

Improved parallelism and scheduling in multi-core software routers

Norbert Egi · Gianluca Iannaccone ·
Maziar Manesh · Laurent Mathy ·
Sylvia Ratnasamy

© Springer Science+Business Media, LLC 2011

Abstract Recent technological advances in commodity server architectures, with multiple multi-core CPUs, integrated memory controllers, high-speed interconnects, and enhanced network interface cards, provide substantial computational capacity, and thus an attractive platform for packet forwarding. However, to exploit this available capacity, we need a suitable software platform that allows effective parallel packet processing and resource management. In this paper, we at first introduce an improved forwarding architecture for software routers that enhances parallelism by exploiting hardware classification and multi-queue support, already available in recent commodity network interface cards. After evaluating the original scheduling algorithm of the widely-used Click modular router, we propose solutions for extending this scheduler for improved fairness, throughput, and more precise resource management. To illustrate the potential benefits of our proposal, we implement and evaluate a few key elements of our overall design. Finally, we discuss how our improved forwarding architecture and resource management might be applied in virtualized software routers.

N. Egi (✉) · L. Mathy
Computing Department, Lancaster University, Lancaster, UK
e-mail: n.egi@lancaster.ac.uk

L. Mathy
e-mail: l.mathy@lancaster.ac.uk

G. Iannaccone · M. Manesh · S. Ratnasamy
Intel Labs, Berkeley, USA

G. Iannaccone
e-mail: gianluca.iannaccone@intel.com

M. Manesh
e-mail: maziar.h.manesh@intel.com

S. Ratnasamy
e-mail: sylvia.p.ratnasamy@intel.com

Keywords Software-based router · Forwarding architectures · Parallel processing · Resource management · Scheduling

1 Introduction

Recent advances in server technology promise significant speedups to applications amenable to parallelization. Router workloads appear ideally suited to exploit these advances, which has led to a renewed interest in the applicability of software routers. Software routers offer several advantages, like a familiar programming environment and ease of extensibility, which offers the potential to serve as a single platform onto which one can consolidate many network functions typically implemented by various middleboxes. Examples are the single server-based routers running open-source router software as offered by Vyatta [1], which have even been touted as cheaper alternatives to other commercial routers. However, the limitation of software routers has always been performance. Which is why the possibility of leveraging recent server advances to further the reach of software routers has been of high interest [2–5].

The recent trajectory of server advances has been of growing the available resource capacity through increased parallelism. Multi-core CPUs increase the available CPU resources; multiple memory controllers integrated in the CPU sockets increase the available memory capacity; multiple I/O links and PCIe buses do the same for I/O and—importantly, high-speed point-to-point interconnects offer high-capacity access between components.

However, precisely because the underlying hardware is parallel, achieving high performance will rely greatly on our ability to distribute packet-processing well across the available resources (not unlike routing across a network—the more alternate routes, the more a routing algorithm has to work to make sure the network capacity is well exploited). Traditionally, this is the job of the scheduler to decide which set of tasks are assigned to which cores and when (if at all) tasks must be moved between cores.

One of the major points of software routers is that they can simultaneously support very different forms of packet processing, with different resource characteristics, on a single hardware architecture. For example, one can imagine having multiple customers—for one we are performing encryption, and for the other deep packet inspection. Since certain operations are more expensive (in terms of the server resources they consume) than others we want to make sure that each customer receives a fair access to the server resources. In general, this requires proper abstractions for flexible and efficient flow differentiation as well as resource allocation. While we also need scheduling techniques that make good use of the resources so as to achieve high performance, and can also support various scenarios of fair resource sharing and isolation.

In this paper, we observe that neither traditional forwarding architectures nor traditional schedulers are ideally suited to our task, for the following reasons: (1) Traditional forwarding architectures are not capable of taking advantage of recent server advances (e.g. multiple I/O links, hardware classification, and multi-queueing in the network interface cards), and thus are limited in exploiting the highest performance

achievable. (2) Schedulers are CPU centric. This can be problematic since one of the distinguishing features of a packet-processing workload is that it stresses more than just the CPU. This becomes even more likely as the number of cores increases resulting in bottleneck potentially being any of the CPU, the memory, or the I/O, and hence being “fair” in terms of CPU may not mean an adequate solution. (3) Schedulers fail to take into account the heterogeneity in workloads. (4) Schedulers also fail to consider the basic question of whether we can accurately measure the resource consumption for different packets/flows in a dynamically changing workload. We cannot hope to allocate resources appropriately if we cannot first accurately account for their consumption.

To this aim, we approach the limitations from two directions. On the one hand, we propose an improved forwarding architecture that outperforms, in terms of throughput and latency, the traditional way of packet forwarding, while it also provides additional advantages for resource management and scheduling. We then focus on understanding the limitations of traditional schedulers and on quantifying their impact, thus identifying the requirements for an improved scheduler design.

The paper is organized into the following sections. In Sect. 2, we discuss the limitations of traditional forwarding architectures and propose an architecture that overcomes these limitations. Section 3 focuses on the resource management, while Sect. 4 describes how the proposed forwarding architecture and resource management extensions might be applied in virtualized software routers. Finally, Sect. 5 concludes the paper.

2 Forwarding architectures

In this section, at first, we analyze the shortcomings of the traditional forwarding architecture widely used in commodity software routers and propose an improved forwarding architecture that outperforms, in terms of throughput and latency, this traditional way of packet forwarding, while it also provides additional advantages for resource management. Our proposed architecture builds on the observation that recent network interface cards (NICs) [6] provide multiple receive and transmit queues to support server virtualization, but these multi-queue NICs can also be used to achieve better parallelism, and hence improve resource management and scheduling in the forwarding engine of software routers.

Before going into the details of the forwarding architectures, we briefly overview the server architecture and software environment we used to implement our extensions and carry out our measurements.

2.1 Server architecture

The hardware environment we use for our study is a dual-socket 2.8 GHz Intel Nehalem server [7], since it implements the most recent advances in server architecture. Figure 1 illustrates the high-level architecture of our server. Multiple processing cores¹ are arranged in “sockets”; our server has two sockets with four cores per

¹We use the terms “CPU,” “core” and “processor” interchangeably.

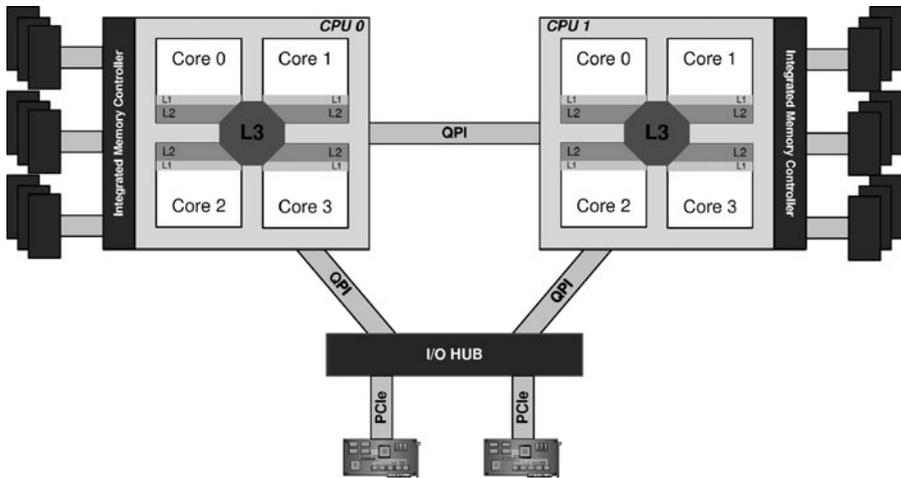


Fig. 1 Server architecture based on point-to-point inter-socket links and integrated memory controllers

socket. All cores in a socket share the same 8 MB L3 cache, while every core also has on its own a 256 KB L2 and a 64 KB L1 cache. A memory controller is integrated within each socket and connects to a portion of the overall memory space via a memory bus. The use of multiple sockets with integrated memory controllers means that memory accesses are non-uniform (NUMA). Dedicated high-speed point-to-point links serve to connect the two sockets directly, as well as to connect each socket to the I/O hub. Finally, the I/O hub connects to the NICs via a set of PCIe buses. Our server has 2 PCIe1.1 \times 8 slots which we populate with 2 NICs, each holding two 10 Gbps ports [6]. These network cards can provide us upto 32 transmit (Tx) and 64 receive (Rx) queues, while they support both the Receive Side Scaling (RSS) [8] as well as the Virtual Machine Device queue (VMDq) [9] mechanisms for distributing the arriving packets into the desired Rx queues.

These servers represent the next-generation replacement for the widely deployed Xeon servers, i.e., these servers conform to the informal notion of a “commodity” server.

Our server runs Linux 2.6.24.7 with the Click Modular Router [10] in polling mode—i.e., the CPUs poll for incoming packets rather than being interrupted. We use Click to perform the packet forwarding and processing functions of our router as it has been shown to offer a good tradeoff between ease of programming and performance.

Click is a modular software architecture that offers a flexible approach for implementing software routers on Linux and FreeBSD. We use Click only in the Linux kernel-mode due to the need for the best performance we can get from our architecture. From a high-level perspective the Click implementation consists of packet processing elements and packet queues, connected in a data-flow like configuration. Elements include the code for performing a packet processing function. A number of elements can be connected in any structure and order needed. Typically, a configuration will have a set of input, forwarding, and output elements, each separated by packet queues.

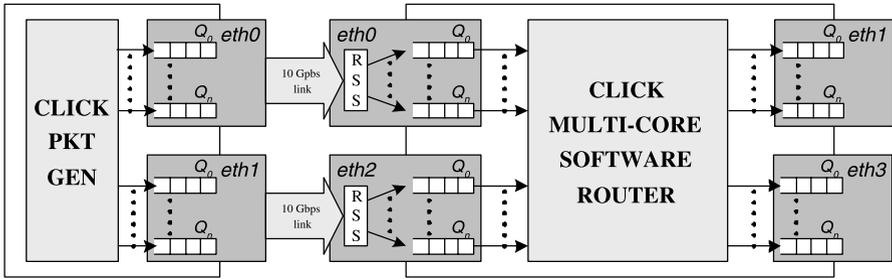


Fig. 2 Experimental configuration

In addition, Click supports multi-threading inside the Linux kernel [11], that is, different elements in the same configuration can be scheduled on separate cores, providing an excellent platform for developing and experimenting with our forwarding and resource management methods. To further justify our choice of Click over native Linux, we refer to the results reported in [12] that show how limited the standard Linux kernel is, as far as its ability of effectively parallelizing the packet forwarding process, and then of exploiting the overall capacity of multi-core systems.

