

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Design and implementation of distributed resource management for time-sensitive applications

This is a pre print version of the following article:

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1615499> since 2016-12-30T16:31:54Z

Published version:

DOI:10.1016/j.automatica.2015.09.015

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Design and Implementation of Distributed Resource Management for Time-Sensitive Applications [★]

Georgios C. Chasparis ^a, Martina Maggio ^b, Enrico Bini ^b, Karl-Erik Årzén ^b

^aSoftware Competence Center Hagenberg GmbH, Softwarepark 21, A-4232 Hagenberg, Austria

^bDepartment of Automatic Control, Lund University, Sweden

Abstract

In this paper, we address distributed convergence to *fair* allocations of CPU resources for time-sensitive applications. We propose a novel resource management framework where a centralized objective for fair allocations is decomposed into a pair of performance-driven recursive processes for updating: (a) the allocation of computing bandwidth to the applications (*resource adaptation*), executed by the resource manager, and (b) the service level of each application (*service-level adaptation*), executed by each application independently. We provide conditions under which the distributed recursive scheme exhibits convergence to solutions of the centralized objective (i.e., fair allocations). Contrary to prior work on centralized optimization schemes, the proposed framework exhibits adaptivity and robustness to changes both in the number and nature of applications, while it assumes minimum information available to both applications and the resource manager. We finally validate our framework with simulations using the TrueTime toolbox in MATLAB/Simulink.

Key words: Resource Management; Distributed Optimization; Real-Time Systems

1 Introduction

The current trend in embedded computing demands that the number of applications sharing the same execution platform increases. This is due to the increased capacity of the new hardware platforms, e.g., through the use of multi-core techniques. An example includes the move from federated to integrated system architectures in the automotive industry [9].

In such scenarios, the need for better mechanisms for controlling the rate of execution of each application becomes apparent. To this end, virtualization or resource reservation techniques [1, 17] are used. According to these techniques, each reservation is viewed as a *virtual*

processor (or *platform*) executing at a fraction of the speed of the physical processor, i.e., the *bandwidth* of the reservation. An orthogonal dimension along which the performance of an application can be tuned is the selection of its *service level*. It is assumed that an application is able to execute at different service levels, where a higher service level implies a higher quality-of-service (QoS). Examples include the adjustable video resolutions and the adjustable sampling rates of a controller.

Typically this problem is solved by using a *resource manager* (RM), which is in charge of: (a) *assigning virtual processors to the applications*, (b) *monitoring the use of resources*, and (c) *assigning the service level to each application*. The goal of the RM is to maximize the overall delivered QoS. This is often done through *centralized optimization* and the use of *feedback* from the applications.

RM's that are based on the concept of feedback, monitor the progress of the applications and adjust the virtual platforms based on measurements [10, 21]. In these early approaches, however, quality adjustment was not considered. Instead, reference [8] proposed an inner loop to control the resource allocation nested within an outer loop that controls the overall delivered quality.

Optimization-based resource managers have also re-

[★] The research leading to these results was supported by the Linneaus Center LCCC, the Swedish VR project n.2011-3635 "Feedback-based resource management for embedded multi-core platform", and the Marie Curie Intra European Fellowship within the 7th European Community Framework Programme. An earlier version of parts of this paper appeared in [6] and its implementation framework appeared in [16].

Email addresses: georgios.chasparis@scch.at (Georgios C. Chasparis), martina.maggio@control.lth.se (Martina Maggio), bini@control.lth.se (Enrico Bini), karlerik@control.lth.se (Karl-Erik Årzén).

ceived considerable attention [15, 19]. These approaches, however, rely on the solution of a centralized optimization that determines *both* the amount of assigned resources and the service levels of all applications [3, 19, 20]. In the context of networking, reference [13] models the service provided by a set of servers to workloads belonging to different classes as a utility maximization problem. However, there is no notion of adjustment of the service level of the applications.

An example of a combined use of optimization and feedback was developed in the ACTORS project [2, 3]. In that project, applications provide a table to the RM describing the required amount of CPU resources and the expected QoS achieved at each supported service level [2, 3]. In the multi-core case, applications are partitioned over the cores and the amount of resources is given for each individual partition. Then, the RM decides the service level of all applications and how the partitions should be mapped to physical cores using a combination of Integer Linear Programming (ILP) and first-fit-decrease (FFD) for bin packing.

On-line centralized optimization schemes have several weaknesses. First, the complexity of the solvers used to implement the RM (such as ILP solvers) grows significantly with the number of applications. It is impractical to have a RM that optimally assigns resources at the price of a large consumption of resources by the RM itself. Second, to enable a meaningful formulation of a cost function in such optimization problems, the RM must compare the quality delivered by different applications. This comparison is unnatural because the concept of quality is extremely application dependent. Finally, a proper assignment of service levels requires application knowledge. In particular, applications must inform the RM about the available service levels and the expected consumed resources at each service level, increasing significantly communication complexity.

To this end, distributed optimization schemes have recently attracted considerable attention. Reference [22] considered a cooperative game formulation for job allocation to service providers in grid computing. Reference [23] proposed a non-cooperative game-theoretic formulation to allocate computational resources to a given number of tasks in cloud computing. Tasks have full knowledge of the available resources and try to maximize their own utility function. Similarly, in [11] the load balancing problem is formulated as a non-cooperative game.

Contrary to the grid computing setup of [22] or the load balancing problem of [11, 23], this paper addresses a lower-level resource allocation problem, that is, *the establishment of fair allocations of CPU bandwidth among time-sensitive applications which adjust their own service levels*. Contrary to the cloud computing setup of [23], a game-theoretic formulation may not easily be

motivated practically when addressing such lower-level (single node) resource allocation problems. Instead, we propose a distributed optimization scheme, according to which a centralized objective for fair allocations is decomposed into a pair of performance-driven recursive processes for updating: (a) the allocation of computing bandwidth to the applications (*resource adaptation*), executed by the RM, and (b) the service level of each application (*service-level adaptation*), executed by each application independently. We provide conditions under which the distributed recursive scheme exhibits convergence to fair allocations).

The proposed scheme introduces a design technique for allocating computing bandwidth to *time-sensitive applications*, i.e., applications whose performance is subject to strict time deadlines, such as multimedia and control applications. In particular, the proposed scheme: (a) exhibits linear complexity with the number of applications, (b) drops the assumption that the RM has knowledge of application details, and (c) exhibits adaptivity and robustness to the number and nature of applications. This paper extends the theoretical contributions of [6] by addressing global convergence and asynchronous updates. Furthermore, reference [16] presents the full implementation framework in Linux.

The paper is organized as follows. Section 2 provides the overall framework, while Section 3 presents the distributed scheme for resource allocation. Section 4 presents the convergence behavior for the synchronous and asynchronous case. Section 5 presents technical details required for the derivation of the main results in Section 4. Section 6 provides selective simulations. Finally, Section 7 presents concluding remarks.

Notation:

- $\Pi_{[a,b]}$ is the projection onto the set $[a, b]$.
- For some finite sequence $\{x_1, x_2, \dots, x_n\}$ in \mathbb{R} , define $\text{col}\{x_1, x_2, \dots, x_n\}$ to be the column vector in \mathbb{R}^n with entries $\{x_1, x_2, \dots, x_n\}$.
- For any $x \in \mathbb{R}$, define the operator $[x]_-$ as follows:

$$[x]_- \triangleq \begin{cases} x, & x \leq 0 \\ 0, & x > 0 \end{cases}$$

- For any $x \in \mathbb{R}^n$ and set $A \subset \mathbb{R}^n$, define $\text{dist}(x, A) \triangleq \inf_{y \in A} \|x - y\|$, where $\|\cdot\|$ denotes the Euclidean norm.
- For some finite set A , $|A|$ denotes the cardinality of A .

2 Framework & Problem Formulation

2.1 Resource manager & applications

The overall framework is illustrated in Figure 1. A set \mathcal{I} of n (time-sensitive) applications are sharing the same




Fig. 1. Resource management framework.

CPU platform. Let i be a representative element of this set. Since we allow applications to dynamically join or leave, the number n may not be constant over time.

The resources are managed by a RM that allocates resources through a Constant Bandwidth Server (CBS) [1] with period P_i and budget Q_i . Hence, application i is assigned a *virtual platform* with bandwidth $v_i = Q_i/P_i$ corresponding to a fraction of the computing power (or speed) of a single CPU. Obviously, not all virtual platforms v_i are feasible, since their sum cannot exceed the number κ of available CPU's. Formally, we define the set of **feasible virtual platforms**, (v_1, \dots, v_n) , as

$$\mathcal{V} \doteq \left\{ \mathbf{v} = (v_1, \dots, v_n) \in [0, 1]^n : \sum_{i=1}^n v_i \leq \kappa \right\}. \quad (1)$$

In this study, the main concern is the computation of the allocation \mathbf{v} in real time such that a centralized objective is achieved. However, we will not be concerned with the exact mapping of this allocation onto the available cores. Such mapping can be performed by a standard first-fit-decrease algorithm. Furthermore, in practice, more constraints might be present, especially if applications are single-threaded (i.e., they may only run on a single core). In this case, the above feasibility constraint will be a relaxed version of the original problem, however, the forthcoming analysis can be modified in a straightforward manner to incorporate additional constraints on \mathcal{V} .

