

# A Fluid-Flow Interpretation of SCED Scheduling

Jörg Liebeherr

Department of Electrical and Computer Engineering  
University of Toronto

**Abstract**—We show that a fluid-flow interpretation of Service Curve Earliest Deadline First (SCED) scheduling simplifies deadline derivations for this scheduler. By exploiting the recently reported isomorphism between min-plus and max-plus network calculus and expressing deadlines in a max-plus algebra, deadline computations no longer require explicit pseudo-inverse computations. SCED deadlines are provided for latency-rate as well as a class of piecewise linear service curves.

## I. INTRODUCTION

Service Curve Earliest Deadline First (SCED) [5] offers an alternative viewpoint on the design of packet scheduling algorithms. The usual approach is to first design a scheduling algorithm and then analyze its properties. SCED proceeds in the reverse order in that it provides a mechanism to realize a scheduling algorithm with given properties. The properties, such as guarantees on rate or delays, are expressed in terms of the concept of *service curves* of the network calculus [4]. Given a service curve, SCED computes deadlines for arriving traffic and transmits traffic in the order of deadlines. As long as no deadline is violated, the scheduler is guaranteed to satisfy the service curve guarantees. The SCED framework in [5] is completed by *schedulability conditions* that predict whether given service curves can be met at a transmission link with bounded (not necessarily fixed) capacity.

Since service curves are arbitrary non-negative increasing functions, SCED has a great deal of flexibility for offering different service guarantees to traffic flows. For example, it does not share the well-known drawback of weighted fair scheduling algorithms [6] when providing low delays to low-bandwidth traffic. (To achieve low delays, fair schedulers must increase the guaranteed rate of a flow).

Guarantees in SCED provide lower bounds on the service. SCED+ [3] is an extension to guarantee bounds on the delay jitter. Scheduling algorithms inspired by SCED, such as Hierarchical Fair Service Curve (HFSC) [12] are widely deployed in the network stack of current operating systems [2], [9].

Even though SCED appears to be an ideal vehicle for interpreting scheduling algorithms in the context of the network calculus, it has not played a major role in recent network calculus research. A closer inspection of SCED provides clues that may offer an explanation for the lack of interest. First, the original formulation of SCED in [5] assumes that all packets have the same size. The condition on equal packet sizes is relaxed in [7, Sec. 2.3.2] by adding an additional ‘packetizer’ service element. The analysis in [5] also requires that a packet can depart in the same time slot where it arrives. This amounts

to an assumption of ‘cut-through’ switching in a network, whereas most networks perform ‘store-and-forward’ switching. Finally, since SCED operations are described within the framework of the min-plus calculus, deadlines are expressed in terms of a pseudo-inverse of a traffic function, which is not very intuitive.

In this paper, we will show that the above issues can be resolved when describing SCED in terms of the max-plus network calculus. By adopting a fluid-flow interpretation of SCED operations, we can exploit the recently established duality between min-plus and max-plus network calculus [8] for an analysis of SCED. Since max-plus expressions are more convenient for computing timestamps, we use them for deadline computations in SCED. For schedulability conditions, we work with min-plus expressions, since the corresponding max-plus conditions become unwieldy.

In Sec. II we briefly discuss network calculus concepts used in this paper. In Sec. III we discuss SCED operations in terms of max-plus algebra expressions. In Sec. IV we derive schedulability conditions for the fluid-flow SCED scheduler. In Sec. V we address the computation of SCED deadlines in fluid-flow SCED. In Sec. VI we address deadline computations in a packet-level system.

## II. DUALITY OF MIN-PLUS AND MAX-PLUS NETWORK CALCULUS

The continuous-time min-plus network calculus conducts an analysis of network elements within a  $(\mathcal{F}_o, \wedge, \otimes)$  dioid algebra, where  $\mathcal{F}_o$  is the set of left-continuous, non-decreasing functions  $F : \mathbb{R} \rightarrow \mathbb{R}_o^+ \cup \{+\infty\}$ , with  $F(t) = 0$  if  $t \leq 0$ , the  $\wedge$ -operation is a pointwise minimum, and  $\otimes$  is the min-plus convolution, which is defined as  $F \otimes G(t) = \inf_{0 \leq s \leq t} \{F(s) + G(t-s)\}$  for two functions  $F, G \in \mathcal{F}_o$ . The cumulative amount of arrivals and departures at a network element in the time interval  $[0, t)$  is given by  $A(t)$  and  $D(t)$ , respectively, with  $A, D \in \mathcal{F}_o$ . The available service at a network element is expressed in terms of a function  $S \in \mathcal{F}_o$ , referred to as *minimum service curve*, which satisfies  $D(t) \geq A \otimes S(t)$  for all  $t$ . When arrivals are bounded by a function  $E \in \mathcal{F}_o$ , such that  $E(s) \geq A(t+s) - A(t)$  for all  $s$  and  $t$ , we say that  $E$  is a *traffic envelope* for  $A$ .

