Challenge-Aware Traffic Protection in Wireless Mobile Backhaul Networks

Javier Martín-Hernández[1], Christian Doerr[1], Johannes Lessmann[2], and Marcus Schöller[2]

[1] TU Delft, Department of Telecommunication, Delft, Netherlands
{J.MartinHernandez,C.Doerr}@tudelft.nl
[2] NEC Europe Ltd., Heidelberg, Germany
{Johannes.Lessmann,Marcus.Schoeller}@neclab.eu

Abstract. To protect active traffic against link or node failures in multi-hop communications networks, several so-called protection schemes have been introduced in the past. The most established ones are path, segment, node and link protection. However, these schemes are limited as challenges are modelled abstractly whereas challenges in real networks can have very different characteristics. Thus, we propose to explicitly take the high impact challenges by introducing a risk-group concept into the multi-path placement scheme, which provides an evaluation of the likelihood of a challenge to simultaneously affect two network elements. We have implemented and evaluated this new methodology in simulations and show that it outperforms the original scheme.

Keywords: wireless mesh, wireless backhaul, challenge awareness.

1 Introduction

In the face of exploding user traffic in cellular networks, increasing the spectral efficiency to achieve higher access capacities is widely considered to be only possible via SDMA (space division multiple access). To this end, very large numbers of small cell base stations will be brought up in the near future. As an example, Picochip, a femtocell maker, claims that London needs to install 70,000 femtocells by 2015 to provide decent 4G LTE mobile services [1]. Since wired backhaul will not be generally available for all small cells, wireless backhaul networks will gain importance.

A drawback of a wireless backhaul network is the instability of its links. Even highly directed microwave links such as used for carrier-grade mobile backhaul networks are affected by bad weather conditions and link quality may degrade or a link might fail completely [2]. To avoid re-routing latency, a common approach is to proactively establish backup routes which are used in case the primary route fails. In the routing community, this is called multi-path routing, in the telecom world, this is known as protection. Given a source node, a destination node and a primary path between them, protection can happen at multiple levels of granularity. Figure 1 shows different widely established protection schemes.

Z. Becvar et al. (Eds.): NETWORKING 2012 Workshops, LNCS 7291, pp. 35–42, 2012.

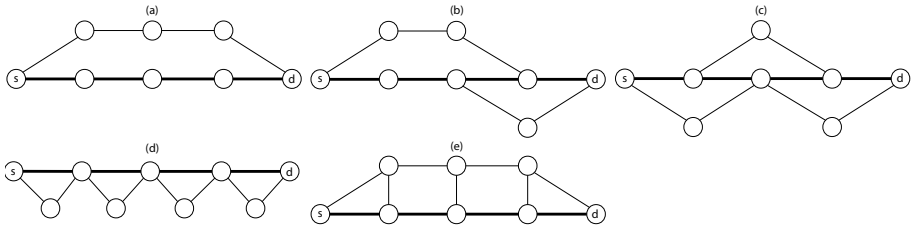

Fig. 1. Different protection schemes. a) path protection. b) segment protection. c) node protection. d) link protection. e) rope-ladder protection.

Rope-ladder protection (RLP) as introduced in [3] (cf. also Figure 1e)) combines the advantages of path, node and link protection by constructing two node-disjoint paths between s and d (i.e. "ropes") and connecting each node on the primary path with a node on the backup path (via "rungs"). As shown in [3], this increases path diversity and path lifetime while reducing loss gaps.

One of the major novelties of this paper over [3] is that it introduces the notion of *challenges* and thus challenge-awareness. A challenge to a wireless backhaul communication network could be a thunderstorm, a congestion hotspot or a virus attack, for example. Many protection schemes are either completely challenge-unaware or designed to meet only one particular challenge. In this paper, a path construction algorithm is proposed that can optimize the protection structure with respect to the high-impact challenges. Based on expert risk assessment, we apply the Shared-Risk Link Group concept known from the optical networks space [4] to the used protection scheme. A special entity of the network management system translates a challenge into a so-called *risk group*. Based on this, it steers the rope-ladder construction process such as to be maximally robust towards the high impact challenges. The mentioned special entity is the Graph Explorer, introduced in [5]. The Graph Explorer (GE) is a general tool that can explore a large set of properties in multi-hop networks. For the sake of this paper, we will use the capabilities of GE to compute the impact of challenges on the network and to steer the rope-ladder construction process accordingly. This will be described in detail in Section 3.