Figure 2 illustrates the topology of our experimental configuration. In our experiments, we used two servers, one acted as traffic generator while the other as the actual router. With regard to the generated traffic, we use a synthetic workload of min-sized (64 B) packets as it stresses the system the most.

For performance measurements, we instrument our server with a proprietary performance tool similar to Intel VTune [13] and somewhat to Oprofile [14].

2.2 Traditional forwarding architecture

Figure 3 illustrates the packet forwarding configuration used in traditional software routers with single-queued NICs [15]. In this configuration, a separate input process (IN_x) and output process (OUT_x) is needed to move the packets between the Rx and Tx rings and the packet processing part (PP_x) of the software. Besides polling the packets in, the input process also classifies the packets and determines to which flows the packets belong to, and places them into the appropriate down-stream queue where the packet waits to be processed by the packet processing and forwarding function(s).

After a packet gets scheduled to be processed by the next “stage” (i.e., PP_x), which in the simplest case includes manipulating the TTL and checksum field of the IP header and modifying the link-level header, it gets enqueued into the second down-stream queue where it again waits until all the packets in front of it are removed by the output process, and the output process gets scheduled to dequeue the given packet and move it to the output queue of the card.

The disadvantages of this configuration can easily be spotted: (1) to forward a packet it needs to be processed by three separate software sections, thus it waits three times to be scheduled, resulting in high latency; (2) besides high latency, the throughput is also going to be unavoidably below what might be achieved by the given hardware architecture: on the one hand, because scheduling and context-switching

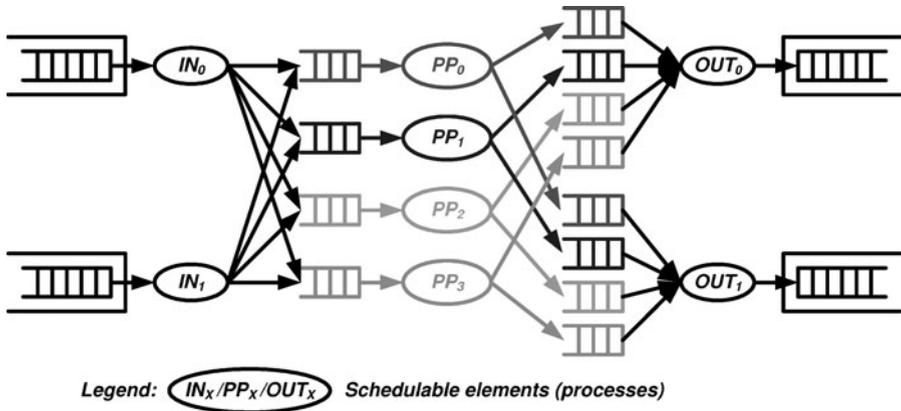


Fig. 3 Traditional software routers' forwarding architecture

of the three different processes (input, packet processing/forwarding, and output) are introducing extra cycles needed to forward a packet, while on the other hand for load-balancing purposes the three different processes handling a packet might be executed on different cores, and as a consequence the packets have to “switch cores” either via the shared last level cache (L3 in our architecture) or even worse, via the main memory; (3) in the case of flow-based resource management, as classification of the packets only happens after a packet has gone through the expensive input operations, and the input process has no control over which flow's packet to poll in next, it is practically impossible to guarantee a flow-based service, while the configuration is also exposed to DoS attacks; (4) resources need to be allocated very carefully to the schedulable processes, in order to avoid under-utilization by allocating resources in an unbalanced way to the processes, resulting in unnecessarily high packet loss. For example, allocating too many resources to the input processes results in a high number of packets which pile up and eventually are dropped from the queues downstream of the input processes; (5) last but not least, contention over the input and output ports occurs when separate cores want to access the same port at the same time.

2.3 Proposed forwarding architecture

A packet-processing workload involves moving lots of packets from input to output ports besides also undertaking different packet processing functions required. The question is how should we distribute this workload among the available cores to make optimum use of our multi-core architecture. With this aim, the three main objectives of our proposed architecture are basically: (1) to entirely bind all the operations needed to forward a packet to a single CPU core; (2) to achieve better parallelism in the execution; (3) and to create a configuration that allows us to more accurately measure the cost of forwarding and thus to more easily and accurately distribute the computational load among all the resources in the system. We discuss the first two objectives in more detail in this section, while the third one in Sect. 3.

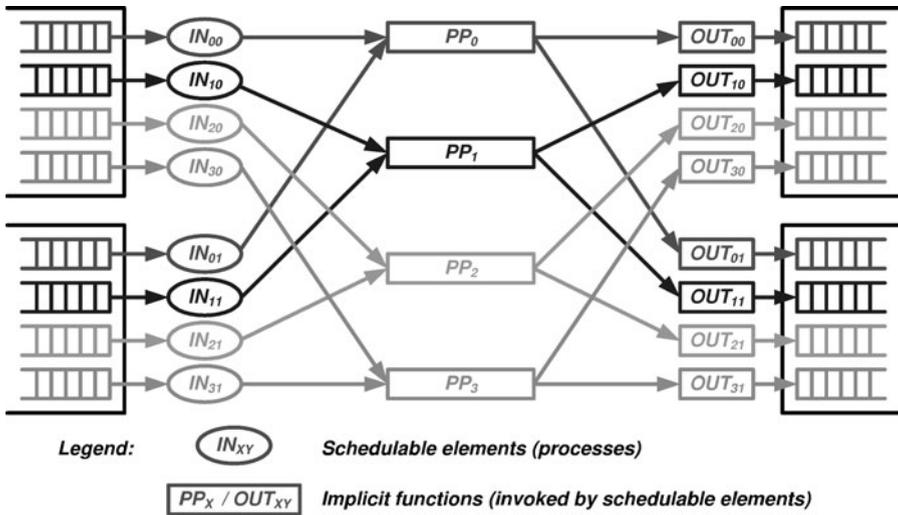


Fig. 4 Proposed software router forwarding architecture

We propose two broad changes to face the previously described challenges and to overcome the drawbacks of traditional forwarding architectures. First, we extend the Click architecture to leverage the multiple hardware queues available on modern NICs. For this, we developed a lock-free device driver for 10 Gbps multi-queue NICs and extended the Click modular router with multi-queue support both on the input (PollDevice) and output (ToDevice) side. Our multi-queued Click extension allows us to bind our extended polling and sending elements to a particular HW queue (as opposed to a particular port). As shown in [16], the use of multiple Rx and Tx queues and how the forwarding paths are distributed across the cores is key to improve both the throughput as well as the latency of our software router, while it also helps to perform more accurate resource management, as it will be demonstrated in the remainder of the paper. The second change is that we extended the ToDevice element so that it can receive packets from an up-stream Click element without the need of an internal packet queue in the forwarding configuration and without the need of scheduling it separately, further simplifying the complexity of the forwarding architecture as well as the scheduling computations (i.e., we extended the ToDevice element with a *push* input and took care that the batching of packets still works in the extended version, also).

Figure 4 shows our proposed forwarding architecture we created by using our extended Click elements. In this configuration, we have k input and k output queues, where k represents the number of CPU cores in the system (4 in the example in Fig. 4), or in a flow-based scenario k might represent the number of separate packet flow classes the router provides service for. When packets arrive on a network card, some of their Ethernet and/or IP header fields are looked up and a hash function decides into which queue they are placed. In the most widely-used case (i.e., Receive Side Scaling, RSS [8]), the packets are distributed in the available input queues based on a hash function, which function also ensures that packets belonging to the same

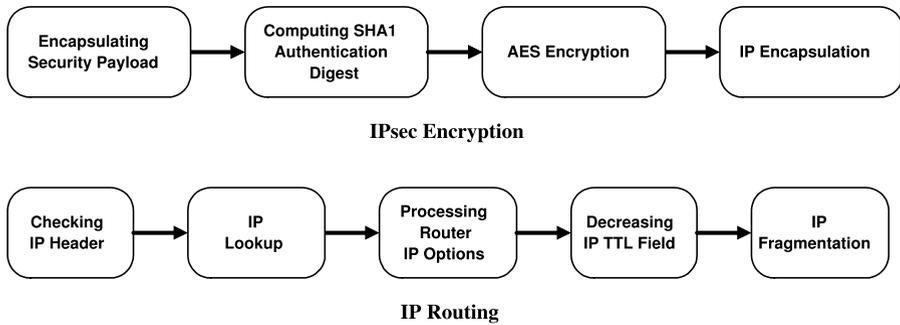


Fig. 5 Examples of basic packet processing functionalities

flow end up in the same input queue, thus avoiding packet reordering within a flow. In addition, packets can also be placed in the queues based on specific values of some header fields, thus providing a flow-based separation of the packets and supporting distinguished (QoS) packet processing. For example, by using VMDq [9] the packets can be classified by their destination MAC address, and a recently released commodity Ethernet Controller chip [17] supports Rx queue assignment based, among others, on L2 Ethertype, L3/L4 5-tuple or even TCP SYN bit fields, and also makes it possible to filter on any 2-byte tuple in the first 64 bytes of the packet.

To forward a packet at the front of a given input queue the multi-queue PollDevice element (i.e., the *IN* element in Fig. 4) needs to be scheduled at first. As a consequence, this input element polls the packet in from the queue and pushes it down for processing to the subsequent packet processing elements (represented by the PP_x rectangle in Fig. 4). After the packet has been processed by these elements, it gets finally pushed to our modified multi-queue ToDevice element, which moves the packet to the associated HW queue's Tx ring, from where the packet is DMAed directly to this hardware queue on the NIC without the interaction of the CPU, before finally transmitted out on the wire. In this configuration we need to only schedule every packet once (via the *IN* element) and after that the packets are forwarded in one step regardless of which output port they are routed to (the arrows among the *IN*, *PP*, and *OUT* elements represent simple function calls). This has a number of advantages that we discuss in parallel with the evaluation of the architecture in Sect. 2.4.

