Leveraging knowledge-as-a-service (KaaS) for QoS-aware resource management in multi-user video transcoding

Luis Costero · Francisco D. Igual · Katzalin
Olcoz · Francisco Tirado

Abstract The coexistence of parallel applications in shared computing nodes, each one featuring different Quality of Service (QoS) requirements, carries out new challenges to improve resource occupation while keeping acceptable rates in terms of QoS. As more application-specific and system-wide metrics are included as QoS dimensions, or under situations in which resource-usage limits are strict, building and serving the most appropriate set of actions (application control knobs and system resource assignment) to concurrent applications in an automatic and optimal fashion becomes mandatory. In this paper, we propose strategies to build and serve this type of knowledge to concurrent applications by leveraging Reinforcement Learning techniques. Taking multi-user video transcoding as a driving example, our experimental results reveal an excellent adaptation of resource and knob management to heterogeneous QoS requests, and increases in the amount of concurrently served users up to 1.24× compared with alternative approaches considering homogeneous QoS requests.

Keywords Resource management · Heterogeneous Quality of Service · Reinforcement Learning · Multi-core architectures · HEVC video transcoding

1 Introduction and motivation

The integration of intelligent policies for resource management and application tuning in shared computing systems is becoming a field of paramount interest to efficiently exploit the potential of the underlying architectures without

Departamento de Arquitectura de Computadores y Automática
Universidad Computense de Madrid
E-mail: {lcostero,figual,katzalin,ptirado}@ucm.es

human intervention. In situations where external limitations in terms of Quality of Service (QoS), tight per-application SLA (Service Level Agreements) or energy consumption are imposed, the development and application of such policies becomes a hurdle difficult to be automatically addressed [14].

Two of the main properties of any generic autonomous system, including resource managers in the fields of Cloud Computing or HPC, are *self-configuration* (ability to adapt to environmental changes) and *self-optimization* (capability to improve performance and reduce overloading or underloading the underlying resources) [10]. Resource managers can be actually considered in terms of the IBM Autonomic Model [7], which encompasses four main generic steps. This sequence of steps is repeated in a control loop, that usually features sensoring and acting capabilities towards the underlying architecture or application. These actions are commonly cast in terms of selecting values for architectural knobs (e.g. core frequency) or application-specific knobs. In addition, these cyclic steps orbit around the existence of a shared Knowledge Base (KB) storing rules that, properly orchestrated, can fulfill the requirements and restrictions imposed without further human intervention.

The development of the KB, however, can become a daunting task when the amount of architectural and application-level knobs increase and their interplay is nontrivial. In stochastic environments such as shared nodes in cloud deployments, in which the application of a given rule does not always yield the same result in terms of performance and/or energy consumption, the creation, maintenance and effective application of the knowledge of the KB is even a more complex task. The challenge is harder in scenarios in which the request arrival rate and its distribution are unknown, or when the throughput or quality attained are content-dependent, and hence unpredictable.

Reinforcement Learning (RL) [12] is a field of Artificial Intelligence that has shown to be effective in problems that feature complex and large state spaces, dynamic environments and without any pre-established knowledge. In this paper, we leverage RL to build, maintain and enrich the Knowledge Base of a centralized resource manager in shared servers in order to attain automatic and efficient resource management and assignation for multiple concurrent applications exhibiting *heterogeneous* QoS demands.

This idea, that can be considered as *Knowledge-as-a-service* (KaaS), is illustrated in this paper by means of a practical yet illustrative example: QoS-aware multi-user video transcoding. In this scenario, multiple users demand the execution of multiple concurrent transcoding processes, each one with different QoS requirements. We consider two types of users with different QoS requirements (regular and premium users), and optional resource-minimization strategies.

Resource allocation and management in shared computing environments has been thoroughly studied in the past from different perspectives. [9] studied the differences between static and dynamic resource allocation for media processing. Dynamic resource provisioning has also been tackled in terms of multiple priorities for heterogeneous QoS requirements in [4, 5]. Based on simulation, [3] studies policies for allocating and deallocating virtual machines to

attend multiple QoS requirements by concurrent applications. Typically, delay time is the metric considered in these works as the only QoS target, ignoring other application-specific quality metrics and architecture-specific considerations (e.g. energy consumption). Many of these works propose predictions or model-based projections for resource usage and QoS violations to apply heuristics and policies for resource management. Reinforcement Learning has been shown to be a useful technique to develop self-adaptive policies for multi-QoS scenarios, with model-based approaches [6] or model-free approaches [8].

Our work introduces the following novelties compared with previous works:

- We integrate RL techniques into a centralized resource manager, and demonstrate its effectiveness to attend multi-user environments with *heterogeneous QoS requirements* within a single compute node.
- We formalize a methodology to build different policies based on RL, that ultimately make it possible to learn different policies serving different QoS requirements in a feasible amount of time.
- We propose strategies to boost learning time based on the construction of the transitions table of the system, used to train the RL system.
- We demonstrate the effectiveness of our proposal on a real-time HEVC encoding application. We deal with extended sets of both application- and system-level knobs. For multi-user video transcoding, we consider quality, throughput and energy consumption as target metrics to be simultaneously improved and/or limited.
- We compare the learning time benefits of our approach against other state-of-the-art RL approaches.
- We propose heuristics to serve different knowledge to individual users in real-time on a shared computing node.
- Reported results are based on an actual multi-core architecture, demonstrating the ability of RL to deal with stochastic (noisy) environments.