Furthermore, each application $i \in \mathcal{I}$ may change its *service level*, s_i . It represents a qualitative indicator of the delivered quality of application i , assuming sufficient amount of resources v_i . Naturally, it can be represented by a real number $s_i \in \mathcal{S}_i \doteq [\underline{s}_i, \infty) \subset \mathbb{R}$, where $\underline{s}_i > 0$ is the minimum possible service level of application i . The domain \mathcal{S}_i inherits the partial ordering from \mathbb{R} , according to which $s_i' \leq s_i''$ implies that the quality delivered at service level s_i' is smaller than or equal to the corresponding quality delivered at s_i'' . The physical interpretation of the service level may only be realized in the context of

a specific application. It may represent any quality indicator of the application, e.g., the inverse of the accuracy of an iterative optimization routine, the details of an MPEG player and the sampling frequency of a controller. We denote $\mathbf{s} \doteq \text{col}\{s_1, \dots, s_n\}$ the service level profile of all applications evolving within $\mathcal{S} \doteq \mathcal{S}_1 \times \dots \times \mathcal{S}_n$.

We implicitly assume here that an application may always increase its service level providing the necessary resources, however, in practice it will always be constrained due to the constraints imposed in v_i . Note, finally, that the service level s_i is an *internal* state of application i , i.e., it can be written/read only by i .

2.2 The matching function

To be able to assess the performance of a time-sensitive application, it is necessary to introduce a *performance function*. The RM is able to measure at any time $t \geq 0$, (a) the *soft-deadline* of each application i , $D_i(t)$, which is the time duration of its last CPU reservation, and (b) the corresponding *job-response time*, $R_i(t)$, which is the time elapsed from the start time to the finishing time of a job during its last reservation. A natural definition of such performance function for time-sensitive applications is the following *matching function*:

$$f_i(t) \doteq \frac{D_i(t)}{R_i(t)} - 1, \quad (2)$$

Note that $f_i \geq -1$, a property that will be used often.

Based on the above definition, we define a *perfect* matching between D_i and R_i to be the situation at which $|f_i| \leq \delta$, for some small $\delta > 0$. This is the case when application i has the correct amount of resources. Instead, a *scarce* matching describes a situation at which $f_i < -\delta$, i.e., when application i does not have enough resources, and an *abundant* matching describes a situation at which $f_i > \delta$, i.e., when application i has more than enough resources.

2.2.1 Nominal matching function

The matching function depends indirectly on the virtual platform v_i and the service level s_i of application i . For a large class of applications, we may derive a nominal representation of the matching function, denoted φ_i , as a function of s_i and v_i as follows:

$$\varphi_i(s_i, v_i) \doteq \beta_i \frac{v_i}{s_i} - 1, \quad (3)$$

for some positive constant β_i . For example, for *multi-media applications*, the soft deadline D_i can be considered constant, while the response time can be defined as $R_i = C_i/v_i$, where $C_i = \alpha_i s_i$ is the execution time per job (at a service level s_i) and v_i is the speed of execution.

Similarly, in *control applications*, $R_i = C_i/v_i$ where C_i denotes nominal time of execution, while the soft deadline D_i is considered inverse proportional to the sampling frequency (or service level) s_i , i.e., $D_i = \alpha_i/s_i$. Both cases lead to a matching function with the form of (3).

It is evident that the nominal matching function (3) satisfies the following properties: For some $s_i, s'_i \in \mathcal{S}_i$ and $v_i, v'_i \in \mathcal{V}$: (P1) $s_i \neq 0 \Rightarrow \varphi_i(s_i, 0) < 0$, that is, the matching must certainly be scarce if no resources are assigned; (P2) $s_i \geq s'_i \Rightarrow \varphi_i(s_i, v_i) \leq \varphi_i(s'_i, v_i)$, if application i lowers its service level, then the performance should not decrease; (P3) $v_i \geq v'_i \Rightarrow \varphi_i(s_i, v_i) \geq \varphi_i(s_i, v'_i)$, if the bandwidth of application i decreases, then the performance should not increase.

2.3 Application weights

The RM may also assign weights to the applications. We introduce the weight $\lambda_i \in (0, 1]$ to represent the *importance* that the RM assigns to application i when adjusting its virtual platform v_i . As we shall see in a forthcoming section, the weights $\{\lambda_i\}$ will determine the direction of adjustment of the virtual platforms $\{v_i\}$ by the RM. The weights $\{\lambda_i\}$ are considered given and determined by the RM.

2.4 Fair allocations & objective

To define *fair allocations*, for each application i , we introduce the following **nominal fairness measure**:

$$\Phi_i(\mathbf{s}, \mathbf{v}) \doteq -(1 - \bar{v}_i)\lambda_i [\varphi_i(s_i, v_i)]_- + \bar{v}_i \sum_{j \neq i} \lambda_j [\varphi_j(s_j, v_j)]_-, \quad (4)$$

where $\bar{v}_i \doteq v_i/\kappa$ is the normalized virtual platform of i over the number of cores.

The function Φ_i captures the deficiency in resources of application i compared to the rest of applications. When application i is not performing well, i.e., $\varphi_i(s_i, v_i) < 0$, and its available resources v_i are small, while the rest of applications are performing well, we should expect large values for Φ_i .

Definition 2.1 (Fair allocation) *For some service level profile $\mathbf{s} \in \mathcal{S}$, a virtual platform profile $\mathbf{v}^* \in \mathcal{V}$ is fair or balanced if $\Phi_i(\mathbf{s}, \mathbf{v}^*) \equiv 0$ for all $i \in \mathcal{I}$.*

According to Definition 2.1, an allocation of virtual platforms \mathbf{v}^* is *fair* for application i only if $v_i^* \neq 0$, since at zero resources $\varphi_i(s_i, 0) < 0$ and $\Phi_i(\mathbf{s}, \mathbf{v}^*) < 0$. Thus, an allocation \mathbf{v}^* is *fair* if either (a) $[\varphi_i(s_i, v_i^*)]_- \equiv 0$ for all i or (b) $[\varphi_i(s_i, v_i^*)]_- < 0$ for all i and the ratio of resources $\bar{v}_i^*/(1 - \bar{v}_i^*)$ coincides with the corresponding ratio of weighted matching functions. Since $[\varphi_i]_- \in [-1, 0]$, case

(b) implies that *the resources are balanced with the negative performances*. For example, if \bar{v}_i^* is large compared to the rest $1 - \bar{v}_i^*$, then $[\varphi_i]_-$ has to be sufficiently negative, i.e., application i should not perform so well compared to the rest. Informally, there could not be application i that monopolizes the resources at a fair allocation when i performs well and the others do not.

The above fairness definition introduces a potential centralized problem for fair allocations.

$$\begin{cases} \min_{\mathbf{s} \in \mathcal{S}, \mathbf{v} \in \mathcal{V}} \sum_{i \in \mathcal{I}} |\Phi_i(\mathbf{s}, \mathbf{v})| \\ \text{s.t.} & \varphi_i(s_i, v_i) = 0, \quad \forall i \in \mathcal{I}. \end{cases} \quad (5)$$

However, neither the RM nor application i has complete knowledge of the details of the nominal matching function $\varphi_i(s_i, v_i)$. Thus, on-line centralized optimization is highly prohibited. Instead, optimization may only be based on measurements collected during run-time.

3 Adjustment Dynamics

In this section, the centralized objective of fair allocations is decomposed into a pair of performance-driven recursive schemes, executed independently by the RM and the applications, thus avoiding the computation and communication complexity of centralized optimization.

3.1 Resource adaptation

The RM updates the bandwidth $\bar{v}_i = v_i/\kappa$, normalized with respect to the number of cores κ . The *unused bandwidth* is $v_r = \kappa - \sum_{i=1}^n v_i$, and its normalized version is $\bar{v}_r = 1 - \sum_{i=1}^n \bar{v}_i$. At time instances $t_k, k = 0, 1, \dots$ the RM assigns resources as follows:

- (1) It measures the matching function $f_i = f_i(t_k)$ for each $i \in \mathcal{I}$, and computes $[f_i(t_k)]_-$.
- (2) It updates the normalized resource allocation vector $\bar{\mathbf{v}} \doteq (\bar{v}_1, \dots, \bar{v}_n)$ as follows:

$$\bar{v}_i(t_{k+1}) = \Pi_{\bar{\mathcal{V}}_i} \left[\bar{v}_i(t_k) + \epsilon F_i(t_k) \right] \quad (6)$$

for each $i = 1, \dots, n$, where $\bar{\mathcal{V}}_i \doteq [0, 1/\kappa]$ and $F_i(t_k)$ is the *observed* fairness measure defined as follows:

$$F_i(t_k) \doteq -(1 - \bar{v}_i(t_k))\lambda_i [f_i(t_k)]_- + \bar{v}_i(t_k) \sum_{j \neq i} \lambda_j [f_j(t_k)]_-.$$

Furthermore, the unused bandwidth is updated according to $\bar{v}_r(t_{k+1}) = 1 - \sum_{i=1}^n \bar{v}_i(t_{k+1})$.

