Instituto de Engenharia de Sistemas e Computadores de Coimbra
Institute of Systems Engineering and Computers
INESC - Coimbra

Dorabella Santos, Teresa Gomes,
Abdulaziz Alashaikh, David Tipper

A Numerical Study of Three Cost Functions for Network Design to Support High Availability Link Upgrade

No. 1 2021

ISSN: 1645-2631

Instituto de Engenharia de Sistemas e Computadores de Coimbra
INESC - Coimbra
Rua Sílvio Lima, Pólo II, 3030-290 Coimbra, Portugal
www.inescc.pt

The work was financially supported by FEDER Funds and National Funds through FCT under project CENTRO-01-0145-FEDER-029312.

A Numerical Study of Three Cost Functions for Network Design to Support High Availability Link Upgrade

Dorabella Santos(1), Teresa Gomes(1,2), , Abdulaziz Alashaikh(3) David Tipper(4)

(1) INESC Coimbra, Department of Electrical and Computer Engineering, University of Coimbra, Polo 2, 3030-290 Coimbra, Portugal.

(2)Department of Electrical and Computer Engineering, University of Coimbra, Polo 2, 3030-290 Coimbra, Portugal.

(3)Department of Computer and Network Engineering, University of Jeddah, Jeddah, Saudi Arabia

(4) University of Pittsburgh, Pittsburgh, PA, USA

e-mail: dsantos@inescc.pt, teresa@deec.uc.pt, asalashaikh@uj.edu.sa, dtipper@pitt.edu

December, 2021

Abstract

To ensure high end-to-end availability in networks, we have considered the *spine* concept, where a set of links are selected to be upgraded to have improved availability at a given cost. The goal of the *spine* is to have an embedded subgraph, at the physical layer, to support high availability between certain end nodes, while providing differentiated classes of resilience with varying levels of end-to-end availability. We present here a numerical analysis of three possible upgrade cost functions and identify one as the more realistic cost function.

Keywords: *optimization, network design, availability, cost function*

1 Introduction

Availability target guarantees can be difficult to ensure using path protection alone [1]. Therefore we assume that the links in the network can be upgraded to have improved availability.

Consider a network represented by a graph $G = (N, E)$, where N is the set of nodes and E is the set of links. Each link $e \in E$ has an associated length ℓ_e and a default availability α_e that depends on its length. Each link can be upgraded to have improved availability $\hat{\alpha}_e$ at a given cost c_e. For each origin-destination node pair (o, d), we intend to ensure that the end-to-end availability via a pair of node-disjoint paths, the active path p_a and the backup path p_b, is at least a target value λ.

The default availability of a path p is given by

$$\mathcal{A}_p = \prod_{e \in p} a_e \tag{1}$$

and the end-to-end availability of a pair of node-disjoint paths between (o, d) is given by

$$\mathcal{A}_{ee} = 1 - (1 - \mathcal{A}_{p_a})(1 - \mathcal{A}_{p_b}) \tag{2}$$

In order to have $\mathcal{A}_{ee} \geq \lambda$ some links of the path pair may need to be upgraded.

2 Cost functions

We assume three costs functions for upgrading links considered in [1]:

$$\begin{aligned} c_{1e} &= (\hat{\alpha}_e - \alpha_e)^2 \cdot \ell_e & (3) \\ c_{2e} &= \left(\frac{\hat{\alpha}_e - \alpha_e}{1 - \alpha_e}\right)^2 \cdot \ell_e & (4) \\ c_{3e} &= -\log\left(\frac{1 - \hat{\alpha}_e}{1 - \alpha_e}\right) \cdot \ell_e & (5) \end{aligned}$$

The upgrade cost of the set of upgraded links $\mathcal{U}$ for each cost function is given by

$$c_i = \sum_{e \in \mathcal{U}} c_{ie} \quad i = 1, 2, 3 \tag{6}$$

The cost function c_1 reflects by how much the link availability was improved. If the links already have high availability α_e, then $\widehat{a_e} \approx a_e$ and so the cost will be small. The cost function c_2 reflects not only how much the link availability was improved, but also on how available it already is. This means that improving a link with a higher availability is more costly than improving one with lower availability by the same amount, which is more realistic. The cost function c_3 reflects that the increase in cost is exponentially slower, the higher the link availability α_e is. The latter function is the one that seems more realistic.

3 Improving link availability - a numerical study

Consider a origin-destination node pair (o, d) connected via a pair of node-disjoint paths, the active path p_a and the backup path p_b.

3.1 Improving path availability

Similarly to [1], consider that p_a has a default availability (given by the product of its links' availabilities) of 0.99 and that p_b has a default availability of 0.90. We consider that it is possible to improve the availability of p_a and/or p_b by some factor ϵ in the following way: (i) either by improving $\mathcal{A}_{p_a}$ by ϵ for which we have $\mathcal{A}_{ee} = 1 - [1 - (\mathcal{A}_{p_a} + \epsilon)](1 - \mathcal{A}_{p_b})$; (ii) by improving $\mathcal{A}_{p_b}$ by ϵ for which we have $\mathcal{A}_{ee} = 1 - (1 - \mathcal{A}_{p_a})[1 - (\mathcal{A}_{p_b} + \epsilon)]$; (iii) or by improving both for which we have $\mathcal{A}_{ee} = 1 - [1 - (\mathcal{A}_{p_a} + \epsilon/2)][1 - (\mathcal{A}_{p_b} + \epsilon/2)]$.

