

Network-Aware Container Scheduling in Multi-Tenant Data Center

Leonardo R. Rodrigues,[◇] Marcelo Pasin,^{*} Omir C. Alves Jr.,[◇]
Charles C. Miers,[◇] Mauricio A. Pillon,[◇] Pascal Felber,^{*} Guilherme P. Koslovski[◇]
Graduate Program in Applied Computing – Santa Catarina State University – Joinville – Brazil[◇]
University of Neuchâtel (UniNE) - Institut d’informatique* – Switzerland

Abstract—Network management on multi-tenant container-based data center has critical impact on performance. Tenants encapsulate applications in containers abstracting away details on hosting infrastructures, and entrust data center management framework with the provisioning of network Quality-of-Service requirements. In this paper, we propose a network-aware multi-criteria container scheduler to jointly process containers and network requirements. We introduce a new Mixed Integer Linear Programming formulation for network-aware scheduling encompassing both tenants and providers metrics. We describe two GPU-accelerated modules to address the complexity barrier of the problem and efficiently process scheduling requests. Our experiments show that our scheduling approach accounting for both network and containers outperforms traditional algorithms used by containers orchestrators.

I. INTRODUCTION

Container-based virtualization offers a lightweight mechanism to host and manage large-scale distributed applications for big data processing, edge computing, stream processing, among others. Multiple tenants encapsulate applications’ environments in containers, abstracting away details of operating systems, library versions, and server configurations. With containers, data center (DC) management becomes application-oriented [1] in contrast to server-oriented when using virtual machines. Several technologies are used to provide connections between containers, such as virtual switches, bridges, and overlay networks [2]. Yet, containers are a catalyst for network management complexity. Network segmentation, bandwidth reservation, and latency control are essential requirements to support distributed applications, but container management frameworks still lack appropriate tools to support Quality-of-Service (QoS) requirements for network provisioning [1].

We argue that container networking must address at least three communication scenarios, despite the orchestration framework used by the DC: highly-coupled container-to-container communication, group-to-group communication, and containers-to-service communication. Google’s Kubernetes offers a viable solution to group network-intensive or highly coupled containers, by using *Pods*. A pod is a group of one or more containers with shared storage and network, and pods must be provisioned on a single host. Because the host bus conducts all data transfers within a pod, communication latency is more constant, increasing the network throughput, achieving values superior to default network switching technologies. However, for large-scale distributed applications,

multiple pods must be provisioned and eventually allocated on distinct servers.

This paper advances the field on network-aware container scheduling, a primary management task on container-based DCs [1], by jointly allocating compute and communication resources to host network-aware requests. The network-aware scheduling is analogous to the virtual network embedding (VNE) problem [3]. Given two graphs, the first one representing user requested containers and all corresponding communications requirements, and the second denoting the DC hosting candidates (servers, virtual machines, links, and paths), one must find a map for each vertex and edge from the request graph to a corresponding vertex and edge on DC graph. Vertices and edges carry weights representing process and bandwidth and constraints. The combined scheduling of containers and network QoS requires a multi-criteria decision, based on conflicting constraints and objectives.

We formally define in this paper the scheduling problem encompassing the network QoS as a Mixed Integer Linear Programming (MILP). We later propose two Graphics Processing Unit (GPU)-accelerated multi-criteria algorithms to process large-scale requests and DC topologies.

The paper is organized as follows. §II describes the problem formulation, while §III defines an optimal MILP for joint container and network QoS requirements allocation. Following, §IV presents the evaluation of the proposed MILP highlighting the efficiency and limitations of network-aware scheduling. Then, §V describes the implementation of two GPU-accelerated algorithms to speed up the scheduling process, and both algorithms are compared with traditional approaches in §VI. Related work is reviewed in §VII and §VIII concludes.