- (3) It computes the original bandwidths by setting $v_i(t_{k+1}) = \kappa \bar{v}_i(t_{k+1})$.
- (4) It updates the time index $k \leftarrow k + 1$ and repeats.

Note that according to the definition of $F_i(t_k)$, if there is a deficiency of resources for i , i.e., $F_i(t_k) > 0$, then \bar{v}_i will increase, otherwise it will decrease. We consider a *constant* step size $\epsilon > 0$, since it provides an adaptive response to changes in the number of applications. In some cases, we will use vector notation, denoting $\bar{\mathbf{v}} \doteq \text{col}\{\bar{v}_1, \dots, \bar{v}_n\}$ which evolves over $\bar{\mathcal{V}} \doteq \bar{\mathcal{V}}_1 \times \dots \times \bar{\mathcal{V}}_n$.

Recursion (6) for the adjustment of resources was motivated by the standard *replicator dynamics* (cf., [24, Chapter 3]) and in particular the discrete-time equivalent (namely *reinforcement learning*) introduced in [7]. Note that the RM time complexity is linear with respect to the number of applications, as demonstrated in [16].

3.2 Service level adaptation

The RM provides information to each application i through an observation signal $Y_i(t_k)$, $k = 0, 1, \dots$, that captures its performance. Applications are designed to adjust their service levels based on $Y_i(t_k)$ as follows:

$$s_i(t_{k+1}) = \Pi_{S_i} [s_i(t_k) + \epsilon Y_i(t_k)], \quad i \in \mathcal{I}. \quad (7)$$

A natural selection for the observation signal is to set $Y_i(t_k) \equiv f_i(t_k)$, i.e., the observed matching function. In this scenario, the application i will increase its service level if $f_i(t_k) > 0$, otherwise it will decrease it. Alternative observation terms can also be defined with similar properties as demonstrated in [6].

4 Convergence

In this section, a characterization of the convergence properties of the proposed distributed scheme is provided in case of (a) *synchronous* applications' updates, and (b) *asynchronous* applications' updates. Asynchronous updates constitute a form of perturbation of the nominal synchronous behavior which may alter significantly the performance of the scheme. Perturbations due to measurement noise are not present, since the RM has direct access to the response time of each application. However, internal uncertainties of an application may result in small deviations from its nominal matching function. Due to the small probability density of such events, we will not discuss robustness with respect to such uncertainties, i.e., **for the remainder of the paper**, we consider $f_i(t) \equiv \varphi(s_i(t), v_i(t))$, where the nominal matching function satisfies (3).

4.1 Feasibility

The first property of the proposed adjustment process is the feasibility of the resulting virtual platforms.

Proposition 4.1 (Feasible allocations) *For sufficiently small step size $\epsilon = \epsilon(n) > 0$, the update recursion*

of projected virtual platforms (6) leads to a sequence of virtual platforms $\{\mathbf{v}(t_k)\}$ which satisfies $\mathbf{v}(t_k) \in \mathcal{V}$ for all $k = 0, 1, \dots$ as long as $\mathbf{v}(t_0) \in \mathcal{V}$.

Proof. See Appendix A. \square

4.2 Minimum guarantees

The adjustment process guarantees *starvation avoidance*, i.e., a positive amount of resources (at least $\epsilon > 0$) to all applications with non-zero weight. Furthermore, it guarantees a *balance* condition, according to which, in overloaded CPU's, no application is able to monopolize resources.

Before stating formally these observations, define:

- $L \doteq \sup_{i \in \mathcal{I}, k \in \mathbb{N}} |F_i(t_k)| < \infty$,
- $\lambda \doteq \min_{i \in \mathcal{I}} \lambda_i > 0$.

Proposition 4.2 (Starvation avoidance) *There exists $\epsilon^* = \epsilon^*(n) < 1/(L+1)\kappa$ with $\epsilon^* \rightarrow 0$ as $n \rightarrow \infty$, such that for any step size $\epsilon \leq \epsilon^*$, $\inf_{k \in \mathbb{N}} \bar{v}_i(t_k) > \epsilon$ for all i .*

Proof. See Appendix B. \square

Proposition 4.3 (Balance) *Pick $0 < \zeta \leq 1/\kappa$ such that $\max_{i \in \mathcal{I}} \{\beta_i \kappa \zeta / \underline{s}_i - 1\} < 0$. For any $\epsilon = \epsilon(\zeta) < \zeta/L$, there exists a number of applications $n^* = n^*(\zeta)$ such that, for any set of applications \mathcal{I} with $|\mathcal{I}| \geq n^*$ and for any $i \in \mathcal{I}$, the following hold:*

- (1) *if $\bar{v}_i(t_0) > \zeta$, then $\bar{v}_i(t_k) \leq \zeta$ after a finite k ;*
- (2) *if $\bar{v}_i(t_0) \leq \zeta$, then $\bar{v}_i(t_k) \leq \zeta$ for all $k = 1, 2, \dots$.*

Also, as $\zeta \rightarrow 0$, $n^(\zeta) \rightarrow \infty$ and $\zeta n^*(\zeta) \rightarrow c$, for some positive constant c .*

Proposition 4.3 states that if we pick ζ such that $\bar{v}_i \leq \zeta$ implies negative matching function for all i , and we consider a sufficiently large number of applications $n \geq n^*$, then all applications will end up with a virtual platform less than ζ within finite time. Informally, when the CPU is overloaded, no application can monopolize the available resources. **Proof.** See Appendix C. \square

4.3 Synchronous convergence

In the forthcoming convergence analysis, we will consider either one of the following hypotheses:

- (H1) Let $\beta_i / \underline{s}_i < 1$ for all i .

(H2) Let the number of applications n be sufficiently large such that, there exists $0 < \zeta \leq 1/\kappa$ satisfying properties (1) and (2) of Proposition 4.3 for $n^*(\zeta) \leq n$.

Hypothesis (H1) corresponds to the case where the applications are highly demanding, while (H2) corresponds to the case where the assigned resources is small compared to the number of applications.

The asymptotic behavior of recursions (6)–(7) can be associated with the limit points of the following collection of (nonlinear) ordinary differential equations (ODE):

$$\begin{pmatrix} \dot{s}_i(\tau) \\ \dot{\bar{v}}_i(\tau) \end{pmatrix} = \begin{pmatrix} \varphi_i(s_i(\tau), \kappa \bar{v}_i(\tau)) \\ \Phi_i(\mathbf{s}(\tau), \kappa \bar{\mathbf{v}}(\tau)) \end{pmatrix} + \mathbf{z}_i(\tau), \quad i \in \mathcal{I}, \quad (8)$$

as the step size ϵ approaches zero, where τ refers to the time-index of the ODE. The vector $\mathbf{z}_i(\tau)$ represents the vector of minimum length required to drive $\bar{v}_i(\tau)$ back to \bar{V}_i and $s_i(\tau)$ back to \mathcal{S}_i . Define $(\mathbf{s}^{\tau_0}(\cdot), \bar{\mathbf{v}}^{\tau_0}(\cdot))$ to be the solution of the ODE (8) starting at $(\mathbf{s}(\tau_0), \bar{\mathbf{v}}(\tau_0))$.

Consider also the linear-time interpolation $(s_{i,\epsilon}(t), \bar{v}_{i,\epsilon}(t))$ of $\{(s_i(t_k), \bar{v}_i(t_k))\}_k$, defined as follows: $s_{i,\epsilon}(t) = s_i(t_k)$, and $\bar{v}_{i,\epsilon}(t) = \bar{v}_i(t_k)$, for every $t_k \leq t < t_{k+1}$. Introduce also the vector notation $\mathbf{s}_\epsilon \doteq \text{col}\{s_{i,\epsilon}\}_i$ and $\bar{\mathbf{v}}_\epsilon \doteq \text{col}\{\bar{v}_{i,\epsilon}\}_i$.

Theorem 4.1 (Synchronous convergence) *The following hold:*

(1) *If either (H1) or (H2) applies, the ODE (8) exhibits stationary points $(\mathbf{s}^*, \bar{\mathbf{v}}^*)$, which satisfy:*

$$s_i^* = \underline{s}_i, \quad \text{and} \quad \begin{cases} \Phi_i(\mathbf{s}^*, \kappa \bar{\mathbf{v}}^*) = 0, & \text{or} \\ \Phi_i(\mathbf{s}^*, \kappa \bar{\mathbf{v}}^*) > 0, \bar{v}_i^* = 1/\kappa \end{cases} \quad \forall i. \quad (9)$$

(2) *If either (i) $\beta_i \kappa / \underline{s}_i \rightarrow 0$ for all i , or (ii) $n \rightarrow \infty$, then (a) any stationary point of the ODE (8) satisfies:*

$$s_i^* = \underline{s}_i, \quad \bar{v}_i^* \rightarrow \min \left\{ \frac{1}{\kappa}, \frac{\lambda_i}{\sum_j \lambda_j} \right\}, \quad \forall i \in \mathcal{I}. \quad (10)$$

(b) $(s_i(t_k), \bar{v}_i(t_k)) \rightarrow (s_i^*, \bar{v}_i^*)$ as $k \rightarrow \infty$ and $\epsilon \rightarrow 0$.¹

Proof. The proof is an immediate implication of a series of propositions presented in detail in Section 5.1. \square

In other words, Theorem 4.1 states that stationary points of the ODE (8) are *fair allocations* (except for trivial cases where a virtual platform is limited by the

¹ By $x(t) \rightarrow A$ for a set A , we mean $\lim_{t \rightarrow \infty} \text{dist}(x(t), A) = 0$.




