

On the Dominance of Minimum-Parallelism Multiprocessor Supply*

Kecheng Yang and James H. Anderson

Department of Computer Science, University of North Carolina at Chapel Hill

Abstract

Many approaches have been proposed to enable disparate real-time software components to share a physical multiprocessor platform by giving each component the “illusion” of executing on a dedicated virtual platform. Such an illusion is supported by specifying a supply interface that indicates how computation time is made available to a component over time. A number of approaches for defining such interfaces have been proposed: so many that sifting through them all can be confusing for the practitioner. In the case of soft real-time applications, one particular proposed interface—minimum-parallelism (MP) supply—has been shown to enable the co-scheduling of different components with no utilization loss. In the case of hard real-time applications, it follows from prior work that MP supply easily dominates other choices if the simplifying assumption is made that supply is allocated on different processors using a common, synchronized allocation period. The main contribution of this paper is to show that the dominance of MP supply is retained if this simplifying assumption is removed. This result suggests that MP supply should be the focus in future work on real-time multiprocessor virtualization.

1 Introduction

Open-systems [5] frameworks allow separate software components to execute together on a common hardware platform, with each component having the “illusion” of executing on a dedicated virtual platform. Providing such an illusion can ease software-development efforts, not only when mixing different applications, but also when integrating separately developed components of the same application. In domains where real-time constraints exist, *temporal isolation* among components should be ensured, *i.e.*, it should be possible to validate the timing constraints of each component independently. Therefore, a specification of the computing capacity allocated to a component is needed.

In early work in this direction pertaining to uniprocessor platforms, Shin and Lee [17] proposed a virtual processor (VP) model called the *periodic resource (PR)* model, which allows the considerable body of work on periodic task scheduling [13] to be exploited in reasoning about the allocation of processor time to components. In the PR model, a VP is specified by the parameters (Π, Θ) , with the interpretation that Θ time units of processor time is guaranteed to the supported component every Π time units.

While this simple model sufficed in the uniprocessor case, it is inadequate in the multiprocessor case, because

the important issue of *parallelism* is ignored. To deal with this issue, Shin *et al.* [16] proposed extending the PR model by adding an additional parameter. Specifically, under their *multiprocessor periodic resource (MPR)* model, the supply allocated to a component is specified by (Π, Θ, m') , with the interpretation that Θ time units of processor time is guaranteed to the component every Π time units with at most m' VPs providing allocation in parallel. That is, the new parameter m' specifies the *maximum degree of parallelism*. In the MPR model, all VPs allocated to a component are required to have a common period Π that is strictly synchronized.

A key characteristic of the MPR model is its flexibility. For example, consider a component that is to be allocated 80% of the capacity of a quad-core machine. The supply interface for that component could be defined as $(100, 320, 4)$, meaning that every 100 time units, the component receives 320 units of processing time on up to four processors. Such a specification does not indicate the precise manner in which processing time is allocated. For example, the component could be allocated 80% of the capacity of each processor, or 100% of three processors and 20% of the fourth, among other choices. Which choice is best?

MP form. In the example just discussed, the second-listed choice is known as *minimum-parallelism (MP) form*. Under MP form, each component is allocated at most one partially available processor, with all other processors allocated to it being fully available. MP form was first proposed by Leontyev and Anderson [9] to support soft real-time container hierarchies, which allow components to include sub-components, which in turn can include their own sub-components, *etc.* Assuming MP form, they showed that container hierarchies with an unlimited number of levels can be supported with bounded deadline tardiness and no utilization loss. In work directed at hard real-time systems, Xu *et al.* [18] observed that, by enforcing MP form in the context of the MPR model, per-component schedulability can be improved.

Because this improvement in schedulability was considered in the context of the MPR model, a common, synchronized allocation period was assumed to be used on all processors allocated to a component. In practice, however, situations exist in which such an assumption may be problematic. A good example of this can be seen in recent work of Durrieu *et al.* [6], who considered a flight management system implemented on a multicore platform wherein clocks on different processors “do not drift [but] have unpredictable initial offsets.” In the future, the assumption of tight synchrony may become even more problematic, as manycore platforms evolve in which core counts soar into the hundreds if not thousands. Similar observations have been made by Lipari and Bini [12] and Bini *et al.* [3], who suggested

*Work supported by NSF grants CPS 1239135, CNS 1409175, and CPS 1446631, AFOSR grant FA9550-14-1-0161, ARO grant W911NF-14-1-0499, and funding from General Motors.

	Common Period	Different Periods
Synchronous	$\Pi^* \leq \Pi_{min}$	$\Pi^* \leq \frac{\Pi_l(1-w_l)w_l}{2(1-w^*)w^*}$
Concrete Asynchronous	$\Pi^* \leq \frac{\Pi_l(1-w_l)w_l}{2(1-w^*)w^*}$	$\Pi^* \leq \frac{\Pi_l(1-w_l)w_l}{2(1-w^*)w^*}$
Non-Concrete Asynchronous	$\Pi^* \leq \Pi_{min}$	$\Pi^* \leq \frac{1}{2}\Pi_{min}$

Table 1: Period restrictions for achieving the dominance of MP form. The notation used is explained in later sections.

generalizing the MPR model so that the VPs allocated to a single component may have different periods with different initial phasings. Does MP form still retain its advantages over other supply forms in the hard-real-time case under this more general notion of VP allocation?

Contributions. In this paper, we answer this question in the affirmative by showing that MP form dominates all other supply forms in the context of these cases: VPs are synchronous, concrete asynchronous, or non-concrete asynchronous (these terms are defined in Sec. 2). In each of these cases, we consider two sub-cases: requiring a common period for all VPs, and allowing such periods to differ. The prior work noted above by Xu *et al.* [18] on the MPR model implies that MP form dominates all other forms in the case of synchronous VPs with a common period. For each other case, we show that an arbitrary component is always dominated by an MP-form component of the same bandwidth (*i.e.*, total processor capacity—see Sec. 2), provided its period is defined properly. The required periods are summarized in Table 1, and explained later. Additionally, in all six cases, we show that an MP-form component can never be dominated by a non-MP-form component of the same bandwidth, regardless of how periods are defined. As hinted at by the various expressions in Table 1, the issue of MP dominance under the considered cases is not as straightforward as one might think at first glance, and many subtleties arise.

Organization. In the following sections, we introduce our system model (Sec. 2), provide some preliminary properties and theorems (Sec. 3), show the dominance of MP form for non-concrete asynchronous VPs (Sec. 4) and synchronous and concrete asynchronous VPs (Sec. 5), show that MP form cannot be dominated by any other form (Sec. 6), discuss related work (Sec. 7), and conclude (Sec. 8).

2 System Model

We consider a compositional system executing upon a physical multiprocessor platform with identical processors. Each component is provided processor time by a set of VPs, each defined according to the PR model, as defined in Sec. 1.

2.1 Periodic Resource Model

Recall from Sec. 1 that under the PR model [17] a VP Γ_i is characterized by two parameters (Π_i, Θ_i) , which indicate that Γ_i supplies Θ_i units of processor time every Π_i time units, where $0 < \Theta_i \leq \Pi_i$. In this paper, we assume contin-




Figure 1: Worst-case supply of Γ_i (adapted from [17]).

uous time, thus Π_i and Θ_i are real numbers. The *bandwidth* of the VP Γ_i is given by $w_i = \Theta_i/\Pi_i$. Note that, for any Π_i , $\Gamma_i = (\Pi_i, \Pi_i)$ defines a VP corresponding to a dedicated physical processor that is always available.

The *supply bound function (SBF)* of the VP Γ_i , denoted $Z(t, \Gamma_i)$, indicates the *minimum* processor time Γ_i can supply during any time interval of length t . Shin and Lee [17] have shown that $Z(t, \Gamma_i)$ can be defined as

$$Z(t, \Gamma_i) = \begin{cases} 0 & \text{if } t'_{\Gamma_i} < 0 \\ \left\lfloor \frac{t'_{\Gamma_i}}{\Pi_i} \right\rfloor \cdot \Theta_i + \epsilon_{\Gamma_i} & \text{if } t'_{\Gamma_i} \geq 0 \end{cases} \quad (1)$$

where

$$t'_{\Gamma_i} = t - (\Pi_i - \Theta_i), \quad (2)$$

$$\epsilon_{\Gamma_i} = \max \left(t'_{\Gamma_i} - \Pi_i \left\lfloor \frac{t'_{\Gamma_i}}{\Pi_i} \right\rfloor - (\Pi_i - \Theta_i), 0 \right). \quad (3)$$

This definition reflects the worst-case scenario illustrated in Figure 1.

2.2 VPs in a Component

We consider a component \mathcal{C} that consists of a set of VPs, denoted $\mathcal{C} = \{\Gamma_i\}$, where $\Gamma_i = (\Pi_i, \Theta_i)$ for $1 \leq i \leq |\mathcal{C}|$. The supply of a component is the sum of the supply of all VPs in this component.

Since $\Gamma_i = (\Pi_i, \Pi_i)$ indicates a dedicated processor regardless of the value of Π_i , we let p denote the number of such dedicated processors and do not bother to specify their periods. Thus, we alternatively denote the component \mathcal{C} by $\mathcal{C} = (p, \mathcal{T})$, where $\mathcal{T} = \{\Gamma_i \mid \Gamma_i \in \mathcal{C} \wedge 0 < w_i < 1\}$. It is clear that

$$|\mathcal{C}| = p + |\mathcal{T}|. \quad (4)$$

We define the *bandwidth* of component \mathcal{C} as

$$\text{bw}(\mathcal{C}) = \sum_{\Gamma_i \in \mathcal{C}} w_i. \quad (5)$$

The bandwidth $\text{bw}(\mathcal{C})$ indicates the total processor share allocation to which \mathcal{C} is entitled. *Minimum-parallelism (MP)* form, mentioned in Sec. 1, is defined as follows.

Def. 1. A component $\mathcal{C} = (p, \mathcal{T})$ is in MP form if and only if $|\mathcal{T}| \leq 1$.

