

Scheduling Storms and Streams in the Cloud

Javad Ghaderi
Columbia University
New York, NY
jghaderi@columbia.edu

Sanjay Shakkottai
University of Texas
Austin, TX
shakkott@austin.utexas.edu

R Srikant
University of Illinois
Urbana, IL
rsrikant@illinois.edu

ABSTRACT

Motivated by emerging big *streaming* data processing paradigms (e.g., Twitter Storm, Streaming MapReduce), we investigate the problem of scheduling graphs over a large cluster of servers. Each graph is a job, where nodes represent compute tasks and edges indicate data-flows between these compute tasks. Jobs (graphs) arrive randomly over time, and upon completion, leave the system. When a job arrives, the scheduler needs to partition the graph and distribute it over the servers to satisfy load balancing and cost considerations. Specifically, neighboring compute tasks in the graph that are mapped to different servers incur load on the network; thus a mapping of the jobs among the servers incurs a cost that is proportional to the number of “broken edges”. We propose a low complexity randomized scheduling algorithm that, without service preemptions, stabilizes the system with graph arrivals/departures; more importantly, it allows a smooth trade-off between minimizing average partitioning cost and average queue lengths. Interestingly, to avoid service preemptions, our approach does not rely on a Gibb’s sampler; instead, we show that the corresponding limiting invariant measure has an interpretation stemming from a loss system.

Keywords

Graph Partitioning, Dynamic Resource Allocation, Markov Chains, Probabilistic Algorithms

1. INTRODUCTION

In recent years, a new computing model – stream processing – is gaining traction for large-scale cloud computing systems. These systems [15, 25, 29, 32] are driven by real time and streaming data applications. For instance, consider the computation needed to answer the question: *How many times does the hashtag “#sigmetrics2015” appear in Twitter over the next two hours?* The key feature here is that the data is not (yet) in a database; instead it is appearing as and when people tweet this hashtag. Applications of such stream computing are in many domains including social network analytics and e-commerce.

To address such stream processing, the emerging computation model of choice is that of graph processing. A computation is represented by a graph, where nodes in the graph represent either data sources or data processing (and operate sequentially on a stream of atomic data units), and edges in the graph correspond to data flows between nodes. To execute such computations, each node of a graph is mapped

to a machine (server/blade) in a cloud cluster (data center), and the communication fabric of the cloud cluster supports the data flows corresponding to the graph edges. A canonical example (and one of the early leaders in this setting) is Twitter’s Storm [29], where the (directed) graph is called a “topology”, an atomic data unit is a “tuple”, nodes are called “spouts” or “bolts”, and tuples flow along the edges of the topology. We refer to [2] for additional discussion.

From the cloud cluster side, there are a collection of machines interconnected by a communication network. Each machine can simultaneously support a finite number of graph nodes. This number is limited by the amount of resources (memory/processing/bandwidth) that is available at the machine; in Storm, these available resources are called “slots” (typically order of ten to fifteen per machine). Graphs (corresponding to new computations) arrive randomly over time to this cloud cluster, and upon completion, leave the cluster. At any time, the scheduling task at the cloud cluster is to map the nodes of an incoming graph onto the free slots in machines to have an efficient cluster operation. As an example, the default scheduler for Storm is round-robin over the free slots; however, this is shown to be inefficient, and heuristic alternatives have been proposed [2].

In this paper we consider a queueing framework that models such systems with graph arrivals and departures. Jobs are graphs that are *dynamically* submitted to the cluster and the scheduler needs to partition and distribute the jobs over the machines. Once deployed in the cluster, the job (a computation graph) will retain the resources for some time duration depending on the computation needs, and will release the resources after the computation is done (i.e., the job departs). The need for efficient scheduling and dynamic graph partitioning algorithms naturally arises in many parallel computing applications [12, 13]; however, the theoretical studies in this area are very limited. To the best of our knowledge, this is the first paper that develops models of dynamic stochastic graph partitioning and packing, and the associated low complexity algorithms with provable guarantees for graph-based data processing applications.

From an algorithmic perspective, our low complexity algorithm has connections to the Gibbs sampler and other MCMC (Monte Carlo Markov Chain) methods for sampling probability distributions (see for example [3]). In the setting of scheduling in wireless networks, the Gibb’s sampler has been used to design CSMA-like algorithms for stabilizing the network [17, 19, 23, 26]. However, unlike wireless networks where the solutions form independent sets of a graph, there is no natural graph structure analog in the graph partition-

ing. The Gibbs sampler can still be used in our setting by sampling *partitions* of graphs, where, each site of the Gibbs sampler is a unique way of partitioning and packing a graph among the machines in the cloud cluster. The difficulty, however, is that there are an exponentially large number of graph partitions, leading to a correspondingly large number of queues. The second issue is that a Gibbs sampler potentially can interrupt ongoing service of jobs. The analog of a service interruption in our setting is the migration of a job (graph) from one set of machines to another in the cloud cluster. This is an expensive operation that requires saving the state, moving and reloading on another set of machines.

A novelty of our algorithm is that *we only need to maintain one queue for each type of graph*. This substantial reduction is achieved by developing an efficient method to explore the space of solutions in the scheduling space. Further, our low complexity algorithm performs updates at appropriate time instances without causing service interruptions. In summary, our approach allows a smooth trade-off between minimizing average partitioning cost and average queue sizes, by using only a small number of queues, with low complexity, and without service interruptions. As it will become clear later, the key ingredient of our method is to minimize a modified energy function instead of the Gibbs energy; specifically, the entropy term in the Gibbs energy is replaced with the relative entropy with respect to a probability distribution that arises in loss systems.

1.1 Related Work

Dynamic graph scheduling occurs in many computing settings such as Yahoo!’s S4 [25], Twitter’s Storm [29], IBM’s InfoSphere Stream [15], TimeStream [22], D-Stream [32], and online MapReduce [6]. Current scheduling solutions in this dynamic setting are primarily heuristic [2, 16, 24].

The static version of this problem (packing a collection of graphs on the machines on a one-time basis) is tightly related to *the graph partitioning problem* [4, 31], which is known to be hard. There are several algorithms (either based on heuristics or approximation bounds) available in the literature [1, 8, 14, 16, 30].

More broadly, dynamic bin packing (either scalar, or more recently vector) has a rich history [5, 7], with much recent attention [11, 27, 28]. Unlike bin packing where single items are placed into bins, our objective here is to pack graphs in a dynamic manner.

1.2 Main Contributions

We study the problem of partitioning and packing graphs over a cloud cluster when graphs arrive and depart dynamically over time. The main contributions of this work can be summarized as follows.

- **A Stochastic Model of Graph Partitioning.** We develop a stochastic model of resource allocation for graph-based applications where either the computation is represented by a graph (Storm [29], InfoSphere Stream [15]) or the data itself has a graph structure (GraphLab [13], Giraph [12]). Most efforts have been on the systems aspects, while employing a heuristic scheduler for graph partitioning and packing. One of the contributions of this paper is the model itself which allows an analytical approach towards the design of efficient schedulers.
- **Deficiencies of Max Weight-type Algorithms.** The

dynamic graph partitioning problem can be cast as a network resource allocation problem; to illustrate we describe a frame-based Max Weight algorithm that can jointly stabilize the system and minimize packing costs. However, such Max Weight-type solutions have two deficiencies:

- (1) they involve periodically solving the static graph partitioning problem (NP-hard in general); thus there is little hope that this can be implemented in practice,
- (2) they require periodic reset of the system configuration to the Max Weight configuration; this interrupts a significant number of ongoing computations or services of the jobs in the system and require them to be migrated to new machines (which is expensive).

- **Low Complexity Algorithms without Service Interruptions.** We develop a new class of low complexity algorithms, specifically targeted for the stochastic graph partitioning problems, and analytically characterize their delay and partitioning costs. In particular, the algorithms can converge to the optimal solution of the static graph partitioning problem, by trading-off delay and partitioning cost (a tunable parameter). Equally important, this class of algorithms do not interrupt the ongoing services in the system. The algorithms rely on creating and removing *templates*, where each template represents a unique way of partitioning and distributing a graph over the machines. A key ingredient of the low complexity algorithms is that the decision to remove or add templates to the system is only made at the instances that a graph is submitted to the cluster or finishes its computation; thus preventing interruption of ongoing services.

1.3 Notations

Some of the basic notations used in this paper are the following. $|S|$ denotes the cardinality of a set S . $A \setminus B$ is the set difference defined as $\{x \in A, x \notin B\}$. $\mathbb{1}\{x \in A\}$ is the indicator function which is 1 if $x \in A$, and 0 otherwise. $\mathbb{1}_n$ is the n -dimensional vector of all ones. \mathbb{R}_+ denotes the set of real nonnegative numbers. For any two probability vectors $\pi, \nu \in \mathbb{R}^n$, the total variation distance between π and ν is defined as $\|\pi - \nu\|_{TV} = \frac{1}{2} \sum_{i=1}^n |\pi_i - \nu_i|$. Further, the Kullback–Leibler (KL) divergence of π from ν is defined as $D_{KL}(\pi \|\nu) = \sum_i \pi_i \log \frac{\pi_i}{\nu_i}$. Given a stochastic process $z(t)$ which converges in distribution as $t \rightarrow \infty$, we let $z(\infty)$ denote a random variable whose distribution is the same as the limiting distribution. Given $x \in \mathbb{R}^n$, $x_{\min} = \min_i x_i$, $x_{\max} = \max_i x_i$.

2. SYSTEM MODEL AND DEFINITIONS

Cloud Cluster Model and Graph-structured Jobs: Consider a collection of machines \mathcal{L} . Each machine $\ell \in \mathcal{L}$ has a set of slots m_ℓ which it can use to run at most $|m_\ell|$ processes in parallel (see Figure 1). These machines are inter-connected by a communication network. Let $M = \sum_\ell |m_\ell|$ be the total number of slots in the cluster.

There is a collection of *jobs types* \mathcal{J} , where each job type $j \in \mathcal{J}$ is described by a *graph* $\mathcal{G}_j(V_j, E_j)$ consisting of a set of nodes V_j and a set of edges E_j . Each graph \mathcal{G}_j represents how the computation is split among the set of nodes V_j . Nodes correspond to computation with each node requiring a slot on some machine; edges represent data flows between these computations (nodes).

Job Arrivals and Departures: Henceforth, we use the word *job* and *graph* interchangeably. We assume graphs of type j arrive according to a Poisson process with rate λ_j , and will remain in the system for an exponentially distributed amount of time with mean $1/\mu_j$. Node of the graph must be assigned to an empty slot on one of the machines. Thus a graph of type \mathcal{G}_j requires a total number of $|V_j|$ free slots ($|V_j| < M$). For each graph, data center needs to decide how to partition the graph and distribute it over the machines.

Queueing Dynamics: When jobs arrive, they can either be immediately served, or queued and served at a later time. Thus, there is a set of queues $\mathbf{Q}(t) = (Q^{(j)}(t) : j \in \mathcal{J})$ representing existing jobs in the system either waiting for service or receiving service. Queues follow the usual dynamics:

$$Q^{(j)}(t) = Q^{(j)}(0) + H^{(j)}(0, t) - D^{(j)}(0, t), \quad (1)$$

where $H^{(j)}(0, t)$ and $D^{(j)}(0, t)$ are respectively the number of jobs of type j arrived up to time t and departed up to time t .

Job Partition Cost: For any job, we assume that the cost of data exchange between two nodes that are inside the same machine is zero, and the cost of data exchange between two nodes of a graph on different machines is one. This models the cost incurred by the data center due to the total traffic exchange among different machines. Note that this model is only for keeping notation simple; in fact, if we make the cost of each edge different (depending for instance on the pair of machines on which the nodes are assigned, thus capturing communication network topology constraints within the cloud cluster), there is minimal change in our description below. Specifically, we only need to redefine the appropriate cost in (2), and the ensuing analysis will remain unchanged.

