Ref. Ares(2019)1458898 - 04/03/2019

LOw Power Heterogeneous Architecture
for NExt Generation of SmaRt Infrastructure and Platforms
in Industrial and Societal Applications

Truck Use Case 3

Co-funded by the Horizon 2020
Framework Programme of the European Union

DELIVERABLE NUMBER	D7.6
DELIVERABLE TITLE	Truck Use Case 3
RESPONSIBLE AUTHOR	Luca Scanavino (CSI)

GRANT AGREEMENT N.	688386
PROJECT REF. NO	H2020- 688386
PROJECT ACRONYM	OPERA
PROJECT FULL NAME	LOw Power Heterogeneous Architecture for Next Generation of SmaRt Infrastructure and Platform in Industrial and Societal Applications
STARTING DATE (DUR.)	01/12/2015
ENDING DATE	30/11/2018
PROJECT WEBSITE	www.operaproject.eu

WORKPACKAGE N. | TITLE	WP7 | Hardware/Software integration and Validation in Real Life workload
WORKPACKAGE LEADER	CSI Piemonte
DELIVERABLE N. | TITLE	D7.6 | Truck use case 3
RESPONSIBLE AUTHOR	Luca Scanavino (CSI)
DATE OF DELIVERY (CONTRACTUAL)	30/11/2018 (M36)
DATE OF DELIVERY (SUBMITTED)	15/02/2019 (M39)
VERSION | STATUS	V2.0 | Final Version
NATURE	R(Report)/ D(Demonstrator)
DISSEMINATION LEVEL	PU(Public)
AUTHORS (PARTNER)	Luca Scanavino (CSI)

VERSION	MODIFICATION(S)	DATE	AUTHOR(S)
0.1	First draft	18/10/2018	Luca Scanavino (CSI)
0.2	HPE Contributions	19/10/2018	Gallig Renaud (HPE)
0.3	TESEO internal review	02/11/2018	Roberto Peveri (TESEO)
0.4	IBM internal review	04/11/2018	Joel Nider (IBM)
1.0	Final Version	06/11/2018	Luca Scanavino (CSI)
1.1	In response to reviewer comments, we reviewed Executive Summary, 1.2.1 and 1.3. In addition, we added also 3.1, 3.2, 3.3	25/01/2019	Luca Scanavino (CSI)
1.2	HPE Contributions	25/01/2019	Gallig Renaud (HPE)
2.0	Final Version	31/01/2019	Luca Scanavino (CSI)

PARTICIPANTS		CONTACT
STMICROELECTRONICS SRL		Giulio Urlini Email: Giulio.urlini@st.com
IBM ISRAEL SCIENCE AND TECHNOLOGY LTD		Joel Nider Email: joeln@il.ibm.com
HEWLETT PACKARD CENTRE DE COMPETENCES (FRANCE)		Gallig Renaud Email: gallig.renaud@hpe.com
NALLATECH LTD		Craig Petrie Email: c.petrie@molex.com
ISTITUTO SUPERIORE MARIO BOELLA		Olivier Terzo Email: terzo@ismb.it
TECHNION ISRAEL INSTITUTE OF TECHNOLOGY		Dan Tsafrir Email: dan@cs.technion.ac.il
CSI PIEMONTE		Vittorio Vallero Email: vittorio.vallero@csi.it
NEAVIA TECHNOLOGIES		Stéphane Gervais Email: s.gervais@lacroix.fr
CERIOS GREEN BV		Frank Verhagen Email: frank.verhagen@certios.nl
TESEO SPA		Stefano Serra Email: stefano.serra@eiffage.com
DEPARTEMENT DE L'ISERE		Olivier Latouille Email: olivier.latouille@isere.fr

ACRONYMS LIST

Acronym	Description
EE	Energy Efficiency
FPGA	Field Programmable Gate Array
GPU	Graphics Processing Unit
HPE	Hewlett Packard Enterprise
IGN	Institut Géographique National
SOTA	State of the Art
UC	Use Case

LIST OF FIGURES

Figure 1 - MICMAC on low power server....8
Figure 2 – PhotoScan Pro on low power server with GPU 9
Figure 3 – Truck Use Case – EE orthomap trend....10
Figure 4 – Truck Use Case – EE radio trend 10
Figure 5 - ANN energy efficiency per thread 11
Figure 6 – Space occupied before OPERA....12
Figure 7 – Space occupied with OPERA 13
Figure 8 – Truck baseline comparison 15

LIST OF TABLES

Table 1 - Truck Use Case – EE trend 9
Table 2 – OPERA improvements 14

EXECUTIVE SUMMARY

Position of the deliverable in the whole project context

This deliverable describes activities of the Truck Use Case during the third cycle of the Opera project, specifically the period between M29 and M36.