This behavior is shown in Fig. 1, where the x-axis shows the ϵ values and the y-axis shows the $\mathcal{A}_{ee}$ values. Note that improving the active path availability (blue line) has the highest impact on the end-to-end availability $\mathcal{A}_{ee}$, since it has a higher default availability already. Note that improving the backup path availability (orange line) has the lowest higher impact on $\mathcal{A}_{ee}$, while the combination of both (green line) is in-between. We conclude that it is more efficient to improve the availability of the path with highest availability (usually the active path).

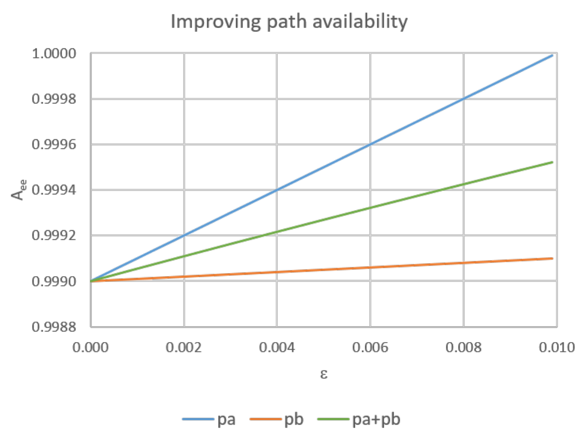

Figure 1: Impact of path availability improvement on the end-to-end availability

3.2 Improving availability in all links

To achieve the improvement ϵ in $\mathcal{A}_{ee}$ and see the impact in the cost, we assume that either all links of p_a or all links of p_b were improved, or that all links of the path pair have been improved. To do so, we assume all links has the same normalized length $\ell = 1$ and so all links have the same default availability α. To illustrate this assume that p_b has always one more link than p_a: p_a has n links and p_b has $n+1$ links.

Then if all links of p_a are upgraded to have improved availability $\hat{\alpha}$, we have:

$$\hat{\alpha}^n = \mathcal{A}_{p_a} + \epsilon = \alpha^n + \epsilon \quad (7)$$

$$\hat{\alpha} = (\mathcal{A}_{p_a} + \epsilon)^{1/n} = (\alpha^n + \epsilon)^{1/n} \quad (8)$$

The same reasoning can be done for p_b:

$$\hat{\alpha} = (\mathcal{A}_{p_b} + \epsilon)^{1/(n+1)} = (\alpha^{n+1} + \epsilon)^{1/(n+1)} \quad (9)$$

and for p_a and p_b with $\epsilon/2$:

$$\hat{\alpha}_a = (\mathcal{A}_{p_a} + \epsilon/2)^{1/n} = (\alpha^n + \epsilon/2)^{1/n} \quad (10)$$

$$\hat{\alpha}_b = (\mathcal{A}_{p_b} + \epsilon/2)^{1/(n+1)} = (\alpha^{n+1} + \epsilon/2)^{1/(n+1)} \quad (11)$$

Fig. 2 illustrates the impact of ϵ on the cost functions c_1 (top), c_2 (middle) and c_3 (bottom) for three sets: Set 1 when $n = 1$ (left), Set 2 when $n = 2$ (center) and Set 3 when $n = 3$ (right). The default $\mathcal{A}_{pa}$ and $\mathcal{A}_{pb}$ values have been assumed to be 0.99 and 0.90, respectively. The blue line refers to all links of p_a being improved to have an improvement on $\mathcal{A}_{p_a}$ of ϵ, the orange line all links of p_b being improved to have an improvement on $\mathcal{A}_{p_b}$ of ϵ, and the grey line all the links of the path pair being improved to have an improvement on $\mathcal{A}_{p_a}$ and $\mathcal{A}_{p_b}$ of $\epsilon/2$.

