Deploying a Smart Queuing System on Edge with Intel OpenVINO Toolkit

Rishit Dagli

Thakur International School

Süleyman Eken (✉ suleyman.eken@kocaeli.edu.tr)

Kocaeli University <https://orcid.org/0000-0001-9488-908X>

Research Article

Keywords: Smart queuing system, edge computing, edge AI, soft computing, optimization, Intel OpenVINO

Posted Date: May 17th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-509460/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Deploying a Smart Queuing System on Edge with Intel OpenVINO Toolkit

Rishit Dagli · Süleyman Eken*

Received: date / Accepted: date

Abstract Recent increases in computational power and the development of specialized architecture led to the possibility to perform machine learning, especially inference, on the edge. OpenVINO is a toolkit based on Convolutional Neural Networks that facilitates fast-track development of computer vision algorithms and deep learning neural networks into vision applications, and enables their easy heterogeneous execution across hardware platforms. A smart queue management can be the key to the success of any sector. In this paper, we focus on edge deployments to make the Smart Queuing System (SQS) accessible by all also providing ability to run it on cheap devices. This gives it the ability to run the queuing system deep learning algorithms on pre-existing computers which a retail store, public transportation facility or a factory may already possess thus considerably reducing the cost of deployment of such a system. SQS demonstrates how to create a video AI solution on the edge. We validate our results by testing it on multiple edge devices namely CPU, Integrated Edge Graphic Processing Unit (iGPU), Vision Processing Unit (VPU) and Field Programmable Gate Arrays (FPGAs). Experimental results show that deploying a SQS on edge is very promising.

Keywords Smart queuing system · edge computing · edge AI · soft computing · optimization · Intel OpenVINO

R. Dagli
Thakur International School Mumbai, India
E-mail: rishit.dagli@gmail.com

*Corresponding author
S. Eken
Kocaeli University, Department of Information Systems Engineering
E-mail: suleyman.eken@kocaeli.edu.tr

1 Introduction

We can define a queue as an arrangement of people or vehicles waiting in line for their turn to get a service or move forward in an activity while queuing is the act of taking place in such an arrangement [1]. In daily activities, queuing always occurred whenever people have interest in mutual service at a same particular period and there are limited capabilities to satisfy or provide service to all the interested individual at once. For example, a queue of people at ticket windows or vehicles waiting in line at the toll. The reason for queuing is to accomplish or get the service intended in a fair and organized manner. Queue management systems (QMS) are the go-to solution for handling queues, allowing for easy management and a streamlined experience while reducing wait times and increasing efficiency.

Often in retail stores or transportation queues people end up wasting a lot of time, due to improper or manual queue management. People can lose time in transportation because of unwanted waiting or wait long behind a billing counter. Retail stores often having a lot of customers makes it further difficult to efficiently manage queues in these scenarios increasing the wait time drastically for a customer creating a bad experience. The queue management ideas could also be helpful in transportation scenarios and manage queues. This would help reduce waiting times in public transportation too.

Apart of the above-mentioned scenarios the ideas of using queue management could particularly be also helpful in allotting workers efficiently in a factory environment. This could help to ensure more efficient manufacturing in factories. Queue management strategies would further be helpful in many such scenarios (see Appendix A for use case requirements).

The edge system we propose also reduces the power consumption by a great amount as it does not have to send the video streaming data to a cloud-based server greatly saving on the power costs. Since the proposed system no longer requires to send or receive data from a server a cloud-based system, we further make this system easily accessible and implementable as all the processing is done on the device itself. This also saves costs of maintaining a server and ensuring up-time. The proposed system performing the computation on the device itself also provides a lower latency as the algorithms are now being run on the device itself [2]. This is particularly helpful on a retail or transportation scenario where low latency in the application would be required.

Thus, we propose this system which is easily accessible and implementable, requires very low deployment costs, consumes lower power, does not have a huge impact of network and also gives a lower latency. Further this SQS can also run on existing hardware with little or no changes. However, to make the model updates to facilitate better inferences from real-time data, the system does use network. We do this not by sending the complete video stream or data but instead uploading a locally trained small piece of the model if and when network becomes available. We do so to ensure privacy by not sending the video stream itself but by sending a transformed version of it i.e. a small sized locally trained model. This is then sent to a managed server which uses this

to improve on or update the model making it better at performing inferences. This considerably reduces the impact on network too. This updated model is then sent back to the edge devices if network is available and the latest version of model is deployed [3].

The remainder of this article is organized as follows. In Section 2, related works are given. In Section 3, proposed SQS on edge is presented. In Section 4, we compare deep networks for smart queuing of different videos and provide performance analysis of various inference implementations. The last section concludes the article and gives future works.

2 Related Works

Nowadays, companies rely a lot on computer vision, especially for established use cases such as detection, classification, and recognition with video surveillance being one of the applications. Companies working on developing these systems look for easy to use and deploy frameworks to facilitate the transition to a final product they can offer their clients. One such solution is the Intel OpenVINO Toolkit [4] which leverages image processing, machine learning and deep learning to solve computer vision problems while allowing cross-platform deployment. It comes with a variety of built-in and optimized algorithms for the wide range of Intel hardware (CPUs, FPGAs, Movidius Neural Compute Stick, and etc.). One of its key strengths is allowing developers to use the same application programming interface and execute their code heterogeneously on Intel accelerators. Deploying SQS on edge with Intel OpenVINO Toolkit has important use cases requiring fast deep learning model inference. So, we firstly give queue and related terms and technologies then present real-time object detector for deployment on edge in next paragraphs.