II. PROBLEM FORMULATION

A. DC Resources and Tenants Requests

Data center resources (bare metal or virtualized) are represented by $G^s(N^s, E^s)$, where N^s denotes the physical servers, and E^s contains all physical links composing the network topology. A vector is associated with each physical server $u \in N^s$, representing the available capacities ($c_u^s[r]; r \in R$) where R denotes resources as RAM and CPU. In addition, bw_{uv}^s represents the available bandwidth between physical servers u and v . Thus, a tenant request is given by $Req(N^c, E^c)$, with N^c being a set of containers and E^c

Notation	Description
$G^s(N^s, E^s)$	DC graph composed of N^s servers and E^s links.
$c_u^s[r]$	Resource capacity vector of server $u \in N^s$.
P^s	All direct paths (physical and logical) on DC topology.
bw_{uv}^s	Bandwidth capacity between servers u and v , $uv \in E^s$.
$Req(N^c, E^c)$	Request, composed of N^c containers and E^c links.
$c_i^{min}[r], c_i^{max}[r]$	Minimum and maximum resources capacities for container $i \in N^c$.
$bw_{ij}^{min}, bw_{ij}^{max}$	Minimum and maximum bandwidth requirement between containers i and j , $ij \in E^c$.
$pod_g \subset N^c$	Set of containers $i \in N^c$ composing a pod $g \in G$.

TABLE I

NOTATION USED ALONG THIS PAPER: i AND j ARE USED FOR INDEXING CONTAINERS, WHILE u AND v ARE USED FOR DC SERVERS.

the communication requirements among them. Also, as in Kubernetes, each container is associated with a pod.

Containers from a pod must be hosted by the same physical server (sharing the IP address and port space). A group of pods G is defined in a tenant's request, and a container $i \in N^c$ is connected to a pod group $g \in G$, indicated by $i \in pod_g$. Instead of requesting a fixed configuration for each QoS requirement, containers are specified as minimum and maximum intervals. For a container i , the minimum and maximum values for any $r \in R$ are respectively defined as $c_i^{min}[r]$ and $c_i^{max}[r]$. The same rationale is applied to containers interconnections (E^c): minimum and maximum bandwidth requirements are given by bw_{ij}^{min} and bw_{ij}^{max} .

A container orchestration framework has to determine whether to accept or not a tenant request. The allocation of containers onto a DC is decomposed into nodes and links assignments. The mapping of containers onto nodes is given by $\mathcal{M}_c : N^c \mapsto N^s$, while the mapping of networking links between containers onto paths is represented as $\mathcal{M}_{ec} : E^c \mapsto P^s$. Table I summarizes the notation used in this paper.

B. Objectives

Energy consumption. To reduce energy consumption, we pack containers in as few nodes as possible, allowing to power off the unused ones. We call this technique *consolidation*, and we reach it by minimizing the DC fragmentation, defined as the ratio of the number of active servers (e.g., those hosting containers) to the total number of DC resources. Server fragmentation is given by $\mathcal{F}(N^s) = |N^{s'}|/|N^s|$, while the same rationale is applied for links, $\mathcal{F}(E^s) = |E^{s'}|/|E^s|$, where $|N^{s'}|$ and $|E^{s'}|$ denote the number of active servers and links, respectively.

Quality-of-Service. A container can be successfully executed with any capacities configuration in the intervals specified as minimum and maximum. However, optimal performance is reached when the maximum values are used. In this sense, utility functions can be applied for representing the improvement on container's configuration. In short, the goal is to maximize Eq. (1) and (2) for each container $i \in N^c$, where $c_{iu}^a[r]$ and bw_{ijuv}^a represent the capacity effectively allocated for vertices and edges, respectively.

$$U(i) = \frac{\sum_{r \in R} \frac{c_{iu}^a[r]}{c_i^{max}[r]}}{|R|}; u = \mathcal{M}_c(i) \quad (1)$$

$$U(ij) = \frac{\sum_{uv \in \mathcal{M}_{ec}(ij)} bw_{ijuv}^a}{bw_{ij}^{max}} \quad (2)$$

III. OPTIMAL MILP FOR JOINT CONTAINER AND NETWORK QoS ALLOCATION

A. Variables and Objective Function

A set of variables (Table II) are proposed to find a solution for joint allocation of containers and bandwidth requirements, as well as to achieve maps $\mathcal{M}_c : N^c \mapsto N^s$, and $\mathcal{M}_{ec} : E^c \mapsto P^s$. The binary variable x_{iu} accounts the mapping of containers on servers. The containers' connectivity (x_{lijuv}) applies the same rationale. For identifying the amount of resources allocated to a container $i \in N^c$, the float vector c_i^a is introduced. Bandwidth allocation follows the same principle and is accounted by float variable bw_{ij}^a .