Functions in the max-plus network calculus designate the time of an arrival or departure event for a given number of bits. We refer to these functions as space-domain functions. The continuous-space version uses a  $(\mathcal{T}_o, \vee, \otimes)$  dioid, where  $\mathcal{T}_o$  is the set of right-continuous, non-decreasing functions  $F : \mathbb{R} \rightarrow \mathbb{R}_o^+ \cup \{-\infty\} \cup \{+\infty\}$ , with  $F(\nu) = -\infty$  if  $\nu < 0$

and  $F(\nu) \geq 0$  if  $\nu \geq 0$ . The  $\vee$ -operation is a pointwise maximum, and  $\overline{\otimes}$  is the max-plus convolution, with  $F \overline{\otimes} G(\nu) = \sup_{0 \leq \kappa \leq \nu} \{F(\kappa) + G(\nu - \kappa)\}$  for two functions  $F, G \in \mathcal{F}_o$ . Arrivals and departures are described by functions  $T_A \in \mathcal{T}_o$  and  $T_D \in \mathcal{T}_o$ . Here,  $T_A(\nu)$  is the arrival time of bit  $\nu$ , where bit values are allowed to be real numbers. A minimum service curve is a function  $\gamma_S \in \mathcal{T}_o$  such that  $T_D(\nu) \leq T_A \overline{\otimes} \gamma_S(\nu)$  for all  $\nu$ , and a traffic envelope  $\lambda_E \in \mathcal{T}_o$  for an arrival time function  $T_A$  satisfies  $\lambda_E(\mu) \leq T_A(\nu + \mu) - T_A(\nu)$  for all  $\nu$  and  $\mu$ .

As shown in [8], there exists an isomorphism between the min-plus and max-plus network calculus via the pseudo-inverse functions

$$\begin{aligned} F^\downarrow(y) &= \inf \{x \mid F(x) \geq y\} = \sup \{x \mid F(x) < y\}, \\ F^\uparrow(y) &= \sup \{x \mid F(x) \leq y\} = \inf \{x \mid F(x) > y\}, \end{aligned}$$

where  $F^\downarrow$  is referred to as lower pseudo-inverse and  $F^\uparrow$  as upper pseudo-inverse. The pseudo-inverses have the following properties:

- $F \in \mathcal{F}_o \Rightarrow F^\uparrow \in \mathcal{T}_o$ .
- $F \in \mathcal{T}_o \Rightarrow F^\downarrow \in \mathcal{F}_o$ .
- $F^\downarrow$  is left-continuous and  $F^\uparrow$  is right-continuous.
- $F$  is left-continuous  $\Rightarrow F = (F^\uparrow)^\downarrow$ .
- $F$  is right-continuous  $\Rightarrow F = (F^\downarrow)^\uparrow$ .

With the pseudo-inverses, we can map operations between the min-plus and max-plus network calculus by

- $(F \wedge G)^\uparrow(\nu) = F^\uparrow \vee G^\uparrow(\nu)$ .
- $(F \otimes G)^\uparrow(\nu) = F^\uparrow \overline{\otimes} G^\uparrow(\nu)$ .
- $(F + G)^\uparrow(\nu) = \inf_{0 \leq \kappa \leq \nu} \max\{F^\uparrow(\kappa), G^\uparrow(\nu - \kappa)\}$ .

For mapping in the other direction we have

- $(F \vee G)^\downarrow(t) = F^\downarrow \wedge G^\downarrow(t)$ .
- $(F \overline{\otimes} G)^\downarrow(t) = F^\downarrow \otimes G^\downarrow(t)$ .
- $\left(\inf_{0 \leq s \leq t} \max\{F(s), G(t - s)\}\right)^\downarrow = F^\downarrow(t) + G^\downarrow(t)$ .

With this, we can set  $A \equiv T_A^\downarrow$  and  $D \equiv T_D^\downarrow$ , as well as  $T_A \equiv A^\uparrow$  and  $T_D \equiv D^\uparrow$ . Service curves and traffic envelopes are related as follows:

- $D(t) \geq A \otimes S(t), \forall t \Rightarrow T_D(\nu) \leq T_A \overline{\otimes} S^\uparrow(\nu), \forall \nu$ .
- $E(s) \geq A(t + s) - A(t), \forall t, s \Rightarrow E^\uparrow(\mu) \leq T_A(\nu + \mu) - T_A(\nu), \forall \nu, \mu$ .
- $T_D(\nu) \leq T_A \overline{\otimes} \gamma_S(\nu), \forall \nu \Rightarrow D(t) \geq A \otimes \gamma_S^\downarrow(t), \forall t$ .
- $\lambda_E(\mu) \leq T_A(\nu + \mu) - T_A(\nu), \forall \nu, \mu \Rightarrow \lambda_E^\downarrow(s) \geq A(t + s) - A(t), \forall t, s$ .

With our convention to use  $S$  and  $E$  for service curves and traffic envelopes in the min-plus network calculus, and  $\gamma_S$  and  $\lambda_E$  in the max-plus network calculus, we can set  $S \equiv \gamma_S^\downarrow$  and  $\gamma_S \equiv S^\uparrow$ , as well as  $E \equiv \lambda_E^\downarrow$  and  $\lambda_E \equiv E^\uparrow$ .

As argued in [8], there is no isomorphism when the min-plus network calculus is defined in discrete time ( $t \in \mathbb{Z}$ ) or the max-plus calculus is defined in discrete space ( $\nu \in \mathbb{Z}$ ). It also does not exist for a packet-level characterization of traffic. Hence, to exploit the above relationships within SCED, we must resort to a fluid-flow description of traffic, where time and space are expressed by non-negative real numbers.

