

Internet Research Task Force
Internet-Draft
Intended status: Informational
Expires: July 19, 2011

W. Tavernier, Ed.
Ghent University - IBBT
D. Papadimitriou
Alcatel-Lucent Bell
D. Colle
Ghent University - IBBT
January 15, 2011

Learning Capable Communication Network (LCCN) problem statement
draft-tavernier-irtf-lccn-problem-statement-01

Abstract

Operational procedures and protocols of today's communication networks typically use explicitly defined mechanisms and representations to reach the goals associated to their design. This practice results into numerous protocols having a restricted space for (self-)adaptability, flexibility, and sensitivity respective to their network context (e.g. network traffic conditions, failure conditions, etc). On the other hand, a wide spectrum of learning and optimization techniques is available such that networks could learn and optimize their behavior in the running context. This document describes the opportunities and challenges for a Learning Capable Communication Network (LCCN).

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on July 19, 2011.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

Internet-Draft

LCCN problem statement

January 2011

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Learning opportunities	4
2.1.	Availability of network data and statistics	4
2.2.	Availability of processing capacity	5
3.	The learning process	5
4.	Architectural implications	7
4.1.	From a pre-defined open-loop control towards a self-adaptive closed-loop control	7
4.2.	The integration of learning capability	9
4.3.	Coexistence with current networking protocols, mechanisms and practices	10
4.4.	Complexity/control vs. performance/labour trade-off measurability	10
5.	Applicability	11
5.1.	Functional domains	11
5.2.	Scope with respect to the hourglass model	11
5.3.	Existing work	12
6.	Research directions	13
6.1.	Relation to existing research domains	13
6.2.	Experimental research objectives	14
7.	IANA Considerations	14
8.	Security Considerations	14
9.	Conclusions	15
10.	Acknowledgements	15
11.	Informative references	15
	Authors' Addresses	16

1. Introduction

As currently instantiated, the Internet hour-glass model drives a top-down approach. Current communication networks typically operate with an explicit internal representation of themselves, their network knowledge, and their global goals. Routers follow explicitly (pre-)defined behavior, persistently decide and uniformly execute. Global Internet behavior is evaluated and configuration is when the evaluation indicates that the networking systems are not accomplishing what they were intended to, or when better functionality or performance is expected.

In several Internet areas, this operational model shows its limits. Inter-domain routing protocols such as BGP are increasingly impacted by topology and policy dynamics, delaying their convergence due to inherent exploration properties. Network management becomes more and more complex, as networks do not automatically take into account network traffic statistics and other dynamic properties. Several efforts have been undertaken to overcome the increasing number of issues. However, improvement of the routing system to accommodate various scales of challenges in network efficiency, further complicates its operation ([\[I-D.ietf-idr-bgp-issues\]](#)). Further patching the inter-domain routing system and equipment will result into more operational complexity.

In this document, we suggest an alternative (bottom-up) approach to the Internet routing and forwarding system operation. Compared to current routed networks requiring explicit specification of expected behavior, self-adaptive systems could dynamically modify or adjust their behavior to varying network conditions in order to tune their operation, optimize their overall performance and even add functionalities through closed-loop adaptive control.

We see three main drivers for the development of Learning Capable Communicatio Networks (LCCN): i) the availability of network-related data, ii) the wide range of possible learning paradigms that can be

borrowed from domains such as Artificial Intelligence (AI), machine learning, and bio-inspired learning, and iii) the increased CPU capacity available at both forwarding and control plane level, allowing for background monitoring, learning and optimization in routers.

The structure of this document is as follows. In [Section 2](#), we describe the opportunities for communication networks to learn how to improve their performance. The next section ([Section 3](#)) gives a more formal but broad definition of the concept of learning. [Section 4](#) provides a first set of architectural implications of adding learning capability to communication networks. The applicability domain of

LCCNs is covered in [Section 5](#), and possible research directions are described in [Section 6](#). Concluding remarks and future work are indicated in [Section 9](#).