There are two main parameters of queue management, which are the number of channels (or servers) and the number of phases of service. Each parameter can take two values: single (one), or multi (several). Different combinations of channels and phases give four distinct types of queue management: (i) Single-channel, single-phase (e.g. an automated car wash), (ii) single-channel, multi-phase (e.g. retail banking), (iii) multi-channel, single-phase (e.g. airline ticket counter), and (iv) multi-channel, multi-phase (e.g. a laundromat with several washers and dryers) [5]. There are also different types of queues. Ranging from structured queues with fixed and predetermined locations, such as at supermarket checkouts, to unstructured ones with no specific rules for position arrangement. Kiosk based queuing systems where clients use a terminal to select a service upon arrival usually appear in environments where medical and banking services are offered, as well as in telecommunication companies and government institutions. To eliminate or at least reduce waiting time and efficiency, modern companies also implement mobile-based queuing systems, customers can use their devices to check the current status of the queue and join it at their convenience only showing up when it is their turn [6]. Each has its own advantages and drawbacks. Moreover, to manage a queue, we

need to understand the queue discipline (the order of servicing: First in - first out (FIFO), Last in - first out (LIFO), Service in random order (SIRO), and priority selection [7].

QMS are critical components in any sector of business. It is important for the service provider to provide efficient queuing system to maintain high customer satisfaction. Besides, QMS can help record, predict, and calculate the statistic of the queuing pattern such as customers' arrival rate, queuing behavioral over time, and average waiting time which can help in decision making. Queuing systems, even smart ones, whereby a client gets a digital ticket from a machine or online and waits for a turn, face many limitations in terms of creating an improved user experience [8]. There are also many reasons to use a QMS. A proper queue management strategy shortens on-site wait times and reduces walkaways. By allowing customers to use their time in a more efficient manner, a QMS builds up customer experience. By improving customer engagement, a QMS turns visitors into customers, and customers into promoters. The QMS is a set of tools and sub-systems for multiple industries including banking/insurance [9], healthcare [10], retails [11], transportation [12], manufacturing [13], government, and telecommunication [14]. Ibrahim et al. [9] employed factor analysis procedure to identify the underlying structure among the explored electronic service quality attributes. Good queue management is one the composite dimensions of electronic service quality that the UK customers' perceptions of their bank actual performance. Rezaee et al. [10] presented a congestion management protocol for healthcare wireless sensor networks (HWSNs). Their protocol includes active queue management scheme to separate virtual queues on a single physical queue to store the input packets from each child node. Berman and Larson [11] modeled facility having workers who are cross-trained to do front room and back room operations in retail sector. Switching from back room to front room is done using queues. Schwarz and Epp [12] proposed a new discrete-time approach for the steady-state analysis of circulating vertical conveyor systems bulk service queues. Rüttiman and Stöckli [13] distinguished the classic traditional batch and queue manufacturing for optimal scheduling sequences. In the light of these studies, it is possible to say that having good queue management procedures in place can make a real difference to any customer-facing environment. COVID-19 has really highlighted the need for effective queue management.

Many computer vision applications rely heavily on Object detection (OD) and object recognition (OR), which creates a need for well-optimized algorithms. However, it is common for object detection CNNs to face some hurdles while moving from development to deployment unless these difficulties are accounted for and addressed as pointed out by Kozlov and Osokin [15]. In their work, they follow up by using a Single Shot MultiBox Object Detection (SSD) to detect vehicles and pedestrians with Intel OpenVINO Toolkit. Osokin [16] also worked on multi-person pose estimation architecture on edge devices and optimized the popular method OpenPose [17]. A new improved framework for neural networks was proposed in the work by Kozlov et al. [18]. Its success is thanks to the use of modern network compression methods, such as sparsity,

quantization, and binarization. Using these techniques results in lighter models that can be run on cheap general-purpose computers. Currently, a PyTorch version of their framework Neural Network Compression Framework (NNCF) is available as a part of OpenVINO Training Extensions ¹. Targeting OpenVINO as well, Kustikova et al. implement a solution for semantic segmentation of on-road images [19]. and another to benchmark inference performance of deep learning models on various types of hardware [20] (CPUs, integrated graphics and embedded devices) ^{2,3}. Similarly, Kljucaric and George [21] used Chinese character recognition to evaluate the performance of current CNN architectures and frameworks on different hardware types (FPGAs, GPUs, and CPUs). Mathew et al. [22] as a continuation of the work on DetectNet Deep Learning Model for lung nodule detection trained on the LIDC dataset, they ported it from NVIDIA to Intel OpenVINO. Castro-Zunti et al. [23] modified a single shot detector (SSD) and used it for license plate segmentation and recognition on an embedded system while providing plate processing time comparisons between a variety of hardware types. regarding performance, Gorbachev et al. [24] suggest a new software solution for holistically analyze model performance on various hardware. Mittal and Bhushan [25] propose an efficient workflow for OpenVINOs. While Jin and Finkel [26] evaluate the performance of 14 deep learning inference models using OpenVINO. When these studies in the literature are examined, it is important to implement an image-based smart queue system. So, we propose an OpenVINO-based SQS, which accelerates the processing speed of the model and meets the real-time applications.

