

Emulation of burst-based adaptive link rates in NetFPGA towards green networking

Shahul Hamead HAJA MOINUDEEN*, Mirnalinee THANGANADAR THANGATHAI,
Kavi Priya DHANDAPANI

Department of Computer Science and Engineering, SSN College of Engineering, Chennai, India

Received: 27.06.2019

Accepted/Published Online: 02.01.2020

Final Version: 08.05.2020

Abstract: In recent times, energy consumption in communication media has been increasing drastically. In the literature, energy-saving techniques that enable network devices to enter sleep state or limit the data rate have been proposed to reduce energy costs. In our earlier work, we proposed an energy-saving technique called burst-based adaptive link rate (BBALR), the simulation of which assures increased energy savings. In this paper, we have emulated the hardware implementation of BBALR and compared its performance with the outputs of other prominent energy-saving policies based on dynamic link rate adaption. The energy savings are mapped from the measured sleep time and reference power values. We have used NetFPGA as the testbed, which is a research platform for building real-time network hardware prototypes.

Key words: Network prototype, green networks, hardware implementation, NetFPGA, policy-based forwarding, dynamic link rates, energy-efficient forwarding, emulation

1. Introduction

The Global e-Sustainability Initiative (GeSI) has been researching on cutting down the CO₂ emissions caused by the ICT sector towards building a sustainable society. The key points of the GeSI SMARTer 2030 report, which is related to our work, are summarized in this section. It is estimated that 12.8 Gt of CO₂ emissions can be cut down by defining energy-efficient practices that the ICT sector needs to incorporate. Predictions say that the energy savings that could be obtained from the ICT sector will be higher than those of the other existing sectors. It could be observed that the energy consumption of the operation of ICT devices is 9.7 times the energy consumption of their manufacturing (Figure 1). The ratio of operating costs to the footprint keeps increasing consistently on a scale of more than 1.5. The energy consumption of ICT is distributed across end-user devices, data centers, and networks with proportions of 47.2%, 28.8%, and 24%, respectively. By 2030, every end-device like IoT and mobile devices will have access to the Internet. Hence, the scope for saving energy costs of the networking plane will exist as a research challenge at device and data center levels. Through energy savings, society can witness better benefits by saving 330 trillion liters of water, etc. Other detailed explanations can be found in the GeSI reports [1, 2].

We have gone through the surveys of green communications. In recent surveys [3], energy-efficient approaches have been classified into node-level and network-level techniques. For this work, we have confined our study to node-level techniques. Link rate adaption is a prominent technique applied at node level to save

*Correspondence: shahulhameadh@ssn.edu.in

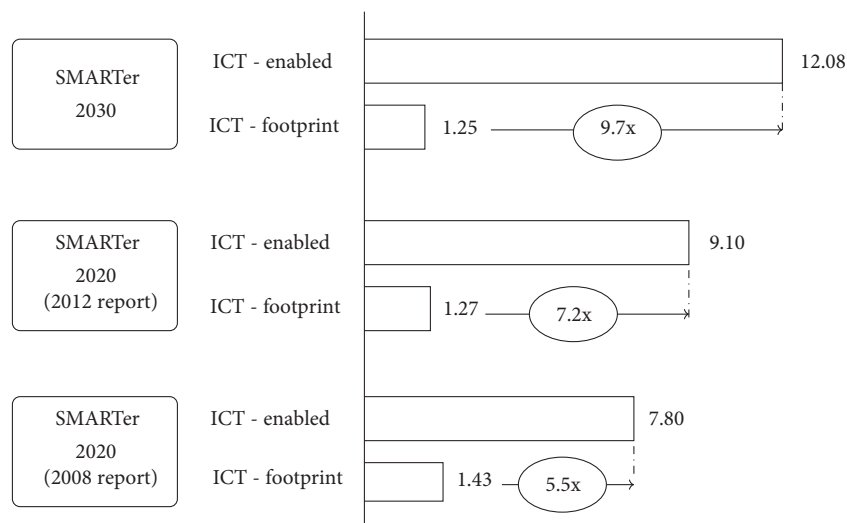


Figure 1. Stipulated energy consumption of ICT sector in 2030.

energy. Cutting-edge networking paradigms like Green-SDN [4] are undergoing integration of link rate adaption with flow-based routing. In the earlier surveys, techniques like reengineering the hardware design and dynamic link rates were forecast as the future of green networks [5, 6]. Link rate adaption was proposed as an energy-saving technique in NICs, switches, and routers in the last decade [5]. It is now being successfully incorporated in emerging networking paradigms.

In our work, we have emulated the real-time hardware implementation of the latest technique named burst-based adaptive link rate (BBALR). We have compared the performance with the results of the other techniques of adaptive link rate (ALR) and burst transmission (BTR).

The system is implemented on a hardware platform known as NetFPGA, which supports rapid prototyping of networking techniques [7, 8]. We have experimented with the BBALR technique to observe the energy savings through the transmission of TCP packets between two hardware units. The benefits yielded through this work would help extend this idea to real routers and switches. The organization of the paper is as follows. Section 2 summarizes the techniques of dynamic link rates. Section 3 specifies the system model adopted for this work. Section 4 introduces the NetFPGA platform and discusses the design and the hardware emulation of the system. Section 5 compares the performance of the system with the existing techniques.

2. Related work

The line card determines the speed of a router and is therefore one of the major components in a router. However, it is also the most power-consuming component since it has packet processors. A router may contain one or more line cards. Dynamic link rates are implemented by imposing sleep on the line cards and activating them only when required. In this section, three techniques that dynamically change the states of line cards are briefly summarized.

2.1. Burst transmission techniques

In this section, we will cover the burst transmission techniques studied in [7, 8]. This is a method that coalesces the packets and sends them as a burst instead of sending them one by one. While the packets are being collected in the buffer, line cards are put into sleep to save energy on the operation. Once the line cards are turned on,

the collected packets are serviced as a burst. The energy is consumed only when packets are processed in active mode. In Figure 2, after an active period, sleep time starts at t_s . At a regular interval, the refreshing of (t_r) happens. During this, the condition to switch to an active state is verified. If the constraints do not hold, the line card remains in a sleep state. Otherwise, wake-up (t_w) happens, which changes the state of the line card to be active. The condition to switch the line cards to the active state is based on three policies, namely burst-full, timeout, and hybrid. In all these policies all the line cards behave synchronously. i.e. they are in active mode together or in sleep mode together. Sleep mode is also called low power idle mode (LPI).




Figure 2. The active and idle state of burst transmission technique.

2.1.1. Burst-full Policy

The size of the burst is predefined with a threshold. When the line cards sleep, the received packets are stored in the buffer. For every arriving packet, the size of the queue is verified with the burst size threshold. Once the queue size meets the burst size, the line cards are woken up from sleep and switch their state to the active mode for servicing the burst. Once the packets are forwarded, the line card goes to sleep mode again to save energy.

2.1.2. Timeout policy

The previous policy would lead to unwanted delays if the traffic rate is slow. In this case, the burst will not be full, which would, in turn, postpone the wake-up event. There is a possibility that TCP would go for retransmission as the packets are not delivered to the destination within the stipulated timeout. To handle this scenario, a timeout value is set, which wakes up the line cards from sleep although the burst is not full.

