

Routing-Led Placement of VNFs in Arbitrary Networks

Joseph Billingsley*, Ke Li*, Wang Miao*, Geyong Min* and Nektarios Georgalas†

* Department of Computer Science, University of Exeter
Exeter, UK

Email: {jb931, k.li, wang.miao, g.min}@exeter.ac.uk

† Research and Innovation, British Telecom
Martlesham, UK

nektarios.georgalas@bt.com

Abstract—The ever increasing demand for computing resources has led to the creation of hyperscale datacentres with tens of thousands of servers. As demand continues to rise, new technologies must be incorporated to ensure high quality services can be provided without the damaging environmental impact of high energy consumption. Virtualisation technology such as network function virtualisation (NFV) allows for the creation of services by connecting component parts known as virtual network functions (VNFs). By optimising the placement and routing of VNFs this technique can be used to maximally utilise available datacentre resources, to maintain a high quality of service whilst minimising energy consumption. Current research on this problem has focussed on placing VNFs and considered routing as a secondary concern. In this work we argue that the opposite approach, a routing-led approach is preferable. We propose a novel routing-led algorithm and analyse each of the component parts over a range of different topologies on problems with up to 16000 variables and compare its performance against a traditional placement based algorithm. Empirical results show that our routing-led algorithm can produce significantly better solutions to large problem instances on a range of datacentre topologies.

Index Terms—Network Function Virtualisation, Multi-Objective Optimisation, Routing-Led VNF Placement.

I. INTRODUCTION

The internet and the services it has enabled have changed how we work [1], travel [2], socialise [3] and even how we engage with politics [4], government [5] and the rest of society [6]. The infrastructure that provides these services is invisible to most people. One critical part of this infrastructure is the datacentre which provide the computing power needed to provide services. Demand for datacentre resources has increased year-on-year and this trend is expected to continue [7]. New technology such as augmented and virtual reality [8], the internet of things [9] and smart cities [10] may have further positive impacts on society and businesses but will demand more from the underlying infrastructure. Existing techniques cannot meet these new demands without significant and damaging environmental impacts [11]. Instead, we must develop solutions that can efficiently use existing resources to allow for the continued development of high quality services

whilst minimising the energy consumption of the datacentres that provide them.

Resolving these two conflicting objectives requires a comprehensive solution that integrates recent hardware and software innovations. Innovations in the structure of the datacentre aim to improve its scalability. These techniques include new network topologies, the arrangement of switches and servers that allow servers to communicate with each other, that facilitate the addition of equipment to the datacentre. Innovation is also occurring at the software level. Previously, common tasks such as firewalls and deep packet inspection would be executed by specialised equipment known as network functions. In a virtualised datacentre these tasks are provided by Virtual Network Functions (VNFs), software running on virtual machines. VNFs can be created, moved and destroyed far faster than physical hardware.

In a modern virtualised datacentre, a service is provided by placing and connecting sequences of VNFs to form a service. The quality of a service (QoS) is determined by metrics such as latency - the expected time for a packet to traverse the service - and packet loss - the proportion of packets that do not complete the service due to an overuse of certain servers or switches. The QoS of each service must be maximised whilst also minimising the datacentre energy consumption. A solution to this problem defines the placement of VNFs and routes between VNFs to form a set of services that meet these conflicting objectives simultaneously. This is known as the VNF placement problem (VNFPP) and has been shown to be NP-Hard [12].

Due to the NP-Hardness of the problem, and the very high number of servers in modern datacentres, several heuristic and metaheuristic solutions have been proposed [12]–[21]. Typically, these approaches place a set of VNFs and then find routes between VNFs. These algorithms implicitly optimise routes by finding improving placements. In this work we categorise these approaches as *placement-led* optimisation techniques. However, key metrics such as latency, packet loss and energy are often minimised by using the shortest, least congested routes available [21]. Hence we should allow the

routes to dictate the placement. We categorise these as *routing-led* optimisation of the VNFPP.

To be able to solve large scale problems, a routing-led algorithm must be able to efficiently find nearby servers that can accommodate each VNF in the service. Some existing works can be classified as routing led. Some use knowledge of the structure of the graph to simplify placement [21], [22] which prevents the solution from solving arbitrary graphs. Others have considered all possible servers and selected the most suitable [18], [19] but this strategy is unsuitable for large datacentres. The work by Goureb *et al* [20] considered three possible routes between the first and last VNFs of a service and placed VNFs only on these routes. This does not guarantee a VNF will be placed when there is a server with sufficient capacity.

This paper proposes a novel metaheuristic to tackle the VNFPP at scale on arbitrary graphs. Our major contributions are as follows:

- We evaluate three strategies to select suitable servers in a scalable routing-led placement algorithm
- We propose a routing-led genetic algorithm approach to the VNFPP that uses these strategies to optimise large scale problems
- We compare our algorithm against a representative placement-led algorithm and evaluate their effectiveness on multiple network topologies with up to 16,000 servers.

The remainder of this paper is organised as follows. Section II reviews the current work on the VNFPP and on comparable metaheuristics in particular. Section III provides a formal definition of the VNFPP and the multiobjective formulation used in this work. Section IV describes the genetic algorithm and operators considered in this work including alternative server location strategies. In Section V the effectiveness of each component of the algorithm is tested and our solution is evaluated against a comparable placement-led approach. Finally, Section VI concludes this paper and outlines some potential future directions.