The rest of the paper is structured as follows. Section 2 proposes efficient mechanisms to build ad-hoc policies leveraging Reinforcement Learning. Section 3 motivates the use of multi-user video transcoding as an illustrative example of such techniques and reports the attained results for different policies designed to adapt multi-user video transcoding to different levels of QoS and resource usage. Section 4 introduces heuristics to leverage the existence of different policies in the KB, and to apply them in scenarios with heterogeneous QoS requirements. In Section 5, we report the observed results and benefits of such heuristics. Section 6 closes the paper with some final remarks.

2 Policy design

A *Markov Decision Process* (MDP) is a framework for modelling decision making problems in stochastic environments. A MDP=$(\mathcal{S}, \mathcal{A}, \mathcal{P}, \mathcal{R})$ is defined by a finite set of states the system can be at each moment of the execution, $\mathcal{S}$, a finite set of actions, $\mathcal{A}$, that can be applied to the system and can

produce a change in the state the system is, a set of probability functions $\mathcal{P} = \{P_a(s, s') : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \to [0, 1]\}$ which determine the probability of moving from one state (s) to another (s') after taking an action (a), and a set of reward functions $\mathcal{R} = \{R_a(s, s') : \mathcal{S} \times \mathcal{S} \times \mathcal{A} \to \mathbb{R}\}$ which define how good or bad was the transition $s \to s'$ due to action a. The final goal of a MDP is to find a policy $\pi(s)$ which maximizes the expected accumulated reward when transitioning through the system. Note that this definition entails that if the system is provided with different reward functions, the obtained policies will be different. However, in most real problems, determining the probabilities or rewards that define a specific MDP is not an easy task, being in most of the cases unknown, or estimated from noisy observations.

Reinforcement Learning (RL) is an area of Machine Learning which can tackle the problem of finding an optimal policy $\pi(s)$ for a MDP where the probabilities or rewards functions are unknown. Specifically, Q-Learning has been proved as a valid RL algorithm able to find the optimal policy following dynamic programming techniques. Q-Learning is a model-free algorithm which can handle problems with unknown stochastic transitions based on an infinite-time exploration of the transitions between states when different actions are applied based on a partially-random policy of choosing actions. However, due to the infinite-nature of the formulation of the algorithm, the optimality of the obtained policy $\pi'(s)$ will be ultimately based on the time the algorithm has been exploring the system and, in essence, the distribution of the number of times each pair *state/action* has been explored. In real-time problems, where the definition of the state comes from real measurements of the environment, this exploration time depends on the frequency the system provides the different metrics, producing long time training sessions. Even more, if the final goal is to train the system several times to obtain different behaviours (changing the reward functions each time), the total training time can extend to unacceptable times. In the following, we address both questions, namely: *(i)* how to define the system to obtain different learned behaviours, and *(ii)* how to boost the learning time providing the transition probabilities to the Q-Learning system.

2.1 Learning different policies

Traditionally, training a Q-Learning system is a trial and error process in which the designer of the experiment modifies and tunes the state and reward function definitions until the obtained policy meets the desired behaviour. Here, we propose efficient mechanisms to define the state and reward functions to obtain different policies with different behaviours with minimum effort.

State definition: Instead of having a unique and complete definition of a state, we propose to divide the states space into different independent subspaces (i.e. the metrics required to build one subspace are not used to create any other subspace). This state definition allows us to have a more fine-grained control on the behaviour of the system, and provides a direct way to check how

a change in the definition of the problem affects to each subspace independently from the others:

$$s = (s_1, \ldots, s_n), \text{ with } s \in \mathcal{S} = \mathcal{S}_1 \times \ldots \times \mathcal{S}_n, \ s_i \in \mathcal{S}_i$$

Reward definition: Similarly to the state definition, instead of having a unique definition of reward function, we propose the use of multiple sub-reward functions, each one valuing the goodness of each sub-state; they can be combined into an unique reward through the use of different coefficients:

$$R(s) = R(s_1, \ldots, s_n) = \alpha_1 * R_1(s_1) + \ldots + \alpha_n * R_n(s_n), \text{ with } s_i \in \mathcal{S}_i, \ \alpha_i \in \mathbb{R}$$

Given this definition of a reward function, training the system to obtain different policies boils down into tuning each sub-reward to obtain the desired behaviour, by means of a tuning-and-test cycle as described in Section 3.2. This tuning process can be carried out in two different ways based on the desired behaviour:

(a) Modifying the coefficients: Assuming all sub-reward functions have the same range (i.e., all functions produce values in the same interval), each coefficient α_i represents the importance of each sub-state in the problem. Modifying these coefficients allows us to give more or less importance to each sub-state, and consequently to the metrics used to build it.

(b) Modifying the reward functions: Each sub-reward function represents how the system will behave respectively to each sub-state (and, therefore, to each metric used to define each sub-state). The goal of modifying the definition of a sub-reward is not to change the importance given to a set of metrics as before, but instead to modify the behaviour of the system w.r.t. to these metrics. For instance, changing one sub-reward function can imply modifying the behaviour from maximizing certain metric to minimize it.