2.1.3. Hybrid Policy

During high traffic, the burst would become full and the packets will get discarded if the timeout policy is imposed. The hybrid policy is a combination of burst-full and timeout policies. A hybrid policy verifies that at least one of the conditions is satisfied. The burst will get serviced either on burst-full or on timeout. Whichever one happens first will wake up the line cards from sleep.

2.2. Adaptive link rates

Adaptive link rates were discussed in [9–11]. In this policy, the link rates are determined by continuously monitoring the queue size. The number of line cards getting switched to active mode is dependent on the dynamic traffic rate. Assume a router has two line cards. A threshold can be fixed to determine the number

of line cards required. If the queue size is less than the threshold, one line card can serve; otherwise, both the line cards can serve. In this policy, at least one line card should be in active mode all the time. Queue size is continuously observed for the change of link rates. If the traffic rate is inconsistent, the line cards undergo frequent switching between active and sleep states. A sudden spike of the link rate would process the packets faster and delete more packets in the queue at the same speed. This can again force the system to reduce the link rate as the queue size is less. Continuation of these operations with a single threshold in a cyclic manner would affect the stability of the system. To avoid this problem, a dual-threshold policy is used. In this policy, the thresholds are maintained as $0 < \text{low threshold} < \text{high threshold} < \text{maximum queue length}$. A high link rate is preferred when the size of the queue reaches the high threshold. A low link rate is preferred only when the size of the queue is reduced to a low threshold. A spike from low to high deletes the packets in the queue in a faster manner but the stability of the system is preserved as the subsequent change of link rate happens only when the size of the queue is less than a low threshold. Whenever a line card is switched to an active state, an additional amount of energy is consumed, known as the switching cost. With the dual-threshold policy, the switching cost is reduced.

2.3. Burst-based adaptive link rates

Burst transmission techniques save assured energy by imposing forced sleep to all the line cards until the wake-up constraint becomes true, but all the line cards are switched to active mode although a high link rate is not required. Adaptive link rates have the advantage of turning on only the required number of line cards to an active state, but adaptive link rates are not intended for burst-based transmission that guarantees assured energy savings. Burst-based adaptive link rate is a hybrid technique that collectively ensures the benefit of both [12]. In this technique, all the line cards sleep, as happens in burst transmission. At the time of wake-up, a decision is made to determine the number of line cards that are required for the current burst. Instead of turning on all the line cards, only the required numbers are switched on. In this technique, a hybrid policy is preferred for wake-up, which is mentioned in Section 2.1.3. Burst-based adaptive link rates guarantee an assured sleep with appropriate policies at respective states. An old hybrid technique can also be seen in [13].

3. System model of BBALR

In this section, the system model of BBALR, whose hardware design is the crux of this work, is presented. The state diagram is depicted in Figure 3. The logical conditions for state transitions are outlined in Table 1. Initially, the system is in sleep state S_0 , where all the line cards observe sleep. During this state, the received packets are stored in the buffer. The listeners that identify the occurrence of burst-full and timeout states run concurrently when the system sleeps. When the current time t meets timeout t_{max} , the system enters state S_1 . Alternatively, the current size of queue p might be the same as burst size p_{max} , which makes the system enter state S_2 . In rare cases, both conditions may become true at the same time, which would lead the system to enter state S_3 . States S_1 , S_2 , and S_3 represent active states and S_0 represents the sleep state. Edges a_1 , a_2 , and a_3 represent wake-up events and b_1 , b_2 , and b_3 represent switch-off events. At S_0 , all the line cards sleep. The policy adopted in this state is low power idle (LPI). Since S_1 represents timeout, the queue size would not be the same as burst size. Hence, the adaptive link rate (ALR) policy is preferred at state S_1 . At S_2 , the burst-full state happens and the queue size must be equal to the burst size threshold. Hence, all the line cards are expected to be turned for a higher link rate (HLR). Since S_3 is an event for which both constraints hold, HLR is preferred because of the burst-full property. When the burst gets serviced fully, i.e. $p = 0$, the state transition happens to S_0 from the rest of the states, which resets the time as $t = 0$.

Table 1. Logical conditions for state transitions.

Present state	Logical conditions	Action	Policy followed
s_0	$t < t_{max}$ and $p < p_{max}$	x_0	LPI
s_0	$t = t_{max}$ and $p < p_{max}$	a_1	ALR
s_0	$t < t_{max}$ and $p = p_{max}$	a_2	HLR
s_0	$t = t_{max}$ and $p = p_{max}$	a_3	HLR
s_1	$t \leq t_{max}$ and $p < p_{max}$	x_1	ALR
s_1	$p = 0$	b_1	LPI
s_2	$t < t_{max}$ and $p \leq p_{max}$	x_2	HLR
s_2	$p = 0$	b_2	LPI
s_3	$t \leq t_{max}$ and $p \leq p_{max}$	x_3	HLR
s_3	$p = 0$	b_3	LPI




Figure 3. State diagram of the system.

3.1. Energy consumption model

As the system state of a line card toggles between active and idle states, the energy consumption model accumulates the energy measurement from these states. Static cost (SC) is the stable power consumption of a line card irrespective of its sleep state. SC remains the same for both active and sleep states. Link cost (LC) is the dynamic measurement of power in an active period for various traffic rates when the device transmits packets at different link rates. When a line card is in the sleep state, LC remains zero until it is woken up. Switching cost (S_{WC}) is the power required to change the line card from the active state to the sleep state. Projected power consumption (PPC) is expressed as the summation of all these three costs, i.e. $PPC = SC + LC + S_{WC}$. These parameters are measured in watt-hours. The link utilization factor (LUF) is defined as the ratio of the proportion of links consumed in the active state to the total links (active + idle) available for a fixed unit of time. To synchronize with the PPC, measured in kWh, we have fixed the time duration as one hour. The value of LUF is measured in the range of 0 to 1. Link rate ratio (LRR) is a measurement that depicts the number of times the other technique measures LUF to the LUF of BBALR. More information about the energy model of BBALR can be found in [12].

4. System design and hardware implementation

We have identified NetFPGA as a suitable platform to craft real-time implementation of the system. Several NetFPGA implementations have been studied from previous works [14–25]. The NetFPGA board is physically connected to the motherboard through the PCI express slot and it can be programmed with hardware logic. The software application from a higher layer can access the hardware through the NetFPGA device driver. In our implementation, we have used the NetFPGA 1G CML board, which has a Xilinx Kintex-7 FPGA, 32-bit microcontroller, 4.5 MB SRAM, 512 MB DDR3 DRAM, and 4 Ethernet ports, namely NF0, NF1, NF2, and NF3. Figure 4 illustrates the high-level architecture of NetFPGA inspired by online content.¹ We have programmed the board with Reference NIC hardware logic through a programmable USB port. In the reference design of NIC, the memory cells are reserved for input and output queues. More details about reference NIC architecture can be found online.²