Concrete vs. non-Concrete. We consider the possibility that the VPs in a component are *asynchronous*, meaning that they can have different phases—a VP Γ_i with a *phase* of ϕ_i is initialized to begin at time ϕ_i , *i.e.*, its first allocation of Θ_i time units occurs within the interval $[\phi_i, \phi_i + \Pi_i)$, its second within $[\phi_i + \Pi_i, \phi_i + 2\Pi_i)$, and so on. As it turns out, the results we obtain depend on whether phases are *known* or *unknown* prior to runtime. In the first case, we say that the VPs are *concrete* asynchronous, and only a particular phase for each VP needs to be considered in schedulability (supply) analysis. In the second case, we say that the VPs are *non-concrete* asynchronous, and the worst case among all possible phases must be considered in schedulability (supply) analysis. *Synchronous* VPs can be considered as a special case of concrete asynchronous VPs where all phases are required to be zero. In this paper, we consider all of the three phasing assumptions regarding VPs: they can be synchronous, concrete asynchronous, or non-concrete asynchronous.

2.3 Parallel Supply Function

The SBF definition in (1) for the PR model hinges only on considering uniprocessor supply allocations. In the multiprocessor case, however, SBFs must also address the important issue of parallelism. Various multiprocessor SBFs have been proposed. The most expressive of these considered to date is the *parallel supply function (PSF)*, proposed by Bini *et al.* [2]. The PSF describes the supply of a component \mathcal{C} by a set of functions, $\{Y_j(t, \mathcal{C}) \mid j \in \mathbb{Z}^+\}$, where each function $Y_j(t, \mathcal{C})$ is defined as follows.

Def. 2. $Y_j(t, \mathcal{C})$ denotes the *minimum* supply of \mathcal{C} during any time interval of length t with a degree of parallelism at most j .

We illustrate the above definition with the following example, and refer readers to the work of Bini *et al.* [2] for a more formal treatment.

Ex. 1. (Adapted from [12].) Let Γ_1 , Γ_2 , and Γ_3 be three VPs that compose \mathcal{C} . Assume that the processor time they make available within the time interval $[0, 11)$ is shown in Figure 2, where the gray boxes represent available processor time. Suppose that all three VPs are fully available at or after time 11. Then, $[0, 11)$ is the interval of length 11 that provides the minimum supply at every degree of parallelism. In this case, $Y_1(t, \mathcal{C}) = 10$ because there are 10 time units in $[0, 11)$ during which at least one VP provides available processor time. $Y_2(t, \mathcal{C}) = 16$ because all three VPs provide available processor time simultaneously only in $[4, 5)$, so $Y_2(t, \mathcal{C})$ is one less than the total available processor time in $[0, 11)$. This total available time is given by $Y_3(t, \mathcal{C}) = 17$.

In this paper, we use PSF functions to describe *exact* lower bounds on supply in order to compare the supply of different components *exactly*. That is, for any j and $t \geq 0$, there exists a possible scenario in which, over some interval of length t , the supply provided by \mathcal{C} with a degree of parallelism at most j is exactly $Y_j(t, \mathcal{C})$.




Figure 2: Example illustrating parallel supply (adapted from [12]).

By Def. 2, we have the following property.

$$(\forall \mathcal{C}, \forall j \geq 1, \forall t \geq 0 :: Y_j(t, \mathcal{C}) \leq jt) \quad (6)$$

Also, By Lemma 1 in [2], the following properties hold.

$$(\forall \mathcal{C}, \forall j \geq 1, \forall t \geq 0 :: Y_j(t, \mathcal{C}) \leq Y_{j+1}(t, \mathcal{C})) \quad (7)$$

$$(\forall \mathcal{C}, \forall j \geq |\mathcal{C}|, \forall t \geq 0 :: Y_j(t, \mathcal{C}) = Y_{j+1}(t, \mathcal{C})) \quad (8)$$

In accordance with Def. 2, $Y_\infty(t, \mathcal{C})$ represents the minimum supply that \mathcal{C} is guaranteed to provide during any time interval of length t with *no constraint on the degree of parallelism*. By Def. 2, $Y_\infty(t, \mathcal{C}) = Y_{|\mathcal{C}|}(t, \mathcal{C})$, because there are at most $|\mathcal{C}|$ dedicated or non-dedicated resources that can provide supply in parallel in \mathcal{C} .

3 Preliminaries

In this section, we provide a condition for establishing the superiority of MP form. This condition will allow us to conclude that MP form dominates other forms. Dominance is defined with respect to component supply based on PSF:

Def. 3. A component \mathcal{C}' *dominates* another component \mathcal{C} if and only if $(\forall j \geq 1, \forall t \geq 0 :: Y_j(t, \mathcal{C}) \leq Y_j(t, \mathcal{C}'))$ holds.

By Def. 3, in order to show the dominance of an arbitrary component \mathcal{C}' over another arbitrary component \mathcal{C} , we must consider all relevant PSF functions. However, the following theorem shows that it suffices to consider only two specific PSF functions.

Theorem 1. Let \mathcal{C} be an arbitrary component, and let \mathcal{C}^* be a component in MP form. If $(\forall t :: Y_\infty(t, \mathcal{C}) \leq Y_\infty(t, \mathcal{C}^*))$ holds, then \mathcal{C}^* dominates \mathcal{C} .

Proof. Let $\mathcal{C} = (p, \mathcal{T})$ and $\mathcal{C}^* = (p^*, \mathcal{T}^*)$. Because \mathcal{C}^* has p^* dedicated processors,

$$(\forall 1 \leq j \leq p^*, \forall t \geq 0 :: Y_j(t, \mathcal{C}^*) = jt). \quad (9)$$

On the other hand, for \mathcal{C} , by (6), we have

$$(\forall 1 \leq j \leq p^*, \forall t \geq 0 :: Y_j(t, \mathcal{C}) \leq jt). \quad (10)$$

By (9) and (10),

$$(\forall 1 \leq j \leq p^*, \forall t \geq 0 :: Y_j(t, \mathcal{C}) \leq Y_j(t, \mathcal{C}^*)). \quad (11)$$

Because \mathcal{C}^* is in MP form, $|\mathcal{T}| \leq 1$, and by (4), $|\mathcal{C}^*| = p^* + |\mathcal{T}^*| \leq p^* + 1$. Therefore, by (8),

$$(\forall j \geq p^* + 1, \forall t \geq 0 :: Y_j(t, \mathcal{C}^*) = Y_\infty(t, \mathcal{C}^*)). \quad (12)$$




Figure 3: The graph of $Z(t, \Gamma_i)$, as an illustration of Properties 1, 2, and 3.

On the other hand, for \mathcal{C} , by (7),

$$(\forall j \geq p^* + 1, \forall t \geq 0 :: Y_j(t, \mathcal{C}) \leq Y_\infty(t, \mathcal{C})). \quad (13)$$

Now, by (12), (13), and $Y_\infty(t, \mathcal{C}) \leq Y_\infty(t, \mathcal{C}^*)$ (from the statement of the theorem), we have

$$(\forall j \geq p^* + 1, \forall t \geq 0 :: Y_j(t, \mathcal{C}) \leq Y_j(t, \mathcal{C}^*)). \quad (14)$$

By (11), (14), and Def. 3, \mathcal{C}^* dominates \mathcal{C} . \square

Before endeavoring to use Theorem 1 to establish the dominance of MP form, we first provide several useful properties concerning the supply function $Z(t, \Gamma_i)$ of an arbitrary VP Γ_i . Property 1 directly follows from the definition of $Z(t, \Gamma_i)$ as given by (1)–(3). Property 2 is established in Lemma 1 in [17], and Property 3 is established in [7]. The intuition behind these properties is illustrated by the graph of $Z(t, \Gamma_i)$ shown in Figure 3.

Property 1. $Z(t, \Gamma_i) = 0$ for $0 \leq t \leq 2(\Pi_i - \Theta_i)$.

Property 2. $Z(t, \Gamma_i) \geq \max\{(t - 2(\Pi_i - \Theta_i))w_i, 0\}$.

Property 3. $Z(t, \Gamma_i) \leq \max\{(t - (\Pi_i - \Theta_i))w_i, 0\}$.

We state two more properties below, in which an alternate definition of $Z(t, \Gamma_i)$ is indirectly considered that is based on the following function f :

$$f(x, \Gamma_i) = \left\lfloor \frac{x}{\Pi_i} \right\rfloor \cdot \Theta_i + \max\left(x - \Pi_i \left\lfloor \frac{x}{\Pi_i} \right\rfloor - (\Pi_i - \Theta_i), 0\right). \quad (15)$$

Note that, by (1) (2) and (3),

$$Z(t, \Gamma_i) = f(t'_{\Gamma_i}, \Gamma_i), \text{ if } t'_{\Gamma_i} \geq 0. \quad (16)$$

When Γ_i is fixed, i.e., Π_i and Θ_i are constants, the following properties apply to $f(x, \Gamma_i)$. These properties can be seen intuitively by considering the graph of $f(x, \Gamma_i)$, which is similar to that of $Z(t, \Gamma_i)$ as illustrated in Figure 3. Property 5 can be seen by observing that the slope of any two points in the graph of $f(x, \Gamma_i)$ is at most one.

Property 4. $f(x, \Gamma_i)$ is monotonically increasing for non-negative x , i.e., $f(x_1, \Gamma_i) \leq f(x_2, \Gamma_i)$ if $0 \leq x_1 \leq x_2$.




Figure 4: Illustration of Claim 1.

Property 5. For any $x, y \geq 0$, $f(x + y, \Gamma_i) \leq f(x, \Gamma_i) + y$, which also implies $f(x - y, \Gamma_i) \geq f(x, \Gamma_i) - y$, provided that $x - y \geq 0$ holds.

We also utilize the two straightforward claims below.

Claim 1. The supply of a VP Γ_i can be zero within any time interval of length $\Pi_i - \Theta_i$, regardless of how the interval aligns with the VP's periods of allocation.