In Fig. 5, we show two practical examples of what the packet processing part (PP_x) of the forwarding path might consist of in the case of IPsec encryption and IP routing.

2.4 Evaluation

In order to illustrate the importance and advantages of our forwarding architecture, we created some simple experiments shown in Figs. 6, 7, 8, and 9. In all of these experiments, we performed only minimal forwarding, that is, traffic arriving at port i is directly forwarded to a pre-determined port j —there is no routing-table lookup nor any other form of packet processing. This minimal forwarding configuration is impor-

Fig. 6 Effect of packets switching cores VS they are processed all the way by the same core

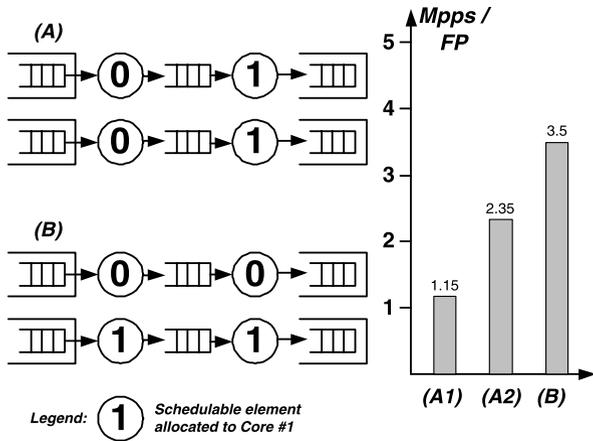
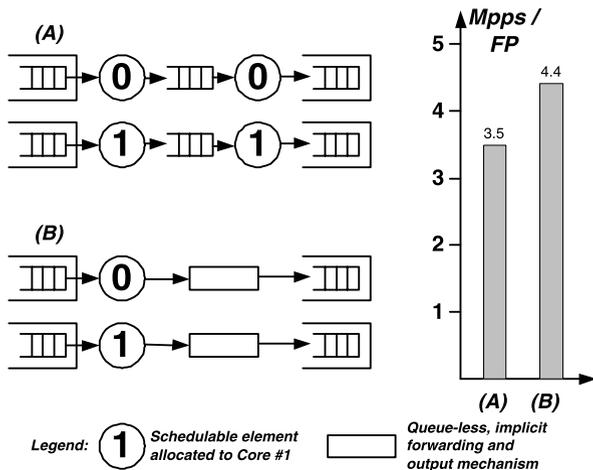


Fig. 7 2-phase VS 1-phase forwarding mechanism

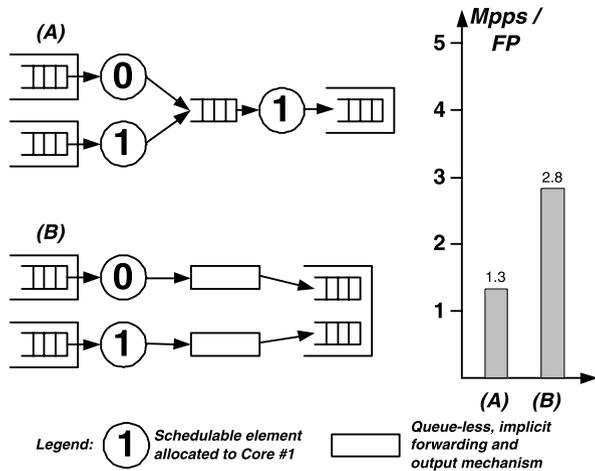


tant, because it only uses the minimal subset of operations that *any* packet-processing application incurs, and because there is no additional processing, the difference in forwarding rates of the different experiments clearly represent the impact of the changes we have applied to the forwarding paths. (Note, that we have labeled every element in Figs. 6, 7, 8, and 9 with the identity of the core processing that element.)

In [16], we have demonstrated how the input and output processing parts of the forwarding paths should be allocated to cores, and how hardware multi-queueing should be exploited for improved performance and parallelism. For completeness, we summarize and expand on these principles here.

With the experiments shown in Fig. 6, we demonstrate the significant performance loss that occurs when multiple CPU cores are involved in the forwarding of the packets. In both scenarios (A and B), we had two forwarding paths, but in scenario A a separate core is polling the packets from the input port than the one that is pushing them out to the output port.

Fig. 8 Advantage of multiple transmit (Tx) queues



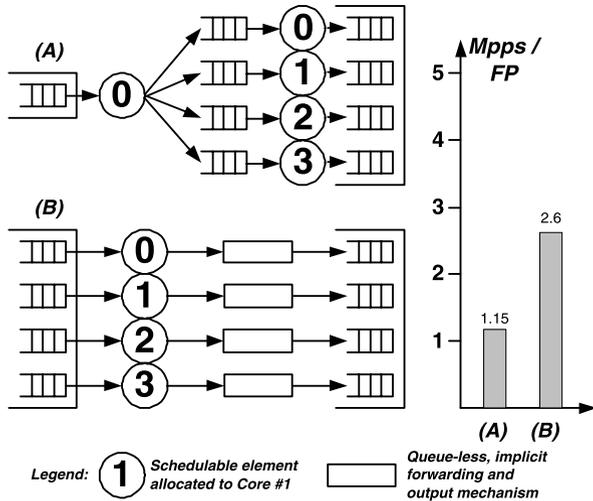
In scenario *A*, we ran two experiments, in the first case (*A1*) the two cores were residing on two separate sockets, thus not sharing any common cache, while in the second case (*A2*) they were on the same socket sharing the same L3 cache. In scenario *B*, we allocated the cores such that the same core is doing all the work for one forwarding path, while the second core for the other forwarding path, thus keeping the packets in the same L1 cache. As the results on the bar-graph show, the throughput is very low when the packets have to cross cores via the main memory (*A1*), it increases to approximately twice as high when packets cross via the L3 cache (*A2*), while it more than triples when the packets are forwarded exclusively by a single core (*B*).

With the experiments in Fig. 7, we demonstrate the performance gain that comes from using our 1-phase forwarding mechanism (experiment *B*) as opposed to the traditional 2-phase mechanism (experiment *A*). There are two factors that contribute to the increased performance. On the one hand, we have to schedule every packet only once instead of twice, while on the other, packets do not have to be enqueued into and then dequeued from an internal software queue. Both of these factors also contribute to a lower latency, which unfortunately we are not able to measure accurately due to the lack of an accurate latency measuring device.²

The experiment in Fig. 7 showed the performance improvement gained by enforcing packets to stay on the same core and by using the proposed 1-phase forwarding mechanism. However, this configuration seems unrealistic in the sense that every forwarding path has only one output port where its packets can leave the router. To overcome this, we have to make sure that every core has parallel access to every output port, which can simply be achieved by using multiple output queues on every port. More precisely, we have to allocate k queues on every output port, where k represents the number of cores in the system. The k queues are needed to avoid con-

²With Click, we were able to measure the average time a packet spends in the internal queue under high load, which is approx. 8600 cycles (3 μ s), but we could not measure the latency gain of reducing the scheduling iterations from two to one per packet.

Fig. 9 Advantage of multiple receive (Rx) queues



attention among different cores accessing the same output port, and to make sure that whichever core has polled a packet in is able to push it out without the need of an internal queue and a second scheduling phase. In Fig. 8, we show how important it is to have multiple output queues in order to avoid contention and to be able to keep the packets on the same core for all the time. As the bar-graph shows, the performance improvement is well above 100%.

As just illustrated, having multiple output queues helps to keep packets on the same core and to avoid output port contention, but the input ports are still unshared, which can lead to load-balancing problems. For example, some cores might be under-utilized while others might be overloaded due to unbalanced arrival traffic. In addition, if we have more cores than ports the packet processing has to be pipelined to exploit all the available cores, like in Fig. 9(A), which can result in core switching and a difficult resource allocation problem. To overcome these problems, we also need multiple input queues. By allocating k queues on every input port we can make sure that every core has parallel access to every input port, and thus it significantly simplifies the resource management. The results in Fig. 9 demonstrate how important multiple input queues are for high-performance packet forwarding. That is, the throughput, in terms of packet rate, more than doubles in our simple configuration, due to the fact that the load of forwarding the packets can more efficiently be distributed among the available cores, and packets do not have to cross cores; they are processed entirely by the same core. This also enables us to use our modified output element, further reducing the overhead of packet forwarding in software.

2.5 Dealing with overloaded Tx queues

As demonstrated in this section, our proposed forwarding architecture enables full parallelization of the resources, while it also improves performance by simplifying the basic structure of the forwarding paths, made possible by hardware multi-queues. However, the architecture in this form, as a consequence of the modified output element, is not equipped with the ability to deal with overloaded output ports. That is,

in the case a Tx queue on the network interface card is full, we cannot push any more packets to it, but because we have the packet already polled out from the Rx queue of the incoming interface, we cannot use that buffer to store the packet until the Tx queue gets, at least partially, emptied. The only thing we can do in such a case, is to throw the packet away (i.e., free its socket buffer in the memory). To overcome this problem, we propose to integrate an auxiliary packet queue into the output element and only enqueue packets in it which could not be moved to the hardware Tx queue (i.e., to the queue's Tx ring, from where the packets are going to be DMA-ed, without the interaction of the CPU, to the Tx queue on the card). This extension differs from the original forwarding mechanism (as shown in Fig. 7(A)) in the sense that the queue is only used when the output is overloaded and packets cannot be moved directly to it. However, as we have removed the output element from the scheduling chain, we have to make sure that whenever there are packets in this queue, the system makes sure that those packets are removed as soon as there is space in the Tx queue of the interface. To this end, we exploit the timer function available in the kernel-mode Click.

The flowcharts in Fig. 10 illustrate the process of enqueueing and removing packets from the auxiliary queue. That is, packets are enqueue whenever a packet cannot be pushed to the Tx ring belonging to the overloaded Tx queue, while packets are removed from the queue either when the timer expires or when a new packet is about to be pushed to output queue and there is space in it. In this latter case, the packets residing in the auxiliary queue are moved to the output at first in order to maintain packet order, and the newly forwarded packet is only pushed to the output when all the packets from the auxiliary queue have been removed; if this is not possible, it gets enqueue into the auxiliary queue as well.