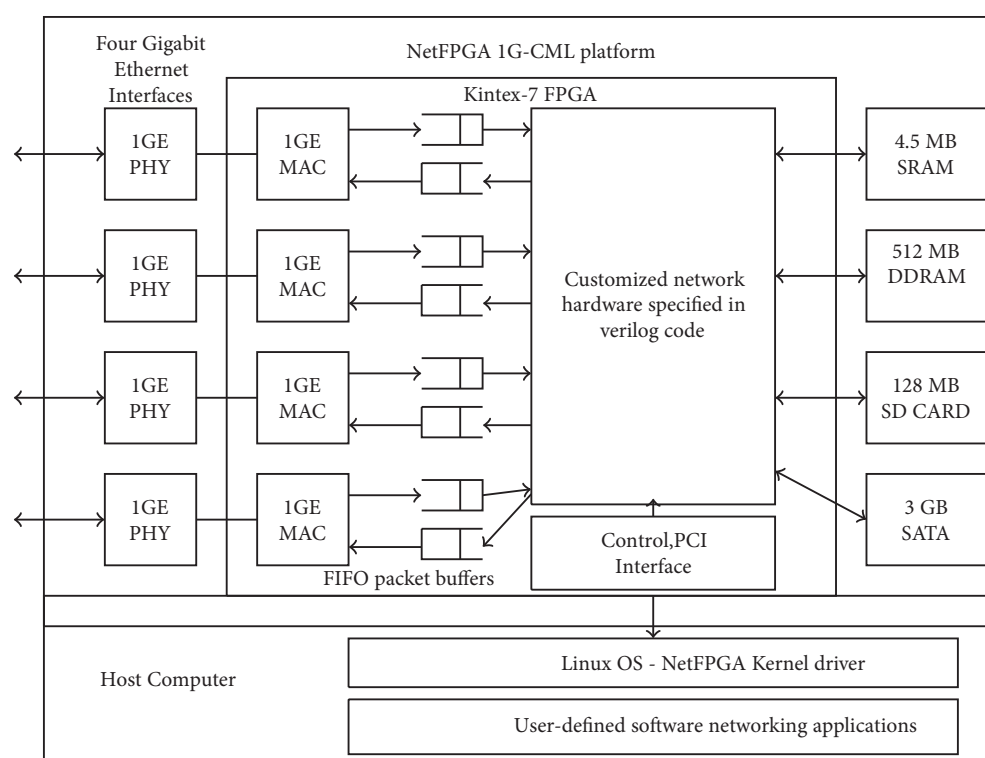


Figure 4. NetFPGA 1G CML architecture.

The implementation design (see Figure 5) encompasses modules at node and board levels. The packet generator module is implemented at the application level, which creates packets in PCAP format with the specified arrival rate λ , i.e. the packets are generated continuously with a mean arrival rate of $1/\lambda$. The generated packets are stored in the queue. The policy base contains energy-saving ideas that are defined as wake-up policies. Generally, policies are composed of queue and clock parameters. The grammar for a policy S is defined as follows:

¹University of Cambridge (2017). NetFPGA Summer Course. Website https://www.cl.cam.ac.uk/research/srg/netos/projects/netfpga/workshop/summer-school-2017/slides/2017_summer_camp_day1_part1.pdf [Accessed 20 November 2018].

²NetFPGA (2014). NetFPGA 1G CML Reference Manual. Website <https://github.com/NetFPGA/NetFPGA-public/wiki/NetFPGA-1G-CML-Reference-NIC> [accessed 21 December 2018].

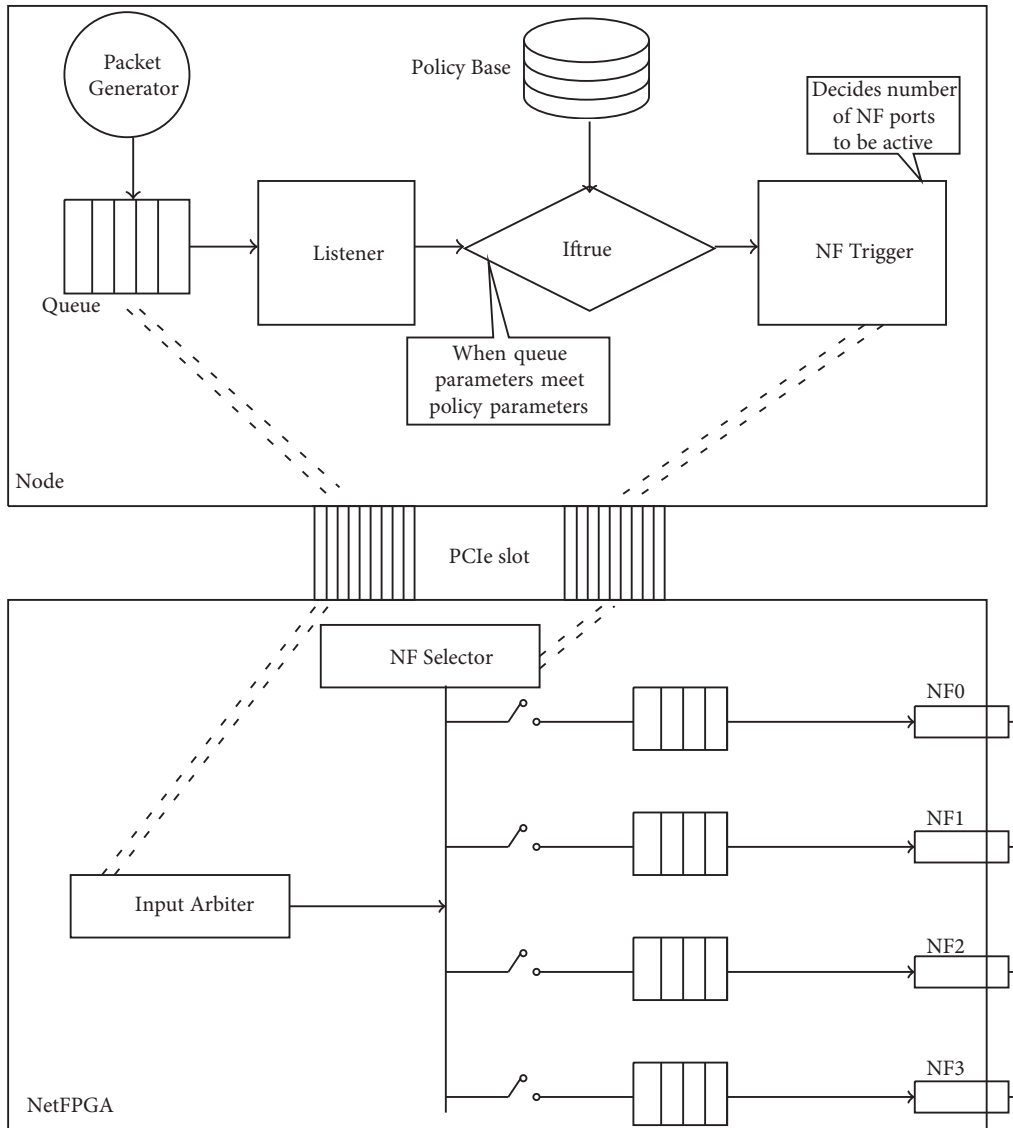


Figure 5. System design.