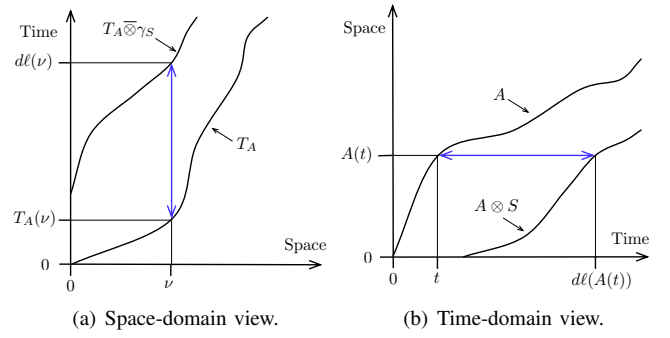


Fig. 1. SCED Deadlines.

### III. FLUID-FLOW SCED

The objective of SCED is a scheduling mechanism that can realize any given minimum service curve. The basic idea is to assign arriving traffic a deadline equal to the latest departure time permitted by the given service curve. As long as all traffic departs before the expiration of the assigned deadlines, the service curve is guaranteed to hold.

We first discuss the deadline assignment from the perspective of the max-plus algebra. We consider a fluid-flow version of SCED, where each bit value  $\nu \in \mathbb{R}_o^+$  is assigned a deadline  $dl(\nu)$ . All traffic is transmitted in the order of deadlines. The deadline assignment is illustrated in Fig. 1(a). Bit  $\nu$  with arrival time  $T_A(\nu)$  is assigned the deadline

$$dl(\nu) = T_A \overline{\otimes} \gamma_S(\nu), \quad (1)$$

where  $\gamma_S$  is a max-plus minimum service curve. This gives the equivalency

$$T_D(\nu) \leq dl(\nu) \iff T_D(\nu) \leq T_A \overline{\otimes} \gamma_S(\nu). \quad (2)$$

Hence, if all traffic departs by its deadline,  $\gamma_S$  is a minimum service curve. Conversely, if  $\gamma_S$  is a minimum service curve, then there is no deadline violation.

The deadline assignment is more intricate when we describe it in terms of min-plus network calculus expressions. The deadline assignment is sketched in Fig. 1(b) for a continuous arrival function  $A$ . The deadline of an arrival just before time  $t$ ,<sup>1</sup> denoted by  $dl(A(t))$ , is set to the time after  $t$  when  $A \otimes S$  has caught up to  $A(t)$ . If the departures at time  $dl(A(t))$ , given by  $D(dl(A(t)))$ , are at least  $A \otimes S(dl(A(t)))$ , then  $S$  satisfies the service curve requirement  $D(dl(A(t))) \geq A \otimes S(dl(A(t)))$ . The computation of the deadline involves the computation of an inverse. More precisely, since neither  $A$  nor  $A \otimes S$  are continuous or strictly increasing, the deadline requires to take a pseudo-inverse. By choosing the upper pseudo-inverse, we recover the deadline from (1), since

$$\begin{aligned} dl(\nu) &= T_A \overline{\otimes} \gamma_S(\nu) \\ &= A^\uparrow \overline{\otimes} S^\uparrow(\nu) \\ &= (A \otimes S)^\uparrow(\nu) \\ &= \sup \{ \tau \mid A \otimes S(\tau) \leq \nu \}. \end{aligned}$$

<sup>1</sup>Note that  $A(t)$  does not include arrivals that occur at time  $t$ .

Then, the deadline for  $A(t)$  is given by

$$d\ell(A(t)) = \sup\{\tau \mid A \otimes S(\tau) \leq A(t)\}.$$

Note that the computation of the pseudo-inverse for an arrival time  $t$  requires to compute  $A \otimes S(\tau)$  for values  $\tau > t$ , which appears to assume knowledge of future arrivals (after time  $t$ ). Fortunately, this is not the case, since for  $\tau > t$ ,

$$A \otimes S(\tau) = \inf_{0 \leq s \leq \tau} \{A(s) + S(\tau - s)\} \leq A(t)$$

if and only if

$$\inf_{0 \leq s \leq t} \{A(s) + S(\tau - s)\} \leq A(t).$$

Despite the additional complexity of deadline computations in a min-plus setting, all discussions of SCED in the literature [1], [3], [5], [7], [10] have chosen a min-plus formulation. Interestingly, the computations in these works use the lower pseudo-inverse for the computation of deadlines.

#### IV. SCHEDULABILITY CONDITION OF FLUID-FLOW SCED

In this section we derive a schedulability condition that determines whether a SCED scheduler at a link with variable transmission rate can support a set of min-plus or max-plus service curves for a set of flows. We will work within the min-plus network calculus, since the max-plus version of the schedulability condition is generally not useful for practical computations (as shown below).

We consider a set  $\mathcal{N}$  of flows. Let  $A_j$  and  $D_j$  denote the time-domain arrival and departure functions of flow  $j \in \mathcal{N}$ . The functions  $T_{A_j} = A_j^\downarrow$  and  $T_{D_j} = D_j^\downarrow$  denote the space-domain formulations of arrivals and departures. Arrivals of flow  $j$  are constrained by a traffic envelope, which is denoted by either  $E_j$  or  $\lambda_j = E_j^\downarrow$ .