For this reason, the content of the document is closely related to the deliverables D7.4 *Truck Use Case 1* and D7.5 *Truck Use Case 2* in which we described not only the activities and the results we accomplished during the first and second phases of Opera Project (M11 – M28) but also other information and contributions that are important for the last one. In this case we do not repeat those points, but we refer to them where necessary and only to promote the understanding.

Therefore, we can consider as valid the considerations about the "Position of the deliverable within OPERA Project" reported in D7.4, because we have the same target, the same partners, the same strategy and the same relationship with the other work packages and other deliverables.

Description of the deliverable

Aiming at guaranteeing coherence and consistency with the other Tasks and Deliverables strictly related to D7.6, we chose the following structure:

- Evolution over the I° and II° Cycles – a summary of the previous phases;
- III° cycle description– Results of the use case during the last phase;
- Outputs and Conclusions – Outcomes due to OPERA project and final remarks.

The main and important inputs for D7.6 are contained in deliverables D7.4 and D7.5 which address these topics:

- The description of OPERA Project targets for each phase
- The configuration of OPERA infrastructure set up during the first and second phases
- The measurements in different condition to implement the orthophoto processing

In these elements, we find the experience and the knowledge of the first two phases, but there are also other inputs that can be taken into account: the technological improvements and the measurements indications provided by WP4 and WP6, in particular the D6.8 where it's reported MICMAC porting activity.

In addition, we also report the improvements provided by OPERA in terms of energy, time and space and how the truck services changed thanks to this project.

In the last chapter (Results & Conclusions), we highlight the improvement in terms of EE comparing the baseline with the best OPERA configuration and the technological items and features that were realized within project effort and we can find in the best configuration.

TABLE OF CONTENTS

EXECUTIVE SUMMARY 5
1 EVOLUTION OVER I° AND II° CYCLE 7
1.1 MEASUREMENT METHODOLOGY 7
1.2 MEASUREMENTS 7
1.2.1 Baseline 7
1.2.2 MICMAC on low power server 8
1.2.3 PhotoScan Pro on low power server without and with GPU 8
1.3 ENERGY EFFICIENCY TREND 9
2 III° CYCLE DESCRIPTION 11
2.1 MICMAC PORTING 11
2.2 REAL-LIFE SERVICE 12
2.3 DEMONSTRATOR 13
3 RESULTS & CONCLUSIONS 14
3.1 BASELINE COMPARISON 15
3.2 2013 SOTA COMPARISON 15
3.3 OPERA OUTCOMES 15
REFERENCES 17

1 EVOLUTION OVER I° AND II° CYCLE

In this section we report the main topics and outcomes met during the first two cycles to appreciate the gradual evolution and the crucial decisions taken in that period, to better understand in the next chapters the last phase activities and the results obtained thanks to OPERA.

Specifically, we report the most significant configurations set up during the first two cycles to highlight and to compare their performance. In addition, we describe also the methodology used to take measurements with the aim to guarantee congruence and uniformity.

1.1 MEASUREMENT METHODOLOGY

It is very important to define policies and rules for measurements, because during OPERA project we set up different configurations and for each one we had to compare consistent values.

In particular, as described in [17], we followed this pattern:

1. We considered always the same cluster of photos to obtain the same orthophoto
2. We measured the duration to obtain the orthophoto
3. We measured the energy consumption to obtain the orthophoto

Thanks to duration and energy consumption, as described in [17], we can obtain two energy efficiency parameters, the first one referred to the orthophoto elaboration and the second one referred to radio services:

$$EE_{Orthomap} = \frac{Reference\ ortho\ map\ (AU)}{Energy\ used\ creating\ the\ map\ (KWh)}$$

$$EE_{truck\ radio} = \frac{period\ full\ operational\ (AU)}{Energy\ used\ during\ period\ (KWh)}$$

We have two different parameters because the state-of-the-art systems for elaborating orthomap provide also radio services. In this way we can focus the improvement only on the first one over the project because that is main goal of OPERA.

