



(51) International Patent Classification:

H04L 12/24 (2006.01) G06F 15/16 (2006.01)  
H04L 12/70 (2013.01)

(21) International Application Number:

PCT/IN2018/050431

(22) International Filing Date:

02 July 2018 (02.07.2018)

(25) Filing Language:

English

(26) Publication Language:

English

(71) Applicant: **TELEFONAKTIEBOLAGET LM ERICSSON (PUBL)** [SE/SE]; SE-164 83, Stockholm (SE).

(72) Inventor; and

(71) Applicant (for SC only): **KATIYAR, Nitin** [IN/IN]; D-101, N D Sepal, Somasandra Palya, HSR Layout, Sec-2, Bangalore 560102 (IN).

(72) Inventors: **AJMERA, Vishal Deep**; B-004, Sobha Garnet Apts., Ibbalur, Bangalore 560103 (IN). **GUPTA, Keshav**; BAGMANE WORLD TECHNOLOGY SEZ, BANGALORE 560048 (IN). **THOMAS, Anju**; C-104, VRR Lakeview Apts, Doddanakundi, BANGALORE 560037 (IN).

(74) Agent: **SINGH, Manisha**; LEXORBIS, 709/710, Tolstoy House 15-17, Tolstoy Marg, New Delhi 110 001 (IN).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, ST, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

(54) Title: SOFTWARE SWITCH AND METHOD THEREIN




Fig. 3

(57) Abstract: A software switch (150) and a method performed by the software switch (150) are disclosed. The software switch (150) receives, from a node (110) deploying a virtual machine (130, 131, 132), a request for a virtual port to be polled by the virtual machine (130, 131, 132). The request comprises a Central Processing Unit "CPU" identity identifying a CPU (140, 141, 142) on which the virtual machine (130, 131, 132) executes, and wherein the request comprises an indication of a clock frequency at which the CPU (140, 141, 142) is set to operate at. The software switch (150) determines a number of packets in a queue associated with the virtual port. The software switch (150) adjusts the clock frequency of the CPU (140, 141, 142) based on the number of packets in the queue. A corresponding computer program (603) and a computer program carrier (605) are also disclosed.



**Declarations under Rule 4.17:**

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

**Published:**

- *with international search report (Art. 21(3))*

## SOFTWARE SWITCH AND METHOD THEREIN

## TECHNICAL FIELD

Embodiments herein relate to computer systems, such as networks of computer  
5 resources, cloud systems or the like. In particular, a method and a software switch e.g. for  
managing power consumption of the software switch are disclosed. A corresponding  
computer program and a computer program carrier are also disclosed.

## BACKGROUND

A known cloud system, such as a telecommunications cloud, may deploy virtual  
10 network functions (VNFs). The cloud system may then be set up as follows. A master node  
manages resources of the cloud system. The resources may be processing resources, such as  
Central Processing Units (CPUs) and processor kernels, memory resources, such as hard  
drives and volatile memory, network resources, such as wireless or wired network cards,  
and the like. The master node sets up one or more host machines by assembling resources of  
15 the cloud system. Furthermore, in a particular host machine the master node may arrange for  
one or more virtual machines to be executed. The virtual machine is typically an execution  
environment, such as a container, a virtual java machine, an operating system or the like.

In order for the virtual machines to be able to communicate with each other and  
entities outside the host machine, the master node instructs a software switch to map one or  
20 more physical ports to a virtual port which is associated with a particular virtual machine. In  
this manner, the master node sets up the particular virtual machine to continuously poll the  
virtual port for packets. The software switch polls its ports and pushes packets to a  
respective queue for each virtual port.

However, since the virtual machine continuously polls its respective queue, the  
25 CPU(s), allocated to it by the master node, will be very busy. Disadvantageously, power  
consumption of the host machine, on which the virtual machine executes, is high due to the  
continuous polling. A further disadvantage is that the continuous polling causes CPU(s) and  
other surrounding electronics to heat up, whereby power consumption for cooling thereof  
may be unnecessarily high.

## 30 SUMMARY

An object may be to more efficiently manage power consumption in a cloud system  
of the above mentioned kind.

According to an aspect, the object is achieved by a method, performed by a software  
switch. The software switch receives, from a node deploying a virtual machine, a request for

a virtual port to be polled by the virtual machine. The request comprises a CPU identity identifying a CPU on which the virtual machine executes, and the request comprises an indication of a clock frequency at which the CPU is set to operate at. The software switch determines a number of packets in a queue associated with the virtual port. The software switch adjusts the clock frequency of the CPU based on the number of packets in the queue.

According to another aspect, the object is achieved by a software switch configured for receiving, from a node deploying a virtual machine, a request for a virtual port to be polled by the virtual machine. The request comprises a CPU identity identifying a CPU on which the virtual machine executes, and the request comprises an indication of a clock frequency at which the CPU is set to operate at. The software switch is configured for determining a number of packets in a queue associated with the virtual port. The software switch is configured for adjusting the clock frequency of the CPU based on the number of packets in the queue.

Thanks to that the request comprises the CPU identity, the software switch is made aware of which CPU the virtual machine is executing on. Hence, by use of the identity, the software switch may adjust the clock frequency of the CPU on which the virtual machine executes.

Further, thanks to that the software switch determines the number of packets in the queue associated with the virtual port polled by the virtual machine, the software switch may adjust the clock frequency based on the number of packets in the queue. The number of packets in the queue may be indicative of load on the virtual machine.

As an example, the clock frequency may be decreased when the load appears to be low and the clock frequency may be increased when the load appears to be high as described in more detail herein. Accordingly, as the clock frequency increases power consumption of the CPU increases and as the clock frequency decreases the power consumption of the CPU decreases. Hence, power consumption efficiently managed, e.g. by taking load on the virtual machine into account. In particular, it may be expressed as that the CPU of the virtual machine continuously polls the queue while the clock frequency of the CPU is adjusted based on traffic load on the virtual port.

An advantage is that power consumption of the host machine and/or the cloud system of the above mentioned kind is efficiently managed. Furthermore, e.g. in case the power consumption is reduced, undesired heating of the CPU is also avoided. Moreover, in case the power consumption is reduced, environmental friendliness of the cloud system may

be improved thanks to the software switch described herein. Additionally, the embodiments herein may also reduce power needed for cooling of host machines.

Thus, the embodiments herein may provide a method and software switch, which more efficiently manage power consumption in a cloud system of the above mentioned kind.

## 5 BRIEF DESCRIPTION OF THE DRAWINGS

The various aspects of embodiments disclosed herein, including particular features and advantages thereof, will be readily understood from the following detailed description and the accompanying drawings, which are briefly described in the following.

Figure 1 is a schematic overview of an exemplifying system in which embodiments  
10 herein may be implemented.

Figure 2 is an illustration depicting a queue of packets destined to a virtual port.

Figure 3 is a flowchart illustrating embodiments of the method in the software switch.