Notation	Type	Description
x_{iu}	Bool	Container $i \in N^c$ is mapped on server $u \in N^s$.
x_{lijuv}	Bool	Connection $ij \in E^c$ is mapped on link $uv \in E^s$.
$c_{iu}^a[r]$	Float	Resource ($r \in R$) capacity vector allocated to container $i \in N^c$ on server $u \in N^s$.
bw_{ijuv}^a	Float	Bandwidth allocated to connection $ij \in E^c$ on link $uv \in E^s$.
f_u	Bool	Server $u \in N^s$ is hosting at least one container.
fl_{uv}	Bool	Link $uv \in E^s$ is hosting at least one connection.

TABLE II

MILP VARIABLES FOR MAPPING CONTAINERS AND VIRTUAL LINKS ATOP A MULTI-TENANT DC.

The objectives (§II-B) are reached by the minimization of Eq. (3). Two additional binary variables are used to identify if DC resources are hosting at least one container or link, f_u and fl_{uv} . Value 1 is set just for active servers, as given by $f_u \geq \frac{\sum_{i \in N^c} x_{iu}}{|N^c|}; \forall u \in N^s$. Physical links follow the same idea, $fl_{uv} \geq \frac{\sum_{ij \in E^c} x_{lijuv}}{|E^c|}; \forall uv \in E^s$. Finally, the importance level of each term is defined by setting α .

$$\begin{aligned} \text{minimize} : & \alpha \left(\sum_{i \in N^c} (1 - U(i)) + \sum_{ij \in E^c} (1 - U(ij)) \right) \\ & + (1 - \alpha) \left(\sum_{u \in N^s} \frac{f_u}{|N^s|} + \sum_{uv \in E^s} \frac{fl_{uv}}{|E^s|} \right) \end{aligned} \quad (3)$$

B. Constraints

DC Capacity, QoS Constraints and Integrity of Pods. A DC server $u \in N^s$ must support all hosted containers, as indicated by Eq. (4), while the bandwidth of link $uv \in E^s$ must support all containers transfers allocated to it, as given by Eq. (5). Eq. (6) guarantees the allocation of a resources capacities from min-max intervals for a containers $i \in N^c$. The same rationale is applied for $ij \in E^c$ on Eq. (7).

$$c_u^s[r] \geq \sum_{i \in N^c} c_{iu}^a[r]; \forall u \in N^s; \forall r \in R \quad (4)$$

$$bw_{uv}^s \geq \sum_{ij \in E^c} bw_{ijuv}^a; \forall uv \in E^s \quad (5)$$

$$c_i^{min}[r] \times x_{iu} \leq c_{iu}^a[r] \leq c_i^{max}[r] \times x_{iu} \quad (6)$$

$$\forall i \in N^c; \forall u \in N^s; \forall r \in R$$

$$bw_{ij}^{min} \times xl_{ijuv} \leq bw_{ijuv}^a \leq bw_{ij}^{max} \times xl_{ijuv} \quad (7)$$

$$\forall ij \in E^c; uv \in E^s$$

$$x_{iu} = x_{ju}; \forall g \in G; \forall i, j \in pod_g; \forall u \in N^s \quad (8)$$

Finally, containers are optionally organized in pods. For guaranteeing the integrity of pods specifications, Eq. (8) indicates that all resources from a pod ($i, j \in pod_g$) must be hosted by the same server ($u \in N^s$).

Binary and Allocation Constraints. A container must be hosted by a single server ($\sum_{u \in N^s} x_{iu} = 1; \forall i \in N^c$), while each virtual connectivity between containers is mapped to a path between resources hosting its source and destination as given by $\sum_{v \in N^s} xl_{ijvu} + \sum_{v \in N^s} xl_{ijuv} = x_{iu} + x_{ju}; \forall u \in N^s; \forall ij \in E^c$. However, on large scale DC topologies, servers are interconnected by multiple paths composed of at least one switch hop. In order to keep the model realistic with current DCs, we rely on network management techniques, such as SDN [4] to control the physical links usage and populate the E^s with updated information and available paths.