We consider a work-conserving link with a time-variable transmission rate. We assume that the transmissions of the link can be bounded by a strict service curve  $C \in \mathcal{F}_o$  [7], defined by the property that for any time interval  $(s, t]$  with positive backlog,

$$\sum_j (D_j(t) - D_j(s)) \geq C(t - s).$$

With a constant-rate link, we have  $C(t) = ct$  for some  $c > 0$ . In the case of packet-level traffic, the transmission of a packet is never interrupted, even if a packet arrives with a shorter deadline than the packet in transmission. This is referred to as *non-preemptive* scheduling. In contrast, with *preemptive* scheduling, the link always transmits traffic with the earliest deadline.

We are interested in deriving a condition that can determine whether a SCED scheduler is able to guarantee service curves  $S_j$  or  $\gamma_{S_j} = S_j^\uparrow$  for each flow  $j \in \mathcal{N}$ . The deadline assignment for each  $\nu \geq 0$  is such that

$$d\ell_j(\nu) = T_{A_j} \bar{\otimes} \gamma_{S_j}(\nu) = (A_j \otimes S_j)^\uparrow(\nu). \quad (3)$$

#### A. Preliminary Results

We first present preliminary results that will aid in the derivation of the schedulability condition. We define  $A_j^{<t}(\tau)$  as

$$A_j^{<t}(\tau) = \sup\{\nu \mid 0 \leq T_{A_j}(\nu) < \tau \text{ and } d\ell_j(\nu) < t\},$$

which are the arrivals from flow  $j$  in the time interval  $[0, \tau)$  with a deadline less than  $t$ . We will use the shorthand  $A_j^{<t}(s, t) = A_j^{<t}(t) - A_j^{<t}(s)$ . Now, let  $t^*$  be the last time before  $t$  ( $t^* \leq t$ ) when the link does not have any backlog from traffic with a deadline before time  $t$ . Also, let  $\ell(t^*) \geq 0$  be the untransmitted portion of the packet that is in transmission at time  $t^*$ . This packet has a deadline greater than or equal to  $t$ . We next relate the function  $A_j^{<t}$  to deadline violations in SCED.

**Lemma 1.** *If a non-preemptive SCED scheduler experiences a deadline violation by time  $t$ , then*

$$\sum_{j \in \mathcal{N}} A_j^{<t}(t^*, t) + \ell(t^*) > C(t - t^*).$$

*Proof.* Let us first ignore that packet transmissions cannot be preempted. If we have a deadline violation by time  $t$ , the amount of traffic with a deadline before  $t$  exceeds the transmission capacity of the link. Since there is no traffic at the link at time  $t^*$  with a deadline before  $t$ , we can ignore all arrivals and transmissions before  $t^*$ . Then, the arrivals from flow  $j$  in  $[t^*, t)$  with a deadline before  $t$  are given by  $A_j^{<t}(t^*, t)$ . The least available transmission capacity of the link in  $[t^*, t)$  is given by  $C(t - t^*)$ . Therefore, a deadline violation by  $t$  implies that  $\sum_{j \in \mathcal{N}} A_j^{<t}(t^*, t) > C(t - t^*)$ . Without packet preemption, the remaining part of the packet in transmission at time  $t^*$ ,  $\ell(t^*)$ , is added to the workload that must be transmitted before  $t$ , which yields the claim.  $\square$

The next lemma provides an interesting property of the function  $A_j^{<t}$ .

**Lemma 2.** *For all  $t \geq 0$ , we have  $A_j^{<t}(t) = A_j \otimes S_j(t)$ .*

*Proof.* Setting  $t = \tau$  in the definition of  $A_j^{<t}(\tau)$ , we get

$$A_j^{<t}(t) = \sup\{\nu \mid T_{A_j} \bar{\otimes} \gamma_{S_j}(\nu) < t\},$$

since  $F \bar{\otimes} G(\nu) \geq F(\nu)$  for  $F, G \in \mathcal{T}_o$ . Writing deadlines in terms of the min-plus algebra, we obtain

$$\begin{aligned} A_j^{<t}(t) &= \sup\{\nu \mid (A_j \otimes S_j)^\uparrow(\nu) < t\} \\ &= ((A_j \otimes S_j)^\uparrow)^\downarrow(t) \\ &= A_j \otimes S_j(t), \end{aligned}$$

where the second line uses the lower-pseudo inverse, and the last line follows from  $F = (F^\uparrow)^\downarrow$  if  $F \in \mathcal{F}_o$ .  $\square$

#### B. Main Result

In this subsection we prove schedulability conditions for a fluid-flow SCED scheduler.

**Theorem 1.** A non-preemptive SCED scheduler with a set  $\mathcal{N}$  of flows as discussed at the beginning of this section guarantees the service curves  $\{S_j\}_{j \in \mathcal{N}}$  if for all  $t \geq 0$

$$\sum_{j \in \mathcal{N}} E_j \otimes S_j(t) \leq [C(t) - \ell_{\max}]^+, \quad (4)$$

where  $\ell_{\max}$  is the maximum packet size and  $[x]^+ = \max\{x, 0\}$ .

Note that the condition requires that  $S_j(t) = 0$  for  $t \leq C^\uparrow(\ell_{\max})$ . This can be ensured by ‘appending’ a delay element with service curve  $\delta_{C^\uparrow(\ell_{\max})}$  to a given service curve  $\gamma$  via  $\gamma \otimes \delta_{C^\uparrow(\ell_{\max})}$ .

