# A Link State Advertisement Protocol for Optical Transparency Islands

Shovan Das, Reza Roshani Tabrizi, Paolo Monti, Marco Tacca, and Andrea Fumagalli
The Open Networking Advanced Research (OpNeAR) Lab
Erik Jonsson School of Engineering and Computer Science
The University of Texas at Dallas, Richardson, TX 75080, USA
Email:{skd043000, rxr058100, paolo, mtacca, andreaf}@utdallas.edu

***Abstract*— Plug and play optical (PPO) nodes enable fast, reconfigurable, and flexible ad hoc deployment of optical networks. Once plugged, the PPO nodes provide all-optical circuits between client nodes to alleviate the electronic processing bottleneck of high speed networks. To offer these wavelength routing functionalities to client nodes the PPO nodes must self-adjust to possible changes of the optical physical topology and fiber propagation characteristics. To discover such changes the PPO nodes must make use of a link state advertisement (LSA) protocol that is scalable in the number of plugged PPO nodes.**

**This paper describes a scalable LSA protocol with constrained message flooding to match the limited propagation reach of the optical signal, i.e., the PPO node transparency island (TI). Scalability is thus achieved naturally by limiting the link state advertisement scope to only those PPO nodes that need to receive the link updates. As discussed in the paper, the proposed protocol appears to be a viable solution when the TI size is relatively small, e.g., in optical networks without signal regeneration.**

## I. Introduction

Today technology in fiber optics communications has the potential to facilitate end-to-end data exchange in the multi-gigabit transmission range. Optical circuits, or lightpaths, can be established in the network to provide transparent channels between end node pairs. Electronic processing of transmitted data along a lightpath is not required, thus avoiding a potential electronic processing bottleneck when high transmission rates are required. This is often referred to as optical transparency [1].

In some areas, the deployment of optical networks may be facilitated by the use of self-configurable plug and play optical (PPO) nodes [2], [3]. Similarly to wireless ad hoc solutions, PPO networking could allow fast, reconfigurable and flexible deployment of optical resources to best fit application requirements, especially in the metro and access area. It could also simplify the complex procedures for the design, installation, and maintenance of today's optical networks, as no-human intervention would be required to perform these tasks.

The key components of the PPO node are [4]: (*i*) an on-board miniature optical transmission laboratory, or mini-lab, (*ii*) real-time transmission models, and (*iii*) a service channel interface for network management and control. The real-time transmission models are used at the PPO node to process the measurements produced by the on-board mini-lab [5]. Their objective is to estimate the maximum transmission rate and span that are permitted when creating lightpaths from the PPO node to other PPO nodes available in the network. Lightpath rate and span are both limited by the physical constraints of the optical medium. As a consequence, the PPO node is able to establish lightpaths to reach only a subset of other PPO nodes, and these nodes are said to belong to the PPO node's *transparency island* (TI) [6]. Note that each PPO node has its own TI, and that some TI(s) may overlap. As explained later, the existence of the TI makes it possible to naturally limit the amount of local control information required at the PPO node.

Once connected, a PPO node must cooperate with other already existing PPO nodes, and determine how to use available optical resources, such as fibers and wavelengths, to provide requested lightpaths to the connected clients, e.g., end users, routers. To do so, the PPO node must 1) discover resources and detect changes in the optical data plane, and 2) solve the routing and wavelength assignment (RWA) problem [1]. Finding a solution to the former problem and assessing the solution scalability is the focus of this paper.

Discovery of resources and detection of changes in networking are often accomplished using link state advertisement (LSA) protocols. A well know example is the open shortest path first (OSPF) protocol [7], [8], whereby link state information entries are flooded across the network. To provide a scalable solution, flooding of LSA entries is constrained within areas, which are defined when the network is designed. These pre-defined areas do not change in time, and each network node belongs to one area only[1], i.e., areas do not overlap. Therefore, LSA protocols, based on preassigned areas or subnetworks, do not fit the ad hoc networking nature of PPO nodes, whereby nodes can be added and removed many times and optical resources may become available/unavailable during the network lifetime. For this reason, the TI-LSA protocol is introduced.