IV. EVALUATION OF THE OPTIMAL MILP FOR NETWORK-AWARE CONTAINERS SCHEDULING

The MILP scheduler and a discrete event simulator were implemented in Python 2.7.10 using CPLEX optimizer (v12.6.1.0). For composing the baseline was used the native algorithms offered by containers orchestrators, Best Fit (BF) (binpacking) and Worst Fit (WF) (spread). As BF and WF natively ignore the network requirements, we included a shortest-path search after the allocation of servers to host containers for conducting a fair comparison.

A. Metrics and MILP Parametrization

The MILP objective function, Eq. (3), is composed of terms to represent the tenant's perspective (the utility of network allocation and the queue waiting time) and the DC fragmentation (the provider's perspective). Although a minimum value is requested for each container parameter, the optimal utility function expects the allocation of maximum values ($\mathcal{U}(\cdot) = 1$). The MILP-based scheduler is guided by the α value to define the importance of each term composing the objective function. For demonstrating the impact of defining α , we evaluated 3 configurations $\alpha = 0; 0.5; 1$. Configurations with $\alpha = 0$ and $\alpha = 1$ define the baseline for comparisons: by setting $\alpha = 0$ the MILP optimizes the problem regarding the fragmentation perspective only, while $\alpha = 1$ represents the opposite; more importance is given to containers and network utilities.

B. Experimental Scenarios

1) *DC Configuration:* A Clos-based topology (termed Fat-Tree) is used to represent the DC [5], [6]. The k factor guides the topology indicating the number of switches, links, and hosts used to compose the DC. A fat-tree build with k -port switches supports up to $k^3/4$ servers. The DC is configured with $k = 4$, and composed of homogeneous servers equipped with 24 cores and 256 GB RAM, while the bandwidth capacity for all links is defined as 1 Gbps.

2) *Requests:* A total of 200 requests is submitted with resources specifications based on uniform distributions for containers capacities, submission time, and duration. Each request is composed of 5 containers with a running time up to 200 events from a complete execution of 500 events. For composing the pods, up to 50% of containers from a single requested are grouped in pods. For the network, the bandwidth requirement between a pair of containers is configured up to 50 Mbps, besides requests with 1 Mbps requirement representing applications without burdensome network requirements. The values for CPU and RAM configuration are uniformly distributed up to 2 and 4, respectively.

C. Results and Discussion

Table III and Figures 1 and 2 present results for utility of network and container requests, provisioning delays, and DC network fragmentation, respectively.

BF and WF algorithms have a well-defined pattern for all network utility metric. For requests with low network requirements (up to 1Mbps), both algorithms tend to allocate the maximum requested value for network QoS. An exception is observed for BF with network-intensive requests (up to 50Mbps) as the algorithm gives priority to minimum requested values for consolidating requests on DC resources. With regarding the network-aware MILP scheduler, even for requests with $\alpha = 0$ focusing on decreasing the DC fragmentation, the scheduler allocated maximum values for network requests, following the BF and WF algorithms. However, the impact of α parametrization is perceived for network-intensive requests. The MILP configuration with $\alpha = 0.5$ shows that the algorithm can jointly consider requests utility and DC fragmentation. The results in Fig. 2 show that scheduling

Algorithm	α	Bandwidth	$\mathcal{U}(ij)$	$\mathcal{U}(i)$
MILP	0	1 Mbps	22.68%	99.90%
		50 Mbps	7.86%	66.29%
	0.5	1 Mbps	26.78%	99.90%
		50 Mbps	86.67%	97.21%
	1	1 Mbps	38.28%	99.90%
		50 Mbps	93.56%	97.80%
WF	-	1 Mbps	100%	98.03%
		50 Mbps	100%	99.98%
BF	-	1 Mbps	100%	97.20%
		50 Mbps	100%	99.46%

TABLE III
LINK AND CONTAINER UTILITIES FOR MILP, BF, AND WF.




Fig. 1. Requests delay when executing the MILP-based scheduler.




Fig. 2. DC network fragmentation when executing the MILP-based scheduler.

network-intensive requests increases the network DC fragmentation. The provisioning delays (Figure 1) explain this fact: the MILP scheduler decreases the queue waiting time for network-intensive requests when compared to BF and WF.