Contributions to the literature with the paper can be listed as follows:

- We propose an efficient and smart way to perform queue management for different use-cases. We majorly focus on deploying the SQS on an edge device.
- We use Power Proxying that can proxy a variety of network protocols, increasing the standby time of the host without compromising its active connections.
- We also propose a way to update the model making the model better as we get more and more real-life data without leveraging security or network. We provide ways with which apart of queuing, complex analysis and finding out trends for consumer patterns in a retail store could also be implemented. This would enable owners to derive insights based on consumer patterns such as where should an advertisement or the billing counter be placed for efficient management of crowd.
- We validate the above-mentioned ideas and claims by testing them on multiple edge devices with Intel’s OpenVINO Toolkit to perform optimization on the model for it to be able to run on edge and deploying it to multiple hardware and edge machines easily. Further we also profile our application and find hot-spots in it with OpenVINO Toolkit.

¹ https://github.com/openvc/openvino_training_extensions/tree/develop/pytorch_Toolkit/nncf

² <https://github.com/itlab-vision/openvino-dl-benchmark>

³ <http://hpc-education.unn.ru/dli>

3 Proposed Architecture

Proposed SQS demonstrates how to create a video AI solution on the edge. The system detects people in a specified area and accordingly detects the number of people in a queue. It would then also notify whether a person would need to change the queue to reduce congestion. If there are more than a certain number of people in the watched area, a warning is given as the queue is full and move to the next queue. In this paper we use various SSD models [27] to detect people in an image using a single deep neural network. This eliminates proposal generation and subsequent pixel or feature re-sampling stage and encapsulates all computation in a single network. SSD has comparable accuracy to methods that utilize an additional object proposal step and is much faster, while providing a unified framework for both training and inference.

We choose multiple models from Intel’s OpenVINO Model Zoo [28] and TensorFlow Model Garden [29]. We test these pre-trained models for use-cases on basis of model load time, inference time on one frame of the video and FPS (Frames Per Second) of the inference stream. We also perform a preliminary comparison of the models based on theoretical parameters specifically billions floating-point operations (GFLOPs) and amount of memory consumption before performing analysis on practical parameters mentioned above. We majorly perform the theoretical analysis to make sure that the edge device would not run out of memory and the model should occupy disk space such that it can be used with a particular edge device.

We perform various optimizations over the source model from Intel’s OpenVINO Model Zoo, TensorFlow Model Garden or transfer learned models. We do these optimizations to be able to run these models on edge devices which have limitations on processing, memory, power-consumption, network usage, and model storage space. While performing the optimizations due to the limitations of edge devices on which the SQS would majorly be deployed, we majorly care about model load time, per frame inference time and model size while performing the optimizations. We describe the process we use to perform the optimization in greater detail in later parts of this paper.

We then make further comparisons between these models and make experiments for various model to find a suitable configuration with the OpenVINO DL Workbench to find a model with lower size and acceptable throughput. We also make use of VTune profiler to help us find and solve any hot-spots in the application that may be faced during run-time. This helps us improve the overall performance and throughput of the application.

To be able to deploy these models on multiple devices we run them on the Inference Engine which provides a common API to deliver inference solutions on the platform of our choice: CPUs, iGPUs, FPGAs or VPUs. To use these, we first convert these models into an acceptable IR (Intermediate Representation) format. It also enables us to use hardware specific APIs to further improve the performance of the models.

We then use Intel’s DevCloud to make and test real time deployments for multiple hardware. We also test for the primary use-cases and scenarios where

the SQS may be used namely in Retail stores, Manufacturing factories and in Public transports. We do so on multiple hardware, CPUs, GPUs (Graphic Processing Unit), FPGAs (Field Programmable Gate Arrays) and VPUs (Vision Processing Unit) to find out the best possible hardware and comparing the different hardware for each of these use cases. With the Intel DevCloud we can make these deployments without using a physical device by submitting a job.

Fig. 1 shows the main steps of the SQS. Fig. 2 illustrates the user workflow for code development, job submission and viewing results. It mainly includes three sub-steps: (i) person detection, (ii) submission of the job to multiple devices on the DevCloud, and (iii) inferencing video output.

Fig. 1: Main steps of the SQS

To demonstrate the following ideas, we perform inference some open-source video clips pertaining to the primary use-case mentioned earlier in this paper, we also run these on multiple devices. The results of these are in next section.

4 Experimental Results

4.1 Experimental setup

We first optimize the model so it can be run on edge devices. This is necessary to reduce cost, latency, especially for deployment on hardware with limited capabilities, be it processing, memory, power consumption, networking, or storage space. The optimization also takes into account more powerful

Fig. 2: SQS run life-cycle

and specialized hardware with accelerators [30]. Hardware and software requirements are as follows:

- *Hardware:* This project makes the use of Intel DevCloud to test on CPU, GPU, FPGA and VPU. So, no specific hardware is required.
- *Software:* Intel® Distribution of OpenVINO™ Toolkit 2019 R3 release. Python > 3.5

4.1.1 Freezing models

We particularly need to employ model freezing strategies while using TensorFlow pre-trained models or on which we perform transfer learning. While using TensorFlow models all trainable parameters are represented as variables in the graph. To get the best possible model, the workflow demands the freezing of these variables between each layer, whose values usually change after each step of the training process. This is achievable through the use of TensorFlow’s functions when using transfer learning from a pre-trained model. We start by defining some necessary classes:

- *NodeDef* to hold one single value or operation from the network.
- *Checkpoint* to take a periodic snapshot of the variables as the model trains.
- *GraphDef* to define the whole executable graph including all the NodeDefs within it. During execution, the values of any node variables are accessible from the checkpoint files [31].
- *Saver* to save the variables in a checkpoint file. They are then modified to be constants using the ‘freeze_graph’ class, which gives us a new GraphDef with the constants in return.