1.2 MEASUREMENTS

1.2.1 Baseline

The baseline configuration is the infrastructure used before OPERA Project, as described in [1], that consists of:

- (2) HP ProLiant DL380 G5
- (1) Mac book Pro 15
- PhotoScan Pro as elaboration software

In [1], we measured the energy consumed by baseline configuration to complete the orthophoto elaboration: 4200 Wh. Thanks to this last value is possible to measure the Baseline $EE_{Orthomap}$:

Baseline $EE_{Orthomap}$ = 1/(4,2 kWh) = 0,23 $kW^{-1}h^{-1}$

The original infrastructure hosted not only the orthophoto elaboration service but also other services (mainly radio communications), for this reason, in [3] we defined a further EE (Baseline $EE_{Truck\ Radio}$), to distinguish the two parameters because OPERA project intervenes to improve only orthophoto elaboration:

$$\text{Baseline } EE_{Truck\ Radio} = 11h / (11h * 2*0{,}18\ kW) = 2{,}77\ kW^{-1}$$

It's important to highlight that the second value is multiplied by 2 because we involve two servers to provide radio services, as reported in [1].

During OPERA project, we set up different configurations to elaborate the orthophoto, using different hardware and software as reported in [3]e [4]. The Baseline $EE_{Orthomap}$ represents the comparison element to evaluate the results that we obtained during the project for each different configuration. In the last chapter (Results & Conclusions), we'll compared this value (Baseline $EE_{Orthomap}$) with the OPERA configuration that achieved the best results in terms of EE, highlighting the improvement gained during the project and the technological element developed in OPERA.

1.2.2 MICMAC on low power server

At the beginning of the project, we selected MICMAC, an open source software, to elaborate the orthophoto because it satisfied the requirements in terms of quality (it has the same ground resolution as PhotoScanPro) and of portability (its code is open source and written in a language compatible with FPGA card).

For this reason, we set up the configuration reported in the following figure and described in [3]:

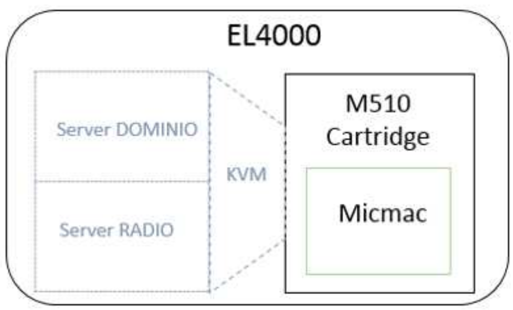

Figure 1 - MICMAC on low power server

We installed two virtual machines to host the other services and MICMAC on the cartridge m510 and we obtained these values:

$$\text{MICMAC } EE_{Orthomap} = 1/(2{,}22\ kWh) = 0{,}45\ kW^{-1}h^{-1}$$

$$\text{MICMAC } EE_{Truck\ Radio} = 11h / (11h * 2*0{,}025\ kW) = 20\ kW^{-1}$$

1.2.3 PhotoScan Pro on low power server without and with GPU

During the second cycle, as declared in [4] and [5], the MICMAC porting on FPGA card is more difficult than we anticipated, because it is necessary to complete a proper audit and analysis, that require a large effort, due to the fact that we encountered technical issues with the software.

For this reason, to guarantee that OPERA project can introduce a heterogeneous architecture to processing the orthophoto, we defined a new strategy. In particular, we substituted MICMAC with PhotoScan Pro and the FPGA card with a GPU card (NVidia P4), as described in [4]:

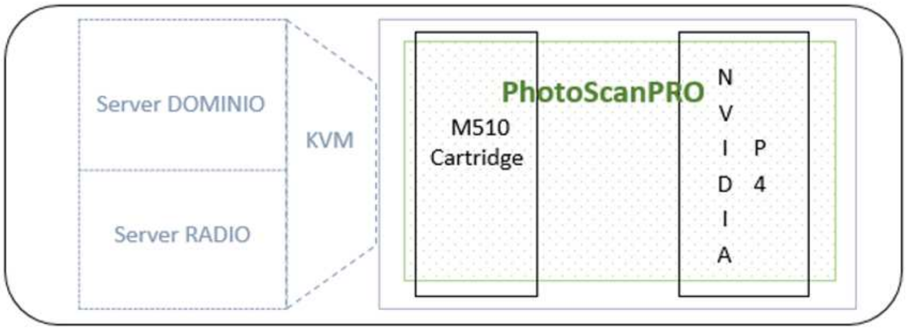

Figure 2 – PhotoScan Pro on low power server with GPU

About this configuration we evaluated the energy efficiency in two different conditions, the first one without the GPU card:

$$\text{PhotoScanPRO } EE_{orthomap} = 1/(1{,}03\ kWh) = 0.97\ kW^{-1}h^{-1}$$

$$\text{PhotoScanPRO } EE_{truck\ radio} = 11h\ /\ (11h * 2*0{,}025\ kW) = kW^{-1}$$

and the second one with the GPU card:

$$\text{PhotoScanPRO + GPU } EE_{orthomap} = 1/(0{,}87\ kWh) = 1.14\ kW^{-1}h^{-1}$$

$$\text{PhotoScanPRO + GPU } EE_{truck\ radio} = 11h\ /\ (11h * 2*0{,}025\ kW) = 20\ kW^{-1}$$

1.3 ENERGY EFFICIENCY TREND

In this section we summarize the outcomes in terms of energy efficiency parameters and we show them graphically.

In the following table, we report the values described in the previous paragraph:

Configuration	$EE_{orthomap}$ ($kW^{-1}h^{-1}$)	EE_{radio} (kW^{-1})
Baseline	0,23	2,77
MICMAC on low power	0,45	20
PhotoScan Pro on low power	0,97	20
PhotoScan Pro on low power + GPU	1,14	20

Table 1 - Truck Use Case – EE trend

The EE graphical trends are shown in the following figures, where we can see that during the project we performed a linear improvement regarding the orthophoto elaboration. Instead, about the other services, after the first improvements adopting virtual machines we have not intervened further:

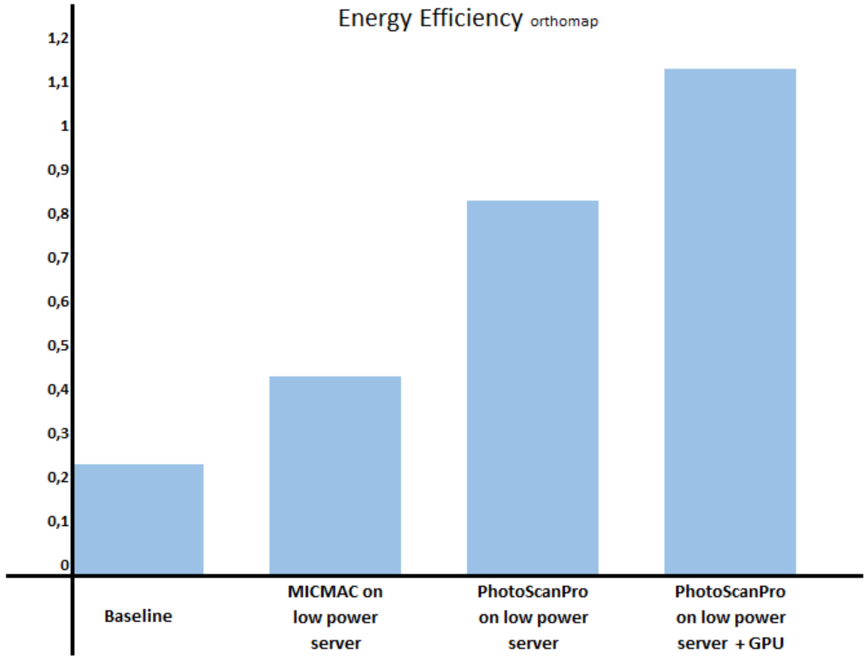

Figure 3 – Truck Use Case – EE orthomap trend

We can see that during the project we improved the orthophoto elaboration gradually, the value improved by a factor of 5 times if we compare the baseline and the last configuration.