This paper is organized as follows. Section 2 discusses previous research in the domain of multipath QoS routing or protection and its relation to our proposition. Section 3 introduces the rope-ladder construction with the help of GE. Section 4 presents our performance evaluation. Section 5 gives a short summary.

2 Related Work

As mentioned previously, work in the context of our paper is discussed as multi-path routing as well as protection schemes.

There is quite a number of works in the domain of multipath routing protocols. An overview of this diverse field can be found in [6]. Different multi-path routing

protocols focus on a number of aspects like load balancing, bandwidth bundling, security, congestion control or even security (sending packets of a sensitive flow via different paths to make eavesdropping more difficult). One goal is obviously resilience to failures. A further classification refers to the independence of the individual paths of the multipath. To increase security or robustness, the paths should be as independent as possible which gives rise to node-disjoint or link-disjoint multipaths. Some protocols do not make any statement or assumption about the path independence.

In the telecom world, particularly in the area of optical networks, related work is known as protection schemes. Terms like path, node or link protection are commonly used in traffic engineering technologies such as MPLS and do not need further discussion here. Interesting to mention, however, is that the optical networks community has introduced a concept called shared-risk link group (SRLG) [4]. A SRLG contains all links in a network that are susceptible to the same risk. The typical use case would be two optical fibers which share a common duct. If the duct is destroyed, likely not only one but both fibers share its fate, leading to disruption of traffic through both fibers. The concept of SRLG is very generic, allowing to capture arbitrary risks. Similarly, shared-risk node groups capture risks impacting one or multiple nodes. In [7], SRLG and SRNG are combined into shared-risk resource groups. Probably most of the SRLG related propositions occupy themselves with finding SRLG (SRNG/SRRG) diverse paths (path protection). However, they do not compare different protection schemes. In this paper, we will compare path and rope-ladder protection in the face of SRLGs.

3 Constructing Challenge-Aware Protection Schemes

In this section, we describe how risk-aware protection schemes between a source and a destination node are constructed. First, we will introduce risk group models of three different challenges. Then we describe protection construction process based on the challenges which are to be considered.

3.1 Challenge Model and Risk Groups

A challenge is an event which occurs in the network and which threatens the network's normal operation. Examples for such challenges in wireless networks include for example adverse weather conditions, virus attacks, failures of software components, equipment theft or network overload.

As indicated previously, the optical networks domain has introduced the idea of shared-risk link groups (SRLG). Here, we apply this concept to multi-path protection. A challenge C is modelled in terms of such a risk group as a set of network elements failing simultaneously. The risk group of a challenge C, denoted as RG_C, is defined in terms of the *logical vicinity* of the protected elements (i.e. nodes and links): given that element e_1 is in risk group RG_C, element e_2 is also in RG_C if its logical vicinity to e_1 – denoted $v(RG_C, e_1, e_2)$ – is above a certain threshold τ. Formally, $n_1 \in RG_C \Rightarrow \forall n_2, v(RG_C, e_1, e_2) > \tau : e_2 \in RG_C$.

The impact of all challenges has been modelled with a step function (corresponds with setting $\tau = 0$) in our simulations. In other words, any link or node affected by the challenge C fails reducing its bandwidth to zero, hence $RG_C = C$. The impact function τ can easily be extended to complex and realistic scenarios. The logical vicinity function $v(RG_C, e_i, e_j)$ of all considered challenges needs to be defined manually by a network expert and can span one or more arbitrary dimensions. For instance, in areal challenges (e.g. a storm cell) the logical vicinity of two elements correlates with the geographical vicinity of the elements, whereas logical vicinity of a challenge exploiting a software bug correlates with the vendor ID.