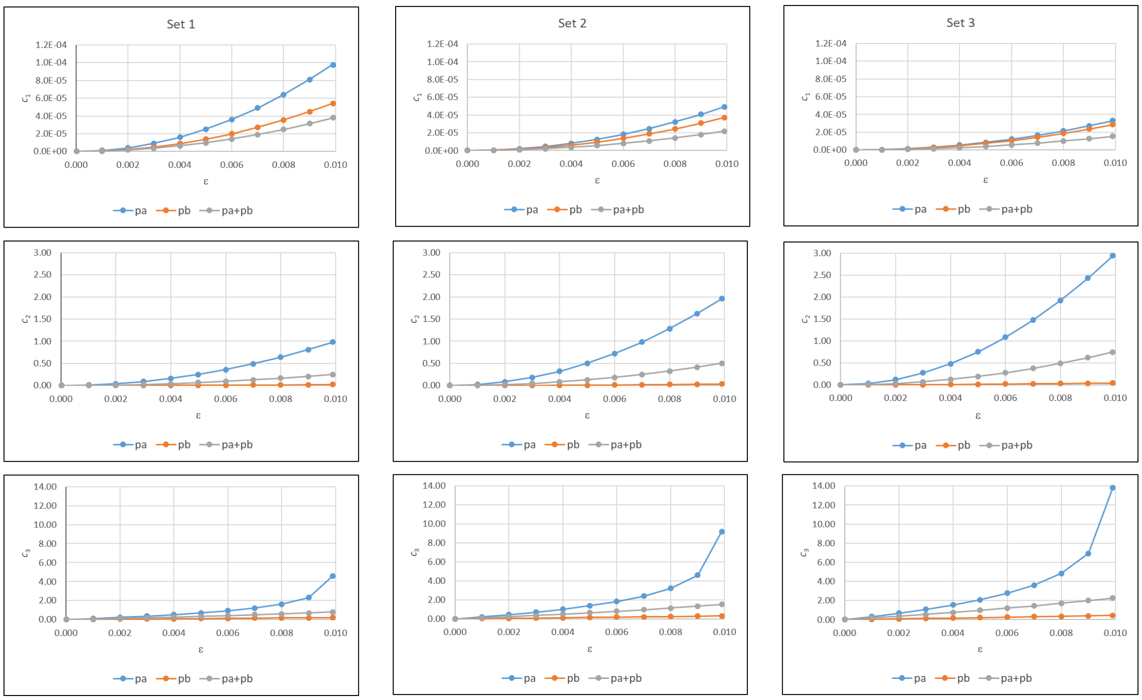

Figure 2: Impact of ϵ on the cost functions c_1 (top), c_2 (middle) and c_3 (bottom) when improving all links of p_a or p_b, or improving all links of both paths

In turn, Fig. 3 illustrates the impact of the achieved $\mathcal{A}_{ee}$ (for each value of ϵ in the previous figure) on the cost functions c_1 (top), c_2 (middle) and c_3 (bottom) for Set 1 to Set 3.

Note that the cost function c_1 is the one that is least realistic, in the sense that as the number of links grow from Set 1 to Set 3, the cost in Fig. 2 decreases since it reflects only the variation in each link availability improvement. Since the default path availabilities are fixed, the more links exist, the higher each link default availability needs to be and so the variation in improvement is not significant. However, c_2 also reflects on how available the link already was and so is much more informative and realistic. Note that the cost function grows as n grows from Set 1 to Set 3, which is also observed in c_3, but in a more significant manner making this cost function the most desirable and realistic.

In all three functions, it is clear that improving the availability of p_a is the most costly, but as seen in Fig. 3 is the most effective in improving $\mathcal{A}_{ee}$. Fig. 2 shows that improving the availability of p_b is the least costly for c_2 and c_3, but as seen in Fig. 3 is the least effective having achieved $\mathcal{A}_{ee} < 0.9992$. The improvement in both paths is a compromise since it achieves $\mathcal{A}_{ee} < 0.9996$