2.2 Boosting learning time

Classical Q-Learning formulations are based on a table combining all the actions and states, and representing the expected rewards obtained at each pair. This table, originally empty, is updated at the same time the system explores the different transitions. Implicitly, at the same time the table is updated, the system is learning the unknown probabilities between states and actions ($\mathcal{P}$).

Although at first glance the bottleneck of the algorithm seems to be the number of the iterations the system needs to perform in order to explore all the transitions enough number of times, in real-world problems it is limited by the frequency at which the actions can be applied and the metrics needed to build the states can be measured. For example, in a situation in which the system is applied to a video encoding process at 24 frames per second, and a transition occurs between frames, the speed of the exploration is limited to 24 transitions per second. In the case when multiple policies are required, and consequently, one training session per policy is needed, this limitation in the speed of the algorithm can make the problem unfeasible.

However, if the probability set between states ($\mathcal{P}$) is known a priori, the algorithm does not need to wait for actual readings of the required environmental metrics to determine movements between states; contrary, it can simulate the transitions based on the information provided by $\mathcal{P}$. Following this idea, we propose an *offline learning process* which dramatically reduces the amount of time needed to obtain each policy: (1) First, all combinations between states and actions are explored enough number of times to build a transition table ($\mathcal{P}'$) that stores the probabilities of moving from one state to the others when applying a specific action. (2) Once $\mathcal{P}'$ has been built, each training process can proceed with the classical Q-Learning formulation. However the state is determined based on the information provided by $\mathcal{P}'$, not by observation. The process of building $\mathcal{P}'$ can be carried out independently of the learning process, or can be extracted from one previous training process storing explicitly the probabilities at the same time the classical Q-Learning algorithm explores the different transitions of the system. The obtained $\mathcal{P}'$ table will be valid for all the subsequent learning processes, unless the transitions between states change due to external factors (for example, processor operating frequency can be altered by changes in the temperature), or there are changes in the definition of states or actions. Note that $\mathcal{P}'$ is used only in the learning process, but once the system has finished to learn, the states and transitions are obtained directly from measurements of the system, and not from $\mathcal{P}'$.

3 Policy extraction for heterogeneous QoS

3.1 A case study: video transcoding with heterogeneous QoS requirements

In previous works [2], authors have showed that a multi-agent Q-Learning approach can be applied to a real-world HEVC (High Efficiency Video Coding) [11] encoding application to obtain simultaneous real-time encoding processes (that is, with a tight lower bound in throughput of 24 frames per second –FPS–), and meeting at the same time constraints in terms of quality, bandwidth, and power consumption. In that work, number of threads (N_{th}), core operating frequency (*freq*) and the value of the QP encoding parameter [11] (higher QP values imply lower encoding quality) were considered as dynamic knobs, and their values were periodically tuned while throughput, power consumption and quality were constantly monitored. However, in that scenario, a *single* common policy was applied to all the videos concurrently processed by the system. We report next the necessary changes to consider multiple heterogeneous QoS demands, and hence multiple policies.

Consider a similar formulation of the problem in which each state is decomposed into three different sub-states: *throughput* (in terms of FPS), *quality* (PSNR, measured in dB) and *power consumption* (W). Consider also a reward function composed by three different sub-rewards functions, each one associated to a specific sub-state, and its coefficients as described before. In this scenario, and assuming the state definition stays constant, this simple formu-
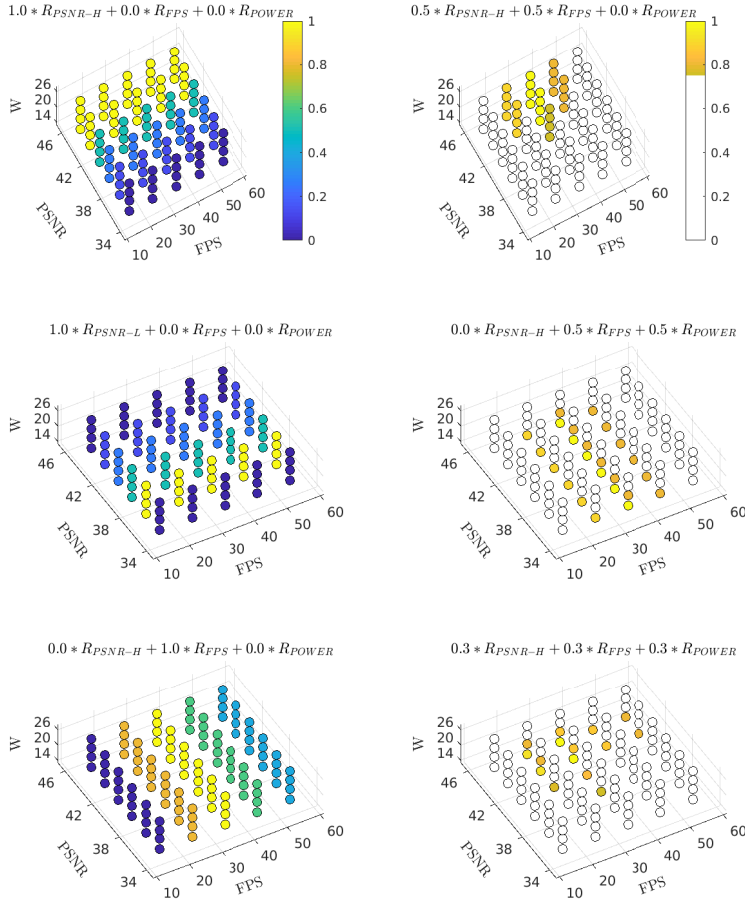
Fig. 1 Rewards obtained in the different states for different combinations of sub-reward definitions (left) and coefficients (right)

lation offers 9 different dimensions (3 reward definitions $\times$ 3 coefficient values) the designer of the experiment can tune to obtain different policies.