[2](#). Learning opportunities

[2.1](#). Availability of network data and statistics

Hosts communicate by sending packets between each other via transit network nodes. As such, a communication network is loaded with packets corresponding to network traffic flows between given network source and destination nodes. Many techniques exist to gather statistics about the resulting traffic flows crossing routers.

- o Online statistical counters measure properties of transiting traffic in a router using counters, for example the number of packets per destination prefix or used packet size distribution curves
- o Traffic sampling: instead of counting certain traffic characteristics, unmodified traffic is captured for some time interval. This sample is then used to derive certain characteristics, using e.g. the setting proposed in [[Estan04](#)]) by means of sample-and-hold technique.

Unfortunately, the resulting statistical data is rarely used to directly improve the routing and/or forwarding decision of network nodes (referring to the active self-adaptive closed control loop in

[Section 4.1](#)). However, it is clear that network operation could benefit from taking these statistics automatically into account to allow for traffic spreading and network load balancing, ordering of prefix updates in traffic-informed re-routing decisions, and so on. To a lesser extent (since the routing system is deterministically adaptive to topological and/or policy changes), this observation also applies to routing information exchanges.

Not only the statistics of network traffic are valuable but also the behavioral aspects of the network itself possibly contain usable information for increasing the performance of the network. Statistics about node or link failures can help network recovery mechanisms to fine tune their operation based on the specific statistical context of the running network. Convergence behavior of routing protocols in the specific running context can be monitored such as to reduce the time of transient loops. In brief, the specific running conditions of communication networks possibly hide (statistical) information, which are currently (largely) unused by current Internet protocols; nevertheless, providing an opportunity to

better analyze the behavior of the network behavior depending on the context it is running within.

[2.2](#). Availability of processing capacity


The possibility of maintaining network statistics is not only dependent on the network conditions and environment themselves, but also on the physical feasibility of monitoring and storing them over longer periods.

Supported by Moore's law, we observe that processing power is increasing over last years, either in pure clock frequency of CPU, or in the occurrence of combinations of multiple CPU's on one chip. In combination with the high increase in line card speeds (up to 100 Gbps), the possibility of capturing useful network statistics in background seems within reach.

[3](#). The learning process

Many research fields study the concept of learning from various view points. In the context of LCCNs, learning algorithms correspond to

the (broad) class of algorithms that discover the relationship between system variables (i.e. input, output and hidden variables) from data samples of its environment (obtained by means of measurement/monitoring). More formally, the learning process consists of the following steps (see Figure 1).