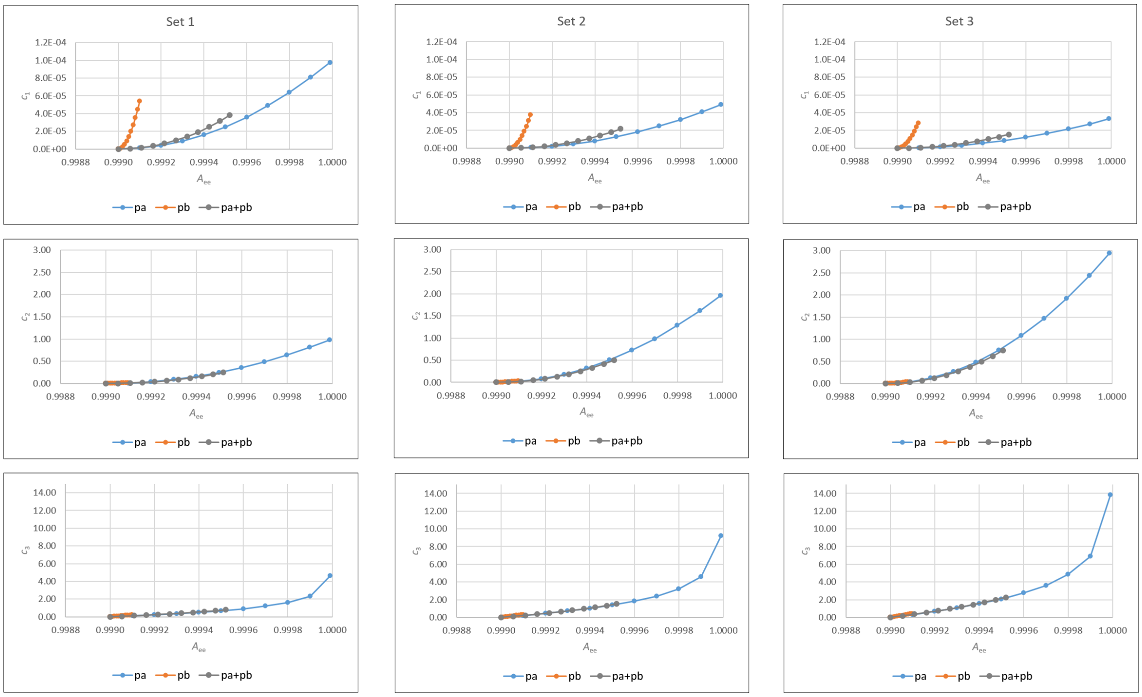

Figure 3: Impact of A_{ee} on the cost functions c_1 (top), c_2 (middle) and c_3 (bottom) when improving all links of p_a or p_b, or improving all links of both paths

but being significantly less costly than the improvement in p_a.

We have also obtained computational results considering the default path availability values $\mathcal{A}_{pa} = 0.999$ and $\mathcal{A}_{pb} = 0.990$ with ϵ ranging from 0.0001 to 0.0010. The results show a similar behavior as can be seen in Fig 7 and Fig. 8 in the Appendix.

3.3 Improving availability in one link

With the same assumptions as above, we now consider that only one link of p_a is upgraded to have improved availability $\hat{\alpha}$:

$$\mathcal{A}_{p_a} + \epsilon = \hat{\alpha} \cdot \alpha^{n-1} \tag{12}$$

$$\alpha^n + \epsilon = \hat{\alpha} \cdot \alpha^{n-1} \tag{13}$$

$$\hat{\alpha} = \frac{\alpha^n + \epsilon}{\alpha^{n-1}} \tag{14}$$

$$\hat{\alpha} = \alpha + \frac{\epsilon}{\alpha^{n-1}} \tag{15}$$

The same reasoning can be done for p_b:

$$\hat{\alpha} = \alpha + \frac{\epsilon}{\alpha^n} \tag{16}$$

and for p_a and p_b with $\epsilon/2$:

$$\hat{\alpha}_a = \alpha + \frac{\epsilon/2}{\alpha^{n-1}} \tag{17}$$

$$\hat{\alpha}_b = \alpha + \frac{\epsilon/2}{\alpha^{n}} \tag{18}$$

Fig. 4 illustrates the impact of ϵ on the cost functions c_1 (top), c_2 (middle) and c_3 (bottom) for Set 1 (left), Set 2 (center) and Set 3 (right). The default $\mathcal{A}_{pa}$ and $\mathcal{A}_{pb}$ values are one again assumed to be 0.99 and 0.90, respectively. The blue line refers to one link of p_a being improved to have an improvement on $\mathcal{A}_{p_a}$ of ϵ, the orange line one link of p_b being improved to have an improvement on $\mathcal{A}_{p_b}$ of ϵ, and the grey line one link of each path being improved to have an improvement on $\mathcal{A}_{p_a}$ and $\mathcal{A}_{p_b}$ of $\epsilon/2$.

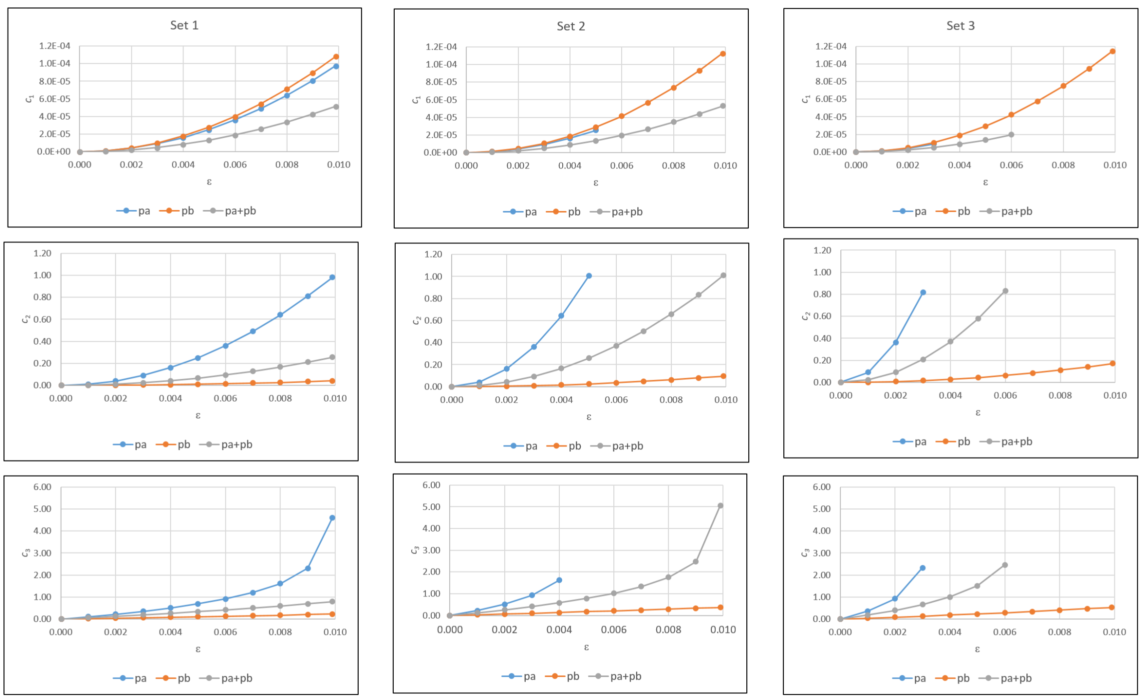

Figure 4: Impact of ϵ on the cost functions c_1 (top), c_2 (middle) and c_3 (bottom) when improving one link of p_a or p_b, or improving one link of each path