This claim is different from Property 1. In order to have a supply of zero within a time interval of length up to $2(\Pi_i - \Theta_i)$, as stated in Property 1, the interval must have a specific alignment with respect to the periods of allocation of Γ_i as shown in Figure 1. However, according to this claim, the supply within *any* time interval of length $\Pi_i - \Theta_i$ can be a zero. Figure 4 shows the only two possibilities that can occur: the considered interval is either included within a single period of allocation, or spans two such periods. In either situation, supply within the interval can be zero.

Claim 2. Let $\mathcal{C}^* = (p^*, \mathcal{T}^*)$ be a component in MP form. If $|\mathcal{T}^*| = 0$, then $Y_\infty(t, \mathcal{C}^*) = t \cdot p^*$. If $|\mathcal{T}^*| = 1$, then letting Γ^* denote the lone VP in \mathcal{T}^* , $Y_\infty(t, \mathcal{C}^*) = t \cdot p^* + Z(t, \Gamma^*)$

This claim follows directly from the definitions above.

4 Non-Concrete Asynchronous

In this section, we consider the case of *non-concrete asynchronous* VPs. In order to apply Theorem 1 in this case to establish the dominance of MP form, we begin by providing an *exact* calculation of $Y_\infty(t, \mathcal{C})$.

For any time interval of length t , a dedicated resource supplies t time units of processor time, and by (1), a non-dedicated resource Γ supplies at least $Z(t, \Gamma)$ time units. Therefore, with the degree of parallelism unconstrained, a component $\mathcal{C} = (p, \mathcal{T})$ provides a supply of at least $tp + \sum_{\Gamma_i \in \mathcal{T}} Z(t, \Gamma_i)$. Moreover, this minimum does indeed happen, as shown in Figure 5. (Note that the alignment shown in the figure can happen because we are assuming for now that VPs are non-concrete asynchronous.) Thus, for any component $\mathcal{C} = (p, \mathcal{T})$,

$$Y_\infty(t, \mathcal{C}) = tp + \sum_{\Gamma_i \in \mathcal{T}} Z(t, \Gamma_i). \quad (17)$$

In the next two subsections, we establish the dominance of MP form in two steps. First, we consider the case in which all VPs in \mathcal{C} share a common period. Second, we build upon this result by considering the case in which the VPs in \mathcal{C} may have different periods.




Figure 5: Illustration of the worst case of $Y_\infty(t, \mathcal{C})$ for non-concrete asynchronous VPs.

4.1 A Common Period

We first consider the case in which the VPs in \mathcal{C} share a common period Π , i.e., $(\forall \Gamma_i = (\Pi_i, \Theta_i) \in \mathcal{C} :: \Pi_i = \Pi)$ holds. We establish our key proof obligation in Theorem 2 below. The following lemma is used in its proof. Specifically, we use it to show how to combine two VPs “locally” in a way that is in accordance with MP form.

Lemma 1. *Let $\Gamma_i = (\Pi, \Theta_i)$ and $\Gamma_j = (\Pi, \Theta_j)$ be two VPs that are not dedicated processors, and without loss of generality, assume $\Theta_i \leq \Theta_j$, i.e., $0 < w_i \leq w_j < 1$. Then, we have the following three exhaustive cases for $w_i + w_j$ and corresponding conclusions.*

1. *If $0 < w_i + w_j < 1$, then $Z(t, \Gamma_i) + Z(t, \Gamma_j) \leq Z(t, \Gamma_k)$, where $\Gamma_k = (\Pi, \Theta_k)$ and $\Theta_k = \Theta_i + \Theta_j$.*
2. *If $w_i + w_j = 1$, then $Z(t, \Gamma_i) + Z(t, \Gamma_j) \leq t$.*
3. *If $1 < w_i + w_j < 2$, then $Z(t, \Gamma_i) + Z(t, \Gamma_j) \leq t + Z(t, \Gamma_k)$, where $\Gamma_k = (\Pi, \Theta_k)$ and $\Theta_k = \Theta_i + \Theta_j - \Pi$.*

Proof. Figure 6 illustrates the three cases of the lemma. A rigorous proof is rather tedious and mechanical, so we defer it to an appendix. \square

Based on Lemma 1, we prove the following theorem by induction.

Theorem 2. *Given an arbitrary component $\mathcal{C} = (p, \mathcal{T})$ such that $(\forall \Gamma_i \in \mathcal{T} :: \Pi_i = \Pi)$, \mathcal{C} is dominated by the MP-form component $\mathcal{C}' = (p^*, \mathcal{T}^*)$ such that $\text{bw}(\mathcal{C}^*) = \text{bw}(\mathcal{C})$ and $(\forall \Gamma_i \in \mathcal{T}^* :: \Pi_i = \Pi)$.*


Proof. We prove the theorem by induction on $|\mathcal{T}|$.

Base Case: $|\mathcal{T}| \leq 1$. In this case, \mathcal{C} and \mathcal{C}^* are identical, because $\text{bw}(\mathcal{C}^*) = \text{bw}(\mathcal{C})$ and $(\forall \Gamma_i \in \mathcal{T}^* :: \Pi_i = \Pi)$. Therefore, by Def. 3, \mathcal{C}^* dominates \mathcal{C} .


Inductive Step. Suppose the theorem holds for any component \mathcal{C} such that $|\mathcal{T}| \leq k$ where $k \geq 1$. We prove that it also holds for any component \mathcal{C} such that $|\mathcal{T}| = k + 1$.

Because $k \geq 1$, $|\mathcal{T}| = k + 1 \geq 2$. Therefore, \mathcal{T} has at least two VPs that are not dedicated processors. Let Γ_i and Γ_j be two arbitrary such VPs. Without loss of generality, assume $0 < w_i \leq w_j < 1$.


To complete the proof, we show the existence of a component $\mathcal{C}' = (p', \mathcal{T}')$ such that \mathcal{C}' has the same bandwidth



(a) Illustration for Case 1



(b) Illustration for Case 2



(c) Illustration for Case 3

Figure 6: Illustration for the cases in Lemma 1.

and period as \mathcal{C} , but fewer VPs that are not dedicated processors, and $Y_\infty(t, \mathcal{C}) \leq Y_\infty(t, \mathcal{C}')$. \mathcal{C}' is constructed via three cases that hinge on the value of $w_i + w_j$.

Case 1: If $0 < w_i + w_j < 1$, then let $p' = p$ and $\mathcal{T}' = \mathcal{T} \setminus \{\Gamma_i, \Gamma_j\} \cup \{\Gamma'_k\}$ where Γ'_k is a new VP such that $\Pi'_k = \Pi$ and $\Theta'_k = \Theta_i + \Theta_j$. Clearly, $\text{bw}(\mathcal{C}) = \text{bw}(\mathcal{C}')$. Also,

$$\begin{aligned}
 & Y_\infty(t, \mathcal{C}) - Y_\infty(t, \mathcal{C}') \\
 &= \{\text{by (17)}\} \\
 &= (p - p')t + \sum_{\Gamma_l \in \mathcal{T}} Z(t, \Gamma_l) - \sum_{\Gamma_l \in \mathcal{T}'} Z(t, \Gamma_l) \\
 &= Z(t, \Gamma_i) + Z(t, \Gamma_j) - Z(t, \Gamma'_k) \\
 &\leq \{\text{by Lemma 1}\} \\
 &= 0.
 \end{aligned}$$

Case 2: If $w_i + w_j = 1$, then let $p' = p + 1$ and $\mathcal{T}' = \mathcal{T} \setminus \{\Gamma_i, \Gamma_j\}$. Clearly, $\text{bw}(\mathcal{C}) = \text{bw}(\mathcal{C}')$. Also,

$$\begin{aligned}
& Y_\infty(t, \mathcal{C}) - Y_\infty(t, \mathcal{C}') \\
&= \{\text{by (17)}\} \\
& (p - p')t + \sum_{\Gamma_l \in \mathcal{T}} Z(t, \Gamma_l) - \sum_{\Gamma_l \in \mathcal{T}'} Z(t, \Gamma_l) \\
&= -t + Z(t, \Gamma_i) + Z(t, \Gamma_j) \\
&\leq \{\text{by Lemma 1}\} \\
& 0.
\end{aligned}$$

Case 3: If $1 < w_i + w_j < 2$, then let $p' = p + 1$ and $\mathcal{T}' = \mathcal{T} \setminus \{\Gamma_i, \Gamma_j\} \cup \{\Gamma_k\}$ where Γ_k is a new VP such that $\Pi_k = \Pi$ and $\Theta_k = \Theta_i + \Theta_j - \Pi$. Clearly, $\text{bw}(\mathcal{C}) = \text{bw}(\mathcal{C}')$. Also,

$$\begin{aligned}
& Y_\infty(t, \mathcal{C}) - Y_\infty(t, \mathcal{C}') \\
&= \{\text{by (17)}\} \\
& (p - p')t + \sum_{\Gamma_l \in \mathcal{T}} Z(t, \Gamma_l) - \sum_{\Gamma_l \in \mathcal{T}'} Z(t, \Gamma_l) \\
&= -t + Z(t, \Gamma_i) + Z(t, \Gamma_j) - Z(t, \Gamma_k) \\
&\leq \{\text{by Lemma 1}\} \\
& 0.
\end{aligned}$$

In all three cases, the following two expressions hold.

$$\text{bw}(\mathcal{C}') = \text{bw}(\mathcal{C}) = \text{bw}(\mathcal{C}^*) \quad (18)$$

$$Y_\infty(t, \mathcal{C}) \leq Y_\infty(t, \mathcal{C}') \quad (19)$$

Also, in Cases 1 and 3, we have $|\mathcal{T}'| = |\mathcal{T}| - 1$, while in Case 2, we have $|\mathcal{T}'| = |\mathcal{T}| - 2$, so $|\mathcal{T}'| \leq |\mathcal{T}| - 1 = (k + 1) - 1 = k$. Therefore, by (18) and by the inductive hypothesis, \mathcal{C}' is dominated by \mathcal{C}^* . Hence, by Def. 3,

$$Y_\infty(t, \mathcal{C}') \leq Y_\infty(t, \mathcal{C}^*). \quad (20)$$

By (19) and (20), $Y_\infty(t, \mathcal{C}) \leq Y_\infty(t, \mathcal{C}^*)$. Also, since \mathcal{C}^* is in MP form, by Theorem 1, \mathcal{C}^* dominates \mathcal{C} . \square

The above theorem shows that, given a bandwidth and a common period shared by a set of asynchronous VPs, a component's supply is maximized when it is in MP form.

