



HAL
open science

Reconfiguration de chaînes de fonctions de services sans interruption

Adrien Gausseran, Andrea Tomassilli, Frédéric Giroire, Joanna Moulierac

► To cite this version:

Adrien Gausseran, Andrea Tomassilli, Frédéric Giroire, Joanna Moulierac. Reconfiguration de chaînes de fonctions de services sans interruption. CORES 2019 - Rencontres Francophones sur la Conception de Protocoles, l'Évaluation de Performance et l'Expérimentation des Réseaux de Communication, Jun 2019, Saint Laurent de la Cabrerisse, France. hal-02118989

HAL Id: hal-02118989

<https://hal.science/hal-02118989>

Submitted on 3 May 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Reconfiguration de chaînes de fonctions de services sans interruption[†]

A. Gausseran¹ et A. Tomassilli¹ et F. Giroire¹ et J. Moulrierac¹

¹Université Côte d'Azur, I3S (UNS-CNRS), Inria Sophia Antipolis, France

Ce travail vise à montrer l'utilité de reconfigurer des chaînes de fonctions de services (SFC) dans le but d'améliorer le coût opérationnel du réseau. Nous proposons un modèle d'optimisation basé sur le mécanisme *make-before-break*, dans lequel l'ancien chemin n'est détruit que quand le nouveau chemin est complètement opérationnel, ceci afin de ne pas avoir d'interruption du trafic. Nous montrons qu'avec notre approche, nommée *break-free*, le coût opérationnel du réseau est réduit tout en augmentant le taux d'acceptation des SFC.

Mots-clefs : Reconfiguration, SDN, SFC, NFV

1 Introduction & State of the art

Network flows are often required to be processed by an ordered sequence of network functions. This notion is known as Service Function Chaining (SFC) [8]. The application of SFCs often requires the usage of two different independent technologies, Software Defined Networking (SDN) and Network Function Virtualization (NFV), which are complementary.

SDN aims at simplifying network management by decoupling the control plane from the data plane, making it easier to manipulate network devices through a software program [4]. As a consequence, the network becomes programmable. With the NFV paradigm, network functions can be implemented in software, becoming Virtual Network Functions (VNFs) that can be instantiated and scaled on-demand without having to install new equipment and executed on generic-purpose servers located in small cloud nodes. Since traffic is dynamic, the allocation of a demand is performed individually without having full knowledge of the incoming traffic. This may lead to a fragmented network which uses more resources than necessary. Also, it may lead to a higher blocking probability even though there are enough resources to serve new demands. Therefore, we must take it into consideration and adjust network configurations in response to changing network conditions. Thus, our problem is to reroute traffic flows through the network and to improve the mapping of network functions to nodes in the presence of dynamic traffic. Our objective is to minimize the network operational cost, defined as the sum of the link bandwidth cost to route the demands and the cost for all the installed VNFs in the network.

Rerouting demands and migrating VNFs may take several time steps. If during this time, traffic is interrupted, it may have a non-negligible impact on the QoS. To tackle this issue, our strategy performs the reconfiguration by using a two-phase approach. First, a new route for the transmission is established while keeping the initial one enabled (i.e., two redundant data streams are both active in parallel). When the network is at the new state, the transmission moves to the new route and the resources used by the initial one are released. This strategy is often referred to as *make-before-break*.

State of the art. Paris *et al.* [7] study the problem of online SDN controllers to decide when to perform flow reconfigurations, but network function requirements are not considered in their work. The closest study to our work is from Liu *et al.* [5]. They consider the problem of optimizing VNF deployment and readjustment to efficiently orchestrate dynamic demands. But an important unaddressed issue concerns the

[†]This work has been supported by the French government through the UCA JEDI (ANR-15-IDEX-01) and EUR DS4H (ANR-17-EURE-004) Investments in the Future projects, and by Inria associated team EfDyNet.

revenue loss of an operator due to the QoS degradation occurring when demands are reconfigured. Duong *et al.* [2] apply this reconfiguration strategy in optical networks using column generation, but they only allow one request to be moved at a given time.

