Maximum Entropy (MaxEnt) Routing and Spectrum Assignment for Flexgrid-based Elastic Optical Networking

CONFERENCE PAPER · MARCH 2014

DOI: 10.1364/OFC.2014.Th4E.7

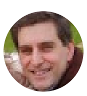

Michael Charles Parker

University of Essex

Andrew Lord

British Telecom

Maximum Entropy (MaxEnt) Routing and Spectrum Assignment for Flexgrid-based Elastic Optical Networking

Paul Wright[1], Michael C. Parker[2], Andrew Lord[1]
1: BT, Adastral Park, Martlesham Heath, Ipswich, IP5 3RE, UK
2: Lexden Technologies Ltd., 67 Temple Road, Epsom, Surrey, KT19 8EY, UK
paul.3.wright@bt.com

Abstract: A maximum entropy approach to routing and spectral allocation in elastic optical networks is implemented using a genetic algorithm and operated on a real network topology. This approach avoids fragmentation problems and increases network utilization.
OCIS codes: (060.4250) Networks; (060.4510) Optical Communications

1. Introduction

Flexgrid is emerging as an important new paradigm in the design of elastic optical networks (EONs) [1], as a means to more efficiently exploit spectrum resources, enable high-capacity superchannels, and allow greater flexibility in the choice of advanced modulation formats. Currently, there are two fundamental (but related) issues affecting flexgrid architectures, which could impede the deployment and successful operation of such networks: optical fragmentation and equipment complexity. The first is due to the arrival of dynamically varying and random demands of differing bandwidth requirements. The difficulty of predicting the size and source-destination pairing of such dynamic demands, means that the optical spectrum tends to become fragmented into narrow bandwidth portions which are potentially too small to be usefully utilized; i.e. only the incidence of a particular (and hence, unlikely) set of demands would ever perfectly exploit the free spectral gaps and lead to most efficient utilization of the network.

A related consequence of the fragmentation problem is that flexgrid photonic equipment has to work that much harder in order to efficiently deal with a fragmented spectrum. Devices such as sliceable bandwidth variable transceivers (S-BVTs), wavelength-selective switches (WSSs) and wavelength converters are being developed to deal with such spectral fragmentation; however, they end up being technically complex. This increases initial deployment costs (CapEx), since this complex technology is more expensive.

We have previously defined a Shannon entropy-based fragmentation metric [2], which accurately measures the fragmentation of a flexgrid optical spectrum. Such an entropic methodology now suggests a maximum entropy (MaxEnt) approach to flexgrid network management, which automatically solves the problems outlined above.

2. MaxEnt Operation of a Flexgrid EON

Up to now, the conventional approach is to keep the optical fibre spectrum as defragmented as possible, so as to operate in the minimum entropy (MinEnt) regime, allowing the easier introduction of new demands into the optical spectrum. For example, in the context of dynamic routing and spectral assignment (RSA), introducing a new source-destination pair demand requires the introduction of a new channel within the optical spectrum; whilst tearing down the demand necessitates the removal of a particular channel bandwidth from the optical spectrum. In optimizing the management of such a MinEnt photonic network, minimized fragmentation is required, with established channels being re-packed together as former demands are removed or reduced, so as to simplify the setting-up of new demands by ensuring that large, contiguous spectrum blocks are available.

Whereas the typical assumption for maximum efficiency means that demands are packed together as closely as possible (MinEnt), we now describe the opposite operation mode, based on the MaxEnt approach. In this case, source-destination pair bandwidth demands are located as far apart from one another as possible across the optical spectrum. Rather than establishing, shifting and tearing-down demands as bandwidth varies between different network nodes, we instead consider having a set of perpetual optical demands between all network node-pairs, whose bandwidths narrow or widen according to the load between any pair of nodes. In having perpetual optical channels which can increase or decrease in bandwidth as required over time, it is clear that having all channels maximally spread-out (fragmented) across the optical spectrum is the most desirable state to have, so that the unfilled intervening spectral resource between channels can be easily exploited at any time (with minimal negotiation) should additional bandwidth capacity be required between a node-pair. Apart from not requiring the ongoing establishment and tearing-down of data channels, operating the fibre's spectrum in a maximally fragmented mode also means that optical guard-bands play a less critical role; with the overall spectral efficiency of operation therefore potentially increasing.

3. Determining a MaxEnt RSA using Genetic Algorithm Optimization

We employ a genetic algorithm (GA) approach to generate the basic MaxEnt EON network design. In order to use a GA algorithm, a way of encoding the network state (the network configuration vector) needs to be defined which represents the routing and spectrum assignment for all demands in the network. This process starts by pre-computing k shortest paths between all pairs of nodes where k is fairly small (e.g. 8) and there is a limit to the end-to-end path length (e.g. 1400 km). The selected route for each node-pair can now be recorded by an integer between 1 and k. Assuming 5000 GHz of C-band spectrum on a single fibre, and an underlying flexgrid with 1-GHz granularity (as supported by the latest WSS devices), the wavelength allocation for each node-pair demand can be identified by an integer between 1 and 5000. The route/spectrum assignment for each node-pair can therefore be defined by two integers and subsequently, for an N node network, the configuration can be defined by a series of $N*(N-1)$ integers.