In this paper, we have modelled three different classes of challenges: (i) a flash crowd event at a congestion hotspot, (ii) a heavy rain cell, and (iii) a virus targeting a firmware bug. A single **congestion hotspot** is defined by a static area, e.g., a train station or a stadium, where huge numbers of users can cause overload situations. All received connection requests are legitimate but cannot be satisfied by the system simultaneously; such events are called *flash crowd events* in contrast to denial of service attacks which are of malicious nature. The logical vicinity function is defined by the area the C affected elements are located at. The second condidered disruption is an areal **thunderstorm cell** moving randomly across the graph, producing a large set of independent thunderstorm challenges $\mathcal{C} = \{C_1, C_2 ... C_k\}$. The logical vicinity function of a single thunderstorm challenge C_i (for $i \leqslant k$) provides that elements e_1 and e_2 appear within the same risk group if a circular rain area with a radius $r(C_i)$ and an given epicentre $\varepsilon(C_i)$ overlaps both elements. Often, logical vicinity will be related to geographic positioning (e.g. distance to the epicentre) but other environmental characteristics may define this function, too. The last attack we considered is a generic virus attacking one **firmware vulnerability** of exposed graph elements, producing a set of firmware challenges $\mathcal{C}$. If a single firmware challenge C_i is launched against the mesh network, all nodes using the targeted firmware version $f(i)$ are threatened and hence share the same risk group.

3.2 The Construction Process

In order to construct challenge-aware rope-ladders, we have combined the RLP scheme with the Graph Explorer as introduced in [5]. The Graph Explorer is a tool to assess various metrics of a network in the face of an arbitrary number of challenges occurring simultaneously in this network. We use this tool to calculate the risk groups before starting to place the rope-ladder structure. The construction of a rope-ladder is divided in three sequential steps as follows.

Step A: Placement of the Primary Path. As the risk groups depend on the links and nodes of the primary path, the choice of the primary path is a crucial one in our process. An intuitive approach would choose the primary path to be the shortest path from source to destination. However this can lead to fragile backup paths crossing high risk elements. Thus, we propose that the process should iterate over all paths up to a maximum stretch with respect to the shortest path, providing a set $\mathcal{P}$ of paths. Eventually, the primary path which

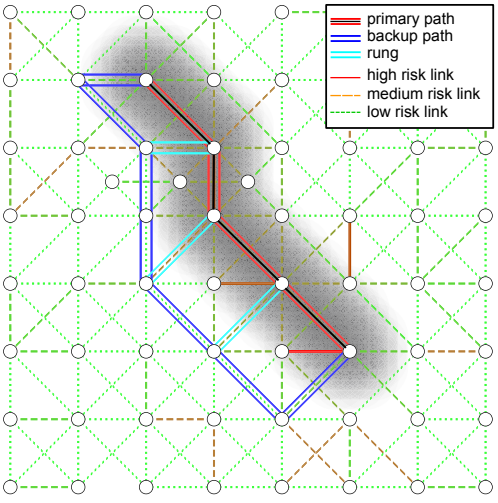

Fig. 2. A challenge-aware RLP scheme for one random Access Point and a Relay node on the G_B topology. The effect of the thunderstorm is displayed as a shadow covering the susceptible links, subject to the condition that the storm hitting the primary path. The darker the shadow, the higher the likelihood for a link to fail together with the primary path.

leads to the less risky backup path is selected as the primary path (as described in the last step).

Step B: Calculation of Link Weights. The input to this step are all the primary paths $\mathcal{P}$ provided by step A, and the high-impact set of challenges $\mathcal{C}$ from which the network should be protected. Depending on the chosen vicinity function, additional information must be made available such as firmware ID, the frequency allocation plan, etc. Multiple risk groups can be added into a unified risk group (URG) by merging the link weights of different challenge types. This merging function must be determined during the network manager's risk assessment process, and it should account for the respective occurrence probability of the different challenges types. In our simulations we assumed that all challenge types are both independent and equally probable. Hence all weights belonging to the same challenge type were further normalized to sum the complete probability. The output of step B are multiple *weight clouds* $\mathcal{W}$, i.e. sets of node and link weights representing the risk group memberships with respect to each primary path. This weight cloud calculation process is computed such that the weight of element e will increment for every time that e shared a challenge in $\mathcal{C}$ with any of the elements in $P \in \mathcal{P}$. An intuitive visual representation of each primary path's risk group is the union of all the challenge instances $\mathcal{C}$ (e.g. thunderstorms in this example) that intersect with the primary path by at least one link or node. The storm's link weights associated with the shortest path can be illustrated as the *cloud* shown in Figure 2.