II. LITERATURE REVIEW

The VNF placement problem requires solving two problems simultaneously: VNFs must be placed and routes must be found to form services.

In many works linear programming techniques have been used [23]–[30] to solve both of these problems simultaneously. However, current implementations can only consider problem instances with hundreds of servers. For larger problem instances heuristic and metaheuristic approaches are more suitable. The work by Luizelli *et al* [12] proposed to combine Variable Neighbourhood Search (VNS) with a linear programming solver. VNS selects a subset of variables that are considered by the linear programming solver whilst the other variables remain fixed. This approach solves the largest problem instances we are aware of in the literature but it is still time consuming requiring 10,000 seconds (~2.5 hours) to solve a problem instance with 1000 servers. Cao *et al* [13] used NSGA-II [31] and MOGA [14] with a binary matrix solution

representation which represented both placement and routing instructions. In this solution representation most of the search space is infeasible and hence custom initialisation and repair operators were used to repair invalid solutions.

Some works do not consider the routing directly and instead consider it as a function of the placement. Leivadreas *et al* [15] use a Tabu search metaheuristic to place each VNF. The shortest routes between VNFs in a service are then used. Similarly, Rankothge *et al* [16] considered problems with up to 400 VNFs, using a Genetic Algorithm with problem specific operators to place VNFs and a depth first search heuristic to find routes between VNFs in a service. In [17], Jiang *et al* used a local search heuristic to find improving solutions that only require one VNF migration. A greedy search heuristic then considers new positions for the VNF and favours positions that can be accessed with low congestion routes. In each of these works, the placement determines the routing. The algorithm implicitly optimises routes by finding improving placements.

An alternative approach is to allow the routing to decide the placement of VNFs. One way to apply this principle is to represent the possible placements of VNFs as nodes in a multistage graph where each edge has some cost assigned to it determined by the route between the two current and next VNF. A solution to the problem is a path through the graph. Bari *et al* [18] first proposed this representation and used a dynamic programming heuristic to solve VNFPP problem instances, aiming to minimise the ‘traffic forwarding cost’ calculated based on the distance between two nodes and the quantity of traffic that was sent between them. Lange *et al* [19] used the same representation and costs with the Pareto Simulated Annealing algorithm and modified the neighbourhood generating step. In their variant, neighbouring solutions are generated by probabilistically exploring the multistage graph with a greater chance of selecting paths that minimise cost. Notably, both of these works consider every possible server at each step. On large problems with many servers this can be a time consuming operation.

More recently, Gouareb *et al* [20] utilised a similar routing led approach in their work. The authors assumed the starting and end points of the service are fixed and found the three shortest routes between these two points. They then optimise the placement of the VNFs for each service whilst ensuring they can only be placed on nodes on these paths. Our earlier work on the VNF Placement Problem used a similar approach [21]. We utilised the layout of the fat tree topology to have a greater likelihood of placing VNFs on the nearest servers. However, this technique is topology specific and will not apply to arbitrary graphs.

Most existing work on the VNF placement problem is limited to solving small scale problem instances. Scalable heuristic and metaheuristic approaches that consider the placement directly have been proposed which implicitly optimise routes between VNFs. Existing routing-led solutions to the VNF placement problem require fixed starting or end points and use techniques that have not been shown to scale to large problems. This work provides a scalable routing-led solution to

