

Towards a Real-Time IoT: Approaches for Incoming Packet Processing in Cyber-Physical Systems*

Ilja Behnke^{a,*}, Christoph Blumschein^a, Robert Danicki^a, Philipp Wiesner^a, Lauritz Thamsen^b, Odej Kao^a

^a*Technische Universität Berlin, Distributed and Operating Systems, Berlin, Germany*

^b*University of Glasgow, School of Computing Science, Glasgow, United Kingdom*

Abstract

Embedded real-time devices for monitoring, controlling, and collaboration purposes in cyber-physical systems are now commonly equipped with IP networking capabilities. However, the reception and processing of IP packets generates workloads in unpredictable frequencies as networks are outside of a developer's control and difficult to anticipate, especially when networks are connected to the internet. As of now, embedded network controllers and IP stacks are not designed for real-time capabilities, even when used in real-time environments and operating systems.

Our work focuses on real-time aware packet reception from open network connections, without a real-time networking infrastructure. This article presents two experimentally evaluated modifications to the IP processing subsystem and embedded network interface controllers of constrained IoT devices. The first, our software approach, introduces early packet classification and priority-aware processing in the network driver. In our experiments this allowed the network subsystem to remain active at a seven-fold increase in network traffic load before disabling the receive interrupts as a last resort. The second, our hardware approach, makes changes to the network interface controller, applying interrupt moderation based on real-time priorities to minimize the number of network-generated interrupts. Furthermore, this article provides an outlook on how the software and hardware approaches can be combined in a co-designed packet receive architecture.

Keywords: real-time, embedded systems, network stacks, iot

1. Introduction

With the introduction of the Internet of Things (IoT) in manufacturing industries and automotive systems IP networks have entered the domain of real-time embedded systems. Technologies such as 5th generation (5G) mobile networks have been driving the integration of such systems in business networks, command and control infrastructures, and machine-to-machine (M2M) communication [11, 36, 17]. Furthermore, commands that are being sent over IP networks are subject to real-time requirements in applications like remote control of industrial machines, remote surgery, and autonomous routing of logistic robots [32, 40, 33]. To this end, recent research has investigated industrial network architectures in the Industrial IoT (IIoT) and Time Sensitive Networking (TSN) [38, 10, 67]. However, the impact of IP networking on the real-time behavior of embedded systems has largely been ignored.

Embedded real-time devices need to be designed holistically and take hardware cost, energy efficiency and robustness into account, while computing power is typically

constrained [65]. At the same time, when devices are used as controlling units in cyber-physical systems, they demand predictable and limited execution times [1, 47]. Input devices like sensors or communication interfaces have always undermined this predictability, as the used Interrupt Requests (IRQs) preempt running real-time process regardless of their priorities [37]. Yet, short Interrupt Service Routines (ISRs) and controllable IRQ frequencies enabled real-time system developers to incorporate these into worst case execution time analyses, making external interrupts manageable.

In this regard, Network Interface Controllers (NICs) act similarly to other I/O devices. When a network packet is received, the content is written to memory and an IRQ is triggered. The preempting ISR then pre-handles the packet before notifying the operating system. Then, a usually highly prioritized network task resumes packet processing in order of entry and forwards it to a waiting application socket. All these tasks are performed independently of the presumed priority of the packet and create a timing overhead proportional to the rate of incoming packets [5, 48]. In common IP networks this rate is not predictable by a real-time applications developer. Yet, IoT use-cases may contain real-time requirements and traditional IP networking without real-time specific protocols or hardware at hand [64].

*This is the preprint of an article to be published in the Journal of Systems Architecture (<https://doi.org/10.1016/j.sysarc.2023.102891>). CC BY-NC-ND

*Corresponding author

Email address: i.behnke@tu-berlin.de (Ilja Behnke)

Evaluations on common off-the-shelf IoT microcontrollers showed that network controllers and IP stack implementations are not prepared for to real-time environments [5]. The tested microcontrollers were overloaded already with 1000 packets per second with no apparent mitigation reactions other than network task and system shut down. In light of this problem, developers must resort to disabling interrupts during critical executions or require separate resources for networking and processing in resource-constrained microcontroller environments. However, these solutions are not practical in scenarios where IP networks are used to control cyber-physical systems conveying soft real-time messages that are relevant for baseline functionality. For example, complex industrial actuators with internal feedback control implemented on constrained embedded devices need to meet local real-time requirements independently of the network load while prioritizing the reception of their real-time messages among the received traffic.

While there is extensive research on enabling technologies in the areas of communication and integration in IoT, device architectures have received less attention [3, 1, 53]. Recent approaches focus on blocking or processing traffic before it arrives at the device [43, 50, 22]. Previous works failed to include network-specific factors and focused on interrupt management only [21, 46]. Past work has proposed to do packet classification in the networking driver as early as possible [13] and identify a priority for each UDP-packet by receiving process and defer the subsequent packet processing [35, 34] based on the assigned priority. Furthermore, solutions exist using specialized networking technologies such as TSN [68, 55, 58] or Software Defined Networking (SDN) [39, 23, 16]. While these could be used to mitigate the problem of high packet rates, they require specialized cooperative networking hardware and a rigorous, inflexible real-time network architecture [63, 55]. To the best of our knowledge, there is no published research on extending NICs and network stack implementations to facilitate connected real-time embedded systems in common IP networks.

Contributions

This is an extended discussion of the work first presented at the International Symposium on Real-Time Distributed Computing 2022 [8] and also includes material previously displayed at the poster session of the Symposium on Applied Computing 2022 [6].

With this article, we focus on the real-time implications of IP packet reception in real-time systems by proposing two different network subsystem modifications. Both approaches differentiate incoming packets by their priority. The priority of packets is determined by the receiving task priority and real-time requirements in the Real-Time Operating System (RTOS). The approaches regard constrained IoT devices with simple NICs and network stacks, not able to participate in real-time specific networking via Time-Sensitive or Software Defined Networking solutions.

The devices communicate using unmodified IP networking.

The first implementation, a **software modification**, reduces the impact of best-effort packet processing on the real-time behavior in IoT devices. This is achieved using an early priority-dependent demultiplexing scheme for incoming packets and subsequent aperiodic per-flow scheduling in the network driver. The modification protects real-time embedded systems against network-induced system overloads while optimizing for low-latency processing of high-priority IP flows. This is achieved by strictly controlling the best-effort performance of low-priority flows.

The second approach addresses the problem with a **hardware modification** of network interface controllers. By mapping IP flows to the priorities of receiver processes, the priority space of an RTOS can be extended to include the moderation of network-generated interrupts. By applying different interrupt moderation parameters of priority-based receive queues and dropping unregistered flows, the amount of interrupts is reduced while high-priority packets are still handled immediately.

Both presented approaches include experimental evaluations under varying network loads on prototypical implementations including real hardware and simulations.

Lastly, we present an outlook on how our hardware and software approaches can be combined effectively to a **unified design**.

Outline

Section 2 provides the relevant background information. Section 3 presents the related work. Section 4 presents the preliminary considerations made for the two approaches. Section 5 presents the software-based modification of the IP subsystem and 6 discusses its evaluation. Section 7 and 8 do the same for the hardware-based adaptation of a NIC. Section 9 gives an outlook towards a combined packet receive architecture. Section 10 concludes this article.

2. Background

This section gives a short introduction to the reception of network packets in RTOSs, interrupt handling, and interrupt moderation.

2.1. Direct Memory Access

In order to relieve the CPU from actively pulling data from or pushing data to the NIC memory, Direct Memory Access (DMA) has been established as a common feature in peripheral hardware. With it, the NIC can asynchronously place received packet data into main memory at previously assigned locations and afterwards only needs to notify the CPU about the arrival of new packets via interrupt.