4.2 Different Periods

We now shift our focus by considering components that consist of a set of asynchronous VPs that may have different periods. Specifically, we consider a component $\mathcal{C} = (p, \mathcal{T})$, where for any two VPs $\Gamma_i, \Gamma_j \in \mathcal{T}$, $\Pi_i \neq \Pi_j$ may hold. We investigate whether such a component \mathcal{C} is dominated by a component in MP with the same bandwidth.

Towards this end, let \mathcal{C}^* be a component in MP form such that $\text{bw}(\mathcal{C}) = \text{bw}(\mathcal{C}^*)$. To begin, note that if $\text{bw}(\mathcal{C})$ is an integer, then \mathcal{C}^* clearly dominates \mathcal{C} , because \mathcal{C}^* has only dedicated processors that provide supply constantly. In the rest of this section, we consider the more interesting case wherein $\text{bw}(\mathcal{C})$ is not an integer. In this case, because \mathcal{C}^* is




Figure 7: Illustration of the counterexample in Sec. 4.

in MP form, $|\mathcal{T}^*| = 1$. Let $\Gamma^* = (\Pi^*, \Theta^*)$ denote the lone VP in \mathcal{T}^* .

It is easy to see that, if \mathcal{C}^* is to dominate \mathcal{C} , then the period Π^* generally will be dependent on the periods of the VPs in \mathcal{C} . In particular, if Π^* is selected to be very large in comparison to the periods of the VPs in \mathcal{C} , then Γ^* may be unable to guarantee any supply over relatively long intervals in which the VPs in \mathcal{C} do. One obvious conjecture is that \mathcal{C}^* will dominate \mathcal{C} as long as $\Pi^* \leq \min\{\Pi_i \mid \Gamma_i \in \mathcal{T}\}$ holds. However, the following counterexample shows that this conjecture is not true.

Counterexample. Consider a component \mathcal{C} with these two VPs: $\Gamma_1 = (11, 1.1(1 + \vartheta))$ and $\Gamma_2 = (10, 10 - \vartheta)$, where ϑ is an arbitrary small positive real number, i.e., $\vartheta \rightarrow 0^+$. An MP-form component \mathcal{C}^* with the same bandwidth also has two VPs: a dedicated processor and $\Gamma^* = (10, 1)$. Note that, in this setting, $\Pi^* \leq \min\{\Pi_i \mid \Gamma_i \in \mathcal{T}\}$ holds. As illustrated in Figure 7, $Y_\infty(22, \mathcal{C}) = 1.1(1 + \vartheta) + 22 - 4\vartheta = 23.1 - 2.9\vartheta$, while $Y_\infty(22, \mathcal{C}^*) = 1 + 22 = 23$. Because $\vartheta \rightarrow 0^+$, $23.1 - 2.9\vartheta > 23$. That is, $Y_\infty(22, \mathcal{C}) > Y_\infty(22, \mathcal{C}^*)$, which implies that \mathcal{C}^* does not dominate \mathcal{C} .

Despite the negative implications of this counterexample, we show next that \mathcal{C}^* does indeed dominate \mathcal{C} if Π^* is further restricted.

Theorem 3. \mathcal{C} is dominated by the MP-form component \mathcal{C}^* as defined above as long as $\Pi^* \leq \frac{1}{2} \min\{\Pi_i \mid \Gamma_i \in \mathcal{T}\}$.

Proof. Given \mathcal{C} , we first construct a new component \mathcal{C}' such that $p' = p$ and $|\mathcal{T}'| = |\mathcal{T}|$. Each $\Gamma'_i = (\Pi'_i, \Theta'_i) \in \mathcal{T}'$ is constructed from the VP $\Gamma_i = (\Pi_i, \Theta_i) \in \mathcal{T}$ by defining $\Pi'_i = \Pi^*$ and $\Theta'_i = \Theta_i \frac{\Pi^*}{\Pi_i}$. These definitions imply

$$w'_i = \frac{\Theta'_i}{\Pi'_i} = \frac{\Theta_i}{\Pi_i} = w_i. \quad (21)$$

By Property 2,

$$\begin{aligned}
Z(t, \Gamma'_i) &\geq \max\{(t - 2(\Pi'_i - \Theta'_i))w'_i, 0\} \\
&= \{\text{by (21) and because } \Pi'_i = \Pi^*\} \\
&\quad \max\{(t - 2\Pi^*(1 - w_i))w_i, 0\} \\
&\geq \{\text{because } \Pi^* \leq \frac{1}{2} \min\{\Pi_i \mid \Gamma_i \in \mathcal{T}\}\} \\
&\quad \max\{(t - \Pi_i(1 - w_i))w_i, 0\}.
\end{aligned}$$

On the other hand, by Property 3,

$$\begin{aligned} Z(t, \Gamma_i) &\leq \max\{(t - (\Pi_i - \Theta_i))w_i, 0\} \\ &= \max\{(t - \Pi_i(1 - w_i))w_i, 0\}, \end{aligned}$$

from which we can conclude the following.

$$(\forall i : 1 \leq i \leq |\mathcal{T}| = |\mathcal{T}'| :: Z(t, \Gamma_i) \leq Z(t, \Gamma'_i)) \quad (22)$$

Also, $p = p'$, and therefore, by (17)

$$Y_\infty(t, \mathcal{C}) \leq Y_\infty(t, \mathcal{C}'). \quad (23)$$

Because $(\forall \Gamma'_i :: \Pi'_i = \Pi^*)$ holds, and by (21), $\text{bw}(\mathcal{C}') = p' + \sum_{\Gamma'_i \in \mathcal{T}'} w'_i = p + \sum_{\Gamma_i \in \mathcal{T}} w_i = \text{bw}(\mathcal{C}) = \text{bw}(\mathcal{C}^*)$ holds, \mathcal{C}' is a component in which all VPs share the same period, and its bandwidth equals the bandwidth of the MP-form component \mathcal{C}^* . Therefore, by Theorem 2, \mathcal{C}^* dominates \mathcal{C}' . By Def. 3, this implies

$$Y_\infty(t, \mathcal{C}') \leq Y_\infty(t, \mathcal{C}^*). \quad (24)$$

By (23) and (24), $Y_\infty(t, \mathcal{C}) \leq Y_\infty(t, \mathcal{C}^*)$ holds, so by Theorem 1, the MP-form component \mathcal{C}^* dominates \mathcal{C} . \square

In some cases, the dominance of \mathcal{C}^* over \mathcal{C} can be established with a weaker restriction on the period Π^* . The following theorem gives such a case; note that harmonic and weakly harmonic¹ periods satisfy the condition given in this theorem.

Theorem 4. *For the component $\mathcal{C} = (p, \mathcal{T})$ defined above, let $\Pi_{\min} = \min\{\Pi_i \mid \Gamma_i \in \mathcal{T}\}$. If the condition $(\forall \Gamma_i \in \mathcal{C} :: \Pi_i = \Pi_{\min} \vee \Pi_i \geq 2\Pi_{\min})$ holds, then \mathcal{C} is dominated by the MP-form component \mathcal{C}^* as defined above if Π^* is set equal to Π_{\min} .*

Proof. We construct \mathcal{C}' in the same way as in the proof of Theorem 3 such that $\Pi'_i = \Pi^* = \Pi_{\min}$ and $\Theta'_i = \Theta_i \frac{\Pi^*}{\Pi_i}$. Given the statement of Theorem 4, we have for each i , $\Pi_i = \Pi^* = \Pi'_i$ or $\Pi_i \geq 2\Pi_{\min} = 2\Pi^* = 2\Pi'_i$. In the former case, $Z(t, \Gamma_i) = Z(t, \Gamma'_i)$ holds; in the latter case, $Z(t, \Gamma_i) \leq Z(t, \Gamma'_i)$ can be shown to follow from Properties 2 and 3 using the same reasoning as in the proof of Theorem 3. Thus, we can establish (22) in the context of this new theorem, and then show exactly as done in the proof of Theorem 3 that \mathcal{C}^* dominates \mathcal{C} . \square

5 Synchronous and Concrete Asynchronous

In this section, we consider components consisting of VPs with specified phases, *i.e.*, both *concrete asynchronous* and *synchronous* VPs. These VPs may have either a common period or different periods.

The case of synchronous VPs and a common period is highly related to the MPR model [17], as that model enforces both of these requirements. The following theorem is easily implied by prior work on the MPR model [18] that shows that, by enforcing MP form, a component abstracted

¹The smallest period divides any larger period.




Figure 8: Illustration of the counterexample in Sec. 5.

by the MPR model achieves its maximum supply.

Theorem 5. *(Follows from [18]) If $\mathcal{C} = (p, \mathcal{T})$ is a synchronous component and $(\forall \Gamma_i \in \mathcal{T} :: \Pi_i = \Pi)$ holds, then it is dominated by the MP-form component $\mathcal{C}' = (p^*, \mathcal{T}^*)$, where $\text{bw}(\mathcal{C}^*) = \text{bw}(\mathcal{C})$ and $(\forall \Gamma_i \in \mathcal{T}^* :: \Pi_i = \Pi)$.*

Because synchronous VPs are a special case of concrete asynchronous VPs where all VP phases happen to be zero, one might expect that Theorem 5 can be extended to concrete asynchronous VPs, and speculate that an arbitrary concrete asynchronous component with a common period is dominated by the MP-form component of the same bandwidth and period. However, this is unfortunately not true.