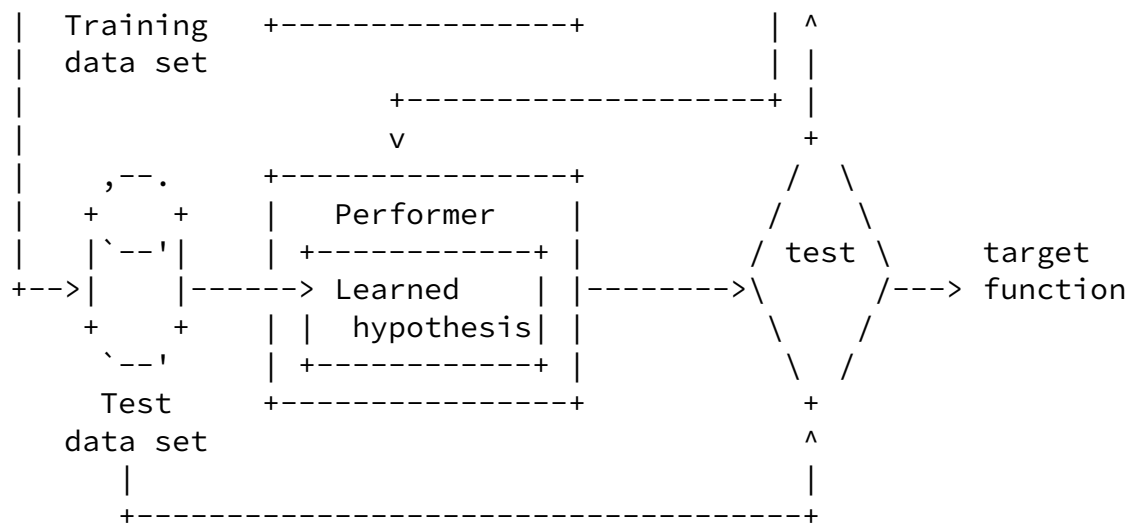


Figure 1

- o Step 0: Choose training and test data sets associated to a given (sequence of) event(s) observed in the system's running environment
- o Step 1: Training (learner): learn an hypothesis h (model), function of the input (training data set) that approximates at best output y (symbolic = classification, numeric = regression). Knowledge: use prior "knowledge" stored in Knowledge Information Base (KIB) to learn h
- o Step 2: Testing (performer): evaluate learned model using test data set

4. Architectural implications

The control of dynamic systems such as communications networks and routers in particular, can be explained as an interactive cycle referred to as the control loop. The coming sections explain the difference of existing communications networks and routers, with the control loop of LCCNs.

4.1. From a pre-defined open-loop control towards a self-adaptive closed-loop control

The configuration and operation of existing communication networks typically consist of a set of components and algorithms acting in a relatively small space of states, transitions and optimization steps. Let's take as example routers: they distribute topology and/or distance information from which they compute (e.g. shortest) routing paths. Using this information, they derive entries looked up to forward packets based on incoming packets' destination address. When a topological or distance change occurs, routing updates are timely disseminated in the network such that each router achieves a coherent full view of the new network topology and/or distances and can re-compute new routing paths taking into account this new state of the network. While these procedures might seem effective at first sight, they are mostly pre-determined and inflexible with respect to the environment they are running in.

Indeed, routers are agnostic to traffic characteristics and to statistics of network failures. This situation occurs because these techniques have been developed in the early days of packet communication networks. At that time, computational and memory resources were scarce, and the resulting techniques needed to act sparingly with the available resources. Moreover, most of these techniques aim to automate manual procedures used to configure or operate communication networks. As such, routers forward packets based on their destination address by applying pre-determined decision rules and execution procedures.

While many engineering disciplines, such as the automotive or bio-industry, have adopted learning techniques to improve the performance of their operational control loops, in computer networking, their application has been restricted mainly to passive applications leading to open-loop control procedures. Examples of such applications are: time series models to analyze and predict network traffic data, anomaly detection techniques to check networks for strange events, or statistical models which try to detect Shared Risk Link Groups (SRLG). Most of the applications of learning techniques are used as interesting side information in the context of network operation. They help network managers to understand and predict the

analyzed and upon which models can be learned, iii) infer rules/decisions from the performed/learned analysis such that the learned model can influence the operation of the network and iv) an execution phase.

[4.2.](#) The integration of learning capability

While it is premature (and part of the research work) to detail the implications on the Internet architecture, the design of a control system incorporating learning capability would benefit from the following design principles.

- o **Adaptability:** modular instead of relying on unified and monolithic approach in order to ensure gradual development (e.g. access vs core router)
- o **Segmentability:** rely on relative local view rather than a network global view in order to ensure scalability, robustness, and resiliency
- o **Sizeability:** inherits distributed properties and capabilities of routing system (e.g. intra- vs inter-domain) in order to ensure organic deployment --instead of a uniform and ubiquitous plane construction

Taking these principles into account, the resulting architecture should specify: i) expected behavior of the self-adaptive closed-loop process, ii) its components, and iii) the interfaces with existing routers' components and between learning-capable routers of a network (both intra- and inter-domain). The resulting closed-loop adaptive control includes a learning component that is either an upfront step or an online process, a feedback phase, and interactions with router/network control.