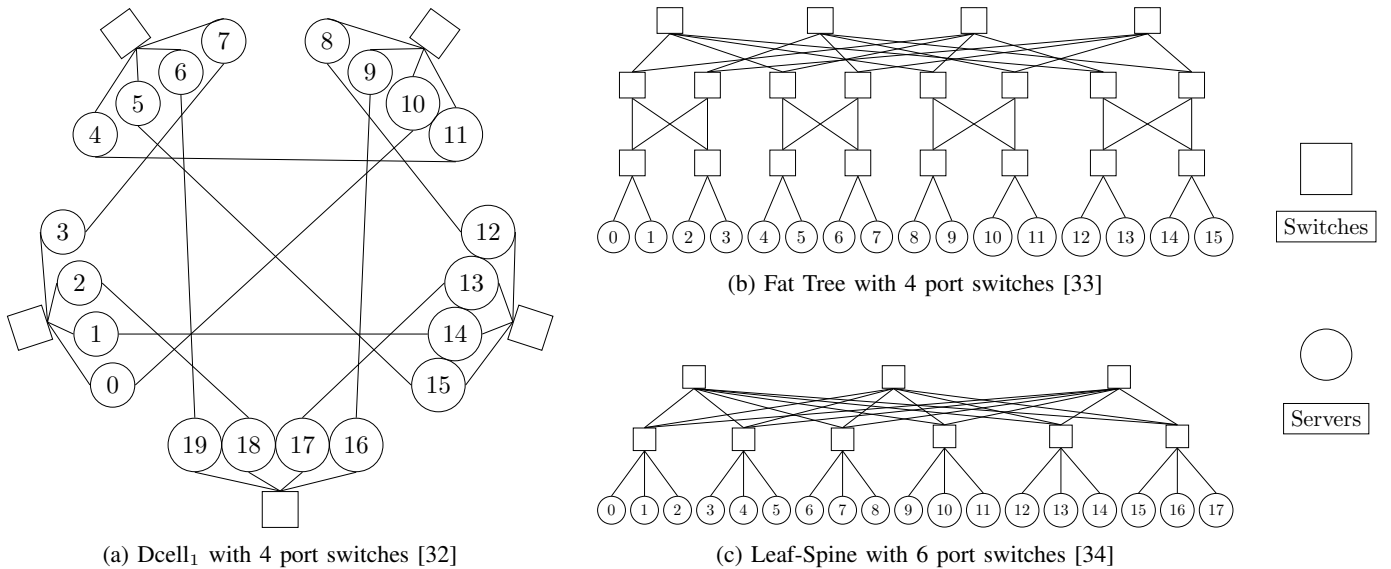


Fig. 1: Common datacentre topologies

the VNF placement problem without the limitations of earlier works.

III. PROBLEM FORMULATION

In the VNF placement problem, the network is represented as a graph, $G = (N, E)$ where N denotes the set of network nodes and E represents the set of links connecting the nodes. The node in the graph can be either a physical network device, e.g. switches, router, gateway, or a server capable of running virtual machines. Let S_i be the i th server in the topology. Multiple VNFs form a service chain to provide a service. Let $V_{i,j}$ denote the j th VNF in the i th service chain. Multi-path routing strategies, e.g. equal-cost multi-path routing (ECMP), are used in datacentre networks to improve transmission efficiency and reliability. Let a path be a sequence of graph nodes connected by edges. A solution to the VNF placement problem is an assignment of VNFs to server nodes and a collection of paths for each service. Let $P_{V_i}^k$ be the k th path of the i th service in the solution. Let A_{S_i} denote the VNFs assigned to the i th server. Finally, let C_{S_i} be the available capacity of the i th server and $C_{V_{i,j}}$ be the required capacity of the j th VNF of the i th service.

The placement and routing parts of the solution are connected by their constraints: (1) the required capacity of the VNFs assigned to a server cannot exceed the available capacity and (2) an instance of each VNF in the service must be visited in the route. Formally these constraints can be defined as:

- 1) $\sum_{v \in A_i} C_v \leq C_{S_i}$
- 2) $\forall v \in V_i, \exists s \in P_{V_i}^k$ s.t. $A_s \supseteq v$

Several models exist to evaluate the quality of a solution to the VNFPP. In this work we will use a queueing theory model proposed in [21] and extended in [35] that estimates the expected latency and packet loss of each service and the total energy consumption of the datacentre. As these objectives can

conflict, we formulate this as a multi-objective optimisation problem with three objectives: the average service latency, the average service packet loss and the total energy consumption of the datacentre.

IV. ROUTING-LED OPTIMISATION OF THE VNFPP

A routing-led optimisation of the VNFPP aims to place sequential VNFs of a service on nodes with short paths between them. In this work, we design a set of operators to be used with multiobjective genetic algorithm algorithms. We first design a solution representation and operators that optimise the placement of the first VNFs of each service. Next we propose three routing-led server selection strategies that use information about the network topology to place the remaining VNFs in each service. Finally we find a set of routes between the VNFs for a service based on the ECMP routing protocol to form a complete solution.

A. Service Origin Operators

The building block hypothesis of genetic algorithms [36] indicates that good solutions are constructed from good ‘building blocks’ and hence dependent components should be exchanged during crossover. In the VNFPP, a single service can maximise its expected service quality by selecting paths that are not being utilised by other services. However, due to limited datacentre resources and energy concerns, it is not typically possible for each service to have dedicated resources. Further, it is likely for two services that start from close servers will use the same nodes and hence share resources. Good solutions to the VNFPP will minimise resource conflicts between services whilst balancing service quality and energy consumption. We should therefore design our solution representation and operators so that the distance between services is maintained during crossover.

In hierarchical networks such as the Fat Tree or Leaf-Spine datacentre topologies (see Fig. 1) servers can be encoded as a string whilst maintaining the relative distances between servers. However, this only applies to a subset of topologies. For example there is no string encoding of the Dcell representation that maintains the relative distances between all servers. Variable length string (VLS) representations and operators also use strings to encode solutions but are designed to find and exchange related groups of variables during optimisation [37]. Rather than specifying a suitable solution representation for each topology a VLS representation may find suitable groups of services during the optimisation.

As we are considering a range of topologies, we experiment with both of these solution representations. In the fixed length string representation, each server is represented by an array that contains the ID of the services that should originate from there. In the VLS representation, each character is a tuple of the service ID and the server ID it should originate from. For a fair comparison we use similar genetic operators for both representations: single point crossover and uniform mutation for the fixed length string representation and messy crossover [37] for the VLS. Uniform initialisation is used in both cases. Server IDs are assigned as shown in Fig. 1. Notably this arrangement of IDs ensures some distance information is maintained in the fixed length string representation.

B. Server Selection Strategies

To evaluate a solution, the remaining VNFs in the service must be placed. Each VNF should be placed on the nearest available server with sufficient capacity. This operation must be efficient as it will be performed many times per solution evaluation. We evaluate three different approaches to this problem a simple breadth first search algorithm, a cached variant which stores information on distance and a novel spanning tree based strategy.