Counterexample. Consider a non-MP-form component \mathcal{C} that has two VPs, $\Gamma_1 = (\Pi, \vartheta)$ and $\Gamma_2 = (\Pi, \vartheta)$, with a common period and arbitrarily small budget, *i.e.*, $\vartheta \rightarrow 0^+$. Suppose these two VPs have different phases, $\phi_1 = 0$ and $\phi_2 = \frac{1}{2}\Pi$, as shown in Figure 8. Observe that any time interval of length $\frac{3}{2}\Pi$ must include exactly one period of allocation of Γ_1 or Γ_2 . Therefore, $Y_1(\frac{3}{2}\Pi, \mathcal{C}) \geq \vartheta$. In contrast, consider the MP-form counterpart of \mathcal{C} : $\mathcal{C}^* = \{\Gamma^*\}$, where $\Gamma^* = (\Pi, 2\vartheta)$. By the worst case illustrated in Figure 8, $Y_1(t, \mathcal{C}^*) = 0$ holds for any $t \leq 2(\Pi - 2\vartheta)$. Because $\vartheta \rightarrow 0^+$, we have $\frac{3}{2}\Pi < 2(\Pi - 2\vartheta)$. Therefore, $Y_1(\frac{3}{2}\Pi, \mathcal{C}^*) = 0 < \vartheta \leq Y_1(\frac{3}{2}\Pi, \mathcal{C})$, which implies that \mathcal{C}^* does not dominate \mathcal{C} .

Nonetheless, we provide next a theorem that shows that a component in non-MP-form will still be dominated by an MP-form component of the same bandwidth, provided the period of the latter is properly selected. Furthermore, the required period selection is valid not only for concrete asynchronous VPs with a common period, but also for synchronous or concrete asynchronous VPs with different periods. The theorem is stated assuming concrete asynchronous VPs, a category that subsumes these other possibilities.

Theorem 6. *If $\mathcal{C} = (p, \mathcal{T})$ is a concrete asynchronous component, then it is dominated by the MP-form component $\mathcal{C}' = (p^*, \mathcal{T}^*)$, where $\text{bw}(\mathcal{C}^*) = \text{bw}(\mathcal{C})$, provided the following condition holds: if $|\mathcal{T}^*| = 1$, then the period of the*

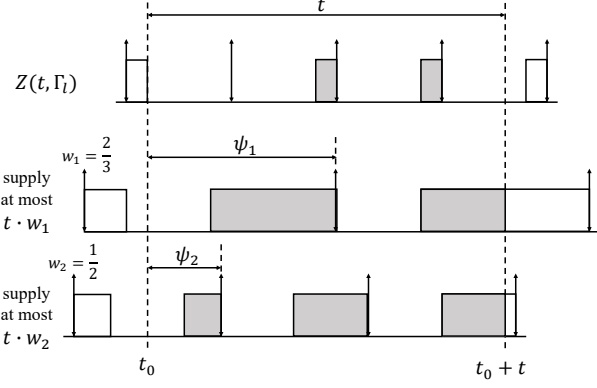


Figure 9: A possible scenario for any concrete phases.

lone VP in \mathcal{T} must satisfy

$$\Pi^* \leq \frac{\Pi_l(1 - w_l)w_l}{2(1 - w^*)w^*}, \quad (25)$$

where l is defined by

$$\Pi_l - \Theta_l = \min\{\Pi_i - \Theta_i \mid \Gamma_i \in \mathcal{T}\}. \quad (26)$$

Proof. Because \mathcal{C}^* is in MP form, $|\mathcal{T}| \leq 1$ holds. If $|\mathcal{T}| = 0$ holds, then \mathcal{C}^* has dedicated processors only. Because $\text{bw}(\mathcal{C}^*) = \text{bw}(\mathcal{C})$ is assumed, this clearly implies that \mathcal{C}^* dominates \mathcal{C} . In the rest of the proof, we focus on the more interesting case wherein $|\mathcal{T}^*| = 1$ holds. In this case, $\text{bw}(\mathcal{C}^*)$ is not integral, so $\text{bw}(\mathcal{C})$ is also not integral. This implies that $|\mathcal{T}| > 0$ holds. We now show that $Y_\infty(t, \mathcal{C}) \leq Y_\infty(t, \mathcal{C}^*)$ holds by considering two cases.

Case 1: $t \leq \Pi_l - \Theta_l$. By Claim 1 and (26), any VP Γ_i in \mathcal{T} can provide zero supply within any time interval of length t , where $t \leq \Pi_l - \Theta_l$. Within any such time interval, the p dedicated processors of \mathcal{C} provide supply continually. Because \mathcal{C}^* is in MP form, $p \leq p^*$ holds. Therefore, $Y_\infty(t, \mathcal{C}) = t \cdot p \leq t \cdot p^* \leq Y_\infty(t, \mathcal{C}^*)$.

Case 2: $t > \Pi_l - \Theta_l$. Let Γ_l be a VP such that $\Pi_l - \Theta_l = \min\{\Pi_i - \Theta_i \mid \Gamma_i \in \mathcal{T}\}$. Then, the allocations described next and illustrated in Figure 9 are *possible* for any concrete VP phases (i.e., synchronous or concrete asynchronous). Let t_0 be a time instant such that Γ_l gets its minimal supply $Z(t, \Gamma_l)$ within the time interval $[t_0, t_0 + t)$. For any other VP Γ_j , where $j \neq l$, let ψ_j denote the distance from t_0 to the start of its next allocation period, i.e., the next allocation period of Γ_j at or after time t_0 starts at time $t_0 + \psi_j$. (Note that the value of ψ_j will depend on the phase of Γ_j .) In this *possible* allocation sequence, if Γ_j has an allocation period that includes t_0 (as depicted), then assume that it provides a supply of $(\Pi_j - \psi_j) \cdot w_j$ time units within that allocation period before t_0 , i.e., in $[t_0 - (\Pi_j - \psi_j), t_0)$. Regardless of whether Γ_j has an allocation period that includes t_0 , assume that it provides supply as late as possible in each of its allocation periods beyond time t_0 . It is easy to show that, in this situation, each Γ_j provides a supply of at most $t \cdot w_j$ time units during $[t_0, t_0 + t)$. By Def. 2, the PSF functions

capture the *minimum* allocation that can occur, which is upper bounded by that demonstrated in the *possible* allocation sequence just discussed. Therefore, we have

$$\begin{aligned} & Y_\infty(t, \mathcal{C}) \\ & \leq t \cdot p + Z(t, \Gamma_l) + \sum_{\Gamma_j \in \mathcal{T} \wedge j \neq l} t \cdot w_j \\ & \leq \{\text{by Property 3}\} \\ & \quad t \cdot p + \max\{w_l \cdot (t - (\Pi_l - \Theta_l)), 0\} + \sum_{\Gamma_j \in \mathcal{T} \wedge j \neq l} t \cdot w_j. \\ & \leq \{\text{by our assumption in Case 2 that } t > \Pi_l - \Theta_l \text{ holds}\} \\ & \quad t \cdot p + w_l \cdot (t - (\Pi_l - \Theta_l)) + \sum_{\Gamma_j \in \mathcal{T} \wedge j \neq l} t \cdot w_j. \\ & = \{\text{rearranging}\} \\ & \quad t \cdot \left(p + \sum_{\Gamma_i \in \mathcal{T}} w_i \right) - w_l \cdot (\Pi_l - \Theta_l) \\ & = \{\text{by (5) and the definition of } w_l\} \\ & \quad t \cdot \text{bw}(\mathcal{C}) - \Pi_l(1 - w_l)w_l. \end{aligned} \quad (27)$$

By Claim 2 and our assumption that $|\mathcal{T}^*| = 1$ holds, we have

$$\begin{aligned} & Y_\infty(t, \mathcal{C}^*) \\ & = t \cdot p^* + Z(t, \Gamma^*) \\ & \geq \{\text{by Property 2}\} \\ & \quad t \cdot p^* + \max\{w^* \cdot (t - 2(\Pi^* - \Theta^*)), 0\} \\ & \geq \{\text{because } \max\{x, y\} \geq x\} \\ & \quad t \cdot p^* + w^* \cdot (t - 2(\Pi^* - \Theta^*)) \\ & = \{\text{rearranging and using the definition of } w^*\} \\ & \quad t \cdot (p^* + w^*) - 2\Pi^*(1 - w^*)w^* \\ & = \{\text{by (5)}\} \\ & \quad t \cdot \text{bw}(\mathcal{C}^*) - 2\Pi^*(1 - w^*)w^*. \end{aligned} \quad (28)$$

By (27) and (28),

$$\begin{aligned} & Y_\infty(t, \mathcal{C}) - Y_\infty(t, \mathcal{C}^*) \\ & \leq \{\text{because } \text{bw}(\mathcal{C}) = \text{bw}(\mathcal{C}^*)\} \\ & \quad 2\Pi^*(1 - w^*)w^* - \Pi_l(1 - w_l)w_l \\ & \leq \{\text{by (25)}\} \\ & \quad 0. \end{aligned}$$

That is, $Y_\infty(t, \mathcal{C}) \leq Y_\infty(t, \mathcal{C}^*)$ for $t > \Pi_l - \Theta_l$.

Combining Cases 1 and 2, we have $Y_\infty(t, \mathcal{C}) \leq Y_\infty(t, \mathcal{C}^*)$ for any $t \geq 0$. Also, \mathcal{C}^* is in MP form. Thus, by Theorem 1, \mathcal{C}^* dominates \mathcal{C} . \square

6 Indomitability of MP Form

Although we have shown that an arbitrary component can always be dominated by a component in MP form with the same bandwidth, this result requires restrictions on the pe-

riod of the MP-form component in some cases. This raises the question of whether the dominance is really due to the definition of MP form or just side effect of the period restrictions. In this section, we address this question. We show that an MP-form component can never be dominated by a non-MP-form component of the same bandwidth, regardless of any restrictions that may be applied to the non-MP-form component.

The following theorem holds, regardless of whether the VPs are synchronous, concrete asynchronous, or non-concrete asynchronous.

Theorem 7. *Given an MP-form component \mathcal{C}^* and an arbitrary non-MP-form component \mathcal{C} such that $\text{bw}(\mathcal{C}^*) = \text{bw}(\mathcal{C})$ holds, \mathcal{C} does not dominate \mathcal{C}^* , no matter how $\{\Pi_i \mid \Gamma_i \in \mathcal{C}\}$ is defined.*

Proof. Let p and p^* denote the number of dedicated processors in \mathcal{C} and \mathcal{C}^* , respectively. Because \mathcal{C}^* is in MP form and $\text{bw}(\mathcal{C}) = \text{bw}(\mathcal{C}^*)$ holds, we have $p \leq p^*$. We consider the two cases $p < p^*$ and $p = p^*$ separately below.