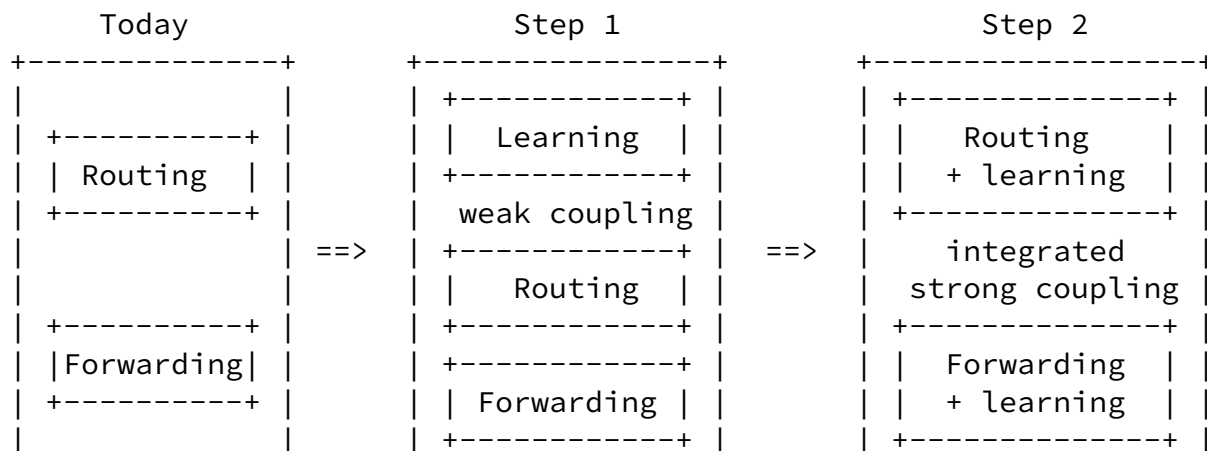


Figure 3

Including learning capabilities into current Internet router architectures can follow a phased approach. Internet routers typically consist of two functional components: i) a forwarding component which takes care of processing and forwarding packets according to pre-configured forwarding tables, and ii) a routing component which takes care of distributing topology/distance information, computing (shortest) routing paths using this information, and storing resulting entries into routing tables. Forwarding table entries are subsequently derived from routing table entries. As a first integration step, a new functional component comprising learning capability could be included. The new component would then be weakly coupled to the existing forwarding and routing components. This implies that the routing and/or forwarding component can be enhanced by of the learning component. These functionalities could be called via pre-defined interfaces between the components. While this is an overlaid but modular build-up of a router, integration of learning capability can go one step further. Indeed, in a next phase, instead of a separate learning component, the learning functionality could be tightly integrated into the routing and forwarding components themselves. This implies that the routing and forwarding processes themselves comprise a learning cycle (a self-adaptive closed-loop control). It is clear that both the phasing and the detailed specification of the architecture is an important challenge in the design of LCCNs.

[4.3.](#) Coexistence with current networking protocols, mechanisms and practices

The roll-out of learning capability into communication networks preferably allows to coexist with well-functioning existing network protocols and mechanisms. This means that LCCNs should not enforce the networking environment to use them or adapt to them, even though they could improve the resulting network performance or solve a number of issues. As such, a transition path towards communication networks including more learning-capability becomes possible without introducing abrupt transition paths.

[4.4.](#) Complexity/control vs. performance/labour trade-off measurability

The implications of using LCCNs should be addressed by determining the relative complexity and understandability they introduce. This does not mean that complex (or black box) LCCN approaches are out of scope, it implies that the additional complexity and understandability resulting from the introduction of this control component should be measurable or can be at least characterized. Measurability (and associated metrics) is an integral part of the

Tavernier, et al.

Expires July 19, 2011

[Page 10]

Internet-Draft

LCCN problem statement

January 2011

investigation work. The assesment should allow users of LCCNs to decide on the level of control vs. performance they are willing to give up/gain. In this context, the analogy can be made with manual configuration of static routing tables vs. running automated shortest path protocols. It is clear that a certain level of control is given up by allowing automated routing protocols to configure routing tables. However, the resulting configuration is verifiable (by routing table inspection), the used algorithm (e.g. Dijkstra shortest path calculation) is known, and the resulting reduction in manual intervention is clear. On the other hand, the more laborous manual configuration allows for setups that are sometimes more tuned to specific traffic patterns (e.g. avoiding bottlenecks) than shortest path-protocols. In most scenario's, the trade-off is clear for network operators: larger networks typically use automated routing protocols for the population of routing tables, whereas smaller, specialised network setups sometimes result into manually configured routing tables. A similar type of trade-off is desired for LCCNs.