The memory assignment usually happens in so-called Buffer Descriptor (BD) rings, as outlined in Figure 1. In




Figure 1: **BD Ring**: The data structure allows simultaneous access from the NIC adding entries and the CPU checking off processed descriptors.

memory such a ring comprises an array of Buffer Descriptors, interpreted as a ring buffer. Each BD contains a memory pointer to the respective buffer and some meta-data for cooperation. The latter typically includes an ownership bit, indicating whether the CPU or NIC is obligated to go on with processing, and a length field indicating how far into the buffer data shall be sent or has been received respectively. This way, both actors can track their current working position(s) individually.

2.2. Receive Path in IP Networking

At a high level, the Receive (RX) path is organized into subsequently executed stages as seen in Figure 2. Upon packet reception, the NIC transfers the packet content to a previously prepared memory location via DMA, marks the corresponding BD entry and triggers an interrupt. The network driver, handling the interrupt, acknowledges the DMA-operation and exchanges the received frame buffer with a newly allocated one. From here, protocol processing can commence disregarding the already finished MAC-layer operations.

The different network stack implementations vary in their set of features. The LWIP network stack is widely spread among embedded applications and sets its focus on memory efficiency [14]. It covers the majority of commonly used and necessary protocols from Layer 2 up to Layer 4 like ARP, IP, TCP, UDP, DNS, and DHCP. It offers different APIs for efficient access, multi-threading, and an implementation of the Berkeley Socket API. Internally, LWIP does not represent a complete data frame but stores a subset of data in `pbuf` structures. Other stack implementations like FreeRTOS+TCP¹ offer the complete frame and stick to the Berkeley sockets API while being thread-safe. Using sockets, application tasks can register to receive and transmit through their desired ports.

¹https://www.freertos.org/FreeRTOS-Plus/FreeRTOS_Plus_TCP/




Figure 2: **RX-Path**: Packets are handled by OS data structures, the network driver, and the networking stack implementation before they can be accessed by the receiving application.

2.3. Interrupt Scheduling

IRQ and task priorities form two independent priority spaces. Hence, hardware interrupts, such as those triggered by incoming network packets, introduce some challenges to scheduling as they might take over CPU resources at any time. Yet, when systematically tamed to known minimum inter-arrival times and Worst Case Execution Times (WCETs), an integration into the considerations of a schedulable task set becomes possible. More problematic is the triggered execution of ISRs in an elevated IRQ-context. It may be either completely uninterruptible itself, or only by another higher priority IRQ-source (“interrupt nesting”).

To minimize the worst-case latency incurred by priority inversion situations between interrupts and high-priority tasks, a widespread programming best practice is to reduce the work done in an ISR to a minimum, only unblocking a deferred Interrupt Service Task (IST) that then does the actual processing. This compromises on interrupt handling performance for better scheduler control. Therefore it is often weighed by the driver developer how much additional latency is acceptable until an ISR/IST split is introduced.

2.4. Aperiodic Scheduling

Received network packets generate workload that can be characterized as an aperiodic task inside the RTOS. One approach to integrate aperiodic tasks into fixed-priority scheduling uses so-called server tasks [60]. To the scheduler these behave as ordinary prioritized tasks. Opposed to other tasks they have no individual objective. Instead, they use their budget to serve the execution of aperiodic jobs. Due to their limited budget in each period they can be easily included into scheduling considerations.

A very simple yet effective aperiodic events server is the deferrable server [62], which we also utilize in our real-time receive architecture. It has a limited CPU-time budget to serve aperiodic events. When the budget is depleted, it pauses execution. At the end of each period, the server budget gets restored to the initial amount. A big advantage is the simplicity of the mechanism and therefore of an implementation for that server scheme. Yet, the deferral of budget consumption incurs on a higher worst-case processing demand than one budget per period, due to possible back-to-back execution patterns.

Let p be the server period and e its execution budget for a period. There may arrive jobs just before the end

of a period consuming the whole capacity e of the server for this period. With the start of the next period and the consequent budget replenishment, another duration e may be serviced to jobs. In the worst case, we need to expect one extra execution budget. Thus, the highest possible server demand $d(\Delta)$ inside an arbitrary interval can be indicated as

$$d(\Delta) = e \cdot \left(\left\lceil \frac{\Delta}{p} \right\rceil + 1 \right)$$

Note that at least for a small period p , the CPU bandwidth $\frac{d(\Delta)}{\Delta}$ still approaches the theoretical server optimum $\frac{e}{p}$.

2.5. Interrupt Moderation

To decrease the performance impact of incoming packets, high performance NICs employ interrupt moderation techniques. Instead of sending an interrupt for each received data frame, the NIC delays the delivery of an interrupt in order to receive and coalesce additional packets [42]. Different strategies to realize the delay exist.

A simple approach is to use a *packet counter* that triggers an interrupt and resets once a certain number of packets have arrived. This leads to a constant and homogeneous reduction of interrupts but also introduces the possibility of starving packets and very unpredictable packet delays. To have control over the time packets reside in memory unnoticed, different types of delay timers are applied:

The *absolute timer* begins a countdown once a packet has been received and only triggers an interrupt once reaching zero. All packets received in this time frame are announced by this interrupt and do not reset the timer. The obvious disadvantage of this approach is the high latency the first packet of each countdown experiences. In low traffic scenarios, this is highly inefficient.

To this end, *packet timers* can be introduced. Instead of having a relatively long countdown timer to trigger an interrupt for multiple packets, the counter is a lot smaller and resets with each incoming packet. In low traffic scenarios this leads to smaller delays while interrupts can be entirely impeded under high traffic. The mostly applied solution is therefore a combination of a longer absolute timer and a shorter packet timer. The tuning of the specific parameters is highly dependent on the expected load and subject to research in high performance computing [45].

3. Related Work

The introduction of unpredictability in real-time environments through interrupts has been a long-standing research topic. In the following, we present past approaches to mitigate interrupt impact as well as approaches towards real-time aware network processing.

3.1. Interrupt Management in Real-Time Systems

The Advanced Interrupt Controller [21] monitors the priority of the currently running process to determine if an interrupt should be triggered or held back by comparing it to the interrupt priority. A simple extension of the interrupt controller unifies the priority spaces of attached interrupts and operating system processes. However, this does not facilitate for the circumstances around network packets since interrupt priorities of all packets are the same and different packets cannot trigger interrupts of different priorities.

Prominent work regarding the unification of priority spaces is the approach implemented in the *Sloth* OS [24] [25]. Sloth implements a general abstraction for software threads and ISR, abolishing their distinction. Instead of using a software scheduler for threads, every control flow is designed as a thread-related system call using the hardware interrupt system. By letting the hardware manage all control flows, context switches have less overhead and — which is more interesting here — ISR and (other) threads preempt each other in accordance with their priorities. While this abolishes the problem of priority inversion, high packet loads still lead to high interrupt frequencies, impacting real-time tasks.

The priority inversion impact of interrupts in real-time systems has been identified and tackled by Amiri et. al. by employing priority inheritance protocols for interrupt service threads [2]. This approach however only works for the schedulable part of interrupt handling of device drivers.

Using interrupt moderation to relieve the CPU in high traffic scenarios is a method studied mainly for high throughput devices, as systems connected to Gigabit Ethernet networks are subject to potentially millions of packets per second [19]. However, some work also exists studying embedded devices running the Linux kernel. Spanos et al. evaluate the performance implications of advanced interrupt handling techniques in the "New API" Linux device driver extension [59].