The TI-LSA protocol is based on the OSPF LSA flooding principle, which is adapted to take advantage of the PPO node's TI. The number of flooded TI-LSA entries is limited by constraining the advertisement within the optical reach of the PPO node, i.e., the boundaries of the PPO node's TI. In other words, each TI-LSA entry reaches the PPO nodes

This research was supported in part by NSF Grant No. ANI-0435393 and the Italian Ministry of University (MIUR) (contract # RBNEO1KNFP).

[1]With the exception of area border gateways.

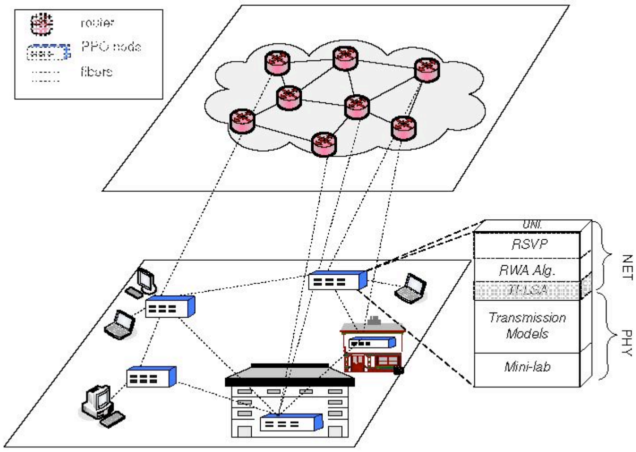


Fig. 1. PPO node enabled network

that require that information, without unnecessarily flooding other PPO nodes that will not make use of that entry. While the TI may resemble the concept of optimal partitioning of domains into OSPF areas [9], [10], a substantial difference between the two is that areas are manually predetermined, while TI(s) are self detected. The TI(s) may change over time as they automatically adapt to the changes of both the optical network topology and the fiber transmission characteristics, e.g., changes in temperature and aging. And so does the constrained flooding of the TI-LSA entries. While delivering to each PPO node all the TI-LSA entries that may be required to intelligently use the available optical resources, the TI-LSA protocol permits to realize networks with a number of PPO nodes that is virtually infinite, thanks to its TI constrained flooding.

A particular challenge in designing the TI-LSA protocol is that TI(s) of distinct PPO nodes may be different and may overlap partially. As explained in Section III-B some substantial protocol changes are then required when compared to conventional single area LSA (SA-LSA) protocols, e.g., the OSPF protocol implemented using a single area. The payoff, as discussed in the paper, is that by taking into account the PPO node's TI the proposed TI-LSA protocol constitutes a scalable solution to the problem of topology discovery and update when the TI size is relatively small. This is indeed the case in optical networking when sophisticated and costly signal regeneration techniques, e.g., 3R [11], are not a viable option. This observation suggests that PPO networking may be well suited in the access and metro area, where both inexpensive equipment on the one hand, and ad hoc deployment on the other may be premium features.

Before describing the TI-LSA protocol, the envisioned PPO node network architecture is defined in the next section.

## II. The PPO Node Network Architecture

Fig. 1 depicts a network with both client (routers) and PPO nodes. The fiber bandwidth is wavelength multiplexed to offer multiple data channels and one dedicated low-speed service channel for management and control signaling. Via a user-network interface (UNI) defined over the service channel, client nodes may request PPO nodes to create lightpaths to form a desired virtual topology. Upon reception of a data lightpath request to connect two client nodes, the PPO node must first determine if it possible to establish that lightpath and meet the transmission rate requirement specified by the client. To perform this task, the PPO node must be aware of the physical topology, fiber transmission impairments (e.g., power loss, polarization mode dispersion) and available data wavelengths. Note that changes of the physical topology may be frequent as PPO nodes are plugged and unplugged dynamically as needed. Once the physical topology, fiber transmission impairments, and available wavelengths are known to the source PPO node, known RWA algorithms can be applied to choose both path and wavelength to establish the requested lightpath. Conventional reservation protocols (e.g., RSVP [12], [13]) may then be used to establish the requested lightpath and allow client nodes to exchange data directly over that lightpath.