Templates: An important construct in this paper is the concept of template. Observe that for any graph \mathcal{G}_j , there are several ways (exponentially large number) in which it can be *partitioned* and *distributed* over the machines (see Figure 1). A *template* corresponds to one possible way in which a graph \mathcal{G}_j can *partitioned* and *distributed* over the machines (see Figure 1). Rigorously, a template A for graph \mathcal{G}_j is an injective function $A : V_j \rightarrow \bigcup_{\ell \in \mathcal{L}} m_\ell$ which maps each node of \mathcal{G}_j to a unique slot in one of the machines. We use $\mathcal{A}^{(j)}$ to denote the set of all possible templates for graph \mathcal{G}_j . Tying back to the cost model, for $A \in \mathcal{A}^{(j)}$, let $b_A^{(j)}$ be the cost of partitioning \mathcal{G}_j according to template A , then

$$b_A^{(j)} = \sum_{(x,y) \in E_j} \mathbb{1}\{A(x) \in m_\ell, A(y) \in m_{\ell'}, \ell \neq \ell'\}. \quad (2)$$

Configuration: While there are an extremely large number of templates possible for each graph, only a limited number of templates can be present in the system at any instant of time. This is because each slot can be used by at most one template at any given time.

To track the collection of templates in the system, we let $C^{(j)}(t) \subset \mathcal{A}^{(j)}$ to be the set of existing templates of graphs \mathcal{G}_j in the system at time t . The system configuration at each time t is then defined as

$$\mathbf{C}(t) = \left(C^{(j)}(t); j \in \mathcal{J} \right). \quad (3)$$




Figure 1: Illustrative templates for partitioning and distributing a five-node graph in a cluster of 3 servers, with each server having 4 empty slots. In this stylized example, all the edges have unit “breaking” costs, i.e., two connected nodes being scheduled on different servers incurs a unit cost. The cost of partitioning the graph according to these templates is as follows: $b_{\text{Template 1}} = 1$, $b_{\text{Template 2}} = 2$, $b_{\text{Template 3}} = 3$, and $b_{\text{Template 4}} = 4$.

By definition, there is a template in the system corresponding to each job that is being served on a set of machines. Further, when a new job arrives or departs, the system can (potentially) create a new template that is a pattern of *empty slots* across machines that can be “filled” with a specific job type (i.e., one particular graph topology). We call the former as *actual templates*, and the latter as *virtual templates*. Further, when a job departs, the system can potentially destroy the associated template.

The set of all possible configurations is denoted by \mathcal{C} . Note that this collection is a union of the actual and virtual templates. Mathematically, $C^{(j)} = C_a^{(j)} \cup C_v^{(j)}$ where $C_a^{(j)}$ is the set of templates that contain *actual* jobs of type j and $C_v^{(j)}$ is the set of *virtual* templates, i.e., templates that are reserved for jobs of type j but currently do not contain any such jobs.

System State and Updates: Finally the system state at each time is then given by:

$$\mathbf{S}(t) = (\mathbf{Q}(t), \mathbf{C}(t)). \quad (4)$$

It is possible that $|C^{(j)}(t)| > Q^{(j)}(t)$ in which case not all the templates in $C^{(j)}$ are being used for serving jobs, these unused templates are the virtual templates.

Define the operation $\mathbf{C} \oplus A^{(j)}$ as adding a feasible template A for graphs of type \mathcal{G}_j to the configuration \mathbf{C} , thus A will be added to $C^{(j)}$ while $C^{(j')}$ remains unchanged for $j' \neq j$. Define $\mathcal{A}^{(j)}(\mathbf{C})$ as the set of possible templates that can be used for adding a graph \mathcal{G}_j when the configuration is \mathbf{C} . Clearly, $A \in \mathcal{A}^{(j)}(\mathbf{C})$ must be an injective function that maps graph G_j to the available slots that *have not* been used by the current templates in the system configuration, i.e.,

$$A : V_j \rightarrow \left(\bigcup_{\ell \in \mathcal{L}} m_\ell \right) \setminus \left(\bigcup_{j' \in \mathcal{J}} \bigcup_{A' \in C^{(j')}} A'(V_{j'}) \right).$$

3. PROBLEM FORMULATION

Given any stationary (and Markov) algorithm for scheduling arriving graphs, the system state evolves as an irreducible and aperiodic Markov chain. Our goal is to minimize the average partitioning cost, i.e.,

$$\begin{aligned} & \text{minimize} && \mathbb{E} \left[\sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{A}^{(j)}} x_A(\infty) b_A^{(j)} \right] & (5) \\ & \text{subject to} && \text{system stability} \end{aligned}$$

where $x_A(\infty)$ is a random variable denoting the fraction of time that a template A is used in steady state. The system stability in (5) means that the average delay (or average queue size) remains bounded. There is an inherent tradeoff between the average delay and the average partitioning cost. For more lenient delay constraints, the algorithm can defer the scheduling of jobs further until a feasible template with low partitioning cost becomes available.

Throughout the paper, let $\rho_j = \lambda_j / \mu_j$ be the load of graphs of type \mathcal{G}_j .

DEFINITION 1 (CAPACITY REGION). *The capacity region of the system is defined as*

$$\Lambda = \left\{ z \in \mathbb{R}_+^{|\mathcal{J}|} : \exists \pi \text{ s.t. } z_j = \sum_{C \in \mathcal{C}} \pi(C) |C^{(j)}|, \right. \\ \left. \sum_{C \in \mathcal{C}} \pi(C) = 1, \pi(C) \geq 0 \right\},$$

where $|C^{(j)}|$ denotes the number of templates of graph \mathcal{G}_j in configuration C .

By the definition, any load vector $z \in \Lambda$ can be supported by a proper time-sharing among the configurations, according to π . Equivalently, for any $z \in \Lambda$, there exists an $\mathbf{x} = [x_A : A \in \cup_j \mathcal{A}^{(j)}]$ such that

$$z_j = \sum_{A \in \mathcal{A}^{(j)}} x_A; \quad j \in \mathcal{J},$$

where x_A is the average fraction of time that template A is used, given by

$$x_A = \sum_{C \in \mathcal{C}} \pi(C) \mathbb{1}(A \in C^{(j)}); \quad A \in \mathcal{A}^{(j)}, \quad j \in \mathcal{J}.$$

It follows from standard arguments that for loads outside Λ , there is no algorithm that can keep the queues stable. Given the loads $\rho = [\rho_j : j \in \mathcal{J}]$, we define an associated static problem.

DEFINITION 2 (STATIC PARTITIONING PROBLEM).

$$\min_{\mathbf{x}} \quad G(\mathbf{x}) := \sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{A}^{(j)}} x_A b_A^{(j)} \quad (6)$$

$$\text{subject to:} \quad \sum_{A \in \mathcal{A}^{(j)}} x_A \geq \rho_j; \quad j \in \mathcal{J} \quad (7)$$

$$x_A \geq 0; \quad A \in \cup_j \mathcal{A}^{(j)} \quad (8)$$

$$\left[\sum_{A \in \mathcal{A}^{(j)}} x_A; \quad j \in \mathcal{J} \right] \in \Lambda \quad (9)$$

The constraints (7)-(9) are the required stability conditions. In words, given ρ_j graphs of type \mathcal{G}_j , for $j \in \mathcal{J}$, the static partitioning problem is to determine how to partition and distribute the graphs over the servers so as to minimize the total partitioning cost. For the set of supportable loads ($\rho \in$

Λ), the static problem is feasible and has a finite optimal value.

If the loads ρ_j 's are known, one can solve the static partitioning problem and subsequently find the fraction of time $\pi(C)$ that each configuration C is used. However, the static partitioning problem is a hard combinatorial problem to solve.

In the next sections, we will describe two approaches to solve the dynamic problem (5) that could converge to the optimal solution of the static partitioning problem, at the expense of growth in delay. The inherent tradeoff between delay and partitioning cost can be tuned in the algorithms. First, we describe a high complexity frame-based algorithm (based on traditional Max Weight resource allocation). Then, we proceed to propose our low complexity algorithm which is the main contribution of this paper.

4. HIGH COMPLEXITY FRAME-BASED ALGORITHM

The first candidate for solving the dynamic graph partitioning problem is to use a Max Weight-type algorithm, with a proper choice of weight for each configuration. However changing the configuration of the system can potentially interrupt a significant number of ongoing services of the jobs in the system. Such service interruptions are operationally very expensive as they incur additional delay to the service or require the storage and transfer of the state of interrupted jobs for future recovery. Hence, to reduce the cost of service interruptions, one can reduce the frequency of configuration updates. In particular, we describe a Frame-Based algorithm which updates the configuration once every T time units. As expected, a smaller value of T could improve the delay and the partitioning cost of the algorithm at the expense of more service interruptions. The description of the algorithm is as follows.

Algorithm 1 Frame-Based Algorithm

- 1: The configuration is changed at the epochs of cycles of length T . At the epoch of the k -th cycle, $k = 0, 1, \dots$, choose a configuration $C^*(kT)$ that solves

$$\max_{C \in \mathcal{C}} \sum_{j \in \mathcal{J}} \sum_{A \in C^{(j)}} \left(\alpha f(Q^{(j)}(kT)) - b_A^{(j)} \right). \quad (10)$$

If there are more than one optimal configuration, one of them is chosen arbitrarily at random. $\alpha > 0$ is a fixed parameter and f is a concave increasing function.

- 2: The configuration $C^*(kT)$ is kept fixed over the interval $[kT, (k+1)T)$ during which jobs are fetched from the queues and are placed in the available templates. It is possible that at some time, no jobs of type j are waiting to get service, in which case some of the templates in $C^*(kT)$ might not be filled with the actual jobs. These are virtual templates which act as place holders (tokens) for future arrivals.
-

The algorithm essentially needs to find a maximum weight configuration at the epochs of cycles, where the weight of template A for partitioning graph \mathcal{G}_j is

$$w_A^{(j)}(t) = \alpha f(Q^{(j)}(t)) - b_A^{(j)}.$$

The parameter α controls the tradeoff between the queue

size and the partitioning cost of the algorithm. For small values of α , the algorithm defers deploying the graphs in favor of finding templates with smaller partitioning cost. For larger values of α , the algorithm gives a higher priority to deployment of job types with large queue sizes.

The optimization (10) is a hard combinatorial problem, as the size of the configuration space \mathcal{C} might be exponentially large, thus hindering efficient computation of the max weight configuration. Theorem 1 below characterizes the inherent tradeoff between the average queue size and the average partitioning cost.

THEOREM 1. *Suppose $\rho(1+\delta^*) \in \Lambda$ for some $\delta^* > 0$. The average queue size and the average partitioning cost under the Frame-Based algorithm are*

$$\mathbb{E}\left[\sum_{j \in \mathcal{J}} f(Q^{(j)}(\infty))\right] \leq \frac{(B_1 + B_2 T) + (1 + \delta^*)G(x^*)/\alpha}{\delta^* \rho_{\min}} + B_1 T$$

$$\mathbb{E}\left[\sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{A}^{(j)}} x_A(\infty) b_A^{(j)}\right] \leq G(x^*) + \alpha(B_1 + B_2 T)$$

where x^* is the optimal solution to the static partitioning problem, $\rho_{\min} = \min_j \rho_j$, and B_1, B_2 are constants.

Hence, as $\alpha \rightarrow 0$, the algorithm yields an α -optimal partitioning cost, and an $O\left(f^{-1}\left(\frac{1}{\alpha}\right)\right)$ queue size. Also as expected, infrequent configuration updates could increase the delay and partitioning cost by multiples of T . The proof of Theorem 1 follows from standard Lyapunov arguments and can be found in the appendix.