Fig. 2. Visualization of asynchronous applications.

size of the core). Furthermore, when the CPU is overloaded (i.e., either due to (i) a high demand, or (ii) a large number of applications), then *the unique fair allocation is a global attractor of the distributed process.*

4.4 Asynchronous convergence

So far, we have implicitly assumed that the adjustment dynamics (6)–(7) have synchronized clocks. However, the virtual platform, v_i , indirectly determines application i 's speed of execution. Hence, the update rate of the service level s_i varies over time.

Under asynchronous updates, the asymptotic allocation of virtual platforms may not necessarily be fair to all applications. Consider, for example, the case where an application i does *not* update its service level, while all other applications do. Then, under limited available resources, application i will retain a sufficiently negative matching function f_i , while the matching functions of all other applications steadily approach zero. This situation may lead to application i getting asymptotically a larger virtual platform independently of its weight λ_i .

To address asynchronous updates, we first introduce the following notation, also visualized in Figure 2.

- t denotes the actual run time;
- t_k^i denotes the update instances of application i ;
- t_m denotes the update instances of the RM;
- $\bar{k}(t, i) \doteq \{k \in \mathbb{N} : t_k^i \leq t < t_{k+1}^i\}$ denotes the most recent to t update index of application i ;
- $\bar{m}(t) \doteq \{m \in \mathbb{N} : t_m \leq t < t_{m+1}\}$ denotes the most recent to t update index of the RM;
- $\psi_i(m) \doteq \max\{m' \leq m : \exists k \text{ s.t. } t_{m'}^i \leq t_k^i < t_m\}$ denotes the most recent update of the RM after which the last update of application i occurred. For example, in Figure 2, $\psi_i(\bar{m}(t)) = \bar{m}(t_k^i)$. We set $\psi_i(m) = 0$ if there exists no k such that $t_k^i < t_m$.
- $N_i(k) \doteq \bar{m}(t_{k+1}^i) - \bar{m}(t_k^i)$ is the number of times that the RM has updated within $[t_k^i, t_{k+1}^i)$, i.e., between two consecutive updates of application $i \in \mathcal{I}$.

Throughout this section, we also admit the *design assumption* that $1 \leq N_i(k) \leq \bar{N}$ for all $i \in \mathcal{I}$ and $k = 0, 1, \dots$, and for some $\bar{N} \in \mathbb{N}$. In other words, (a) the frequency at which the RM updates is larger than the frequency of every application i , and (b) each application i

updates with positive frequency. Without loss of generality, we will also assume that $t_0 = 0$, i.e., the RM starts updating first.

Given the above notation, the update recursion of each application i can be written as follows:

$$s_i(t_{k+1}^i) = \Pi_{S_i}[s_i(t_k^i) + \epsilon Y_i(t_{\bar{m}(t_k^i)}^i)], \quad (11)$$

since the computation of the observation signal $Y_i(\cdot)$ of each application i is based upon the most recent performance measurements taken by the RM at $\bar{m}(t_k^i)$.

Theorem 4.2 (Asynchronous convergence) *Let application i 's actual recursion be given by (11), with*

$$Y_i(t_k^i) = N_i(k) \cdot Y_i'(t_k^i), \quad (12)$$

for each $i \in \mathcal{I}$, where $Y_i'(\cdot) \doteq f_i(\cdot)$. Then, the conclusions of Theorem 4.1 continue to hold.

Proof. The proof will be shown in Section 5.2. \square

5 Technical Derivation

In this section, the technical proofs of Theorems 4.1–4.2 are presented through a series of propositions.

5.1 Synchronous convergence (Theorem 4.1)

The proof of Theorem 4.1 is an immediate implication of the following steps: (a) derivation of the *ODE approximation* of the adjustment dynamics (6)–(7), provided by Proposition 5.1, (b) characterization and stability analysis of its stationary points, provided by Propositions 5.2–5.3.

5.1.1 ODE approximation

We begin by establishing a connection between the asymptotic properties of the recursions (6)–(7) with the locally asymptotically stable sets of the ODE (8).

Proposition 5.1 (Synchronous ODE) *Consider the recursions (6)–(7), according to which both the RM and the applications update synchronously at fixed time instances t_k , $k = 1, 2, \dots$. If A is a locally asymptotically stable set in the sense of Lyapunov² for the ODE (8), then, for any initial condition $(\mathbf{s}(t_0), \bar{\mathbf{v}}(t_0))$ in the domain of attraction of A , $(\mathbf{s}(t_k), \bar{\mathbf{v}}(t_k)) \rightarrow A$ as $k \rightarrow \infty$ and $\epsilon \rightarrow 0$.*

² See [14, Definition 3.1].

Proof. The observation signal of the overall recursion is uniformly bounded, and the vector field of the ODE (8) is a continuous function on its domain. By [12, Theorem 1.1] (which shows convergence of Euler's method), we have that for every $\tau > 0$:

$$\lim_{\epsilon \rightarrow 0} \sup_{k=0, \dots, \lfloor \tau/\epsilon \rfloor} \|(\mathbf{s}_\epsilon(t_k), \bar{\mathbf{v}}_\epsilon(t_k)) - (\mathbf{s}^{\tau_0}(\tau_k), \bar{\mathbf{v}}^{\tau_0}(\tau_k))\| = 0,$$

where $(\mathbf{s}^{\tau_0}(\tau_0), \bar{\mathbf{v}}^{\tau_0}(\tau_0)) = (\mathbf{s}(t_0), \bar{\mathbf{v}}(t_0))$ and $\tau_k \doteq \epsilon k$. Given that A is locally asymptotically stable and the initial condition $(\mathbf{s}(t_0), \bar{\mathbf{v}}(t_0))$ lies within the region of attraction of A , the conclusion follows in a straightforward manner by the fact that any solution of the ODE (8) with initial condition $(\mathbf{s}(t_0), \bar{\mathbf{v}}(t_0))$ converges to A . \square

5.1.2 Stationary points

In this section, we characterize the stationary points of the ODE (8). In general, if an allocation $(\mathbf{s}^*, \bar{\mathbf{v}}^*)$ exists such that $\phi_i(s_i^*, \kappa \bar{v}_i^*) \equiv 0$ for all i , then such allocation will be a stationary point of the ODE (8) and a *fair allocation* according to Definition 2.1. In situations though where the CPU is overloaded, there might not be such allocations. The following proposition provides a characterization of the stationary points in such cases.

Proposition 5.2 (Stationary points) *Under either (H1) or (H2), the ODE (8) exhibits stationary points satisfying (9). Furthermore, as either (i) $\beta_i/\underline{s}_i \rightarrow 0$ or (ii) $n \rightarrow \infty$, any stationary point satisfies (10).*

Proof. If hypothesis (H1) is satisfied, then $\phi_i(s_i, \kappa \bar{v}_i) < 0$ for all $(\mathbf{s}, \bar{\mathbf{v}})$ and $\tau \geq 0$. In this case, any stationary point $(\mathbf{s}^*, \bar{\mathbf{v}}^*)$ satisfies (9) which equivalently implies that: $s_i^* \equiv \underline{s}_i$ and

$$\bar{v}_i^* = \min \left\{ \frac{1}{\kappa}, \frac{\lambda_i \phi_i(s_i^*, \kappa \bar{v}_i^*)}{\sum_{j \in \mathcal{I}} \lambda_j \phi_j(s_j^*, \kappa \bar{v}_j^*)} \right\}. \quad (13)$$

The mapping defined by the second expression of \bar{v}_i^* is well defined since $\phi_i(s_i^*, \kappa \bar{v}_i^*) < 0$ for all i . Furthermore, according to Brouwer's fixed point theorem [4, Corollary 6.6], it exhibits at least one fixed point since it is a continuous mapping defined on a compact set. The possibility that $\bar{v}_i^* = 0$ is excluded by Proposition 4.2. Finally, under hypothesis (H1), if we take $\beta_i/\underline{s}_i \rightarrow 0$ for all i , then $\phi_i(s_i, v_i) \rightarrow -1$ for all i , which further implies property (10).