The PPO node sub-layers are shown in Fig. 1. First, the PPO node must characterize and monitor the key transmission parameters of the outgoing fiber links connecting to neighboring PPO nodes. This task is performed in cooperation with the neighboring PPO nodes, using the on-board mini-lab [5] and a Hello protocol [7]. The mini-lab measures the fiber key transmission parameters and detects their changes over time. In addition, the Hello protocol detects the loss/repair of the service channel connecting adjacent PPO nodes. For example, a fiber cut is detected by both the mini-lab, as a loss of a data plane link, and the Hello protocol, as a loss of the control plane link.

Changes of the fiber transmission parameters may positively or negatively affect the optical signal quality in the data plane. Based on their impact on the signal quality they are classified as `perf-up`, `link-up`, `link-down`, or `perf-down` events. When their impact on the signal quality cannot be determined, they are classified as `perf-unknown` events. To gain a comprehensive view of the physical topology, fiber measurements and detected fiber link changes are flooded to other PPO nodes in the form of TI-LSA entries. The TI-LSA entries may also be used to flood information about the availability of wavelengths. Each PPO node combines the TI-LSA entries received from the neighboring PPO nodes to build its own link state database. Based on its link state database and by using real-time transmission models [5], the PPO node can swiftly determine whether a requested lightpath can be established.

As already mentioned, in this architecture, the PPO node is only required to deal with lightpaths that can be established while producing acceptable optical signal quality at the receiver. In practical terms, lightpaths that span across too many fibers (without signal regeneration) may not be created as their resulting signal quality does not yield the desired bit error rate. This observation leads to the conclusion that the PPO node's link state database may be limited to represent a subset of the entire physical topology, i.e., its own TI.

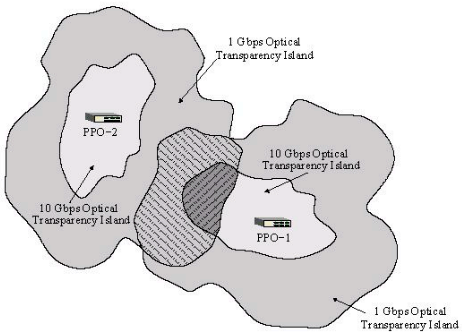


Fig. 2. Optical transparency islands

## III. TI-LSA Protocol

This section provides the formal definitions of TI and the description of the TI-LSA protocol.

### A. *Transparency Island Definition*

Fig. 2 provides a sketch of the optical TI(s) of two PPO nodes. The TI reach is dependent on the employed transmission rate and might vary over time due to changes of both the topology, and the fiber parameter values. Plugging of new PPO nodes and/or deployment of new fibers may enable some PPO nodes to reach remote PPO nodes that were previously unreachable. Alternatively, unplugging of a PPO node and/or a fiber cut may deprive some PPO nodes of optical resources, thus reducing the size of the TI. A formal definition of the TI (at a given time) is given next.

The PPO node network is modeled as a graph $G(\mathcal{N}, \mathcal{A})$, where $\mathcal{N}$ is the set of PPO nodes in the network and $\mathcal{A}$ is the set of uni-directional links connecting the PPO nodes. A PPO node with ID $i$ is denoted as $N_i$ and a link connecting PPO nodes $N_i$ and $N_v$ is denoted as $l_{(i,v)}$. Let $C^{(r)}_{(i,j)}$ be the set of all links and PPO nodes that can be used to establish a simple lightpath[2], operating at transmission rate $r$, connecting PPO node $N_i$ to $N_j$. Set $C^{(r)}_{(i,j)}$ defines a subgraph of $G(\mathcal{N}, \mathcal{A})$. Note that lightpaths containing loops are not allowed. Let $TI^{(r)}_i$ be the transparency island associated with node $N_i$, for lightpaths operating at transmission rate $r$. $TI^{(r)}_i$ is a subgraph of $G$, defined by:

$$TI^{(r)}_i = \cup_{j \in \mathcal{N}} \left( C^{(r)}_{(i,j)} \right). \qquad (1)$$

$TI^{(r)}_i$ contains all PPO nodes and links that can be used to establish a lightpath at transmission rate $r$ that originates at PPO node $N_i$.