5. LOW COMPLEXITY ALGORITHMS WITHOUT SERVICE INTERRUPTIONS

In this section, we develop a low complexity algorithm that can be used to solve (5) without interrupting/migrating the ongoing services. Before describing the algorithm, we first introduce a (modified) *weight* for each template. Given the vector of queue sizes $\mathbf{Q}(t)$, and a concave increasing function $f: \mathbb{R}_+ \rightarrow \mathbb{R}_+$, the weight of template $A \in \mathcal{A}^{(j)}$, $j \in \mathcal{J}$, is defined as

$$\tilde{w}_A^{(j)}(t) = \alpha f^{(j)}(\mathbf{h} + \mathbf{Q}(t)) - b_A^{(j)}, \quad (11)$$

where $f^{(j)}: \mathbb{R}_+^{|\mathcal{J}|} \rightarrow \mathbb{R}_+$ is

$$f^{(j)}(x) = \max\left\{f(x_j), \frac{\epsilon}{8M} f(x_{\max})\right\}; \quad x_{\max} = \max_{j \in \mathcal{J}} x_j, \quad (12)$$

where $\mathbf{h} = h\mathbf{1}_{|\mathcal{J}|}$, and $\alpha, h \in \mathbb{R}_+$, and $\epsilon \in (0, 1)$ are the parameters of the algorithm.

At the instances of job arrivals and departures, the algorithm makes decisions on the templates that are added to/removed from the system configuration. It is important that the addition/removal of templates by the algorithm *does not* disrupt the ongoing service of existing jobs in the configuration.

The low complexity algorithm is a randomized algorithm in which the candidate template to be added to the configuration is chosen randomly among the set of feasible templates. In particular, the following *Random Partition Procedure* is used as a subroutine in our low complexity algorithm.

Algorithm 2 Random Partition Procedure

Input: current configuration \mathcal{C} , and a graph $\mathcal{G}(V, E)$, $V = \{v_1, \dots, v_{|V|}\}$

Output: a virtual template A for distributing \mathcal{G} over the machines.

```

1:  $k \leftarrow 1$ 
2: slot-available  $\leftarrow 1$ 
3: while  $k \leq |V|$  and slot-available do
4:   if there are no free slots available on any of the machines then
5:     slot-available  $\leftarrow 0$ ;  $A \leftarrow \emptyset$ 
6:   else
7:     place  $v_k$  uniformly at random in one of the free slots
8:      $A(v_k) =$  index of the slot containing  $v_k$ 
9:      $k \leftarrow k + 1$ 
10:  end if
11: end while

```

When a random template is generated according to Random Partition Procedure, the decision to keep or remove the template is made probabilistically based on the weight of the template. The description of the low complexity algorithm (called *Dynamic Graph Partitioning (DGP) algorithm*) is as follows. In the the description, $\beta > 0$ is a fixed parameter.

Algorithm 3 Dynamic Graph Partitioning (DGP)

Arrival instances. Suppose a graph (job) \mathcal{G}_j arrives at time t , then:

- 1: This job is added to queue $Q^{(j)}$.
- 2: A virtual graph \mathcal{G}_j is randomly distributed over the machines, if possible, using *Random Partition Procedure*, which creates a virtual template $A^{(j)}$ for distributing a graph \mathcal{G}_j over the machines with some partitioning cost $b_A^{(j)}$. Then, this virtual template is added to the current configuration with probability $\frac{\exp(\frac{1}{\beta} \tilde{w}_A^{(j)}(t^+))}{1 + \exp(\frac{1}{\beta} \tilde{w}_A^{(j)}(t^+))}$, otherwise, it is discarded and the configuration does not change. The virtual templates of type j leave the system after an exponentially distributed time duration with mean $1/\mu_j$.
- 3: If there is one or more virtual templates available for accommodating graphs of type \mathcal{G}_j , a job from $Q^{(j)}$ (e.g., the head-of-the-line job) is placed in one of the virtual templates chosen arbitrarily at random. This converts the virtual template to an actual template.

Departure instances. Suppose a departure of a (virtual or actual) template $A^{(j)}$ occurs at time t , then:

- 1: If this an actual template, the job departs and queue $Q^{(j)}(t^+)$ is updated.
 - 2: A virtual template of the same type $A^{(j)}$ is added back to the configuration with probability $\frac{\exp(\frac{1}{\beta} \tilde{w}_A^{(j)}(t^+))}{1 + \exp(\frac{1}{\beta} \tilde{w}_A^{(j)}(t^+))}$.
 - 3: If a virtual template for accommodating a graph \mathcal{G}_j is available in the system, and there are jobs in $Q^{(j)}(t^+)$ waiting to get service, a job from $Q^{(j)}(t^+)$ (e.g. the head-of-the line job) is placed in one of the virtual templates chosen arbitrarily at random. This converts the virtual template to an actual template.
-

To simplify the description, we have assumed that the system starts from empty initial configuration and empty queues but this is not necessary for the results to hold. We emphasize that the DGP algorithm does not interrupt the ongoing services of existing jobs in the system. The following theorem states our main result regarding the performance of the algorithm.

THEOREM 2. *Suppose $\rho(1+\delta^*) \in \Lambda$ for some $0 < \delta^* < 1$. Consider the Dynamic Graph Partitioning (DGP) algorithm with function*

$$f(x) = \log^{1-b}(x); \quad b \in (0, 1),$$

and parameters

$$\alpha \leq \beta < 1; \quad \epsilon \leq \delta^*; \quad h \geq \exp\left(C_0 \frac{1}{\beta} \left(\frac{1}{\epsilon}\right)^{\frac{2-b+1/b}{1-b}}\right),$$

where C_0 is a large constant independent of all these parameters. Then the average queue size and the average partitioning cost under the DGP algorithm are

$$\begin{aligned} \sum_{j \in \mathcal{J}} \mathbb{E} \left[f(Q_j(\infty)) \right] &\leq \frac{2}{\rho_{\min} \delta^*} \left(\hat{K}_2 + \hat{K}_3 - \frac{\beta}{\alpha} \log \gamma_{\min} + \right. \\ &\left. \frac{1}{\alpha} (1 + \delta^*/2) G(x^*) + \frac{\epsilon}{\alpha} b_{\max} \right), \\ \mathbb{E} \left[\sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{A}^{(j)}} x_A(\infty) b_A^{(j)} \right] &\leq G(x^*) + \alpha (\hat{K}_2 + \hat{K}_3) \\ &- \beta \log \gamma_{\min} + \epsilon b_{\max}, \end{aligned}$$

where x^* is the optimal solution to the static partitioning problem, $\rho_{\min} = \min_j \rho_j$, $\hat{K}_2 \leq f'(h)(M + \sum_j \rho_j)$ and $\hat{K}_3 \leq f(M+h)M$, and γ_{\min} , and b_{\max} are constants.

We would like to point out that in the above theorem the bounds are explicit for any choices of $\alpha, \beta, \epsilon, h$. The constant γ_{\min} is $\min_C \gamma_C$ for a distribution γ to be defined in (14) and has a loss-system interpretation (see Step 1 in the Proof of Theorem 2), and b_{\max} is the maximum partitioning cost of any job type (which is obviously less than M^2).

The parameter h is called the *bias* and adds an offset to the queues to ensure the algorithm operates near the optimal point at (effectively) all times. The parameter β has the similar role as the temperature in Gibbs sampler. As $\beta \rightarrow 0$, in steady state, the algorithm generates configurations that are closer to the optimal configuration, however at the expense of growth in queue sizes. We refer to Section 6 for the proof and also more insight into the operation of the algorithm.

The following corollary gives an interpretation of the result for a particular choice of the parameters.

COROLLARY 1. *Choose $\alpha = \beta^2$, $h = \exp\left(\left(\frac{1}{\beta}\right)^{1/(1-b)}\right)$, $\epsilon = \beta^{b^2/4}$, then as $\beta \rightarrow 0$,*

$$\begin{aligned} \sum_{j \in \mathcal{J}} \mathbb{E} \left[f(Q_j(\infty)) \right] &\leq \Theta\left(\left(\frac{1}{\beta}\right)^2\right), \\ \mathbb{E} \left[\sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{A}^{(j)}} x_A(\infty) b_A^{(j)} \right] &\leq G(x^*) + \Theta(\beta^{b^2/4}). \end{aligned}$$

The corollary above demonstrates how the choice of β controls the tradeoff between approaching the optimal partitioning cost and the queueing performance.

Remark 1. Comparison with CSMA: In the setting of scheduling in wireless networks, the Gibbs sampler has been used to design CSMA-like algorithms for stabilizing the network [17, 19, 23, 26]. Our algorithm is different from this line of work in three fundamental aspects:

(i) Not relying on Gibbs sampler: Unlike wireless networks where the solutions form independent sets of a graph, there is no natural graph structure analog in the graph partitioning. The Gibbs sampler (and CSMA) can still be used in our setting by sampling *partitions* of graphs, where, each site of the Gibbs sampler is a unique way of partitioning and packing a graph among the machines. The difficulty, however, is that there are an exponentially large number of graph partitions, leading to a correspondingly large number of queues for each type of graph. A novelty of our algorithm is that *we only need to maintain one queue for each type of graph*. This substantial reduction is achieved by using *Random Partition Procedure* for exploring the space of solutions. This leads to minimizing a modified energy function instead of the Gibbs energy; specifically, the entropy term in the Gibbs energy is replaced with the relative entropy with respect to a probability distribution that arises in an associated loss system (see Step 1 in Section 6).

(ii) No service interruptions: Our low complexity algorithm performs updates at appropriate time instances without causing service interruptions.

(iii) Adding bias to the queues: The queue-based CSMA algorithms are concerned with stability which pertains to the behavior of the algorithm for large queue sizes. This is not sufficient in our setting because we are not only concerned with stability, but more importantly with the optimal (graph partitioning) cost of the system. The bias h boosts the queue sizes artificially to ensure that the system operates effectively near the optimal point at *all queue sizes*. Without the bias, when the queue sizes are small, the optimal cost of the algorithm could be far from optimal.

Remark 2. An Alternative Algorithm: An alternative description of the algorithm is possible using a dedicated Poisson clock for each queue (independent of arrivals) where the template decisions are made at the ticks of the dedicated clocks. We have presented this alternative algorithm in the appendix.

6. PROOFS

In this section, we present the proof of Theorem 2. Before describing the proof outline, we make the following definition.

Definition: $\overline{\text{DGP}}(\tilde{W})$. Consider the dynamic graph partitioning algorithm with fixed weights $\tilde{W} = [\tilde{w}_A^{(j)}; A \in \mathcal{A}^{(j)}, j \in \mathcal{J}]$, namely, when weights are not chosen according to (11) but they are simply some fixed numbers all the time. With minor abuse of notations, we use $\overline{\text{DGP}}(\tilde{W})$ to denote this algorithm that uses weights \tilde{W} all the time. Description of $\overline{\text{DGP}}(\tilde{W})$ is exactly the same as the dynamic partitioning algorithm, except that at arrival/departure instance at time t , the decision to add/keep a virtual template $A^{(j)}$ is made according to probability $\frac{\exp(\frac{1}{\beta} \tilde{w}_A^{(j)})}{1 + \exp(\frac{1}{\beta} \tilde{w}_A^{(j)})}$, independently of $Q(t)$.

Proof Outline. The proof of Theorem 2 has three steps:

Step 1: We analyze the steady-state distribution of configurations under $\overline{\text{DGP}}(\tilde{W})$ with fixed weights \tilde{W} , and show

that for small values of β , $\overline{\text{DGP}}(\tilde{W})$ will generate configurations which are “close” to the max weight configuration, when the template weights are per \tilde{W} .

Step 2: We show that when weights are chosen according to (11), although the weights $\tilde{W}(t)$ are time-varying, the distribution of configurations in the system will be “close” to the corresponding steady-state distribution of $\overline{\text{DGP}}(\tilde{W}(t))$, for *all times* t long enough. We show that such “time-scale decomposition” holds under the suitable choice of the bias h and the function f .

Step 3: Finally, we stitch the dynamics of queues and configurations together through Lyapunov optimization method to compute the queuing and partitioning cost of our algorithm.

Step 1: Steady-State Analysis of $\overline{\text{DGP}}(\tilde{W})$

Under $\overline{\text{DGP}}(\tilde{W})$, the configuration of the system evolves as a “time-homogeneous” Markov chain over the state space \mathcal{C} . Note that from the perspective of evolution of configuration in the system, we do not need to distinguish between virtual and actual templates, since transition rates from any configuration C do not depend on whether the templates in C are actual or virtual. To see this, consider any virtual template of graphs \mathcal{G}_j in $C(t)$. No matter if the virtual template is filled with an actual job or not, the residual time until the departure of this template is still exponential with rate μ_j , due to the memoryless property of exponential distribution and because both virtual templates and jobs have exponential service times with the same mean $1/\mu_j$. The following proposition states the main property of $\overline{\text{DGP}}(\tilde{W})$.