The issue of DoS attacks in industrial IoT environments has been addressed by Niedermaier et al. [50]. A dual microcontroller architecture is proposed to separate networking tasks from critical real-time processes. The presented setup highlights the disproportionate processing requirements of IP networking on microcontrollers. While this does mitigate the effects of DoS attacks, it is not obvious how time critical packets are separated from attack packets.

The PIERES tool [7] is a small framework running on real-time IoT devices, designed to analyze the real-time behavior under different network loads and hardware configurations. This playground allows developers and researchers to perform network interrupt experiments on real-time embedded systems with different network interface controller implementations, load generators and timing utilities.

3.2. Real-Time Aware Packet Processing

The network stack architecture *Lazy Receiver Processing* (LRP) introduced an important and much used approach still interesting today [13]. It improves performance, stability and fairness on server systems with high incoming network throughput. The processing of newly arriving packets can cause persistent crowding of the application processes so that they are not able to receive data. The packets then have to be discarded while the applications continue to starve. This situation is prevented by early multiplexing of incoming packets, prioritized execution of the rest of the protocol stack in the context of the receiving process, and thus enabling early discarding of packets on congested paths. This approach would benefit from a pre-sorting of packets by the hardware and the general interrupt decrease attained from our multiqueue NIC. The decrease in throughput as the packet rate increases can thus be prevented or mitigated. By consistently differentiating different network flows, they take an elegant approach that may also efficiently improve real-time behavior in IoT devices.

Building atop the idea of LRP, Lee et al. investigated on reducing the impact of Low Priority (LP)-packets on the real-time behavior of a network-independent task by introducing port-based prioritization of protocol processing [35, 34]. In order to achieve this, they classify a UDP-packet by its port in a “top half” interrupt handler. By looking up a special port-priority-table, whose data is sourced by all bound UDP-sockets, their top half infers a priority for each packet. In Linux’s `softirq` scheduling entity belonging to the kernel, they introduce a gate functionality: Packets are only ever processed as long as their priority is higher than the current active priority of the system. Otherwise their processing is delayed until at least the next regular `softirq` invocation. They consequently show how their modification is able to reduce a long-running critical task lateness measurably. However their implementation is restricted by the inappropriate scheduling behavior of the `softirq`-handler in Linux, which is not preemptable even by the most critical processes and gets rescheduled in similar way as polling, adding unnecessary high network latency once packets aren’t processed eagerly anymore. Moreover, their work only considers UDP-packets. Finally, when considering overloading scenarios a mere flow differentiation and prioritization is not sufficient for protecting execution guarantees, since packets may also arrive at a highly prioritized task port in high quantities.

The time-predictable IP stack `tpIP` [54] addresses the challenge of real-time communication in cyber-physical systems. To enable timing predictability and WCET analysis the proposed stack uses polling functions in the socket API with non-blocking read and write operations. While focusing on timing analysis and predictability, no measures are taken towards processing performance, interrupt scheduling or the issue of traffic overloads.

Strategies presented in [12] deal with the detection and mitigation of network packet overloads in real-time systems. The *Burst Mitigation* approach, limits the amount of IRQs that may get processed in a time slice, effectively applying a deferrable server scheduling scheme which considers each IRQ a standard-sized job. While the work does not consider differentiating mitigation measures over different packet flows, the evaluation already hints the practicality of simple mitigation techniques that can be used beneficially in our approach. The *Queue mitigation* enables back-propagation of frame queue stagnation to the IRQ, making it adopt the priority of the network task with some delay. This has great effect when the network task is only equipped with medium priority and no particular small scheduling latency is required. Two other approaches deal with dynamic schemes, taking advantage of the critical task slack time. The evaluation indicates the effectiveness of the static mitigation approaches. The dynamic ones however strongly depend on the cooperativeness of the critical task and its execution time consistency.

The Linux kernel provides advanced networking capabilities for routing and traffic control (tc) [27]. The latter provides mechanisms to control IP traffic such as traffic shaping and forwarding. Using queuing disciplines (qdiscs), incoming and outgoing packets can be queued for each network device. Traffic rates can be moderated and packets can be assigned priorities on the basis of meta information. Similar to the approaches presented in this article, packets can be matched to certain descriptions and inserted to a qdisc depending on the packet classification. Recent work also shows that the novel Linux kernel features eBPF and EDT can be used to improve the scalability of filters in Linux-based networking [4]. Due to the packet-proportional overhead, RTOSs do not commonly provide similar tools.

3.3. Specialized Hardware

Network Interface Controllers with multiple transmit and receive queues have been introduced by Intel as early as 2007. The goal is to make use of multicore systems by parallelizing network load on the different queues. The trend is to increase the number of queues to facilitate cloud computing as Zhu et. al. showed in 2020 [69]. Multiqueuing in general exists for different high-throughput I/O devices but to our best knowledge is not a common sight among real-time or embedded architectures.

Some modern NICs support serving received packets into multiple BD rings [57]. For once, this is useful to distribute packet reception work over multiple CPU cores in high-performance scenarios. The assignment of packets to these queues can also result from a hardware classification based on packet header fields, which may be configurable by the driver [56, 66]. However, this is a feature only found in advanced NIC hardware [26].

Loom [61] is a multiqueue NIC design that moves per-flow scheduling decisions from the software network stack into the NIC. This way, high throughput and homogeneous

policy enforcement can be guaranteed while also providing isolation in multi-tenant cloud data centers.

In [41] Lonardo et al. present an application specific NIC design run on FPGAs for high energy physics experiments. The design allows a remote DMA to CPU and GPU memories, relieving the OS from data transfer management, allowing real-time processing on data received over the network.

3.4. Industrial Networking

Low latency, predictability and reliability are longstanding requirements in industry automation. In the 2000s field bus technologies were moved to work via Ethernet connections. Industrial Ethernet encompasses all usage of Ethernet in an industrial setting. The target applications typically have both latency and reliability requirements and this drives the design of protocols away from traditional Ethernet approaches for collision detection and avoidance. The most common traditional real-time Ethernet protocols are EtherNet/IP [9], PROFINET and EtherCAT [28]. EtherNet/IP uses the Common Industrial Protocol (CIP) over Ethernet to controlling real-time devices on the network. PROFINET prioritizes traffic on one physical Ethernet network based on real-time classes. Particularly, the PROFINET-IRT class caters for stringent real-time requirements, leveraging globally synchronized time-triggered switches and interfaces and supporting sub-millisecond cycles. EtherCAT is particularly relevant in industrial control and automation due to its speed and determinism. It is specified to offer cycle times of less than $100\mu s$ by bypassing transfer processing on the slave device and traffic prioritization.

In 2011, the first part of what was to be known as Audio/Video Bridging (AVB) was published in a series of IEEE Standards for switched Ethernet [29, 30, 52, 51]. With AVB, later renamed to TSN a network can be configured to provide Quality of Service (QoS) guarantees that allow real-time traffic to flow through the network with extremely low packet loss and with predictable and bounded latency. The need for TSN has been reinforced by Industry 4.0 requirements. Several standards exist to guarantee time synchronization, packet delivery and maximum latency in industrial environments where contracts can be made between senders and receiver [15]. Approaches to minimize interrupt loads are not part of proposed TSN protocols.

In this work, we assume the absence of an industrial networking infrastructure. Hence, aperiodic and unexpected bursts of traffic are possible. Compared to the referred works making similar assumptions, there are none that address the vertical discrimination of packet handling based on flow priorities. The following sections will describe and evaluate such approaches.




Figure 3: **Data flow:** Path of received messages through the regarded device.

4. Preliminary Considerations

This section highlights considerations concerning the environment, devices, and data flow for both presented approaches.