In turn, Fig. 5 illustrates the impact of the achieved $\mathcal{A}_{ee}$ (for each value of ϵ in the previous figure) on the cost functions c_1 (top), c_2 (middle) and c_3 (bottom) for Set 1 to Set 3.

Once again, Fig. 4 shows that c_1 has a undesirable behavior. So let us focus only on cost functions c_2 and c_3. Note that in this case, because we are only improving a single link in the path(s), the link needs to be upgraded to a higher value to ensure the desired outcome in $\mathcal{A}_{p_a}$ and/or $\mathcal{A}_{p_b}$. So it may not be possible choose ϵ up to 0.010, since this may cause $\hat{\alpha} > 1$ which does not make sense, which is why in Fig. 4 some curves stop at a certain value of ϵ. Despite that it maintains that improving the availability of p_a is then most costly, of p_b is the least costly and of both paths is in the middle. As confirmed by Fig. 5, it is also true that improving the

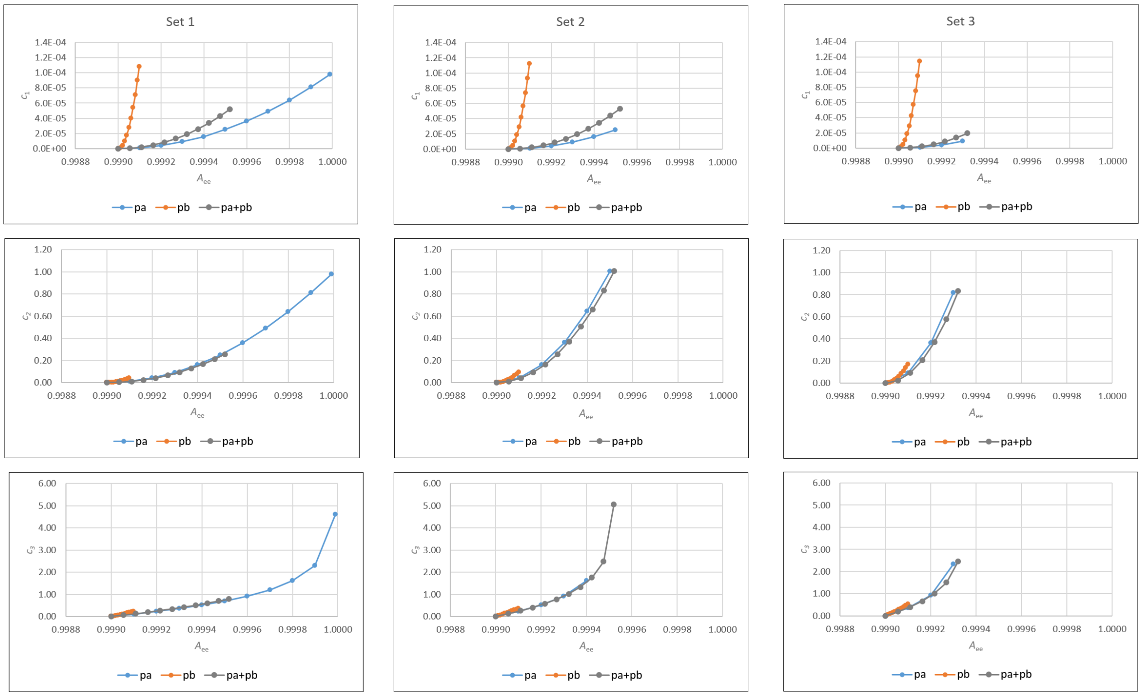

Figure 5: Impact of A_{ee} on the cost functions c_1 (top), c_2 (middle) and c_3 (bottom) when improving one link of p_a or p_b, or improving one link of each path

availability of p_a is the most effective.

We have also obtained computational results considering the default path availability values $\mathcal{A}_{pa} = 0.999$ and $\mathcal{A}_{pb} = 0.990$ with ϵ ranging from 0.0001 to 0.0010. The results show a similar behavior as can be seen in Fig 9 and Fig. 10 in the Appendix.

4 Conclusions

We have seen that an acceptable cost function should not only reflect the improvement in the link availability, but should also reflect how available the link already was, such as c_2 and c_3. Moreover, to have a more desirable behaviour the cost function should exhibit a significant increase the larger the improvement and also the higher the default availability, such as c_3. In other words, it is more costly to improve more available links than less available ones. Therefore, c_3 is the most desirable cost function presented.