3 Resource management

In the previous section, we showed the throughput and latency gains one can achieve with our proposed forwarding architecture. In this section, we are going to discuss the other advantages the architecture has to offer from the perspective of resource management.

The objective of proper resource management in software routers is to provide flexible and efficient flow differentiation, resource allocation, and scheduling in order to ensure fair resource sharing, isolation, as well as high performance. In Sect. 2.4, we demonstrated that we can achieve efficient flow differentiation by using hardware classification and multi-queueing available in current NICs. In this section, we show how the current Click scheduler and resource allocator might be improved by providing a more predictable system behavior, a guaranteed service (i.e., a guaranteed packet rate and/or bandwidth), higher per-flow and overall throughput and better fairness amongst the forwarding paths competing for the same resources.

3.1 Resource management in click

In Click, the base unit of scheduling are the *Tasks*. A *Task* is considered to be a chain of elements, which practically compose a part of a forwarding path (FP), that are

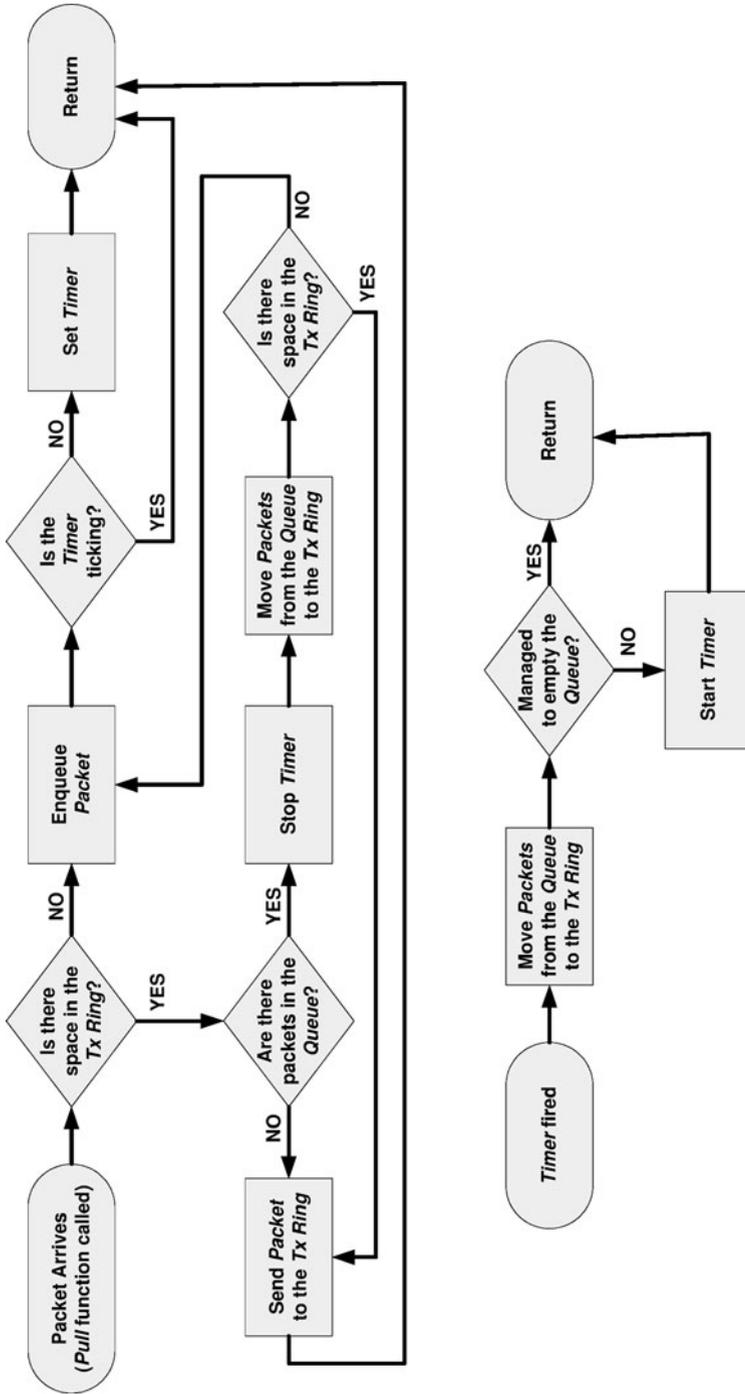


Fig. 10 Flowchart demonstrating the proposed operation of the queue integrated into the modified ToDevice element

traversed by a packet within a single scheduling step and it starts with an explicitly schedulable element. These explicitly schedulable elements have a so called *run_task* function that is called upon scheduling, after which the element executes its packet processing function(s) and hands the packet over to the next element downstream in the router configuration. Thus, these other elements are scheduled implicitly, by calling either their *push* or *pull* functions (i.e., they do not have any *run_task* function). A packet is processed and handed over to the next element until it either ends up in a Click queue inside the configuration or in one of the output ports' (or output queues') Tx ring. In these cases, the CPU is released and another element within the same thread, with a *run_task* function, at the front of a *Task*, is scheduled. These schedulable elements are responsible for either the input (e.g., *PollDevice*, *FromDevice*) or output processing (e.g., *ToDevice*) or for unqueueing the packets from an internal Click queue (e.g., *Unqueue*, *RatedUnqueue*, etc.). Note that in our proposed forwarding architecture the *PollDevice* element (i.e. IN_{XY} in Fig. 4) is the only schedulable element.

The CPU scheduler currently implemented in Click is based on the Proportional-Share (PS) Stride scheduling algorithm [18]. The implemented scheduler holds the following, mainly positive, characteristics that also apply to most PS schedulers in general:

- Each task reserves a given amount a resources (represented in *Tickets* in the Stride scheduling algorithm as well as in Click) and it is guaranteed to receive at least this amount of resources when it is not idle. In general, PS schedulers, resources reserved for a *Task* usually mean a given number of CPU cycles-per-second. In current Click, the reserved resources mean the number of cycles it takes to execute a *Task*. Hence, because in Click the *Tasks* are non-preemptive they can hold the CPU, and thus put load on the other resources (e.g., memory and I/O), as long as they need, regardless of how long it takes to execute them. As a consequence, *Tasks* in Click with the same amount of tickets are scheduled equally often, but because of their different resource needs they consume different amounts of the different resources. This way of scheduling results in an unfair resource usage, for which we propose a solution below.
- An idle *Task* cannot “save tickets” to use it when it becomes active.
- For work-conserving purposes, tickets unused by idle *Tasks* are distributed among the active *Tasks* within the same thread. These active *Tasks* are not charged for these extra resources.
- To do a context-switch between *Tasks* is an expensive operation, therefore it is recommended to avoid scheduling an idle *Task* if possible. To this aim, Click implements an *adjust_ticket* mechanism, that dynamically changes the actual ticket value of a *Task* between 1 and its allocated ticket value (1024 by default) in correlation with the number of packets processed recently by the *Task* in question.

All the *Tasks* within the same kernel thread are co-scheduled relative to each other's share of the CPU (i.e., a *Task*'s *Ticket* relative to all the other *Tasks*' *Tickets* that run within the same kernel thread). That is, if all the *Tasks* have the same *Ticket* they are scheduled equally often, and in practice this resembles to that of a Round Robin scheduler. However, if a *Task* has different amount of *Tickets* it will be scheduled in proportion to its *Ticket* value relative to other *Tasks* *Ticket* value. That is, if

one *Task* has twice as many *Tickets* as another one it will be scheduled twice as often. This method is realized by using a virtual time line, and the time a *Task* should be scheduled to run again is calculated in the following manner:

$$NEXT_TIME = CURRENT_TIME + (1/TICKET * CONSTANT_NUMBER)$$

The *Tasks* are scheduled in a serial order, always running the one with the lowest virtual time value. *Tasks* running on separate CPUs (i.e., within separate kernel threads) have no influence on each other, as they are put into separate working queues and are scheduled separately.³

As far as resource allocation is concerned, Click currently provides two mechanisms for that. On the one hand, it lets the user to statically define which *Task* should be run within which kernel thread. This method is useful for experimentation, but it is very limited when it comes to dynamic resource management in real forwarding scenarios. On the other hand, Click also provides a method where the recent CPU usage of every *Task* is measured, and based on this measurement the *Tasks* (if needed) are redistributed among the available threads for load-balancing purposes.

Although this latter method provides more sophisticated resource management than the static method, it still lacks a number of significant features, such as, it does not take into account the cache hierarchy, it only measures the recent cycle usage of the *Tasks* also including in the measurements the cycles used when the *Task* was practically idle and consumed only a few cycles, while it does not maintain any statistics about the real cost of forwarding a packet by any given *Task*. In addition, it does not use any other metric, besides CPU cycles, to describe a *Task*.

3.2 Resource guarantees

In order to guarantee that all the forwarding paths get access to the amount of resources they are entitled to and no more,⁴ the total cost of each forwarding path must be monitored and the CPU scheduling parameters tuned accordingly. Note that we need to monitor the cost of each forwarding path regularly to ensure that the current CPU scheduling adjustments reflect the recent resource usage of all the forwarding paths, including the ones where the per-packet processing cost can vary significantly depending on the character of the packet in question. To this end, we extended Click's original CPU scheduler by adding the ability to measure the cost of each forwarding path and adjusting the weight of each of them according to their costs, thus achieving fair CPU scheduling.⁵ In addition, we can also change the quantity of CPU resources

³Of course, *Tasks* running on separate CPU cores might have effect on each other's operation in the case the cores, the *Tasks* are run on, share some other resource (e.g., cache memory, network interface card, etc.) but such factors are out of focus of the Click scheduler. To deal with cache contention one might incorporate a method similar to that of proposed in [19] into the Click scheduler, but for now such modifications are out of the focus of this paper.

⁴Unless, of course, there are more resources available, in which case *Tasks* can use more as long as they do not adversely affect the operation of others. This way of scheduling is called work-conserving scheduling.