Step C: Placement of the Backup Structure. A set of backup paths $\mathcal{B}$ is found by iterating over all URG pairs $\{P, W\}$ offered by step B. A shortest path algorithm determines the backup path with the least weight which does not exceed an arbitrary stretch limit. The basic idea of our approach is that the backup path circumvents the *cloud* in Figure 2 and stays out of it for as long as possible, hence minimizing the link weights that will be crossed. An enhancement to this step currently under investigation as part of our ongoing research is the use of a risk threshold heuristic which is acceptable for the backup path. The process would only iterate over the shortest paths until a backup path is found with acceptable risk, improving the overall time performance. Finally, once all $\{P, W\}$ pairs have been processed, the backup path with the lowest weight in $\mathcal{B}$ is selected to form the rope-ladder; the rungs are determined afterwards (the rungs are cross-connects from the primary path to the backup path to minimize the loss gap for real time application flows - see [3] for details).

4 Simulation Results

This section describes the simulation scenario built to evaluate the performance of challenge-aware RLP as introduced in Section 3, followed by a qualitative analysis of the simulation results.

4.1 Scenario Description

The selected application is a unicast VoIP application, simulating a G.711 VoIP codec over 1Mb/s duplex connections. To simulate the effect of a challenge on a voice stream, a voice call is held between two random nodes for an arbitrary time span of 3 minutes. This data flow is established via a primary path following a RLP scheme. One minute into the call, an instance of a challenge occurs; causing the bandwidth of the affected links to be reduced to zero for the duration of the challenge, virtually disconnecting them. As soon as the links become unavailable, the central routing protocol will divert in-flight packets and adapt routing tables to the RLP scheme backup path through the rung which is closest to the challenge. The challenge remains in place for one minute, after which all the links in the network are restored to their initial state.

4.2 Simulation Results

First, we measure the packet loss that different protection schemes suffer by a storm cell occurring. The percentage of packets lost by an oblivious RLP scheme (i.e. a rope ladder uninformed about possible challenges during construction) is 10.3%, approximately two times the percentage of packets lost by the RG_{Storm} aware RLP scheme, which lost 4.8% of packets as shown in the leftmost chart of Figure 3(a). Secondly we evaluate the gap size, measured as the the maximum difference in sequence numbers between two consecutive received frames. Given that the routing is controlled by a central authority, the delay induced

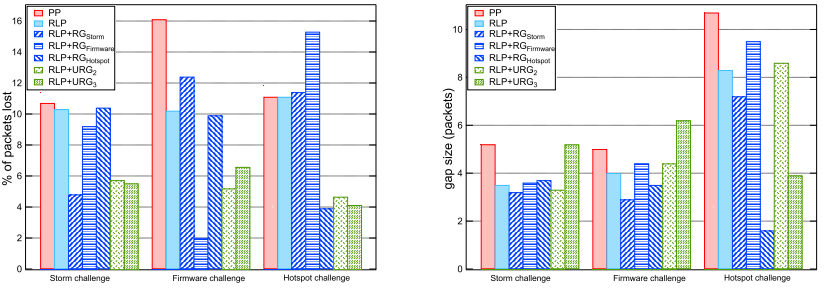

(a) Percentage of packets lost (b) Average gap size

Fig. 3. (a) Percentage of packets lost and (b) average gap size of protection schemes. The six sets of charts display the gap size effects of *Storm* (left), $Firmware$ (middle), and $Hostspot$ (right) challenges on G_B. Each set displays seven bars, corresponding to the protection schemes: oblivious PP, oblivious RLP, RG_{Storm}-aware, $RG_{Firmware}$-aware, $RG_{Hotspot}$-aware, and URG_2-aware and URG_3-aware.

by a challenge message propagation is dismissed, i.e., the routing tables are instantly updated across the network. This effect works in favour of PP schemes, by ignoring the propagation delay involved in route table synchronization. Nevertheless the flows' gap size effect of rerouted in-flight packets is still noticeable. All three simulation scenarios displayed in Figure 3 illustrate RLP challenge aware schemes suffering smaller packet loss and gap sizes than PP oblivious schemes.