PROPOSITION 1. *Consider the $\overline{\text{DGP}}(\tilde{W})$ with fixed weights $\tilde{W} = [\tilde{w}_A^{(j)}; A \in \mathcal{A}^{(j)}, j \in \mathcal{J}]$. Then in steady state, the distribution of configurations π will solve the following optimization problem*

$$\max_{\pi \in \mathbb{R}_+^{|\mathcal{C}|}, \sum_{C \in \mathcal{C}} \pi_C = 1} \mathbb{E}_\pi \left[\sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{C}^{(j)}} \tilde{w}_A^{(j)} \right] - \beta D_{KL}(\pi \parallel \gamma), \quad (13)$$

where $D_{KL}(\cdot \parallel \cdot)$ is the KL divergence of π from the probability distribution γ , where

$$\gamma_C = \frac{1}{Z_\gamma} \left(\sum_{\ell} |m_\ell| - \sum_j |C^{(j)}| |V_j| \right)! \prod_j \rho_j^{|C^{(j)}|}, \quad C \in \mathcal{C} \quad (14)$$

and Z_γ is the normalizing constant.

Before describing the proof of Proposition 1, we briefly highlight the main features of $\overline{\text{DGP}}(\tilde{W})$ algorithm:

- (i) The algorithm does not interrupt the ongoing services of existing jobs in the system and does not require dedicated computing resources.
- (ii) The algorithm is different from Gibbs sampler as it does not maximize the Gibbs energy. The entropy term $H(\pi)$ in the Gibbs energy has been replaced by the relative entropy $D_{KL}(\pi, \gamma)$.
- (iii) The distribution γ has the interpretation of the steady-state distribution of configurations in an *associated loss system* defined as follows: at arrival instances, the arriving graph is randomly distributed over the machines if possible (according to Random Partition Procedure),

otherwise it is dropped; at the departure instances, the job (and hence its template) leaves the system.

PROOF OF PROPOSITION 1. Consider the maximization problem

$$\begin{aligned} & \max_{\{\pi(C)\}} F^{(\beta)}(\pi) \\ & \text{subject to} \quad \sum_{C \in \mathcal{C}} \pi(C) = 1 \\ & \quad \quad \quad \pi(C) \geq 0, \quad \forall C \in \mathcal{C}. \end{aligned}$$

with function $F^{(\beta)}(\pi)$ as in (13), which is

$$\begin{aligned} F^{(\beta)}(\pi) &= \sum_C \pi(C) \sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{C}^{(j)}} \tilde{w}_A^{(j)} - \beta \sum_C \pi(C) \log \pi(C) \\ & \quad + \beta \sum_C \pi(C) \log \gamma(C). \end{aligned}$$

Notice that $F^{(\beta)}(\pi)$ is strictly concave in π . The lagrangian is given by $L(\pi, \eta) = F^{(\beta)}(\pi) + \eta(\sum_C \pi(C) - 1)$ where $\eta \in \mathbb{R}$ is the lagrange multiplier. Taking $\partial L / \partial \pi(C) = 0$ yields

$$\pi(C) = \exp(-1 + \frac{\eta}{\beta}) \gamma(C) \exp(\frac{1}{\beta} \sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{C}^{(j)}} \tilde{w}_A^{(j)}); \quad C \in \mathcal{C},$$

which is automatically nonnegative for any η . Hence, by KKT conditions (π^*, η^*) is the optimal primal-dual pair if it satisfies $\sum_C \pi^*(C) = 1$. Thus the optimal distribution π^* is

$$\pi^*(C) = \frac{1}{Z_\beta} \gamma(C) \exp(\frac{1}{\beta} \sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{C}^{(j)}} \tilde{w}_A^{(j)}). \quad (15)$$

where Z_β is the normalizing constant.

Next we show that the $\overline{\text{DGP}}(\tilde{W})$ algorithm indeed produces the steady-state distribution (15) with the choice of γ in (14), by checking the detailed balance equations. Consider a template $A^{(j)}$ for graphs of type \mathcal{G}_j . The detailed balanced equation for the pair C and $C \oplus A^{(j)}$, such that $C \oplus A^{(j)} \in \mathcal{C}$, is given by

$$\pi(C \oplus A^{(j)}) \mu_j \frac{1}{1 + e^{\frac{1}{\beta} \tilde{w}_A^{(j)}}} = \pi(C) \frac{\lambda_j}{|\mathcal{A}^{(j)}(C)|} \frac{e^{\frac{1}{\beta} \tilde{w}_A^{(j)}}}{1 + e^{\frac{1}{\beta} \tilde{w}_A^{(j)}}}.$$

The left-hand-side is the departure rate of (virtual or actual) template $A^{(j)}$ from the configuration $C \oplus A^{(j)}$. The right-hand-side is the arrival rate of (actual or virtual) graphs \mathcal{G}_j to the configuration C that are deployed according to template $A^{(j)}$ chosen uniformly at random from $\mathcal{A}^{(j)}(C)$ (Recall that Random Partition Procedure used in the algorithm selects a template $A \in \mathcal{A}^{(j)}(C)$ uniformly at random). Thus the detailed balanced equation is simply

$$\pi(C \oplus A^{(j)}) = \pi(C) \frac{\rho_j}{|\mathcal{A}^{(j)}(C)|} e^{\frac{1}{\beta} \tilde{w}_A^{(j)}}. \quad (16)$$

Noting that

$$|\mathcal{A}^{(j)}(C)| = \binom{\sum_{\ell} |m_\ell| - \sum_j |C^{(j)}| |V_j|}{|V_j|} |V_j|!,$$

it is then easy to see that (15) with γ as in (14), indeed satisfies the detailed balance equations, and the normalizing condition that $\sum_C \pi(C) = 1$. This concludes the proof. \square

The parameter β has the similar role as the temperature in Gibbs sampler. As $\beta \rightarrow 0$, in steady state, the $\overline{\text{DGP}}(\tilde{W})$

algorithm generates configurations that are closer to the optimal configuration with maximum weight

$$\tilde{W}^* = \max_{C \in \mathcal{C}} \sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{C}^{(j)}} \tilde{w}_A^{(j)}. \quad (17)$$

The following corollary contains this result.

COROLLARY 2. *Let $F^{(\beta)}(\pi^{*(\beta)})$ be the optimal objective function in (13). The algorithm $\overline{\text{DGP}}(\tilde{W})$ is asymptotically optimal in the sense as $\beta \rightarrow 0$, $F^{(\beta)}(\pi^{*(\beta)}) \rightarrow \tilde{W}^*$. Moreover, for any $\beta > 0$,*

$$\mathbb{E}_{\pi^{*(\beta)}} \left[\sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{C}^{(j)}} \tilde{w}_A^{(j)} \right] \geq \max_{C \in \mathcal{C}} \sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{C}^{(j)}} \tilde{w}_A^{(j)} + \beta \min_{C \in \mathcal{C}} \log \gamma_C$$

PROOF OF COROLLARY 2. Let \tilde{C}^* be the maximizer in (17). As a direct consequence of Proposition 1,

$$\mathbb{E}_{\pi^{*(\beta)}} \left[\sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{C}^{(j)}} \tilde{w}_A^{(j)} \right] - \beta D(\pi^{*(\beta)} \parallel \gamma) \geq \tilde{W}^* - \beta D(\delta_{\tilde{C}^*} \parallel \gamma).$$

Since $D(v \parallel \gamma) \geq 0$, for any distribution v ,

$$\begin{aligned} \mathbb{E}_{\pi^{*(\beta)}} \left[\sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{C}^{(j)}} \tilde{w}_A^{(j)} \right] &\geq \tilde{W}^* - \beta D(\delta_{\tilde{C}^*} \parallel \gamma) \\ &= \tilde{W}^* + \beta \log \gamma_{\tilde{C}^*} \\ &\geq \tilde{W}^* + \beta \min_{C \in \mathcal{C}} \log \gamma_C. \end{aligned}$$

□

Step 2: Time-Scale Decomposition for $\text{DGP}(\tilde{W}(t))$.

Recall that $\overline{\text{DGP}}(\tilde{W}(t))$ denotes the algorithm that uses the weight $\tilde{W}(t)$ at all times $s \geq 0$. With minor abuse of notation, we use $\text{DGP}(\tilde{W}(t))$ to denote the Dynamic Graph Partitioning algorithm (Section 5) and its associated time-inhomogeneous Markov chain over the space of configurations \mathcal{C} . The weights $\tilde{W}(t)$ are time-varying (because of the queue dynamics), however the $\text{DGP}(\tilde{W}(t))$ algorithm can still provide an adequately accurate approximation to the optimization (17) at each time, for proper choices of function f and the bias h .

Roughly speaking, for the proper choices of f and h , $f(h + Q^{(j)}(t))$ will change adequately slowly with time such that a time-scale separation occurs, i.e., convergence of Markov chain $\overline{\text{DGP}}(\tilde{W}(t))$ to its steady state distribution $\pi_t^{(\beta)}$ will occur at a much faster time-scale than the time-scale of changes in $f(h + Q^{(j)}(t))$ (and thus in the weights). Hence, the probability distribution of configurations under $\text{DGP}(\tilde{W}(t))$ will remain “close” to $\pi_t^{(\beta)}$ (the steady state distribution of configurations under $\overline{\text{DGP}}(\tilde{W}(t))$). The proof of such a time-scale separation follows from standard arguments in e.g., [9, 10, 23].

We first uniformize (e.g. [18, 21]) the continuous Markov chain $\mathbf{S}(t) = (\mathbf{C}(t), \mathbf{Q}(t))$ by using a Poisson clock $N_\xi(t)$ of rate

$$\xi = 2 \left(\sum_j \lambda_j + M \sum_j \mu_j \right). \quad (18)$$

Let $\mathbf{S}[k] = (\mathbf{C}[k], \mathbf{Q}[k])$ be the corresponding jump chain of the uniformized chain. Note that $\mathbf{S}[k]$ is discrete time and at each index k , either a graph \mathcal{G}_j arrives with probability $\frac{\lambda_j}{\xi}$, or a (virtual/actual) template of type j leaves the

system with probability $\frac{|C^{(j)}| \mu_j}{\xi}$, or $\mathbf{S}[k]$ remains unchanged otherwise. The following proposition states the main “time-scale decomposition” property with respect to the associated jump chain (which can be naturally mapped to the original Markov chain).