It is difficult to obtain an exact expression for the cost function. Increasing link availability can be done by increasing the mean time between failures (more robust links, etc.) or the mean time to repair (better maintenance teams, prioritizing maintenance on these links, etc.), which involves several factors to be accounted for. Although c_3 is a simplistic approach, it allows us to study the impact of link availability improvement on the cost, in a satisfactory manner.

We note that to achieve a target path availability, improving only on a single link is too great an effort and is not as effective as improving on all links. We have used toy examples to illustrate this behaviour. In practice, the links have different lengths and the longer links are less available than the shorter ones. Moreover, usually it is a set of links that are upgraded instead of a single link or all links.

References

[1] A. Alashaikh, D. Tipper, and T. Gomes, "Embedded network design to support availability differentiation," *Annals of Telecommunications*, vol. 74, no. 9-10, pp. 605–623, 2019.

Appendix

We have also obtained computational results considering the default path availability values $\mathcal{A}_{pa} = 0.999$ and $\mathcal{A}_{pb} = 0.990$ with ϵ ranging from 0.0001 to 0.0010.

Improving path availability

In Fig. 6, we show how the path availability is improved by factor ϵ. Note that the results are similar to those in Fig. 1.

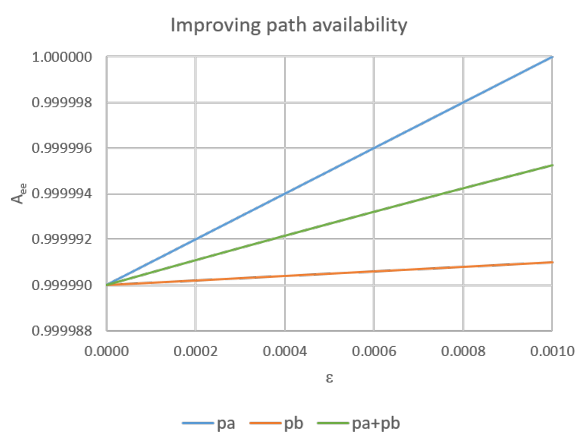

Figure 6: Impact of path availability improvement on the end-to-end availability

Improving availability in all links

Fig. 7 illustrates the impact of ϵ on the cost functions c_1 (top), c_2 (middle) and c_3 (bottom) for three Set 1 (left), Set 2 (center) and Set 3 (right). The blue line refers to all links of p_a being improved to have an improvement on $\mathcal{A}_{p_a}$ of ϵ, the orange line all links of p_b being improved to have an improvement on $\mathcal{A}_{p_b}$ of ϵ, and the grey line all the links of the path pair being improved to have an improvement on $\mathcal{A}_{p_a}$ and $\mathcal{A}_{p_b}$ of $\epsilon/2$. The results are similar to that shown in Fig. 2.

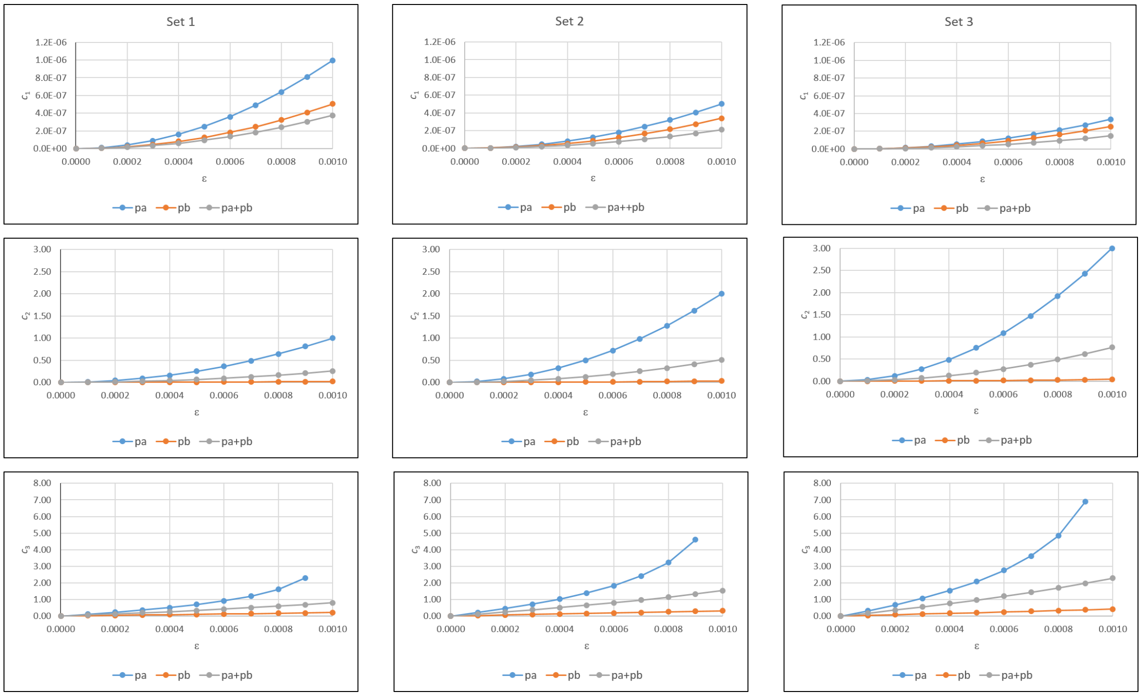

Figure 7: Impact of ϵ on the cost functions c_1 (top), c_2 (middle) and c_3 (bottom) when improving all links of p_a or p_b, or improving all links of both paths