Case 1: $p < p^*$. By Claim 1, regardless of the VPs' phases, the supply of each VP $\Gamma_i \in \mathcal{T}$ can be zero for any time interval of length t such that $0 < t \leq \Pi_i - \Theta_i$, so $Y_\infty(t, \mathcal{C}) = t \cdot p$ for any t such that $0 < t \leq t_s$, where $t_s = \min\{\Pi_i - \Theta_i \mid \Gamma_i \in \mathcal{T}\}$. On the other hand, $Y_\infty(t, \mathcal{C}^*) \geq t \cdot p^*$ for any $t > 0$. Thus, for any t such that $0 < t \leq t_s$, we have $Y_\infty(t, \mathcal{C}) = t \cdot p < t \cdot p^* \leq Y_\infty(t, \mathcal{C}^*)$, i.e., $Y_\infty(t, \mathcal{C}) < Y_\infty(t, \mathcal{C}^*)$. Note that the stated range for t is not vacuous. This is because \mathcal{C} is not in MP form, which implies that $|\mathcal{T}| > 0$ holds, and hence that $t_s > 0$ holds as well. Because $Y_\infty(t, \mathcal{C}) < Y_\infty(t, \mathcal{C}^*)$ holds, by Def. 3, \mathcal{C} does not dominate \mathcal{C}^* .

Case 2: $p = p^*$. In this case, we have $|\mathcal{T}^*| = 1$, because if $|\mathcal{T}^*| = 0$ holds, then either \mathcal{C} is also in MP form or $\text{bw}(\mathcal{C}) > \text{bw}(\mathcal{C}^*)$, neither of which is allowed by the statement of the theorem. Let Γ^* denote the lone VP in \mathcal{C}^* and let w^* denote its bandwidth. Then, $w^* = \sum_{\Gamma_i \in \mathcal{T}} w_i$, since $\text{bw}(\mathcal{C}^*) = \text{bw}(\mathcal{C})$. Also, because \mathcal{C} is not in MP form, by Def. 1, both $|\mathcal{T}| \geq 2$ and $(\forall \Gamma_i \in \mathcal{T} :: w_i > 0)$ hold. Therefore, $(\forall \Gamma_i \in \mathcal{T} :: w_i < w^*)$. Letting $w_{\max} = \max\{w_i \mid \Gamma_i \in \mathcal{T}\}$, this implies

$$w_{\max} < w^*. \quad (29)$$

Let δ be the greatest common divisor of the values in $\{\Pi_i \mid \Gamma_i \in \mathcal{T}\}$. Then, the processor-time allocation illustrated in Figure 10, where every VP provides $\delta \cdot w_i$ time units of processor time at the end of every *aligned* time window of δ time units, is possible regardless of any assumptions regarding the VPs' phases. This is because, in this schedule, each VP Γ_i is allocated Θ_i time units within *any* time interval of length Π_i . Such an allocation satisfies the specification of Γ_i regardless of how phases are defined. Under this allocation pattern, each VP other than the one with the maximum bandwidth w_{\max} provides all of its supply in parallel with that maximum-bandwidth VP. Furthermore, with the depicted allocations, the minimum supply during any time interval of length t with a degree of parallelism of *one*




Figure 10: Illustration of Case 2 of Theorem 7.

is

$$\left\lfloor \frac{t}{\delta} \right\rfloor \delta w_{\max} + \max\left\{t - \left\lfloor \frac{t}{\delta} \right\rfloor \delta - (1 - w_{\max})\delta, 0\right\} \leq t \cdot w_{\max}.$$

Because the PSF functions, by Def. 2, capture the worst case among all possible allocation scenarios,

$$Y_1(t, \mathcal{T}) \leq t \cdot w_{\max}. \quad (30)$$

Therefore, given that \mathcal{C} has p dedicated processors,

$$Y_{p+1}(t, \mathcal{C}) = tp + Y_1(t, \mathcal{T}) \leq t(p + w_{\max}). \quad (31)$$

On the other hand, for \mathcal{C}^* , for any $t \geq 2(\Pi^* - \Theta^*) = 2\Pi^*(1 - w^*)$, by (17)

$$\begin{aligned} & Y_{p^*+1}(t, \mathcal{C}^*) \\ &= \{\text{by (8) and because } \mathcal{C}^* \text{ is in MP form}\} \\ & Y_\infty(t, \mathcal{C}^*) \\ &= \{\text{by (17)}\} \\ & tp^* + Z(t, \Gamma^*) \\ &\geq \{\text{by Property 2, and since } t \geq 2(\Pi^* - \Theta^*)\} \\ & tp^* + w^*(t - 2(\Pi^* - \Theta^*)) \\ &= \{\text{rearranging and using } w^* = \Theta^*/\Pi^*\} \\ & t(p^* + w^*) - 2\Pi^*w^*(1 - w^*). \end{aligned}$$

Because $p = p^*$ holds in Case 2,

$$Y_{p+1}(t, \mathcal{C}^*) \geq t(p + w^*) - 2\Pi^*w^*(1 - w^*). \quad (32)$$

By (31) and (32), for any $t \geq 2\Pi^*(1 - w^*)$,

$$Y_{p+1}(t, \mathcal{C}^*) - Y_{p+1}(t, \mathcal{C}) \geq t(w^* - w_{max}) - 2\Pi^*w^*(1 - w^*).$$

Hence, by (29), for any $t > \frac{2\Pi^*w^*(1-w^*)}{w^*-w_{max}} > 2\Pi^*(1 - w^*)$, $Y_{p+1}(t, \mathcal{C}) < Y_{p+1}(t, \mathcal{C}^*)$.

Thus, by Def. 3, \mathcal{C} does not dominate \mathcal{C}^* . Note that the above argument is valid regardless of the definition of $\{\Pi_i \mid \Gamma_i \in \mathcal{C}\}$. \square

Theorem 7 shows that, no matter how the periods of a non-MP-form component are defined, it cannot dominate any component in MP form with the same total bandwidth.

7 Related Work

In work on uniprocessors, Mercer *et al.* [14] proposed a mechanism that abstracts the notion of a processor capacity reservation as a uniprocessor with reduced speed. Abeni and Buttazzo [1] proposed the *constant bandwidth server (CBS)*. Lipari and Baruah [10] extended CBS to a hierarchical scheduling framework. Mok *et al.* [15] proposed the *bounded delay partition*, based upon which Lipari and Bini [11] derived the “best” server parameters for a given application. The PR model proposed by Shin and Lee [17] was extended by Easwaran *et al.* [7] to allow VPs to have relative deadlines different from periods.

In work on multiprocessors, Leontyev and Anderson [9] initially proposed MP form to schedule each component using at most one partially available processor in soft real-time systems. Shin *et al.* [16] proposed the MPR model. Easwaran *et al.* [8] derived a cluster-based hierarchical scheduler by applying the MPR model. Burmyakov *et al.* [4] extended the MPR model by providing information of resource allocation at each degree of parallelism. Xu *et al.* [18] extended the MPR model to the DMPR model by requiring VPs to be allocated in MP form, and proposed a cache-aware analysis framework.

In much of the just-cited work, a supply-bound function is provided to characterize the minimum resource allocation of a component, in order to perform schedulability analysis. Furthermore, Bini *et al.* [3] proposed the multi supply function (MSF) to provide a supply-bound function for each VP. Subsequently, Bini *et al.* [2] proposed PSFs, which are strictly more powerful than MSFs. PSFs provide a supply-bound function for each degree of parallelism. To the best of our knowledge, PSFs are the most expressive means of characterizing resource-allocation supply on multiprocessors.

8 Conclusion

We studied processor allocations to components comprised of multiple VPs, which may be synchronous, concrete asynchronous, or non-concrete asynchronous. We showed that any arbitrary component is always dominated by an MP-form component of the same bandwidth, provided the pe-

riod used in defining the MP-form component meets certain requirements. We also showed that a component in MP form can never be dominated by any non-MP-form component of the same bandwidth, regardless of how periods are defined.

MP supply form has additional advantages that are beyond the scope of this paper. For example, task migration costs will tend to be less when fewer processors are used, and dedicated processors are easier to allocate than non-dedicated ones. Also, synchronization protocols can more easily cope with only one partially available processor than multiple ones, which may lose supply at different times. We intend to consider such additional advantages in future work. Additionally, the period restrictions derived in this paper are not known to be tight. That is, a violation of these restrictions does not necessarily imply non-dominance. We defer the study of tight period restrictions to future work.

Acknowledgements: We are grateful to Enrico Bini and Linh Thi Xuan Phan for helpful discussions concerning the results in this paper.