1) *Simple BFS*: A breadth first search (BFS) will progressively consider the closest nodes. The search procedure can stop as soon as a server is found with enough capacity and has a worst case time complexity of $O(|N| + |E|)$.

2) *Cached BFS*: A BFS will expand the same servers many times during optimisation. Instead a BFS can be performed from every server before the optimisation begins and the sequence of nearest servers can be stored for each server. This has a preprocessing time complexity of $O((|N| + |E|) \cdot |S|)$ operations and a memory complexity of $O(|S|^2)$. Subsequent operations have a worst case time complexity of $O(|S|)$.

3) *Spanning Tree Search*: On very large networks executing a BFS for every service is wasteful and the preprocessing time and memory requirements of the ‘Cached’ algorithm may be very high. We propose an alternative technique that combines the simple search and cached approaches to find feasible solutions quickly with lower preprocessing time and memory complexity.

In this technique the network topology graph is first converted into a spanning tree ST . A table is constructed for each node in the spanning tree which tracks the distance to, and the

Algorithm 1 Update spanning tree search tables

```

1: procedure UPDATE( $ST, start, capacity$ )
2:   let  $Q$  be a queue
3:   label  $start$  as discovered
4:    $Q.enqueue((start, start, capacity, 0))$ 
5:   while  $Q$  is not empty do
6:      $(v, p, c, d) \leftarrow Q.dequeue()$ 
7:      $v.rowFor(p).cap \leftarrow c$            ▷ Set row capacity
8:      $v.rowFor(p).dist \leftarrow d$        ▷ Set row distance
9:     find the best ( $b$ ) and second best ( $s$ ) rows in  $v$ 
10:    if the best row has changed then
11:      for  $n$  in  $ST.adjacentEdges(v)$  do
12:        if  $n$  is discovered then
13:          continue
14:        label  $n$  as discovered
15:        if  $n = b.edge$  then
16:           $Q.enqueue(n, v, b.cap, b.dist+1)$ 
17:        else
18:           $Q.enqueue(n, v, s.cap, s.dist+1)$ 
19:      else if the second best row has changed then
20:        if  $b.edge$  is not discovered then
21:          label  $b.edge$  as discovered
22:           $Q.enqueue(n, v, s.cap, s.dist+1)$ 

```

capacity of, the nearest server with sufficient capacity. The tables are initialised by performing a BFS from each server over the spanning tree. When a node is added to the search front, the algorithm checks whether the source server is better than the server currently recorded for the current edge and updates the edge if this is true. If a table is not updated the search does not need to propagate to the children of the node. This has a worst case time complexity of $O((|N| + |E|) \cdot |S|)$.

To place a VNF the placement algorithm greedily takes the nearest edge with sufficient capacity. As a spanning tree is fully connected and has no loops, exactly one edge for each node will lead to a server with the maximum capacity and each edge will lead to a distinct server. Additionally, as the algorithm always takes steps towards the closest server, the placement step will never revisit a node. Hence it is clear that the algorithm will always find a server that can accommodate the current VNF if such a server exists.

Once a VNF has been placed we must update all tables that reference the changed server to correct their capacity. To do so we propagate information on the current best server the node is aware of from the changed node to adjacent nodes recursively. Similar to the construction process we can perform this efficiently with a BFS as a node must only propagate information if its best or second best row was changed as a result of the update. As the ‘best’ server for a table has the highest capacity and is the closest, tables can have different best nodes which minimises the number of updates needed. In particular, most messages will not propagate past the graphs

TABLE I: Time taken preprocessing, preparing a solution for evaluation and total for 10,000 solutions

	S	Preprocessing (s)		Mean Solution Preparation (ms)			Total for 10000 Solutions (s)		
		Cached	Spanning	Simple	Cached	Spanning	Simple	Cached	Spanning
Dcell	420	0.00	0.00	0.07	0.02	0.14	0.69	0.20	1.38
	930	0.00	0.00	0.17	0.05	0.29	1.71	0.52	2.94
	8190	373.00	0.00	1.76	0.62	2.09	17.55	379.22	20.90
	17556	3586.00	0.00	3.99	1.19	4.69	39.88	3597.86	46.94
	33306	—	1.00	8.06	—	9.30	80.61	—	93.97
	74256	—	5.00	21.68	—	22.96	216.83	—	234.60
Fat Tree	432	0.00	0.00	0.27	0.02	0.10	2.70	0.21	1.00
	1024	0.00	0.00	1.14	0.05	0.23	11.37	0.53	2.34
	8192	209.00	0.00	35.75	0.76	2.16	357.48	216.56	21.64
	16000	1361.00	0.00	83.99	1.17	3.55	839.86	1372.65	35.54
	27648	—	0.00	168.30	—	6.93	1682.96	—	69.31
	65536	—	2.00	403.05	—	18.88	4030.51	—	190.80
Leaf-Spine	392	0.00	0.00	0.16	0.02	0.09	1.58	0.19	0.86
	1152	0.00	0.00	0.67	0.06	0.26	6.66	0.63	2.59
	8192	111.00	0.00	12.26	0.60	1.93	122.59	117.00	19.30
	15488	717.00	0.00	36.93	1.16	3.86	369.25	728.62	38.61
	32768	—	0.00	143.32	—	9.74	1433.18	—	97.40
	64082	—	1.00	503.74	—	22.39	5037.42	—	224.87

central nodes, the nodes which minimise the greatest distance to all other vertices [38]. Due to its distance to other nodes, a central node will hold a reference to the best servers in the datacentre. Hence an update will only propagate past the node if it was from one of the best servers in the datacentre. In the worst case, placement will traverse the diameter of the spanning tree. If an update does pass a central node all tables will need to be updated as all tables will hold a reference to the changed server. Hence the worst case time complexity for the expansion is $O(d + |N| + |E|)$ where d is the diameter of the network. Pseudocode for the update step is given in Alg. 1.