Figure 4 and Figure 5 are flowcharts illustrating an example of a method in the  
15 software switch.

Figure 6 is a block diagram illustrating an example of the software switch.

## DETAILED DESCRIPTION

Throughout the following description, similar reference numerals have been used to denote similar features, such as nodes, actions, modules, circuits, parts, items, elements,  
20 units or the like, when applicable. In the Figures, features that are optional and appear in some embodiments are indicated by dashed lines.

**Figure 1** depicts an exemplifying **system 100** in which embodiments herein may be implemented. In this example, the system 100 is a telco cloud. In other examples, the system 100 may be a computer system, a data center, a cluster of computers, a data system, a server  
25 system, a hardware machine system or the like.

The system 100 may be said to comprise a **node 110**, which may be responsible for various tasks, such as resource allocation, deployment of virtual machines (VMs) etc. Resource allocation may e.g. refer to which hardware resources, such as memory, processing units, network capabilities that may be allocated, e.g. used, by which VM. In  
30 related literature, the node 110 may be referred to as an orchestrator, a master node, a central managing node, a host machine managing node or the like.

The system 100 may further be said to comprise a **host machine 120**, such as a logical host machine assembled from resources of the system 100 according to known manners, a hardware host machine, a computer, a server or the like.

The host machine 120 comprises one or more **virtual machines**, denoted VM, **130, 131, 132**. The node 110 may have scheduled the virtual machines 130, 131, 132 to be executed on the host machine 120.

Each VM 130, 131, 132 is associated with a respective set of **CPUs 140, 141, 142**. This means that each VM 130, 131, 132 may be executed by one or more CPUs. As shall be understood from the description below, each VM 130, 131, 132 may typically have a respective dedicated CPU that may be responsible for continuously polling for packets.

Moreover, the host machine 120 comprises a software (SW) switch 150. The software switch 150 may manage one or more virtual ports 151 and one or more physical ports 152.

The node 110 may **communicate 160** with the software switch 150 and each of the virtual machines 130, 131, 132. In this manner, requests, messages, responses, indications and generally information may be exchanged between the node 110 and/or the software switch 150 and/or one or more of the virtual machines 130, 131, 132.

Furthermore, the software switch 150 may **communicate 170** with each of the virtual machines 130, 131, 132. In this manner, requests, messages, responses, indications and generally information may be exchanged between the software switch 150 and/or one or more of the virtual machines 130, 131, 132.

**Figure 2** illustrates an exemplifying **queue 200**, which may be associated with a virtual port which e.g. is polled by the virtual machine 130. In this example, the number of **packets Pkt 1, Pkt 2, Pkt 3** in the queue is three (3) and a size of the queue is five (5) as derived from the **indices 1, 2, 3, 4 and 5** to the left of the queue. The size of the queue thus refers to a maximum length, or a maximum number of packets that may be put in the queue. Sometimes, the number of packets in the queue may be referred to as  $q\_len$ , which then in this example would be  $q\_len = 3$ .

**Figure 3** illustrates an exemplifying method according to embodiments herein when implemented in the system 100 of Figure 1. The software switch 150 may thus perform a method for managing power consumption of the host machine 120 and/or a virtual machine 130 and/or at least one virtual machine 130, 131, 132.

One or more of the following actions may be performed in any suitable order.

**Action A010**

The software switch 150 receives, from a node 110 deploying a virtual machine 130, 131, 132, a request for a virtual port to be polled by the virtual machine 130, 131, 132.

5 The request comprises a CPU identity (ID) identifying a CPU 140, 141, 142 on which the virtual machine 130, 131, 132 executes. By means of the CPU identity, the software switch 150 may identify which CPU(s) are associated with the virtual port.

10 The request comprises an indication of a clock frequency at which the CPU 140, 141, 142 is set to operate at. The indication of the clock frequency may be interpreted by the software switch as a current clock frequency, e.g. being a maximum possible clock frequency of the CPU on which the virtual machine executes.

When the node 110 deploys the virtual machine 130, 131, 132, the node 110 may previously have taken a decision to schedule the virtual machine for execution by the host machine 120.

**Action A020**

15 The software switch 150 may store the indication of the clock frequency, e.g. in a memory 602 of Figure 6. In this manner, the software switch 150 may obtain knowledge about e.g. the maximum possible clock frequency when desired as explained herein.

**Action A030**

20 The software switch 150 may poll physical ports of the software switch 150 to obtain one or more packets destined to the virtual machine 130, 131, 132.

**Action A040**

25 The software switch 150 may increment a low load counter, after each time the polling A030 of the physical port has been performed and when the length of the queue is less than a third threshold value for indicating low load since last time the adjusting A080 of the clock frequency was performed. In Figure 5, the third threshold may be exemplified by `q_low_load_threshold`.

**Action A050**

30 The software switch 150 may determine, i.e. based on so called forwarding processing on said one or more received packets, the virtual port to be polled by the virtual machine 130, 131, 132. Forward processing may refer to that the software switch 150 may inspect said one or more received packets, or packet headers, to determine the destination port, e.g. the virtual port, to forward to. Hence, said destination port is to be polled by the virtual machine 130, 131, 132.

**Action A060**

The software switch 150 may push the obtained one or more packets to the queue of the determined virtual port.

**Action A070**

5 The software switch 150 determines a number of packets in a queue associated with the virtual port, i.e. the number of packets that are currently present in the queue. Any packets in the queue may be destined to the virtual port.

**Action A080**

10 The software switch 150 adjusts the clock frequency of the CPU 140, 141, 142 based on the number of packets in the queue.

In some embodiments, the action A080 of adjusting the clock frequency may comprise adjusting the clock frequency of the CPU 140, 141, 142 based on the number of packets in the queue by decreasing the frequency when the number of packets in the queue is less than a first threshold value for indicating low load at the virtual machine 130, 131, 15 132, respectively, and by increasing the frequency when the number of packets in the queue is greater than a second threshold value for indicating high load at the virtual machine.

In this manner, the clock frequency may be set to a higher frequency - when load at the virtual machine is higher than the second threshold value – relatively a lower frequency to which the clock frequency may be set when the load at the virtual machine is lower than 20 the first threshold value.

The first threshold may be different from the second threshold. Alternatively, the first threshold may be the second threshold. Expressed differently, the second threshold may be the first threshold. Further, it can be expressed as that the first threshold comprises the second threshold.

25 Furthermore, the action A080 of adjusting the clock frequency by decreasing the frequency may comprise decreasing the frequency until the clock frequency is less than or equal to a first clock frequency indicating a lower limit for the clock frequency. In this manner, it may be ensured that the clock frequency may be set to at least a lower frequency limit, given by the first clock frequency, irrespectively of how small or low the load at the 30 virtual machine 130, 131, 132 may become. The first clock frequency may be set depending on application and requirements of the system 100. Generally, it may be noted that the smaller or lower the first clock frequency is, the smaller or lower the power consumption may be. At the same time, processing capacity, e.g. in terms of flops per second, may be smaller, the smaller or lower the first clock frequency may become. Hence, the first clock

frequency may be set while making a trade-off between maximum achievable power reduction and minimum required processing capacity. The first clock frequency may thus be seen as a lowest allowable frequency, which may be set according to prevailing circumstances.