Figure 1 shows different combinations of sub-rewards and coefficient values. Each dot in the plot represents a state value; its color represents the reward given to that state (higher is better). On the left, three different examples of sub-reward functions are shown: (1) R_{PSNR-H} which gives maximum reward to the states with maximum quality; (2) R_{PSNR-L} which minimizes quality, but ensures a minimum quality (giving a reward of 0 to the states below the threshold); and (3) R_{FPS} which aims at obtaining real-time encoding processes giving no reward to the states with throughput below 24 FPS, maximum reward to the states between 30 and 40 FPS, and a decreasing reward to

the states above. Observe how modifications in the sub-reward functions alter which are the better solutions, or goal states (yellow points in the figure). This fact can be clearly seen in the first two plots: while the first one maximizes quality, changing the definition in the other one turns the states with minimum quality, but above the constraint (36dB in the plot), into the goal states. On the right, different combinations of the same sub-reward functions with different coefficients are shown. For the sake of clarity, only the states with reward ≥ 0.75 are colored. Here, R_{POWER} minimizes the power consumption. Observe how small changes in the way functions are combined dramatically alter the goal states of the system. For example, results vary from maximizing quality and ensuring real-time throughput on the top, to restricting the goal space to those with minimum power, meeting real time requirements in the middle, and a combination of both behaviours on the bottom. This figure shows how our proposal offers a plethora of different combinations the designer of the experiment can test, and how a huge number of different policies can be obtained following the methodology described next.

3.2 Methodology to extract multiple policies

Consider a scenario in which a video provider needs to attend multiple video encoding requests from different users with different requirements: *regular* users which require a minimum of quality, and *premium* users which need encoded videos with maximum quality, with real-time results for both. In this scenario, two different policies are desirable: one which maximizes quality for premium users (π^P), and another which guarantees a minimum quality for regular users (π^R). Additionally, when the server load increases, it would be desirable to apply different policies to minimize the resources used by each user, satisfying a minimum quality to each type of user, and still maintaining real-time throughput (24 FPS) (policies π_P and π_R).

Following the previous ideas, we propose a simple and concise methodology to generate the different policies spending a reasonable amount of time and effort in the process. In a first step, the definition of the different states and rewards is done. Each substate needs to be discretized based on expert knowledge of the problem, while the reward functions are defined based on the goal the policy has to achieve. In this step, coefficients are set together with each reward function. In a second step, a detailed simulation of the reward functions for the different states is performed to check if the goal states are the right ones (similar to the space exploration shown in Figure 1). If the states with higher reward are not the desired ones, the previous step is repeated until the reward functions and coefficients are properly tuned. In the last step, the system is trained with the previous rewards and coefficients. If the obtained results are not the desirable ones, the process is repeated until the states, rewards and coefficients are tuned. Although this process can traditionally last for an unacceptable amount of time, note that leveraging the construction of $\mathcal{P}'$, as described in Section 2.2, makes this process relatively fast and easy.
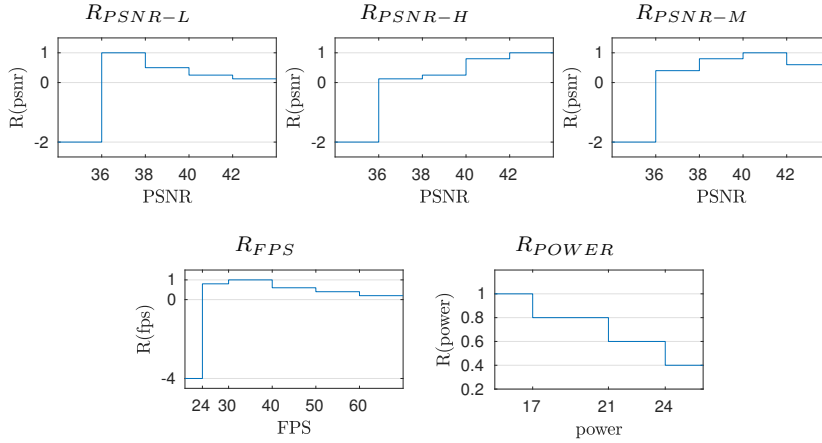
Fig. 2 Sub-reward functions used for the different policies: three sub-rewards for different levels of quality (above), and the functions for real-time encoding and power (below)

$$\pi^R : 0.7 * R_{PSNR-L} + 0.1 * R_{POWER} + 0.5 * R_{FPS}$$
$$\pi^P : 0.7 * R_{PSNR-H} + 0.0 * R_{POWER} + 0.5 * R_{FPS}$$
$$\pi_R : 0.7 * R_{PSNR-L} + 0.5 * R_{POWER} + 0.5 * R_{FPS}$$
$$\pi_P : 0.7 * R_{PSNR-M} + 0.5 * R_{POWER} + 0.5 * R_{FPS}$$