C. Routing

To simplify the generation of solutions, the expansion step outputs a sequence of VNF placements but does not dictate how packets should be routed between them. As in ECMP, we split traffic over the set of shortest routes between nodes. This information can be stored efficiently by creating a set of condensed routing tables for each node in the datacentre topology. Each node stores the range of server IDs that are on the shortest path of each edge. A server ID can appear in multiple ranges if multiple shortest routes exist. The set of paths for the solution can be found by recursively following the instructions in the routing tables.

V. PERFORMANCE EVALUATION

Each algorithm was evaluated on three popular datacentre topologies: Fat Tree [33], Leaf-Spine [34] and Dcell [32]. Each of these datacentre topologies were designed for different use cases. The Fat Tree topology is designed to allow for equal

bandwidth on all routes in large datacentres using commodity switches. The Leaf-Spine topology is intended for datacentres with high inter-server traffic. The Dcell topology is designed to be robust against failures. Fat Tree and Leaf-Spine are hierarchical datacentre topologies whilst Dcell is a server centric topology. Groups of 12 tests were run in parallel using a CPU with 12 threads and 6 cores with a 3.4GHz base clock and 32Gb of RAM.

A. Server Selection Strategies

To evaluate the efficiency of each server selection strategy 1000 different solutions were generated for each network topology at six different scales from ~ 400 servers to $\sim 70,000$ servers. Services were added to each problem instance until at least 60% of available capacity was required to place all services. The average preprocessing and processing times are listed in Table I. The largest instances of the test problem for the ‘cached’ selection strategy required more memory than was available for testing and have been omitted.

As the cached strategy does not directly use the network topology after preprocessing, solution preparation time is comparable over each datacentre topology. Additionally the cached technique places services 4-5x faster than the other strategies. However, on large scale problems the preprocessing step is resource intensive, in particular memory intensive, which makes it unsuitable for large problem instances.

In contrast, the simple strategy is significantly affected by the choice of network topology. The simple selection strategy is able to solve large problem instances on the Dcell network topology in an acceptable amount of time but is an order of magnitude slower on both hierarchical topologies. As the

TABLE II: Mean hypervolume for each algorithm after 10000 evaluations

	$ S $	Placement Led			Fixed Length String			Variable Length String		
		IBEA	NSGA-II	MOEA/D	IBEA	NSGA-II	MOEA/D	IBEA	NSGA-II	MOEA/D
Dcell	420	† 0.298	† 0.298	† 0.300	0.755	0.766	0.742	† 0.706	† 0.699	† 0.691
	930	† 0.302	† 0.303	† 0.301	0.745	0.751	† 0.725	† 0.715	† 0.709	† 0.708
	8190	† 0.337	† 0.337	† 0.337	0.733	0.739	† 0.728	0.742	† 0.728	0.737
Fat Tree	432	† 0.419	† 0.417	† 0.403	0.661	0.667	† 0.639	† 0.621	† 0.602	† 0.606
	1024	† 0.457	† 0.457	† 0.453	0.695	0.701	† 0.670	† 0.670	† 0.651	† 0.658
	8192	† 0.511	† 0.511	† 0.511	0.728	0.734	† 0.719	0.734	† 0.711	0.726
	16000	—	—	—	0.749	0.749	† 0.740	† 0.741	† 0.717	† 0.733
Leaf-Spine	392	† 0.597	† 0.593	† 0.573	0.823	0.832	0.813	† 0.784	† 0.756	† 0.764
	1152	† 0.573	† 0.574	† 0.571	0.804	0.809	† 0.778	† 0.777	† 0.752	† 0.768
	8192	† 0.562	† 0.562	† 0.562	0.779	0.784	† 0.768	0.781	† 0.749	† 0.767
	15488	—	—	—	0.791	0.792	† 0.778	† 0.785	† 0.740	† 0.772

The highest mean hypervolumes before rounding are highlighted. † Denotes highlighted result is significantly better than the best value according to a Wilcoxon’s ranked sum test, $P < 0.05$. * Indicates which setting was used for significance tests where multiple best values exist.

simple selection strategy will stop as soon as it finds a server with sufficient capacity it benefits from densely connected server centric topologies where most steps will lead to a server. On hierarchical topologies such as Leaf-Spine and Fat Tree a BFS originating from a server will visit all switches in the network before it considers most of the servers.

The spanning strategy produces consistent results across each topology and can efficiently scale to large problem instances. However, it underperforms when compared to the cached strategy on smaller problems. As the spanning strategy must perform a BFS to update node tables after placing each VNF, its performance is comparable to the simple server selection strategy on problems with densely connected servers such as the Dcell. However, on hierarchical problems the spanning strategy can stop earlier more often as it is likely to stop at a central node.