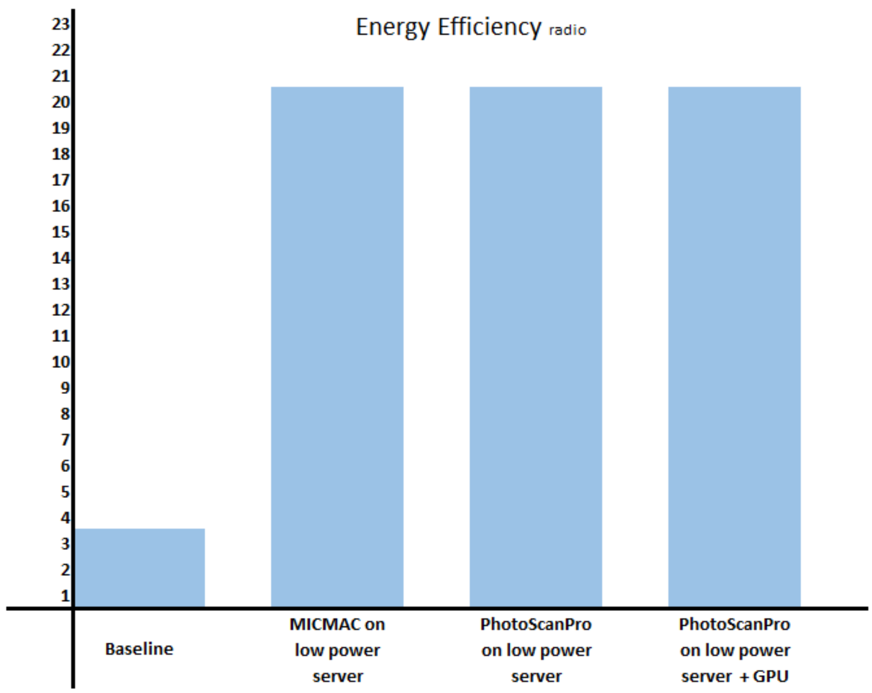

Figure 4 – Truck Use Case – EE radio trend

With the use of low power server, EE_{radio} improved 7x. As it's not an OPERA target, we didn't intervene on that, as written before, but we limited our intervention to moving these services to virtual machines, with the aim to isolate the first energy efficiency parameter that is the core for this use case.

2 III° CYCLE DESCRIPTION

In this chapter, we describe the activities carried out during the last cycle taking into account the change in strategy on heterogeneous architecture reported in [4], where we declared to focus our effort on two topics:

- refining the MICMAC porting
- introducing the configuration PhotoScan Pro on low power server + GPU card

In addition, we also report the information about the demonstrator.

2.1 MICMAC PORTING

As declared in [4], HPE and Nallatech worked during this cycle to demonstrate that an FPGA card can contribute to both reducing power consumption and enhancing the processing performance to obtain an orthophoto.

Specifically, they worked on MICMAC, the Open Source product developed by IGN for orthophoto elaboration. CSI selected this software because fitted both with the requirements declared by Nallatech:

- It's written in C/C++
- It's compatible with OpenCL

and the Protezione Civile requirement:

- it guarantees a good ground accuracy (1 cm)

After a deep analysis about the MICMAC code, as described in [5], some issues were found:

- inefficient and non-scalable parallelization method
- huge source code
- lack of comments

Due to these issues, it's very hard to adapt for offloading to an accelerator such as an FPGA, for these reasons and after profiling the software, Nallatech and HPE focused their effort on TAPIOCA, that is the module where it spends more of the time to complete the task, and specifically on ANN algorithm.

In [5], Nallatech and HPE estimated 4-thread version of the FPGA ANN implementation is predicted to be delivering a better efficiency per second versus the 32 threads version of the CPU (M510 Cartridge), as shown in the following figure:

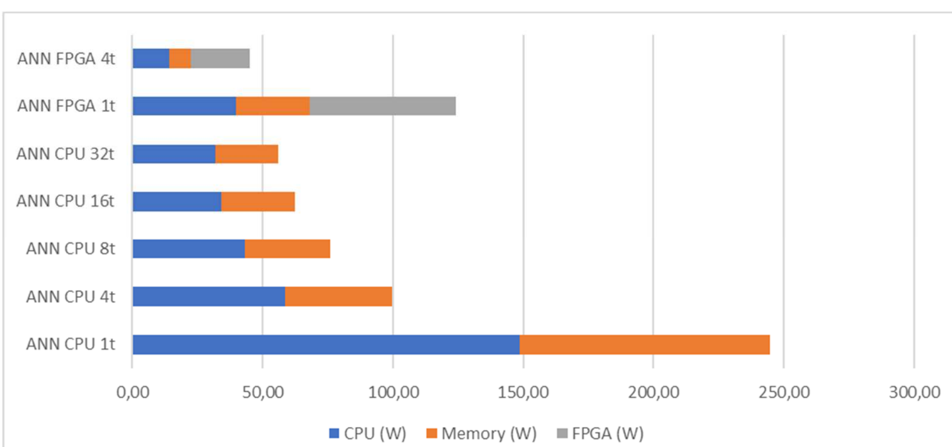

Figure 5 - ANN energy efficiency per thread

The code porting introduced a remarkable x3 speedup and a drop of power consumption estimated at 15%.

It is possible to find more details about this activity and its measurements in [5].

2.2 REAL-LIFE SERVICE

In the previous cycle, we demonstrated that the heterogeneous architecture (low power server + GPU) required less energy, processing time and occupied space than the SOTA solution to complete the same orthophoto, as described in [4].

For this reason, during this last phase of the project the truck operator replaced the SOTA solution with OPERA solution. To achieve this target, he needed to consolidate and to test properly not only the orthophoto elaboration but also the other services hosted in the two virtual machines (Server Radio and Server Dominio):

- Radio services
- Items monitoring
- Radios Monitoring

There has not yet been an opportunity to use the new services in real-life situation, but the operator completed about 40 tests about radio service and monitoring and two completed cycles to produce the orthophoto, that means to take pictures from the drone and elaborate the set of photos with the heterogeneous architecture. The outcomes of these test were all positive.

After this additional experience, the operator had a very good opinion about the new infrastructure, he appreciated directly the lower power consumption (5 times). In addition, we have also a specific value about the saved occupied space, we saved 3 rack unit as shown in the two following figures:

Figure 6 – Space occupied before OPERA

Figure 7 – Space occupied with OPERA

In the last figure, it's highlighted the saved space thanks to OPERA, because another item (out of the scope) was in maintenance.

But most of all, the operator appreciates having halved the orthophoto processing time as reported in [4].

Thanks to these three elements, Protezione Civile and CSI are setting up a new truck smaller than the current one, in this way it will be possible to intervene faster (the new vehicle is not so heavy and bulky) providing the same service (orthophoto elaboration and radios).

2.3 DEMONSTRATOR

Today, Protezione Civile has a single vehicle which needs to be always ready to intervene in case of disasters. For this reason, the demonstrator is hosted within HPE premises located in Grenoble.

HPE set up a laboratory to show the main outcome realized during the project for this use case:

- MICMAC porting, running ANN offloaded to FPGA to elaborate on a limited cluster of photos as reported in [5]
- PhotoScan Pro full elaboration on low power server and GPU, based on the full cluster of images

3 RESULTS & CONCLUSIONS

At the end of the project, we demonstrate the feasibility of what we declared in the project proposal, because, despite some issues, we realized not only a partial code porting of MICMAC on FPGA card, but also that thanks to that it's possible to save energy and time using heterogeneous architecture as reported in [5].

We didn't complete the porting, because it requires a lot of effort and the realization of a product (an FPGA card with MICMAC on board) that isn't the aim of the project. In any case, we gained experience by that, specifically it's very important to consider the effort to analyse and to intervene on the code, because at the beginning of the project we underestimated it.

In addition, to provide a further evidence that the heterogeneous architecture consumes less energy suppling the same services, we set up an alternative solution (low power server + GPU card) that brought such good results that CSI changed the SOTA solution with this one.