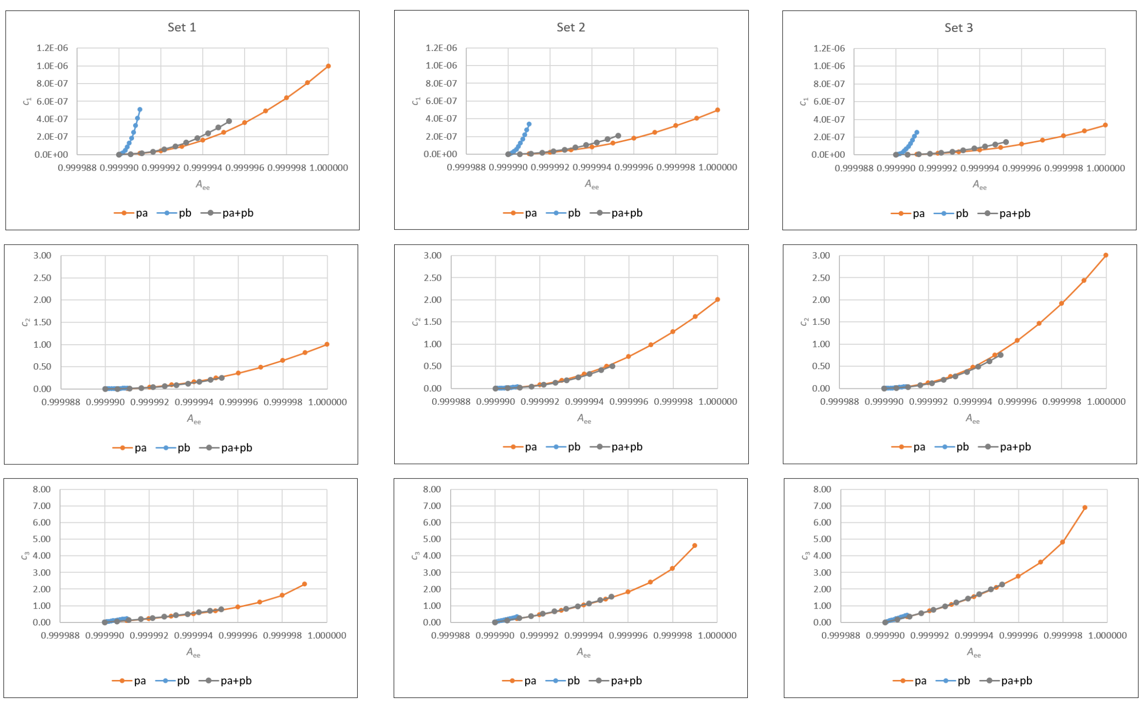

Figure 8: Impact of A_{ee} on the cost functions c_1 (top), c_2 (middle) and c_3 (bottom) when improving all links of p_a or p_b, or improving all links of both paths

Fig. 8 illustrates the impact of the achieved $\mathcal{A}_{ee}$ (for each value of ϵ in the previous figure) on the cost functions c_1 (top), c_2 (middle) and c_3 (bottom) for Set 1 to Set 3. The results are similar to that in Fig. 3.