2 Problem Statement and Modeling

We consider that the flows are splittable as it is frequent to have load balancing in networks [1]. We consider the SFCs one by one, and find a route which minimizes the *additional network operational cost* to be paid. We then reconfigure the network with a *make-before-break* mechanism when one of the following conditions holds: periodically, after a given period of time; when the set of requests has changed significantly (after a given number of SFC arrivals and departures); when a request arrives and cannot be accepted with the current provisioning and routing solution.

We compare the results of **Break-Free** with **Breaking-Bad** that breaks the flows before rerouting them, implying packet losses and QoS degradation. When a reconfiguration has to be carried out, **Breaking-Bad** considers basically a static setting with the requests present in the network and finds an optimal Routing & Provisioning solution without considering the current setting.

Presentation of Break-Free: reconfiguration without interruption. Our idea consists in improving the current solution, getting as close as possible to the optimal one while being achievable by a make-before-break reconfiguration. To this end, we set a number of intermediate reconfiguration steps, T , and the goal of the optimization is to find a routing with minimal cost, which can be reached from the current routing using T reconfiguration steps. At time 0, the configuration is the current one (i.e., paths and function locations for all the SFCs). Then, at each step of reconfiguration, a valid solution (respecting the link and node capacities) is computed, and some SFCs may have two active configurations. As a single step of reconfiguration may not be enough, the ILP has several intermediate reconfiguration steps. The objective function is to minimize the network operational cost of the final state. The value of T is an important parameter. Indeed, a value too small may lead to models with poor solutions, while a value too large to models with prohibitive execution times. The model is based on the notion of layered graph with $|c_d|$ layers (where $|c_d|$ denotes the number of VNFs in the chain) as presented in [3]. The variables are the following ones.

- $\varphi_{uv,i}^{d,t} \geq 0$ is the amount of flow on Link (u, v) in Layer i at Time step t for Demand d .
- $\alpha_{u,i}^{d,t} \geq 0$ is the fraction of flow of Demand d using Node u in Layer i at Time step t .
- $x_{uv,i}^{d,t} \geq 0$ is the maximum amount of flow on Link (u, v) in Layer i at Time steps t and $t - 1$ for Demand d .
- $y_{u,i}^{d,t} \geq 0$ is the maximum fraction of flow of demand d using Node u in Layer i at Time steps t or $t - 1$.
- $z_{u,f}^T \in \{0, 1\}$, where $z_u^f = 1$ if function f is activated on Node u at Time step T in the final routing.

The initial configuration is given, meaning that, for each SFC $d \in D$, the variables $\varphi_{uv,i}^{d,0}$ and $\alpha_{u,i}^{d,0}$ are known. $c_{u,f}$ is the installation cost of the function $f \in F$ which also depends on the node u . Each demand $d \in D$ is modeled by a quadruple with v_s the source, v_d the destination, c_d the ordered sequence of network functions that need to be performed, and bw_d the required units of bandwidth. The parameter β specifies the balance between bandwidth and VNFs activation cost.

Objective:

$$\min \sum_{d \in D} \sum_{(u,v) \in E} \sum_{i=0}^{|c_d|} bw_d \cdot \varphi_{uv,i}^{d,T} + \beta \cdot \sum_{u \in V} \sum_{f \in F} c_{u,f} \cdot z_{u,f}^T$$

Flow conservation constraints. For each SFC $d \in D$, node $u \in V$, Time step $t \in \{1, \dots, T\}$.

$$\sum_{(u,v) \in \omega^+(u)} \varphi_{uv,0}^{d,t} - \sum_{(v,u) \in \omega^-(u)} \varphi_{vu,0}^{d,t} + \alpha_{u,0}^{d,t} = \begin{cases} 1 & \text{if } u=v_s \\ 0 & \text{else} \end{cases}$$

$$\sum_{(u,v) \in \omega^+(v)} \varphi_{uv,|c_d|}^{d,t} - \sum_{(v,u) \in \omega^-(v)} \varphi_{vu,|c_d|}^{d,t} - \alpha_{u,|c_d|-1}^{d,t} = \begin{cases} -1 & \text{if } v=v_d \\ 0 & \text{else} \end{cases}$$

Reconfiguration de chaînes de fonctions de services sans interruption

$$\sum_{(u,v) \in \omega^+(u)} \varphi_{uv,i}^{d,t} - \sum_{(v,u) \in \omega^-(u)} \varphi_{vu,i}^{d,t} + \alpha_{u,i-1}^{d,t} - \alpha_{u,i}^{d,t} = 0 \quad \text{for } 0 < i < |c_d|$$