The set of parameters is defined as $P = \{p_1, p_2, p_3 \dots\}$. These parameters represent the system attributes. They can be dynamic in nature. The set of relational operators is defined as $R = \{<, <=, ==, >, >=, !=\}$. The set of logical operators is defined as $L = \{!, \wedge, \vee, \oplus\}$. The set of thresholds is defined as $\Gamma = \{\tau_1, \tau_2, \tau_3 \dots\}$. The system parameters are continuously compared with threshold values to verify that the policy constraints are true. A constraint is composed as $c \rightarrow p \diamond \tau$, where $p \in P, \diamond \in R, \tau \in \Gamma$. A set of constraints is defined as $C = \{c_1, c_2, c_3 \dots\}$. A policy could be a single constraint or it could be composed of multiple, which is expressed as $S \rightarrow c | s \Delta c$, where $s \in S, \Delta \in L, c \in C$.

The wake-up policy for the burst-based adaptive link rate is as follows: Let t be the current time of the system. Let ρ be the current queue size. Let β be the burst-full threshold. Let Υ be the timeout threshold. Let the burst-full constraint be $C_b = (\rho == \beta)$. Let the timeout constraint be $C_t = (t == \Upsilon)$. Let the policy of the system be $S_{sys} = C_b \vee C_t$, i.e. wake-up happens when at least one of the constraints is true.

The listener module verifies whether the dynamic parameters of the system satisfy the constraints of policy. The NF Trigger module determines the number of ports to be selected for transmitting the packets. The link rate gets dynamically decided based on the policy that is imposed. In order to switch between low rates and high rate in the burst-based adaptive link rates policy, the number of packets in the burst is expected to be greater than the threshold and it can be less than the threshold for high to low. Once a wake-up happens, the packets are copied from the node memory to the NetFPGA memory. A portion of NetFPGA memory is reserved for the input arbiter in the reference NIC design. A module called Packet distributor transmits the packets from the Input arbiter through the selected NF ports in a cyclic manner.

The NetFPGA ports are directly connected to each other. Data transmission occurs through both the ports if a high link rate is preferred, or else only one port is used for transmission. The policy dynamically mandates the number of ports that should be used. The energy savings obtained through the BBALR policy between the two hardware units can be studied and analyzed for speculating energy savings in any larger topology.

5. Results and discussion

In this work, we have done the real-time implementations for burst-based adaptive link rate, burst transmission, and adaptive link rate policies. The comparison of the three is discussed in the following section. We have measured the performance of the implementation via various parameters. Though the objective is to measure energy savings, we have measured QoS parameters as well to demonstrate that the functional objective is not compromised. A couple of parameters measure energy whereas some of them measure QoS. A few parameters even act as a balancing trade-off between energy and QoS. The performance parameters and the context of the usage are given in Table 2.

The packet generator module is common for all the policies. It follows the Poisson process with a mean interarrival time ($1/\lambda$) where λ is the mean arrival rate. In our implementation, we have specified $1/\lambda$ in microseconds. For high arrival rates, we have $1/\lambda$ as low as 10 microseconds. For low traffic, $1/\lambda$ gets delayed with a maximum value of 10000 microseconds. We have measured the performance of the system for high traffic, medium traffic, and low traffic. The mean arrival time of high traffic ranges from 10 to 100 microseconds with an increment of 10 microseconds at each measurement. Medium traffic's interarrival time ranges from 100 to 1000 microseconds with an increment of 100 microseconds. The low traffic ranges between 1000 and 10000 microseconds with an increment of 1000 microseconds at each measurement. In all these graphs (Figures 6, 7, 8, 9, 10, 11, 12, 13, and 14), the x-axes represent the mean interarrival time in microseconds. The burst threshold for the hybrid technique has been set as 150 milliseconds, whereas, for others, it is limited to 30 milliseconds. Similarly, the timeout has been set as 50 milliseconds for the hybrid one and 10 milliseconds for the rest of the techniques.

The performance of the hybrid policy is compared with other policies. We have used acronyms in the graphs of BTR for burst transmission rates, ALR for adaptive link rates, and BBALR for burst-based adaptive link rate (hybrid) policies. Performance measurements in the graphs are depicted for 60 seconds. Parameters like throughput are measured directly from the input arbiter part of NetFPGA. Every parameter is divided into three columns of high traffic rate (a), medium traffic rate (b), and low traffic rate (c). Because of heavy traffic, the throughput is slightly unstable as the service gets delayed, whereas for medium and low traffic, the curve is stabilized to the exponential distribution. The throughput is dependent on the arrival rate, which follows the Poisson process. Hence, the interarrival time follows an exponential distribution, which could be observed from the graphs in Figure 6.

Table 2. Performance parameters and context.

Parameter	Rationale	Type
Throughput	Data transmission in KBps obtained in spite of discrete idle states.	QoS
Average number of packets in queue	Mean queue length measured out of alternating active and sleep states. It would not exceed the maximum capacity of the burst.	QoS
Waiting time	Mean waiting time of packets in milliseconds to get serviced as burst (longer waiting time affects the QoS).	QoS
Sleeping time	Duration of idle time in milliseconds, which is directly mapped to energy savings as $LC = 0$.	Energy
Switching time	Time is taken to toggle from an active state to sleep state and vice versa. Consumes the power of active state during the course of switching in milliseconds.	Energy
Wake-up count	Reason for wake-up. The number of times timeout occurred vs. the number of times burst-full occurred. Frequent wake-up consumes more switching costs.	Energy and QoS
Projected power consumption (PPC)	The major parameter discussed in Section 3.1. The reference watts are from Guo et al. [26] and are applied in active and sleep states to get the PPC. It is expressed in watt hours.	Energy
Stipulated energy savings	The percentage of the duration of sleep in the total implementation period.	Energy
Mean service length	Mean duration of the active period in milliseconds to deliver the packets as the burst.	QoS
Link utilization factor (LUF)	The ratio of usage of the ports to the total availability. It is discussed in Section 3.1.	Energy and QoS
Link rate ratio	The ratio of LUF of existing technique vs. LUF of BBALR. It is discussed in Section 3.1.	Energy and QoS

In Figure 7, the mean number of packets in a queue follows exponential distribution if the waiting time is fixed. If the service happened due to a timeout event, the waiting times for different traffic rates remain the same. This can be seen with interarrival time greater than 400 microseconds. For high traffic rate, waiting time is reduced due to burst-full wake-ups, which in turn lead to a constant mean number of packets in the queue. This can be observed for mean arrival time less than 400 microseconds (Figure 8). Based on the policy and dynamic parameters of the system, the NF ports are selected for transmission. The selected NF ports are woken from sleep mode to transmit the packets waiting in queue. For low traffic rates, only one port is used. The dynamic usage of ports has been verified with the Wireshark tool.

The sleeping time is computed from the idle time of the NF ports (Figure 9). As the mean arrival time between the packets gets delayed, the sleep time increases. Since energy savings are directly proportional to the amount of sleep, this can be understood as energy savings increasing with respect to the quantity of sleep. Due to the enforcement of compulsory sleeping time, the mean queue length and waiting time increase up to five times without compromising the throughput. It could be noted that throughput obtained for BBALR is nearly the same as that for the other techniques (Figure 6).

Switching time (Figure 10) is the time to prepare the ports for transmission; therefore, it could be perceived as the time between idle mode and active mode. For better performance, the total switching time is




Figure 6. Comparison of throughput in KBps for different traffic rates.

