (e.g., a PC or LAN switch port) could initiate a request to change the link data rate.
- Only existing link data rates would be considered (i.e., 10, 100, 1000 and 10000 Mb/s). For 1 Gb/s NICs, 10 and 100 Mb/s data rate would be achievable. For 10 Gb/s NICs, possibly only a 1 Gb/s lower data rate may be achievable.
- Labeling programs such as Energy Star could require use of ALR technology to ensure significant market penetration.

## Outstanding Questions

Some open questions are:

- What information about ALR capability should be exchanged during link set-up?

- How fast must the data rate switching time be for ALR to be most useful?
- What should the control mechanism be? Possible mechanisms are auto-negotiation and MAC frame handshake.
- What should the policy for use of the ALR mechanism be? Possible policies include reactive (e.g., based on buffer threshold and measured link utilization) and predictive (e.g., based on previous link utilization).
- How much can be similar – in terms of ALR mechanisms and policies – between 1 Gb/s copper and 10 Gb/s fiber Ethernet?
- What are the market considerations such as design and implementation cost, and inducements such as Energy Star specification for use of ALR?

### Data needs

More data is needed to substantiate and guide the design of ALR. This includes additional traffic characterization on Ethernet links and more measurements of power consumption of NICs at different data rates**.**

### Possibilities to use existing hardware

It may be that enough savings could be obtained by utilizing the existing auto-negotiation mechanism in existing NICs to merit pursuing this in parallel to a more comprehensive ALR solution. A decrease in link data rate is already implemented by some NICs used in notebook computers. These NICs drop the link data rate from 1 Gb/s to 100 or 10 Mb/s when the notebook computer enters a low-power sleep state (e.g., Microsoft Windows standby). Work focused on the use of existing hardware is unlikely to affect the content of any resulting standard, so this topic can be pursued separately.

### Status

Christensen and Nordman [1] presented the ALR concept in a tutorial at the July 2005 IEEE 802.3 plenary meeting held in San Francisco. Several companies expressed interest in exploring the concept further. The March 2006 plenary seems to be an appropriate timeframe to propose a study group for ALR. Companies that have contributed to discussions to date on this include Broadcom, Cisco, Intel and Force10 Networks.

### References


[1] K. Christensen and B. Nordman, "Reducing the Energy Consumption of Networked Devices", IEEE 802.3 tutorial, July 19, 2005 (San Francisco). URL: http://www.csee.usf.edu/~christen/energy/ieee_tutorial.pdf.

[2] C. Gunaratne, K. Christensen, and B. Nordman, "Managing Energy Consumption Costs in Desktop PCs and LAN Switches with Proxying, Split TCP Connections, and Scaling of Link Speed," *International Journal of Network Management*, Vol. 15, No. 5, pp. 297-310, September/October 2005.

[3] M. Gupta and S. Singh, "Greening of the Internet", *Proceedings of ACM SIGCOMM*, pp. 19-26, August 2003.

[4] M. Gupta, S. Grover, and S. Singh, "A Feasibility Study for Power Management in LAN Switches", *Proceedings of the 12th IEEE International Conference on Network Protocols*, pp. 361-371, October 2004.

[5] A. Odlyzko, "Data Networks are Lightly Utilized, and Will Stay That Way", *Review of Network Economics*, Vol. 2, No. 3, pp. 210-237, September 2003.

[6] K. Roth, F. Goldstein, and J. Kleinman, "Energy Consumption by Office and Telecommunications Equipment in Commercial Buildings; Volume I: Energy Consumption Baseline", Arthur D. Little Reference No. 72895-00, January 2002.


### Contact

The primary contact for this pre-proposal is:


Mike Bennett
Lawrence Berkeley National Laboratory
MJBennett@lbl.gov
510 486 7913

## Appendix A – Frequently Asked Questions (FAQ) for Adaptive Link Rate

### 1) What is ALR?

Adaptive Link Rate (ALR) is a proposal that Ethernet links dynamically change their data rates in response to traffic levels. For high levels of traffic, a high link data rate is necessary and must be used. For low levels of traffic, a low data rate can be sufficient and should be used. It is well known that most Ethernet links – especially PC to LAN switch links – have very low traffic levels most of the time. ALR could apply to 10 Gb/s in the future or to 1 Gb/s currently. ALR does not propose that new data rates be invented, but rather that existing 10, 100, 1000, and 10000 Mb/s data rates be used.

### 2) What is the need for ALR?

The need for ALR is part of a growing need to reduce the energy costs associated with operating residential and commercial computing equipment. ALR can be a first step to address direct energy use of network links and equipment. ALR could be incorporated into future EPA Energy Star guidelines.