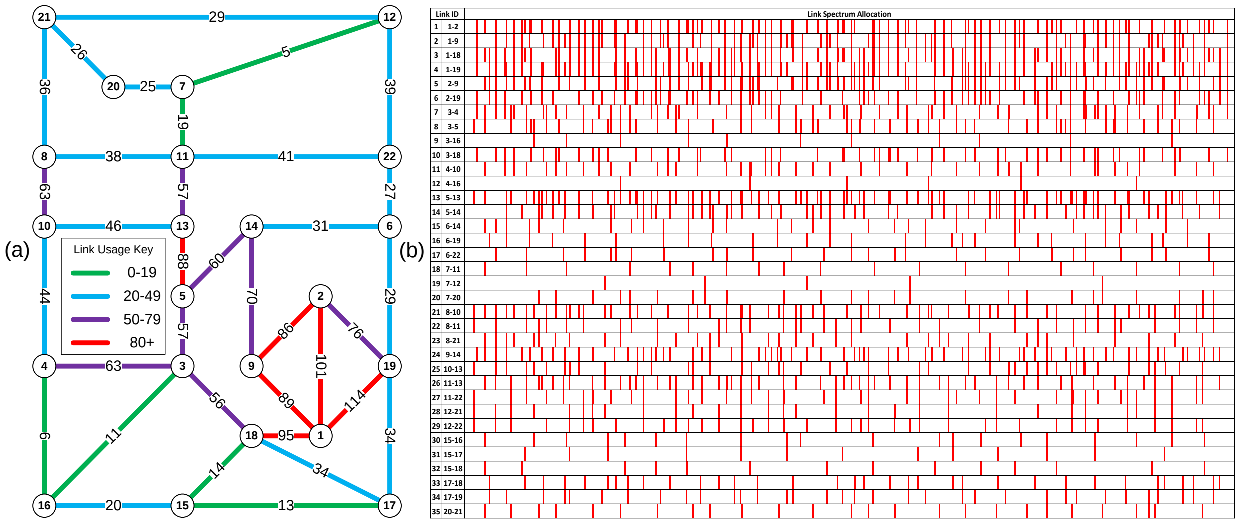

Figure 1: (a) 22 Node BT Reference Network showing MaxEnt link utilization in the single fibre scenario; (b) Link usage diagram showing the centre-frequency placements of node-pair demands across all links in the network in the single fibre scenario.

The GA algorithm is initialized by randomly creating the first generation of network configurations by randomly filling several thousand network configuration vectors. To generate each successive generation, the following process is followed: 1) Validate each network configuration vector to check that there is no signal overlap on any link in the network, ignoring any configuration where an overlap is present. 2) For validated configurations, the Network Shannon Entropy (NSE) is calculated, by summing the individual link entropies using the formula [2]:

$$H_{frag} = -\sum_{i=1}^{P} \frac{D_i}{D} \ln\left(\frac{D_i}{D}\right) \tag{1}$$

where for P blocks of used/unused spectrum, D_i is the number of slots in the current block and D is the total number of flexgrid slots in the entire spectrum band. The NSE calculation is used to quantify how good a solution is (i.e. how separated the signals are from one another) with larger values indicating that the network configuration shows greater signal separation, i.e. nearer to the MaxEnt ideal. 3) Once the NSE is calculated for each valid configuration within a generation, configurations with a NSE that exceeds a certain value (such as being above the average entropy for the generation) are bred together by means of randomly choosing whether to splice or mutate the values within the network configuration vector to generate a new generation of several thousand candidates with the better traits of the current generation. 4) Repeat steps 1 to 3 on the new generation until the average NSE of successive generations has stabilized. 5) The network configuration with the largest NSE in the final generation is then selected as the solution to the problem and used to configure the network nodes and for allocating future demands.

Additional fibre links (i.e. space-division multiplexing or SDM) added between network nodes to ease hot spots and bottlenecks are straightforwardly dealt with, by modeling the increased capacity simply as additional (contiguous) spectrum, but with a periodic property analogous to the free-spectral range characteristic of arrayed-waveguide gratings (AWGs), such that continuity with the spectral paths with other non-SDM links is maintained.

4. Results of Applying the MaxEnt Approach to a Real (BT) Network Topology

The approach described in the previous section has been applied to the 22-node BT reference network shown in Fig.1(a), which therefore has 231 unique node pairings. Two scenarios were considered: the first, where each link is served by a single fibre (SF) pair; and the second, where each link is served by a dual fibre (DF) pair. Each run of the MaxEnt GA algorithm on a network of this size took around 8 hours to compute on a standard single-core PC. The final SF MaxEnt solution showing the spread of demands across all links is shown in Fig.1(b).