In summary, it is evident that network QoS must be considered by the scheduler to decrease the queue waiting time and to reserve utility’s dynamic configurations. Moreover, the results obtained from MILP configured with $\alpha = 0.5$ demonstrated the real trade-off between fragmentation and utility, or in other words, provider’s and tenant’s perspectives.

V. GPU-ACCELERATED HEURISTICS

Although MILP is efficient to model and highlights the impact of network-aware scheduling, solving this problem is known to be computationally intractable [3] and practically infeasible for large-scale scenarios. Therefore, we developed two GPU-accelerated multi-criteria algorithms to speed up the joint scheduling of containers and network with QoS requirements. We selected two multi-criteria algorithms: Analytic Hierarchy Process (AHP) and Technique for Order Preference by Similarity to Ideal Solution (TOPSIS), chosen due to their multidimensional analysis, being able to work with several servers simultaneously. Also, AHP and TOPSIS provide a structured method to decompose the problem and to consider trade-offs in conflicting criteria. Following the notation used to express the MILP (Table I), both algorithms analyze the same set of criteria c_u^s for a given server u . In addition, the sum of all bandwidth capacity bw_{uv}^s with source on u (given by bw_u^s) and the current server fragmentation (f_u) are accounted and included on c_u^s capacity vector. The multi-criteria algorithms analyzed all variables described in Section II-B as attributes.

A. Weights Distribution

AHP and TOPSIS algorithms are guided by a weighting vector to define the importance of each criteria. While the MILP has α to indicate the importance level of each term in the objective function, the multi-criteria function decomposes α into a vector $W = \{\alpha_0, \alpha_1, \dots, \alpha_{|R|-1}\}$; $\sum_{i \in R} \alpha_i = 1$. Tab. IV presents different W compositions to the MILP objective.

Scenario	CPU	RAM	Fragmentation	Bandwidth
Flat	0.25	0.25	0.25	0.25
Clustering	0.17	0.17	0.5	0.16
Network	0.17	0.17	0.16	0.5

TABLE IV

WEIGHTING SCHEMA FOR AHP AND TOPSIS. THE FLAT CONFIGURATION IS EQUIVALENT TO $\alpha = 0.5$ IN MILP, WHILE CLUSTERING AND NETWORK REPRESENTS $\alpha = 0$ AND $\alpha = 1$, RESPECTIVELY.

The multi-criteria analysis with clustering configuration optimizes the problem aiming at DC consolidation (equivalent to $\alpha = 0$ on MILP formulation) through the definition of high importance level (50%) to fragmentation criteria, while the other criteria share equally the last 50%. In other hand, the execution with network configuration ($\alpha = 1$ from MILP formulation), the bandwidth criteria receive a higher importance level (50%) while the other criteria share equally the last 50%. This configuration makes the scheduler select servers that have the highest residual bandwidth. Finally, the flat configuration sets the same importance weight for all criteria (following the $\alpha = 0.5$ rationale on MILP).

B. AHP

The AHP is a multi-criteria algorithm that hierarchically decomposes the problem to reduce the complexity, and performs a pairwise comparison to rank all alternatives [7]. In short, the hierarchical organization is composed of three main levels. The objective of the problem is configured at the top of the hierarchy, while the set of criteria is placed in the second level, and finally, in the third level represents all the viable alternatives to solve the problem.

In our context, the selection of the most suitable DC to host a container is performed in steps. In the first step two vectors (M_1 and M_2) are built combining all criteria and alternatives (second and third level of AHP hierarchy) applying the weights defined in Table IV. In other words, $M_1[v] = W[v]; \forall v \in R$ while $M_2[v \times |N^s| + u] = c_u^s[v]; \forall u \in N^s; \forall v \in R$. The representation based on a vector was chosen to exploit the Single Instruction Multiple Data (SIMD) GPU-parallelism. Later the pairwise comparison is applied for all elements into the hierarchy. If $M_1[v \times |R| + u] > 0$, the value $M_1[v \times |R| + u] - M_1[i \times |R| + u]$ is attributed to; In addition, if the cell value is < 0 , $\frac{1}{M_1[v \times |R| + u] - M_1[i \times |R| + u]}$ is set; and 1 otherwise. The same rationale is applied for M_2 , indexed by $v \times |N^s|^2 + i \times |N^s| + u$. Later, both vectors are normalized. At this point, the algorithm calculates the local rank of each element in the hierarchy (L_1 and L_2), as described in Eqs. (9) and (10), $\forall u, v \in R; \forall i, j \in N^s$. Finally, the global priority (PG) of the alternatives is accounted to guide the host selection, as given by $PG[v] = \sum_{x \in N^s} P_1[v] \times P_2[v \times |N^s| + x]$.