### B. *Protocol Description*

This section describes the TI-LSA protocol used to gather link state information at every PPO node, i.e., building and maintaining $TI\_DB^{(r)}_i$ at PPO node $N_i$ for one single rate $r$. For the first version of the TI-LSA protocol, only one data rate is supported. For simplicity, in the remainder of the paper, the index $r$ is not used, e.g., $TI\_DB^{(r)}_i$ is denoted by $TI\_DB_i$. $TI\_DB_i$ is a subgraph of $G(\mathcal{N}, \mathcal{A})$ and it is the transparency island associated with node $N_i$ that is built based on — possibly inaccurate — link state information available at PPO node $N_i$. $TI_i$ is the transparency island of PPO node $N_i$ when $N_i$ has complete and accurate knowledge of the PPO layer topology.

The design of the TI-LSA protocol is based on the following assumptions:

- PPO node $N_i$ is the *owner node* of its outgoing links, i.e., $N_i$ is in charge of advertising the status of its outgoing links,
- the owner node detects the status of its outgoing link in a finite duration of time, i.e., each outgoing link status change is detected by the combined use of a Hello protocol [7] and measurements produced by the on-board mini-lab,
- links $l_{(i,v)}$ and $l_{(v,i)}$ are subject to the same performance changes, e.g., the two links go down or come up at once,
- all transmitted control messages are received without error within a finite duration of time,
- all exchanged control messages are processed within a finite duration of time,
- a lightpath can be established only when the bit error rate performance is satisfactory,
- it is assumed that the wavelength chosen for the lightpath does not affect the lightpath performance, i.e., if a lightpath can be established on a simple path $p \in C^{(r)}_{(i,j)}$ using a wavelength, then it can be also established using any other wavelength along the same simple path $p$,
- if a lightpath can be established on a simple path $p \in C^{(r)}_{(i,j)}$, then any lightpath routed using a sub-paths of $p$ can be established too.

Recall that the objective of the TI-LSA protocol is to inform $N_i$ of any status change that may occur on the links that belong to $TI_i$, i.e., the TI-LSA protocol must guarantee that $TI\_DB_i = TI_i$ within a finite duration of time. This objective is accomplished by transmitting messages that contain link state update(s), i.e., LSU-ENTRY(s). An LSU-ENTRY is a field of the flooded message that contains information regarding one link. For example, an LSU-ENTRY regarding link $l_{(i,v)}$ may contain the fiber power loss and current dispersion profile, the available wavelengths, etc. By the simple fact of receiving an LSU-ENTRY about link $l_{(i,v)}$, PPO node $N_j$ can conclude that $N_i$ and $N_v$ are active and running. Several LSU-ENTRY(s) (local or remote) can be grouped together by the PPO node to form a single advertisement message defined as LSU-PACKET. The PPO node can also remove any unnecessary LSU-ENTRY from a LSU-PACKET to reduce the signaling overhead.

A number of events can trigger PPO node $N_i$ to flood the status of one or more of its outgoing links to all the neighboring PPO nodes. These events are:

[2]It is assumed that if a lightpath can be established, then its performance in terms of bit error rate is satisfactory.

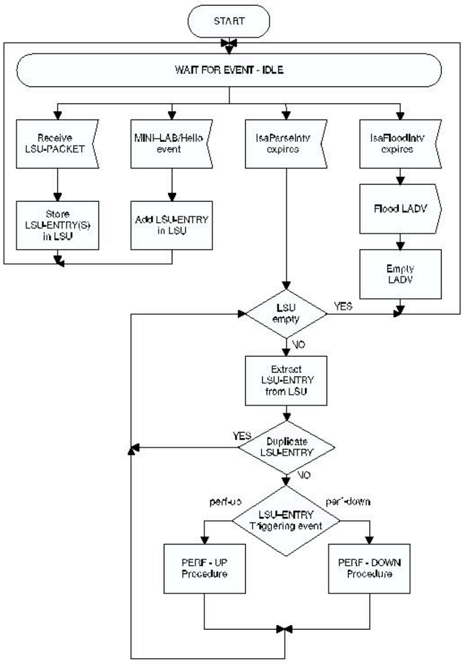


Fig. 3. TI-LSA protocol flowchart

- `perf-up`, `link-up`, `perf-unknown`, i.e, a link may improve its performance, come up or change its performance in an unknown way, respectively,
- `perf-down`, `link-down`, i.e., a link can deteriorate its performance, or go down, respectively.