References

- [1] L. Abeni and G. Buttazzo. Integrating multimedia applications in hard real-time systems. In *19th RTSS*, 1998.
- [2] E. Bini, M. Bertogna, and S. Baruah. Virtual multiprocessor platforms: Specification and use. In *30th RTSS*, 2009.
- [3] E. Bini, G. Buttazzo, and M. Bertogna. The multi supply function abstraction for multiprocessors. In *15th RTCSA*, 2009.
- [4] A. Burmyakov, E. Bini, and E. Tovar. The generalized multiprocessor periodic resource interface model for hierarchical multiprocessor scheduling. In *20th RTNS*, 2012.
- [5] Z. Deng and J.W.S. Liu. Scheduling real-time applications in an open environment. In *18th RTSS*, 1997.
- [6] G. Durrieu, M. Faugere, D.G. Girbal, S. and Perez, C. Pagetti, and W. Puffitsch. Predictable flight management system implementation on a multicore processor. In *Embedded Real Time Software*, 2014.
- [7] A. Easwaran, M. Anand, and I. Lee. Compositional analysis framework using EDP resource models. In *28th RTSS*, 2007.
- [8] A. Easwaran, I. Shin, and I. Lee. Optimal virtual cluster-based multiprocessor scheduling. *Real-Time Systems*, 43(1), 2009.
- [9] H. Leontyev and J. Anderson. A hierarchical multiprocessor bandwidth reservation scheme with timing guarantees. In *20th ECRTS*, 2008.
- [10] G. Lipari and S. Baruah. A hierarchical extension to the constant bandwidth server framework. In *7th RTAS*, 2001.
- [11] G. Lipari and E. Bini. Resource partitioning among realtime applications. In *15th ECRTS*, 2003.
- [12] G. Lipari and E. Bini. A framework for hierarchical scheduling on multiprocessors: from application requirements to run-time allocation. In *31th RTSS*, 2010.
- [13] C. Liu and J. Layland. Scheduling algorithms for multiprogramming in a hard real-time environment. *JACM*, 30:46–61, 1973.
- [14] C.W. Mercer, S. Savage, and H. Tokuda. Processor capacity reserves: Operating system support for multimedia applications. In *Proceedings of IEEE International Conference on Multimedia Computing and Systems*, 1994.
- [15] A.K. Mok, X. Feng, and D. Chen. Resource partition for real-time systems. In *7th RTAS*, 2001.
- [16] I. Shin and I. Easwaran, A. and Lee. Hierarchical scheduling framework for virtual clustering of multiprocessors. In *20th ECRTS*, 2008.
- [17] I. Shin and I. Lee. Periodic resource model for compositional real-time guarantees. In *24th RTSS*, 2003.
- [18] M. Xu, L.T.X. Phan, O. Sokolsky, S. Xi, C. Lu, C. Gill, and I. Lee. Cache-aware compositional analysis of real-time multicore virtualization platforms. In *34th RTSS*, 2013.

Appendix: Proof of Lemma 1

In this appendix, we formally prove Lemma 1, which is restated below.

Lemma 1. *Let $\Gamma_i = (\Pi, \Theta_i)$ and $\Gamma_j = (\Pi, \Theta_j)$ be two VPs that are not dedicated processors, and without loss of generality, assume $\Theta_i \leq \Theta_j$, i.e., $0 < w_i \leq w_j < 1$. Then, we have the following three exhaustive cases for $w_i + w_j$ and corresponding conclusions.*

1. *If $0 < w_i + w_j < 1$, then $Z(t, \Gamma_i) + Z(t, \Gamma_j) \leq Z(t, \Gamma_k)$, where $\Gamma_k = (\Pi, \Theta_k)$ and $\Theta_k = \Theta_i + \Theta_j$.*
2. *If $w_i + w_j = 1$, then $Z(t, \Gamma_i) + Z(t, \Gamma_j) \leq t$.*
3. *If $1 < w_i + w_j < 2$, then $Z(t, \Gamma_i) + Z(t, \Gamma_j) \leq t + Z(t, \Gamma_k)$, where $\Gamma_k = (\Pi, \Theta_k)$ and $\Theta_k = \Theta_i + \Theta_j - \Pi$.*

Proof. We consider the three cases of the lemma individually.

Case 1: In this case,

$$\Theta_k = \Theta_i + \Theta_j, \quad (33)$$

so

$$\Theta_i \leq \Theta_j < \Theta_k, \quad (34)$$

because $\Theta_i \leq \Theta_j$ is assumed by the statement of the lemma. By (2), (33), and (34),

$$t'_{\Gamma_i} \leq t'_{\Gamma_j} < t'_{\Gamma_k}. \quad (35)$$

In the next paragraph, we dispense with all possibilities that occur when at least one of t'_{Γ_i} and t'_{Γ_j} is negative.

First, if $t'_{\Gamma_i} \leq t'_{\Gamma_j} < 0$, then by (1), $Z(t, \Gamma_i) + Z(t, \Gamma_j) = 0 \leq Z(t, \Gamma_k)$. Second, if $t'_{\Gamma_i} < 0 < t'_{\Gamma_j}$, then by (1), $Z(t, \Gamma_i) + Z(t, \Gamma_j) = 0 + Z(t, \Gamma_j) \leq Z(t, \Gamma_k)$, where the last inequality follows by (35) and Property 4. In the rest of the proof for Case 1, we focus on the remaining possibility for this case, namely, $0 \leq t'_{\Gamma_i} \leq t'_{\Gamma_j}$, which by (35), implies

$$0 \leq t'_{\Gamma_i} \leq t'_{\Gamma_j} < t'_{\Gamma_k}. \quad (36)$$

Applying (16) to $Z(t, \Gamma_i)$, $Z(t, \Gamma_j)$, and $Z(t, \Gamma_k)$, respectively, we have the following.

$$\begin{aligned} Z(t, \Gamma_i) &= \{\text{by (16)}\} \\ &\quad f(t'_{\Gamma_i}, \Gamma_i) \\ &\leq \{\text{by (36) and Property 4}\} \\ &\quad f(t'_{\Gamma_k}, \Gamma_k) \\ &= \{\text{by (15)}\} \\ &\quad \left\lfloor \frac{t'_{\Gamma_k}}{\Pi} \right\rfloor \Theta_i + \\ &\quad \max \left(t'_{\Gamma_k} - \Pi \left\lfloor \frac{t'_{\Gamma_k}}{\Pi} \right\rfloor - (\Pi - \Theta_i), 0 \right) \end{aligned} \quad (37)$$

Similarly, for the same reasons,

$$\begin{aligned} Z(t, \Gamma_j) &\leq \left\lfloor \frac{t'_{\Gamma_k}}{\Pi} \right\rfloor \Theta_j + \\ &\quad \max \left(t'_{\Gamma_k} - \Pi \left\lfloor \frac{t'_{\Gamma_k}}{\Pi} \right\rfloor - (\Pi - \Theta_j), 0 \right). \end{aligned} \quad (38)$$

By (15) and (16),

$$\begin{aligned} Z(t, \Gamma_k) &= \left\lfloor \frac{t'_{\Gamma_k}}{\Pi} \right\rfloor \Theta_k + \\ &\quad \max \left(t'_{\Gamma_k} - \Pi \left\lfloor \frac{t'_{\Gamma_k}}{\Pi} \right\rfloor - (\Pi - \Theta_k), 0 \right). \end{aligned} \quad (39)$$

For notational simplicity, we introduce the two terms below.

$$\Phi = t'_{\Gamma_k} - \Pi \left\lfloor t'_{\Gamma_k} / \Pi \right\rfloor \quad (40)$$

$$\begin{aligned} \Delta &= \max(\Phi - (\Pi - \Theta_i), 0) + \max(\Phi - (\Pi - \Theta_j), 0) \\ &\quad - \max(\Phi - (\Pi - \Theta_k), 0) \end{aligned} \quad (41)$$

Now, by (45), (46) and (47), we have

$$\begin{aligned} &Z(t, \Gamma_i) + Z(t, \Gamma_j) - Z(t, \Gamma_k) \\ &\leq \left\lfloor \frac{t'_{\Gamma_i}}{\Pi} \right\rfloor (\Theta_i + \Theta_j - \Theta_k) + \Delta \\ &= \{\text{by (33)}\} \\ &\quad \Delta. \end{aligned} \quad (42)$$

Given the derivation above, we can complete the proof by showing $\Delta \leq 0$. This result is implied by the following claim.

Claim 3. *In Case 1, $\Delta \leq 0$.*

Proof. By (40), $0 \leq \Phi < \Pi$. Also, by (34), $\Pi - \Theta_k < \Pi - \Theta_j \leq \Pi - \Theta_i$. Thus, we have four cases for the value of Φ : Φ lies within $[0, \Pi - \Theta_k)$, or $[\Pi - \Theta_k, \Pi - \Theta_j)$, or $[\Pi - \Theta_j, \Pi - \Theta_i)$, or $[\Pi - \Theta_i, \Pi)$. In each case, it is straightforward to eliminate the max functions in (41) and show that $\Delta \leq 0$ holds.

Case 1.1: In this case, $\Phi \in [0, \Pi - \Theta_k)$, which implies $\Delta = 0$.

Case 1.2: In this case, $\Phi \in [\Pi - \Theta_k, \Pi - \Theta_j)$, which implies $\Delta = 0 - (\Phi - (\Pi - \Theta_k)) \leq 0$, because $\Phi \geq \Pi - \Theta_k$ holds in this case.

Case 1.3: In this case, $\Phi \in [\Pi - \Theta_j, \Pi - \Theta_i)$, which implies $\Delta = (\Phi - (\Pi - \Theta_j)) - (\Phi - (\Pi - \Theta_k)) = \Theta_j - \Theta_k < 0$, by (34).

Case 1.4: In this case, $\Phi \in [\Pi - \Theta_i, \Pi]$, which implies $\Delta = (\Phi - (\Pi - \Theta_i)) + (\Phi - (\Pi - \Theta_j)) - (\Phi - (\Pi - \Theta_k)) = \Phi - \Pi + \Theta_i + \Theta_j - \Theta_k < 0$, by (33) and the fact that $\Phi < \Pi$ holds in this case. \square

Claim 3 and (42) together imply $Z(t, \Gamma_i) + Z(t, \Gamma_j) \leq t + Z(t, \Gamma_k)$, as required.

Case 2: In this case, $w_i + w_j = 1$. By Property 3, $Z(t, \Gamma_i) \leq \max\{(t - (\Pi - \Theta_i))w_i, 0\} \leq tw_i$. Similarly, $Z(t, \Gamma_j) \leq tw_j$. Thus, $Z(t, \Gamma_i) + Z(t, \Gamma_j) \leq t(bw_i + bw_j) = t$.

Case 3: In this case,

$$\Theta_k = \Theta_i + \Theta_j - \Pi, \quad (43)$$

so

$$\Theta_k < \Theta_i \leq \Theta_j, \quad (44)$$

since $\Theta_j < \Pi$ holds and $\Theta_i \leq \Theta_j$ is assumed by the statement of the lemma. By (2), (43), and (44), $t'_{\Gamma_k} < t'_{\Gamma_i} \leq t'_{\Gamma_j}$. In the next paragraph, we dispense with all possibilities that occur when at least one of t'_{Γ_k} , t'_{Γ_i} , and t'_{Γ_j} is negative.