Fig. 3 Reward functions defined for each policy

Back in our scenario, the sub-rewards shown in Figure 2 where used to generate the policies: three different reward functions which provide three different levels of quality (R_{PSNR-L}, R_{PSNR-M}, and R_{PSNR-H}), a reward function which guarantees real-time encoding (R_{FPS}), and a reward function which minimizes power (R_{POWER}). To guarantee a minimum of quality, the functions shown in Figure 1 were slightly modified to give a negative reward to those states below the threshold (PSNR < 36). Similarly, the definition of R_{FPS} gives a negative reward to the states below real-time; in this case, a lower reward is given to ensure that videos are always encoded in real-time, not having processes with high quality but throughput below 24 FPS. The maximum reward of R_{FPS} is not given to 24 FPS, but to the next state, due to the fact that, being close to 24 FPS will produce a higher amount of frames being encoded below the threshold due to the variability on the content between frames. To estimate the power consumption of each application, the model $P = \text{nth} * (\alpha * \text{freq}^2 + \beta) + \gamma$ was used, setting experimentally the parameters for our specific platform, obtaining a root mean squared error of $0.97W$ and a maximum error of $2.4W$, which are negligible in our machine with a maximum energy consumption of $125W$.

	N. frames	Learning Time
MonoAgent	$n \times 3000$	$n \times 17h$
MultiAgent	$n \times 500$	$n \times 3h$
This work	$500 + (n-1) \times 500$	$3h + (n-1) \times 1min$

Table 1 Learning time to obtain n different policies by different approaches. The Mono- and Multi-agent approaches are those described in [2]. Learning time has been calculated assuming a learning rate $\approx 24fps$

We identify π^R as our base policy, used to tune and polish the actions and state definitions, at the same time the transition table $\mathcal{P}'$ is recorded. Once this policy is created, other policies can be easily derived in a reasonable time by means of the described methodology. Our policies, detailed in Figure 3, were defined based on the following ideas:

- For defining π^R, a combination of a function which minimizes PSNR but ensures a minimum quality (R_{PSNR-L}), and a function which ensures real-time encoding (R_{FPS}) were used. Additionally, a reward to minimize power consumption was incorporated with a small coefficient (R_{POWER}).
- The major difference between π^R and π^P is the reward function used to evaluate the quality. The former minimizes quality, and the latter maximizes it (R_{PSNR-H}). To achieve high quality videos without violations of throughput, the reward which minimizes power was totally removed. As real-time encoding is mandatory in both cases, this term was not modified.
- If resource minimization is desired for a regular user, π_R achieves that goal by increasing the coefficient of the function which minimizes the power consumption respectively to its counterpart π^R (from 0.1 to 0.5).
- The design decisions to create π_P are similar to the ones used to create π_R, but in this case, because obtaining high quality videos is a resource-hungry process, the sub-reward function associated with PSNR was changed to obtain still high quality videos, but lower quality than in π^P (R_{PSNR-M}).

Following the ideas described above to boost learning time, once the first policy (π^R) was defined, the training time to obtain the remaining policies was reduced from days to hours, as shown in Table 1 when compared against other traditional approaches described in [2, 8]. Note that, in this example, all learning times were extracted using the same machine and setup; also, observe that our approach inherits the advantages of the Multi-agent implementation (in terms of a reduction in the number of necessary frames to converge from 3000 to 500 compared with a Mono-agent approach), and adds additional gains as the number of desired policies is increased. In this case, adding a new policy is translated into roughly one extra minute of computing time. In the case of the traditional Mono- or Multi-agent approaches, each new policy would require a complete learning process, adding 17h and 3h per policy, respectively.

Table 2 Average knob values learnt by the system for Regular (top) and Premium (bottom) users with and without resource minimization, and output metrics for the different videos used to validate the system

	N_{th}		Freq		QP		Quality		Δ	
	π^R	π_R	π^R	π_R	π^R	π_R	π^R	π_R	π^R	π_R
v1	4.2	3.3	1.6	1.7	35.2	36.0	39.4	39.1	6.0	2.4
v2	3.1	2.9	1.4	1.5	37.0	36.9	37.8	37.8	0.4	0.4
v3	3.1	2.6	1.3	1.5	37.0	37.0	38.0	38.0	0.3	0.3
v4	3.3	2.9	1.3	1.4	37.0	37.0	37.1	37.1	0.4	0.4
avg.	3.4	2.9	1.4	1.5	36.6	36.7	38.1	38.0	1.8	0.9

	N_{th}		Freq		QP		Quality		Δ	
	π^P	π_P	π^P	π_P	π^P	π_P	π^P	π_P	π^P	π_P
v1	4.6	3.6	1.8	1.6	24.9	33.4	43.7	40.2	2.0	6.3
v2	4.2	3.3	1.7	1.4	24.9	32.2	43.2	40.2	2.3	4.2
v3	4.2	3.1	1.6	1.3	23.7	33.0	43.4	39.9	0.6	0.7
v4	4.7	3.1	1.7	1.3	22.1	33.2	43.2	39.0	0.7	0.5
avg.	4.4	3.3	1.7	1.4	23.9	33.0	43.4	39.8	1.4	2.9