If, instead, hypothesis (H2) is satisfied, then, by Proposition 4.3, there exists a finite k^* , such that, $\bar{v}_i(t_k) \leq \zeta$ for all $k > k^*$. By convergence of Euler's method, $\phi_i(s_i^{\tau_0}(\tau), \kappa \bar{v}_i^{\tau_0}(\tau)) < 0$ for all $\tau \geq \tau_{k^*} \doteq \epsilon k^*$ and all i . Thus, the fixed-point property (13) also applies. Furthermore, if $n \rightarrow \infty$, then by Proposition (4.3), $\zeta \rightarrow 0$,

and $\phi_i(s_i^{\tau_0}(\tau), \kappa \bar{v}_i^{\tau_0}(\tau)) \rightarrow -1$ uniformly on $\tau \geq \tau_{k^*}$ and i , which implies (10). \square

5.1.3 Global Asymptotic Stability (GAS)

The following proposition provides a strong convergence property of the ODE (8).

Proposition 5.3 (GAS) *If either (i) $\beta_i/s_i \rightarrow 0$ or (ii) $n \rightarrow \infty$, then the unique stationary point of the ODE (8), satisfying property (10), is globally asymptotically stable in the sense of Lyapunov.*

Proof. See Appendix D. \square

5.2 Asynchronous convergence (Theorem 4.2)

The proof of Theorem 4.2 is a direct implication of establishing *equivalence* between the synchronous and asynchronous update recursions satisfying property (12). In particular, we define *equivalence* between two (deterministic) update recursions as follows.

Definition 5.1 (Equivalent updates) *Two update recursions of the form (11) and observation signals $\{Y_i(t_k^i)\}$ and $\{Y_i'(t_k^i)\}$, $i \in \mathcal{I}$, are equivalent if the corresponding linear-time interpolations of the updated variables, $s_{i,\epsilon}(\cdot)$ and $s_{i,\epsilon}'(\cdot)$, respectively, satisfy*

$$\limsup_{\epsilon \rightarrow 0} \limsup_{t \geq 0} |s_{i,\epsilon}(t) - s_{i,\epsilon}'(t)| = 0.$$

In other words, two deterministic update recursions of the form (11) are considered equivalent when the corresponding linear-time interpolations approach each other uniformly in time as ϵ approaches zero.

We introduce the following *fictitious* recursion for each i ,

$$s_i'(t_{m+1}) = \Pi_{S_i} [s_i'(t_m) + \epsilon Y_i'(t_{\psi_i(m)})], \quad (14)$$

for all $m \geq 0$. This update is synchronized with the time index of the RM and $Y_i'(\cdot) \doteq f_i(\cdot)$. Note that the fictitious observation signals are defined at times $\{t_{\psi_i(m)}\}$, i.e., at the most recent update of the RM prior to the last update of application i . Since the RM starts updating first, we also set $Y_i'(t_{\psi_i(m)}) \equiv 0$ for all m such that $\psi_i(m) = 0$, since no performance measurements are available at time $t = 0$.

The following proposition shows that, if we pick appropriately the observation signals of the (actual) asynchronous update (11), then the asynchronous update becomes equivalent with the synchronous update of (7).

Proposition 5.4 (Equivalence) *For each application $i \in \mathcal{I}$, assume that its actual update recursion is given by (11), where $Y_i(t_k^i)$ satisfies property (12). Then, the following statements hold:*

- (1) *The fictitious synchronous update (14) is equivalent with the asynchronous update (11).*
- (2) *The fictitious synchronous update (14) is equivalent with the synchronous update (7).*

Proof. See Appendix E. \square

6 Experimental Evaluation

6.1 Simulation framework

To test the assignment of virtual platforms v_i and service levels s_i , the resource management framework was implemented in TrueTime [5]. TrueTime is a MATLAB/Simulink-based tool, embedded within Simulink, that allows the simulation of tasks executing within real-time kernels. TrueTime implements virtual processors through the Constant Bandwidth Server (CBS) [1]. Also, it is possible to adjust the CPU time allocated to the running applications (the bandwidth v_i), as needed by our RM. Moreover, TrueTime offers the ability to simulate memory management and protection, therefore being a perfect match to simulate our resource management framework.


A TrueTime kernel simulates a single CPU that hosts the execution of the RM and the CBS servers (virtual processors) on top of which the applications are running. The RM observes the matching function, f_i , of each application i and computes the new reservation v_i according to (6). *Observe* in this case means that the RM is able to read the start and finishing time of each job and it computes the matching function according to (2). Then, it updates the virtual platforms and communicates to the applications the observations $Y_i(t_k^i) \equiv f_i(t_k^i)$ according to (12).

It is here assumed that applications are composed of some time sensitive portions of code, called *jobs*. For example, in a media encoder/decoder a job may be the encoding/decoding of an MPEG frame. Applications are requested to inform the RM about the desired duration of each job. Below we report a template of the job code. To ease the presentation, we omit some implementation details, which can be found in [16].


```

1 | /* j is the job index*/
2 | id = signal_job_start(j);
3 | adjust = get_performance(j);
4 | /* body of the job. If service aware, it should
5 |    modify its resource requirement by adjust */
6 | do_work(/* parameters */);
7 | signal_job_end(j);

```



(a) RM deals with asynchronicity.



(b) RM does not deal with asynchronicity.

Fig. 4. TrueTime simulation of three applications that asynchronously update their service levels.




Fig. 3. TrueTime simulation of five applications with different weights that update their service levels.

As the application runs, it is asked to signal the start and the end of a job. This signaling actions are per-

formed by invoking respectively `signal_job_start` and `signal_job_end`, providing as parameter the index j of the job. Within the job, the first action is the invocation of the function `get_performance`. This function, which is computed by the monitoring infrastructure, returns a measurement of the service level adjustment required to achieve a perfect matching between the service level and the virtual platform. Jobs are assumed to be periodic.

To simulate service-aware applications, we developed a synthetic test application, which performs some computation depending on the service level s_i . All jobs of the application have deadline D_i and are executed in a forever loop. The execution requirement of each job is a linear function $a_i s_i + b_i$ of the service level. Hence, applications with a large a_i are more service-sensitive than applications with a_i close to zero. All applications parameters (D_i , ϵ , a_i , b_i), which determine the application behavior and its capacity to adapt, are set at start time. This enables, for example, the coexistence of fully service-aware applications together with service-unaware ones (with $a_i = 0$).

6.2 Synchronous updates

In the first scenario, we consider five applications running in a single core and updating synchronously with the RM according to (7). We pick $\lambda_1 = 0.9$, $\lambda_2 = 0.7$, $\lambda_3 = 0.5$, $\lambda_4 = 0.3$ and $\lambda_5 = 0.1$. Each application i has $a_i = 20$, $b_i = 200$ and job deadline $D_i = 1$ msec.

The initial service levels for the five applications is set to 10 and the applications update their service levels every time they perform a new job. Finally, we restrict the maximum assignable bandwidth by the RM to 90% to simulate the fact that the operating system should have some space to execute on the same hardware.

Figure 3 reports the quantities measured during the experiment. All applications are gradually reducing their service levels as expected due to the negative matching function f_i . According to Theorem 4.1, app_1 should receive a larger virtual platform compared to the rest of the applications due to its larger weight. The final allocation of virtual platforms may not correspond exactly to the values assigned by condition (10), since the conditions partially hold at the beginning of the simulation when f_i are significantly less than 0. Note though that the relative importance of the applications is preserved due to the synchronous updates.

6.3 Asynchronous updates

In this scenario, we investigate the effect of the asynchronicity in the final allocation of virtual platforms. We consider three different applications starting at the same time, with weights $\lambda_1 = 0.1$, $\lambda_2 = 0.5$ and $\lambda_3 = 0.8$. Each application uses resources according to $a_i = 40$, $b_i = 100$. Also they all have a job deadline $D_i = 10$ msec. Application app_1 updates its service level after completing 10 jobs, while app_2 and app_3 update their service levels after completing 1 job, i.e., asynchronous updates are introduced. The initial service levels of the three applications are set equal to 10 and the upper and lower bounds on the service levels are set equal to 0 and 20, respectively.

Figure 4a reports the quantities measured during the experiment, when applications apply the scheme prescribed in the hypotheses of Theorem 4.2 for asynchronous updates. Each application employs a constant step-size sequence of $\epsilon = 0.03$. It can be noticed that the RM is able to maintain a virtual platform allocation that is consistent with the weights of the applications, while driving all matching functions to zero.

To strengthen the motivation for asynchronicity management, Figure 4b shows the same simulation when applications do not employ the adjusted observation signal introduced in Equation (12), and instead they employ the originally introduced observation signal of the synchronous scheme (7). Due to the slow update rate of app_1 , it maintains a high service level for longer period, which subsequently leads to maintaining a smaller matching function f_1 for longer period. Thus, even though the weight of app_1 is smaller compared to the rest of applications, the RM favors this application significantly by gradually providing more resources. This simulation demonstrates that the original scheme of Section 3 may not be fair to applications

under asynchronous updates, and it was the main motivation for the development of the updated dynamics of Theorem 4.2.

7 Conclusions

We proposed a distributed management framework for allocating CPU resources to time-sensitive applications. Given that future computing systems will have to accommodate large number of applications of different demand levels, resource allocation should not be independent of the applications' performance (a notion captured through the matching function in this paper). Furthermore, given that resources are always finite, applications with higher flexibility in adjusting their demands (or service level) should decrease their service levels in overload cases. This paper proposed a distributed scheme that incorporates both these two elements, i.e., both measurements of the applications' performance, and applications' service-level adjustment.