First, if $t'_{\Gamma_k} < t'_{\Gamma_i} \leq t'_{\Gamma_j} < 0$, then by (1), $Z(t, \Gamma_i) + Z(t, \Gamma_j) = 0 \leq t + 0 = t + Z(t, \Gamma_k)$. Second, if $t'_{\Gamma_k} < t'_{\Gamma_i} < 0 < t'_{\Gamma_j}$, then by (1), $Z(t, \Gamma_i) + Z(t, \Gamma_j) = 0 + Z(t, \Gamma_j) \leq 0 + t = t + 0 = t + Z(t, \Gamma_k)$. Third, if $t'_{\Gamma_k} < 0 \leq t'_{\Gamma_i} \leq t'_{\Gamma_j}$, then we have $t'_{\Gamma_k} < 0$, which by (2), implies $t - (\Pi - \Theta_k) < 0$, and hence, $t < \Pi - \Theta_k$. By the statement of the lemma (and in particular, Case 3), $\Pi - \Theta_k = 2\Pi - \Theta_i - \Theta_j \leq 2(\Pi - \Theta_i)$. Thus, we have $t < 2(\Pi - \Theta_i)$, which by Property 1, implies $Z(t, \Gamma_i) = 0$. Also, by (1), $Z(t, \Gamma_j) \leq t$, and because $t'_{\Gamma_k} < 0$, $Z(t, \Gamma_k) = 0$. Therefore, $Z(t, \Gamma_i) + Z(t, \Gamma_j) = Z(t, \Gamma_j) \leq t = t + Z(t, \Gamma_k)$.

Next, we focus on the remaining possibility in Case 3, namely, $0 \leq t'_{\Gamma_k} < t'_{\Gamma_i} \leq t'_{\Gamma_j}$. Applying (16) to $Z(t, \Gamma_i)$, $Z(t, \Gamma_j)$, and $Z(t, \Gamma_k)$, respectively, we have the following.

$$\begin{aligned} Z(t, \Gamma_i) &= f(t'_{\Gamma_i}, \Gamma_i), \\ &= \{\text{by (15) and } \Pi_i = \Pi\} \\ &\quad \left\lfloor \frac{t'_{\Gamma_i}}{\Pi} \right\rfloor \cdot \Theta_i + \\ &\quad \max \left(t'_{\Gamma_i} - \Pi \left\lfloor \frac{t'_{\Gamma_i}}{\Pi} \right\rfloor - (\Pi - \Theta_i), 0 \right) \end{aligned} \quad (45)$$

$$\begin{aligned} Z(t, \Gamma_j) &= f(t'_{\Gamma_j}, \Gamma_j) \\ &= \{\text{rearranging}\} \\ &\quad f(t'_{\Gamma_i} + (t'_{\Gamma_j} - t'_{\Gamma_i}), \Gamma_j) \\ &\leq \{\text{by Property 5; note that } t'_{\Gamma_j} - t'_{\Gamma_i} \geq 0\} \\ &\quad f(t'_{\Gamma_i}, \Gamma_j) + (t'_{\Gamma_j} - t'_{\Gamma_i}) \end{aligned}$$

$$\begin{aligned} &= \{\text{by (2) and } \Pi_j = \Pi_i = \Pi\} \\ &\quad f(t'_{\Gamma_i}, \Gamma_j) + \Theta_j - \Theta_i \\ &= \{\text{by (15) and } \Pi_j = \Pi\} \\ &\quad \left\lfloor \frac{t'_{\Gamma_i}}{\Pi} \right\rfloor \cdot \Theta_j + \Theta_j - \Theta_i + \\ &\quad \max \left(t'_{\Gamma_i} - \Pi \left\lfloor \frac{t'_{\Gamma_i}}{\Pi} \right\rfloor - (\Pi - \Theta_j), 0 \right) \end{aligned} \quad (46)$$

$$\begin{aligned} Z(t, \Gamma_k) &= f(t'_{\Gamma_k}, \Gamma_k) \\ &= \{\text{rearranging}\} \\ &\quad f(t'_{\Gamma_i} - (t'_{\Gamma_i} - t'_{\Gamma_k}), \Gamma_k) \\ &\geq \{\text{by Property 5; note that } t'_{\Gamma_i} - t'_{\Gamma_k} \geq 0\} \\ &\quad f(t'_{\Gamma_i}, \Gamma_k) + (t'_{\Gamma_i} - t'_{\Gamma_k}) \\ &= \{\text{by (2) and } \Pi_i = \Pi_k = \Pi\} \\ &\quad f(t'_{\Gamma_i}, \Gamma_k) + \Theta_i - \Theta_k \\ &= \{\text{by (43)}\} \\ &\quad f(t'_{\Gamma_i}, \Gamma_k) + \Pi - \Theta_j \\ &= \{\text{by (15) and } \Pi_k = \Pi\} \\ &\quad \left\lfloor \frac{t'_{\Gamma_i}}{\Pi} \right\rfloor \cdot \Theta_k + \Pi - \Theta_j + \\ &\quad \max \left(t'_{\Gamma_i} - \Pi \left\lfloor \frac{t'_{\Gamma_i}}{\Pi} \right\rfloor - (\Pi - \Theta_k), 0 \right) \end{aligned} \quad (47)$$

For notational simplicity, we introduce the two terms below.

$$\Phi' = t'_{\Gamma_i} - \Pi \left\lfloor t'_{\Gamma_i} / \Pi \right\rfloor \quad (48)$$

$$\begin{aligned} \Delta' &= \max(\Phi' - (\Pi - \Theta_i), 0) + \max(\Phi' - (\Pi - \Theta_j), 0) \\ &\quad - \max(\Phi' - (\Pi - \Theta_k), 0) \end{aligned} \quad (49)$$

Now, by (45), (46) and (47), we have

$$\begin{aligned} &(Z(t, \Gamma_i) + Z(t, \Gamma_j)) - (t + Z(t, \Gamma_k)) \\ &\leq \left\lfloor \frac{t'_{\Gamma_i}}{\Pi} \right\rfloor (\Theta_i + \Theta_j - \Theta_k) + \Theta_j - \Theta_i - \Pi + \Theta_j - t + \Delta' \\ &= \{\text{rearranging and by (2) and (43)}\} \\ &\quad \left\lfloor \frac{t'_{\Gamma_i}}{\Pi} \right\rfloor \Pi + 2\Theta_j - \Theta_i - \Pi + \Delta' - (t'_{\Gamma_i} + \Pi - \Theta_i) \\ &= \{\text{rearranging}\} \\ &\quad \left\lfloor \frac{t'_{\Gamma_i}}{\Pi} \right\rfloor \cdot \Pi - t'_{\Gamma_i} - 2(\Pi - \Theta_j) + \Delta' \\ &= \{\text{by (48)}\} \\ &\quad \Delta' - \Phi' - 2(\Pi - \Theta_j). \end{aligned} \quad (50)$$

Given the derivation above, we can complete the proof

by showing $\Delta' - \Phi' - 2(\Pi - \Theta_j) \leq 0$. This result is implied by the following claim.

Claim 4. In Case 3, $\Delta' - \Phi' - 2(\Pi - \Theta_j) < 0$.

Proof. By (48), $0 \leq \Phi' < \Pi$. Also, by (44), $\Pi - \Theta_j \leq \Pi - \Theta_i < \Pi - \Theta_k$. Thus, we have four cases for the value of Φ' : Φ' lies within $[0, \Pi - \Theta_j)$, or $[\Pi - \Theta_j, \Pi - \Theta_i)$, or $[\Pi - \Theta_i, \Pi - \Theta_k)$, or $[\Pi - \Theta_k, \Pi)$. In each case, it is straightforward to eliminate the max functions in (49) and show that $\Delta' - \Phi' - 2(\Pi - \Theta_j) < 0$ holds.

Case 3.1: In this case, $\Phi' \in [0, \Pi - \Theta_j)$, which implies $\Delta' - \Phi' - 2(\Pi - \Theta_j) = -\Phi' - 2(\Pi - \Theta_j) < 0$.

Case 3.2: In this case, $\Phi' \in [\Pi - \Theta_j, \Pi - \Theta_i)$, which implies $\Delta' - \Phi' - 2(\Pi - \Theta_j) = -3(\Pi - \Theta_j) < 0$.

Case 3.3: In this case, $\Phi' \in [\Pi - \Theta_i, \Pi - \Theta_k)$, which implies

$$\begin{aligned}
& \Delta' - \Phi' - 2(\Pi - \Theta_j) \\
&= \{\text{by (49)}\} \\
& \quad \Phi' - 3(\Pi - \Theta_j) - (\Pi - \Theta_i) \\
&< \{\text{in this case, } \Phi' < \Pi - \Theta_k \text{ holds}\} \\
& \quad (\Pi - \Theta_k) - 3(\Pi - \Theta_j) - (\Pi - \Theta_i) \\
&= \{\text{rearranging}\} \\
& \quad -2\Pi + 2\Theta_j + (\Theta_i + \Theta_j - \Pi - \Theta_k) \\
&= \{\text{by (43)}\} \\
& \quad -2\Pi + 2\Theta_j \\
&< \{w_j < 1 \text{ in the lemma statement implies } \Theta_j < \Pi\} \\
& \quad 0.
\end{aligned}$$

Case 3.4: In this case, $\Phi' \in [\Pi - \Theta_k, \Pi)$, which implies

$$\begin{aligned}
& \Delta' - \Phi' - 2(\Pi - \Theta_j) \\
&= \{\text{by (49)}\} \\
& \quad -3\Pi + \Theta_i + 3\Theta_j - \Theta_k \\
&= \{\text{rearranging}\} \\
& \quad -2\Pi + 2\Theta_j + (\Theta_i + \Theta_j - \Pi - \Theta_k) \\
&= \{\text{by (43)}\} \\
& \quad -2\Pi + 2\Theta_j \\
&< \{w_j < 1 \text{ in the lemma statement implies } \Theta_j < \Pi\} \\
& \quad 0.
\end{aligned}$$

□

Claim 4 and (50) together imply $Z(t, \Gamma_i) + Z(t, \Gamma_j) \leq t + Z(t, \Gamma_k)$, as required. □