4.1.2 Optimizing models

As mentioned earlier we first optimize the model so that it can be deployed on edge devices. In case of TensorFlow models we first freeze the models too. We use multiple techniques to do so, the majorly used ones are:

- *Operation Fusing*: Many Convolutional Neural Network (CNN) include layers that can be represented as a sequence of linear operations such as addition, multiplication or matrix multiplications. These layers can be fused into previous convolution or fully connected dense layers. However, this technique is not applicable when a convolutional layer comes after a add operation due to paddings.
- *Quantized Distillation*: By improving the training process using distilling [32]. This is achievable in the form of a distillation loss from a teacher network and inserted into a smaller student network, we quantize the weights of the latter to a limited set of levels [33]. The result is a significantly smaller model size.
- *Mixed precision Quantizing*: Here we quantize different layers with different bit-widths. For mixed precision quantization, we assume that we have the flexibility to choose different precisions for different layers of a network. Mixed precision computation is widely supported by hardware platforms such as CPUs, FPGAs, and dedicated accelerators. Thus, we choose precision for each layer, for most of the models tested we use Differentiable Neural Architecture Search [34]. This allows us to reduce the model size and also receive lower latency predictions.

4.1.3 Converting models to an intermediate representation

To use these models with Intel’s OpenVINO Toolkit they first need to be converted to an Intermediate Representation (IR) which is OpenVINO Toolkit’s format of graph representation and its own operation set. A graph here is represented with two files: an XML file and a binary file [35]. The XML file describes a network topology using <layer> tag for an operation node and <edge> tag for a data-flow connection. Each operation can also have a fixed number of attributes. Constant values like convolutional weights or biases are stored in the accompanying binary file. This Intermediate Representation can be used to run the model on edge with OpenVINO’s Deep Learning Inference Engine [36] which is a unified API to allow high performance inference on many hardware types (CPUs, GPUs, FPGAs, Neural Compute Stick, Raspberry Pi etc.) making the deployment procedure a lot easier. We can also make modifications or changes in the workflow for a specific device type with the Deep Learning Inference Engine.

After configuring the OpenVINO Model Optimizer for a particular framework, it provides with some easy to use framework specific Python scripts which one can directly use to convert a model to Intermediate Representation (IR). The OpenVINO model optimizer performs some basic model optimiza-

tion and converts it to Intermediate Representation. We also make modifications and extend the Model Optimizer with custom primitives to support more customization as mentioned earlier.

We also tested for these use-cases with some pre-trained models from the Intel OpenVINO Model Zoo. The Intel OpenVINO Model Zoo also serves as a source of finding open-source pre-trained models converted in the IR format. One can directly use models from the Model Zoo and use them with Intel OpenVINO.

4.1.4 Post training optimization

We then use post training quantization to convert the models into a more hardware-friendly representation by applying specific methods that do not require re-training of the model. Before quantization we also align ranges of output activation of convolutional layers in order to reduce the quantization error. To perform post training quantization, we use the following methods.

- *8-bit quantization*: We use this vanilla quantization method that automatically inserts FakeQuantize operations into the model graph based on the specified target hardware and initializes them from the calibration dataset. We opportunistically replace 32-bit floating point (FP32) computations with 8-bit integers (INT8) and transform the FP32 computational graph [37]. These seem to greatly improve the performance of the model. We then also adjust biases of convolutional and fully-connected layers based on the quantization error of the layer.
- *Parzen Estimator*: We then use a Tree-structured Parzen Estimator (TPE) Approach, a sequential model-based optimization method to perform global optimization of the post training quantization parameters. We do this by first defining a hyper-parameter search space and create an objective function which takes in hyper-parameters and outputs a score that it needs to maximize, in this case the score we used was the appropriate queue detection. We then get some observations which are randomly selected from the set of hyper-parameters. We then make two quantiles x_1 and x_2 where we divide all the good scores into one quantile. We then model these using the Parzen Estimator which is the average of the kernels centered on this data. We then draw hyper parameters and evaluate them in terms of both quantile densities. These hyper parameters are then evaluated back on the objective function [38].
- *Per-channel quantization*: We also make use of Per-channel quantization for Convolutional and Fully-Connected layers. Our experiments show that per-channel quantization is needed to compensate for the accuracy drop resulting from quantization, asymmetric per-layer quantization seems to work best for this, while at the same time, it works as a good baseline for post-training quantization of weights and activations. This results in an accuracy that is very close to what floating-point offers for all networks [39]. We perform asymmetric quantization according to the following formula:

$$\begin{aligned} \text{input low}' &= \min(\text{input low}, 0) \\ \text{input high}' &= \max(\text{input high}, 0) \end{aligned}$$

So,

$$\begin{aligned} \text{input high}'' &= \frac{ZP - \text{levels} + 1}{ZP} \times \text{input low}' \\ \text{input low}'' &= \frac{ZP}{ZP - \text{levels} + 1} \times \text{input high}' \end{aligned}$$

Thus, we use the following process for post-training quantization to speed up the model inference times drastically and reduce the model sizes by a great amount. This particularly helps in ensuring that retail stores or factories can use this with the devices already in use by them. Furthermore, as we will mention we have also tried these algorithms on specialized embedded systems too.