⁵Note that the resource consumption for forwarding a packet can only be measured when the packet exits the forwarding path and control is returned to the Click scheduler.

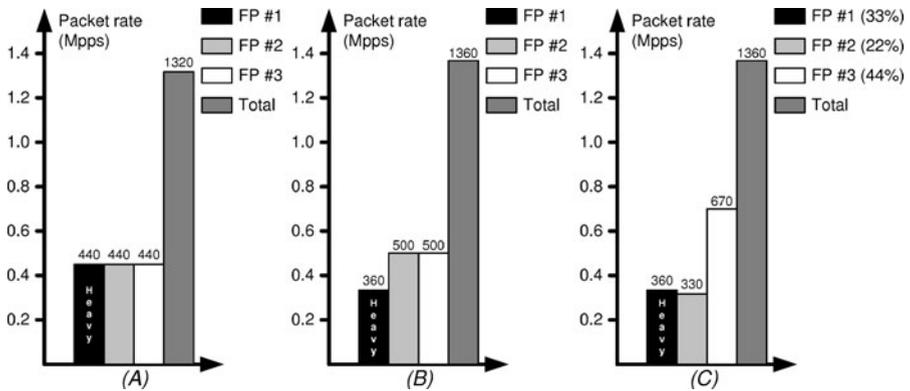


Fig. 11 Resource usage with the current Click scheduler

that the given FPs are entitled to. The following pseudo-code describes the procedure the *Tickets* are calculated based on the *Tasks*' cost and weight values.

Procedure of calculating the Tickets of Tasks based on their cost and weight values:

VARIABLES:

N	The number of <i>Tasks</i> in the thread
W_i	Weight of <i>Task</i> i ($i = 1, \dots, N$) (defined by the administrator) $W_i = 1.0$ by default
C_i	Cost of <i>Task</i> i ($i = 1, \dots, N$) (measured on by the system)
T_i	Ticket value of <i>Task</i> i (to be calculated)
$DTV = 1024$	Default <i>Ticket</i> Value

begin

$C_{total} = 0$

for all i in N do

$C_{total} += C_i$

for all i in N do

$T_i = DTV * C_{total} / (N * C_i) * W_i$

end

In an attempt to mitigate the different forwarding costs associated with different packets, whilst the sampling frequency is tunable, we sampled at a rate of one in every twenty packets to balance the sampling overhead against the fidelity of the measurements. To illustrate the scheduling properties of the configurable virtual forwarding plane, we ran an experiment using a configuration with three FPs sharing the resource of a single CPU, where two of the FPs were identical and the third one included a few expensive Click elements to increase the cost of forwarding a packet. Results in Fig. 11(A) show that when using the default Click scheduler we achieved the same forwarding rate for each of the FPs (approximately 440 Kpps), meaning that the CPU was not shared equally as the more expensive FP managed to forward the same number of packets per second as the others. Figure 11(B) illustrates the throughput of the FPs when our scheduler extension was used to load-balance the FPs' CPU usage. The

results show that load-balancing works properly, since the packet rate of the expensive FP decreased while that of the other two FPs increased; it takes approximately 40% more CPU cycles to forward a packet by FP1 than by FP2 and FP3. Figure 11(C) shows that weighted fair-sharing is working in our extended scheduler as well. In this experiment, we left the share of the expensive FP at 1/3 of the total, while moved 1/3 of FP2's resources to that of FP3, thus FP3 was entitled to twice as much resources (i.e., 44%) than FP2 (i.e., 22%). As we can see, FP1 kept its rate where it was in Fig. 11(B), but FP3 increased by about 50%, while FP2 dropped approximately with the same rate.

3.3 Scheduling based on different resources

In previous work [2, 3, 5], we have shown that in multi-core architectures the bottleneck of packet processing can as easily be at the memory or I/O as at the CPUs, and hence resource allocation needs to be able to handle the case where any of these is the scarce resource and not the CPU. Therefore, we find it necessary to not only measure the CPU cycles used by the *Tasks*, but also other metrics describing the utilization of resources, like memory and I/O bandwidth, cache miss rates at any level of the cache hierarchy, instructions-per-second, etc. and then use these metrics for resource allocation and scheduling calculations. To this end, we use performance monitoring counters to determine the utilization of the resources. These hardware counters are commonplace on modern processors, they are very low overhead counters and were originally introduced for profiling and performance analysis and are therefore suitable for our purposes, that is, to measure the cost of forwarding a packet on multiple resources. In our hardware architecture, every Hyper-Thread (there are two Hyper-Threads per core when Symmetric Multi-Threading is enabled) possesses 3 fixed and 4 customizable counters, but there are over 500 events these counters can measure, using time-division multiplexing. Scheduling memory or I/O resources is not as simple as scheduling CPU cycles. However, in scenarios where the memory or I/O has become the bottleneck, we can adjust the CPU scheduler's parameters so as to ensure a fair memory access rate by controlling the forwarding paths' CPU resources, for example, if a forwarding path is scheduled to use CPU resources less often it will access the memory and/or I/O less often, also.

One additional great advantage of our proposed forwarding architecture is that it makes it possible to measure the resource utilization of packet forwarding very accurately and cost effectively. This comes from the fact that every packet is forwarded from the input to the output port in one "go" and on a single core, that is, we have to read and store the values of the counters only once per packet and only on a single CPU. To get the correct resource utilizations we recommend to reset the counters just before a *Task* gets scheduled and read them just after the *Task* finishes running or take the difference of two consecutive readings taking care of occasional counter overflows.

To demonstrate the resource usage of different scheduling principles (i.e., default Click, fair CPU, and fair memory scheduling),⁶ we at first measured the resource

⁶Scheduling based on other metrics might be used as well.

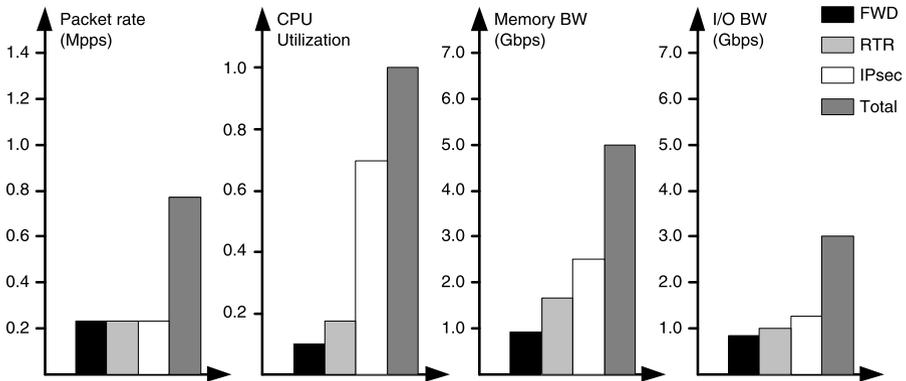


Fig. 12 Resource usage with the current Click scheduler

utilization of different workloads. In the second step, we calculated the resource usage of the demonstrated scheduling principles based on the previously measured resource utilizations. The workloads we use for evaluations in the remainder of this section are the followings:

- (1) *Minimal forwarding (FWD)*: This is the same application that we used in Sect. 2.4, that is, traffic arriving at port i is directly forwarded to a pre-determined port j —there is no routing-table lookup nor any other form of packet processing.
- (2) *IP routing (RTR)*: We implement full IP routing including checksum calculations, updating headers and performing a longest-prefix-match lookup of the destination address in an IP routing table. For this latter, we use Click’s implementation of the popular D-lookup algorithm [20] and, in keeping with recent reports, a routing table size of 256 K entries. For synthetic input traffic, we generate packets with random destination addresses so as to stress cache locality for IP lookup operations.
- (3) *IPsec packet encryption (IPsec)*: Every packet is encrypted using AES-128 encryption, as is typical in VPNs.
- (4) *CRC calculation (CRC)*: 32 bit CRC is calculated over the whole packet and appended to the end of it.

Our selection represents commonly-deployed packet-processing applications that are fairly diverse in their computational needs. For example, minimal forwarding stresses memory and I/O; IP routing additionally references large data structures; encryption is CPU-intensive; while CRC is memory-intensive.

Figure 12 shows the achieved packet rates and resource utilization with 64 B packets of three different forwarding paths/workloads (FWD, RTR, and IPsec) co-scheduled on the same CPU by the default Click scheduler. As this default scheduler does not take into account any cost parameters, the FPs are considered to be equal and are scheduled equally often, thus resulting in the same throughput. However, as every FP uses different amounts of resources their individual resource consumption are significantly unbalanced and unfair. To overcome this fairness issue, Figs. 13 and 14 show the packet rates and resource consumptions in the case of a fair CPU and a fair

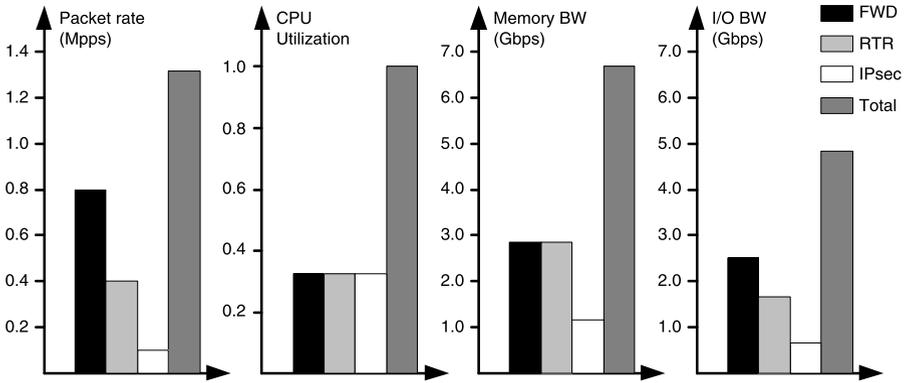


Fig. 13 Resource usage with fair CPU scheduling

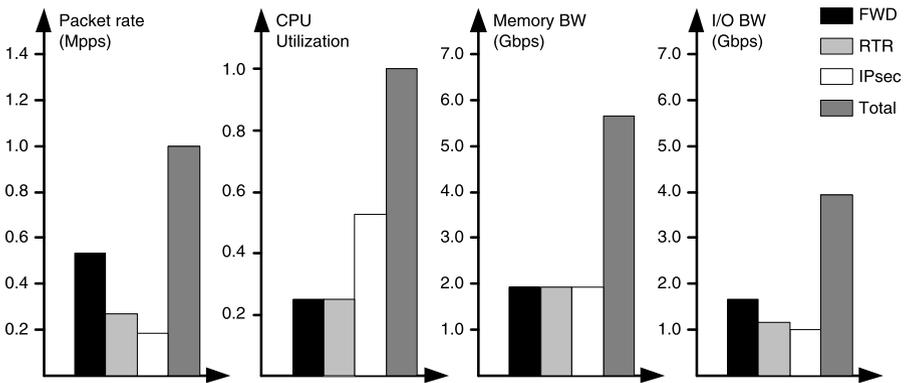


Fig. 14 Resource usage with fair memory scheduling

memory scheduler, respectively. As the graphs show, the packet rate of the cheaper workloads (i.e., FWD and RTR) increases while for IPsec it decreases accordingly to their CPU and memory needs.

