

Network Engineering Department, Universitat Politècnica de Catalunya

---

# An Energy-friendly Scheduler for Edge Computing System

**Alejandro Llorens-Carrodegua**  
([alejandro.llorens@entel.upc.edu](mailto:alejandro.llorens@entel.upc.edu))

Stefanos G. Sagkriotis


Cristina Cervelló-Pastor


Dimitrios P. Pezaros

November 24, 2021

- 1 Introduction
- 2 Problem statement and system model
- 3 Proposed scheduling solution
- 4 Evaluation and results
- 5 Conclusion

- 1 Introduction
- 2 Problem statement and system model
- 3 Proposed scheduling solution
- 4 Evaluation and results
- 5 Conclusion






+

Instantaneous  
response

Robustness

Availability




)

+

Instantaneous  
response

Robustness

Availability



)

+

Instantaneous  
response  
Robustness  
Availability

+

Inter-operative **cluster** compounding  
that enable failure-recovery and  
accumulative processing capacity

**Network Function  
Virtualization**

**Edge and Fog  
Computing**



**Network Function  
Virtualization**

**Edge and Fog  
Computing**





Allow fast deployment of applications

Better resource utilization and scalability


Cost reduction







Deploying and managing services in a cluster can be a laborious and error-prone task



Deploying and managing services in a cluster can be a laborious and error-prone task

An scheduler integrating energy measurements in the placement decisions, and considering all the cluster's nodes is required

- 1 Introduction
- 2 Problem statement and system model**
- 3 Proposed scheduling solution
- 4 Evaluation and results
- 5 Conclusion



A controller node is deployed alongside a set of computing nodes ( $N_c$ ) (cluster of SBCs)



A controller node is deployed alongside a set of computing nodes ( $N_c$ ) (cluster of SBCs)

Incoming requests are scheduled according to their deadline and resources requirements

A controller node is deployed alongside a set of computing nodes ( $N_c$ ) (cluster of SBCs)

Incoming requests are scheduled according to their deadline and resources requirements

Controller and computing nodes are energy-constrained devices with fixed computational resources

A controller node is deployed alongside a set of computing nodes ( $N_c$ ) (cluster of SBCs)

Incoming requests are scheduled according to their deadline and resources requirements

Controller and computing nodes are energy-constrained devices with fixed computational resources

+

The scheduler place events by considering remaining battery of nodes to increase the lifetime and resilience of the cluster

A controller node is deployed alongside a set of computing nodes ( $N_c$ ) (cluster of SBCs)

Incoming requests are scheduled according to their deadline and resources requirements

Controller and computing nodes are energy-constrained devices with fixed computational resources

+

The scheduler place events by considering remaining battery of nodes to increase the lifetime and resilience of the cluster

The controller not only must be capable of coordinating and managing the deployed services and tasks, but also monitor the battery status of each node, including its own battery levels

Notation	Description
P	Set of deployable units of computing nodes
N	Set of physical nodes where events can be scheduled
$N_c$	Set of computing nodes ( $N_c \subseteq N$ )
S	Network service request arriving to controller node
F	Sequence of VNFs compounding a network service request
T	Task request arriving to controller node
p	Each virtual node created on the physical nodes to run the events
n	Each physical node where virtual nodes are created
$p_{f_i}$	Indicates the virtual node where function $f_i$ is running
$p_T$	Indicates the virtual node where task T is running
$f_i$	Each network function forming part of a network service
$t_r$	Running time of an event before considering it completed
$t_s$	Starting time of an event when being processed in the selected node
$t_c$	Completion time of an event in the selected node
$t_e$	Execution time of an event in the assigned node
$t_a$	Arrival time of an event request in the controller node
$t_t$	Total time of an event in the system
$t_w$	Waiting time of an event in the priority queue
d	Deadline for processing a given event

$$t_e = t_c - t_s$$

$$t_t = t_c - t_a$$

$$t_w = \begin{cases} t_s - t_{a_s=1} & \text{if the event is a service} \\ t_s - t_{a_T} & \text{if the event is a task} \end{cases}$$

- 1 Introduction
- 2 Problem statement and system model
- 3 Proposed scheduling solution**
- 4 Evaluation and results
- 5 Conclusion