With preemptive scheduling we set  $\ell_{\max} = 0$ . We point out that the schedulability condition of preemptive SCED when expressed in the max-plus algebra [8, see Corollary 12.5] is

$$\inf_{\nu_1, \dots, \nu_N} \max_{j=1, \dots, N} E_j^\uparrow \otimes S_j^\uparrow(\nu_j) \geq \frac{\nu}{C}, \quad \forall \nu \geq 0.$$

Clearly, this condition is not useful for practical schedulability tests.

*Proof.* According to (2), the deadline assignment from (3) does not result in a deadline violation if and only if  $\gamma_{S_j} = S_j^\downarrow$  is a max-plus service curve for flow  $j$  ( $T_{D_j} \leq T_{A_j} \otimes \gamma_{S_j}$ ). By the duality properties,  $S_j$  is then a min-plus service curve. We will show that a deadline violation implies that (4) does not hold. Hence, if (4) holds, there cannot be a deadline violation.

Assume that there is a deadline violation before  $t$ , and let  $t^*$  be as defined above. Each flow  $j \in \mathcal{N}$  satisfies

$$A_j^{<t}(t^*) = A_j(t^*) \quad (5)$$

This is a consequence from the fact that earlier arrivals of flow  $j$  have an earlier deadline. Since there are arrivals after time  $t^*$  with a deadline before  $t$ , the deadlines of all arrivals before  $t^*$  must be less than  $t$ . We now derive for flow  $j \in \mathcal{N}$

$$\begin{aligned} A_j^{<t}(t^*, t) &= A_j \otimes S_j(t) - A_j(t^*) \\ &= \inf_{0 \leq s \leq t} \{A_j(s) + S_j(t-s)\} - A_j(t^*) \\ &\leq \inf_{t^* \leq s \leq t} \{A_j(s) + S_j(t-s)\} - A_j(t^*) \\ &= \inf_{0 \leq s \leq t-t^*} \{A_j(t^* + s) - A_j(t^*) + S_j(t-t^* - s)\} \\ &\leq \inf_{0 \leq s \leq t-t^*} \{E_j(s) + S_j(t-t^* - s)\} \\ &= E_j \otimes S_j(t-t^*). \end{aligned}$$

In the first step, we use Lemma 2 and (5). The second step simply expands the convolution. The third step relaxes the infimum by restricting its range, and the fourth step makes a change of variable. The inequality in the fifth step follows since  $E$  is an envelope, that is,  $E_j(s) \geq A_j(t^* + s) - A_j(t^*)$ , which yields the convolution in the last step.

By Lemma 1, since  $E_j \otimes S_j(t) \geq A_j \otimes S_j(t)$  and  $\ell_{\max} \geq \ell(t^*)$ , a deadline violation before  $t$  implies that

$$\sum_{j \in \mathcal{N}} E_j \otimes S_j(t) + \ell_{\max} > C(t),$$

which contradicts (4). Thus, we cannot have a deadline violation. Hence, the functions  $S_j$  are minimum service curves.  $\square$

The following condition, which follows directly from Theorem 1, is useful when no information is available on the arrivals.

**Corollary 1.** Under the assumptions of Theorem 1, the SCED scheduler guarantees service curves  $\{S_j\}_{j \in \mathcal{N}}$  if for all  $t \geq 0$

$$\sum_{j \in \mathcal{N}} S_j(t) \leq [C(t) - \ell_{\max}]^+,$$

Since the available transmission capacity of the link in a time interval  $(s, t]$  may exceed  $C(t-s)$ , the condition in Theorem 1 is a sufficient condition. On the other hand, if the link is a fixed-rate work-conserving link with exact service curve  $C(t) = ct$ , such that  $\sum_j D_j(t) = (\sum_j A_j) \otimes C(t)$  for all  $t \geq 0$ , and arrivals of flows may saturate their envelopes, that is,  $A_j(t) = E_j(t)$ , we can provide a necessary condition for guaranteeing minimum service curves  $\{S_j\}_{j \in \mathcal{N}}$ , which are close to (4).

**Theorem 2.** Consider a SCED scheduler that offers an exact service curve  $C(t) = ct$ . Assume that the arrivals from each flow  $j \in \mathcal{N}$  can saturate its envelope  $E_j$ . If SCED ensures each flow  $j \in \mathcal{N}$  a minimum service curve  $S_j$ , then, for all  $t \geq 0$ ,

$$\sum_{j \in \mathcal{N}} E_j \otimes S_j(t) \leq ct. \quad (6)$$

For preemptive scheduling, the condition in (6) is necessary and sufficient. Since, for non-preemptive scheduling, the condition in (4) is not always necessary (e.g., if there is only one flow), reducing the difference between (4) and (6) requires knowledge of the number of flows and the service curves of each flow.

*Proof.* Suppose that (6) does not hold for some value of  $t$ . Let the arrivals saturate their envelopes, that is  $A_j(\tau) = E_j(\tau)$  for all  $0 \leq \tau \leq t$  for each  $j \in \mathcal{N}$ . Since, by assumption, each of the  $\{S_j\}_{j \in \mathcal{N}}$  is a minimum service curve, we have for each  $j \in \mathcal{N}$  that

$$D_j(t) \geq E_j \otimes S_j(t).$$

Summing over all flows and using the violation of (6), we get

$$\sum_{j \in \mathcal{N}} D_j(t) \geq \sum_{j \in \mathcal{N}} E_j \otimes S_j(t) > ct.$$

However, this is not possible since the aggregate departures from all flows in the interval  $[0, t]$  cannot exceed  $ct$ .  $\square$

## V. COMPUTATIONS OF SCED DEADLINES

As we have seen, by avoiding the need to compute pseudo-inverses, SCED deadlines in a max-plus setting are conceptually simpler and more intuitive than presented in the literature on SCED [1], [3], [5], [7], [11].