The final solution for both scenarios showed a range of different link utilizations, with some links being heavily used (e.g. 114 node-pair channels for SF, 103 for DF) and others being lightly used (5 for both SF and DF) in terms of the number of demands passing through them. The link figures and colouring in Fig.1(a) illustrates how this usage varies across the network in the SF scenario. However, both scenarios show a very similar link usage as can be seen in Fig.2(a), and with an almost identical average number of node-pair channels (46.9 for SF, 46.6 for DF). The reason for this is attributed to the fact that as extra fibres are modeled for additional spectrum capacity, maximal entropy is still achieved by spreading almost the same set of demands further apart across the doubled contiguous DF spectrum. The minor differences are attributed to different GA optimization pathways for the two scenarios.

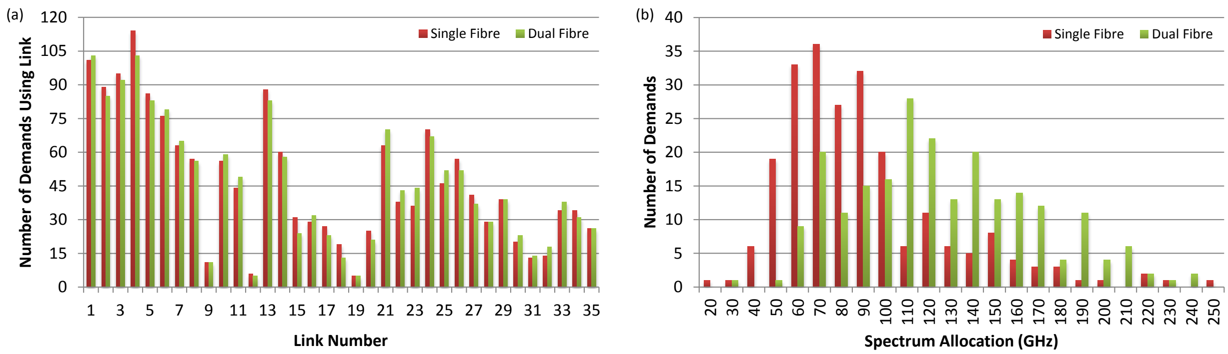

Figure 2: Results from MaxEnt-GA on the BT Reference Network (a) Link usage of solutions; (b) Distribution of allocated spectrum per demand.

The high utilization usage of some links in the network shows that the SF solution may not be able to deliver the required bandwidth flexibility/growth for some demands, and indicates that a multiple fibre solution may be required. Fig.2(b) shows the distribution of the widths of the allocated spectrum slots after the MaxEnt GA algorithm has been run, with the red and green columns showing the results of the SF and DF scenarios respectively. In the SF scenario the minimum allocation is 16 GHz with the mode at 60-70 GHz; and for the DF scenario the minimum is 25 GHz with the mode at 110-120 GHz. We note for any given demand, its allocated spectrum is shared with the demand's neighboring allocations, assuming a superchannel approach where demands can dynamically flex with traffic demands although ultimately, spectrum is limited by the most utilized link that a demand passes through.

Although the modal allocation of both these scenarios provides useful spectrum bands, especially where shorter paths can make use of modulation formats with higher spectral efficiency [3], the rather low spectrum allocation for a few of the demands does indicate that specific links (hot spots) in the network could benefit from additional fibre pairs. The MaxEnt GA solution of Fig.1(b) has only been optimized in terms of maximizing the NSE without considering the size of the spectrum allocations this produces. Further MaxEnt optimization can take place to take account of actual and predicted node-pair traffic levels, so that larger demands are given wider spectrum allocations.

5. Conclusions

In this paper we have presented a new methodology for designing an EON using a basic MaxEnt approach where each unique node-pair within a network is pre-allocated a route and spectrum assignment and using BVTs to dynamically serve traffic between the demands. We have also shown how the MaxEnt problem can be solved using a GA-based optimization, and analyzed the results of applying it to a real (BT) network topology. The results show that the MaxEnt approach produces an insightful network design, and indicate that a multiple fibre solution may be required at certain key links in the network in order to achieve the required bandwidths between the network nodes.

6. Acknowledgements

The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7/2007-2013 under grant agreement n° 317999 IDEALIST project.

7. References

[1] O. Gerstel, M. Jinno, A. Lord, S.J.B. Yoo, "Elastic Optical Networking: A New Dawn for the Optical Layer?", IEEE Communications Magazine, February 2012

[2] P. Wright, M.C. Parker, A. Lord, "Simulation Results of Shannon Entropy based Flexgrid Routing and Spectrum Assignment on a Real Network Topology", Paper We.2.E.4, *ECOC'13*, London, UK, September 2013

[3] D.J. Ives, S.J. Savory, "Fixed versus Flex Grid with Route Optimised Modulation Formats and Channel Data Rates of 400Gbits and Above", Paper P.5.11, *ECOC'13*, London, UK, September 2013