5 It may here be noted that the expressions “clock frequency”, “first clock frequency” and “second clock frequency” are distinguished, i.e. different, from each other and refers to different frequency values.

Moreover, the action A080 of adjusting the clock frequency by increasing the frequency may comprise increasing the clock frequency until the clock frequency is equal to  
10 a second clock frequency indicating an upper limit for the clock frequency. In this manner, it may be ensured that the clock frequency may be set to at most an upper frequency limit, given by the second clock frequency, irrespectively of how great the load at the virtual machine 130, 131, 132 may become. The second clock frequency may usually be set to a maximum possible operating frequency of the CPU, possibly while applying some margin.  
15 The second clock frequency may thus be seen as a maximum possible frequency, which may be set according to e.g. prevailing CPU specifications.

In some embodiments, the adjusting A080 of the clock frequency may be performed at regular intervals.

In these embodiments, the adjusting of the clock frequency may be performed at  
20 regular intervals based on a number of times the action A030 of polling the physical ports has been performed. In more detail, a reduction of the clock frequency may done at regular interval, while e.g. an increase of the clock frequency may preferably be performed as soon as the second threshold is exceeded. In this manner, a risk of dropped packets is reduced.

Alternatively, the adjusting of the clock frequency may be performed at regular  
25 intervals according to a time interval, such as 10 ms, 20 s, 4 min or the like depending on application.

Figure 4 and Figure 5 illustrate a further exemplifying method according to the embodiments herein.

Turning to **Figure 4**, one or more of the following actions may be performed.

#### 30 **Action 410**

The node 110 launches, such as deploys or the like, a VM with a CPU ID CPUID#1 in the host machine 120.

As an example, during VM creation, the node 110 may need to enable a power management feature according to the embodiments herein for a virtual port created in the

software switch 150, wherein the virtual port shall be polled by the VM, being created and/or deployed.

The node 110 may, e.g. when requesting the virtual port, specify the physical core ID, as an example of the aforementioned CPU identity, corresponding to the virtual port that would be used for exchanging, i.e. sending and/or receiving, packets between VM and the software switch 150.

In case, the CPU identity is set to -1, the embodiments herein may preferably not be invoked.

#### **Action 420**

The software switch 150 may store, such as record, register or the like, the clock frequency. The clock frequency may be received by the software switch in connection with the request for the virtual port as in action A010. In an example, it may be that the clock frequency is predefined, such as given to the software switch 150 at configuration thereof. Further, it may be that the clock frequency is received by the software switch 150 from the node 110 prior to action A010.

#### **Action 430**

The software switch 150 may poll physical ports, such as physical network interface cards (NICs) continuously according to the clock frequency, i.e. according to a default or initial clock frequency.

In this action, the software switch 150 may further increment a global counter, e.g. `poll_iteration` as shown in Figure 4.

#### **Action 440**

The software switch 150 may check if the global counter `poll_iteration` is greater than, or equal to, a polling threshold `g_cnt` for when to regularly adjust, in particular reduce, the clock frequency as in action A030.

#### **Action 450**

If the global counter `poll_iteration` is greater than, or equal to, the polling threshold `g_cnt`, the software switch 150 may reset the global counter `poll_iteration`.

#### **Action 460**

The software switch 150 may get, such as read, obtain, fetch or the like, a low load counter `low_load_cnt` for each virtual port of the software switch 150.

#### **Action 470**

The software switch 150 may check if the low load counter `low_load_cnt` is greater than, or equal to, a predefined value, `p_cnt`, as an example of the third threshold value for

indicating low load since last time the adjusting A080 of the clock frequency was performed, e.g. a reduction of the clock frequency was made.

**Action 480**

If the low load counter `low_load_cnt` is greater than, or equal to, the predefined value, `p_cnt`, the software switch 150 may reduce the clock frequency.

**Action 490**

The software switch 150 may reset the low load counter `low_load_cnt` irrespectively of outcome of action 470. This may mean that action 490 may be performed before any one or more of actions 450, 460, 470 and 480, but after action 440.

Hence, according to actions 440 to 490, the software switch 150 may check the low load counter `low_load_cnt` for each of the virtual ports once the `poll_iteration` counter exceeds `g_cnt`. If the low load counter `low_load_cnt` for a virtual port is greater than a predefined value, e.g. `p_cnt`, the software switch 150 may trigger a reduction of the clock frequency of the associated CPU. The reduction of the clock frequency may in turn yield a reduction of power consumption by that CPU. Finally, the software switch 150 may reset the low load counter `low_load_cnt` for the virtual port associated with the CPU.

In view of the description of the actions in Figure 4, it may be noted that one or more of actions 450 to 490 may be performed for one or more of the virtual port of the software switch 150, i.e. at those ports at which it is desired to adjust the clock frequency of the corresponding CPUs.

Now with reference to **Figure 5**, one or more of the following actions may be performed for every polling for which the global counter `poll_iteration` is less than the polling threshold `g_cnt`.

**Action 500**

The software switch 150 may check if any packet (`pkt`) is received.

**Action 510**

When action 500 results in that the software switch 150 determines that a packet has been received, the software switch 150 may forward the packet to an appropriate virtual port, according to a mapping table, if the appropriate virtual port is identified based on the received packet.

**Action 520**

The software switch 150 may determine, such as estimate, calculate and the like, a number of packets `q_dep`, or `q_len`, in the queue of the virtual port.

**Action 530**

Subsequent to action 520, the software switch 150 may check if the number of packets  $q_{dep}$  is greater than, or equal to, the second threshold value  $q_{high\_threshold}$  for indicating high load at the virtual machine.

**5 Action 540**

If the checking of action 530 results in that the software switch 150 finds that the number of packets  $q_{dep}$  is greater than, or equal to, the second threshold value, the software switch 150 may increase the clock frequency towards the default clock frequency, or the software switch 150 may reset the clock frequency to the default clock frequency.

10 Furthermore, the software switch 150 may reset the low load counter  $low\_load\_cnt$  when the number of packets  $q_{dep}$  is greater than, or equal to, the second threshold value.

Subsequent to action 540, the software switch 150 may proceed at action 430 of Figure 4.

**Action 550**

15 If the checking of action 530 results in that the software switch 150 finds that the number of packets  $q_{dep}$  is less than the second threshold value, the software switch 150 may check if the number of packets  $q_{dep}$  is less than, or equal to, the first threshold value  $q_{low\_threshold}$  for indicating low load at the virtual machine.

**Action 560**

20 The software switch 150 may increment the low load counter  $low\_load\_cnt$  when the number of packets  $q_{dep}$  is less than, or equal to, the first threshold value.