As the preprocessed datacentre topology can be used for many tests, the cached selection strategy may be a suitable choice for small problem instances, problems where many function evaluations will be used or situations with very large amounts of resources. The simple strategy may be a good choice for placement algorithms designed to solve server centric topologies. For general purpose problems the spanning tree algorithm produces consistently good results and can scale to large problem instances.

B. Routing-Led Optimisation

The routing-led algorithms were next evaluated against a representative placement-led one. To provide a fair comparison, genetic operators were selected that are similar to those used in the fixed length string (FLS) routing-led algorithm. In the initialisation step, the initial solutions are placed as in the FLS initialisation operator but the remaining VNFs are assigned to subsequent servers in the string. The same mutation and crossover operators are used for the placement-

led algorithm as in the routing-led algorithms. Without a repair mechanism the placement-led algorithm, mutation and crossover rarely generate feasible solutions. A repair mechanism is applied that moves VNFs on overfilled servers to the nearest server in the string that can accommodate it. Paths are constructed as in the FLS operators but VNFs are selected by their distance in the string.

Each algorithm was used to solve 30 problem instances for each datacentre topology at different scales. The operators were used with the multiobjective evolutionary algorithms: NSGA-II [31], IBEA [39], and MOEA/D [40]. For each problem instance a set of services was generated that require at least 60% of the available capacity. In both the FLS and VLS representation, the cached server selection strategy was used to prepare a solution. The hypervolume (HV) metric was used to compare the overall results. As larger networks will host more services and hence use more energy we consider the expected energy contribution of each service rather than the total energy in the HV calculation. The nadir and utopian points were estimated using the worst and best values for each objective for all solutions and all objectives were normalised before calculating the HV. Table II shows the mean hypervolume from 30 problem instances after 10,000 iterations and Table III shows the execution time, excluding preprocessing. Memory constraints prevented the problem for being run on larger problem instances. Results on the largest instances of the Dcell and for the larger instances for the placement-led algorithm are not available as these tests did not terminate in an acceptable amount of time.

In all instances the routing-led algorithms significantly outperformed the placement-led algorithm. These results support the argument that routes should lead the placement of VNFs. Placing VNFs in a routing-led strategy increases the chance that the VNFs in each service are near to each other. A placement-led strategy does not have this benefit and requires

TABLE III: Mean execution time (s) for each algorithm and 10000 evaluations

	S	Placement Led			Fixed Length String			Variable Length String		
		IBEA	NSGA-II	MOEA/D	IBEA	NSGA-II	MOEA/D	IBEA	NSGA-II	MOEA/D
Dcell	420	† 14	† 15	† 20	† 8	† 7	† 10	† 8	† 6	† 10
	930	† 63	† 67	† 98	* 22	22	† 32	† 24	22	† 38
	8190	† 9577	† 11073	† 16437	2826	† 4093	† 4295	† 3726	3054	† 6544
Fat Tree	432	† 36	† 43	† 47	† 9	* 7	† 16	† 9	7	† 13
	1024	† 143	† 172	† 144	† 23	21	† 54	† 25	20	† 50
	8192	† 6461	† 6209	† 5542	555	† 717	† 1526	† 630	460	† 1734
	16000	—	—	—	† 2983	† 4229	† 4799	† 2696	† 2650	† 7009
Leaf-Spine	392	† 19	† 23	† 25	† 7	* 5	† 11	† 7	5	† 10
	1152	† 115	† 123	† 102	22	21	† 48	† 25	23	† 48
	8192	† 3972	† 3925	† 3547	476	† 744	† 1174	† 634	451	† 1671
	15488	—	—	—	† 2161	† 3082	† 3630	† 2209	† 1591	† 6607

The lowest mean times are highlighted. † Denotes highlighted result is significantly better than the best value according to a Wilcoxon’s ranked sum test, $P < 0.05$. * Indicates which setting was used for significance tests where multiple best values exist.

the evolutionary algorithm to improve routes via improving placements.

Notably the FLS representation significantly outperformed the VLS representation. As a string representation can still maintain some information on the distance between servers, it is likely that the FLS may result in good initial building blocks for the crossover. In comparison, the uniform initialisation used for the VLS does not reflect the relationship between servers immediately.

Results from Table III show the routing-led algorithm strategy was significantly faster than the placement-led strategy in this instance. As the placement-led algorithm has not been optimised for this problem it is not a fair comparison. However, the routing-led algorithms are three orders of magnitude faster than other metaheuristics that have considered problems of this scale in the literature (22-38 seconds vs 10,000 seconds for problems with ~1000 servers [12]) and the problems considered are significantly larger than ones in the literature so far. As the field has not yet converged on a problem definition for the VNFPP it is not possible to compare the quality of solutions directly.

Notably, Dcell requires an order of magnitude longer to optimise than other topologies. It is possible that this is due to the arrangement of the topology information in memory. The Fat Tree and Leaf-Spine topologies can be laid out in memory such that nearby servers in the topology are also near to each other in memory. This is not true for a Dcell topology as a one dimensional representation cannot fully capture information on the distance between servers (as discussed with respect to strings in Section IV). As a result, memory accesses in the Dcell topology are less predictable than memory accesses in hierarchical topologies. This factor should be considered in the design of future models and metaheuristics.