For the purposes of this article we consider an IoT device to be an embedded system running an RTOS and a small number of processes with fixed priorities managed by a preemptive scheduler. The embedded device is a constrained IoT node in an IP network serving real-time applications as well as best-effort networking. No specialized real-time network protocols or infrastructure are available to the device. In the network stack, a driver controls DMA transfers, establishes cache coherency and passes packet buffers to a networking task by means of a queue. The device contains a baseline embedded NIC used to connect to an IP network which itself might be connected to the Internet.

The proposed approaches make use of IP flow information for real-time aware packet prioritization. Packets are received by real-time processes on the embedded system.

Definition 1. Let \mathcal{F} be the set of IP flows received by the regarded system. An IP flow $f \in \mathcal{F}$ is a sequence of IP packets arriving at the regarded device and for the purposes of this article characterized by the tuple $(Src, Dst_{port}, P, t_P)$.

- Src is the source node, identified by its IP address.
- Dst_{port} is the destination port number.
- P is the priority of the flow as observed by the regarded device.
- t_P is the minimum expected period of the flow, meaning the interarrival time between two packets in the flow.

The priority of a flow is a parameter assigned by the receiving system, as the real-time implications to this system need to be considered. The period of a flow depends on the sender, application, and network infrastructure.

Data Flow. Each task listens on a separate socket for messages. Each message is delivered using one IP packet. The tasks process the message and produce an output or control a physical actuator. Figure 3 shows the resulting data flow model.

4.1. Priority Inversion in the RX Path

We observe a problem of incorrect priority enforcement in the networking subsystem receive path. Generally, an incoming packet is used by (at most) one specific task. Yet, as long as its purpose is unknown, we have to schedule the processing of each packet equally. For the practical implementation of the networking subsystem in an RTOS, that means assigning a fixed priority to the protocol processing server task, inevitably creating a priority inversion situation:

- Using a high priority, as usually found in current embedded frameworks, high priority tasks may starve in favor of low priority packets.
- Using a low priority, a receiving high priority task may wait for the end of execution of a medium priority tasks.

The above formulated priority inversion cannot be completely eliminated by its inherent nature. On the one hand, we do not want to spend computing resources until knowing whether it will be worth it considering the current scheduling situation. On the other hand, incorporating the use of a particular packet requires prior protocol handling.

As already shown in [12, 5, 49], a first angle to prevent receive packet induced overloads is the introduction of budget enforcement to the entire RX-path. However, when limiting the reception of incoming packets, one might include soft real-time traffic that, while not being critical to the integrity of the controlled cyber-physical system, is relevant for baseline functionalities. This converts the catastrophic situation of potential system failure into a tenable but still undesirable situation of reduced network communication liveness, trading availability for robustness. Note that we are expecting multiple IP flows with different levels of timeliness requirements and packet frequencies enabling a different grasp on the problem.

5. The Software Approach: Priority-Aware Scheduling of Incoming Packets

The first of the two presented approaches is entirely software based. By adapting common light weight IP stacks, we aim to introduce real-time aware packet processing and mitigate the unpredictability of IP networking.

5.1. Requirements

A modification of the packet reception subsystem for real-time IoT devices should satisfy the following requirements.

- 1.1 *Priority Inversion Mitigation.* Packets of high-priority flows must be processed and delivered before those of low-priority flows.

- 1.2 *Overload Protection.* Packet floods must not lead to system overloads. In case of a packet flood, low-priority flows must be throttled in favor of high-priority processes. A general rate limitation must protect the real-time properties of running tasks.

- 1.3 *Performance Retention.* The approach should not introduce a longer packet processing delay. However, predictability and overload protection must be prioritized.

5.2. Overview

The proposed architecture is designed around a data structure of differentiated flow queues, which replaces the simple frame queue. Each flow defines a priority and a period, to affect the further processing of its packets. For the prioritization of processing, we add a priority manipulation mechanism to the protocol processing task. In order to gain a maximal advantage from scheduling the subsequent processing stage, the driver is modified to do only the necessary work of classifying incoming packets to flows by their header entries. The remaining activity is then executed on packet retrieval by the scheduled protocol processing task as presented in Figure 4.




Figure 4: **Architecture overview:** Packets are classified early and enqueued by their flow, with individual periodic capacity restrictions applied. Further processing is scheduled by inherited priority.

The proposed architecture combines three concepts:

1. *Soft Early Demultiplexing* into receiver-centric flows.
2. *Prioritized Protocol Handling* based on these flows.
3. *Rate Limitation* applied per flow as well as overall.

While these are not novel ideas independently, we argue that only in this combination they exhibit properties making for a viable solution to the discussed problem:

- *Early Demultiplexing* is necessary for differentiating flows on an End-to-End basis, without reliance on network QoS and as a result satisfy *Requirement 1.1*.
- Proper *prioritization* facilitates best-effort communication processes that utilize background resources on the same system and is necessary to satisfy *Requirements 1.2 and 1.1*.
- *Rate limitation* as a last resort protects the system from being vulnerable to unexpectedly high traffic in High Priority (HP) flows satisfying *Requirement 1.2*.

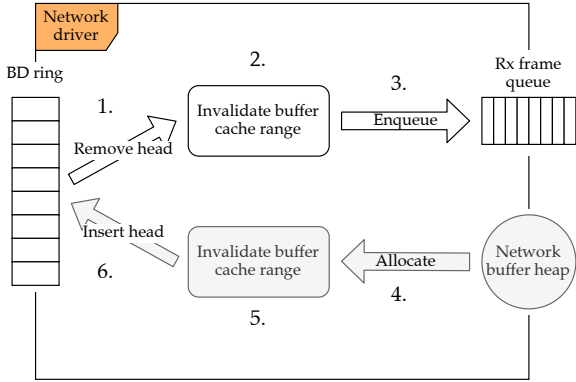


Figure 5: **Receive activity in the original driver:** Once an interrupt occurs, the driver code moves a packet buffer from the BD ring to a simple queue, and fills the vacant position with a newly allocated one. Due to cache coherency requirements of the memory system, the buffer caches have to be invalidated for both the retrieved and the replacement buffer.

Thus, being able to fully defend the considered system against flooding induced overload, while at the same time ensuring high connectivity for particular well-behaved HP-flows even in scenarios with overall high incoming traffic, and handling LP-flows with best-effort resources, this combination forms an IoT real-time aware overload protection.

In the following subsections we introduce the concept of each of the three basic building blocks of our architecture and discuss relevant implementation aspects.

5.3. Soft Early Demultiplexing

In order to minimize the effort spent until after classification, we employ Early Demultiplexing [13]. By peeking into key header entries, a packet is assigned to its eventual receiver process.

The benefit of demultiplexing performed in software depends heavily on the amount of work that can be saved by mere demultiplexing compared to full protocol processing. Since the packet scheduling in our architecture can only influence the processing that follows after Early Demultiplexing, the achievable degree of partial network liveness in overload scenarios depends on its quick execution.

Starting from the existing driver receive path, depicted in Figure 5, we introduce two changes: Packet classification and lazy cache invalidation.

5.3.1. Packet Classification

The classification differentiates incoming packets into flows defined by the protocols ARP, ICMP, TCP and UDP. While the former two form a single flow, the latter are further differentiated by local port numbers in order to implement the receiver task association.