Events `perf-up`, `link-up`, and `perf-unknown` are grouped together, as they require the TI-LSA protocol to perform the same set of procedures. The same apply to events `perf-down` and `link-down`. Note that the plugging (unplugging) of a PPO node is detected in the form of a set of `link-up` (`link-down`) events.

The following description applies to PPO node $N_i$. The flowchart of the TI-LSA protocol is shown in Fig. 3. Upon reception of an LSU-PACKET at an incoming interface, every LSU-ENTRY in the LSU-PACKET is temporarily stored in set LSU. An LSU-ENTRY can also be created and stored in the LSU set in response to an event generated by either the mini-lab or the Hello message exchange. Before processing the received LSU-ENTRY(s) in LSU, the PPO node waits for a counter (`lsaParseIntv`) to expire. The `lsaParseIntv` counter is used to make it possible to receive multiple LSU-PACKET(s) and to process all their LSU-ENTRY(s) at once.

The first step when processing a received LSU-ENTRY is to verify that the LSU-ENTRY is not a duplicate, i.e., the LSU-ENTRY was not already received. This can be done by assigning a sequence number to the LSU-ENTRY. If the LSU-ENTRY is not a duplicate, the second step is to check which event triggered the flooding, e.g., `perf-up` or `perf-down`, etc. Depending on the triggering event, the LSU-ENTRY is handled following one of two possible procedures (Figs. 4 and 6). The outcome of these procedures is to determine whether or not the LSU-ENTRY has to be included in the local $TI\_DB_i$ and further advertised to neighboring PPO nodes. If the LSU-ENTRY has to be further flooded, it is stored into the link advertisement (LADV) database. When a counter (`lsaFloodIntv`) expires, all LSU-ENTRY(s) in LADV are moved into an LSU-PACKET and flooded at once.

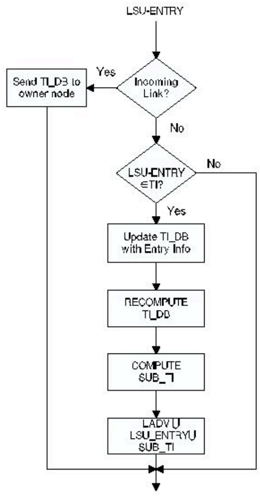


Fig. 4. `perf-up`, `link-up`, and `perf-unknown` flowchart

Fig. 4 shows the flowchart of the procedure used when the LSU-ENTRY is triggered by one of the following events: `perf-up`, `link-up`, and `perf-unknown`. Let the LSU-ENTRY be associated with link $l_{(k,v)}$. First, $N_i$ must determine whether the LSU-ENTRY is associated with an incoming link, i.e., $v = i$. If so, $N_i$ must send its entire transparency island database ($TI\_DB_i$) to the owner of the link, i.e., PPO node $N_k$. Otherwise, if $l_{(k,v)}$ belongs to $TI_i$[3] then both $l_{(k,v)}$ and $N_v$ are added to $TI\_DB_i$. Then, $SUB\_TI^{(i)}_{(k,v)}$ is computed as follows. $SUB\_TI^{(i)}_{(k,v)}$ is a subgraph of $TI\_DB_i$ and it is defined as the set of all links and nodes that may be used to establish lightpaths starting at $N_i$, using $l_{(k,v)}$, and terminating at $N_j$, $\forall N_j \in \mathcal{N}$. The LSU-ENTRY(s) associated with all the links in $SUB\_TI^{(i)}_{(k,v)}$ are then added to LADV for flooding. Note that this step is peculiar of the TI-LSA protocol and it is not required in the conventional OSPF. An example is used next to clarify the importance of this step.

Fig. 5 shows the importance of calculating $SUB\_TI$ and adding all LSU-ENTRY(s) associated with links in $SUB\_TI$ to LADV with an example. In the example, it is assumed that lightpaths cannot exceed 4 hops due to signal quality requirements. Thus a generic link $l_{(u,v)}$ belongs to $TI\_DB_i$ if it is possible to establish at least one lightpath (which does not require more than 4 hops) from $N_i$ to $N_v$ using $l_{(u,v)}$. Assume that links $l_{(c,d)}$ and $l_{(d,c)}$ are up only after time $t_0$.