VI. CONCLUSIONS AND FUTURE WORK

Existing metaheuristic works on the VNF placement problem typically aim to place VNFs without considering the datacentre topology. An alternative approach places VNFs in a service so that they are joined by short routes. We classify this group of algorithms as routing-led techniques. Existing works that meet this classification require fixed starting and ending locations or only worked for specific topologies. In this work we proposed a new routing-led optimisation algorithm that does not have these limitations. The algorithm was tested against multiple topologies with datacentres containing up to 16,000 servers and 1000s of services. The routing-led algorithm was found to be significantly faster and produce better results than a comparable placement based algorithm on all test instances and shown to scale to problems much larger than have previously been considered.

The algorithm is able to solve large scale problem instances efficiently on most of the topologies considered. Future work could consider cache-friendly or cache oblivious evaluation techniques to allow the optimisation process to efficiently solve large scale graphs with unpredictable memory accesses. Additionally alternative solution representations that may better capture information on service quality should be considered.

ACKNOWLEDGMENT

K. Li was supported by UKRI Future Leaders Fellowship (Grant No. MR/S017062/1) and Royal Society (Grant No. IEC/NSFC/170243). J. Billingsley was supported by EPSRC Industrial CASE and British Telecom (Grant No. 16000177).

REFERENCES

- [1] OECD, “New skills for the digital economy,” 2016.