3.3 Experimental setup and obtained results

The described framework and techniques have been implemented in a real server using MAMUT, a centralized resource manager described in [2]. The HEVC encoder used was Kvazaar [13], an open-source video encoder able to achieve real-time encoding processes. High Definition/1080p (1080×720 pixels) video sequences were used in all the experiments, mainly extracted from the JCT-VC benchmark [1], using three sequences to train the system, and four different sequences to run the experiments and report results (identified as v1. . . v4). In the following, the reported results correspond to average values obtained after 5 repetitions of each experiment.

Table 2 reports the behaviour of each described policy applied to each video, showing the average knob values set for number of threads (N_{th}), frequency (in GHz) and QP, and the output metrics recorded: quality (PSNR measured in dB), and real time throughput violations (measured as the percentage of time the video has been encoded below 24 FPS (-Δ-)). First, observe how changes in one reward function can produce opposite behaviours. Consider, for example, policies π^R and π^P: by modifying exclusively the reward function in charge of quality, the obtained PSNR changes drastically (38.1 to 43.4 dB, respectively). On the contrary, observe how modifications in the coefficients without altering the reward functions can keep the global behaviour of the system intact, but modify the internal actions chosen by the system. Comparing policies π^R and π_R, both achieve comparable quality levels (38.1 vs 38.0 dB), but the number of threads used when the policy π_R is acting decreases down to 1 thread in average (in the case of v1 sequence) with respect to π^R. Regarding policies π^P and π_P, observe how the impact in the number
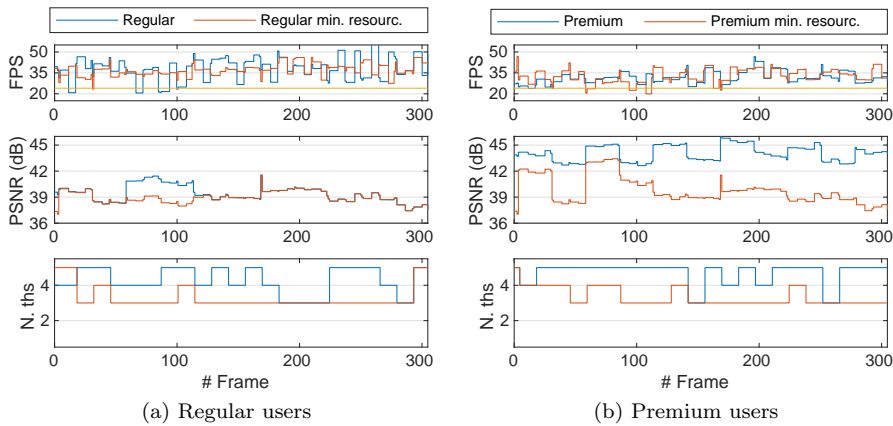
(a) Regular users   (b) Premium users

Fig. 4 System behaviour timelines and metrics obtained when encoding the v1 sequence with all different policies for Regular users (left) and Premium users (right). The yellow line indicates real-time encoding (24 FPS)

of threads is the same as in the other policies, but the change in the reward which affects quality produces slightly lower PSNR.

Figure 4 shows a timeline of the encoding process of the sequence v1 with the different policies. For clarity reasons, only the changes in number of threads is shown, but dynamic adaptation of frequency and QP values are also performed in the process. In general terms, observe how in the left, in the case of regular users, both policies π^R and π_R obtain similar (and low) quality but the number of threads used in each one changes drastically. On the right, both policies obtain higher quality results, but there is a clear difference in quality and resource usage between them as desired in our formulation.

4 Combining multiple policies into a global system

In the previous scenario, if the number of simultaneous requests is large enough, it may be the case that attending all incoming request simultaneously exhausts the available resources in the server, not being able to encode the different sequences on real time. In this situation, it is common to enqueue the requests, and attend them in order as the previous videos finish and enough resources are available. Of course, the time a request is hold in the queue does not depend only on the number of videos being encoded, but also on the type of user who made the request. To reduce the waiting time of each client on the queue, in this section we propose a 3-tier heuristic that, based on the previous obtained policies, is able to reduce the delays thanks to choosing the appropriate policy to apply to each client at each moment. This heuristic is a simple example of how a simple approach can benefit from having different policies to apply, obtaining better results than other static approaches. Moreover, note

---

Algorithm 1: 3TierPol

---

1 begin
2 | cl = queue.first;
3 | newState = currentState;
4 | **if** *(canRunClient(currentState, cl, NReg, NPrem))* **then**
5 | | newState = currentState;
6 | **else if** *(currentState == S0* **and** *canRunClient(S1, cl, NReg, NPrem))* **then**
7 | | newState = S1;   `// Try to move to S1`
8 | **else if** *(cl.type==PREMIUM* **and** *canRunClient(S2, cl, NReg, NPrem))* **then**
9 | | newState = S2;   `// Try to move to S2`
10 | **else**
11 | | queue.insert_front(cl);   `// Not enough resources`
12 | | **return** false;
13 | currentState = newState;   `// Update state`
14 | **if** *(client.type==REGULAR)* **then** nReg++ **else** nPrem++;
15 | **return** runClient(cl);   `// Start encoding request`

---

that the methodology described in the previous sections is also valid for any other approach based on multiple simultaneous policies.