$$L_1[v \times |R| + u] = \frac{\sum_{x \in R} M_1[v \times |R| + x]}{|R|} \quad (9)$$

$$L_2[v \times |N^s| + j] = \frac{\sum_{x \in N^s} M_2[v \times |N^s|^2 + i \times |N^s| + x]}{|N^s|} \quad (10)$$

C. TOPSIS

The Technique for Order Preference by Similarity to Ideal Solution (TOPSIS) is based in the shortest Euclidean Distance from the alternative to the ideal solution [8]. The benefits of this algorithm are three-fold: (i) can handle a large number of criteria and alternatives; (ii) requires a small number of qualitative inputs when compared to AHP; and (iii) is a compensatory method, allowing the analysis of trade-off criteria.

The ranking of the DC candidates is performed in steps. Initially, the evaluation vector M correlates DC resources (N^s) and the criteria elements (R): $M[v \times |N^s| + u] = c_u^s[v]; \forall u \in N^s; \forall v \in R$, which is later normalized. The next step is the application of weighting schema on M values: $M[v \times |N^s| + u] = M[v \times |N^s| + u] \times W[v]; \forall u \in N^s; \forall v \in R$. Based on M , two vectors are then composed with the maximum and minimum values for each criteria, represented by A^+ (the upper-bound solution quality) and A^- (the lower-bound). TOPSIS requires the calculation of Euclidean distances between M and upper- and lower-bounds, composing Ed^+ and Ed^- . Finally, a closeness coefficient array is accounted for all DC servers, $Rank[u] = \frac{Ed^- [u]}{Ed^+ [u] + Ed^- [u]}; \forall u \in N^s$, and afterwards the resulting array is sorted on decreasing order, indicating the selected candidates.

D. GPU Implementation

The AHP and TOPSIS are decomposed in GPU-tailored kernels following a pipeline execution. The first kernel in charge is acquiring DC and network-aware containers requests, while the remaining kernels perform the comparisons using the parallel reduction technique. A special explanation is required for selecting physical paths to host containers interconnections. After the selection of the most suitable server for each pod presented in the tenant's request, the virtual links between the containers must be set. A modified Dijkstra algorithm is used to compute the shortest path that has the maximum available bandwidth between the hosting servers. The modified Dijkstra is implemented as a single kernel to allow multiple executions, where each thread calculates a different source and destination pair. As the links between every two nodes in the DC are undirected, the GPU implementation uses a specific array representation to reduce the total space needed. The main principle of the data structure of this algorithm is that $u < v$ where u is the source and v the destination, and the paths $u \rightarrow v$ and $v \rightarrow u$ are the same.

VI. EVALUATION OF GPU-ACCELERATED HEURISTICS

The GPU-accelerated scheduler and a discrete event simulator were implemented in C++, using GCC compiler *v.8.2.1* and CUDA framework *v.10.1*.

A. Experimental Scenarios

The evaluation considers a DC composed of homogeneous servers equipped with 24 cores, 256 GB RAM and interconnected by a Fat-Tree topology ($k = 20$) and bandwidth capacity of 1 Gbps for all links. A total of 6000 requests were submitted to be scheduled, each composed of 4 containers with

a running time up to 250 events from a complete execution of 500 events. For composing the requests, up to 50% of containers from a single request are grouped in pods, while the bandwidth requirement between a pair of containers is configured up to 50 Mbps (a heavy network requirement).

B. Results and Discussion

Results are summarized by Table V and Figures 3 and 4, showing data for the runtime, utility of network and container requests, provisioning delays correlated to the DC fragmentation and DC network fragmentation, respectively.