In this paper, service-level adjustment is performed by *prescribing* certain dynamics to the applications. Such prescribed response dynamics was the first step towards the development of a fully distributed allocation scheme. In a fully distributed setup, where applications are not prescribed the response dynamics, the question is whether a RM can still be designed that guarantees fair allocation of resources independently of the type of applications and their adjustment dynamics.

A Proof of Proposition 4.1

Let us first consider the unconstrained version of the adjustment dynamics (6), i.e., take $\kappa = 1$. (We will revisit this assumption later.) In this case, the sum of the normalized virtual platforms can be expressed as:

$$\sum_{i=1}^n \bar{v}_i(t_{k+1}) - 1 = \left(\sum_{i=1}^n \bar{v}_i(t_k) - 1 \right) \left(1 + \epsilon \sum_{j=1}^n \lambda_j [f_j(t_k)]_- \right).$$

Given that $-1 \leq [f_i(t_k)]_- \leq 0$, for sufficiently small $\epsilon = \epsilon(n) > 0$, the second term of the right-hand side is positive for all $k = 1, 2, \dots$. If $\sum_{i=1}^n \bar{v}_i(t_k) \leq 1$, then $\sum_{i=1}^n \bar{v}_i(t_{k+1}) \leq 1$. Thus, for a bounded number of applications, we may pick sufficiently small $\epsilon = \epsilon(n)$ such that $\bar{v}_i(t_k) \in [0, 1]$ and $\sum_{i=1}^n \bar{v}_i(t_k) \leq 1$ for all $k = 1, 2, \dots$

We consider now the constrained version of the recursion (6). For some time t , let us assume that $\sum_{i=1}^n \bar{v}_i(t_k) \leq 1$, i.e., the allocation is feasible at time t_k . When we update this allocation using (6), the projection operator is activated only if $\bar{v}_i(t_k) + \epsilon F_i(t_k) > 1/\kappa$ for some i . Given

that $\sum_{i=1}^n \bar{v}_i(t_k) + \epsilon F_i(t_k) \leq 1$ (as we showed for the unconstrained dynamics), this quantity may only reduce after applying the projection operator. Thus, feasibility is also preserved under the constrained recursion.

B Proof of Proposition 4.2

First, note that $\sum_{j=1}^n \lambda_j [f_j(t_k)]_- \geq -n$, for all $k = 0, 1, \dots$. According to (6), the incremental difference of \bar{v}_i for the unconstrained dynamics, satisfies:

$$\Delta \bar{v}_i(k) \doteq \bar{v}_i(t_{k+1}) - \bar{v}_i(t_k) \geq \epsilon (-\lambda_i [f_i(t_k)]_- - n \bar{v}_i(t_k)) \quad (\text{B.1})$$

Define the set $\Gamma_i(\epsilon) \doteq \{\bar{v}_i \in \mathcal{V}_i : \bar{v}_i \in (\epsilon, (L+1)\epsilon]\}$, and pick ϵ sufficiently small such that $\epsilon < 1/(L+1)\kappa$. Since $\sup_{k \in \mathbb{N}} |F_i(t_k)| \leq L$, in order for $\bar{v}_i(t_k)$ to approach zero, there should be a time k^* at which $\bar{v}_i(t_{k^*}) \in \Gamma_i(\epsilon)$. Assuming that $\bar{v}_i(t_k) \in \Gamma_i(\epsilon)$, if we pick ϵ sufficiently small, we have $[f_i(t_k)]_- \leq \beta_i \kappa (L+1)\epsilon / \underline{s}_i - 1$. Thus, the right-hand side of (B.1) further satisfies:

$$\epsilon (-\lambda_i [f_i(t_k)]_- - n \bar{v}_i(t_k)) \geq \epsilon \lambda - (\beta_i \kappa / \underline{s}_i + n)(L+1)\epsilon^2.$$

For a given number of applications n , there exists $\epsilon^* = \epsilon^*(n) < 1/(L+1)\kappa$ with $\epsilon^*(n) \rightarrow 0$ as $n \rightarrow \infty$, such that, if $\epsilon < \epsilon^*$, then the above quantity is strictly positive, i.e., $\Delta \bar{v}_i(k) > 0$. Since under the unconstrained dynamics the lower bound of the projection operator is not reached, the same will also hold for the constrained dynamics. From (B.1), we conclude that if $\bar{v}_i(t_0) > \epsilon$ for all i , then $\inf_{k \in \mathbb{N}} \bar{v}_i(t_k) \geq \epsilon \forall i$.

C Proof of Proposition 4.3

At time instance k , let $\mathcal{I}' \subseteq \mathcal{I}$ be the set of applications with resources greater than ζ , i.e., $\mathcal{I}' \doteq \{i \in \mathcal{I} : \bar{v}_i(t_k) > \zeta\}$. Pick $0 < \zeta \leq 1/\kappa$ such that $\gamma^* \doteq \max_{\mathcal{I} \setminus \mathcal{I}'} \{\beta_i \kappa \zeta / \underline{s}_i - 1\} < 0$, i.e., all applications in $\mathcal{I} \setminus \mathcal{I}'$ have a negative matching function. Pick also $\epsilon < \zeta/L$.

(1) For any $i \in \mathcal{I}'$, the incremental difference of \bar{v}_i at k is defined as $\Delta \bar{v}_i(k) \doteq \bar{v}_i(t_{k+1}) - \bar{v}_i(t_k) = \epsilon F_i(t_k)$, assuming that the projection operator in (6) is not activated. Note that for all $j \in \mathcal{I} \setminus \mathcal{I}'$, $\bar{v}_j(t_k) \leq \zeta$ and

$$\sum_{j \in \mathcal{I} \setminus \mathcal{I}'} \lambda_j [f_j(t_k)]_- \leq \left(\sum_{j \in \mathcal{I} \setminus \mathcal{I}'} \lambda_j \right) \gamma^* \leq |\mathcal{I} \setminus \mathcal{I}'| \gamma^* \lambda.$$

Hence, according to the definition of $F_i(t_k)$, we have:

$$\Delta \bar{v}_i(k) \leq -\epsilon(1 - \bar{v}_i(t_k)) \lambda_i [f_i(t_k)]_- + \epsilon |\mathcal{I} \setminus \mathcal{I}'| \gamma^* \bar{v}_i(t_k) \lambda \leq \epsilon(1 - \zeta) + \epsilon(n - \lfloor (1 - \zeta)/\zeta \rfloor) \zeta \gamma^* \lambda,$$

where the last inequality results from the fact that $-\lambda_i [f_i(t_k)]_- \leq 1$, $\bar{v}_i(t_k) > \zeta$, $1 - \bar{v}_i(t_k) < 1 - \zeta$ and

$|\mathcal{I} \setminus \mathcal{I}'| \geq n - \lfloor (1 - \zeta)/\zeta \rfloor$. For any

$$n \geq n_1^*(\zeta) \doteq \left\lceil \left\lfloor \frac{(1 - \zeta)}{\zeta} \right\rfloor + \frac{-2 + \zeta}{\zeta \gamma^* \lambda} \right\rceil,$$

we have $-\zeta < -\epsilon L \leq \Delta \bar{v}_i(k) \leq -\epsilon < 0$. In this case, the initial assumption that the projection operator in (6) is not activated is also valid. Furthermore, according to [18, Theorem 5.1], the process $\bar{v}_i(t_k)$ will enter $[0, \zeta]$ within finite k .

(2) For any application $i \in \mathcal{I} \setminus \mathcal{I}'$, the unconstrained incremental difference $\Delta \bar{v}_i(k) \doteq \epsilon F_i(t_k)$ at time k satisfies:

$$\Delta \bar{v}_i(k) \leq \epsilon(1 - \bar{v}_i(t_k)) + \epsilon \bar{v}_i(t_k) \gamma^* \lambda (n - \lfloor (1 - \zeta)/\zeta \rfloor - 1),$$

since $-\lambda_i [f_i(t_k)]_- \leq 1$,

$$\sum_{j \neq i} \lambda_j [f_j(t_k)]_- \leq \sum_{j \in \mathcal{I} \setminus \mathcal{I}' \setminus i} \lambda [f_j(t_k)]_- \leq \gamma^* \lambda |\mathcal{I} \setminus \mathcal{I}' \setminus i|,$$

and $|\mathcal{I} \setminus \mathcal{I}' \setminus i| \geq (n - \lfloor (1 - \zeta)/\zeta \rfloor - 1)$. In order for the process $\bar{v}_i(t_k)$ to exit the set $[0, \zeta]$, there should be a time instance k^* at which $\bar{v}_i(t_{k^*}) \in (\zeta - \epsilon L, \zeta]$. For any $\bar{v}_i(t_k) \in (\zeta - \epsilon L, \zeta]$, we have:

$$\Delta \bar{v}_i(k) \leq \epsilon(1 - \zeta + \epsilon L) + \epsilon(\zeta - \epsilon L) \gamma^* \lambda (n - \lfloor (1 - \zeta)/\zeta \rfloor - 1).$$

If the number of applications satisfy:

$$n \geq n_2^*(\zeta) \doteq \left\lceil 1 + \left\lfloor \frac{1 - \zeta}{\zeta} \right\rfloor + \frac{-2 + \zeta - \epsilon L}{(\zeta - \epsilon L) \gamma^* \lambda} \right\rceil,$$

then, $\Delta \bar{v}_i(k) \leq -\epsilon < 0$, which implies that the unconstrained process $\{\bar{v}_i(t_k)\}$ will not exit $[0, \zeta]$ for all future times. The same will hold for the constrained process.