PROPOSITION 2. *Let ν_n denote the (conditional) probability distribution of configuration at index n given the queues $\mathbf{Q}[n]$ under $\text{DGP}(\tilde{W}(\mathbf{Q}[n]))$. Let π_n be the steady state distribution of configurations corresponding to $\overline{\text{DGP}}(\tilde{W}(\mathbf{Q}[n]))$. Given any $0 < \epsilon < 1$, and any initial state $\mathbf{S}[0] = (\mathbf{Q}[0], \mathbf{C}[0])$, there exists a time $n^* = n^*(\epsilon, \beta, \mathbf{S}[0])$ such that for all $n \geq n^*$, $\|\pi_n - \nu_n\|_{TV} \leq \epsilon/16$.*

COROLLARY 3. *Given $0 < \epsilon < 1$, for all $n \geq n^*(\epsilon, \beta, \mathbf{S}(0))$,*

$$\begin{aligned} \mathbb{E}_{\nu_n} \left[\sum_j \sum_{A \in \mathcal{C}^{(j)}} w_A^{(j)}(n) \right] &\geq \beta \min_{C \in \mathcal{C}} \log \gamma_C - \epsilon b_{\max} \\ &\quad + \left(1 - \frac{\epsilon}{4}\right) \max_{C \in \mathcal{C}} \sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{C}^{(j)}} w_A^{(j)}(n). \end{aligned}$$

PROOF. Consider any $n \geq n^*(\epsilon, \beta, \mathbf{S}(0))$. Let $\tilde{W}^*(n) := \max_{C \in \mathcal{C}} \sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{C}^{(j)}} \tilde{w}_A^{(j)}(n)$. First note that from Corollary 2, Proposition 2, and definition of $\|\cdot\|_{TV}$,

$$\begin{aligned} \mathbb{E}_{\nu_n} \left[\sum_j \sum_{A \in \mathcal{C}^{(j)}} \tilde{w}_A^{(j)}(n) \right] &= \mathbb{E}_{\pi_n} \left[\sum_j \sum_{A \in \mathcal{C}^{(j)}} \tilde{w}_A^{(j)}(n) \right] \\ &\quad + \sum_C \left[(\pi_n(C) - \nu_n(C)) \sum_j \sum_{A \in \mathcal{C}^{(j)}} \tilde{w}_A^{(j)}(n) \right] \\ &\geq \tilde{W}^*(n) + \beta \min_{C \in \mathcal{C}} \log \gamma_C - 2 \left(\frac{\epsilon}{16} \right) \tilde{W}^*(n) \\ &= \left(1 - \frac{\epsilon}{8}\right) \tilde{W}^*(n) + \beta \min_{C \in \mathcal{C}} \log \gamma_C. \end{aligned} \quad (19)$$

Next, note that by the definition of $w_A^{(j)}$ (see (11), (12)), for any $j \in \mathcal{J}$, $A \in \mathcal{A}^{(j)}$,

$$w_A^{(j)}(n) \leq \tilde{w}_A^{(j)}(n) \leq w_A^{(j)}(n) + \frac{\alpha \epsilon}{8M} f(Q_{\max}(n) + h),$$

hence for any configuration $C \in \mathcal{C}$,

$$0 \leq \sum_{j \in \mathcal{J}} \sum_{C \in \mathcal{C}^{(j)}} (\tilde{w}_A^{(j)}(n) - w_A^{(j)}(n)) \leq \alpha \frac{\epsilon}{8} f(Q_{\max}(n) + h).$$

Suppose $Q^{(j')}(n) = Q_{\max}(n)$ for some $j' \in \mathcal{J}$. Then for any $A' \in \mathcal{A}^{(j')}$,

$$\begin{aligned} \alpha f(Q_{\max}(n) + h) - b_{A'}^{(j')} &= w_{A'}^{(j')}(n) \\ &\leq \max_{C \in \mathcal{C}} \sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{C}^{(j)}} w_A^{(j)}(n). \end{aligned}$$

Therefore, it follows that

$$\begin{aligned} 0 \leq \sum_{j \in \mathcal{J}} \sum_{C \in \mathcal{C}^{(j)}} (\tilde{w}_A^{(j)}(n) - w_A^{(j)}(n)) &\leq \frac{\epsilon}{8} \max_C \sum_{j \in \mathcal{J}} \sum_{C \in \mathcal{C}^{(j)}} w_A^{(j)} \\ &\quad + \frac{\epsilon}{8} b_{\max}. \end{aligned}$$

Let $W^*(n) := \max_{C \in \mathcal{C}} \sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{C}^{(j)}} w_A^{(j)}(n)$. Using the

above inequality and (19),

$$\begin{aligned}
& \mathbb{E}_{\nu_n} \left[\sum_j \sum_{A \in C^{(j)}} w_A^{(j)}(n) \right] \\
& \geq \mathbb{E}_{\nu_n} \left[\sum_j \sum_{A \in C^{(j)}} \tilde{w}_A^{(j)}(n) \right] - \frac{\epsilon}{8} W^*(n) - \frac{\epsilon}{8} b_{max} \\
& \geq (1 - \frac{\epsilon}{8}) \tilde{W}^*(n) + \beta \log \gamma_{min} - \frac{\epsilon}{8} W^*(n) - \frac{\epsilon}{8} b_{max} \\
& \geq (1 - \frac{\epsilon}{4}) W^*(n) + \beta \log \gamma_{min} - \frac{\epsilon}{8} b_{max}.
\end{aligned}$$

□

PROOF OF PROPOSITION 2. Below we mention a sketch of the proof of the “time-scale decomposition” property for our algorithm.

Let Φ^Q be the infinitesimal generator of the Markov chain $(C(t))$ under $DGP(\tilde{W}(Q))$, for some vector of queues Q . Let $P_\xi^Q = I + \frac{1}{\xi} \Phi^Q$ denote the corresponding transition probability matrix of the jump chain $(C[n])$, obtained by uniformizing $(C(t))$ using the Poisson clock $N_\xi(t)$ of rate ξ in (18). We use $P_\xi^Q(C, C')$ to denote the transition probability from configuration C to configuration C' .

The Markov chain $(C[n])$ is irreducible, aperiodic, and reversible, with the unique steady-state distribution π in (15). In this case, it is well known that the convergence to the steady-state distribution is geometric with a rate equal to the Second Largest Eigenvalue Modulus (SLEM) of P_ξ^Q [3]. Further, using the choice of ξ in (18), $(C[n])$ is a lazy Markov chain because at each jump index n , the chain will remain in the same state with probability greater than $1/2$. In this case, for any initial probability distribution μ_0 and for all $n \geq 0$,

$$\|\mu_0 (P_\xi^Q)^n - \pi\|_{TV} \leq \theta_2^n \frac{1}{2\sqrt{\pi_{min}}}, \quad (20)$$

where θ_2 is the second largest eigenvalue of P_ξ^Q , and $\pi_{min} = \min_C \pi(C)$. Correspondingly, the mixing time of the chain (defined as $\inf\{n > 0 : \|\nu(n) - \pi(n)\|_{TV} \leq \delta\}$) will be less than $\frac{-\log(2\delta\sqrt{\pi_{min}})}{(1-\theta_2)}$.

Lemma 1 below provides a bound on θ_2 and hence on the convergence rate of Markov chain P_ξ^Q .

LEMMA 1. Let $K_0 = \left(\frac{\rho_{min} \wedge 1}{\rho_{max} \vee 1}\right)^M \frac{\wedge_j(\mu_j \wedge \lambda_j)}{|C|(M!)^2}$. Then,

$$\frac{1}{1 - \theta_2} \leq \frac{2\xi^2}{K_0^2} \exp \left[\frac{2(M+1)}{\beta} (\alpha f(Q_{max} + h) + b_{max}) \right]. \quad (21)$$

PROOF OF LEMMA 1. It follows from Cheeger’s inequality [3] that $\frac{1}{1-\theta_2} \leq \frac{2}{\Psi^2(P_\xi)}$ where $\Psi(P_\xi)$ is the conductance of the Markov chain P_ξ^Q . The conductance is further bounded from below as

$$\Psi(P_\xi) \geq 2\pi_{min} \min_{C \neq C'} P_\xi^Q(C, C'). \quad (22)$$

Under $\overline{DGP}(\tilde{W})$, with $\tilde{W} = \tilde{W}(Q)$,

$$\begin{aligned}
\min_{C \neq C'} P_\xi^Q(C, C') &= \frac{1}{\xi} \min_j \frac{\mu_j}{1 + e^{\frac{1}{\beta} \tilde{w}_A^{(j)}}} \wedge \frac{\lambda_j}{|\mathcal{A}^{(j)}(C)|} \frac{e^{\frac{1}{\beta} \tilde{w}_A^{(j)}}}{1 + e^{\frac{1}{\beta} \tilde{w}_A^{(j)}}} \\
&\geq \frac{\wedge_j(\mu_j \wedge \lambda_j) \exp(\frac{1}{\beta} \tilde{w}_{min}) \wedge 1}{\xi M!} \frac{1}{1 + \exp(\frac{1}{\beta} \tilde{w}_{max})} \\
&\geq \frac{\wedge_j(\mu_j \wedge \lambda_j)}{\xi M!} \frac{\exp(-\frac{1}{\beta} b_{max})}{1 + \exp(\frac{\alpha}{\beta} f(Q_{max} + h))}
\end{aligned}$$

Note that the steady state distribution of the jump chain is still $\pi(C) = \frac{\gamma(C)}{Z_\beta} \exp(\frac{1}{\beta} \sum_j \sum_{A \in C} \tilde{w}_A^{(j)})$, for γ defined in (14). Then

$$\begin{aligned}
Z_\gamma Z_\beta &\leq \sum_{C \in \mathcal{C}} \exp\left(\frac{1}{\beta} \sum_j \sum_{A \in C} \tilde{w}_{max}\right) M! (\rho_{max} \vee 1)^M \\
&\leq |\mathcal{C}| \exp\left(\frac{M\alpha}{\beta} f(Q_{max} + h)\right) M! (\rho_{max} \vee 1)^M,
\end{aligned}$$

therefore

$$\pi_{min} \geq K_1 \exp\left(-\frac{M\alpha}{\beta} f(Q_{max} + h) - \frac{M}{\beta} b_{max}\right), \quad (23)$$

where $K_1 = \left(\frac{\rho_{min} \wedge 1}{\rho_{max} \vee 1}\right)^M \frac{1}{|C|M!}$. Hence

$$\Psi(P_\xi^Q) \geq \frac{K_0}{\xi} \exp\left(-\frac{M+1}{\beta} (\alpha f(Q_{max} + h) + b_{max})\right).$$

where $K_0 = K_1 \frac{\wedge_j(\mu_j \wedge \lambda_j)}{M!}$. □

LEMMA 2. For any configuration $C \in \mathcal{C}$, $e^{-\sigma_n} \leq \frac{\pi_{n+1}(C)}{\pi_n(C)} \leq e^{\sigma_n}$, where

$$\sigma_n = \frac{2M\alpha}{\beta} f'\left(f^{-1}\left(\frac{\epsilon}{8M} f(h + Q_{max}(n+1))\right) - 1\right). \quad (24)$$

PROOF OF LEMMA 2. Note that

$$\frac{\pi_{n+1}(C)}{\pi_n(C)} = \frac{Z_n(\beta)}{Z_{n+1}(\beta)} e^{\frac{\alpha}{\beta} \sum_{j, A \in C^{(j)}} f^{(j)}(Q(n+1)+h) - f^{(j)}(Q(n)+h)}.$$

It is easy to show that

$$\frac{Z_n(\beta)}{Z_{n+1}(\beta)} \leq \max_C e^{\frac{\alpha}{\beta} \sum_{j, A \in C^{(j)}} f^{(j)}(Q(n+1)+h) - f^{(j)}(Q(n)+h)}.$$

Let $Q^*(n) := f^{-1}\left(\frac{\epsilon}{8M} f(h + Q_{max}(n))\right) - h$, and define $\tilde{Q}^{(j)}(n) := \max\{Q^*(n), Q^{(j)}(n)\}$. Then,

$$\begin{aligned}
& f^{(j)}(Q(n+1) + h) - f^{(j)}(Q(n) + h) \\
&= f(\tilde{Q}^{(j)}(n+1) + h) - f(\tilde{Q}^{(j)}(n) + h) \\
&\leq f'(\tilde{Q}^{(j)}(n+1) + h - 1) |\tilde{Q}^{(j)}(n+1) - \tilde{Q}^{(j)}(n)| \\
&\leq f'(Q^*(n+1) + h - 1) \\
&= f'(f^{-1}\left(\frac{\epsilon}{8M} f(h + Q_{max}(n+1))\right) - 1)
\end{aligned}$$

where we have used the mean value theorem and the facts that f is a concave increasing function and at each index n , one queue can change at most by one. Therefore,

$$\frac{\pi_{n+1}(C)}{\pi_n(C)} \leq e^{2\frac{M\alpha}{\beta} f'(f^{-1}\left(\frac{\epsilon}{8M} f(h + Q_{max}(n+1))\right) - 1)}.$$

A similar calculation shows that also

$$\frac{\pi_n(C)}{\pi_{n+1}(C)} \leq e^{2\frac{M\alpha}{\beta} f'(f^{-1}\left(\frac{\epsilon}{8M} f(h + Q_{max}(n+1))\right) - 1)}.$$

This concludes the proof. \square

Next, we use the following version of Adiabatic Theorem from [23] to prove the time-scale decomposition property of our algorithm.