# Proposed scheduler

## SOC and capacity-based scheduler (SOCCS)

# Proposed scheduler

## SOC and capacity-based scheduler (SOCCS)



# Proposed scheduler

## SOC and capacity-based scheduler (SOCCS)





# Proposed scheduler

## Scheduler block (II)

# Proposed scheduler

## Scheduler block (II)

# Proposed scheduler

## Scheduler block (I)

$$n_{\text{score}} = \alpha (\text{SOC} = 100) + (1 - \alpha) (1 - \frac{E_{\text{CPU}}}{\text{Usage}_{\text{CPU}_{\text{max}}}})$$

# Proposed scheduler

## Scheduler block (I)

$$n_{\text{score}} = \alpha (\text{SOC} = 100) + (1 - \alpha) (1 - \frac{E_{\text{CPU}}}{\text{Usage}_{\text{CPU}_{\text{max}}}})$$

# Proposed scheduler

## Scheduler block (I)

$$n_{\text{score}} = \alpha (\text{SOC} = 100) + (1 - \alpha) (1 - \frac{E_{\text{CPU}}}{\text{Usage}_{\text{CPU}_{\text{max}}}})$$



- 1 Introduction
- 2 Problem statement and system model
- 3 Proposed scheduling solution
- 4 Evaluation and results**
- 5 Conclusion

Parameter	Values
Number of VNFs in a Service	5 - 10
CPU Capacity per Node (milliCPU)	4000
Memory Capacity per Node (Ki)	7998464
Required CPU per Event (milliCPU)	150 - 250
Required Memory per Event (Ki)	200 - 500
Event Arrival Rates (Events/time-unit)	2 - 12

CPU usage is measured in CPU units and is expressed as an absolute quantity. Thus, 100 milliCPU and 0.1CPU are the same amount of CPU usage in a single-core, dual-core, or 48-core machine.



Remarks:

Number of successfully scheduled events was relatively similar for both criteria

Remarks:

Number of successfully scheduled events was relatively similar for both criteria

Noticeable difference in the number of rejections and deadline violations for high-generation rates

Remarks:

Number of successfully scheduled events was relatively similar for both criteria

Noticeable difference in the number of rejections and deadline violations for high-generation rates

Reduction of 50%, 61% and 66% in the amount of rejected events

Remarks:

Number of successfully scheduled events was relatively similar for both criteria

Noticeable difference in the number of rejections and deadline violations for high-generation rates

Reduction of 50%, 61% and 66% in the amount of rejected events

Reductions are more significant than the increments in the deadline violations, which can be up to 16%, 42% and 49%







# Results

## SOC regression model (II)

# Results

## SOC regression model (II)

# Results

## SOC regression model (II)

SOC regression model for any node  $n$  in the SBC cluster ( $n \in N$ , where  $N = N_c + 1$ ):

$$SOC_{pred\_model} = \theta_0 + \theta_1 pkt_{in,n} + \theta_2 cpu_n + \sum_{i=1}^{N_c} (\theta_{3+i} compute_i) cpu_n^2 + \theta_{4+i} compute_i$$

SOC regression model for any node  $n$  in the SBC cluster ( $n \in N$ , where  $N = N_c + 1$ ):

$$SOC_{pred\_model} = \theta_0 + \theta_1 pkt_{in\_n} + \theta_2 cpu_n + \sum_{i=1}^{N_c} (\theta_{3+i} compute_i) cpu_i^2 + \theta_{4+i} compute_i$$

Model coefficients (  $\theta$  ) are obtained from the available training dataset

SOC regression model for any node  $n$  in the SBC cluster ( $n \in N$ , where  $N = N_c + 1$ ):

$$SOC_{pred\_model} = \beta_0 + \beta_1 pkt_{in\_n} + \beta_2 cpu_n + \sum_{i=1}^{N_c} (\beta_{3+i} compute_i) + \beta_4 compute_n^2$$

Model coefficients (  $\beta$  ) are obtained from the available training dataset

Categorical variables  $compute_i$  represents each compute node in the model ( $i \in N_c$ )

SOC regression model for any node  $n$  in the SBC cluster ( $n \in N$ , where  $N = N_c + 1$ ):

$$SOC_{pred\_model} = \beta_0 + \beta_1 pkt_{in_n} + \beta_2 ctrl + \sum_{i=1}^{N_c} (\beta_{3+i} compute_i) + \beta_4 compute_n^2$$