This means that the low load counter  $low\_load\_cnt$  is typically increased by one. However, depending on application, the low load count  $low\_load\_cnt$  may be increased by two, three, four etc. as deemed appropriate.

25 In view of the actions described above in relation to Figure 5, it may be noted that actions 520 to 560 may be performed for all virtual ports of the software switch 150 with which a virtual machine is associated.

To conclude, by varying the frequency of CPU given to VM as mentioned above, the software switch 150 may manage the power consumption of the CPUs based on the  
30 load/traffic handled by the virtual machines deployed in the host machine 110. Whenever there is low or moderate traffic, the software switch 150 may reduce the clock frequency of the related CPU(s). Thus, an overall power usage by the host machine in the system 100 may be reduced.

Expressed somewhat differently, the embodiments herein may monitor traffic load, e.g. in terms of number of packets in queues, on the virtual ports connected to the virtual machine and utilize this information to tune the clock frequency of the virtual machine's CPU.

5 With reference to **Figure 6**, a schematic block diagram of an example of the software switch 150 of Figure 1 is shown.

The software switch 150 may comprise a **processing unit 601**, such as a means for performing the methods described herein. The means may be embodied in the form of one or more hardware units and/or one or more software units. The term "unit" may thus refer to  
10 a circuit, a software block or the like according to various embodiments as described below.

The software switch 150 may further comprise a **memory 602**. The memory may comprise, such as contain or store, instructions, e.g. in the form of a **computer program 603**, which may comprise computer readable code units.

According to some embodiments herein, the software switch 150 and/or the  
15 processing unit 601 may comprise a **processing circuit 604** as an exemplifying hardware unit, which may comprise one or more processors. Accordingly, the processing unit 601 may be embodied in the form of, or 'realized by', the processing circuit 604. The instructions may be executable by the processing circuit 604, whereby the software switch 150 is operative to perform the methods of Figure 3 and/or Figure 4 and/or Figure 5. As  
20 another example, the instructions, when executed by the software switch 150 and/or the processing circuit 604, may cause the software switch 150 to perform the method according to Figure 3 and/or Figure 4 and/or Figure 5.

In view of the above, in one example, there is provided a software switch 150. Again, the memory 602 contains the instructions executable by said processing circuit 604  
25 whereby the software switch 150 is operative for:

receiving, from a node 110 deploying a virtual machine 130, 131, 132, a request for a virtual port to be polled by the virtual machine 130, 131, 132, wherein the request comprises a Central Processing Unit "CPU" identity identifying a CPU 140, 141, 142 on which the virtual machine 130, 131, 132 executes, and wherein the request comprises an  
30 indication of a clock frequency at which the CPU 140, 141, 142 is set to operate at,

determining a number of packets in a queue associated with the virtual port, and  
adjusting the clock frequency of the CPU 140, 141, 142 based on the length of the queue.

Figure 6 further illustrates a **carrier 605**, or program carrier, which provides, such as comprises, mediates, supplies and the like, the computer program 603 as described directly above. The carrier 605 may be one of an electronic signal, an optical signal, a radio signal and a computer readable medium.

5 In some embodiments, the software switch 150 and/or the processing unit 601 may comprise one or more of a **receiving unit 610**, a **determining unit 620**, an **adjusting unit 630**, a **polling unit 640**, a **pushing unit 650** and an **incrementing unit 690** as exemplifying hardware units. The term “unit” may refer to a circuit when the term “unit” refers to a hardware unit. In other examples, one or more of the aforementioned exemplifying  
10 hardware units may be implemented as one or more software units.

Moreover, the software switch 150 and/or the processing unit 601 may comprise an **Input/Output unit 606**, which may be exemplified by the receiving unit and/or a sending unit when applicable.

Therefore, according to the various embodiments described above, the software  
15 switch 150 and/or the processing unit 601 and/or the receiving unit 610 and/or the processing circuit 604 is configured for receiving, from a node 110 deploying a virtual machine 130, 131, 132, a request for a virtual port to be polled by the virtual machine 130, 131, 132. The request comprises a Central Processing Unit “CPU” identity identifying a CPU 140, 141, 142 on which the virtual machine 130, 131, 132 executes, and wherein the  
20 request comprises an indication of a clock frequency at which the CPU 140, 141, 142 is set to operate at.

The software switch 150 and/or the processing unit 601 and/or the determining unit 610 and/or the processing circuit 604 is configured for determining a number of packets in a queue associated with the virtual port.

25 The software switch 150 and/or the processing unit 601 and/or the adjusting unit 610 and/or the processing circuit 604 is configured for adjusting the clock frequency of the CPU 140, 141, 142 based on the length of the queue.

The software switch 150 may be configured for adjusting the clock frequency of the CPU 140, 141, 142 based on the length of the queue by decreasing the frequency when the  
30 length is less than a first threshold value for indicating low load at the virtual machine 130, 131, 132 and by increasing the frequency when the length is greater than a second threshold value for indicating high load at the virtual machine.

The software switch 150 may be configured for adjusting the clock frequency by decreasing the frequency until the clock frequency is less than or equal to a first clock frequency indicating a lower limit for the clock frequency.

5 The software switch 150 may be configured for adjusting the clock frequency by increasing the frequency until the clock frequency is equal to a second clock frequency indicating an upper limit for the clock frequency.

10 The software switch 150 and/or the processing unit 601 and/or the polling unit 610 and/or the processing circuit 604 may be configured for polling physical ports of the software switch 150 to obtain one or more packets destined to the virtual machine 130, 131, 132.

The software switch 150 and/or the processing unit 601 and/or the determining unit 610 and/or the processing circuit 604 may be configured for determining the virtual port to be polled by the virtual machine 130, 131, 132.

15 The software switch 150 and/or the processing unit 601 and/or the determining unit 610 and/or the processing circuit 604 may be configured for pushing the obtained one or more packets to the queue of the determined virtual port.

The software switch 150 may be configured for adjusting the clock frequency at regular intervals.

20 The software switch 150 may be configured for adjusting the clock frequency at regular intervals based on a number of times the polling of the physical ports has been performed.

The software switch 150 may be configured for adjusting the clock frequency at regular intervals according to a time interval.

25 The software switch 150 and/or the processing unit 601 and/or the determining unit 610 and/or the processing circuit 604 may be configured for incrementing a low load counter, after each time the polling of the physical port has been performed and when the length of the queue is less than a third threshold value for indicating low load since last time the adjusting of the clock frequency was performed.

30 The first threshold may be different from the second threshold. Alternatively, the first threshold may be the second threshold. Expressed differently, the second threshold may be the first threshold. Further, it can be expressed as that the first threshold comprises the second threshold.

As used herein, the term “node”, or “network node”, may refer to one or more physical entities, such as devices, apparatuses, computers, servers or the like. This may

mean that embodiments herein may be implemented in one physical entity. Alternatively, the embodiments herein may be implemented in a plurality of physical entities, such as an arrangement comprising said one or more physical entities, i.e. the embodiments may be implemented in a distributed manner, such as on a cloud system, which may comprise a set  
5 of server machines. In case of a cloud system, the term “node” may refer to a virtual machine, such as a container, virtual runtime environment, a software module or the like. The virtual machine may be assembled from hardware resources, such as memory, processing, network and storage resources, which may reside in different physical machines, e.g. in different computers.