Depending on the used network stack, the lookup from the port to a flow may either be performed using the existent network stack’s list of bound socket control blocks, or else requires an additional data structure managed by the

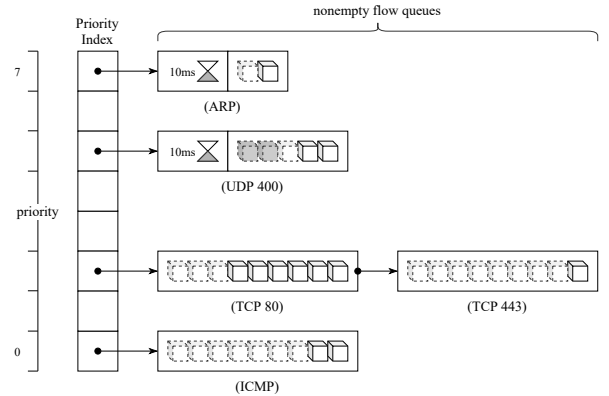


Figure 6: **Differentiated flow queues:** Between reception by the driver and further driver and protocol processing, packets get stored in a queue according to their identified flow. These queues are organized by the flow priority, to facilitate fast retrieval of the highest/lowest prioritized packet.

driver. In our prototype based on FreeRTOS+TCP, the socket managing code in the original network stack can easily be locked in a critical section, leaving the ISR safe to access it.

If a scenario requires anticipating a large number of bound sockets, a sophisticated data structure with better complexity should be employed. However, with only a few sockets bound at any particular point in time, a linear linked list lookup as found in typical embedded network stacks suffices.

Instead of enqueueing every received packet to the same RX frame queue, each packet is inserted into a specific queue according to the result of the classification. Because the packets do not necessarily get processed in bounded time, the network subsystem might experience buffer starvation. To avoid this, buffers of low priority packets are *recycled* when the buffer memory reaches its limit. Buffer recycling here means that the allocated buffer element is freed and a new empty element is appended to the BD ring. The packet is effectively dropped. To this end, the differentiated flow queues are organized in a priority queue structure as depicted in Figure 6.

The priority of a flow is defined by its respective receiver task and the overall priority space is equal to the one used by the RTOS task scheduler. This way, packet processing priorities are inherited according to the real-time considerations made for the running processes starting with the enqueueing of packets. Another feasible option is the utilization of network priorities such as the Differentiated Services Field ², flow priorities from real-time protocols, or traffic classes inside TSN-based networks. Yet, we assume environments that do not necessarily provide a network priority, thus relying on receiver priorities.




Figure 7: **Receive driver activity in our approach:** The driver is separated into two halves. In the *eager driver* (1), a minimal effort is taken to classify each packet into a flow. As part of the scheduled subsequent protocol processing, the *deferred driver* (2) establishes cache coherency and refills the BD ring once the packet is needed.

5.3.2. Lazy Cache Invalidation

On embedded systems that feature CPU-caches, the commonly cache incoherent DMA introduces a significant cost with the obligation to invalidate the transferred buffer cache lines. In our case, the memory architecture requires the network driver to invalidate buffer cache lines prior to and after the processing by the NIC DMA engine, as shown in Figure 5. As cache management noticeably prolongs the execution time of Early Demultiplexing, we incorporate a lazy cache coherency establishment scheme into the driver to retain the highest possible performance as per *Requirement 1.3*. The driver is therefore split into two halves, as depicted in Figure 7.

- (1) **Eager driver:** An immediately processed layer, executed as part of the ISR, classifies and enqueues packets.
- (2) **Deferred driver:** A schedulable layer, executed in the network task according to the packet priority, establishes full cache coherency of received packets and prepares fresh replacement buffers.

Prior to classification, only the first cache lines of the packet buffer containing the relevant header fields are invalidated. Once the packet is chosen to be processed further, the remaining part is invalidated and a fresh packet buffer prepared and appended to the DMA BD ring (cf. Section 2.1). This implies that as the differentiated flow queues fill up with packets, the BD loses free packet buffers, forming a closed pool of packets shared by the BD ring and the differentiated flow queues. To prevent starvation of the BD-ring caused by LP-packets in the differentiated flow queues, the eager driver recycles lowest priority packet buffers once the BD-ring hits a critical threshold (e.g. $\frac{1}{2}$). This can be carried out with little cost, since only the accessed header cache lines have to be invalidated again.

The resulting activity in the eager half driver is depicted in Figure 8. Notable are the three different execution paths that might be taken: If, due to a high BD ring fill state a packet buffer has to be recycled, and the currently considered packet is of lowest priority, it gets recycled in a short-circuiting branch ^(a). A flow queue may decline further packets to prevent overload by this particular flow ^(b). Lastly, if the short-circuit branch was not taken but the BD ring fill state is high, another packet buffer has to be recycled and inserted into the BD ring ^(c).

5.4. Prioritized Protocol Handling

Once the heterogeneous incoming packets are demultiplexed into differentiated flow queues, the protocol processing can be carried out with the receiver priority, as proposed in [13, 35].

We apply a priority inheritance scheme to the single protocol processing task [44]. It allows the task scheduler to preempt the packet processing at any point in time. Additionally, it keeps a low footprint in terms of task resources and can be integrated into embedded network stacks that commonly use a single network task.

To implement this scheme, the priority of the network task has to be moderated depending on the currently processed packet and waiting packets, in order to avoid priority inversion. Consider \mathcal{F} as in Definition 1, $P(f)$ the priority of flow f and $p(f)$ a packet of f . Let further \mathcal{W} be the set of currently waiting packets and \mathcal{E} the set of packets in processing. The priority of the network task ($P_{IP\text{-task}}$) must be assigned as follows.

$$P_{IP\text{-task}} = \max(P(f) : f \in \mathcal{F} \wedge \exists p(f) \in \mathcal{W} \cup \mathcal{E})$$

This assignment implies the network task priority is recomputed every time a packet gets queued or a packet has been processed. On packet reception, the priority needs to be elevated *iff* the respective flow priority is higher than

²RFC 2474

the current priority assigned to the network task. On finished packet processing, the priority needs to be decreased *iff* the priority of the highest priority packet waiting in the differentiated flow queues is lower than the current network task priority. This operation is supported by the ability of the differentiated flow queue data structure to efficiently provide the highest enqueued priority (reconsider Figure 6).

It may appear that by using priority inheritance carried out per packet, we put an overly high computational burden on the fixed-priority task scheduler. Yet, among all possible designs that involve the task scheduler in packet scheduling decisions by correctly signaling the current priority demand at each time, this design has the lowest scheduler data structure manipulation overhead: Another design could use multiple processing tasks with constant priority, to which packets are assigned according to their flow. The unblocking operation triggered when the first packet of a particular priority is enqueued then adds at least the same overhead — the task has to be moved into the priority-respective ready task list, and moved out once blocked again.

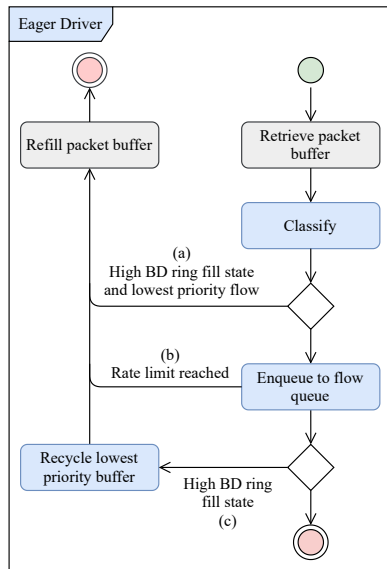


Figure 8: **Eager driver ISR**: Key execution paths that determine whether and when a packet buffer is recycled to save execution time in high-load scenarios.

In order to also have the deferrable parts of the driver processing scheduled according to packet flows, the networking task dequeues a highest priority packet buffer from the differentiated flow queues and executes the second half of the driver before continuing with the regular processing procedure.