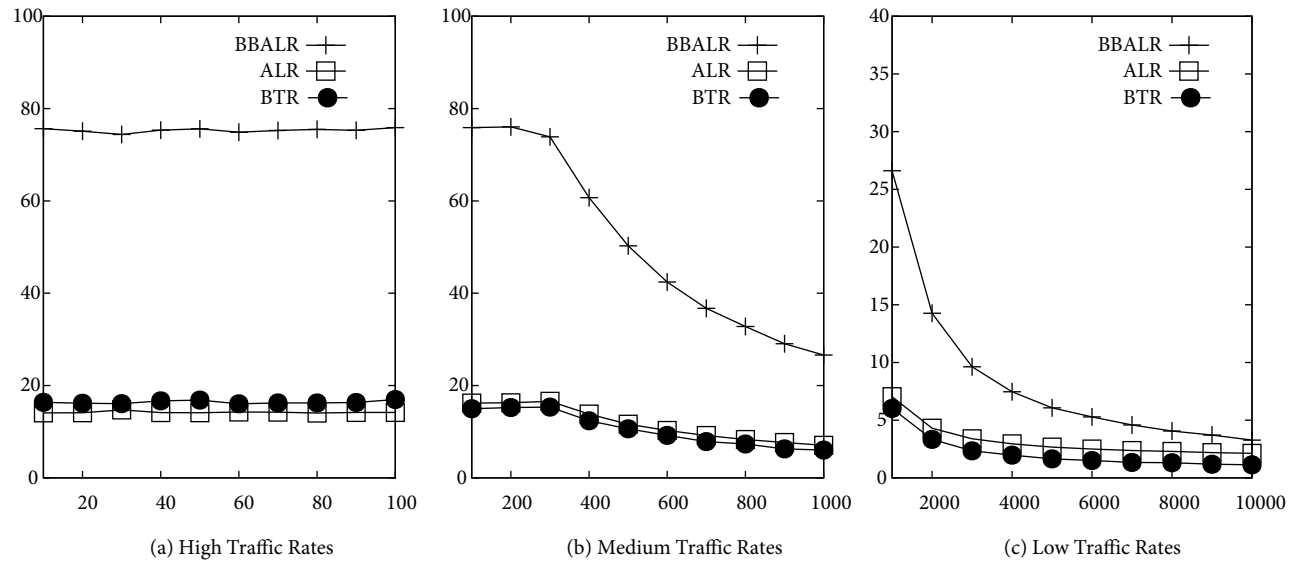


Figure 7. Comparison of the average number of packets for different traffic rates.

expected to be less. Through our experiment, we can see that BBALR performs well with reduced switching count compared to ALR and BTR because of the higher payload per active period. ALR and BTR are measured with a lower mean number of packets in the queue and less waiting time, whereas BBALR undergoes delayed waiting time. On the other hand, because of the higher thresholds, BBALR saves more energy through increased sleep time than the other two techniques in spite of increased waiting time. Throughput is almost the same for all, as the packet generator is kept independent of the policies. In Figure 11 we have illustrated the total number of services after wake-ups. The BBALR policy is applied in the first row, ALR policy is applied in the second row, and BTR policy is applied in the third row. The policies are applied for high, medium, and low traffic rates from left to right. Because of the higher threshold values, which are suitable for energy savings, the




Figure 8. Comparison of waiting time in milliseconds for different traffic rates.




Figure 9. Comparison of sleeping time in milliseconds for different traffic rates.

number of services is less for BBALR, whereas it is high for ALR and BTR. It could be observed that burst-full wake-ups dominate for high traffic rates and time-out wake-ups control the low traffic rates. For medium traffic rates, the initial control starts with burst-full, and in the middle (300–400), the wake-up control is transferred to the time-out constraint. This behavior has been consistently observed for all the policies through which there is a variation in the number of wake-ups.

The performance parameters for different policies are observed in Figure 12, in which the traffic rate varies in the logarithmic scale. As mentioned earlier, throughput is similar for all the policies. For high rates, it is consistent, and it drops slowly with respect to the increase in delay in mean arrival time (Figure 12a). A curve with similar throughput can be seen for the mean number of packets in the queue for BBALR because of energy-friendly thresholds (see Figure 12b). Other policies follow a reduced mean number of packets in the

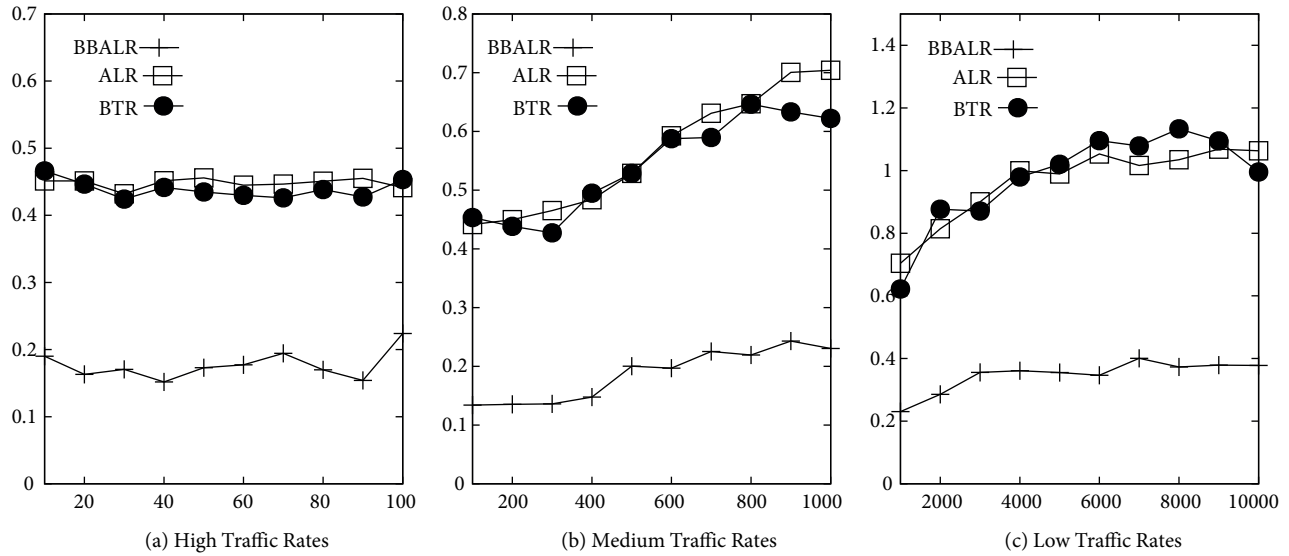


Figure 10. Comparison of switching time in milliseconds for different traffic rates.

queue. The average waiting time (Figure 12c) is inversely proportional to the number of packets in the queue because of the aforementioned reason in Figure 7 and Figure 8. Switching time for the policies increases slowly and steadily, which follows a similar pattern. BBALR improves performance with reduced switching costs (Figure 12d). The sleeping time, an energy-saving factor, starts with similar values for high arrival rates and gains more scope for energy conservation with medium and low arrival rates (Figure 12e).