10 As used herein, the term “unit” may refer to one or more functional units, each of which may be implemented as one or more hardware units and/or one or more software units and/or a combined software/hardware unit in a node. In some examples, the unit may represent a functional unit realized as software and/or hardware of the node.

15 As used herein, the term “computer program carrier”, “program carrier”, or “carrier”, may refer to one of an electronic signal, an optical signal, a radio signal, and a computer readable medium. In some examples, the computer program carrier may exclude transitory, propagating signals, such as the electronic, optical and/or radio signal. Thus, in these examples, the computer program carrier may be a non-transitory carrier, such as a non-transitory computer readable medium.

20 As used herein, the term “processing unit” may include one or more hardware units, one or more software units or a combination thereof. Any such unit, be it a hardware, software or a combined hardware-software unit, may be a determining means, estimating means, capturing means, associating means, comparing means, identification means, selecting means, receiving means, sending means or the like as disclosed herein. As an  
25 example, the expression “means” may be a unit corresponding to the units listed above in conjunction with the Figures.

As used herein, the term “software unit” may refer to a software application, a dynamic link library (DLL), a software component, a software module, a software object, an object according to component object model (COM), a software function, a software engine,  
30 an executable binary software file or the like.

The terms “processing unit” or “processing circuit” may herein encompass a processing unit, comprising e.g. one or more processors, an application specific integrated circuit (ASIC), a field-programmable gate array (FPGA) or the like. The processing circuit or the like may comprise one or more processor kernels.

As used herein, the expression “configured to/for” may mean that a processing circuit is configured to, such as adapted to or operative to, by means of software configuration and/or hardware configuration, perform one or more of the actions described herein.

5 As used herein, the term “action” may refer to an action, a step, an operation, a response, a reaction, an activity or the like. It shall be noted that an action herein may be split into two or more sub-actions as applicable. Moreover, also as applicable, it shall be noted that two or more of the actions described herein may be merged into a single action.

10 As used herein, the term “memory” may refer to a hard disk, a magnetic storage medium, a portable computer diskette or disc, flash memory, random access memory (RAM) or the like. Furthermore, the term “memory” may refer to an internal register memory of a processor or the like.

15 As used herein, the term “computer readable medium” may be a universal serial bus (USB) memory, a digital versatile disc (DVD), a Blu-ray disc, a software module that is received as a stream of data, a Flash memory, a hard drive, a memory card, such as a MemoryStick, a multimedia card (MMC), secure digital (SD) card, etc. One or more of the aforementioned examples of computer readable medium may be provided as one or more computer program products.

20 As used herein, the term “computer readable code units” may be text of a computer program, parts of or an entire binary file representing a computer program in a compiled format or anything there between.

25 As used herein, the expression “transmit” and “send” are considered to be interchangeable. These expressions include transmission by broadcasting, uni-casting, group-casting and the like. In this context, a transmission by broadcasting may be received and decoded by any authorized device within range. In case of uni-casting, one specifically addressed device may receive and decode the transmission. In case of group-casting, a group of specifically addressed devices may receive and decode the transmission.

30 As used herein, the terms “number” and/or “value” may be any kind of digit, such as binary, real, imaginary or rational number or the like. Moreover, “number” and/or “value” may be one or more characters, such as a letter or a string of letters. “Number” and/or “value” may also be represented by a string of bits, i.e. zeros and/or ones.

As used herein, the terms “first”, “second”, “third” etc. may have been used merely to distinguish features, apparatuses, elements, units, or the like from one another unless otherwise evident from the context.

As used herein, the term “subsequent action” may refer to that one action is performed after a preceding action, while additional actions may or may not be performed before said one action, but after the preceding action.

As used herein, the term “set of” may refer to one or more of something. E.g. a set of  
5 devices may refer to one or more devices, a set of parameters may refer to one or more parameters or the like according to the embodiments herein.

As used herein, the expression “in some embodiments” has been used to indicate that the features of the embodiment described may be combined with any other embodiment disclosed herein.

10 Even though embodiments of the various aspects have been described, many different alterations, modifications and the like thereof will become apparent for those skilled in the art. The described embodiments are therefore not intended to limit the scope of the present disclosure.

## CLAIMS:

1. A method, performed by a software switch (150), wherein the method comprises:
  - receiving (A010), from a node (110) deploying a virtual machine (130, 131, 132), a request for a virtual port to be polled by the virtual machine (130, 131, 132),  
5 wherein the request comprises a Central Processing Unit “CPU” identity identifying a CPU (140, 141, 142) on which the virtual machine (130, 131, 132) executes, and wherein the request comprises an indication of a clock frequency at which the CPU (140, 141, 142) is set to operate at,  
determining (A070) a number of packets in a queue associated with the virtual  
10 port, and  
adjusting (A080) the clock frequency of the CPU (140, 141, 142) based on the number of packets in the queue.
2. The method according to claim 1, wherein the adjusting (A080) of the clock frequency  
15 comprises adjusting the clock frequency of the CPU (140, 141, 142) based on the number of packets in the queue by decreasing the frequency when the number of packets in the queue is less than a first threshold value for indicating low load at the virtual machine (130, 131, 132) and by increasing the frequency when the number of packets in the queue is greater than a second threshold value for indicating high load at the virtual  
20 machine.
3. The method according to claim 2, wherein the adjusting (A080) of the clock frequency by decreasing the frequency comprises decreasing the frequency until the clock frequency is equal to a first clock frequency indicating a lower limit for the clock  
25 frequency.
4. The method according to claim 2 or 3, wherein the adjusting (A080) of the clock frequency by increasing the frequency comprises increasing the clock frequency until the clock frequency is equal to a second clock frequency indicating an upper limit for the  
30 clock frequency.
5. The method according to any one of the preceding claims, wherein the method comprises:

polling (A030) physical ports of the software switch (150) to obtain one or more packets destined to the virtual machine (130, 131, 132),

determining (A050) the virtual port to be polled by the virtual machine (130, 131, 132), and

5 pushing (A060) the obtained one or more packets to the queue of the determined virtual port.

6. The method according to any one of the preceding claims, wherein the adjusting (A080) of the clock frequency is performed at regular intervals.

10

7. The method according to claim 6, when dependent on claim 5, wherein the adjusting (A080) of the clock frequency is performed at regular intervals based on a number of times the polling (A030) of the physical ports has been performed.

15

8. The method according to claim 6, wherein the adjusting (A080) of the clock frequency is performed at regular intervals according to a time interval.

9. The method according to any one of claims 5-8, wherein the method comprises, after each time the polling (A030) of the physical port has been performed and when the  
20 number of packets in the queue is less than a third threshold value for indicating low load since last time the adjusting (A080) of the clock frequency was performed, incrementing (A040) a low load counter.