Finally, by defining $n^*(\zeta) \doteq \max\{n_1^*, n_2^*\}$, both statements (1) and (2) will hold for any $n \geq n^*$. Note that $n^* \rightarrow \infty$ and $\zeta n^* \rightarrow c$ as $\zeta \rightarrow 0$, for some $c > 0$.

D Proof of Proposition 5.3

Let $(\mathbf{s}^*, \bar{\mathbf{v}}^*)$ be a stationary point of the ODE (8), where by property (9) satisfies $s_i^* = \underline{s}_i$. Define the function $W(\mathbf{s}, \bar{\mathbf{v}}) \doteq 1/2(\bar{\mathbf{v}} - \bar{\mathbf{v}}^*)^T (\bar{\mathbf{v}} - \bar{\mathbf{v}}^*) \geq 0$. The derivative of W with respect to time τ satisfies:

$$\dot{W}(\mathbf{s}, \bar{\mathbf{v}}) = \sum_{i=1}^n (\bar{v}_i - \bar{v}_i^*)^T \Phi_i(\mathbf{s}, \kappa \bar{\mathbf{v}}). \quad (\text{D.1})$$

Note that at the stationary point and for every i , either one of the following holds: (a) $\bar{v}_i^* = 1/\kappa$ and $\Phi_i(\mathbf{s}^*, \kappa \bar{\mathbf{v}}^*) > 0$, or (b) $\Phi_i(\mathbf{s}^*, \kappa \bar{\mathbf{v}}^*) = 0$. Note that case (a) implies $(\bar{v}_i - \bar{v}_i^*)^T \Phi_i(\mathbf{s}, \kappa \bar{\mathbf{v}}) < 0$ for all $\bar{v}_i < 1/\kappa$ and all $s_i > \underline{s}_i$.

Thus, in case condition (a) is satisfied for some applications, the corresponding additive terms in (D.1) will be strictly negative. Without loss of generality, it suffices to investigate the derivative (D.1) when all applications satisfy condition (b). Furthermore, it suffices to consider the case where $s_i \equiv s_i^* = \underline{s}_i$, since under hypothesis (i) $\phi_i(s_i, v_i) < 0$ for all $s_i \in \mathcal{S}_i$, $v_i \in \mathcal{V}_i$ and i , and under hypothesis (ii) there exists a time τ^* after which $\phi_i(s_i(\tau), \kappa \bar{v}_i(\tau)) < 0$, for all $\tau \geq \tau^*$ (according to Proposition 4.3 and convergence of Euler's method shown in Proposition 5.1). Thus, for all $\tau \geq \tau^*$, we have:

$$\begin{aligned} \dot{W}(\mathbf{s}^*, \bar{\mathbf{v}}) &= \sum_{i=1}^n (\bar{v}_i - \bar{v}_i^*) (\Phi_i(\mathbf{s}^*, \kappa \bar{\mathbf{v}}) - \Phi_i(\mathbf{s}^*, \kappa \bar{\mathbf{v}}^*)) \\ &= - \left(\sum_{j=1}^n \lambda_j \right) \sum_{i=1}^n |\bar{v}_i - \bar{v}_i^*|^2 - \sum_{i=1}^n \frac{\lambda_i \beta_i \kappa}{s_i^*} |\bar{v}_i - \bar{v}_i^*|^2 \\ &\quad - \sum_{i=1}^n (\bar{v}_i - \bar{v}_i^*) \sum_{j=1}^n \frac{\lambda_j \beta_j \kappa}{s_j^*} (\bar{v}_i^* \bar{v}_j^* - \bar{v}_i \bar{v}_j). \end{aligned}$$

We denote by I_1 , I_2 and I_3 the three terms of the r.h.s. of the above expression, i.e., $\dot{W}(\mathbf{s}^*, \bar{\mathbf{v}}) \equiv I_1 + I_2 + I_3$. Note that: $I_1 = -(\sum_{j=1}^n \lambda_j) \|\bar{\mathbf{v}} - \bar{\mathbf{v}}^*\|_2^2$, $|I_2| \leq \max_{i \in \mathcal{I}} \{\beta_i \kappa / s_i^*\} \|\bar{\mathbf{v}} - \bar{\mathbf{v}}^*\|_2^2$ and

$$\begin{aligned} |I_3| &\leq \sum_{i=1}^n |\bar{v}_i - \bar{v}_i^*| \sum_{j=1}^n \frac{\lambda_j \beta_j \kappa}{s_j^*} |\bar{v}_i \bar{v}_j - \bar{v}_i^* \bar{v}_j^*| \\ &\leq \sum_{i=1}^n |\bar{v}_i - \bar{v}_i^*| \sum_{j=1}^n \frac{\lambda_j \beta_j \kappa}{s_j^*} (|\bar{v}_j - \bar{v}_j^*| \bar{v}_i + |\bar{v}_i - \bar{v}_i^*| \bar{v}_j^*) \\ &\leq \sup_{i, \tau \geq \tau^*} \{\bar{v}_i\} \max_{j \in \mathcal{I}} \left\{ \frac{\lambda_j \beta_j \kappa}{s_j^*} \right\} \|\bar{\mathbf{v}} - \bar{\mathbf{v}}^*\|_1^2 + \\ &\quad \max_{j \in \mathcal{I}} \left\{ \frac{\lambda_j \beta_j \kappa}{s_j^*} \right\} \sum_{j=1}^n \bar{v}_j^* \|\bar{\mathbf{v}} - \bar{\mathbf{v}}^*\|_2^2. \end{aligned}$$

Given that $\|\bar{\mathbf{v}} - \bar{\mathbf{v}}^*\|_1 \leq \sqrt{n} \|\bar{\mathbf{v}} - \bar{\mathbf{v}}^*\|_2$, $\sum_{j=1}^n \bar{v}_j^* \leq 1$ and $\lambda_i \leq 1$ for all i , we have

$$|I_3| \leq \max_{j \in \mathcal{I}} \left\{ \frac{\beta_j \kappa}{s_j^*} \right\} \left(n \sup_{i, \tau \geq \tau^*} \{\bar{v}_i\} + 1 \right) \|\bar{\mathbf{v}} - \bar{\mathbf{v}}^*\|_2^2.$$

Under hypothesis (i), i.e., as $\beta_i / \underline{s}_i \rightarrow 0$ for all i , and for some fixed size of applications n , the first term, I_1 , dominates in size the term $I_2 + I_3$ uniformly in time. Since $I_1 < 0$ for any $\bar{\mathbf{v}} \neq \bar{\mathbf{v}}^*$, we have that $\dot{W}(\mathbf{s}^*, \bar{\mathbf{v}}) < 0$ for any $\bar{\mathbf{v}} \neq \bar{\mathbf{v}}^*$. Under hypothesis (ii), i.e., as $n \rightarrow \infty$, Proposition 4.3 implies that $n \sup_{i, \tau \geq \tau^*} \{\bar{v}_i(\tau)\} \leq \zeta n(\zeta)$ approaches a positive constant. Thus, the first term I_1 dominates in size the term $I_2 + I_3$ when $n \rightarrow \infty$. In this case, $\dot{W}(\mathbf{s}^*, \bar{\mathbf{v}}) < 0$ for any $\bar{\mathbf{v}} \neq \bar{\mathbf{v}}^*$ such that $\bar{v}_i \leq \zeta$ for all i . Thus, under either (i) or (ii), and by [14, Theorem 3.2], we conclude that the unique stationary point $(\mathbf{s}^*, \bar{\mathbf{v}}^*)$,

satisfying (10), is globally asymptotically stable.