Model coefficients ( $\beta$ ) are obtained from the available training dataset

Categorical variables  $compute_i$  represents each compute node in the model ( $i \in N_c$ )

Binary indicator  $ctrl$  distinguishes if the node is a controller or not



SOC regression model for any node  $n$  in the SBC cluster ( $n \in N$ , where  $N = N_c + 1$ ):

$$SOC_{pred\_model} = \theta_0 + \theta_1 pkt_{in,n} ctrl + \theta_2 cpu_n + \sum_{i=1}^{N_c} (\theta_3 compute_i) cpu_n^2 + \theta_4 compute_i$$

Model coefficients ( $\theta$ ) are obtained from the available training dataset

Categorical variables  $compute_i$  represents each compute node in the model ( $i \in N_c$ )

Binary indicator  $ctrl$  distinguishes if the node is a controller or not

SOC regression model is transformed regarding the node whose SOC is to be predicted (i.e., a linear model with two predictor variables and a second-order polynomial model based on CPU usage)

Least Loaded Scheduler (LLS)    Kubernetes Scheduler (KS)    SOC and Capacity-based Scheduler (SOCCS)

Least Loaded Scheduler (LLS)    Kubernetes Scheduler (KS)    SOC and Capacity-based Scheduler (SOCCS)

Remarks:

For low-value generation rates, KS rejects some events (i.e., around 8 events for 5 events/time-unit)

Least Loaded Scheduler (LLS)    Kubernetes Scheduler (KS)    SOC and Capacity-based Scheduler (SOCCS)

Remarks:

For low-value generation rates, KS rejects some events (i.e., around 8 events for 5 events/time-unit)

SOCCS outperforms the other algorithms in terms of scheduled VNFs for high-value generation rates

Least Loaded Scheduler (LLS)    Kubernetes Scheduler (KS)    SOC and Capacity-based Scheduler (SOCCS)

Remarks:

For low-value generation rates, KS rejects some events (i.e., around 8 events for 5 events/time-unit)

SOCCS outperforms the other algorithms in terms of scheduled VNFs for high-value generation rates

SOCCS reduces the rejected VNFs w.r.t LLS by 11%, 12% and 12% for high generation rates (i.e., 8, 10 and 12)

Least Loaded Scheduler (LLS)    Kubernetes Scheduler (KS)    SOC and Capacity-based Scheduler (SOCCS)

Remarks:

For low-value generation rates, KS rejects some events (i.e., around 8 events for 5 events/time-unit)

SOCCS outperforms the other algorithms in terms of scheduled VNFs for high-value generation rates

SOCCS reduces the rejected VNFs w.r.t LLS by 11%, 12% and 12% for high generation rates (i.e., 8, 10 and 12)

These reductions are higher in comparison with KS, demonstrated by the values up to 16%, 23% and 18%









Remarks:

The greater the event arrival rate, the higher the battery consumption

**Remarks:**

The greater the event arrival rate, the higher the battery consumption

KS has the highest battery consumption for all the generation rates

Remarks:

The greater the event arrival rate, the higher the battery consumption

KS has the highest battery consumption for all the generation rates

SOCCS guarantees the lowest battery consumption as it always selected the node with the maximum SOC and minimum CPU usage

- 1 Introduction
- 2 Problem statement and system model
- 3 Proposed scheduling solution
- 4 Evaluation and results
- 5 Conclusion

We proposed a scheduling algorithm, i.e. SOCCS, that assigns events to an SBC cluster by considering predicted battery consumption and used resources.

We analyzed the case of using the unused controller resources to deploy certain services and tasks. The obtained results confirmed that such consideration is a convenient criterion in a resource-constrained environment.

The evaluation results revealed that SOCCS outperformed the baseline algorithms with regard to analyzed metrics such as rejected and scheduled events, deadline violations in scheduled events and battery consumption.

Our scheduler reduced the cluster's operation time to process the requested events which indisputably had a lower impact in the battery consumption with regard to the other schedulers.



**Thank you very much  
for your time.**

**3G**

**4G Any question?**

**5G**