At the beginning of the project, we declared in [1] that we wanted to complete orthophoto elaboration within 2 hours and to consuming 0,350 kWh. Even if we didn't achieve these targets, because we need 6 hours and 0,870 kWh, we can consider Truck use case outcomes as a whole positively for two reasons. Firstly, we demonstrated the heterogeneous architecture feasibility and effectiveness. Secondly, if we compare the situation before OPERA and after OPERA, we can obtain the same outcome reducing energy consumption, time and space as shown in the following table:

Parameter	SOTA	OPERA
Occupied space	4 U	1 U
Energy consumption	4.2 kWh	0,87 kWh
Processing time	11 h	6 h

Table 2 – OPERA improvements

The table is taken from [4] to report the differences between state-of-the-art and OPERA solution.

In the following paragraphs, we compare the best OPERA configuration developed for Truck UC with two aspects:

1. The baseline established at the beginning of the project
2. The OPERA target to reduce 100 times the energy consumption of servers as compared to the state of the art in 2013

Finally, we reported in paragraph 3.3 what we demonstrated during the project and which OPERA features are present in the best OPERA configuration for this UC.

3.1 BASELINE COMPARISON

In this paragraph, we compare the baseline described in section 1.2.1 and the best OPERA configuration. To define which is the best OPERA configuration we can use the *Table 1 - Truck Use Case – EE trend*, selecting the solution that achieve the higher $EE_{orthomap}$, from these values we can declare that the best results using the configuration that involves low power server, GPU card and PhotoScan PRO. Specifically, we improved 5 times the $EE_{orthomap}$, as we can see graphically from the following figure

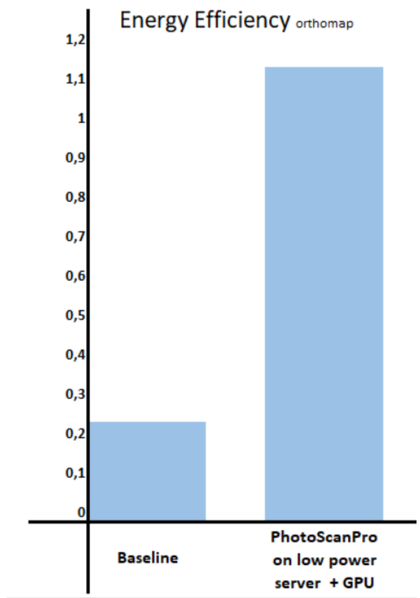

Figure 8 – Truck baseline comparison

3.2 2013 SOTA COMPARISON

In this paragraph we compare the energy consumed by the best OPERA configuration servers with the state of the art in 2013. The original target was to reduce energy consumption 100 times.

To do that we can compare the energy consumption of baseline configuration with the best OPERA configuration (low power server + PhotoScanPRO + GPU), these values are reported in chapter [1]:

Baseline energy consumption = 4,2 kWh

Low power server + PhotoScanPRO + GPU = 0,87 kWh

These values are strictly related to EE so we have the same improvement: 5 times.

At the end of the project, we can see that we are not close to the original target, the main reason is the difficulties about MICMAC porting on FPGA card, because only thanks one or more additional card, properly programmed for this specific operation, it would be possible to reduce the gap.

3.3 OPERA OUTCOMES

During OPERA project, thanks to WP6, we demonstrated the feasibility of MICMAC porting on FPGA card, but due to the issues described before in the best OPERA configuration we don't have this feature but other ones developed within this initiative. Specifically, we achieved these results not only adopting items and technologies available on the market, but also thanks to specific features realized in OPERA. This result has been achieved by adopting state-of-the-art low power servers, which are leveraging high efficiency components. These dense servers are hosted in a EL4000 enclosure, which was only on the drawing board while OPERA officially started. HPE gathered feedback from CSI on the Truck use-case to physically adapt the end product. Depth has been reduced by 10 centimetres and the way Moonshot cartridges are mounted has also been redesigned to better match deployments in space constrained environments (from top loaded to side mount).

OPERA project also required modifications on the power monitoring of the device. A firmware has been developed to enable better granularity in the recording of events. By default, HPE only support a record every 5 seconds to avoid premature failures of the NAND flash associated to the management controller (ILO). This has been set as a customized value for Redfish monitoring, with a minimal rate of 1 second.

REFERENCES

[1] D2.1 *Use Cases and Requirements 1*
[2] D4.1 *Energy efficiency metrics*
[3] D7.4 *Truck Use Case 1*
[4] D7.5 *Truck Use Case 2*
[5] D6.8 - FPGA Integration in Low Power Server