5.5. Rate Limitation

To take advantage of Early Demultiplexing while at the same time keeping the system protected from overload conditions, deterministic mitigation techniques [12]

are applied to all but the low priority best-effort flows. Additionally, the unconditionally executed ISR that demultiplexes incoming packets could incur a high load even if the subsequent scheduling cuts off further processing. Hence, an additional, global rate limitation needs to be present.

To apply the rate limitation, each flow is scheduled by a conceptual aperiodic events server with each incoming packet being modelled as an aperiodic request. In our prototype we use the deferrable server scheme (cf. Section 2.4). Beyond the server capacity, packets are discarded. For the individual flow queues, this happens as part of the inserting operation (reconsider Figure 8), in order to avoid a situation with a paused HP flow queue full of packets blocking all other processing.

To enforce a global rate limitation, once the capacity has been reached in one period, the driver processing switches from ISR-based execution to a polling driver task, staying in this mode until the capacity is not immediately reached at the begin of a period anymore. When not processing packet receive IRQs issued by the NIC, the BD ring is filled until eventually packets are discarded by the NIC.

5.6. Policy Integration

In order to control the scheduling properties *capacity*, *period* and *priority* for a flow, we expose these to the user for each socket via the `setsockopt`-Application Programming Interface (API). Special flows such as for the protocols ARP, ICMP and those that are managed by the network stack, such as DNS and DHCP, can be configured using C macro definitions. Similarly, the scheduling properties for the global rate limitation can be defined.

5.7. Limitations

The ability to proceed with deferred packet processing after a phase of higher system load depends on the number of available packet buffers. As these buffers have to be prepared for immediate DMA operation and therefore a constant amount is dedicated to the lower levels of processing, additional memory might be necessary.

IP fragmentation cannot be dealt with properly in our architecture. To demultiplex fragmented packets, their reassembly had to be done in the ISR, jeopardizing its WCET. This design treats all packet fragments as belonging to a background priority flow. Yet, IP fragmentation is discouraged, as it introduces robustness, reliability and security issues [31, 20].

6. Evaluation of the Software Approach

In this section we present empirical results collected from our prototypical IP stack implementation and subsequently discuss the effectiveness of the approach.

6.1. Test Setup

The test setup contains the FreeRTOS operating system with a modified FreeRTOS+TCP stack running on a Xilinx Zynq-7000 processing system containing a dual-core ARM Cortex A9. Networking is done through a Gigabit-class Ethernet interface controlled by a Marvell 88E1518 Physical Layer (PHY) controller. Notable features are DMA and TX/RX-checksum offloading. Measurements are taken on a single core.

Two methods for measuring the effect on system load under high packet loads were pursued:

1. *Passive*: A background worker carries out CPU intensive work and monitors its performance.
2. *Active*: The software is instrumented to indicate notable events, i.e. task switches, IRQs, and packet processing.

The former is suitable for precisely estimating the average load that a particular scenario puts on the CPU. While the latter introduces some overhead in the range of 1-5% to the processing and misses some of the IRQ switching, it allows us to evaluate the distribution of processing-induced latency.

6.2. Experiment 1: CPU-Time Saved with Early Demultiplexing

In this scenario two UDP sockets are bound, one with a low and one with a high priority receiver process. To not alter the results, the capacity of all flows as well as the overall IRQ limitation is set to infinity.

Multiple system configurations were confronted with a zero-length UDP-packet load of a constant rate for 60 seconds. Through passive measurement performed by a medium-priority task, the average CPU processing time per packet was then calculated (Figure 9). In this experiment we observed that the CPU costs for processing a single packet are rather independent from the magnitude of incoming traffic, staying approximately constant in the range from 10^2 to 10^6 pkt/s.

Results

The results show the difference in processing time between the packet processing paths. When LP packets get no chance to be scheduled, the executed activity is only that of the Early Demultiplexing ISR with an average processing duration of $1.62\mu s$ per packet. Compared to the original stack as a baseline, which needs $12.1\mu s$ to fully process a packet, this results in a speedup of 7.5x. However, due to the short-circuiting logic depicted in Figure 8 ^(a), in this constant LP-flow measurement the packet buffers are discarded without being placed into a flow queue. When disabling the short-circuiting code path, the per-packet processing time increases to $1.75\mu s$, still yielding a seven-fold speedup compared to the full processing in the original stack.




Figure 9: **Processing impact:** CPU time per zero-length UDP packet under loads between 10^2 and 10^6 pkt/s with different configurations.

¹Short-circuiting branch disabled. ²Cache invalidation deferral disabled.

In this scenario, the HP packets are processed the entire network stack and cause a processing time of $12.3\mu s$ per packet, decreasing receive performance by 1.7 % compared to the baseline stack. This already small relative difference would decrease further if the subsequent (obligatory) reception by the receiver task was taken into account.

By modifying the prototype to again eagerly establish cache coherency in the ISR, the time spent for LP packets increases notably to $4.4\mu s$. Hence, we conclude that incorporating a driver deferral mechanism into the architecture is essential to the performance on cached systems.

6.3. Experiment 2: Packet Processing Latency

The second experiment deals with the predictability of packet processing latencies in the modified IP stack. Using the active approach, the reconstruction of precise execution times of each packet is possible. Additionally, this allows us to differentiate between the execution paths of the modified driver. Since we instrumented the ISR entry, some constant IRQ overhead due to context saving is not included in this analysis. Compared to Experiment 1, where the overall impact on CPU time is measured, this experiment measures the duration of the eager driver per packet.

The system was flooded with 10^5 zero-length UDP packets of two different priorities successively. Figure 10 visualizes the distributions of ISR processing duration for specific processing paths. For each distribution, the quantiles 0%, 90%, 99%, 99.9%, 99.99% are visualized as horizontal bars, in order to estimate a probabilistic WCET.

Results

LP packets initially take the fastest path ("regular"), where incoming packets are enqueued without any other processing. Once the BD ring has reached a high fill state, packet buffers have to be recycled. Since the incoming packets are already at the lowest level present in the differentiated flow queues, the short-circuiting path ("short-circuit", ^(a) in Figure 8) is taken.




Figure 10: **Eager driver:** Latency distributions for different execution paths in our modified stack: The horizontal bars indicate the percentiles 0%, 90%, 99%, 99.9%, 99.99%

HP packets in contrast can cause a noticeable increase in ISR processing time. At each occurrence of such a packet, the network task priority has to be increased in order to be scheduled subsequently ("prio+"). In case the BD ring is already filled by previously received LP packets now waiting inside their flow queue, a revocation is needed, adding further processing time ("recycling; prio+"). We also investigated on HP packets that get rejected from their flow queue ("mitigating"), yet they behave similarly as shortcircuited packets.

The results show that the first three execution paths are similarly fast, while the ones that include an increase in priority or recycling operations are more costly. As we discussed in section 5.4, a priority increase can only happen if the flow priority of a received packet is higher than the one of all the currently enqueued ones. Without the network task being active to process packets and lower the highest enqueued priority again, this is only possible once for each flow in a cascade of increasingly prioritized flows. Thus, when the system is flooded for some time and LP packets start building up in their queues, eventually the faster paths of the ISR will be taken.

6.4. Experiment 3: Mitigation and Prioritization

The final experiments show the effect of protocol processing prioritization and rate limitation, applied both for an individual flow and globally. Experiments were conducted for multiple combinations of packet flood rates for a HP- and LP-flow, respectively, over a duration of 3 seconds each. Again, a medium prioritized task measured the CPU load passively, preventing the scheduling of LP packets. Additionally, a receiver task was employed for the HP flow in order to count the packets that arrived at their destination. Figure 11 shows the CPU utilization and the ratio of successfully received HP packets to sent ones, as a function of both packet rates.