[3] The actual algorithm used to determine whether the LSU-ENTRY is to be included in $TI\_DB_i$ is outside the scope of this paper.

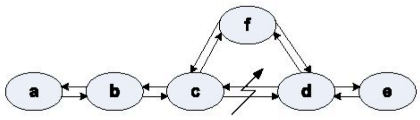


Fig. 5. $SUB\_TI$ example

The TI's before $t_0$ are:

- $TI_a = TI\_DB_a = \{l_{(a,b)}, l_{(b,c)}, l_{(c,f)}, l_{(f,d)}\}$,
- $TI_b = TI\_DB_b = \{l_{(b,a)}, l_{(b,c)}, l_{(c,f)}, l_{(f,d)}, l_{(d,e)}\}$,
- $TI_c = TI\_DB_c = \{l_{(b,a)}, l_{(c,b)}, l_{(c,f)}, l_{(f,d)}, l_{(d,e)}\}$,
- $TI_d = TI\_DB_d = \{l_{(b,a)}, l_{(c,b)}, l_{(f,c)}, l_{(d,f)}, l_{(d,e)}\}$.

At time $t_0$, $l_{(c,d)}$ becomes operational and triggers the following events[4]. When $N_d$ detects the status change of its incoming link $l_{(c,d)}$, it sends an LSU-ENTRY for every link in $TI\_DB_d$ to $N_c$. When $N_c$ receives and processes the LSU-ENTRY(s), the only entry not already present in $TI\_DB_c$ is the one associated with link $l_{(d,f)}$. Therefore $l_{(d,f)}$ is the only link that is added to $TI\_DB_c$. Then $N_c$ computes $SUB\_TI^{(c)}_{(c,d)} = \{l_{(c,d)}, l_{(d,f)}, l_{(d,e)}\}$. For each link in $SUB\_TI^{(c)}_{(c,d)}$ an LSU-ENTRY is stored in LADV and flooded using an LSU-PACKET by $N_c$ to its neighbors. When $N_b$ receives this LSU-PACKET, the only link that must be added to $TI\_DB_b$ is $l_{(d,f)}$. Then $N_b$ computes $SUB\_TI^{(b)}_{(c,d)} = \{l_{(b,c)}, l_{(c,d)}, l_{(d,f)}, l_{(d,e)}\}$. Once again, for each link in $SUB\_TI^{(b)}_{(c,d)}$ an LSU-ENTRY is stored in LADV and flooded using an LSU-PACKET by $N_b$ to its neighbors. When $N_a$ receives the LSU-PACKET, both $l_{(d,f)}$ and $l_{(d,e)}$ are added to $TI\_DB_a$.

Notice now what would happen if a procedure similar to the OSPF flooding mechanism were used in the example instead, i.e., at each PPO node no $SUB\_TI$ is computed and the LSU-ENTRY associated with a link is flooded only if the link is not already present in the PPO node database. When $N_b$ receives the LSU-PACKET from $N_c$, it floods an LSU-PACKET containing information about $l_{(d,f)}$, but not about $l_{(d,e)}$, because the latter is already in $TI\_DB_b$. As a result, $N_a$ does not receive the information about the newly available link $l_{(d,e)}$ and is not able to correctly compute $TI\_DB_a$.

Fig. 6 shows the flowchart of the procedure used when the LSU-ENTRY is triggered by one of the following events: `perf-down` and `link-down`. Let the LSU-ENTRY be associated with link $l_{(k,v)}$. Notice that, whenever there is a `perf-down` and `link-down` type, $TI\_DB_i$ may at most lose some links (and nodes). Therefore, it is not necessary for $N_i$ to compute $SUB\_TI_i$. First, $N_i$ checks whether an entry associated with link $l_{(k,v)}$ is present. If not, no further action is required. Otherwise, $N_i$ recomputes and updates $TI\_DB_i$, and stores an LSU-ENTRY associated with $l_{(k,v)}$ in LADV.

[4]For the sake of simplicity, events that are not relevant to the example are ignored.