E Proof of Proposition 5.4

(1) Let us first consider the unconstrained versions of the actual (11) and fictitious update (14). In this case, the corresponding linear-time interpolations satisfy for any run time $t > 0$:

$$\begin{aligned} s_{i,\epsilon}(t) - s'_{i,\epsilon}(t) &= s_i(t_{\bar{k}(t,i)}^i) - s'_i(t_{\bar{m}(t)}^i) \\ &= \sum_{k=0}^{\bar{k}(t,i)-1} \epsilon Y_i(t_{\bar{m}(t_k^i)}^i) - \sum_{m=0}^{\bar{m}(t)-1} \epsilon Y'_i(t_{\psi_i(m)}) \\ &= \sum_{k=0}^{\bar{k}(t,i)-1} \epsilon Y_i(t_{\bar{m}(t_k^i)}^i) - \sum_{m=\bar{m}(t_{\bar{k}(t,i)-1}^i)+1}^{\bar{m}(t)-1} \epsilon Y'_i(t_{\psi_i(m)}) \\ &\quad - \sum_{k=0}^{\bar{k}(t,i)-2} \epsilon N_i(k) Y'_i(t_{\psi_i(\bar{m}(t_{k+1}^i))}), \end{aligned}$$

where the third summation of the r.h.s. summarizes all the observation terms of the RM up to time index $\bar{m}(t_{\bar{k}(t,i)-1}^i)$, and the second summation summarizes all the remaining terms, i.e., from $\bar{m}(t_{\bar{k}(t,i)-1}^i) + 1$ until $\bar{m}(t) - 1$. For example, in Figure 2, the third summation corresponds to all time indices up to $\bar{m}(t_{\bar{k}-1}^i)$, while the second summation corresponds to the remaining terms up to time $\bar{m}(t) - 1$. Since $1 \leq N_i(k) \leq \bar{N}$, the latter terms may be zero to maximum \bar{N} in numbers. Also, $\psi_i(\bar{m}(t_{k+1}^i)) = \bar{m}(t_k^i)$. Thus, we have:

$$\begin{aligned} s_{i,\epsilon}(t) - s'_{i,\epsilon}(t) &= \sum_{k=0}^{\bar{k}(t,i)-2} \epsilon \left(Y_i(t_{\bar{m}(t_k^i)}^i) - N_i(k) Y'_i(t_{\bar{m}(t_k^i)}^i) \right) + \\ &\quad \epsilon Y_i(t_{\bar{m}(t_{\bar{k}(t,i)-1}^i)}^i) - \sum_{m=\bar{m}(t_{\bar{k}(t,i)-1}^i)+1}^{\bar{m}(t)-1} \epsilon Y'_i(t_{\psi_i(m)}). \end{aligned}$$

From the last expression, we observe that if $Y_i(t_k^i) = N_i(k) Y'_i(t_k^i)$, then, the first term in the r.h.s. becomes identically zero. Given that $\sup_{i, t \geq 0} |Y'_i(t)| \leq \ell$, for some $\ell > 0$, we also have $\sup_{i, t \geq 0} |Y_i(t)| \leq \ell \bar{N}$. Hence,

$$\begin{aligned} |s_{i,\epsilon}(t) - s'_{i,\epsilon}(t)| &\leq \ell \bar{N} + \epsilon \left(\bar{m}(t) - \bar{m}(t_{\bar{k}(t,i)-1}^i) - 1 \right) \ell \\ &\leq \ell (3\bar{N} - 1), \end{aligned}$$

which approaches zero as $\epsilon \rightarrow 0$ uniformly in time. Thus, we showed that $s_{i,\epsilon}(\cdot)$ and $s'_{i,\epsilon}(\cdot)$ are equivalent according to Definition 5.1. Since the corresponding projected versions are simply truncations to the set $\mathcal{S}_i \equiv [\underline{s}_i, \infty)$, the same conclusion applies for the projected versions.

(2) Let $\{s''_i(t_k^i)\}_k$ denote the service level recursion under the synchronous update (7) in order to distinguish it from the actual asynchronous one (11). Let us also denote $s'_{i,\epsilon}(\cdot)$ the corresponding linear-time interpolation. Similarly to the proof of part (1), it suffices to consider the unconstrained recursions. For every i and $t \geq 0$, and since $1 \leq N_i(k) \leq \bar{N}$, we have:

$$\begin{aligned} |s'_{i,\epsilon}(t) - s''_{i,\epsilon}(t)| &= |s'_i(t_{\psi_i(\bar{m}(t))}) - s''_i(t_{\bar{m}(t)})| \\ &\leq \sum_{k=\psi_i(\bar{m}(t))}^{\bar{m}(t)-1} \epsilon |Y'_i(t_k)| \\ &\leq \epsilon(\bar{m}(t) - \psi_i(\bar{m}(t)))\ell \leq \epsilon\bar{N}\ell, \end{aligned}$$

which implies equivalence of $s'_{i,\epsilon}(\cdot)$ and $s''_{i,\epsilon}(\cdot)$.

References

- [1] L. Abeni and G. Buttazzo. Integrating multimedia applications in hard real-time systems. In *Proc. of the 19th IEEE Real-Time Systems Symposium*, pages 4–13, Madrid, Spain, Dec. 1998.
- [2] K.-E. Årzén, V. Romero Segovia, S. Schorr, and G. Fohler. Adaptive resource management made real. In *Proc. 3rd Workshop on Adaptive and Reconfigurable Embedded Systems*, Chicago, IL, USA, Apr. 2011.
- [3] E. Bini, G. C. Buttazzo, J. Eker, S. Schorr, R. Guerra, G. Fohler, K.-E. Årzén, R. Vanessa, and C. Scordino. Resource management on multicore systems: The ACTORS approach. *IEEE Micro*, 31(3):72–81, 2011.
- [4] K. Border. *Fixed Point Theorems with Applications to Economics and Game Theory*. Cambridge University Press, 1985.
- [5] A. Cervin, D. Henriksson, B. Lincoln, J. Eker, and K.-E. Årzén. How does control timing affect performance? Analysis and simulation of timing using Jitterbug and TrueTime. *IEEE Control Systems Magazine*, 23(3):16–30, June 2003.
- [6] G. Chasparis, M. Maggio, K.-E. Årzén, and E. Bini. Distributed management of CPU resources for time-sensitive applications. In *Proc. of the 2013 American Control Conference*, 2013.
- [7] G. Chasparis and J. Shamma. Distributed dynamic reinforcement of efficient outcomes in multiagent coordination and network formation. *Dynamic Games and Applications*, 2(1):18–50, 2012.
- [8] T. Cucinotta, L. Palopoli, L. Abeni, D. Faggioli, and G. Lipari. On the integration of application level and resource level QoS control for real-time applications. *IEEE Transactions on Industrial Informatics*, 6(4):479–491, Nov. 2010.
- [9] M. Di Natale and A. Sangiovanni-Vincentelli. Moving from federated to integrated architectures in automotive: The role of standards, methods and tools. *Proc. of the IEEE*, 98(4):603–620, 2010.
- [10] J. Eker, P. Hagander, and K.-E. Årzén. A feedback scheduler for real-time controller tasks. *Control Engineering Practice*, 8(12):1369–1378, Jan. 2000.
- [11] D. Grosu and A. T. Chronopoulos. Noncooperative load balancing in distributed systems. *Journal of Parallel and Distributed Computing*, 65(9):1022–1034, 2005.
- [12] A. Iserles. *A First Course in the Numerical Analysis of Differential Equations*. Cambridge University Press, 2009.
- [13] B. Johansson, C. Adam, M. Johansson, and R. Stadler. Distributed resource allocation strategies for achieving quality of service in server clusters. In *Proc. of the 45th IEEE Conference on Decision and Control*, pages 1990–1995, Dec. 2006.
- [14] H. Khalil. *Nonlinear Systems*. Prentice-Hall, 1992.
- [15] C. Lee, J. P. Lehoczky, D. Sieworek, R. Rajkumar, and J. Hansen. A scalable solution to the multi-resource QoS problem. In *Proc. of the 20th IEEE Real-Time Systems Symposium*, pages 315–326, Phoenix, AZ, Dec. 1999.
- [16] M. Maggio, E. Bini, G. Chasparis, and K.-E. Årzén. A game-theoretic resource manager for RT applications. In *25th Euromicro Conference on Real-Time Systems, ECRTS13*, Paris, France, July 2013.
- [17] C. W. Mercer, S. Savage, and H. Tokuda. Processor capacity reserves: Operating system support for multimedia applications. In *Proc. of IEEE International Conference on Multimedia Computing and Systems*, pages 90–99, Boston, MA, U.S.A., May 1994.
- [18] M. B. Nevelson and R. Z. Hasminskii. *Stochastic Approximation and Recursive Estimation*. American Mathematical Society, Providence, RI, 1976.
- [19] R. Rajkumar, C. Lee, J. Lehoczky, and D. Siewiorek. A resource allocation model for QoS management. In *Proc. of the IEEE Real Time System Symposium*, 1997.
- [20] M. Sojka, P. Piša, D. Faggioli, T. Cucinotta, F. Checoni, Z. Hanzálek, and G. Lipari. Modular software architecture for flexible reservation mechanisms on heterogeneous resources. *Journal of Systems Architecture*, 57(4):366–382, 2011.
- [21] D. C. Steere, A. Goel, J. Gruenberg, D. McNamee, C. Pu, and J. Walpole. A feedback-driven proportion allocator for real-rate scheduling. In *Proc. of the 3rd Symposium on Operating Systems Design and Implementation*, Feb. 1999.
- [22] R. Subrata, A. Y. Zomaya, and B. Landfeldt. A cooperative game framework for QoS guided job allocation schemes in grids. *IEEE Transactions on Computers*, 57(10):1413–1422, Oct. 2008.
- [23] G. Wei, A. V. Vasilakos, Y. Zheng, and N. Xiong. A game-theoretic method of fair resource allocation for cloud computing services. *The Journal of Supercomputing*, 54(2):252–269, Nov. 2010.
- [24] J. Weibull. *Evolutionary Game Theory*. MIT Press, Cambridge, MA, 1997.