As one might notice, all the resources except the one the scheduling is based on are unbalanced. Although, it is possible to devise a scheduling algorithm that equally balances the load over multiple resources, we find it unnecessary. The main reason for this is that in a real-time scenario there will be always one resource at a given time that is mainly responsible for the contention. Although, this resource might change over time due to the change of the arriving traffic to the FPs, it is simpler and more efficient to equip the scheduler with the ability to schedule the *Tasks* based on different resources, but base the scheduling on only one at a time, instead of all of them, including the ones below their saturation point. That is, in an overload scenario if we schedule and fair-share the saturated resource, we ensure that the *saturated* resource is used equally, which might result in unequal usage of the other resources, but as they are still below their saturation point, this unequal resource usage will not cause any performance drop for any of the FPs.

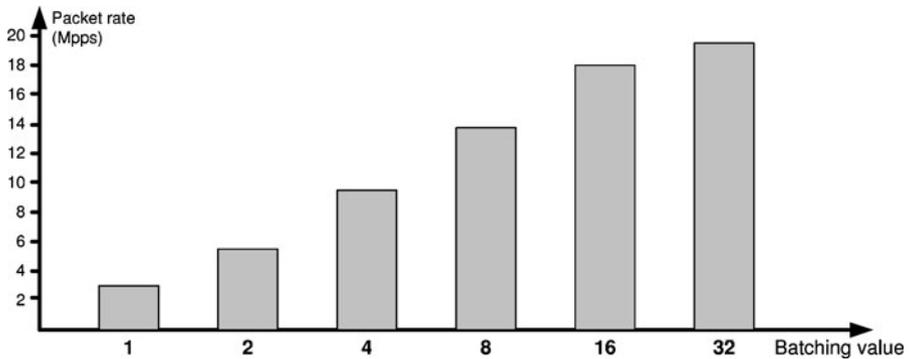


Fig. 15 The advantage of batch processing in software

3.4 Granularity and frequency

In Figs. 13 and 14, we showed the effect of scheduling the *Tasks* based either on their CPU cycle or memory bandwidth consumption. The question “how this might be implemented?” comes naturally. One part is, what we have already described previously, to measure the resource consumption of every *Task*, while the other is to set when the packets should be scheduled. Scheduling a packet happens via the *Task* that is allocated to the hardware input queue the packet resides in. For this latter part, we have two tools available, one is the *Ticket* value of the *Tasks* while the other is the *BURST* value of the schedulable elements standing at the front of the *Tasks* (i.e., *PollDevice*, or *IN_{XY}* in Fig. 4). In the current implementation, both of these parameters are static (i.e., *Ticket* = 1024, *BURST* = 8), but with some simple extensions it is possible to set them dynamically by the scheduler. The *Ticket* value is the parameter that directly represents how often a *Task* should be scheduled relative to the others. The higher this value is the more often the *Task* is going to be scheduled (in a linear proportion). The *BURST* value is the parameter for batching. That is, when a schedulable element is scheduled it processes as many packets as there are available up to the value of the *BURST* parameter. The advantage of batching is well-known; it reduces the overhead of context-switching, and thus improves the performance.

To demonstrate how significant batching is in Click, we ran experiments with minimal forwarding workload and 64 B packets with different batching values. Figure 15 shows the performance gain as the function of the increasing batching value. Besides demonstrating the performance gain, this set of experiments also helps us to determine the exact cost of context-switches. Using linear-regression, we found that it takes approximately 6100 cycles to perform a context switch in Click on our machines, which includes recalculating the virtual time when the currently finished *Task* should be run again, putting the *Task* in the right place of the working queue and fetch the next *Task* from this working queue and run it. The cost of the context-switch is important to know for accurate resource management, as this value has to be included in the cost of forwarding a packet.

However, it is important to note that with batch processing of the packets, we decrease the responsiveness of the system and make the scheduling granularity coarser,

which might not be favorable after a certain point. That is, the administrator of the system needs to determine a maximum value, preferably in terms of CPU cycles, for how long a *Task* can keep the resources when it is scheduled, in order to avoid too coarse switching of *Tasks* resulting in poor responsiveness. Let R represent this maximum value. To this end, we recommend to determine the highest burst rate for every *Task* and use that burst rate every time the *Task* gets scheduled. The maximum allowed burst value of every *Task* can be determined according to the following equation:

$$BURST_i = \left\lfloor \frac{R - C_{CS}}{C_{C_i}} \right\rfloor;$$

where C_{CS} indicates the cost of a context-switch (i.e., 6100 cycles in our case), and C_{C_i} the CPU cost of *Task*_{*i*}.

3.5 Unified input elements

As we just demonstrated in Sect. 3.4, reducing the number of context-switches (i.e., performing batch processing) results in a significant performance improvement, due to the fact that performing a context-switch takes a high number of CPU cycles. To this end, we propose another modification to the forwarding architecture in order to further decrease the number of context-switches and to improve the architecture's performance. By looking at the proposed architecture in Fig. 4, one can see that there are multiple schedulable input elements accessing separate input queues that are executed on the same core (e.g., IN_{00} and IN_{01} , or IN_{10} and IN_{11}). By extending our input element with the ability to poll packets from multiple queues residing on separate network interface cards, we managed to further reduce the number of schedulable elements, and as a consequence the number of context switches needed to forward packets from multiple input ports. With this modification, the reduction in the number of schedulable elements is in the order of X , where X represents the number of different interfaces a unified, multi-input element polls packets from.

Figure 16 shows the forwarding architecture with the unified input elements. In this example, similarly as before, we have two input and two output ports processed by four cores. As a consequence of having two input ports, the number of schedulable elements becomes half of what it was before (see Fig. 4). In order to demonstrate the advantage of our unified input element, we ran experiments comparing its performance with that of the originally proposed element in Sect. 2.3. In these experiments, we had four interfaces being polled by a single core, in the first case using four single-input elements, while in the second case using a single multi-input element. We ran these experiments with increasing batching value and calculated the performance improvement in terms of packet rate increase. Figure 17 shows the increase in packet rate of the configuration with the unified element compared to the one with the single-input elements. As we can see, the performance gain varies between 20% and 30%, depending on the batching value for this four interface scenario. However, with higher number of interfaces, we expect to have a bigger performance improvement, due to the more significant decrease in the number of schedulable elements as well as context-switches per packets processed.

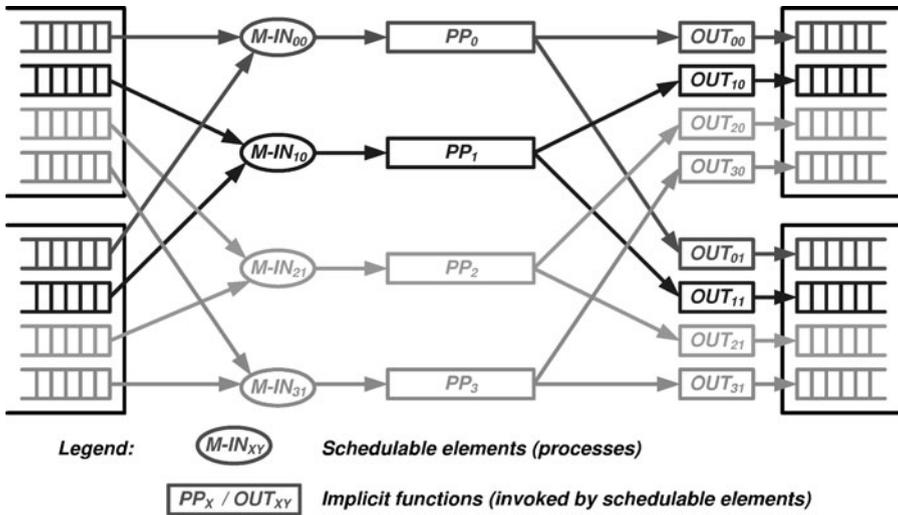
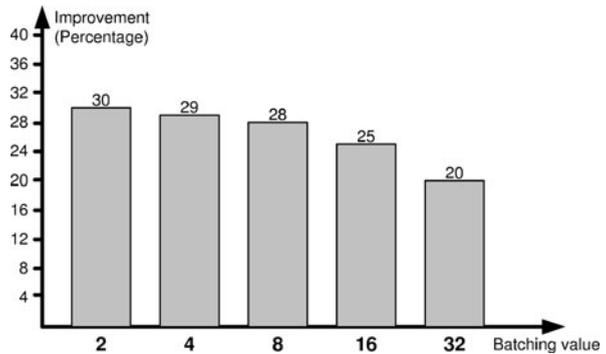


Fig. 16 Proposed forwarding path architecture with unified input elements

Fig. 17 Improvement in percentage with four interfaces polled by one core



In Fig. 18, we illustrate in a simple way how this performance improvement is achieved. The top line represents the time it takes to process eight batches of packets from four interfaces with the configuration using the originally proposed, non-unified input elements. To perform this task, we need four input elements (i.e., PollDevice element), one for every interface. The second line illustrates the time needed to process the same eight batches of packets from four interfaces, but now by the unified input element with multiple inputs, one per interface. In our unified input element to move from one interface to another in order to poll packets in, is a very simple and cheap operation, while to do the same with the single-input elements, it becomes a significantly more expensive operation. That is, it costs us a whole context-switch, which is approximately 6,100 cycles on our experimental machine, as reported before. As the number of packets polled in a batch increases, the difference in performance between the single-input and multi-input configuration decreases. However, as we saw in Fig. 17 even with a batching value of 32, the difference is 20% with only four ports. Now, it can also be easily seen why the performance improvement would increase