4.1.5 Comparing models with Deep Learning Workbench

The OpenVINO Deep Learning Workbench provides interactive user interface (see Fig. 3), simplified process of 8-bit quantization, speeding up convolutional operations using the Winograds minimal filtering algorithms, measuring accuracy of the resulting model [24, 40]. We first import the OpenVINO IR model, select the target hardware available for inference. This then allows us to analyze the theoretical information about the model like GFLOPs, amount of memory consumption and the inference times.

The screenshot shows the OpenVINO Deep Learning Workbench interface. At the top, there's a header with the OpenVINO logo and version information: "Version: 1.0.2205.782d052e". Below the header is a table with columns: Dataset Name, Date, Size, Status, Type, Action, and Select. The table contains two rows: "imagenet100" and "VOC_small". Both rows show a status of "Success" (green checkmark) and a size of "13 Mb" and "8 Mb" respectively. The "Action" column has a trash icon for each row. The "Select" column has radio buttons. At the bottom, there are two buttons: "Import Local Dataset" and "Autogenerate Dataset in ImageNet format".

Dataset Name	Date	Size	Status	Type	Action	Select
imagenet100	05/12/19, 05:46	13 Mb	Success	ImageNet		☐
VOC_small	05/12/19, 05:47	8 Mb	Success	VOC Object Detection		☐

Buttons at the bottom: Import Local Dataset, Autogenerate Dataset in ImageNet format

Fig. 3: Testing multiple models with DL Workbench

We perform various such experiments to find a suitable configuration that can provide the highest throughput with minimal latency. As mentioned earlier we also perform some post training quantization, the Deep Learning Workbench allows us to compare the performance of these quantized models with the original model. We compare the original and quantized model on multiple hardware, CPU, GPU, Intel Neural Compute Stick (Myriad VPU) and FPGAs too. We then select the best configuration for these models and implement them.

Thus, we compare some models: Transfer learned from existing state of the art models (specifically from MobileNet [41, 42], Inception [43] and ResNet [44]), models from the OpenVINO Model Zoo and TensorFlow Model Garden in this manner based on various theoretical parameters such as GFLOPS or memory consumption and different hardware (see Fig. 4). We also perform the comparisons between the model on practical factors like inference time, FPS and model loading time as we mention later in this paper.

Fig. 4: Testing Models with DL Workbench

4.2 Identifying Hotspots in the Application

Analyzing the end application and finding areas in the application that particularly have a negative impact towards the final application, increasing latency or consuming a lot of power is mandatory to put forth a better application to the users. We also want to optimize the code as much as possible to run on edge infrastructure considering the power, memory, network and inference time limitations. We use Intel VTune Amplifier to perform these tasks⁴. Considering the full end-to-end solution, it is an important aspect to optimize the hotspots and instead move towards a holistic solution. The Intel VTune Amplifier also makes it simple to do so and allows us to explore the parts of our end to end application and find the hot spots in it. Finding these hot spots in the application proves to be useful in optimizing the code.

It is also another important aspect to build a solution considering that a problem may not be due to the software implementation but inappropriate allocation of hardware resources or understanding how the underlying hardware works thus building an overall holistic solution giving better performance. For example, if the use scenario is transcoding an efficient Gen GPU implementation would set up concurrent encodes and decodes. On Intel processor graphics hardware, transcode is faster than running encode or decode separately. With

⁴ <https://software.intel.com/sites/default/files/1d/13/8486>

transcode being the overhead of additional synchronization and moving raw frames to/from the CPU can be avoided (as demonstrated in the Intel VTune Whitepaper [45]). Using the VTune amplifier and analyzing the application from it, allowing us to derive insights from it. We get performance information from all compute engines in one dashboard interface and aligned on a common timeline. This allows us to see all the components working together and allow for easier use.

4.3 Testing the Application for Production

We need to test the application and the models for performance in various use cases and hardware. We use a few open-sourced video clips pertaining to the major use cases mentioned for demonstration purposes. We also make use of Intel DevCloud which is a cloud service designed to help developers' prototype and experiment with computer vision applications using the Intel Distribution of OpenVINO Toolkit. It also allows us to test on latest hardware and software too and easily explore real-world workloads [46].

The Intel DevCloud allows us to test on CPUs, GPUs, FPGAs and VPUs, we test on these hardware for the mentioned use cases. While testing out the models we submit it as a job to the Intel DevCloud which then allows us to use the hardware and test out the models. For demonstration purposes as we show here, we test the models on:

- CPU - IEI Tank 870-Q170 edge node with an Intel Core™ i5-6500TE processor
- CPU and integrated GPU (iGPU) – IEI Tank 870-Q170 edge node with an Intel Core™ i5-6500TE processor with HD Graphics 530 integrated GPU
- VPU - IEI Tank 870-Q170 edge node with an Intel Core™ i5-6500TE processor with Intel Neural Compute Stick 2
- FPGA - IEI Tank 870-Q170 edge node with an Intel Core™ i5-6500TE processor with IEI Mustang-F100-A10 FPGA card

Fig. 5 show the comparisons among hardware with Intel DevCloud for the three use cases mentioned earlier for demonstration purposes.