4.1 A 3-tier heuristic leveraging KaaS (3TierPol)

The proposed heuristic is based on a state machine with three different states, namely: *(1)* S0, the initial state. This state is active when there are enough resources for all the clients, so that resource usage reduction is not needed. In this state, policies π^R and π^P are applied to regular and premium users respectively. *(2)* If there are not enough resources to encode all the requests, state S1 will try to reduce the resource consumption of the regular users applying policy π_R. In this state, premium users are still allowed to use as many resources as they require (policy π^P). *(3)* When there are not enough resources for the incoming request in the state S1, and only if the incoming request arrives from a premium user, the heuristic will move to state S2, minimizing the resources for both kind of users, regular and premium. In this state, policies π_R and π_P are used. Algorithm 1 shows a detailed pseudocode of the heuristic.

To determine if a new video can be encoded in a given state, or if the system needs to move to another state in order to free resources, a prediction of the number of cores in use is performed. 3TierPol stores in an internal table the average number of cores used by each policy with the training videos, as an estimation value of the cores used by future videos. Knowing at each moment the policies that are in use (i.e., the current state), the number of regular and premium users being attended, and the incoming user type, a prediction of the total number of cores in use is calculated. If the amount of predicted cores is lower or equal to the number of physical cores of the machine, the request is attended and the video encoding can commence. Else, the heuristic will try to move to the next state. If there is not enough room for the user in any state, the request is pushed in front of the queue again. Inserting the client in the

front (instead of enqueueing it again in the back) allows to attend the users in the arrival order. When a video finishes to be encoded and some resources are free, the heuristic checks if it can move to a previous state that does not minimize the resource usage and provides greater quality.

5 Experimental results

For the sake of realism, we assume that several different videos from different users arriving over time need to be served simultaneously, minimizing the waiting time of each client and meeting requirements in quality (based on the type of user) and throughput ($\geq$24 FPS). Each experiment is determined by the arrival rate (5s, 10s and 15s), and the percentage of premium users (0%, 25%, 50%, 75% and 100%). Each experiment comprises 10 sequences to be encoded, randomly selected, with a duration of 2500 frames each ($\approx$ 100 seconds at 24 FPS). To obtain reliable data, each configuration of frequency and premium/total users relation was explored through 5 different combinations of 10 videos, and each combination was run 3 times, reporting average values. Although three different arrival frequencies where explored in the experiments (5s, 10s and 15s), only the results of 10s are shown in the following, obtaining similar and comparable measurements at the other frequencies.

For the sake of comparison, we have implemented two additional alternatives to the heuristic approach (3TierPol) presented in Section 4. Both alternatives implement a static decision making process, choosing the policy to serve each video based on the kind of user, and not on the environment. The policy utilized does not change during the whole encoding process:

1. 1Pol-Mamut: This strategy corresponds to a traditional Q-Learning system using only one policy (π^R) to encode all the sequences as an extreme case, trading off quality for throughput, as implemented by MAMUT [2].
2. 2Pol: Due to the existence of different user types with different requirements, we have implemented a Q-Learning system using two different tables, depending on the user type attended (π^R and π^P). This policy is the opposite to the previous one: it offers maximum quality to each type of user, without considering decreasing quality to serve more users simultaneously.

In both cases, the algorithm to determine whether a video can be encoded or needs to wait on the queue is the same as that used in 3TierPol.

5.1 Comparative performance discussion

The plot on the top of Figure 5 shows the amount of users attended per minute (on average) based on the number of premium user requests for the three explored approaches. Depending on the number of premium users, two different behaviours are observed: when the amount of premium users is below 50%, and when the amount is greater or equal. On the first group, 3TierPol
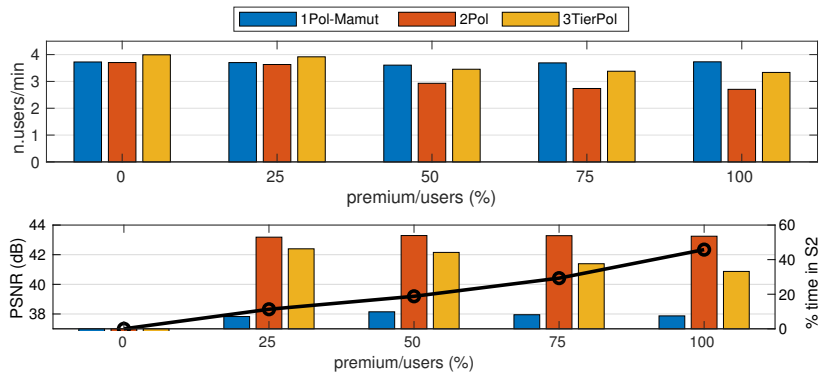
Fig. 5 Users per minute attended by each approach with users requests arriving every 10 seconds (top), and quality obtained by each approach when encoding videos from premium users (bars), and percentage of time 3TierPol is in the S2 state (line)

outperforms the other strategies, and it is able to process more users per minute reducing slightly the quality obtained (0.13 dB for regular and 1.9 dB for premium users in the worst case), see the bottom plot in Figure 5.