NetFPGA implementations of the green reference router have been taken from [26, 28]. We have traced the extent of the duration for which the NF ports are in active mode and sleep mode. Power consumption is a function that takes the length of the active period as the parameter. Energy savings is a function that takes the length of sleeping time as a parameter. From an earlier NetFPGA experiment, watts consumed for 2 ports, 1 port, and no ports at low and high frequencies are listed in Table 3. With the substitution of these aforementioned values, the power consumption for a time duration in watt-hours for different policies could be projected. We have scaled the power-related parameters for 24 hours duration. Similarly, stipulations about energy savings are possible from the mentioned values substituted for the duration of sleep time. Projected power consumption (PPC) is expressed as the sum of static, dynamic, and switching power costs of the components. From Figure 13 it can be observed that BBALR would consume less power than other policies. There is a difference of 40 watt-hours between the minimum power consumption of BBALR and other policies. Similarly, BBALR can save energy at rates double the other policies for low traffic rates.

The link utilization factor (LUF) is depicted in Figure 14b. Usually, the default link rate consumes more power as it does not follow energy policy. We have measured it through the ratio of the length of the active period to the cumulative length of the active and sleep period. As LUF depends on incoming traffic, the curve follows an exponential distribution. LUF is high for faster arrival rates and it drops for lower arrival rates. The amount of times power consumed by an existing policy for a particular traffic rate compared to the power consumed by BBALR for the same traffic rate can be understood as the link rate ratio. LRR (existing policy, BBALR) is the ratio of the LUF of an existing policy to the LUF of BBALR. LRR values can be found in Table 4. Through LRR values, it could be observed that LUFs of all the policies are the same for low traffic rates. There is a slight change in medium traffic. When comparing BBALR, the existing policies double its

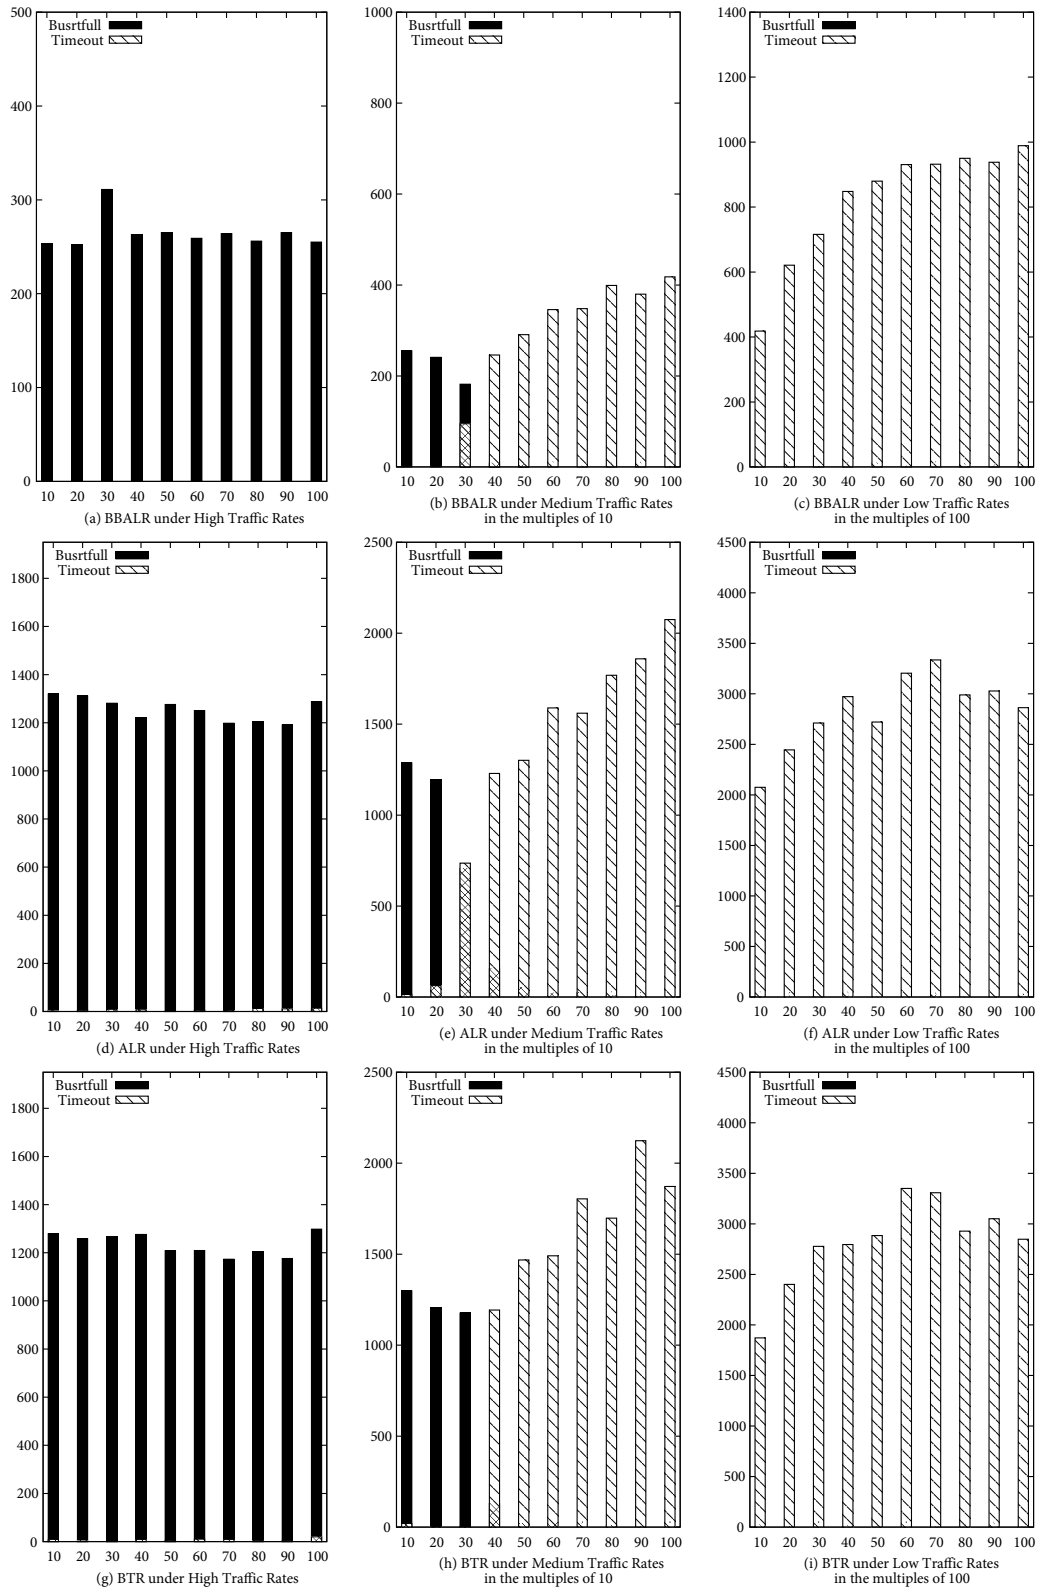


Figure 11. Comparison of wake-up count of policies.

Table 3. Watts measured with respect to the usage of ports [27].