As shown in Fig. 5, the inference time graph FPGAs take the least amount of time to perform inferences. The client needs inference to be performed fast so this FPGA can do a commendable job. We can see that FPGA also gives the highest FPS. The client requires 25 - 30 FPS which can be addressed by an FPGA. Further they are also field programmable and have a high life which is requested by the client. Thus, FPGA would be a good choice for manufacturing use case. Fig. 6 shows smart queuing output on manufacturing video.

CPU has a comparatively lower inference time than a GPU and VPU which is crucial for retail scenario. A CPU can also make inferences on a good enough FPS for a retail store. It would also help save costs and electricity bill

Fig. 5: Comparing models on different hardware

as requested by the client. Thus, CPU would be a good choice for retail use case. Fig. 7 shows smart queuing output on retail video.

VPU provides a very high inference time and inference on very low FPS. A CPU or GPU would ideally be good for this if the budget and power scenario would be ignored. Since the client requires these aspects VPU would be a good choice for transportation use case. Fig. 8 shows smart queuing output on transportation video.

Fig. 6: Output for manufacturing use case

Fig. 7: Output for retail use case

Fig. 8: Output for transportation use case

5 Conclusion

With the proposed system, smart queuing, has been developed by inferencing deep models on various devices (CPUs, integrated GPUs, embedded devices). An open source implementation of OpenVINO Toolkit based SQS is available on GitHub. To demonstrate creating a video AI solution on the edge, we perform inference some open-source video clips pertaining to the use-cases such as retail, manufacturing, and public transportation. Through the use case and relevant client requirements, we suggest a potential hardware type and explain how this hardware would satisfy each of the requirements.

In the future, we plan to increase the number of software and hardware platforms. We will analyse performance for the inference implementations based on the well-known deep learning frameworks. Further we will compare inference performance using various deep learning tools.

Funding

The author received no financial support for the research, authorship, and/or publication of this article.

Conflict of interest

The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Availability of data and material

The codes and input videos are available at: <https://github.com/Rishit-dagli/Smart-Queueing-System-On-Edge>

A video of the demo is available at: <https://www.youtube.com/watch?v=W-DWBYhrwj0>

Authors' contributions

R.D. conceived of the presented idea and performed the computations. S.E. encouraged R.D. to investigate smart queue on edge and supervised the findings of this work. Both of authors discussed the results and contributed to the final manuscript.

References

- [1] Hyun Lee. *The Definitive Guide to Queue Management Systems*. 2019. URL: <https://www.qminder.com/what-is-queue-management-system>.
- [2] Jiao Zhang et al. "Energy-latency tradeoff for energy-aware offloading in mobile edge computing networks". In: *IEEE Internet of Things Journal* 5.4 (2017), pp. 2633–2645.
- [3] Qiang Yang et al. "Federated machine learning: Concept and applications". In: *ACM Transactions on Intelligent Systems and Technology (TIST)* 10.2 (2019), pp. 1–19.
- [4] *Intel Distribution of OpenVINO toolkit*. URL: <https://software.intel.com/en-us/OpenVINO-toolkit>.
- [5] Richelle Adams. "Active queue management: A survey". In: *IEEE Communications Surveys & Tutorials* 15.3 (2012), pp. 1425–1476.
- [6] Neha Titarmare and Ashwini Yerlekar. "A Survey on Patient Queue Management System". In: *International Journal of Advanced Engineering, Management and Science* 4.4 (2018), p. 239985.

- [7] Ali Dada and Frédéric Thiesse. “Sensor applications in the supply chain: the example of quality-based issuing of perishables”. In: *The Internet of Things*. Springer, 2008, pp. 140–154.
- [8] Mohammed Ghazal, Rania Hamouda, and Samr Ali. “An iot smart queue management system with real-time queue tracking”. In: *2015 Fifth International Conference on e-Learning (econf)*. IEEE. 2015, pp. 257–262.
- [9] Essam E Ibrahim, Matthew Joseph, and Kevin IN Ibeh. “Customers’ perception of electronic service delivery in the UK retail banking sector”. In: *International Journal of Bank Marketing* (2006).
- [10] Abbas Ali Rezaee, Mohammad Hossein Yaghmaee, and Amir Masoud Rahmani. “Optimized congestion management protocol for healthcare wireless sensor networks”. In: *Wireless Personal Communications* 75.1 (2014), pp. 11–34.
- [11] Oded Berman and Richard C Larson. “A queueing control model for retail services having back room operations and cross-trained workers”. In: *Computers & Operations Research* 31.2 (2004), pp. 201–222.
- [12] Justus Arne Schwarz and Martin Epp. “Performance evaluation of a transportation-type bulk queue with generally distributed inter-arrival times”. In: *International Journal of Production Research* 54.20 (2016), pp. 6251–6264.
- [13] Bruno G Rüttimann and Martin T Stöckli. “From Batch & Queue to Industry 4.0-Type Manufacturing Systems: A Taxonomy of Alternative Production Models”. In: *Journal of Service Science and Management* 13.2 (2020), pp. 299–316.
- [14] Giovanni Giambene. *Queueing Theory and Telecommunications*. Springer, 2005.
- [15] Alexander Kozlov and Daniil Osokin. “Development of real-time ADAS object detector for deployment on CPU”. In: *Proceedings of SAI Intelligent Systems Conference*. Springer. 2019, pp. 740–750.
- [16] Daniil Osokin. “Real-time 2d multi-person pose estimation on CPU: Lightweight OpenPose”. In: *ArXiv Preprint arXiv:1811.12004* (2018).
- [17] Zhe Cao et al. “OpenPose: realtime multi-person 2D pose estimation using Part Affinity Fields”. In: *ArXiv Preprint arXiv:1812.08008* (2018).
- [18] Alexander Kozlov et al. “Neural network compression framework for fast model inference”. In: *ArXiv Preprint arXiv:2002.08679* (2020).
- [19] Valentina Kustikova et al. “Intel Distribution of OpenVINO Toolkit: A Case Study of Semantic Segmentation”. In: *International Conference on Analysis of Images, Social Networks and Texts*. Springer. 2019, pp. 11–23.
- [20] Valentina Kustikova et al. “DLI: Deep Learning Inference Benchmark”. In: *Russian Supercomputing Days*. Springer. 2019, pp. 542–553.
- [21] Luke Kljucaric and Alan D George. “Deep-Learning Inferencing with High-Performance Hardware Accelerators”. In: *2019 IEEE High Performance Extreme Computing Conference (HPEC)*. IEEE. 2019, pp. 1–7.
- [22] Gina Mathew, S Sindhu Ramachandran, and VS Suchithra. “Lung Nodule Detection from low dose CT scan using Optimization on Intel Xeon