25

10. The method according to any one of the preceding claims, wherein the first threshold is different from the second threshold, or wherein the first threshold is the second threshold.

11. A software switch (150) configured for:

30

receiving, from a node (110) deploying a virtual machine (130, 131, 132), a request for a virtual port to be polled by the virtual machine (130, 131, 132), wherein the request comprises a Central Processing Unit "CPU" identity identifying a CPU (140, 141, 142) on which the virtual machine (130, 131, 132) executes, and wherein the request comprises an indication of a clock frequency at which the CPU (140, 141, 142) is set to operate at,

35

determining a number of packets in a queue associated with the virtual port, and

adjusting the clock frequency of the CPU (140, 141, 142) based on the number of packets in the queue.

- 5 12. The software switch (150) according to claim 11, wherein the software switch (150) is configured for adjusting the clock frequency of the CPU (140, 141, 142) based on the number of packets in the queue by decreasing the frequency when the number of packets in the queue is less than a first threshold value for indicating low load at the virtual machine (130, 131, 132) and by increasing the frequency when the number of packets in the queue is greater than a second threshold value for indicating high load at the virtual machine.
- 10
13. The software switch (150) according to claim 12, wherein the software switch (150) is configured for adjusting the clock frequency by decreasing the frequency until the clock frequency is equal to a first clock frequency indicating a lower limit for the clock frequency.
- 15
14. The software switch (150) according to claim 12 or 13, wherein the software switch (150) is configured for adjusting the clock frequency by increasing the frequency until the clock frequency is equal to a second clock frequency indicating an upper limit for the clock frequency.
- 20
15. The software switch (150) according to any one of claims 11-14, wherein the software switch (150) is configured for:
- 25       polling physical ports of the software switch (150) to obtain one or more packets destined to the virtual machine (130, 131, 132),
- determining the virtual port to be polled by the virtual machine (130, 131, 132),
- and
- pushing the obtained one or more packets to the queue of the determined virtual port.
- 30
16. The software switch (150) according to any one of claims 11-15, wherein the software switch (150) is configured for adjusting the clock frequency at regular intervals.

17. The software switch (150) according to claim 16, when dependent on claim 15, wherein the software switch (150) is configured for adjusting the clock frequency at regular intervals based on a number of times the polling of the physical ports has been performed.

5

18. The software switch (150) according to claim 16, wherein the software switch (150) is configured for adjusting the clock frequency at regular intervals according to a time interval.

10 19. The software switch (150) according to any one of claims 15-18, wherein the software switch (150) is configured for:

incrementing a low load counter, after each time the polling of the physical port has been performed and when the number of packets in the queue is less than a third threshold value for indicating low load since last time the adjusting of the clock frequency was performed.

15

20. The software switch (150) according to any one of claims 11-19, wherein the first threshold is different from the second threshold, or wherein the first threshold is the second threshold.

20

21. A computer program (603), comprising computer readable code units which when executed on a software switch (150) cause the software switch (150) to perform a method according to any one of claims 1-10.

25 22. A carrier (605) providing a computer program (603) according to the preceding claim, wherein the carrier (605) is one of an electronic signal, an optical signal, a radio signal and a computer readable medium.