- [2] Z. Xiang, V. P. Magnini, and D. R. Fesenmaier, "Information technology and consumer behavior in travel and tourism: Insights from travel planning using the internet," *Journal of Retailing and Consumer Services*, vol. 22, pp. 244–249.
- [3] J. A. Bargh and K. Y. McKenna, "The internet and social life," *Annual Review Psychology*, vol. 55, pp. 573–590, 2004.
- [4] H. Farrell, "The consequences of the internet for politics," *Annual review of political science*, vol. 15, 2012.
- [5] A. Chadwick and C. May, "Interaction between states and citizens in the age of the internet: "e-government" in the united states, britain, and the european union," *Governance*, vol. 16, no. 2, pp. 271–300, 2003.
- [6] M. Castells, "The impact of the internet on society: a global perspective," *Ch@nge. 19 Key Essays on How the Internet is Changing our Lives*, pp. 132–133, 2014.
- [7] M. Avgerinou, P. Bertoldi, and L. Castellazzi, "Trends in data centre energy consumption under the European code of conduct for data centre energy efficiency," *Energies*, vol. 10, no. 1470, pp. 1–18, 2017.
- [8] J. Orlosky, K. Kiyokawa, and H. Takemura, "Virtual and augmented reality on the 5g highway," *JIP*, vol. 25, pp. 133–141, 2017.
- [9] M. R. Palattella, M. Dohler, L. A. Grieco, G. Rizzo, J. Torsner, T. Engel, and L. Ladid, "Internet of things in the 5g era: Enablers, architecture, and business models," *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 3, pp. 510–527, 2016.
- [10] H. Chourabi, T. Nam, S. Walker, J. R. Gil-García, S. Mellouli, K. Nahon, T. A. Pardo, and H. J. Scholl, "Understanding smart cities: An integrative framework," in *HICSS'12: Hawaii International International Conference on Systems Science*, 2012, pp. 2289–2297.
- [11] A. Andrae and T. Edler, "On global electricity usage of communication technology: trends to 2030," *Challenges*, vol. 6, no. 1, pp. 117–157, 2015.
- [12] M. C. Luizelli, W. L. da Costa Cordeiro, L. S. Buriol, and L. P. Gaspary, "A fix-and-optimize approach for efficient and large scale virtual network function placement and chaining," *Computer Communications*, vol. 102, pp. 67–77, 2017.
- [13] J. Cao, Y. Zhang, W. An, X. Chen, Y. Han, and J. Sun, "VNF placement in hybrid NFV environment: Modeling and genetic algorithms," in *ICPADS'16: Proc. of the 22nd IEEE International Conference on Parallel and Distributed Systems*, 2016, pp. 769–777.
- [14] C. M. Fonseca and P. J. Fleming, "Multiobjective optimization and multiple constraint handling with evolutionary algorithms. i. A unified formulation," *IEEE Trans. Systems, Man, and Cybernetics, Part A*, vol. 28, no. 1, pp. 26–37, 1998.
- [15] A. Leivadeas, M. Falkner, I. Lambadaris, M. Ibnkahla, and G. Kesidis, "Balancing delay and cost in virtual network function placement and chaining," in *NetSoft'18: 4th IEEE Conference on Network Softwarization and Workshops*, 2018, pp. 433–440.
- [16] W. Rankothge, J. Ma, F. Le, A. Russo, and J. Lobo, "Towards making network function virtualization a cloud computing service," in *IM'15: Proc. of the 2015 IFIP/IEEE International Symposium on Integrated Network Management*, 2015, pp. 89–97.
- [17] J. W. Jiang, T. Lan, S. Ha, M. Chen, and M. Chiang, "Joint VM placement and routing for data center traffic engineering," in *INFOCOM'12: Proc. of the 2012 IEEE Conference on Computer Communications*, 2012, pp. 2876–2880.
- [18] M. F. Bari, S. R. Chowdhury, R. Ahmed, and R. Boutaba, "On orchestrating virtual network functions," in *CNSM'11: Proc. of the 11th International Conference on Network and Service Management*, 2015, pp. 50–56.
- [19] S. Lange, A. Grigorjew, T. Zimmer, P. Tran-Gia, and M. Jarschel, "A multi-objective heuristic for the optimization of virtual network function chain placement," in *ITC'17: 29th International Teletraffic Congress*, 2017, pp. 152–160.
- [20] R. Gouareb, V. Friderikos, and A. Aghvami, "Virtual network functions routing and placement for edge cloud latency minimization," *IEEE Journal on Selected Areas in Communications*, vol. 36, no. 10, pp. 2346–2357, 2018.
- [21] J. Billingsley, K. Li, W. Miao, G. Min, and N. Georgalas, "A formal model for multi-objective optimisation of network function virtualisation placement," in *EMO'19 Evolutionary Multi-Criterion Optimization - 10th International Conference*, 2019, pp. 529–540.
- [22] F. C. T. Chua, J. Ward, Y. Zhang, P. Sharma, and B. A. Huberman, "Stringer: Balancing latency and resource usage in service function chain provisioning," *IEEE Internet Computing*, vol. 20, no. 6, pp. 22–31, 2016.
- [23] G. Miotto, M. C. Luizelli, W. L. da Costa Cordeiro, and L. P. Gaspary, "Adaptive placement & chaining of virtual network functions with NFV-PEAR," *J. Internet Services and Applications*, vol. 10, no. 1, pp. 3:1–3:19, 2019.
- [24] B. Addis, D. Belabed, M. Bouet, and S. Secci, "Virtual network functions placement and routing optimization," in *CLOUDNET'15: Proc. of the 4th IEEE International Conference on Cloud Networking*, 2015, pp. 171–177.
- [25] F. B. Jemaa, G. Pujolle, and M. Pariente, "Analytical models for qos-driven VNF placement and provisioning in wireless carrier cloud," in *MSWiM'16: Proceedings of the 19th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems*, 2016, pp. 148–155.
- [26] M. Gao, B. Addis, M. Bouet, and S. Secci, "Optimal orchestration of virtual network functions," *Computer Networks*, vol. 142, pp. 108–127, 2018.
- [27] D. B. Oljira, K. Grinnemo, J. Taheri, and A. Brunström, "A model for qos-aware VNF placement and provisioning," in *NFV-SDN'17: Conference on Network Function Virtualization and Software Defined Networks*, 2017, pp. 1–7.
- [28] A. Baumgartner, V. S. Reddy, and T. Bauschert, "Combined virtual mobile core network function placement and topology optimization with latency bounds," in *EWSDN'15: Proc. of the 4th European Workshop on Software Defined Networks*, 2015, pp. 97–102.
- [29] A. Alleg, R. Kouah, S. Moussaoui, and T. Ahmed, "Virtual network functions placement and chaining for real-time applications," in *CAMAD'17: Proc. of the 22nd IEEE International Workshop on Computer Aided Modeling and Design of Communication Links and Networks*, 2017, pp. 1–6.
- [30] A. Marotta, E. Zola, F. D'Andreagiovanni, and A. Kassler, "A fast robust optimization-based heuristic for the deployment of green virtual network functions," *J. Network and Computer Applications*, vol. 95, pp. 42–53, 2017.
- [31] K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," *IEEE Trans. Evolutionary Computation*, vol. 6, no. 2, pp. 182–197, 2002.
- [32] C. Guo, H. Wu, K. Tan, L. Shi, Y. Zhang, and S. Lu, "Dcell: a scalable and fault-tolerant network structure for data centers," in *SIGCOMM'08: Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, 2008, pp. 75–86.
- [33] M. Al-Fares, A. Loukissas, and A. Vahdat, "A scalable, commodity data center network architecture," in *SIGCOMM'08 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, 2008, pp. 63–74.
- [34] Cisco, "Spine and leaf architecture: Design overview white paper," 2019. [Online]. Available: <http://www.youtube.com/watch?v=8RrbUyw9uSg>
- [35] J. Billingsley, K. Li, W. Miao, G. Min, and N. Georgalas, "Multi-objective virtual network function placement: Formal model and effective algorithm."
- [36] S. Forrest and M. Mitchell, "Relative building-block fitness and the building block hypothesis," in *FOGA'92: Proceedings of the Second Workshop on Foundations of Genetic Algorithms*, 1992, pp. 109–126.
- [37] D. E. Goldberg, B. Korb, and K. Deb, "Messy genetic algorithms: Motivation, analysis, and first results," *Complex Systems*, vol. 3, no. 5, 1989.
- [38] C. Smart and P. J. Slater, "Center, median, and centroid subgraphs," *Networks*, vol. 34, no. 4, pp. 303–311, 1999.
- [39] Q. Zhang and H. Li, "MOEA/D: A multiobjective evolutionary algorithm based on decomposition," *IEEE Trans. Evolutionary Computation*, vol. 11, no. 6, pp. 712–731, 2007.
- [40] E. Zitzler and S. Künzli, "Indicator-based selection in multiobjective search," in *PPSN'04 8th International Conference on Parallel Problem Solving from Nature*, 2004, pp. 832–842.