- and Core processors with Intel Distribution of OpenVINO Toolkit". In: *TENCON 2019-2019 IEEE Region 10 Conference (TENCON)*. IEEE. 2019, pp. 1783–1788.
- [23] Riel D Castro-Zunti, Juan Yépez, and Seok-Bum Ko. "License plate segmentation and recognition system using deep learning and OpenVINO". In: *IET Intelligent Transport Systems* 14.2 (2020), pp. 119–126.
- [24] Yury Gorbachev et al. "OpenVINO deep learning workbench: Comprehensive analysis and tuning of neural networks inference". In: *Proceedings of the IEEE International Conference on Computer Vision Workshops*. 2019, pp. 0–0.
- [25] Varnit Mittal and Bharat Bhushan. "Accelerated Computer Vision Inference with AI on the Edge". In: *2020 IEEE 9th International Conference on Communication Systems and Network Technologies (CSNT)*. IEEE. 2020, pp. 55–60.
- [26] Zheming Jin and Hal Finkel. "Analyzing Deep Learning Model Inferences for Image Classification using OpenVINO". In: *2020 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. IEEE. 2020, pp. 908–911.
- [27] Wei Liu et al. "Ssd: Single shot multibox detector". In: *European Conference on Computer Vision*. Springer. 2016, pp. 21–37.
- [28] *OpenVINO Model Zoo*. URL: https://github.com/OpenVINOtoolkit/open_model_zoo.
- [29] *TensorFlow Model Garden*. URL: <https://github.com/tensorflow/models>.
- [30] *Optimize machine learning models*. URL: https://www.tensorflow.org/model_optimization.
- [31] Harshvardhan C Takawale and Abhishek Thakur. "Talos app: On-device machine learning using tensorflow to detect android malware". In: *2018 Fifth International Conference on Internet of Things: Systems, Management and Security*. IEEE. 2018, pp. 250–255.
- [32] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. "Distilling the knowledge in a neural network". In: *ArXiv Preprint arXiv:1503.02531* (2015).
- [33] Antonio Polino, Razvan Pascanu, and Dan Alistarh. "Model compression via distillation and quantization". In: *ArXiv Preprint arXiv:1802.05668* (2018).
- [34] Bichen Wu et al. "Mixed precision quantization of convnets via differentiable neural architecture search". In: *ArXiv Preprint arXiv:1812.00090* (2018).
- [35] *Intermediate Representation Used in OpenVINO*. URL: https://docs.openvinotoolkit.org/latest/openvino_docs_MO_DG_IR_and_opsets.html#intermediate_representation_used_in_openvino.
- [36] *Deep Learning Inference Engine Developer Guide*. URL: https://docs.openvinotoolkit.org/2020.4/openvino_docs_IE_DG_Deep_Learning_Inference_Engine_DevGuide.html.

- [37] Aishwarya Bhandare et al. “Efficient 8-bit quantization of transformer neural machine language translation model”. In: *ArXiv Preprint arXiv:1906.00532* (2019).
- [38] James S Bergstra et al. “Algorithms for hyper-parameter optimization”. In: *Advances in Neural Information Processing Systems*. 2011, pp. 2546–2554.
- [39] Raghuraman Krishnamoorthi. “Quantizing deep convolutional networks for efficient inference: A whitepaper”. In: *ArXiv Preprint arXiv:1806.08342* (2018).
- [40] Andrew Lavin and Scott Gray. “Fast algorithms for convolutional neural networks”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2016, pp. 4013–4021.
- [41] Andrew G Howard et al. “Mobilenets: Efficient convolutional neural networks for mobile vision applications”. In: *ArXiv Preprint arXiv:1704.04861* (2017).
- [42] Mark Sandler et al. “Mobilenetv2: Inverted residuals and linear bottlenecks”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2018, pp. 4510–4520.
- [43] Christian Szegedy et al. “Rethinking the inception architecture for computer vision”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2016, pp. 2818–2826.
- [44] Kaiming He et al. “Deep residual learning for image recognition”. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2016, pp. 770–778.
- [45] *Using Intel® VTune™ Amplifier to Optimize Media Video Applications*. URL: <https://software.intel.com/content/www/us/en/develop/download/white-paper-using-vtune-to-optimize-media-video-applications.html>.
- [46] *Intel DevCloud*. URL: <https://software.intel.com/content/www/us/en/develop/tools/devcloud/edge.html>.