PROPOSITION 3. (Adapted from [23]) Suppose

$$\frac{\sigma_n}{1 - \theta_2(n+1)} \leq \delta'/4 \text{ for all } n \geq 0, \quad (25)$$

for some $\delta' > 0$, where $\theta_2(n+1)$ denotes the second largest eigenvalue of $P_\xi^{Q(n+1)}$. Then $\|\pi_n - \nu_n\|_{TV} \leq \delta'$, for all $n \geq n^*(\beta, \delta', \mathbf{S}(0))$, where n^* is the smallest n such that

$$\frac{1}{\sqrt{\pi_{min}(0)}} \exp\left(-\sum_{k=0}^n (1 - \theta_2(k))^2\right) \leq \delta'. \quad (26)$$

In our context, Proposition 3 states that under (25) and (26), after n^* steps, the distribution of the configurations over templates will be close to the desired steady-state distribution. To get some intuition, σ_n has the interpretation of the rate at which weights change, and $1/(1-\theta_2(n+1))$ has the interpretation of the time taken for the system to reach steady-state after the weights change. Thus, condition (25) ensures a *time-scale decomposition* – the weights change slowly compared to the time-scale that the system takes in order to respond and “settle down” with these changed weights.

It remains to show that that our system indeed satisfies the conditions of Proposition 3 as we do next, for the choice of $\delta' = \frac{\epsilon}{16}$. Suppose $f(x) = \log^{1-b}(x)$, for some $0 < b < 1$. Let $y = f(Q_{max}(n+1) + h)$. Obviously $f'(x) \leq 1/x$, so in view of equations (24), (21), (25), it suffices to have

$$\frac{2M\alpha}{\beta} \frac{1}{f^{-1}(\frac{\epsilon}{8M}y) - 1} \exp\left[\frac{4M}{\beta}(b_{max} + \alpha y)\right] \leq \frac{K_0^2 \epsilon}{128\xi^2}.$$

Note that $f^{-1}(x) = \exp(x^{\frac{1}{1-b}})$. Suppose $\alpha \leq \beta$. A simple calculation shows that it suffices to jointly have

$$\begin{aligned} y &\geq \frac{8M}{\epsilon} \log^{1-b} 3, \\ 4M\mu_{max}y - \frac{1}{2}\left(\frac{\epsilon}{8M}y\right)^{\frac{1}{1-b}} &\leq 0, \\ \frac{4M}{\beta}b_{max} - \frac{1}{2}\left(\frac{\epsilon}{8M}y\right)^{\frac{1}{1-b}} &\leq \log\left(\frac{K_0^2 \epsilon}{512\xi^2 M}\right). \end{aligned}$$

In summary, the condition (25) holds if

$$y \geq \left(\frac{1}{\epsilon}\right)^{2-b+\frac{1}{b}} \left(\frac{1}{\beta}\right)^{1-b} C_0 \quad (27)$$

or as a sufficient condition, if

$$h \geq \exp\left(C_0 \frac{1}{\beta} \left(\frac{1}{\epsilon}\right)^{\frac{2-b+\frac{1}{b}}{1-b}}\right) \quad (28)$$

for

$$C_0 \geq 8M\left(8Mb_{max} + 2\left|\log\frac{512\xi^2}{K_0^2}\right| + 2 + (8M)^{2/b}\mu_{max}^{(1-b)/b}\right).$$

Next, we find n^* that satisfies

$$\sum_{k=0}^{n^*-1} (1 - \theta_2(k))^2 \geq -\log\left(\frac{\epsilon}{16}\right) - \frac{1}{2}\log(\pi_{min}(0)).$$

From (23), and since $\alpha \leq \beta$,

$$-\log(\pi_{min}(0)) \leq \log K_1 + M\mu_{max}f(Q_{max}(0) + h) + \frac{M}{\beta}b_{max}.$$

Using Lemma 1, it can be shown that

$$\begin{aligned} &\sum_{k=0}^{n^*-1} (1 - \theta_2(k))^2 \\ &\geq \frac{2\xi^2}{K_0^2} e^{-4\frac{Mb_{max}}{\beta}} \sum_{k=0}^{n^*-1} e^{-4M\mu_{max}f(Q_{max}(k)+h)} \\ &\geq \frac{2\xi^2}{K_0^2} e^{-4\frac{Mb_{max}}{\beta}} \sum_{k=0}^{n^*-1} e^{-4M\mu_{max}f(Q_{max}(0)+h+n^*)} \\ &\geq \frac{2\xi^2}{K_0^2} e^{-4\frac{Mb_{max}}{\beta}} n^* (Q_{max}(0) + h + n^*)^{\frac{-4M\mu_{max}}{\log^b h}}. \end{aligned}$$

For $h \geq \exp((8M\mu_{max})^{1/b})$, it then suffices that

$$\begin{aligned} &\frac{2\xi^2}{K_0^2} e^{-4\frac{Mb_{max}}{\beta}} n^* (Q_{max}(0) + h + n^*)^{-1/2} \geq \\ &\log\left(\frac{16}{\epsilon}\right) + \frac{1}{2}\log K_1 + \frac{M\mu_{max}}{2}f(Q_{max}(0) + h) + \frac{M}{\beta}b_{max} \end{aligned}$$

which is clearly satisfied by choosing $n^* = Q_{max}(0) + h$ for h in (28) and C_0 a large enough constant.

Step 3: Lyapunov Analysis

The final step of the proof is based on a Lyapunov optimization method [20]. We develop the required Lyapunov arguments for $S(k) = (Q(k), C(k))$, i.e., the jump chain of the uniformized Markov chain. Consider the following Lyapunov function

$$V(k) = \sum_{j \in \mathcal{J}} \frac{1}{\mu_j} F(Q^{(j)}(k) + h),$$

where $F(x) = \int_h^x f(\tau) d\tau$. Recall that $f(x) = \log^{1-b} x$. Therefore F is convex, and following the standard one-step drift analysis

$$\begin{aligned} V(k+1) - V(k) &\leq \\ &\sum_{j \in \mathcal{J}} \frac{1}{\mu_j} f(Q^{(j)}(k+1) + h) \left(Q^{(j)}(k+1) - Q^{(j)}(k)\right) = \\ &\sum_{j \in \mathcal{J}} \frac{1}{\mu_j} f(Q^{(j)}(k) + h) \left(Q^{(j)}(k+1) - Q^{(j)}(k)\right) + \\ &\sum_{j \in \mathcal{J}} \frac{1}{\mu_j} \left(f(Q^{(j)}(k+1) + h) - f(Q^{(j)}(k) + h)\right) \left(Q^{(j)}(k+1) - Q^{(j)}(k)\right). \end{aligned}$$

By the mean value theorem, and using the fact that f is a concave increasing function, it follows that

$$\begin{aligned} &|f(Q^{(j)}(k+1) + h) - f(Q^{(j)}(k) + h)| \leq \\ &f'(h) |Q^{(j)}(k+1) - Q^{(j)}(k)| \end{aligned}$$

Recall that $C(k) = (C_a^{(j)}(k), C_v^{(j)}(k))$ where $C_v^{(j)}$ is the set of virtual templates (i.e., the templates that do not contain jobs of type j) and $C_a^{(j)}$ is the set of actual templates.

For notational compactness, let $\mathbb{E}_{\mathbf{S}(k)}[\cdot] = \mathbb{E}[\cdot | \mathbf{S}(k)]$, where $\mathbf{S}(k)$ is the state of the system at each index k . Then

$$\mathbb{E}_{\mathbf{S}(k)} \left[V(k+1) - V(k) \right] \leq f'(h) \sum_j \frac{1}{\mu_j} \left(\frac{\lambda_j}{\xi} + \frac{|C^{(j)}(k)\mu_j}{\xi} \right) + \sum_j \frac{1}{\mu_j} f(Q^{(j)}(k) + h) \left[\frac{\lambda_j}{\xi} - \left(|C^{(j)}(k)| - |C_v^{(j)}(k)| \right) \frac{\mu_j}{\xi} \right],$$

where we have used the fact that at most one arrival or departure can happen at every jump index, i.e., $|Q^{(j)}(k+1) - Q^{(j)}(k)| \in \{0, 1\}$.

Note that clearly the maximum number of templates of any type of jobs that can fit in a configuration is less than M (recall that $M = \sum_{\ell \in \mathcal{L}} |m_\ell|$). Moreover, none of the templates of type j will be virtual if more than M jobs of type j are available in the system, hence,

$$|C_v^{(j)}(k)| f(h + Q^{(j)}(k)) \leq |C_v^{(j)}(k)| f(h + M).$$

and therefore,

$$\begin{aligned} \mathbb{E}_{\mathbf{S}(k)} \left[V(k+1) - V(k) \right] &\leq K_2 + K_3 \\ &+ \frac{1}{\xi} \sum_j f(Q^{(j)}(k) + h) \left(\rho_j - |C^{(j)}(k)| \right) \end{aligned}$$

where $K_2 \leq f'(h)(M + \sum_j \rho_j)/\xi$, and $K_3 \leq Mf(h+M)/\xi$. Therefore, it follows that

$$\begin{aligned} \alpha \mathbb{E}_{\mathbf{S}(k)} \left[V(k+1) - V(k) \right] &+ \frac{1}{\xi} \mathbb{E}_{\mathbf{S}(k)} \left[\sum_{j \in \mathcal{J}} \sum_{A \in C^{(j)}(k)} b_A^{(j)} \right] \leq \\ \alpha(K_2 + K_3) &+ \frac{\alpha}{\xi} \sum_{j \in \mathcal{J}} \rho_j f(Q^{(j)}(k) + h) \\ - \frac{1}{\xi} \sum_{j \in \mathcal{J}} \sum_{A \in C^{(j)}(k)} &\left[\alpha f(Q^{(j)}(k) + h) - b_A^{(j)} \right] \end{aligned}$$

Taking the expectation of both sides with respect to ν_n (distribution of configurations given the queues at $n > n^*$), we get

$$\begin{aligned} \alpha \mathbb{E}_{\mathbf{Q}(k)} \left[V(k+1) - V(k) \right] &+ \frac{1}{\xi} \mathbb{E}_{\mathbf{Q}(k)} \left[\sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{A}^{(j)}} x_A(k) b_A^{(j)} \right] \\ &\leq \alpha(K_2 + K_3) + \frac{\alpha}{\xi} \sum_{j \in \mathcal{J}} \rho_j f(Q^{(j)}(k) + h) \\ - \frac{1}{\xi} \mathbb{E}_{\mathbf{Q}(k)} \left[\sum_{j \in \mathcal{J}} \sum_{A \in C^{(j)}(k)} &\left(\alpha f(Q^{(j)}(k) + h) - b_A^{(j)} \right) \right] \\ &\leq \alpha(K_2 + K_3) - \frac{\beta}{\xi} \log \gamma_{min} + \frac{\alpha}{\xi} \sum_{j \in \mathcal{J}} \rho_j f(Q^{(j)}(k) + h) \\ - \frac{1}{\xi} \left(1 - \frac{\epsilon}{4} \right) W^*(k) &+ \frac{\epsilon}{\xi} b_{max}, \end{aligned} \quad (29)$$

where the last inequality is based on Corollary 3, where

$$W^*(k) = \max_{C \in \mathcal{C}} \sum_{j \in \mathcal{J}} \sum_{A \in C^{(j)}} \left(\alpha f(Q^{(j)}(k) + h) - b_A^{(j)} \right)$$

Notice that equivalently

$$\begin{aligned} W^*(k) &= \max_{\{x_A\}} \sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{A}^{(j)}} x_A \left(\alpha f(Q^{(j)}(k) + h) - b_A^{(j)} \right) \\ \text{subject to } &\left(\sum_{A \in \mathcal{A}^{(j)}} x_A; j \in \mathcal{J} \right) \in \Lambda \\ &x_A \geq 0; \forall A \in \cup_j \mathcal{A}^{(j)} \end{aligned}$$