The original stack was slightly modified to feature an overall ISR rate limitation, in order to allow a meaningful comparison to our approach. It is implemented by switching to polling mode once the capacity is reached for a certain period, similar to the one employed in our prototype.

In this experiment, the limitations is set at 3 packets per 2 milliseconds.

Results

The CPU utilization increases linearly along with both packet rates, until the global limit of 1500 pkt/s is reached. Once the polling mode is active, the CPU load drops noticeably. This can be accounted to the performance improvements gained by switching to a polling-based retrieving activity that handles multiple packets at once. Further increasing the packet rate causes more HP packets to be discarded by the NIC.

For the modified stack, parameter values anticipating a similar worst case CPU utilization were chosen. We configured a high priority flow to allow one packet per millisecond and an unbounded low priority flow. The ISR was limited to processing 7000 packets per second.

The CPU load also increases linearly with both packet rates. As we would expect from the results of the first experiment, the load increases much slower with increasing LP packet rates (notice the denser scale). Above 1000 pkt/s (blue line) of HP packets, the utilization stagnates as processing of further packets is cut off by the flow queue. The additional triggered ISR executions are negligible at this scale. When the sum of both rates exceeds 7000 pkt/s (black line), the CPU utilization also drops with polling activated. Regarding the liveness of the HP flow, we can see how it continuously decreases above the flow-specific rate of 1000 pkt/s. Additionally, the global limitation impacts the HP flow. So, independent of the HP flow rate itself, the communication liveness drops as the system is flooded with LP packets.




Figure 11: **Results:** CPU utilization and HP flow liveness at various packet rates on our modified system versus the original system employing only an overall rate limitation.

The blue and black lines mark flow-specific and overall rate limitations respectively.

When comparing the approach to a simple mitigating stack as a baseline, it becomes clear that it cannot help

with processing higher rates of important packets. Instead, supported by fast Early Demultiplexing and individual prioritization against the remaining tasks, it allows to postpone an overall limitation. This way, a system can sustain a much higher load of less important packets before real-time disturbing effects start to occur.

7. The Hardware Approach: Priority-Aware Interrupt Moderation

Improving the prioritized real-time performance of network drivers still leaves a door open to overwhelm the system with network-generated interrupts. Since the IRQs are triggered outside the operating system’s sphere of influence, smartly moderating these interrupts requires some modification to the generating hardware, in this case the NIC. This section proposes an extension to the receive functionality of NICs to minimize IRQs under high load while maintaining short packet receive delays for critical tasks.

7.1. Requirements

We specify four requirements to the NIC adaptation appropriate for real-time IoT devices with control over physical processes.

- 2.1 *Interrupt Moderation.* The danger of network-generated interrupt floods and unpredictable networking overheads should be mitigated by reducing the number of interrupts triggered by the NIC.
- 2.2 *Packet Prioritization.* Solely limiting the number of interrupts leads to an analogous increase in receive delays as packets are accumulated before a notification occurs. Furthermore, the system can still be flooded with unrelated packets forcing the operating system to handle them. Packets need to be classified and filtered before they reach the operating system.
- 2.3 *NIC Parametrization.* To effectively prioritize and filter packets in different environments, the introduced NIC needs to be configurable. It needs to be possible to affect generated receive delays per process and tune them for the scenario’s real-time requirements.
- 2.4 *Continuous Configurability.* The necessary configuration of the NIC needs to be possible dynamically during runtime to facilitate changes in processes and the environment.

7.2. Overview

The problem of network-generated interrupts affecting system performance can be solved by interrupt moderation (*Requirement 2.1*). However, the techniques presented also have drawbacks. While they increase the overall efficiency of interrupt processing, they also increase the resulting packet delays and make them less predictable since packets are held back for a variable amount of time.

In real-time systems, where process runtimes depend on incoming traffic, the occurrence of inaccuracies must be minimized. Therefore, a potential NIC design must attempt to reduce network overhead while guaranteeing low and constant latency for critical packets. As presented in Section 3.3 some specialized hardware exists running on FPGAs for specific real-time application types as well as multiqueue NICs for better multicore performance in data centers. However, to the best of our knowledge none exist for embedded IoT hardware or real-time processing in general IP networks.

To this end, the interrupt moderation parameters are designed to be reconfigurable. By minimizing the relative packet delay, the time a packet dwells in memory before it is processed is reduced. At the same time, however, this increases the total number of interrupts, which reduces efficiency. This problem is not unique to real-time IoT devices, but is a natural consequence of interrupts. In closed IIoT environments, however, we can take advantage of this fact. Since embedded systems typically perform a fixed set of specific tasks, we can use their metadata to filter and manage incoming packets before interruption at the hardware level. These are the priorities of the protocols or packet-receiving processes and their associated IP flows. Interrupt moderation can thus become a tool to enforce priority-compliant traffic scheduling before they enter the realm of the operating system.

7.3. Hardware Modifications

The modifications made to the NIC concern only the reception of packets and begin after frame validation at the MAC layer. An illustration of the design can be seen in Figure 12. To accommodate incoming packets belonging to different real-time processes, the receive buffer holding packet descriptors of the NIC is divided into multiple queues realized as ring buffers. This way, packet descriptors are assigned to different queues depending on their destination process and its priority, accounting for *Requirement 2.2*.

The metadata of validated packets is compared to a list of registered ports residing in a distribution map on the NIC. Here, packets are assigned to queues which hold packets of one IP flow each. According to the process priority and expected packet load, different interrupt moderation configurations (e.g. delay timers and counter threshold) are applied to them by the operating system.

This way, packets for critical processes trigger interrupts immediately upon reception while less important packets (packets with low priority receiving tasks) are coalesced before one interrupt is triggered for all packets in the respective queue, indicated by the millisecond specifications in Figure 12. Packets with no associated process can be dropped before an interrupt is triggered since these packets would be dropped by the operating system at a later point in any case, but after generating unnecessary ISR and network stack work. This is especially important under high unanticipated traffic loads targeting the




Figure 12: **Multiqueue NIC:** Traffic is organized into different queues with exemplary delay values attached; configurations of queues and mapping are performed by the operating system upon socket binding.

device and potentially leading to a denial of service. Finding appropriate queue configurations for different process priorities is part of the design process of the embedded real-time system as introduced delays need to be part of scenario modeling.

7.3.1. Synthetically Added Bursts

Coalescing packets in the NIC reduces the number of interrupts triggered, ISRs run and context switches performed. However, the amount of data to be processed by the network stack remains unchanged for packets registered for one of the running processes. Depending on the number of packets coalesced, interrupt moderation might lead to an accumulation of network stack workload into bursts. The necessary runtime to process an incoming packet is a lot smaller than the delay introduced by coalescing packets. A meaningful delay through these bursts can hence only happen under extremely high packet rates. We consider this when choosing the queue parameters as follows.

7.3.2. Relevant Parameters

The multiqueue NIC introduces four parameters affecting packet delays and resource utilization as posed in *Requirement 2.3*:

- *Number of queues m .* The number of queues the receive buffer is divided into depends on the number of currently active processes accepting packets and supported protocols.
- *Size of a queue n_q .* The number of elements of a queue q corresponds to its expected packet load, available memory, and moderation parameters.
- *Absolute queue timer values $t_{abs}(q)$.* Periodic duration until an interrupt is triggered by the queue q .
- *Packet timer values $t_{pack}(q)$.* Amount of time after a packet is received by the queue q that triggers an interrupt if not reset by another incoming packet.