[5.](#) Applicability

[5.1.](#) Functional domains

The incorporation of learning component within the router architecture aims to i) enhance Internet functionality in order to cope with known operational challenges such as manageability, and diagnosability, ii) address new challenges such as security and accountability, and iii) improve its performance (in terms of e.g. scalability and availability) by adapting forwarding and routing system decisions. In this context of network quality, we can think of the automated inclusion of network traffic knowledge into the


configuration of routes and resulting forwarding tables.

[5.2.](#) Scope with respect to the hourglass model

Even if learning paradigms can be applied at all levels of the hourglass model, LCCN-related research focuses on the (largest) lower half of the hourglass model ("everything over IP, and IP over everything"). As depicted in Figure 4, the goal of LCCN research is to apply learning capabilities from the transport layer up to the physical layer (including thus also the network and datalink layers).

Whereas learning capability is typically being used at the application layer already, for example by banking applications, large-scale websites such as Amazon or Google, except for TCP, the real networking machinery that is running below is still relying on low-information processes with very limited learning capabilities.

The incorporation of a learning component within wired and wireless communication network systems aims to improve both their operation and performance from the physical network layer up to the TCP/IP layer. TCP can be qualified as an exception in the sense that it incorporates some of the procedures involved in learning processes. Indeed, its transmission window size is adaptively changed during the communication between network end points such as to maximize throughput while keeping the resulting congestion as low as possible. However, it mainly concerns end-to-end learning while learning within the network itself provides additional value (as shown by the work performed e.g. in [[Tavernier10](#)]).



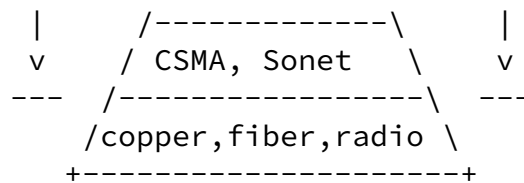


Figure 4

5.3. Existing work

Although the penetration of learning capability in current network protocols is rather low, in several domains some studies have been conducted on the possible value of introducing learning capability or intelligence into the networking mechanisms.

Learning systems have been succesful applied for example in cognitive radio networks and optical networks. Using such systems, wireless network nodes adaptively change their transmission and/or reception parameters to communicate efficiently avoiding interference with other networks and nodes. The adaptive change of these parameters is based on the active monitoring of several factors in the external and internal radio environment, such as radio frequency spectrum, user

behavior and network state. More information about cognitive radio networks can be found in [[Haykin2007](#)].

[Riziotis07] made a survey on the succesful application of computational intelligence techniques in the domain of photonics and optical networks. Tens of studies are cited on the succesfull application of optimization and learning techniques in the design and operation of optical networks. For example in [[Goncalves04](#)], agents make use of Artificial Neural Networks for monitoring an optical link of a network and predicting anomalous situations so that pro-active measures can be taken before faults occur. This technique showed to be significantly less costly compared to providing 1+1 protection on DWDM links.

The insight resulting of bringing together conducted research on learning capability in networked environments can result into a common base of and architecture to further investigate and deploy learning capability into new networked contexts. Such a bottom-up

approach can be valuable as it can give us lessons in common challenges, and ways to tackle them in order to reach a higher level of adoption of LCCNs.

6. Research directions

6.1. Relation to existing research domains

Learning opportunities in communication networks have characteristics that are typical well-suited for research techniques borrowing from (machine) learning, robotics, AI, computational biology, etc.

- o Difficult to explicitly characterize: events cannot be well characterized even when examples are available (inherent complexity in characterizing an event)
- o Correlation: hidden correlations and trends between events within large amounts of associated data
- o Dynamicity: changing conditions over time (in particular, for routing system but also variability of traffic, user expectations and behaviors)
- o Quantity: amount of available data is too large for handling by manual intervention
- o Evolutive: new events are constantly detected/discovered

6.2. Experimental research objectives

Experimental research is a primary goal of the activities to be conducted. The following objectives would be targeted:

- o The production of various studies is stimulated and should enable evaluation of performance and functional improvement resulting from the exploitation of various learning paradigms. A common understanding of these paradigms and their associated capabilities could complement this first step. The resulting bottom-up approach allows to combine insights of several use cases involving

learning in networks to find the common base and best architecture/practices in the development of LCCNs.