No. of active ports	Power consumption (W)	
	125 MHz	62.5 MHz
0	6.994	5.333
1	7.841	6.325
2	8,767	7.302
3	9.760	8.373
4	10.725	9.441

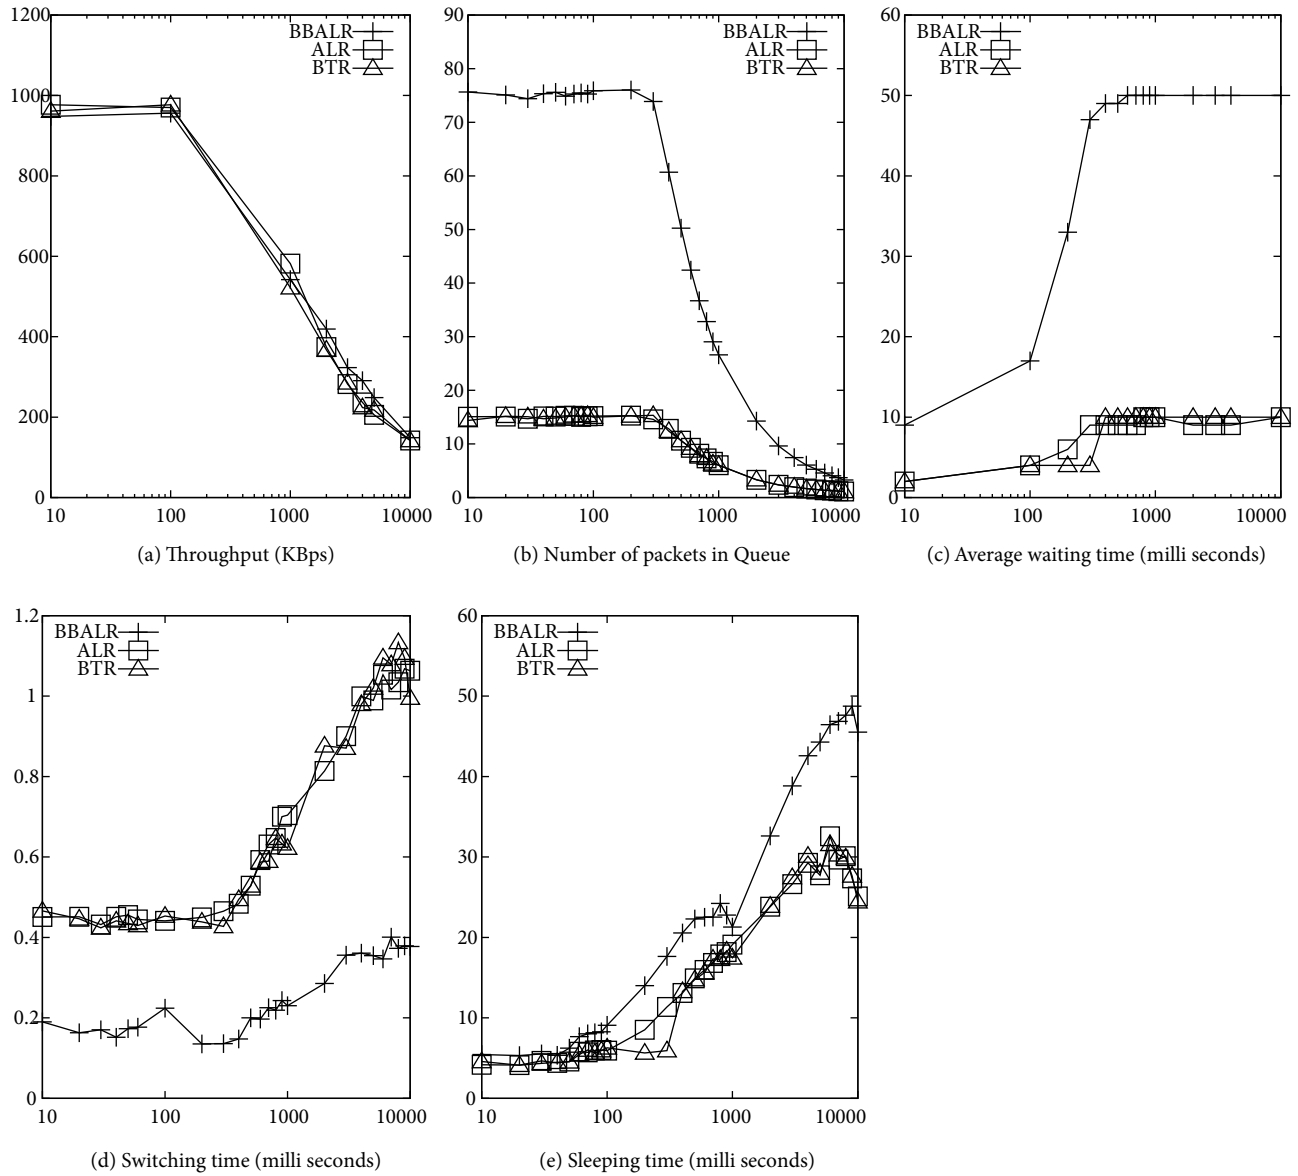


Figure 12. Performance analysis in logarithmic scale.




Figure 13. Projected power consumption in watt-hours per day (a) and stipulated energy savings in percentage (b).

utilization or energy cost for higher traffic rates. The results of the simulation [12] are based on the increase of difference in mean interarrival time whose units are seconds, whereas the results of this work are based on units in microseconds. Because of this change, PPC, LUF, and LRR vary in magnitude with similar behaviors. Another important aspect of observation is the measurement of mean service length (see Figure 14a). It could be seen that BBALR is measured with longer service length than the other policies. Because of burst behavior, the service length is comparatively longer. Once the service is done, the device enters into sleep mode until the constraints of the policy become true, which in turn would benefit uninterrupted longer sleep. The ALR policy undergoes a short nap and experiences frequent switching because of its packet-based transmission policy. Though BTR supports smaller bursts, all the ports are unnecessarily switched to active mode even for low traffic rates, which in turn wastes energy instead of saving it during the active period. Hence, the sleeping quality is maintained in BBALR by skipping unwanted switching costs.

Table 4. Link rate ratio.

Policies/ Traffic	Low	Medium	High
ALR, BBALR	1.033	1.139	1.982
BTR, BBALR	1.033	1.163	1.981

Finally, the difference in output between simulation and implementation is discussed. The results of implementation are based on the interarrival times, whose units are measured in microseconds, whereas the

simulation results are obtained for interarrival times set in seconds. The measures of energy parameters in the simulation remain constant for fast interarrival times of smaller than 2 seconds due to the constraints of the simulation environment. On the other hand, the QoS and energy measurements change proportionally for the implementation even for faster interarrival times of less than 100 microseconds. The variations in output parameters like PPC for BBALR can be seen in the graphs (Figures 12, 13, and 14). Due to the smaller reference power values, no significant variation could be observed in the PPC of the simulation for high traffic rates [12]. The PPC calculations in the simulation use reference voltage values for two cases, staying at the same link rate and switching between different link rates (Table 5). The PPC calculations for emulation use the reference power values in Table 3 for the number of active ports instead of the changes in the link rates. As the reference power values used in implementation is 10 times larger than the values used in simulation and the experimental setup of the implementation uses FPGA, we have preferred using these values in the implementation of our technique. Due to the large reference valuations used in BBALR, variations in PPC can be witnessed.




Figure 14. (a) Comparison of mean service length in milliseconds, (b) comparison of link utilization factor.