**Delay Guarantees:** The max-plus service curve for guaranteeing a delay bound  $d$  for traffic is simply  $\gamma_S(\nu) = d$ . This leads to the deadline computation

$$dl(\nu) = \max_{0 \leq \kappa \leq \nu} \{T_A(\kappa) + d\} = T_A(\nu) + d.$$

That is, the deadline is the sum of the arrival time and the delay bound. This deadline assignment corresponds to that of Earliest-Deadline-First (EDF) scheduling.

**Rate Guarantees:** The computation of SCED deadlines for a rate guarantee with service curve  $\gamma_S(\nu) = \frac{\nu}{R}$  requires a little bookkeeping at the start of a busy period of a flow. Here, a busy period of a flow is a maximal time interval where the flow has a positive backlog  $B$ , with  $B(t) = A(t) - D(t)$ . The following computation assumes that all arrivals occur at the start of or within a busy period, and that the number of busy periods within any finite time interval is finite. This assumption does not hold for general fluid-flow traffic arrivals. In particular, if traffic arrives at a constant rate and is served at the same rate, no backlog builds up, and, hence, there is no busy period. On the other hand, in practical scenarios, where arrivals occur in chunks of arbitrary size and the maximum service rate has an upper bound, any arrival creates a backlog, and, therefore, starts or falls into a busy period.

Consider the arrival time  $T_A(\nu)$  of a bit value  $\nu$ . We suppose the arrival occurs in a busy period that started at time  $t$ , that is,  $t = \sup\{s \leq t \mid A(t) = D(t)\}$ . Let  $\underline{\nu}$  be the bit that started the busy period, with arrival time  $T_A(\underline{\nu}) = t$ , that is,  $\underline{\nu} = \inf\{\kappa \mid T_A(\kappa) \geq t\}$ . For the delay  $W$ , defined as  $W(\nu) = T_D(\nu) - T_A(\nu)$ , we have

$$W(\underline{\nu}) = 0, \quad W(\kappa) > 0, \quad \forall \kappa \in (\underline{\nu}, \nu].$$

(In general, it is possible that  $W(\kappa) > 0$  for all  $\kappa \in (\underline{\nu}, \nu]$  [8, §11.5]. However, since the service curve  $\gamma_S(\kappa) = \frac{\kappa}{R}$  does not allow a delay at the start of a busy period, we get  $B(T_A(\underline{\nu})) = W(\underline{\nu}) = 0$ .) With  $T_D(\kappa) \leq T_A \bar{\otimes} \gamma_S(\kappa)$  for all  $\kappa$ , we therefore have

$$T_A(\kappa) < T_A \bar{\otimes} \gamma_S(\kappa), \quad \forall \kappa \in (\underline{\nu}, \nu]. \quad (7)$$

Under these assumptions, the interval over which  $T_A \bar{\otimes} \gamma_S(\nu)$  is computed can be reduced as stated in the following lemma.

**Lemma 3.** *Given an arrival  $\nu$  at time  $t$  to a network element that offers the service curve  $\gamma_S(\nu) = \frac{\nu}{R}$ . If  $\nu > \underline{\nu}$ , then*

$$T_A \bar{\otimes} \gamma_S(\nu) = \sup_{0 \leq \kappa \leq \underline{\nu}} \left\{ T_A(\kappa) + \frac{\nu - \kappa}{R} \right\}.$$

*Proof.* Since  $T_A$  and  $\gamma_S$  are right-continuous, by [8, Lemma 4.1(9)], there exists a  $\mu \in [0, \nu]$  such that

$$T_A \bar{\otimes} \gamma_S(\nu) = T_A(\mu) + \frac{\nu - \mu}{R}.$$

If  $\mu \in (\underline{\nu}, \nu]$ , we get

$$\begin{aligned} T_A \bar{\otimes} \gamma_S(\nu) &= T_A(\mu) + \frac{\nu - \mu}{R} \\ &< \sup_{0 \leq \kappa \leq \mu} \left\{ T_A(\kappa) + \frac{\mu - \kappa}{R} \right\} + \frac{\nu - \mu}{R} \\ &\leq \sup_{0 \leq \kappa \leq \nu} \left\{ T_A(\kappa) + \frac{\nu - \kappa}{R} \right\} \\ &= T_A \bar{\otimes} \gamma_S(\nu), \end{aligned}$$

In the second line, we used (7), and in the third line, we enlarged the range of the supremum. Obviously, there is a contradiction, and we can conclude that  $\mu \leq \underline{\nu}$ .  $\square$