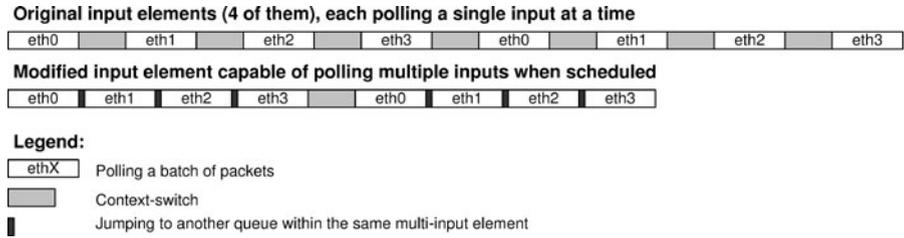


Fig. 18 Illustrating the time it takes to process eight batches of packets

with a higher number of interfaces being polled by a single multi-input element. On the other hand, as we have no control over the arriving traffic patterns, we cannot make sure that there will always be a high number of packets in every queue, or any number of packets at all. Thus, when there are a low number of packets present in the input queue(s) of one (or more) interface(s), the ratio of CPU cycles spent on actual packet processing over the cycles spent with context-switches decreases, resulting in decreased performance with the non-unified configuration. This is also confirmed by the results in Fig. 17, where the lower batching value scenarios could also be interpreted as there had been only a lower number of packets in the input queues of some interfaces.

3.6 Co-scheduling principle

Today’s general purpose operating systems adapt timeslice-based multi-processor scheduling to be aware of resource contention and take steps to mitigate it. This means that timeslices need to start at the same time on each core, thus such schedulers require synchronization across all the cores.

As described above, the Click *Tasks* cannot be preempted, eliminating the possibility of timeslice-based scheduling. This has the advantage of better resource utilization compared to timeslice-based scheduling. Namely, when in the latter case a process terminates earlier than when its timeslice would have finished, the given CPU is going to be idle for the remaining time of the timeslice. In Click’s scheduler, the resources are handed back as soon as a *Task* finishes processing the packets. However, because of the lack of synchronized timeslices across all the cores, it is virtually impossible to have a global scheduler that can accurately co-schedule exact *Tasks* running on separate cores, which might lead to unpredictable system behavior as there is no control over what tasks, with what characteristics are going to run at the same time.

To overcome this unpredictability, we propose to co-schedule *Tasks* with the same or similar characteristics on the same core (*Principle #1*), thus ensuring that the given core is always using approximately the same amount of shared resources (i.e., memory, cache, I/O) regardless which *Task* is running at a given time.

To demonstrate the importance of this co-scheduling principle, we ran three experiments, each with 4 forwarding paths, 2 IPsec and 2 CRC, distributed on 2 cores, each *Task* (FP) being allocated 50% of the CPU resources, emulating a scenario where the CPU is the bottleneck. Figure 19 illustrates the total memory and I/O utilization for the three experiments (as well as the CPU utilization, but because both cores were

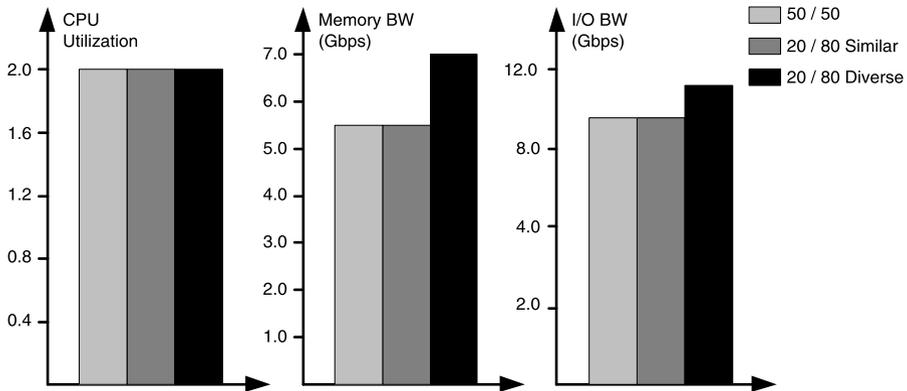


Fig. 19 Resource utilizations when complying and when not complying with Principle #1

running at 100% all the time, this has less relevance). The first bars (called “50/50”) are the baseline numbers, in which experiment the two IPsec FPs were sharing a core, and the two CRC FPs were sharing the other. The arriving packet rate for every FP was identical and consisted of only 1024 B packets. In the second experiment (called “20/80 Similar”), the FP-to-core allocation was the same, but on both cores one of the FPs received only 25% of the other FPs packet rate, resulting in a 20/80 resource usage due to the proportional-share scheduling mechanism. As the results show, the total memory and I/O utilization has not changed, due to the fact that the *Tasks* sharing the same core have the same resource usage characteristics. In the third experiment, we allocated one IPsec and one CRC FP to every core and generated 20% of the traffic on both cores to the IPsec FP, while 80% to the CRC. As the results show, because the CRC workload has different characteristics (higher memory and less CPU needs) it used more memory and because of the higher packet rate more I/O bandwidth as well. This phenomenon can be critical in the case when the CPU is the resource saturating at first, but the memory is near its saturation point as well, and when the arriving traffic pattern changes the bottlenecked resource might fluctuate between CPU and memory.⁷

3.7 Utilizing hyper-threading

Recent Intel processors support the Hyper-Threading (HT) Technology [21] (a.k.a Simultaneous Multi-Threading (SMT)), delivering thread level parallelism for more efficient use of the resources and higher processing throughput. Our server architecture supports two threads per core, thus in total 16 threads. The key behind the improved performance of Hyper-Threading is the ability to fetch instructions from multiple threads in a cycle, thus instructions from more than one thread can be executed in any given pipeline stage at a time. However, to really see improvement in the performance with SMT switched on, we found that careful resource allocation

⁷The saturation point of the resources can be determined by benchmark experiments.

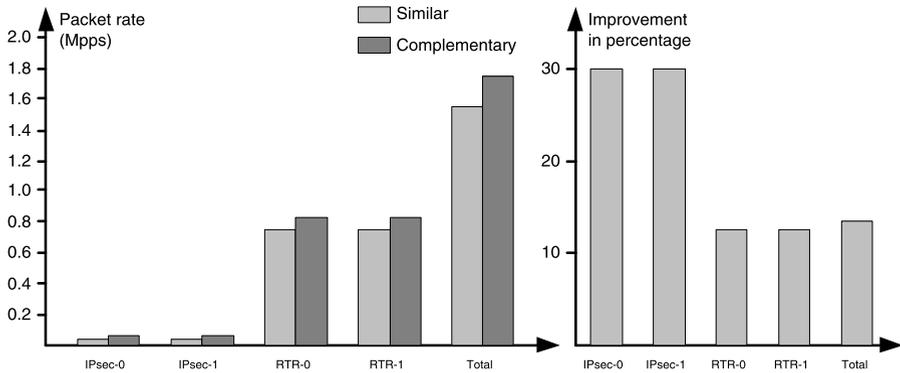


Fig. 20 Performance improvement when complying with Principle #2

is needed in the case of a parallelized software router architecture. Our finding is that *Tasks* with more diverse (or complementary) characteristics executed on adjacent threads⁸ result in more significant performance improvement than *Tasks* with similar characteristics (*Principle #2*). Figure 20 shows the results of our experiments on this. In these experiments, we had 4 FPs, 2 IPsec, and 2 RTR, and 2 cores each with 2 threads, thus 4 threads in total. In the first case (called “Similar”), we allocated the identical FPs to the adjacent threads, while in the second case (called “Complementary”) one thread on every core was executing an IPsec FP, while the other a RTR FP. The left graph in Fig. 20 shows the packet rate for every FP, while the right graph shows the improvement in the packet rate of the *Complementary* experiment over the *Similar* one in percentage. As we can see for the IPsec FPs, the improvement is nearly 30%, for the RTR it is nearly 12%, while the total packet rate improvement is slightly above 12%. This improvement comes from the fact that the adjacent threads were executing workloads with different characteristics (i.e., a CPU-intensive IPsec, and a less CPU-intensive, but more memory-intensive RTR), and while their instructions were stressing separate resources the parallelism provided by Hyper-Threading managed to more efficiently utilize these resources (i.e., while the RTR workload was waiting for data to be fetched from the memory the IPsec workload managed to carry on with execution in parallel).

4 Forwarding architecture’s applicability in virtual routers

Over the last years, several publications appeared indicating the increased interest in router virtualization [5, 22–26], and in the potential that all of this might be carried out on $\times 86$ commodity hardware without moving back in performance compared to commercial routers’ performance.

The benefits of virtual routers (VR) are obvious: a single hardware platform can simultaneously perform independent routing for multiple networks in a manner that

⁸We call the two threads running in parallel on the same core “adjacent threads.” Some literature call them “physical” and “logical” threads.