Node usage over two consecutive time periods. For each $d \in D$, $u \in V$, $t \in \{1, \dots, T\}$ and layer $i \in \{0, \dots, |c_d| - 1\}$.

$$\alpha_{u,i}^{d,t} \leq y_{u,i}^{d,t} \quad \text{and} \quad \alpha_{u,i-1}^{d,t-1} \leq y_{u,i}^{d,t} \quad \text{and} \quad \alpha_{u,i}^{d,t} + \alpha_{u,i-1}^{d,t-1} \geq y_{u,i}^{d,t}$$

Link usage over two consecutive time periods. For each $d \in D$, link $(u, v) \in E$, layer $i \in \{0, \dots, |c_d|\}$, $t \in \{1, \dots, T\}$.

$$\phi_{uv,i}^{d,t} \leq x_{uv,i}^{d,t} \quad \text{and} \quad \phi_{uv,i}^{d,t-1} \leq x_{uv,i}^{d,t} \quad \text{and} \quad \phi_{uv,i}^{d,t} + \phi_{uv,i}^{d,t-1} \geq x_{uv,i}^{d,t}$$

Make Before Break - Node capacity constraints. The capacity of a node $u \in V$ is shared between each layer and cannot exceed C_u considering the resources used over two consecutive time periods. Δ_f represents the number of units of bandwidth required by the function $f \in F$. For each $u \in V$, $t \in \{1, \dots, T\}$.

$$\sum_{d \in D} bw_d \sum_{i=0}^{|c_d|-1} \Delta_{f_i}^{c_d} y_{u,i}^{d,t} \leq C_u$$

Make Before Break - Link capacity constraints. The capacity of a link $(u, v) \in E$ is shared between each layer and cannot exceed C_{uv} considering the resources used over two consecutive time periods. For each $(u, v) \in E$, $t \in \{1, \dots, T\}$.

$$\sum_{d \in D} bw_d \sum_{i=0}^{|c_d|} x_{uv,i}^{d,t} \leq C_{uv}$$

Location constraints. A node may be enabled to run only a subset of the virtual network functions. For each $d \in D$, $u \in V$, layer $i \in \{0, \dots, |c_d| - 1\}$, if the $(i+1)$ -th function of c_d cannot be installed on Node u , we add the following constraint for each time step $t \in \{1, \dots, T\}$.

$$\alpha_{u,i}^{d,t} = 0$$


Functions activation. To know which functions are activated on which nodes in the final routing. For each $u \in V$, function $f \in F$, demand $d \in D$, and Layer $i \in \{0, \dots, |c_d| - 1\}$.

$$\alpha_{u,i}^{d,T} \leq z_{u,f_i}^{c_d}$$

3 Numerical Results

We study the impact of the reconfiguration by comparing the results of `Break-Free` (for different numbers of reconfiguration steps from 1 to 3) with the ones of `Breaking-Bad`, and `No-Reconf` which never reconfigures the SFC. We consider two scenarios: the first one with low traffic in which we plot the network operational cost, and the second one, with high traffic scenario, all the demands cannot be accepted, and we study the acceptance rate of demands. The experiments are conducted on a real-world topology from SNDlib [6]: `tal` (24 nodes, 55 links). We considered 250 demands for each network. The source and destination of each demand are chosen uniformly at random among the nodes. Following [9], the lifetime of a demand is exponentially distributed (rounded to an integral number of time steps). When the lifetime is reached, the demand leaves the network. The volumes of the demands are chosen randomly and are on average 2 times higher in the high-traffic scenario than in the low-traffic one. Also, each demand is associated with an ordered sequence of 2 to 3 functions uniformly chosen at random from a set of 5 different functions. The demands are routed one by one, greedily, when they arrive.