Rewriting the result in Lemma 3 as

$$\begin{aligned} T_A \bar{\otimes} \gamma_S(\nu) &= \sup_{0 \leq \kappa \leq \underline{\nu}} \left\{ T_A(\kappa) + \frac{\nu - \kappa}{R} \right\} + \frac{\nu - \underline{\nu}}{R} \\ &= \max \left[ \sup_{0 \leq \kappa < \underline{\nu}} \left\{ T_A(\kappa) + \frac{\nu - \kappa}{R} \right\}, T_A(\underline{\nu}) \right] + \frac{\nu - \underline{\nu}}{R} \\ &= \max \{ T_A \bar{\otimes} \gamma_S(\underline{\nu}^-), t \} + \frac{\nu - \underline{\nu}}{R}, \end{aligned}$$

where we use the notation  $x^- = \sup_{y < x} y$ , we can construct a deadline assignment for the rate service curve. Let us add an index to the busy periods so that  $t_n$  and  $\underline{\nu}_n$  denote the start time and the first bit of the  $n$ th busy period of a flow. Then the deadline assignment in the  $n$ th busy period is given by

$$dl(\nu) = \max \left( dl(\underline{\nu}_n^-), t_n \right) + \frac{\nu - \underline{\nu}_n}{R}, \quad \nu \in [\underline{\nu}_n, \underline{\nu}_{n+1}). \quad (8)$$

For the computation of the first busy period, we define  $dl(\underline{\nu}_1^-) = -\infty$ . This is consistent with our derivations since  $\underline{\nu}_1 = 0$  and, with  $T_A \bar{\otimes} \gamma_S \in \mathcal{T}_o$ , we get  $T_A \bar{\otimes} \gamma_S(\nu) = -\infty$  for  $\nu < 0$ .

This deadline assignment is easily implemented, since we must only keep track of the arrived bits in the current busy period. Consider the  $n$ th busy period which starts at  $t_n$  with bit value  $\underline{\nu}_n$ . At the begin of a busy period, we take the larger of the current time ( $t_n$ ) and the last assigned deadline ( $dl(\underline{\nu}_n^-)$ ). This value is added to  $\frac{t}{R}$  to obtain the deadline of the  $\mu$ th bit in the busy period. Since  $\nu$  in (8) is equal to  $\nu = \mu + \underline{\nu}_n$  for  $\nu \in [\underline{\nu}_n, \underline{\nu}_{n+1})$ , the resulting deadline is equal to (8). Note that the start time of a busy period is simply the time of an arrival to an empty buffer.

**Latency-Rate Guarantees:** We can combine the deadline assignment of a delay server and a rate server to get the deadline assignment of a latency-rate server. Let  $\gamma_1(\nu) = \frac{\nu}{R}$  and  $\gamma_2(\nu) = d$ , the service curve of a latency-rate server is  $\gamma_1 \bar{\otimes} \gamma_2(\nu) = \frac{\nu}{R} + d$ . Since

$$T_A \bar{\otimes} (\gamma_1 \bar{\otimes} \gamma_2)(\nu) = T_A \bar{\otimes} \gamma_1(\nu) + d,$$

the deadline of  $\nu$  for a latency-rate server is given by  $\bar{dl}(\nu)$  from

$$\bar{dl}(\nu) = dl(\nu) + d, \quad (9)$$

where  $dl(\nu)$  is the deadline for  $\gamma_1$  computed with (8).

**Piecewise linear convex service curve:** The deadline computation with (8) and (9) can be extended to piecewise linear

convex max-plus service curves. A single segment of such a service curve has the form

$$\gamma_S(\nu) = \left[ \frac{\nu}{r} - e \right]^+$$

for some  $r > 0$  and  $e \geq 0$ . For  $e = 0$ , this service curve is obviously a rate server, and the deadline computation from (8) applies. For  $e > 0$ , we essentially have a delay correction with a negative value. Since the earliest deadline of a packet is its arrival time, we compute the deadline as

$$\bar{d}\ell(\nu) = \max\{d\ell(\nu) - e, T_A(\nu)\}. \quad (10)$$

with  $d\ell(\nu)$  from (8). We obtain a piecewise convex max-plus service curve with multiple segments, from

$$\gamma_S(\nu) = \max_{i=1, \dots, N} \left\{ \left[ \frac{\nu}{r_i} - e_i \right]^+ \right\},$$

with  $e_1 < e_2 < \dots < e_N$  and  $R_1 < R_2 < \dots < R_N$ . The deadline for the piecewise linear convex max-plus service curve is computed by  $\bar{d}\ell(\nu) = \max_{i=1, \dots, N} d\ell_i(\nu)$ , where  $d\ell_i$  is the deadline computed for the  $i$ th segment.

**Traffic shaping:** The SCED principle is also applicable to traffic shaping. A max-plus traffic envelope  $\lambda_E \in \mathcal{T}_o$  realizes an exact service curve, with  $T_D(\nu) = T_A \otimes \lambda_E(\nu)$ . Here, the convolution provides the time when the shaper releases bit  $\nu$ . We therefore refer to the max-plus convolution as the release time and denote it by  $r\ell$ , with

$$r\ell(\nu) = T_A \bar{\otimes} \lambda_E(\nu).$$

As an example, the max-plus envelope for a token bucket with rate  $r$  and bucket size  $b$  has the envelope  $\lambda(\nu) = \left[ \frac{\nu}{r} - \frac{b}{r} \right]^+$ . We compute the release times with (8) and (10), where we replace ‘ $d\ell$ ’ by ‘ $r\ell$ ’.

## VI. PACKETIZED SYSTEMS

While a fluid-flow interpretation of SCED is perfectly aligned with network calculus theory, verifying and enforcing deadlines for each (real) value  $\nu$  of a traffic flow is obviously not practical. In a packet system, each packet is assigned a single deadline, and all bits belonging to the same packet receive the same deadline. We now discuss adjustments of SCED for a packet-level system.