Let x^* be the optimal solution to the static partitioning problem. By the feasibility of x^* , $\rho_j \leq \sum_{A \in \mathcal{A}^{(j)}} x_A^*$, for all $j \in \mathcal{J}$, hence

$$\sum_{j \in \mathcal{J}} \rho_j f(Q^{(j)}(k) + h) \leq \sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{A}^{(j)}} x_A^* f(Q^{(j)}(k) + h). \quad (30)$$

Further, by assumption, ρ is strictly inside Λ , thus there exists a δ^* such that $\rho(1 + \delta^*) \in \Lambda$. It is easy to show by the monotonicity of \mathcal{C} (i.e., if $C \in \mathcal{C}$, $C \setminus A \in \mathcal{C}$, for all $A \in C^{(j)}$, $j \in \mathcal{J}$) that, at the optimal solution, the constraint (7) should in fact hold with equality. Hence

$$\left(\sum_{A \in \mathcal{A}^{(j)}} x_A^* (1 + \delta^*) : j \in \mathcal{J} \right) \in \Lambda.$$

Therefore,

$$W^*(k) \geq \sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{A}^{(j)}} \left((1 + \delta) x_A^* \right) \left(\alpha f(Q^{(j)}(k) + h) - b_A^{(j)} \right) \quad (31)$$

For $\epsilon \leq \delta^*$, $(1 - \frac{\epsilon}{4})(1 + \delta^*) \geq 1 + \frac{\delta}{2}$, for any $\delta \in [0, \delta^*]$. Then using (30) and (31) in (29),

$$\begin{aligned} \alpha \mathbb{E}_{\mathbf{Q}(k)} \left[V(k+1) - V(k) \right] &+ \frac{1}{\xi} \mathbb{E}_{\mathbf{Q}(k)} \left[\sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{A}^{(j)}} x_A(k) b_A^{(j)} \right] \\ &\leq \alpha(K_2 + K_3) + \frac{\alpha}{\xi} \sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{A}^{(j)}} x_A^* f(Q^{(j)}(k) + h) \\ &- \frac{1}{\xi} \left(1 + \frac{\delta}{2} \right) \sum_{j \in \mathcal{J}} \sum_{A \in \mathcal{A}^{(j)}} x_A^* \left(-b_A^{(j)} + \alpha f(Q^{(j)}(k) + h) \right) \\ &- \frac{\beta}{\xi} \log \gamma_{min} + \frac{\epsilon}{\xi} b_{max} \\ &= \frac{1}{\xi} \left(1 + \frac{\delta}{2} \right) G(x^*) - \frac{\alpha \delta}{\xi 2} \sum_{j \in \mathcal{J}} \sum_{A \in C^{(j)}(k)} x_A^* f(Q^{(j)}(k) + h) \\ &- \frac{\beta}{\xi} \log \gamma_{min} + \frac{\epsilon}{\xi} b_{max} + \alpha(K_2 + K_3). \end{aligned} \quad (32)$$

It follows from this that the Markov chain $(\mathbf{Q}(k), \mathbf{C}(k))$ is positive recurrent as a consequence of the Foster-Lyapunov theorem, with Lyapunov function $V(\cdot)$. Taking the expectation of both sides of (32) with respect to $\mathbf{Q}(k)$, and then taking summation over $k = 0, \dots, N-1$, and dividing by N , and letting $N \rightarrow \infty$ yields

$$\begin{aligned} \limsup_N \frac{1}{N} \sum_{k=0}^{N-1} \sum_{j \in \mathcal{J}} \mathbb{E} \left[f(Q^{(j)}(k) + h) \right] \\ \leq \frac{\alpha \xi (K_2 + K_3) - \beta \log \gamma_{min} + \epsilon b_{max} + (1 + \delta/2) G(x^*)}{\alpha \rho_{min} \delta / 2} \end{aligned}$$

$$\begin{aligned} \limsup_N \frac{1}{N} \sum_{k=0}^{N-1} \mathbb{E} \left[G(x(k)) \right] \\ \leq (1 + \delta/2) G(x^*) + \alpha \xi (K_2 + K_3) - \beta \log \gamma_{min} + \epsilon b_{max} \end{aligned}$$

where we have used the fact that contribution of queue sizes and costs in $(0, n^*)$ to the average quantities vanishes to zero as $N \rightarrow \infty$. The above inequalities can be independently optimized over $\delta \in [0, \delta^*]$ (the performance of the algorithm is independent of δ). Here we choose $\delta = \delta^*$ in the queue inequality and $\delta = 0$ in the cost inequality. The statement of Theorem then follows using the Ergodic theorem and the fact that the jump chain and the original chain have the same steady-state average behaviour.

7. CONCLUSIONS

Motivated by modern stream data processing applications, we have investigated the problem of dynamically scheduling graphs in cloud clusters, where a graph represents a specific computing job. These graphs arrive and depart over time. Upon arrival, each graph can either be queued or served immediately. The objective is to develop algorithms that assign nodes of these graphs to free (computing) slots in the machines of the cloud cluster. The performance metric for the scheduler (partition graphs and map it to slots) is to minimize the average graph partitioning cost, while keeping the system stable.

We have proposed a novel class of low complexity algorithms which can approach the optimal solution by exploiting the trade-off between delay and partitioning cost, without causing service interruptions. The key ingredient of the algorithms is the generation/removal of random templates from the cluster at appropriate instances of time, where each template is a unique way of partitioning a graph.

8. REFERENCES

- [1] K. Andreev and H. Racke. Balanced graph partitioning. *Theory of Computing Systems*, 39(6):929–939, 2006.
- [2] L. Aniello, R. Baldoni, and L. Querzoni. Adaptive online scheduling in Storm. In *7th ACM international conference on Distributed event-based systems*, pages 207–218, 2013.
- [3] P. Bremaud. *Markov chains: Gibbs fields, Monte Carlo simulation, and queues*, volume 31. Springer, 1999.
- [4] A. Buluç, H. Meyerhenke, I. Safro, P. Sanders, and C. Schulz. Recent advances in graph partitioning. *CoRR*, arXiv:1311.3144, 2013.
- [5] E. G. Coffman Jr, M. R. Garey, and D. S. Johnson. Approximation algorithms for bin packing: A survey. In *Approximation algorithms for NP-hard problems*, pages 46–93. PWS Publishing Co., 1996.
- [6] T. Condie, N. Conway, P. Alvaro, J. M. Hellerstein, K. Elmeleegy, and R. Sears. MapReduce online. In *NSDI*, volume 10, page 20, 2010.
- [7] J. Csirik and G. J. Woeginger. *On-line packing and covering problems*. Springer, 1998.
- [8] A. E. Feldmann and L. Foschini. Balanced partitions of trees and applications. *Algorithmica*, pages 1–23, 2012.
- [9] J. Ghaderi, T. Ji, and R. Srikant. Flow-level stability of wireless networks: Separation of congestion control and scheduling. *IEEE Transactions on Automatic Control*, 59(8):2052 – 2067, 2014.
- [10] J. Ghaderi and R. Srikant. On the design of efficient CSMA algorithms for wireless networks. In *49th IEEE Conference on Decision and Control (CDC)*, pages 954–959, 2010.
- [11] J. Ghaderi, Y. Zhong, and R. Srikant. Asymptotic optimality of BestFit for stochastic bin packing. *SIGMETRICS Perform. Eval. Rev.*, 42(2):64–66, Sept. 2014.
- [12] Giraph. <https://giraph.apache.org/>.
- [13] Graphlab. <http://graphlab.com/>.
- [14] B. Hendrickson and R. W. Leland. A multi-level algorithm for partitioning graphs. *SC*, 95:28, 1995.
- [15] IBM InfoSphere Platform, <http://www-01.ibm.com/software/data/infosphere>.
- [16] J. W. Jiang, T. Lan, S. Ha, M. Chen, and M. Chiang. Joint VM placement and routing for data center traffic engineering. In *Proceedings of IEEE INFOCOM*, pages 2876–2880, 2012.
- [17] L. Jiang and J. Walrand. A distributed CSMA algorithm for throughput and utility maximization in wireless networks. *IEEE/ACM Transactions on Networking (TON)*, 18(3):960–972, 2010.
- [18] S. A. Lippman. Applying a new device in the optimization of exponential queuing systems. *Operations Research*, 23(4):687–710, 1975.
- [19] J. Liu, Y. Yi, A. Proutiere, M. Chiang, and H. V. Poor. Towards utility-optimal random access without message passing. *Wireless Communications and Mobile Computing*, 10(1):115–128, 2010.
- [20] M. J. Neely. *Stochastic network optimization with application to communication and queueing systems*, volume 3. Morgan & Claypool Publishers, 2010.
- [21] M. L. Puterman. *Markov decision processes: discrete stochastic dynamic programming*, volume 414. John Wiley & Sons, 2009.
- [22] Z. Qian, Y. He, C. Su, Z. Wu, H. Zhu, T. Zhang, L. Zhou, Y. Yu, and Z. Zhang. TimeStream: Reliable stream computation in the cloud. In *EuroSys 2013*, pages 1–14, 2013.
- [23] S. Rajagopalan, D. Shah, and J. Shin. Network adiabatic theorem: An efficient randomized protocol for contention resolution. In *ACM SIGMETRICS Performance Evaluation Review*, volume 37, pages 133–144. ACM, 2009.
- [24] M. Rychly, P. Koda, and P. Smrz. Scheduling decisions in stream processing on heterogeneous clusters. In *Eighth International Conference on Complex, Intelligent and Software Intensive Systems (CISIS)*, pages 614 – 619, July 2014.
- [25] S4 distributed stream computing platform, <http://incubator.apache.org/s4>.
- [26] D. Shah and J. Shin. Randomized scheduling algorithm for queueing networks. *The Annals of Applied Probability*, 22(1):128–171, 2012.
- [27] A. Stolyar and Y. Zhong. Asymptotic optimality of a greedy randomized algorithm in a large-scale service system with general packing constraints. *arXiv preprint arXiv:1306.4991*, 2013.
- [28] A. L. Stolyar. An infinite server system with general packing constraints. *Operations Research*, 61(5):1200–1217, 2013.
- [29] Storm: Distributed and fault-tolerant realtime computation, <http://storm.incubator.apache.org>.

- [30] C. Walshaw and M. Cross. Mesh partitioning: a multilevel balancing and refinement algorithm. *SIAM Journal on Scientific Computing*, 22(1):63–80, 2000.
- [31] Graph partition problem on Wikipedia, http://en.wikipedia.org/wiki/Graph_partition.
- [32] M. Zaharia, T. Das, H. Li, S. Shenker, and I. Stoica. Discretized streams: an efficient and fault-tolerant model for stream processing on large clusters. In *Proceedings of the 4th USENIX conference on Hot Topics in Cloud Computing*, pages 10–10. USENIX Association, 2012.