Low-traffic scenario - Network cost. In Figure 1a, we first see that `Break-Free` has similar performances to `Breaking-Bad`. Recall that `Breaking-Bad` provides a lower bound for our algorithm as it interrupts the SFCs. Moreover, `Break-Free` achieves this performance for any number of time steps (even 1). Indeed, when the network is not congested, there is enough capacity to host both the old and new routes and to perform efficient reconfiguration. Reconfiguration leads to a better resource utilization compared to `No-Reconf`. Indeed, reconfiguring regularly the network permits a reduction of 19% of network operational cost while using 22.5% fewer VNFs and 18.5% less link bandwidth. In this scenario there are on average 25 SFCs on the network at the same time and since there is no congestion the executions are fast for the 150 SFCs : from 9 seconds for `No-Reconf` to 12 seconds for `Break-Free` 3 steps.



(a) Low-Traffic scenario - Network operational cost. (b) High-Traffic scenario - Percentage of accepted profit.

FIGURE 1: Results for τ_{a1} for two different cases.

High-Traffic scenario - Acceptance profit rate. The network is congested as there are not enough resources to satisfy all the SFCs. As a consequence, some demands are rejected. We show the percentage of accepted profit achieved by the three algorithms in Figure 1b. We define the profit associated with an accepted demand as the asked volume of bandwidth multiplied by its duration. We show the profit as a percentage in terms of maximum achievable profit. It can be seen that *No-Reconf* and *Break-Free* (with 1-step) lead to equivalent profit (around 81%), while *Break-Free* (with 2 and 3 steps) and *Breaking-Bad* have similar performances (around 90%). For this congested scenario, one step of reconfiguration is not enough as there is not enough space to move the requests, thus the requests are rejected. If we increase the number of steps of reconfiguration for *Break-Free*, we can reach the same performances as *Breaking-Bad*. In this scenario there are on average 35 SFCs on the network at the same time and since there is a lot of congestion the executions are slower and vary more for the 250 SFCs : from 15 seconds for *No-Reconf* to 80 seconds for *Breaking-Bad* to more than 2000 for *Break-Free* 3 steps.

As a conclusion of these results, *Break-Free* allows to lower the network cost and to increase the acceptance rate almost as much as *Breaking-Bad*. Reconfiguration has to be done often to attain a significant gain showing that the make-before-break mechanism is crucial to avoid an impact on the QoS. Further simulation results are given in [3], such as the impact of the frequency of reconfiguration or the impact of setting a time limit for the resolution of the linear program.

4 Conclusion

In this work, we provide a solution, *Break-Free*, to reconfigure a set of requests which have to go through service function chains. The requests are routed greedily when they arrive, leading to a sub-optimal use of network resources, bandwidth, and virtual network functions. *Break-Free* reroutes the requests to an optimal or close to optimal solution while providing a make-before-break mechanism to avoid impacting the rerouted requests. Our algorithm is fast and provides a close to optimal reconfiguration in a few seconds. However, as a future work, we plan to propose algorithms to handle larger instances.

Références

- [1] B. Augustin, T. Friedman, and R. Teixeira. Measuring load-balanced paths in the internet. In *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, 2007.
- [2] H. Duong, B. Jaumard, D. Coudert, and R. Armolavicius. Efficient make before break capacity defragmentation. *IEEE International Conference on High Performance Switching and Routing*, 2018.
- [3] A. Gausseran, A. Tomassilli, F. Giroire, and J. Moulierac. Don't interrupt me when you reconfigure my service function chains. Technical report, Inria, Dec. 2018.
- [4] H. Kim and N. Feamster. Improving network management with software defined networking. *IEEE Com Magazine*, 2013.
- [5] J. Liu, W. Lu, F. Zhou, P. Lu, and Z. Zhu. On dynamic service function chain deployment and readjustment. *IEEE Transactions on Network and Service Management*, 14(3) :543–553, 2017.
- [6] S. Orłowski, R. Wessälly, M. Pióro, and A. Tomaszewski. Sndlib 1.0—survivable network design library. *Wiley Networks*, 2010.
- [7] S. Paris, A. Destounis, L. Maggi, G. S. Paschos, and J. Leguay. Controlling flow reconfigurations in sdn. In *IEEE INFOCOM*, 2016.
- [8] P. Quinn and T. Nadeau. Problem statement for service function chaining. Technical report, 2015.
- [9] S. Saha, W. Tavernier, M. Rost, S. Schmid, D. Colle, M. Pickavet, and P. Demeester. Network service chaining with optimized network function embedding supporting service decompositions. *Computer Networks*, 93 :492–505, 2015.