1/5




Fig. 1

200




Fig. 2

2/5

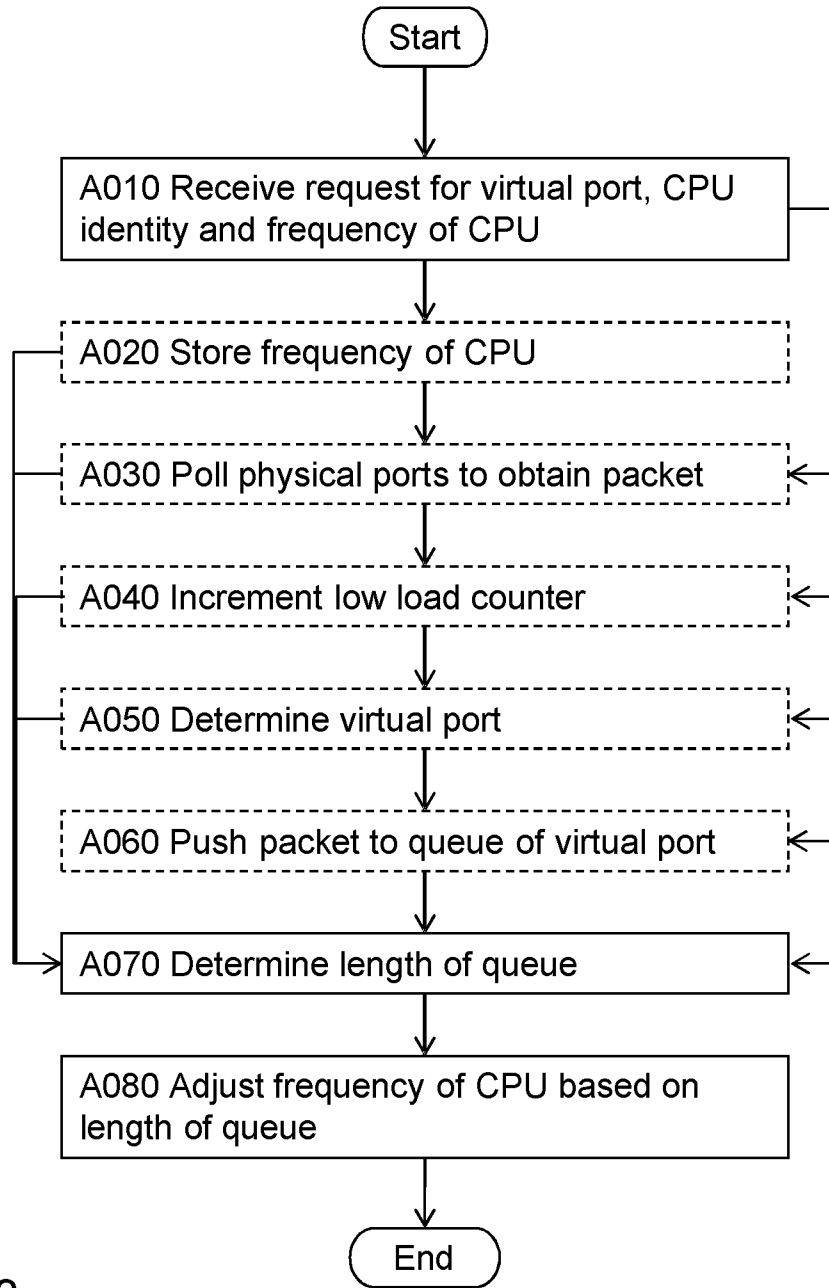


Fig. 3

3/5

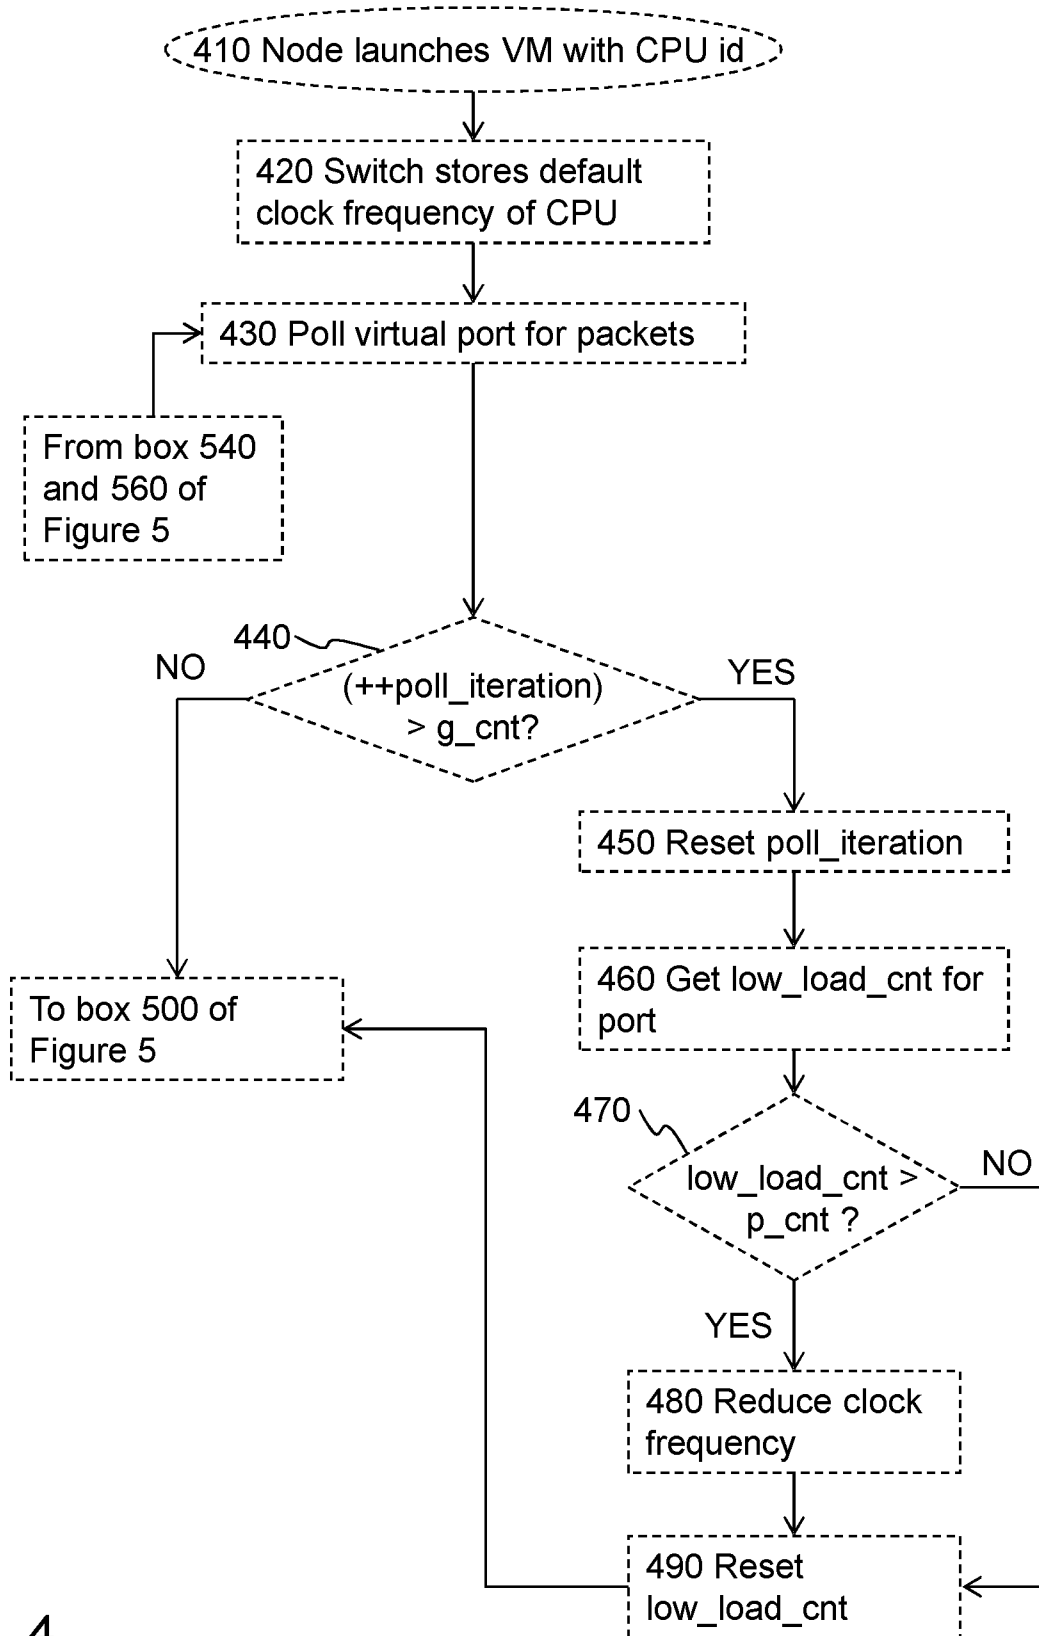


Fig. 4

4/5

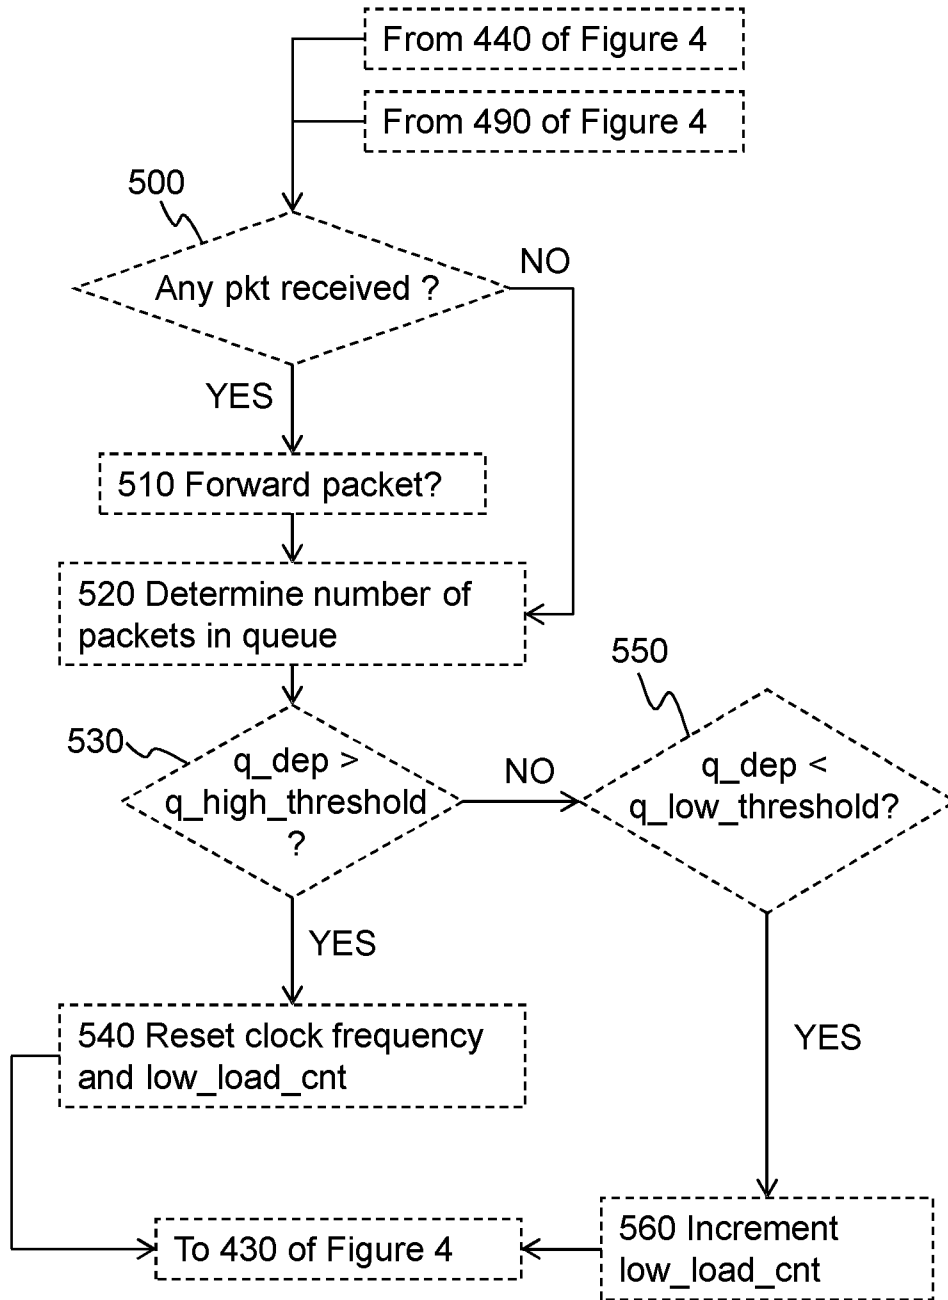


Fig. 5

5/5

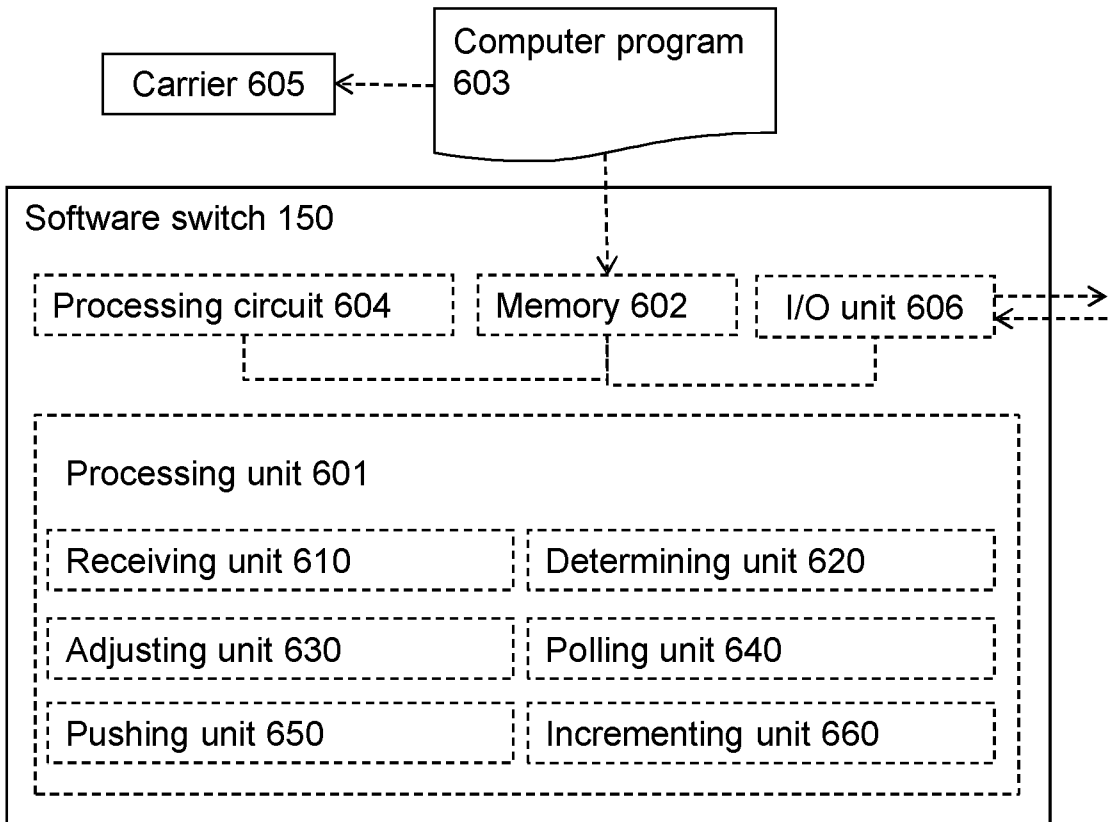


Fig. 6

## INTERNATIONAL SEARCH REPORT

International application No.  
PCT/IN2018/050431

A. CLASSIFICATION OF SUBJECT MATTER  
H04L12/24, H04L12/70, G06F15/16 Version=2018.01

According to International Patent Classification (IPC) or to both national classification and IPC

## B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

H04L, G06F.

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

Databases: TotalPatent One, IPO Internal Database  
Keywords: virtual machine, frequency, power conservation.

## C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	JP2012164032A (MURAKAMI TAKEO ET AL.) 30 August 2012 (30-08-2012) Whole document	1-22
Y	CN106527653A (LI TUO) 22 March 2017 (22-03-2017) Whole document	1-22

Further documents are listed in the continuation of Box C.  See patent family annex.

* Special categories of cited documents:	"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be of particular relevance	"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"E" earlier application or patent but published on or after the international filing date	"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	"&" document member of the same patent family
"O" document referring to an oral disclosure, use, exhibition or other means	
"P" document published prior to the international filing date but later than the priority date claimed	

Date of the actual completion of the international search 15-10-2018	Date of mailing of the international search report 15-10-2018
-------------------------------------------------------------------------	------------------------------------------------------------------

Name and mailing address of the ISA/ Indian Patent Office Plot No.32, Sector 14, Dwarka, New Delhi-110075 Facsimile No.	Authorized officer Sumedh Prabhakar Kadam Telephone No. +91-1125300200
----------------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------