Improving availability in one link

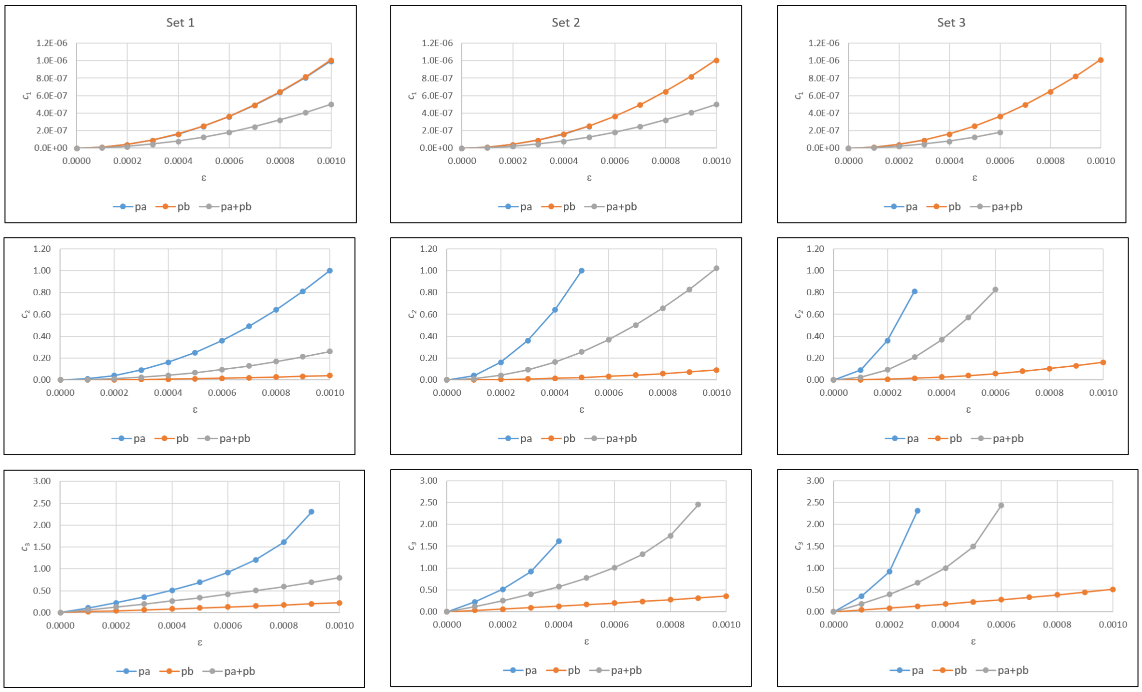

Figure 9: Impact of ϵ on the cost functions c_1 (top), c_2 (middle) and c_3 (bottom) when improving one link of p_a or p_b, or improving one link of each path

Fig. 9 illustrates the impact of ϵ on the cost functions c_1 (top), c_2 (middle) and c_3 (bottom) for Set 1 (left), Set 2 (center) and Set 3 (right). The blue line refers to one link of p_a being improved to have an improvement on $\mathcal{A}_{p_a}$ of ϵ, the orange line one link of p_b being improved to have an improvement on $\mathcal{A}_{p_b}$ of ϵ, and the grey line one link of each path being improved to have an improvement on $\mathcal{A}_{p_a}$ and $\mathcal{A}_{p_b}$ of $\epsilon/2$. The results are similar to those in Fig. 4.

Fig. 10 illustrates the impact of the achieved $\mathcal{A}_{ee}$ (for each value of ϵ in the previous figure) on the cost functions c_1 (top), c_2 (middle) and c_3 (bottom) for Set 1 to Set 3. The results are similar to those in Fig. 5.

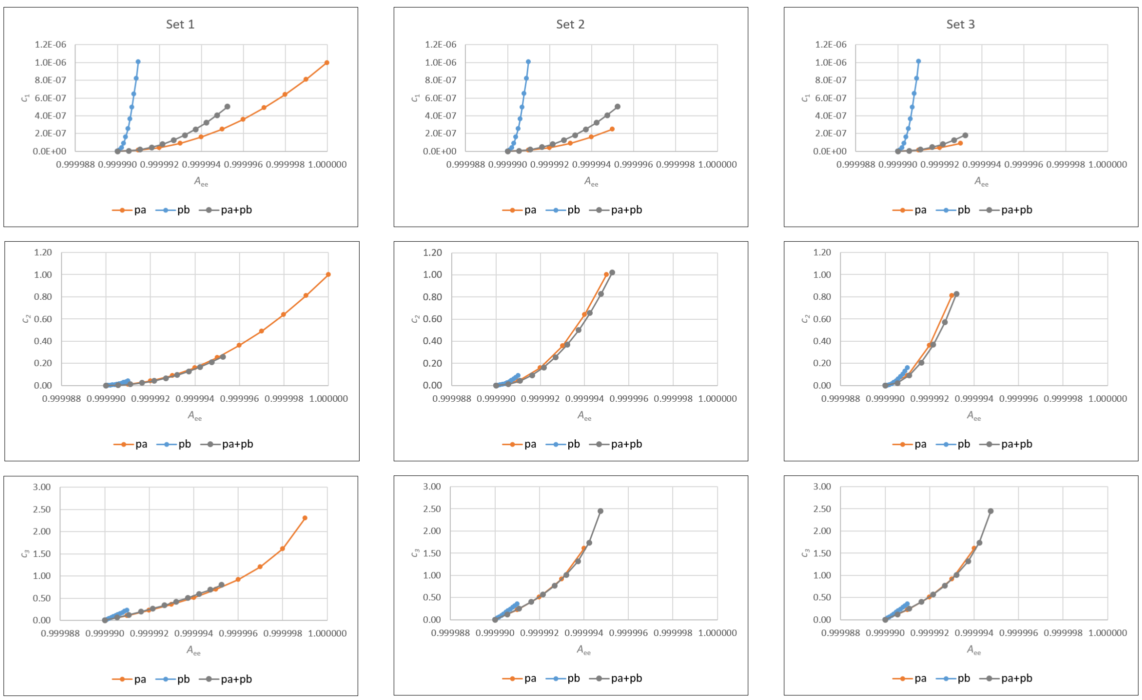

Figure 10: Impact of A_{ee} on the cost functions c_1 (top), c_2 (middle) and c_3 (bottom) when improving one link of p_a or p_b, or improving one link of each path