In order to test the behaviour of RLP schemes subject to challenges not included in the planned risk group RG_C we also measured the performance of a challenge aware RLP scheme against unexpected sets of challenges. The blue striped bars in Figure 3(b) illustrate this effect: challenge aware RLP schemes' performance degrades under the effect of unexpected challenges. The performance of schemes under unexpected attacks may even degrade beyond their oblivious counterparts. Such is the case for RLR RG_{Storm} schemes under Firmware challenges, which lost 4.4% of the voice packets, as opposed to only 4.0% for the oblivious scheme. This adverse effect motivated us to study multi-challenge aware protection schemes through the use of URG. Ideally a multi-challenge aware rope ladder structure can withstand different non-simultaneous attacks without a significant drop in performance. First we define URG_2 as the unified risk group formed by adding the two risk groups with the highest impact out of the storm cell, firmware virus and hotspot. Additionally we define URG_3 as the risk group resulting from the addition of all three risk groups. Weights are consequently normalized, as specified in section 3. Simulations show that the percentage of packets lost by URG_2 and URG_3 schemes under a storm cell challenge are 5.4% and 5.2% respectively (as illustrated Figure 3(a)), representing a significant improvement over the oblivious RLP scheme (10.3%), yet not performing as good as a RG_{Storm} aware scheme (4.8%). Overall, in terms of packet

loss all URG protection schemes' outperform both their oblivious and challenge-aware schemes under the effects of unexpected challenges. On the other hand Figure 3(b) illustrates URG schemes suffering large gap sizes. The protection scheme with the lowest gap size is the one tailored to the challenge, i.e. $RG_{Hotspot}$ with a gap of 1.6 packets. However the gap sizes of URG_2 and URG_3 schemes under flash crowd challenges are 8.6 packets and 3.9 packets respectively, whereas the oblivious RLP scheme only lost 8.3 packets. In conclusion, URG schemes' gap sizes are highly dependant on the type of challenge and considered risk groups; these schemes may underperform their oblivious counterparts when faced to an expected sets of challenges.

5 Conclusion

In this paper, we have presented an algorithm to improve the placement of rope-ladder protection schemes in multi-hop wireless networks. The algorithm is based on the Graph Explorer, a general tool to explore properties and metrics in arbitrary graphs. During the network planing phase, we have employed a risk group approach which makes use of a logical vicinity function that relates each link and node in the network to individual risk groups. During network operation the Graph Explorer assesses possible placements of rope-ladders such as to be maximally robust towards certain challenges. Using simulations, we have evaluated this challenge aware rope-ladder scheme with the original rope-ladder scheme and with path protection. The packet loss rate was reduced by up to 80.4% compared to the oblivious scheme, but more remarkably the number of protection schemes surviving the challenge onset was increased by up to 25%. Focusing on the high impact challenges during the network design stage is critical.

Acknowledgements. The research leading to these results has been funded by the European Commission, under grant agreement no. 224619 (ResumeNet).

References

1. `http://eweekeurope.co.uk/news/london-needs-70000-cells-for-4g-broadband-40779`
2. Jabbar, A., Raman, B., Frost, V.S., Sterbenz, J.P.G.: Weather Disruption-Tolerant Self-Optimising Millimeter Mesh Networks. In: Hummel, K.A., Sterbenz, J.P.G. (eds.) IWSOS 2008. LNCS, vol. 5343, pp. 242–255. Springer, Heidelberg (2008)
3. Lessmann, J., Schöller, M., Zdarsky, F., Banchs, A.: In: Proceedings of the 2010 IEEE WoWMoM, WOWMOM 2010. IEEE Computer Society (2010)
4. Sebos, P., Yates, J., Hjalmtysson, G., Greenberg, A.: In: Optical Fiber Communication Conference, WDD3. Optical Society of America (2001)
5. Doerr, C., Hernandez, J.M.: In: Third International Conference on Dependability, DEPEND (2010)
6. Mueller, S., Tsang, R.P., Ghosal, D.: Multipath Routing in Mobile Ad Hoc Networks: Issues and Challenges. In: Calzarossa, M.C., Gelenbe, E. (eds.) MASCOTS 2003. LNCS, vol. 2965, pp. 209–234. Springer, Heidelberg (2004)
7. Datta, P., Somani, A.K.: Comput. Netw. 52, 2381–2394 (August 2008)