Diving into details of the behaviour of the second group (premium $\geq 50\%$), we observe how the amount of premium users impacts in the performance of the system: as many premium users are attended, more resources are used to encode these videos, and less resources are available to encode new incoming requests. This behaviour is shown in 2Pol and 3TierPol, but not in 1Pol-Mamut. In the case of 1Pol-Mamut, the behaviour is to encode all the videos with the same policy (π^R), used by the others approaches to encode only regular users. On the one hand, this approach uses less resources to process each premium user (because they are considered as regular), but on the other hand, the quality obtained (lower plot on the Figure) is quite lower than the other approaches, obtaining not admissible levels for premium users (up to 5.4 less dB in PSNR). Second, when comparing 2Pol and 3TierPol, we can observe how the latter is able to serve more videos at the same time (up to 1.24×), reducing slightly the obtained quality (a loss of 2.4 dB in the worst case). Finally, observing the psnr obtained in each experiment (lower plot) we can see the internal behaviour of 3TierPol: as the number of premium users increases, 3TierPol needs to move to states that reduce the resource usage at the expense of reducing the quality. This can be seen in the plot at the bottom of Figure 5, showing the percentage of the time 3TierPol is in the S2 state depending on the number of premium users. Observe how the time 3TierPol is at state S2 increases with the number of premium users.

6 Conclusions

Training a Q-Learning system from scratch can be an ardours work. In addition, if different policies with different behaviour for the same system are

needed, training becomes an unfeasible task in practice. To tackle this problems, in this work we have proposed a methodology to obtain new policies with minimum effort, once the definitions have been tuned for a first desired behaviour. To achieve that goal, two different mechanisms were proposed: defining the reward functions to obtain the desired behaviour, and speeding up the learning time of the problem. Together, these techniques make it feasible to explore multiple combinations without much effort and yield dramatic improvements in terms of learning time. We have shown how our proposal can be applied to a real scenario obtaining different policies to encode videos on real-time with different QoS objectives, ensuring a minimum of quality in all cases and varying the resource usage at will of the policies designer.

By means of a realistic scenario where multiple requests have to be attended simultaneously, we have shown how a heuristic approach can handle multiple policies at the same time, being able to decide which policy apply to each user at each moment, obtaining improvements up to $1.24\times$ in the number of users attended per unit of time when comparing with less sophisticated approaches.

References

1. Bossen F, Common H (2013) Test conditions and software reference configurations. JCT-VC Doc
2. Costero L, Iranfar A, Zapater M, Igual FD, Olcoz K, Atienza D (2019) MAMUT: Multi-agent reinforcement learning for efficient real-time multi-user video transcoding. In: Design, Automation Test in Europe Conference
3. Farhad SM, Bappi MSI, Ghosh A (2016) Dynamic Resource Provisioning for Video Transcoding in IaaS Cloud. In: 2016 IEEE 18th Int. Conf. on High Performance Computing and Communications, IEEE, pp 380–384
4. Gao G, Wen Y, Westphal C (2016) Dynamic Resource Provisioning with QoS Guarantee for Video Transcoding in Online Video Sharing Service. In: 2016 ACM Multimedia Conf. - MM '16, pp 868–877
5. Gao G, Wen Y, Westphal C (2019) Dynamic Priority-Based Resource Provisioning for Video Transcoding With Heterogeneous QoS. IEEE Trans on Circuits and Systems for Video Technology 29(5):1515–1529
6. Ho HN, Lee E (2015) Model-based reinforcement learning approach for planning in self-adaptive software system. In: Int. Conf. on Ubiquitous Information Management and Communication, pp 103:1–103:8
7. IBM (2005) An architectural blueprint for autonomic computing. Tech. rep., IBM
8. Iranfar A, Zapater M, Atienza D (2018) Machine learning-based quality-aware power and thermal management of multistream HEVC encoding on multicore servers. IEEE Trans Parallel Distrib Syst 29(10):2268–2281
9. Sembiring K, Beyer A (2013) Dynamic resource allocation for cloud-based media processing. In: ACM Workshop on Network and Operating Systems Support for Digital Audio and Video, ACM Press, pp 49–54

10. Singh S, Chana I (2015) QoS-Aware Autonomic Resource Management in Cloud Computing. ACM Computing Surveys 48(3):1–46
11. Sullivan GJ, Ohm JR, Han WJ, Wiegand T (2012) Overview of the high efficiency video coding (HEVC) standard. IEEE Trans on Circuits and Systems for Video Technology 22(12):1649–1668
12. Sutton RS, Barto AG (1998) Reinforcement learning: An introduction, vol 1. MIT press Cambridge
13. Viitanen M, Koivula A, Lemmetti A, Ylä-Outinen A, Vanne J, Hämäläinen TD (2016) Kvazaar: Open-source HEVC/H.265 encoder. In: 24th ACM Int. Conf. on Multimedia
14. Wang L, Gelenbe E (2018) Adaptive Dispatching of Tasks in the Cloud. IEEE Transactions on Cloud Computing 6(1):33–45