6. Conclusion

The hardware implementation of energy-saving policies for data transmission in wired networks has been performed in the NetFPGA platform. A suitable reference NIC architecture has been programmed to the board. Policies are defined at a high level from the software layer. The performance parameters are measured with different traffic rates for all the policies including burst-based adaptive link rate, which has energy-friendly thresholds. From the experiments, it is observed that the burst-based adaptive link rate harvests cumulative

Table 5. Reference power values used in simulation [12].

Constant link rate and link rate switching	Power (watts)
10 Mbps	0.1
100 Mbps	0.4
1000 Mbps	3.6
10/100	0.8
10/1000	1.05
100/1000	1.6

benefit with a combination of extended sleep and reduced switching costs compared to other policies. The QoS measures are also well maintained in BBALR. The core values that impact the energy measurement through projected power consumption and link utilization factor are measured directly from the hardware. The results of emulation are more insightful than the results of simulation. As future work, the hardware emulation can be extended to a full-length prototype along with its formal verification.

Acknowledgment

The authors would like to thank SSN Management for supporting us financially under the internal funding scheme.

References

- [1] GeSI. Accenture Strategy. ICT Solutions for 21st Century Challenges. SMARTer 2030 - Research Report, 2015.
- [2] GeSI. The Boston Consulting Group. The Role of ICT in Driving Sustainable Future. SMARTer 2020 - Research Report, 2012.
- [3] Maaloul R, Chaari L, Cousin B. Energy energy saving in carrier-grade networks: a survey. *Computer Standards & Interfaces* 2018; 55 (1): 8â 26.
- [4] Tuysuz M, Ankarali Z, Gozuepek D. A survey on energy efficiency in software defined networks. *Computer Networks* 2017; 113 (1): 188â 204.
- [5] Nedeveschi S, Popa L, Iannaccone G, Ratnasamy S, Wetherall D. Reducing network energy consumption via sleeping and rate-adaptation. In: *NSDI'08 Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation*; San Francisco, CA, USA; 2008. pp. 323â 336.
- [6] Zilberman N, Audzevich Y, Covington G, Moore A. NetFPGA SUME: Toward 100 Gbps as research commodity. *IEEE Micro Magazines* 2014; 34: 32â 41.
- [7] Meng J, Ren F, Jiang W, Lin C. Modeling and understanding burst transmission algorithms for energy efficient ethernet. In: *IEEE/ACM 21st International Symposium on Quality of Service(IWQOS)*; Montreal, Canada; 2013. pp. 173â 182.
- [8] Herreria-Alonso S, Rodriguez M, Veiga M, Lopez-Gracia C. A power saving model for burst transmission in energy-efficient ethernet. *IEEE Communications Letters* 2011; 15 (5): 584â 586.
- [9] Zilberman N, Audzevich Y, Kalogeridou G, Bojan N, Zhan J. NetFPPA-rapid prototyping of high bandwidth devices in open source. In: *25th International Conference on Field Programmable Logic and Applications(FPL)*; London, UK; 2015. p. 1.
- [10] Gunaratne C, Christensen K, Suen S, Nordman B. Ethernet Adaptive Link Rate (ALR): Analysis of a buffer threshold policy. *IEEE Globecom* 2008; 57 (1): 1â 6.

- [11] Gunaratne C, Christensen K, Nordman B, Suen S. Reducing the energy consumption of ethernet with Adaptive Link Rate (ALR). *IEEE Transactions on Computers*. 2008; 57 (4): 448â 461.
- [12] Sowmiya N, Hamead H, Mirnalinee TT. Adaptive link rates for burst based transmission towards the prosperity of green networks. *IEEE International Conference on Advanced Networks and Telecommunications System; Bangalore, India; 2017*. p. 1â 6.
- [13] Jin S, Fan R, Yue W. A hybrid energy saving strategy with LPI and ALR for energy efficient ethernet. In: *2nd International Conference on Computer Science and Network Technology; Changchun, China; 2013*. pp. 311â 315.
- [14] Michalski M. The system for delay measurement in ethernet networks on NetFPGA cards. In: *IEEE 15th International Conference on High Performance Switching and Routing; Vancouver, Canada; 2014*. pp. 1â 4.
- [15] MartÃnez-Aguilar R, FernÃndez G. Implementation of stateless routing mechanisms for multicast traffic on NetFPGA card. In: *IEEE Colombian Conference on Communication and Computing; Popayan, Colombia; 2015*. pp. 1â 5.
- [16] Michalski M, Sielach T. The analysis of time reaction in OpenFlow switches in NetFPGA cards and ROFL. In: *2015 International Symposium on Networks, Computers and Communications; Hammamet, Tunisia; 2015*. pp. 1â 5.
- [17] Cao J, Zheng X, Sun L, Jin J. The development status and trend of NetFPGA. In: *2015 International Conference on Network and Information Systems for Computers; Wuhan, China; 2015*. pp. 101â 105.
- [18] Seovic J, Takov S, Velickovic J, Smiljanic A. Network topology exchange between NetFPGA-10G and quagga routing platforms. In: *23rd Telecommunications Forum Telfor; Belgrade, Serbia; 2015*. pp. 63â 66.
- [19] Jin J, Sun L, Guo F, Wang X. Low power design for on-chip networking processing system. In: *28th IEEE International System-on-Chip Conference; Beijing, China; 2015*. pp. 167â 172.
- [20] Chouhan T. Implementation of present cryptographical algorithm for the encryption of messages in NETFPGA 1G. In: *International Conference on Computational Intelligence and Communication Networks; Jabalpur, India. 2015*. pp. 1115â 1119.
- [21] Su T, You L, Wang Q, Hou C. The high speed switching experiment based on NetFPGA SUME. In: *11th International Conference on Computer Science & Education; Nagoya, Japan; 2016*. pp. 652â 657.
- [22] Dorosh S, Debita G, Schauer P. Network hardware analyzer based on NetFPGA 1G. In: *IEEE 26th International Conference on Enabling Technologies: Infrastructure for Collaborative Enterprises; Poznan, Poland; 2017*. pp. 150â 154.
- [23] Cerovic D, Piccolo V, Amamou A, Haddadou K, Pujolle G. Fast packet processing: a survey. *IEEE Communications Surveys & Tutorials* 2018; 20 (4): 3645â 3676.
- [24] Duan T, Lan J, Hu Y, Liu S. A reconfigurable hardware architecture for packet processing. *Chinese Journal of Electronics* 2018; 27 (2): 428â 432.
- [25] Oeldemann A, Wild T, Herkersdorf A. FlueNT10G: A programmable FPGA-based network tester for multi-10-gigabit ethernet. In: *28th International Conference on Field Programmable Logic and Applications; Dublin, Ireland; 2018*. pp. 179â 185.
- [26] Guo F, Ormand O, Collier M, Wang X. Power measurement of NetFPGA based router. In: *IEEE Online Conference on Green Communications; 2012*. pp. 3814â 3827.
- [27] Guo F, Wang X, Song M, Wei Y, Ormond O. Greening the NetFPGA reference router. *Energies* 2016; 9 (7): 500â 521.
- [28] Song T, Jiang Z, Wei Y, Ma X, Ormand O et al. Traffic aware energy efficient router: Architecture, prototype, and algorithms. *IEEE Journal on Selected Areas in Communications* 2016; 34 (12): 3814â 3827.