A Appendix: Client Requirements for Use Cases

Table 1: Use case 1: Manufacturing

<i>Requirement Observed</i>	<i>How does the chosen hardware meet this requirement?</i>
The factory has a vision camera installed at every belt. Each camera records video at 30-35 FPS (Frames Per Second). The client would like the image processing task to be completed five times per second.	This aspect shows that latency is a major concern in this case. Once an FPGA is programmed with the bitstream required for this application can run the model with very high performance and give a very low latency. As an FPGA can run many sections of the chip in parallel and the ability of FPGA to not go off-chip for performing inference from the model would particularly be very useful for this kind of scenario. An FPGA also does not send the output back to the CPU using PCIe bus making the inference a lot faster.
The second issue the client has encountered is that a significant percentage of the semiconductor chips being packaged for shipping have flaws. These are not detected until the chips are used by clients. If these flaws could be detected prior to packaging, this would save money and improve the company's reputation.	To solve this issue the edge system would require to be able to deliver high performance. FPGAs would be a perfect fit for doing so as they provide a very high performance. Another aspect of FPGA which might help in this scenario would be its ability to be reprogrammed on the field, which can help improve inferences.
To be able to detect chip flaws without slowing down the packaging process, the system would need to be able to run inference on the video stream very quickly.	FPGAs can also be used as hardware accelerators speeding up the inference. The parallel processing and no need to go off chip further provide an added advantage in performing the inference faster.
Additionally, because there are multiple chip designs and new designs are created regularly the system would also need to be flexible so that it can be reprogrammed and optimized to quickly detect flaws in different chip designs.	This need of the client makes FPGA align to the required hardware. FPGAs are highly flexible, they are field programmable and can be reprogrammed as needed.
The client would ideally like it to last for at least 5-10 years.	FPGAs have a very long lifespan generally 10 years and thus could be used by the client.
Queue Monitoring Requirements	
Maximum number of people in the queue	2
Model precision chosen (FP32, FP16, or Int8)	FP16

Table 2: Use case 2: Retail

<i>Requirement Observed</i>	<i>How does the chosen hardware meet this requirement?</i>
Most of the store's checkout counters already have a modern computer, each of which has an Intel i7 core processor. Currently these processors are only used to carry out some minimal tasks that are not computationally expensive.	Since, the client already has a lot of CPUs which are not able to carry some computationally expensive tasks. These CPUs with some new ones for the expensive tasks could be used.
The client does not have much money to invest in additional hardware.	Existing CPUs with some new ones can be used by the client can be made use of to reduce the costs.
The client would like to save as much as possible on his electric bill.	CPUs could meet the hardware requirements and also help save on the client's electric bill.
Queue Monitoring Requirements	
Maximum number of people in the queue	2-5
Model precision chosen (FP32, FP16, or Int8)	FP32

Table 3: Use case 3: Transportation

<i>Requirement Observed</i>	<i>How does the chosen hardware meet this requirement?</i>
The CPUs in these machines are currently being used to process and view CCTV footage for security purposes and no significant additional processing power is available to run inference.	A VPU or NCS 2 can be plugged in a USB port and has a convenient plug and play kind of performance. This particularly is favorable as no more additional processing power is available to run inference.
The client's budget allows for a maximum of \$300 per machine.	A VPU or a NCS 2 stick costs almost \$100 opposed to the comparatively higher costs of FPGA, CPU or GPU.
The client would like to save as much as possible both on hardware and future power requirements.	A VPU or NCS 2 stick is designed to run on very low power. The NCS 2 can run on just 1 W of power adhering to the clients needs.
Queue Monitoring Requirements	
Maximum number of people in the queue	7-15
Model precision chosen (FP32, FP16, or Int8)	FP16

Figures

Figure 1

Main steps of the SQS

Figure 2

SQS run life-cycle

OpenVINO

DEEP LEARNING WORKBENCH

Version: 1.0.2205.782d052e

Dataset Name ↑	Date ↑	Size	Status	Type	Action	Select
imagenet100	05/12/19, 05:46	13 Mb		ImageNet		☐
VOC_small	05/12/19, 05:47	8 Mb		VOC Object Detection		☐

+

Import Local Dataset

⚙

Autogenerate Dataset in ImageNet format

Figure 3

Testing multiple models with DL Workbench

Inference History

#	Start Time	Parallel Streams	Batch	Status	Filter
A Baseline	17/07/19, 07:57	Sync	1	✓	☐
B Group Inference	18/07/19, 11:46	Group: 1-8, 1	1-16, 4	✓	☐

Figure 4

Testing Models with DL Workbench

Figure 5

Comparing models on different hardware

Figure 6

Output for manufacturing use case

Figure 7

Output for retail use case

Figure 8

Output for transportation use case