Let  $\ell_n$  denote the size of the  $n$ th packet ( $n \geq 1$ ) of a flow and  $L_n = \sum_{k=1}^n \ell_k$  the cumulative size of the first  $n$  packets, with  $L_0 = 0$ . The bits of the  $n$ th packet cover the range  $L_{n-1} \leq \nu < L_n$ . For a system with packet-level arrivals we have

$$T_A(\nu) = T_A^p(n), \quad \nu \in [L_{n-1}, L_n),$$

where  $T_A^p(n)$  is the arrival time of packet  $n$ . With a service curve for delays,  $\gamma_S(\nu) = d$ , the (fluid-flow) SCED deadline assignment according to (1) is

$$d\ell(\nu) = T_A^p(n) + d, \quad \nu \in [L_{n-1}, L_n),$$

hence a packet-level deadline assignment  $d\ell^p(n) = T_A^p(n) + d$  is congruent with the fluid-flow assignment.

With a rate-guarantee with  $\gamma_S(\nu) = \frac{\nu}{R}$ , fluid-flow SCED assigns each bit value of a packet a different deadline. For a packet-level system, we use [8, Eq. (12.14)] which showed

$$T_A^p \bar{\otimes} \gamma_S(\nu) = \max\{T_A \bar{\otimes} \gamma_S(L_{k-1}^-), T_A^p(k)\} + \frac{\nu - L_{k-1}}{R},$$

for  $\nu \in [L_{n-1}, L_n)$ . Hence, a packet-level deadline assignment

$$d\ell^p(n) = \max\{d\ell^p(n-1), T_A^p(k)\} + \frac{\ell_n}{R}, \quad (11)$$

relates to the fluid-flow assignment  $d\ell(\nu)$  by

$$d\ell(\nu) \leq d\ell^p(n) - \frac{\ell_n}{R}, \quad \nu \in [L_{n-1}, L_n).$$

Therefore, by adjusting the service curve by an additional delay, yielding  $\gamma_{S'}(\nu) = \frac{\nu + \ell_{\max}}{R}$ , where  $\ell_{\max}$  is the maximum packet size of the flow, the deadline  $d\ell' = T_A \bar{\otimes} \gamma_{S'}$  satisfies

$$d\ell'(\nu) \leq d\ell^p(n), \quad \text{if } \nu \in [L_{n-1}, L_n).$$

We see that the packet-level assignment in (11) meets all deadlines of a fluid-flow assignment with the adjusted rate service curve  $\gamma_{S'}$ . The deadline assignment in (11) is of course that of the VirtualClock scheduling algorithm in [13].

## VII. CONCLUSIONS

By resorting to max-plus algebra for the computation of deadlines, we showed that a SCED scheduler can be much simplified. The computation of deadlines for latency-rate and piecewise-linear convex max-plus service curves only requires state information on the start time and the traffic served in the current busy period. A packet-level algorithm for deadline computations emphasized the relationship of SCED for rate-based service curves and VirtualClock scheduling.

## REFERENCES

- [1] C.-S. Chang. *Performance Guarantees in Communication Networks*. Springer Verlag, 2000.
- [2] K. Cho. Managing traffic with ALTQ. In *Proc. USENIX, FREENIX Track*, pages 121–128, June 1999.
- [3] R. L. Cruz. SCED+: efficient management of quality of service guarantees. In *Proc. IEEE Infocom*, pages 625–634, April 1998.
- [4] R. L. Cruz. Quality of service guarantees in virtual circuit switched networks. *IEEE Journal on Selected Areas in Communications*, 13(6):1048–1056, 1995.
- [5] R. L. Cruz, H. Sariowan, and G. C. Polyzos. SCED: A generalized scheduling policy for guaranteeing quality-of-service. *IEEE/ACM Transactions on Networking*, 7(5):669–684, 1999.
- [6] A. Demers, S. Keshav, S. Shenker. Analysis and simulation of a fair queueing algorithm. In *Proc. ACM SIGCOMM*, pages 1–12, Aug. 1989.
- [7] J. Y. Le Boudec and P. Thiran. *Network Calculus*. Springer Verlag, Lecture Notes in Computer Science, LNCS 2050, 2001.
- [8] J. Liebeherr, Duality of the max-plus and min-plus network calculus. *Foundations and Trends in Networking*, 11(3-4):139–282, 2017.
- [9] K. Rechert. Queueing disciplines: traffic control mit Linux. *Linux-Magazin Online (in German)*, February 2005.
- [10] H. Sariowan. *A service curve approach to performance guarantees in integrated services networks*. Ph.D. thesis, University of California San Diego, 1996.
- [11] H. Sariowan, R. L. Cruz, and G. C. Polyzos. Scheduling for quality of service guarantees via service curves. In *Proc. ICCCN*, pages 512–520, Sep. 1995.
- [12] I. Stoica, H. Zhang, T. S. E. Ng. A hierarchical fair service curve algorithm for link-sharing real-time and priority services. In *Proc. ACM SIGCOMM*, pages 249–262, Sep. 1997.
- [13] L. Zhang. VirtualClock: A new traffic control algorithm for packet switching networks. In *Proc. ACM Sigcomm*, pages 19–29, Sep. 1990.