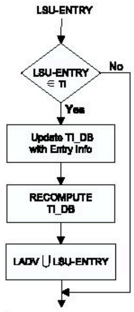


Fig. 6. `perf-down`, `link-down` flowchart

## IV. Emulation Results

The TI-LSA protocol was implemented and tested on the $\Omega$-$E$ network emulator. The $\Omega$-$E$ is a cluster of Linux based PC's, which can be configured to emulate the exchange of control messages in a desired network topology via netfilter/iptables functionalities [14]. Each PC hosts multiple processes, each one emulating one PPO node. Each outgoing (incoming) fiber link of a PPO node is realized as a virtual ethernet interface in the host PC.

The following assumptions are made when running the emulation. For simplicity, every link in the topology is assumed to have the same length and made of the same fiber type. Each fiber carries one wavelength. It is assumed that the determinant factor for the lightpath performance is the optical power budget. The transmission power level and receiver sensitivity are assumed to be the same at all nodes. With these assumptions the lightpath maximum span is the same for all PPO node pairs and can be simply measured in terms of maximum number of hops (links). The maximum number of hops per lightpath used in the emulation is either 2 or 5.

Hello messages are sent every 10 s. If three consecutive Hello messages are lost the corresponding link/node is considered to be down. The flooding interval (`lsaFloodIntv` timer defined in Section III-B) at each PPO node is set to be equal to 60 s. The starting times for both the Hello messages to be sent and the flooding mechanism to begin are randomly chosen in the interval $[0, \texttt{lsaFloodIntv}]$. The random choice avoids synchronization among PPO nodes.

Performance results are collected for both the TI-LSA protocol and the OSPF like single area LSA protocol (SA-LSA). With SA-LSA all the nodes in the network are set to be in the same area/domain. To asses the scalability of the two protocols a number of experiments is performed using network topologies that have the same nodal degree (2.5) and different number of PPO nodes. Each network topology is randomly generated using the Doer and Leslie's formula [15]. The total number of experiments is chosen so that the presented average values have a confidence interval of 10% or better at 90% confidence level.

Fig. 7 reports the average number of LSU-ENTRY(s) processed at each PPO node during a complete start-up phase (i.e., all PPO nodes are turned on at once) till convergence

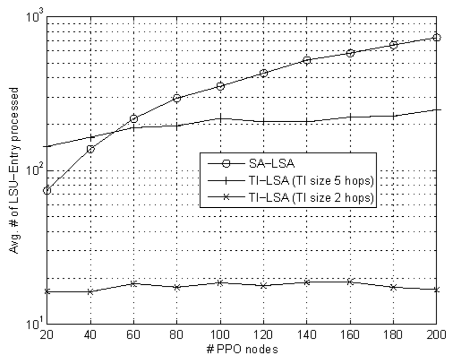


Fig. 7. Average number of LSU-ENTRY(s) processed at each PPO node vs. number of PPO nodes in the network

is reached and every PPO node has completely built its own TI. The average number of LSU-ENTRY(s) is shown as a function of number of PPO nodes in the network. With SA-LSA, the addition of a new PPO node/link is advertised to all the other PPO nodes. As a consequence the average number of processed LSU-ENTRY(s) increases with the number of PPO nodes in the network. Conversely, the TI-LSA protocol is able to scale with the network size. By constraining the advertisement within the boundaries of each PPO node's TI, TI-LSA is able to contain the average number of processed LSU-ENTRY(s) to almost a constant value. As intuition suggests larger TI sizes correspond to an increased average number of processed LSU-ENTRY(s).

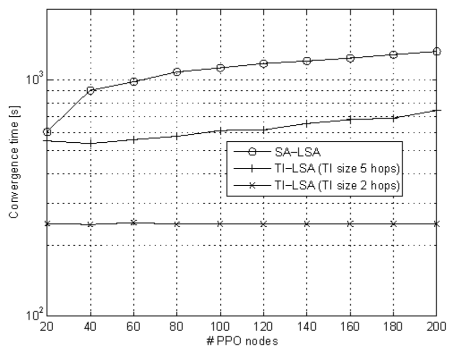


Fig. 8. Convergence time vs. number of PPO nodes in the network

Fig. 8 reports the convergence time of both of the protocols as a function of the number of PPO nodes in the network. This is the time required to transmit all the LSU-ENTRY(s) during the network start-up phase. Results confirm the earlier claim that, contrary to SA-LSA, the TI-LSA protocol performance is only affected by the TI size and not by the overall network size.