Algorithm	Scenario	# Events	Average Runtime (s)	$U(ij)$	$U(i)$
BF	-	2462	79.38	100%	96.89%
WF	-	1007	47.80	100%	99.41%
AHP	Flat	949	9.45	100%	98.22%
	Clustering	936	7.51	100%	99.10%
	Network	928	6.90	100%	98.41%
Topsis	Flat	894	3.67	100%	98.85%
	Clustering	916	3.84	100%	99.01%
	Network	892	3.48	100%	98.94%

TABLE V
RUNTIME, LINK AND CONTAINER UTILITIES FOR BF, WF, AHP AND TOPSIS.




Fig. 3. DC network fragmentation from the GPU-accelerated scheduler.




Fig. 4. DC links fragmentation from the GPU-accelerated scheduler.

Figure 3 shows that the multi-criteria algorithms have a small variation for request delay, grouping the data in high fragmentation percentages, while the WF induces delay in requests regardless the DC fragmentation. In turn, the BF algorithm imposes higher delay to requests resulting in a small fragmentation percentage, below 30% of network fragmentation. WF and BF generate a long requests queue impacting directly in the total computational time needed to schedule all the tenants' requests.

Regarding the container's utility (Table V), the multi-criteria algorithms give priority to schedule requests mixing between the maximum and minimum requirements, increasing the number of containers in the DC. The WF tends to allocate the maximum value for the requests, while the BF tends to

give the minimum values of the requests. While the multi-criteria algorithms increase the number of containers in the DC reducing the total delay, the network fragmentation have similar behavior with the WF algorithm, as shown in Figure 4. Meanwhile the BF keeps the network fragmentation small due to the long delays that it applies in the requests. It is possible to observe that the multi-criteria algorithms present better consolidation results when compared to the WF and BF algorithms, due to their capacity to allocate more requests in the DC keeping the fragmentation similar to WF. It is possible to conclude that the network weighting schema is essential to perform a joint scheduling of container and network requirements. It is important to emphasize: the GPU-accelerated algorithms can schedule the requests with bandwidth requirements atop a large-scale DC in a few seconds. Specifically, TOPSIS outperformed BF, WF, and AHP results.

VII. RELATED WORK

The orchestration and scheduling of virtualized DC is a trendy topic of the specialized literature. MILP techniques offer optimal solutions which are generally used as a baseline for comparisons [4], but the problem complexity and search space often create opportunities for heuristic-based solutions [3].

Guerrero [9] proposed a scheduler for container-based micro-services. The containers workload and the networking details were analyzed to perform the DC load balance. Guo [10] proposed a scheduler to optimize the load balancing and workload through the neighborhood division in a micro-service method. Both proposals were analyzed on small-scale DCs as the problem complexity imposes a barrier on real-scale use. The GenPack [11] scheduler employs monitoring information to define the appropriated group of a container based on the resource usage, avoiding resources disputes among containers. A security-concerned scheduler was proposed by [12], based on bin-packing executing a BF approach. GPU-accelerated algorithms can be applied to speed-up these heuristics reaching large-scale DCs [13].

A joint scheduler based on priority-queue, AHP and Particle Swarm Optimization (PSO) is proposed by [14]. The requests are sorted by their priority level and waiting time, and then the tasks are sent to the AHP to be ranked and then serving as an input to PSO. The results show a reduction on makespan up to 15% when compared to PSO. In addition, [15] proposed a VM scheduler based on TOPSIS and PSO. The scheduler was compared with 5 meta-heuristics using the 4 metrics: makespan, transmission time, cost and resource utilization, achieving an improvement up to 75% when compared to traditional schedulers. Although many multi-criteria solutions appear in the literature, we were unable to find schedulers dealing with containers, pods, and their virtual networks.

Network requirements are disregarded or partially attended by major of reviewed schedulers. Even well-known orchestrators (*e.g.*, Kubernetes) consider the network as second-level and not critical parameters. Containers are used to model large-scale distributed applications, and it is evident that network allocation can impact on applications performance [2].

VIII. CONCLUSION

We investigated the joint scheduling of network QoS and containers on multi-tenant DCs. A MILP formulation and experimental analysis reveal that a network-aware scheduler can decrease DC network fragmentation and processing delays. However, solving a MILP is known to be computationally intractable and practically infeasible for large-scale scenarios. We then developed two GPU-accelerated multi-criteria algorithms, AHP and TOPSIS, to schedule requests on a large-scale DC. Both network-aware algorithms outperformed the traditional schedulers with regard to DC and tenant perspectives. Future work includes the scheduling of batch requests and a distributed implementation for increasing the fault tolerance.