### 3) What is the primary application for ALR?

The primary application for ALR is expected to be for the Ethernet link from the PC to first-level LAN switch. PC links are idle (see question #16) the great majority of the time. ALR is not intended for the data center where link utilization is often high and requirements for very low latencies (and high throughputs) are far greater than for PC links. The number of relevant NICs that could use ALR is well over 100 million in the US alone.

### 4) What is the expected operating cost savings achievable from ALR?

For the individual user, ALR could achieve energy savings of several dollars per year (more than recovering the expected additional cost of adding ALR to a NIC). On a larger scale, ALR has the potential of saving hundreds of millions of dollars per year in energy costs in the USA alone. Companies with thousands of PCs would see a measurable reduction in their yearly energy bill.

### 5) Is ALR technically feasible for 1 Gb/s?

The answer to this question can best be determined as part of a multi-company study work though we believe that ALR is easily feasible for 1 Gb/s. Evidence for this is that existing NICs already have the capability to change their link data rate when a notebook enters a low-power sleep state.

### 6) Is ALR technically feasible for 10 Gb/s on copper medium?

We believe that ALR is readily feasible for 1 Gb/s. ALR for 10 Gb/s is a more challenging technical problem that we believe can be solved. The answer to this question can best be determined as part of a multi-company study.

### 7) How is a link data rate change initiated?

A link data rate change could be initiated by a handshake mechanism (e.g., via newly defined MAC frames). Use of existing auto-negotiation mechanisms can/should also be investigated. When to initiate a link data rate change is a question of policy. Queue thresholds and/or utilization measurements can be used to determine when to increase and decrease data rate such that packet delay is not perceivably affected.

### 8) Can existing auto-negotiation be used to change data rates?

The existing auto-negotiation scheme in 802.3 takes 100's of milliseconds to change data rates at 1 Gb/s link data rate. This is too slow for ALR to be most effective. See also question #10.

### 9) Do both ends of a link have to be ALR compatible?

Yes. Both the NIC in the PC and the port in the switch must support ALR. Auto-negotiation would need to be used to detect and agree on the use of ALR between two ends of a link.

### 10) Can there be different levels of implementation of ALR?

Yes. This is based on time scales of interest. For fast switching (e.g., 1 millisecond) a new handshake mechanism is likely needed. For slow switching (say, 100s of milliseconds to few seconds) one could use existing mechanisms and this is then primarily a "software problem".

### 11) Will ALR result in a user-perceptible performance impact?

No. A user should not be able to perceive ALR if switching time is less than (roughly) 1 millisecond and good policies for switching data rate are established. A goal of ALR must be to have no user performance impact.

### 12) Does ALR have to be a standard?

For fast switching of data rates, standardization of new mechanisms would ensure compatibility and interoperability at all levels, and greatly increase the chances that the installed functionality is actually used.

### 13) Is there any precedent or "competition" for ALR in other link layer technologies?

ADSL2+ supports multiple power states and data rates. ADSL2+ is a last-mile Internet access technology that is standardized by the ITU. A 2002 White Paper by Aware, Inc. describes power management in ADSL2+ (http://www.dslprime.com/a/adsl21.pdf).

### 14) What are the key reference documents for ALR?

The key reference for ALR is the pre-proposal document: *Adaptive Link Rate (ALR) for Full-Duplex Ethernet Pre-Proposal for an IEEE 802.3 Study Group* (http://www.csee.usf.edu/~christen/energy/alr_proposal.pdf).

### 15) Why support an activity to standardize ALR?

ALR offers an opportunity for product differentiation – there is a market for low-power and low-energy devices. ALR also offers an opportunity to improve energy efficiency in a studied manner, rather than possibly being told how to do it by regulatory entities. ALR could conceivably be part of future Energy Star and similar guidelines either in the US and/or abroad.

### 16) What makes a link relatively "idle"?

For purposes of ALR, an idle link in not one with no data being transmitted. Rather, it is one in which the amount of data being transmitted is very low compared to peak capacity. PC to switch links are typically very lightly utilized over the long run as they are primarily intended for high-speed burst capability (e.g., occasional large file transfers). Studies have shown that Ethernet links are typically utilized in the range of a few percent.

### 17) Who are the key contacts for ALR?

The primary contact for ALR in its pre-proposal stage is:

Mike Bennett
Lawrence Berkeley National Laboratory
MJBennett@lbl.gov
(510) 486-7913