## V. Summary

This paper described an LSA protocol with message flooding constrained to match the PPO node's TI. The TI-LSA protocol allows PPO nodes to discover resources and detect changes in the optical data plane and physical topology in a scalable manner with respect to the number of nodes. The number of flooded TI-LSA entries is limited by constraining the advertisement within the optical reach of the PPO node, i.e., the PPO node's TI. In other words, each TI-LSA entry reaches only the PPO nodes that require that information, without unnecessarily flooding other PPO nodes that will not make use of that entry. While the TI may resemble the area concept in the Internet, a substantial difference between the two is that areas are manually predetermined, while TI(s) are self detected. The TI(s) may change over time as they automatically adapt to the changes of both the optical network topology and fiber transmission characteristics. And so does the constrained flooding of the TI-LSA entries. While delivering to each PPO node all the TI-LSA entries that may be required to intelligently use the available optical resources, the TI-LSA protocol permits to realize networks with a number of PPO nodes that is virtually infinite, thanks to its TI constrained flooding.

A number of open challenges remain to be addressed and will have to be studied carefully. For example, solutions that provide end-to-end routing across multiple TI(s) must be identified as current inter-area routing solutions cannot work in conjunction with the TI-LSA protocol in a straightforward way [16], [17].

## References

[1] R. Ramaswami and K. N. Sivarajan, *Optical Networks: a Practical Perspective (Second Edition)*. Morgan Kaufmann Publishers, Inc., 2002.
[2] A. Fumagalli, R. Hui, F. Maloberti, S. Gregori, I. Cerutti, and M. Tacca, "High-Speed Self-Configuring Networks Based on Cost-Effective Plug-and-Play Optical (PPO) Nodes," UTD, Tech. Rep. UTD/EE-06/2004, October 2004.
[3] A. E. Willner, "Future Highly-Efficient Optical Networks," in *Proc. of CLEO*, vol. 2, May 2005, pp. 1052–1054.
[4] I. Cerutti, A. Fumagalli, R. Hui, P. Monti, A. Paradisi, and M. Tacca, "Plug and Play Networking with Optical Nodes," in *Proc. of IEEE ICTON*, vol. 3, June 2006, pp. 133–138.
[5] ——, "Plug and Play Optical (PPO) Nodes: Network Functionalities and Built-In Fiber Characterization Techniques," *J. Opt. Netw.*, to appear, 2007.
[6] P. E. Green, *Fiber Optic Networks*. Prentice-Hall, 1993.
[7] J. Moy, "RFC 2328 - OSPF Version 2," April 1998.
[8] D. Katz, "RFC 3630 - Traffic Engineering (TE) Extensions to OSPF Version 2," September 2003.
[9] R. Rastogi, Y. Breitbart, M. Garofalakis, and A. Kumar, "Optimal Configuration of OSPF Aggregates," *IEEE/ACM Trans. on Net.*, vol. 11, no. 2, pp. 181–194, April 2003.
[10] A. Shaikh, D. Wang, G. Li, J. Yates, and C. Kalmanek, "An Efficient Algorithm for OSPF Subnet Aggregation," in *Proc. of IEEE ICNP*, Washington, DC, USA, 2003.
[11] H.-P. Nolting, "All-Optical 3R-Regeneration for Photonic Networks," in *Proc. of ONDM*, 2003.
[12] B. Braden, "RFC 2205 - Resource ReSerVation Protocol (RSVP) – Version 1 Functional Specification," September 1997.
[13] D. O. Awduche, "RFC 3209 - RSVP-TE: Extensions to RSVP for LSP Tunnels," December 2001.
[14] "The Netfilter.org Project," http://www.netfilter.org/.
[15] M. Doar and I. Leslie, "How Bad is Naive Multicast Routing?" in *Proc. of INFOCOM*, vol. 1, March/April 1993, pp. 82–89.
[16] M. S. Savasini, P. Monti, M. Tacca, A. Fumagalli, and H. Waldman, "Regenerator Placement with Guaranteed Connectivity in Optical Networks," in *Proc. of ONDM*, May 2007.
[17] S. Das, P. Monti, M. Tacca, and A. Fumagalli, "Optical Corridor Routing Protocols," in *Proc. of IEEE ICTON*, July 2007.