- o As different distribution models can be considered for what concerns the distribution of the learning processes (taking into account the various objectives but also constraints resulting from network partition), determining which model best fit Internet evolution is a specific target of this research activity.
- o Iterative cycles of experimentation shall allow to determine suitability of the resulting architecture as well as to determine practical feasibility, applicability and deployability of the concept on a large scale. Documentation of appropriate use cases/ scenarios would complement this work item.

7. IANA Considerations

This memo includes no request to IANA.

8. Security Considerations

It is desirable that LCCNs provide visibility on the possible mis-use of their learning capability. As such, the assesment of their attractiveness for deployability becomes easier.

Beside the research objectives detailed here above, security mechanisms for "communication channels" between learning components and "learning components" themselves shall be considered comprising among others message authentication but also means to prevent e.g. man-in-the-middle and DDoS attacks. In the LCCN context, the question becomes what is sufficient for protecting the Internet against such attacks. Is it sufficient to provide secure communication channels as well as adequate authentication and verification/validation mechanisms for the information exchanged over these channels, or can we rely on learning to determine protecting

decisions systems should take to ensure their own defense against such attacks ? These are security topics that can be further investigated in the context of LCCN research.

9. Conclusions

Current communication networks fail to use network-related statistics which could be valuable to improve their performance. In addition, current networks fail to provide solutions to challenging issues, because they become too complex to operate and manage by manual/open loop procedures. A learning-capable communication network (LCCN) includes a learning component which learns based on the network environment statistics and adapts and optimizes its behavior upon this. This gives new possibilities to improve network efficiency in several domains including network recoverability, accountability, security, scalability, and so on. The challenge (and next steps) of LCCNs lies into: i) developing self-adaptive closed-loop control system relying on learning capability, ii) building and applying it to various network mechanisms and iii) verifying the resulting prototypes in experimental environments.

10. Acknowledgements

This work is supported by the European Commission (EC) Seventh Framework Programme (FP7) ECODE project (Grant No.223936).

11. Informative references

[AI-modern]

Russell, S., "Artificial Intelligence: A Modern Approach", 2003.

[Estan04] Estan, C., "Building a better NetFlow", october 2004.

[Goncalves04]

Goncalves, C., "Applying artificial neural networks for fault prediction in optical network links", december 2007.

[Haykin2007]

Haykin, S., "Cognitive radio: brain-empowered wireless communications", february 2007.

[I-D.ietf-idr-bgp-issues]

Lange, A., "Issues in Revising BGP-4 ([RFC1771](#) to [RFC4271](#))", [draft-ietf-idr-bgp-issues-03](#) (work in

progress), August 2010.

[PRML] Bishop, C., "Pattern Recognition and Machine Learning", october 2003.

[Riziotis07]
Riziotis, C., "Computational intelligence in photonics technology and optical networks: A survey and future perspectives", december 2007.

[Tavernier10]
Tavernier, W., "Using AR(I)MA-GARCH models for improving the IP routing table update", october 2010.

Authors' Addresses

Wouter Tavernier (editor)
Ghent University - IBBT
Gaston Crommenlaan 8 bus 201
Gent, 9050
Belgium

Phone: +32(0)9 331 49 81
Email: wouter.tavernier@intec.ugent.be

Dimitri Papadimitriou
Alcatel-Lucent Bell
Copernicuslaan 50
Antwerpen, 2018
Belgium

Phone:
Email: dimitri.papadimitriou@alcatel-lucent.com

Didier Colle
Ghent University - IBBT
Gaston Crommenlaan 8 bus 201
Gent, 9050
Belgium

Phone: +32(0)9 331 49 70
Email: didier.colle@intec.ugent.be