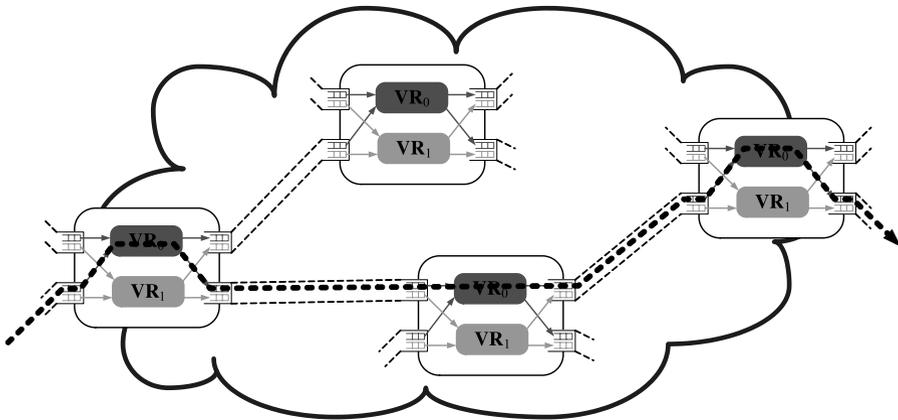


Fig. 21 Multiple Virtual Networks on top of a single physical infrastructure built with Virtual Routers

permits independent management, full administrative control, as well as potentially full customization of the virtual networks. A practical example for this is a number of small businesses within the same building sharing the same router, each managing its own slice of the router. In addition, virtual routers can also be a substrate for innovation in network architecture design, thus enabling the evolution of the Internet. Virtual routers can also induce different service models by allowing the physical network to be decoupled from the providers, and making it possible to virtualize an entire network. For example, Fig. 21 illustrates a fraction of a network that consists of virtual routers shared by two service providers. Each of the providers has a slice of the routers' resource: they access a set of dedicated queues on the network interface cards, and manage a dedicated fraction of the available CPU and memory resources. By ensuring adequate isolation among the virtual routers running on top of the same physical hardware, traffic forwarded by any of the virtual routers has no influence on the operation of the others.

The authors in [3] and [5] propose a VR architecture, where the forwarding planes are separated from the control planes and are run near the physical devices, all executed in the privileged driver domain of the Xen Virtual Machine Monitor [27], while the control planes are running separately in the virtualization domains of Xen as illustrated in Fig. 22.

The advantage of moving all the forwarding paths into a single driver domain near the device drivers is that there is no performance degradation compared to the hardware's native performance, what is usually so significant with network virtualization as demonstrated in [28]. However, with this we lose the isolation and fairness provided by the virtual domains, which we are able to gain back with our proposed forwarding architecture and resource management. To this end, we only have to realize the analogy between the architecture proposed in Sects. 2.3 and 3.5, and the virtualized forwarding architecture introduced in [3] and [5]. That is, consider Fig. 16, where every input and output queue, schedulable input and packet processing elements, as well as every output element designated with a separate shade can be considered as a separate virtual router's forwarding engine. That is, every VR has a slice from every

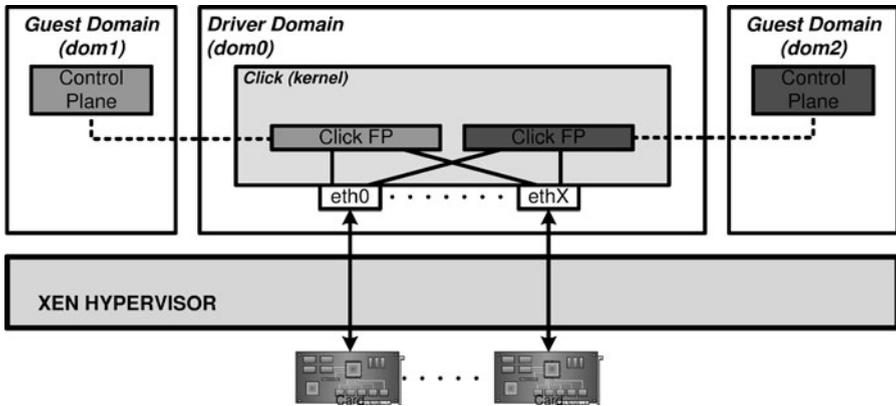


Fig. 22 A well-performing virtual router architecture

input and output port, it can realize any packet processing function inside the PP_x part, and the resource allocation and scheduling can be carried out the same way as if the forwarding paths were paths with guaranteed service, thus ensuring a guaranteed amount of resources as well as isolation among the virtual forwarding engines. In addition, exploiting the VMDq (Virtual Machine Device Queue) [9] support in current NICs further justifies the architecture's applicability in virtual routers. That is, the VMDq mechanism classifies the arriving packets based on their destination MAC address, which address can be used to identify the virtual router the packet flow belongs to and the packet can be placed into the input queue allocated to that given virtual router. The intermittent thick line in Fig. 21 illustrates how a packet flow crosses a network via a set of virtual routers belonging to the same service provider. The packets upon entering the network are placed into the appropriate input queue belonging to the given provider's virtual router on the given ingress router in the network. The packets are polled in by the virtual router's forwarding engine, they are processed by it and their next hop router is determined; after which the proper destination MAC address is set to ensure the packets end up in the input queue of the same provider's virtual router on the next-hop router. Finally, the packets are placed into the dedicated output queue of the virtual router on the interface connected to the next-hop router. The provider's next-hop virtual routers are performing the same operation on the packet flow until they reach their final destination or leave the provider's network.

5 Conclusion

The performance of modern multi-core commodity architectures clearly indicates its viability for high performance packet forwarding. However, as we demonstrated in this paper, a software router platform has to be designed in a way that provides effective parallelism and proper resource management for workloads with diverse resource characteristics.

Our proposed forwarding architecture enables full parallelization of the resources, while it also improves performance by simplifying the basic structure of the forwarding paths, made possible by hardware multi-queues.

Our results in Sect. 3 clearly indicate that we need to pay attention to the utilization of multiple resources, and not only the CPU, when allocating resources to different workloads. For predictable system behavior, we identified *Principle #1*, recommending to co-schedule workloads with similar characteristics within the same thread, while for improved performance co-schedule workloads with different characteristics on adjacent hyper-threads (*Principle #2*). Our proposed architecture is also capable of performing flow-based packet processing, while it can also be used as a forwarding plane of virtualized software routers.

References

1. Vyatta Open Networking (2010) Software-based Routing & Security. <http://www.vyatta.com/products/>
2. Argyraki K, Baset SA, Chun B-G, Fall K, Iannaccone G, Knies A, Kohler E, Manesh M, Nedveschi S, Ratnasamy S (2008) Can software routers scale. In: ACM presto workshop, Aug 2008
3. Greenhalgh A, Handley M, Mathy L, Egi N, Hoerd M, Huici F (2008) Fairness issues in software virtual routers. In: ACM presto workshop, Aug 2008
4. Ye Q, MacGregor MH (2007) Click on a cluster: a viable approach to scale software-based routers. In: IEEE ICC'07
5. Egi N, Greenhalgh A, Handley M, Mathy L, Hoerd M, Huici F (2008) Towards high performance virtual routers on commodity hardware. In: ACM CoNext, Dec 2008
6. Intel 10 Gigabit XF SR Server Adapters. <http://www.intel.com/products/server/adapters/10-gbe-xfsr-adapters/10-gbe-xfsr-adapters-overview.htm>
7. Next Generation Intel Microarchitecture (Nehalem). http://www.intel.com/pressroom/archive/reference/whitepaper_Nehalem.pdf, 2008
8. Receive side scaling on Intel Network Adapters. <http://www.intel.com/support/network/adapters/pro100/sb/cs-027574.htm>
9. Virtual machine device queues. http://www.intel.com/technology/platform-technology/virtualization/VMDq_whitepaper.pdf, 2007
10. Kohler E, Morris R et al (2000) The click modular router. ACM Trans Comput Syst 18(3):263–297
11. Chen B, Morris R (2001) Flexible control of parallelism in a multiprocessor PC router. In: USENIX technical conference
12. Bolla R, Bruschi R (2006) Ip forwarding performance analysis in the presence of control plane functionalities in a pc-based open router. In: Springer distributed cooperative laboratories: networking, instrumentation, and measurements. Springer, Berlin, pp 143–158
13. VTune Performance Analyzer. <http://www.intel.com/cd/software/products/asmo-na/eng/239144.htm>
14. OProfile—A System Profiler for Linux. <http://oprofile.sourceforge.net/news/>
15. Qie X, Bavier A, Peterson L, Karlin S (2001) Scheduling computations on a software-based router. In: Proceedings of ACM SIGMETRICS 2001, Cambridge, MA, USA
16. Dobrescu M, Egi N, Argyraki K, Chun B-G, Fall K, Iannaccone G, Knies A, Manesh M, Ratnasamy S (2009) Routebricks: exploiting parallelism to scale software routers. In: Proc of 22nd ACM symposium on operating systems principles, Big Sky, MT, Oct 2009
17. Intel 82599 10 GbE Controller Datasheet, http://download.intel.com/design/network/datashts/82599_datasheet.pdf, 2010
18. Waldspurger CA, Weihl WE (1995) Stride scheduling: deterministic proportional-share resource management. Tech. rep. MIT/CS/TM-528, MIT Laboratory for Computer Science, June 1995
19. Fedorova MSA, Seltzer M, Small C (2005) Casc: a cache-aware scheduler for multithreaded chip multiprocessors. Sun Labs technical report, March. <http://labs.oracle.com/scalable/pubs/CASC.pdf>
20. Gupta P, Lin S, McKeown N (1998) Routing lookups in hardware at memory access speeds. In: IEEE Infocom

21. Intel Hyper-Threading Technology. <http://www.intel.com/technology/platform-technology/hyper-threading/>
22. Turner J et al (2007) Supercharging PlanetLab—a high performance, multi-application, overlay network platform. In: ACM SIGCOMM, August 2007
23. Wang Y et al (2008) Virtual routers on the move: live router migration as a network-management primitive. In: ACM SIGCOMM, August 2008
24. Anwer MB, Feamster N (2009) Building a fast, virtualized data plane with programmable hardware. In: ACM VISA workshop, Aug 2009
25. Liao Y, Yin D, Gao L (2009) PdP: parallelizing data plane in virtual network substrate. In: ACM VISA workshop, Aug 2009
26. Schaffrath G, Werle C, Papadimitriou P, Feldmann A, Bless R, Greenhalgh A, Kind M, Maennel O, Mathy L (2009) Network virtualization architecture: proposal and initial prototype. In: ACM VISA workshop, Aug 2009
27. Barham P, Dragovic B, Fraser K, Hand S, Harris T, Ho A, Neugebauer R, Pratt I, Warfield A (2003) Xen and the art of virtualization. In: 19th ACM symposium on operating systems principles, October 2003. ACM Press, New York
28. Egi N, Greenhalgh A, Handley M, Hoerd M, Mathy L, Schooley T (2007) Evaluating xen for virtual routers. In: PMECT'07, August 2007