Acknowledgments. The research leading to the results presented here has received funding from UDESC, FAPESC, and the European Union’s Horizon 2020 research and innovation programme under the LEGaTO Project (legato-project.eu), grant agreement No 780681.

REFERENCES

- [1] B. Burns, B. Grant, D. Oppenheimer, E. Brewer, and J. Wilkes, “Borg, omega, and kubernetes,” *Queue*, vol. 14, no. 1, pp. 10:70–10:93, 2016.
- [2] K. Suo, Y. Zhao, W. Chen, and J. Rao, “An analysis and empirical study of container networks,” in *IEEE INFOCOM 2018-IEEE Conf. on Computer Communications*. IEEE, 2018, pp. 189–197.
- [3] M. Rost, E. Döhne, and S. Schmid, “Parametrized complexity of virtual network embeddings: Dynamic & linear programming approximations,” *SIGCOMM Comput. Commun. Rev.*, vol. 49, no. 1, pp. 3–10, Feb. 2019.
- [4] F. R. de Souza, C. C. Miers, A. Fiorese, M. D. de Assunção, and G. P. Koslovski, “Qvia-sdn: Towards qos-aware virtual infrastructure allocation on sdn-based clouds,” *Journal of Grid Computing*, Mar 2019.
- [5] S. Arjun, O. Joon, A. Amit, A. Glen, A. Ashby, B. Roy *et al.*, “Jupiter rising: A decade of clos topologies and centralized control in google’s datacenter network,” in *Sigcomm ’15*, 2015.
- [6] R. Niranjana Mysore, A. Pamboris, N. Farrington, N. Huang, P. Miri *et al.*, “Portland: A scalable fault-tolerant layer 2 data center network fabric,” *SIGCOMM Comput. Commun. Rev.*, vol. 39, pp. 39–50, 2009.
- [7] T. L. Saaty, “Making and validating complex decisions with the AHP/ANP,” *J SYST SCI SYST ENG*, vol. 14, no. 1, pp. 1–36, 2005.
- [8] C.-L. Hwang and K. Yoon, *Multiple Attribute Decision Making*. Lecture Notes in Economics and Mathematical Systems, Springer, 1981.
- [9] C. Guerrero, I. Lera, and C. Juiz, “Genetic algorithm for multi-objective optimization of container allocation in cloud architecture,” *J GRID COMPUT*, vol. 16, no. 1, pp. 113–135, Mar 2018.
- [10] Y. Guo and W. Yao, “A container scheduling strategy based on neighborhood division in micro service,” in *NOMS 2018-2018 IEEE/IFIP Network Operations and Management Symp.* IEEE, 2018, pp. 1–6.
- [11] A. Havet, V. Schiavoni, P. Felber, M. Colmant, R. Rouvoy, and C. Fetzer, “Genpack: A generational scheduler for cloud data centers,” in *IEEE Int. Conf. on Cloud Engineering (IC2E)*, April 2017, pp. 95–104.
- [12] S. Vaucher, R. Pires, P. Felber, M. Pasin, V. Schiavoni, and C. Fetzer, “SGX-aware container orchestration for heterogeneous clusters,” in *IEEE 38th Int. Conf. on Distributed Comp. Systems*, July 2018.
- [13] L. L. Nesi, M. A. Pillon, M. D. de Assunção, C. C. Miers, and G. P. Koslovski, “Tackling virtual infrastructure allocation in cloud data centers: a gpu-accelerated framework,” in *2018 14th Int. Conf. on Network and Service Management (CNSM)*, Nov 2018, pp. 191–197.
- [14] H. B. Alla, S. B. Alla, A. Ezzati, and A. Touhafi, “An efficient dynamic priority-queue algorithm based on ahp and pso for task scheduling in cloud computing,” in *HIS*. Springer, 2016, pp. 134–143.
- [15] N. Panwar, S. Negi, M. M. S. Rauthan, and K. S. Vaisla, “Topsis-pso inspired non-preemptive tasks scheduling algorithm in cloud environment,” *Cluster Computing*, pp. 1–18.