APPENDIX

A. PROOF OF THEOREM 1

The proof is standard and based on a Lyapunov optimization method [20]. Consider a Lyapunov function $V(t) = \sum_{j \in \mathcal{J}} \frac{1}{\mu_j} F(Q^{(j)}(t))$, where $F(x) = \int_0^x f(\tau) d\tau$. Recall that $f: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is a concave increasing function; thus F is convex. Choose an arbitrarily small $u > 0$. It follows from convexity of F that for any $t \geq 0$,

$$\begin{aligned} V(t+u) - V(t) &\leq \sum_{j \in \mathcal{J}} \frac{1}{\mu_j} f(Q^{(j)}(t+u))(Q^{(j)}(t+u) - Q^{(j)}(t)) \\ &= \sum_{j \in \mathcal{J}} \frac{1}{\mu_j} f(Q^{(j)}(t))(Q^{(j)}(t+u) - Q^{(j)}(t)) \\ &\quad + \sum_{j \in \mathcal{J}} \frac{1}{\mu_j} (f(Q^{(j)}(t+u)) - f(Q^{(j)}(t)))(Q^{(j)}(t+u) - Q^{(j)}(t)). \end{aligned}$$

By definition

$$Q^{(j)}(t+u) - Q^{(j)}(t) = H^{(j)}(t, t+u) - D^{(j)}(t, t+u),$$

where for any $0 \leq t_1 \leq t_2$,

$$\begin{aligned} H^{(j)}(t_1, t_2) &= \mathcal{N}_j^H(\lambda_j(t_2 - t_1)) \\ D^{(j)}(t_1, t_2) &= \mathcal{N}_j^D\left(\int_{t_1}^{t_2} C_a(\tau) \mu_j d\tau\right), \end{aligned}$$

where $\mathcal{N}_j^A(z)$ and $\mathcal{N}_j^D(z)$ denote independent Poisson random variables with rate z , for all $j \in \mathcal{J}$. Recall that $C(t) = (C_a^{(j)}(t), C_v^{(j)}(t))$ where $C_v^{(j)}$ is the set of virtual templates (i.e., the templates that do not contain jobs of type j) and $C_a^{(j)}$ is the set of actual templates. It is easy to see that

$$|f(Q^{(j)}(t+u)) - f(Q^{(j)}(t))| \leq f'(0)|Q^{(j)}(t+u) - Q^{(j)}(t)|,$$

by the mean value theorem, and the fact that f is a concave increasing function. For notational compactness, let $\mathbb{E}_{\mathbf{S}(t)}[\cdot] = \mathbb{E}[\cdot | \mathbf{S}(t)]$, where $\mathbf{S}(t)$ is the state of the system at each time t . Clearly the the maximum number of templates that can fit in a configuration is less than M . It is easy to see that

$$\begin{aligned} \mathbb{E}_{\mathbf{S}(kT)} [V(t+u) - V(t)] &\leq \\ &\sum_j \frac{1}{\mu_j} \mathbb{E}_{\mathbf{S}(kT)} \left[f(Q^{(j)}(t))(A^{(j)}(t, t+u) - D^{(j)}(t, t+u)) \right] \\ &\quad + K_2 u + o(u), \end{aligned}$$

where $K_2 = f'(0) \sum_{j \in \mathcal{J}} (\rho_j + M)$. Clearly virtual templates do not exist in the configuration if there are more than M

jobs in the system. Hence it follows that

$$\begin{aligned} &\sum_{j \in \mathcal{J}} \frac{1}{\mu_j} \mathbb{E}_{\mathbf{S}(kT)} \left[D^{(j)}(t, t+u) f(Q^{(j)}(t)) \right] \geq \\ &\sum_{j \in \mathcal{J}} \mathbb{E}_{\mathbf{S}(kT)} \left[|C^{(j)}(t)| f(Q^{(j)}(t)) \right] u - K_3 u - o(u) \end{aligned}$$

for $K_3 = Mf(M)$. Note that the algorithm keeps the configuration fixed over intervals $[kT, (k+1)T)$, i.e., $C(t) = C(kT)$ for $t \in [kT, (k+1)T)$. Let $\Delta_{t,u} := \frac{1}{u} \mathbb{E}_{\mathbf{S}(kT)} [V(t+u) - V(t)]$, then,

$$\begin{aligned} \Delta_{t,u} &\leq \mathbb{E}_{\mathbf{S}(kT)} \left[\sum_{j \in \mathcal{J}} f(Q^{(j)}(t)) (\rho_j - |C^{(j)}(kT)|) \right] \\ &\quad + K_2 + K_3 + o(1) \\ &\leq \sum_{j \in \mathcal{J}} f(Q^{(j)}(kT)) (\rho_j - |C^{(j)}(kT)|) \\ &\quad + K_2^2 T + K_2 + K_3 + o(1) \end{aligned}$$

Taking the limit $u \rightarrow 0$,

$$\frac{d\mathbb{E}_{\mathbf{S}(kT)} [V(t)]}{dt} \leq \sum_{j \in \mathcal{J}} f(Q^{(j)}(kT)) (\rho_j - |C^{(j)}(kT)|) + K_4$$

where $K_4 = K_2^2 T + K_2 + K_3$. Let $\Delta_k = \mathbb{E}_{\mathbf{S}(kT)} [V((k+1)T) - V(kT)]$, then over the k -th cycle

$$\frac{\alpha}{T} \Delta_k + \mathbb{E}_{\mathbf{S}(kT)} \left[\sum_{j \in \mathcal{J}} \sum_{A \in C^{(j)}(kT)} b_A^{(j)} \right] \leq \text{Term}_1 - \text{Term}_2 + \alpha K_4,$$

where

$$\begin{aligned} \text{Term}_1 &= \alpha \sum_{j \in \mathcal{J}} \rho_j f(Q^{(j)}(kT)), \\ \text{Term}_2 &= \sum_{j \in \mathcal{J}} \sum_{A \in C^{(j)}(kT)} (\alpha f(Q^{(j)}(kT)) - b_A^{(j)}). \end{aligned}$$

Let x^* be the optimal solution to the static partitioning problem. The rest of the proof is similar to the Lyapunov analysis in the proof of Theorem 2 (step 3), i.e., for any $0 \leq \delta \leq \delta^*$,

$$\begin{aligned} \text{Term}_1 &\leq \alpha \sum_{j \in \mathcal{J}} \sum_{A \in A^{(j)}} x_A^* f(Q^{(j)}(kT)) \\ \text{Term}_2 &\leq \sum_{j \in \mathcal{J}} \sum_{A \in A^{(j)}} \left((1+\delta)x_A^* \right) (\alpha f(Q^{(j)}(kT)) - b_A^{(j)}) \end{aligned}$$

Putting everything together,

$$\begin{aligned} \frac{\alpha}{T} \Delta_k + \mathbb{E}_{\mathbf{S}(kT)} \left[\sum_{j \in \mathcal{J}} \sum_{A \in C^{(j)}} b_A^{(j)} \right] &\leq (1+\delta)G(x^*) \\ &\quad - \alpha \delta \rho_{\min} \sum_{j \in \mathcal{J}} f(Q^{(j)}(kT)) + \alpha K_4. \end{aligned}$$

Then it follows from the Foster-Lyapunov theorem that the Markov chain $veS(kT)$, $k = 0, 1, 2, \dots$ (and therefore Markov chain $\mathbf{S}(t)$, $t \geq 0$) is positive recurrent. As in the step 3 in the proof of of Theorem 2, we take the expectation from both sides of the above equality with respect to $\mathbf{S}(kT)$, and then sum over $k = 0, \dots, N-1$, divide by N , and let $N \rightarrow \infty$. Then the statement of the theorem follows by choosing $B_1 = K_2 + K_3$ and $B_2 = K_2^2$.

B. AN ALTERNATIVE DESCRIPTION OF DYNAMIC GRAPH PARTITIONING ALGORITHM

The Dynamic Graph Partitioning (DGP) algorithm, as described in Section 5, does not require any dedicated clock as the decisions are made at the instances of job arrival and departure. In this section, we present an alternative description of the algorithm by using dedicated clocks. Each queue $Q^{(j)}$ is assigned an independent Poisson clock of rate $\hat{\lambda}e^{\alpha f^{(j)}(\mathbf{h}+\mathbf{Q}(t))/\beta}$, where $\hat{\lambda}$ is a fixed constant depending on how fast the iterations in the algorithm can be performed. Equivalently, at each time t , the time duration until the tick of the next clock is an exponential random variable with parameter $\hat{\lambda}e^{\alpha f^{(j)}(\mathbf{h}+\mathbf{Q}(t))/\beta}$. This means if $Q^{(j)}$ changes at time $t' > t$ before the clock makes a tick, the time duration until the next tick is reset to an independent exponential random variable with parameter $\hat{\lambda}e^{\alpha f^{(j)}(\mathbf{h}+\mathbf{Q}(t'))/\beta}$. The description of the algorithm is given below.

Algorithm 4 Alternative Dynamic Graph Partitioning (ADGP) Algorithm

At the instances of dedicated clocks.

Suppose the dedicated clock of queue $Q^{(j)}$ makes a tick, then:

- 1: A virtual template $A^{(j)}$ is chosen randomly from currently feasible templates for graph \mathcal{G}_j , given the current configuration, using *Random Partition Procedure*, if possible. Then this template is added to the configuration with probability $e^{-\frac{1}{\beta}b_A^{(j)}}$ and discarded otherwise. The virtual template leaves the system after an exponentially distributed time duration with mean $1/\mu_j$.
- 2: If there is a job of type j in $Q^{(j)}$ waiting to get service, and a virtual template of type j is created in step 1, this virtual template is filled by a job from $Q^{(j)}$ which converts the virtual template to an actual template.

At arrival instances.

- 1: Suppose a graph (job) of type \mathcal{G}_j arrives. The job is added to queue $Q^{(j)}$.

At departure instances.

- 1: At the departure instances of actual/virtual templates, the algorithm removes the corresponding template from the configuration.
 - 2: If this is a departure of an actual template, the job is departed and the corresponding queue is updated.
-

The algorithm will yield average queue size and partitioning cost performance similar to those in Theorem 2. The proof essentially follows the three steps of the proof of Theorem 2. Here, we only describe the main property of the Alternative Dynamic Graph Partitioning algorithm with fixed weights, which we refer to as $\overline{\text{ADGP}}(\tilde{W})$ (the counterpart $\overline{\text{DGP}}(\tilde{W})$ in Section 6).

PROPOSITION 4. *Under $\overline{\text{ADGP}}(\tilde{W})$, the steady state distribution of configurations solves*

$$\max_{\pi} \mathbb{E}_{\pi} \left[\sum_{j \in \mathcal{J}} \sum_{A \in C^{(j)}} \tilde{w}_A^{(j)} \right] - \beta D_{KL}(\pi \parallel \hat{\gamma})$$

for the following distribution $\hat{\gamma}$

$$\hat{\gamma}_C = \frac{1}{\tilde{Z}_{\hat{\gamma}}} \left(\sum_{\ell} |m_{\ell}| - \sum_j |C^{(j)} \cap V_j| \right)! \prod_j \left(\frac{\hat{\lambda}}{\mu_j} \right)^{|C^{(j)}|} \quad (33)$$

where $\tilde{Z}_{\hat{\gamma}}$ is the normalizing constant.

Similarly to the $\overline{\text{DGP}}(\tilde{W})$ algorithm, as $\beta \rightarrow \infty$, the optimizing π converges to $\hat{\gamma}$. The distribution $\hat{\gamma}$ has the interpretation of the steady state distribution of configurations in a loss system with arrival rates $\hat{\lambda}_j = \hat{\lambda}$, $j \in \mathcal{J}$, and service rates $\hat{\mu}_j = \mu_j$, $j \in \mathcal{J}$. In the loss system, when a graph arrives, it is randomly distributed over the machines if possible; otherwise it is dropped. At the departure instances, the job and hence its template leave the system.

PROOF OF PROPOSITION 4. The proof is basically identical to the proof of Proposition 1. The only difference is that the detailed balance equations are given by

$$\pi(C \oplus A^{(j)})\mu_j = \pi(C) \frac{\hat{\lambda}e^{\alpha f(Q^{(j)})/\beta}}{|A^{(j)}(C)|} e^{-\frac{1}{\beta}b_A^{(j)}}$$

for any configuration C and $C \oplus A^{(j)} \in \mathcal{C}$; $A^{(j)} \in \mathcal{A}^{(j)}$, $j \in \mathcal{J}$. Here the LHS is the departure rate of (virtual or actual) template $A^{(j)}$ from the configuration $C \oplus A^{(j)}$. The RHS is the rate at which the (actual or virtual) template $A^{(j)}$ for graphs \mathcal{G}_j is added to configuration C (the *Random partition Procedure* selects a template $A \in \mathcal{A}^{(j)}(C)$ uniformly at random). Thus the detailed balanced equations are given by

$$\pi(C \oplus A^{(j)}) = \pi(C) \frac{\hat{\lambda}/\mu_j}{|A^{(j)}(C)|} e^{\frac{1}{\beta}w_A^{(j)}}. \quad (34)$$

and it is easy to see that (15) with γ replaced with $\hat{\gamma}$ in (33), indeed satisfies the detailed balance equations, with the normalizing condition that $\sum_C \pi(C) = 1$. The fact that that this distribution maximizes the stated objective function follows in parallel with the arguments in the proof of Proposition 1. \square