Additionally, the system introduces one implicit parameter:

- *Maximum expected packet rate $R_{max}(q)$.* The maximum expected packet rate of the flow corresponding to a queue q . Equal to $\frac{1}{t_P(f)}$ as per Definition 1.

The timer values are used to span a time window of how long a packet remains in the queue. Depending on the packet rate, a variable number of packets is then coalesced to be announced by one interrupt. As these parameters have a high impact on the timeliness of incoming traffic and generated workload on the real-time device, the accuracy of their configuration is of high importance. While the number of queues is directly dependent on the current number of active (i.e. socket binding) processes, timer values and queue sizes have to be cautiously chosen. To be able to sensibly choose the parameters knowledge about expected packet rates and slack times of real-time processes is necessary.

The added packet processing time needs to be accounted for when developing an IP-connected real-time system. Process deadlines must allow for enough slack time for the system to handle concurrent packet reception. The higher the slack times are, the more packets can be processed without resulting in deadline misses. The worst-case scenario is subject to high interrupt rates affecting the process with the smallest slack time. For the calculation of appropriate queue parameters this value has to be factored in. The parameters must be chosen respecting the following considerations.

Queue size. Choosing an appropriate queue size affects memory consumption as well as the maximum number of packets that can be coalesced in one interrupt. Applying interrupt moderation generally means holding more unprocessed packets in memory ultimately increasing the demand for the whole system. Memory implications for queue structures behave analogous but in a much smaller scale as only descriptors are held.

The more packets can be held by one queue, the bigger the packet burst to be processed by the IP stack may become. Hence, this parameter also enforces an upper limit for the incoming packet rate per queue as elements are dropped when new packets arrive at a full queue. Additionally, this value has to be kept small enough for the IP stack processing time to be smaller than the minimum slack time when all queues generate a burst at the same time. The resulting delay corresponds to the Worst-Case Packet Processing Delay (WCPD) which depends on the per packet processing time t_{netstack} .

$$\text{WCPD} = t_{\text{netstack}} \sum_{q=0}^{m-1} n_q$$

Absolute timer value. The absolute queue timer realizes the upper latency bound of the interrupt rate window. To minimize the number of interrupts, this parameter needs to be maximized. At the same time, a higher absolute timer value also leads to a higher added latency an incoming packet might experience. Hence, the absolute timer value is limited by the maximum additional delay the underlying process can handle while still meeting its deadline. As a high value increases the burst of packets to be processed under high loads, the chance of the queue filling up increases, leading to packet loss. The maximum expected packet rate R_{max} per queue q needs to be factored in. t_d is the process-specific maximum allowed delay.

$$t_{\text{abs}}(q) \leq \max(t_d(q), t_{\text{qf}}(q)) - \text{WCPD}$$

$$t_{\text{qf}}(q) = \frac{n_q}{R_{\text{max}}(q)}$$

Packet timer value. The packet timer realizes the lower latency bound of the interrupt rate window. Increasing this value generally decreases the number of interrupts as there is more time available for a new packet to arrive and reset the timer. This also means, that this value directly influences the minimum additional latency an incoming packet experiences. At the same time the amount of interrupts is highly dependent on the incoming traffic shape. The worst case in terms of interrupts generated is a packet rate corresponding to the packet timer value (one interrupt per packet). Hence, the available slack time needs to suffice to handle the number of interrupts generated by all packet timers combined when every timer iteration of each queue triggers an interrupt. $t_P(f)$ is the expected packet arrival period as per Definition 1.

$$t_P(f) \leq t_{\text{pack}}(q) \leq t_{\text{abs}}(q)$$

7.3.3. Configuration

As shown in Figure 12, two interfaces are used for the NIC configurations. One for setting the queuing parameters mentioned earlier and a second to write process-to-IP flow mappings to the distribution map. Both configurations are performed when a socket is bound using the

network stack API (see Section 2.2). For this purpose, the socket API is extended with driver calls that make the specific changes. Whenever a new process registers or releases a socket, the operating system transparently adjusts the number of queues and their parameters. The delay times and the size of the queues must be set for specific scenarios. The required information is passed to the driver as socket binding parameters.

As stated in Requirement 2.4, the system must be dynamically tunable at runtime to facilitate changes in processes or IP flows. Since the configuration process is linked to the socket API, the required tuning parameters can be passed at any time by the registering process. In the same way, NIC queues are released when a socket is no longer bound.

Since a network packet does not contain explicit information about the receiving process, a mapping must be made between the packet metadata and the processes. For this purpose, a mapping between IP flows and processes is created and placed on the NIC. In this design, the destination port is used to map a packet to a process. However, it could also be extended to match specific tags. This is be useful when all components of the distributed system are under the control of the developer to add security measures. The map must be on the NIC itself to cause minimal additional distribution delay, and still be configurable by the operating system to reflect the current set of existing processes.

7.3.4. Memory Implications

The presented approach has implications for memory usage on two levels: Firstly, the network buffer on the host system needs to be able to hold packet contents until they are processed, even when multiple packets are coalesced. Depending on the timer values and packet rate this might be a multiple of the usually necessary space. The network buffer resides on system RAM and is accessed by DMA. The second level is the required memory on the NIC. In an example implementation with 32 bit addresses, 32 bit timers, and a conservatively chosen maximum packet queue length of 65,536 KB the on-NIC memory necessary for one table entry is 30 Bytes as broken down in Table 1.

Table 1: Required register memory per queue on the NIC for an example implementation.

Component	Size
<i>port_id</i>	16 bit
<i>base_address</i>	32 bit
<i>buffer_size</i>	16 bit
<i>offset</i>	16 bit
<i>next_base_address</i>	32 bit
<i>packet_timer</i>	32 bit
<i>absolute_timer</i>	32 bit
<i>packet_timer_expiration</i>	32 bit
<i>queue_timer_expiration</i>	32 bit

The *base_address* field contains the RAM address of the beginning of the queue packet memory. The *offset* is incremented for each incoming packet by its size. The *next_base_address* is switched for the *base_address* when an interrupt for the queue occurs to be able to receive new packets while the buffered ones are processed. When reserving one queue for non-transport layer protocols such as ARP (not requiring interrupt moderation fields and hence being 14 Bytes wide), the total memory requirement is

$$m \cdot 30 \text{ B} + 14 \text{ B}$$

for m table entries on the NIC.

8. Evaluation of the Hardware Approach

The proposed NIC extension reduces the number and frequency of interrupts caused by incoming packets. Yet, as packets that belong to registered processes are not dropped, driver and network stack workloads remain in a time shifted manner. We evaluate the resulting timing implications by conducting three sets of experiments: The first explores the ability to reduce interrupts. The second compares the robustness of the real-time system against high traffic loads. The third analyzes the effects of different queue configurations under expected loads.

8.1. Test Setup

As the design proposes changes to hardware but an evaluation on a real IoT device is necessary for plausibility, the evaluation platform comprises of two layers. One layer assuring plausibility by providing a real IoT device running real-time processes (process layer) and one allowing hardware and configuration changes to the NIC (NIC layer). The evaluation setup is depicted in Figure 13.

To evaluate the real-time behavior of a running IoT system, we used an ESP32³ microcontroller for the process layer. It is equipped with a dual-core CPU and widely used for IoT tasks that involve communication via WiFi and Bluetooth. The two cores of the device permit a separation between observed processes and testing system. The observed processes run on the real-time operating system FreeRTOS where scheduling is performed preemptively on basis of process priorities. To generate realistic network loads, traffic traces from common industrial control systems are used.

8.2. NIC Layer: Simulator Implementation

The upper half of Figure 13 illustrates the NIC layer. A traffic generator pre-processes captured network traces and synthetic load patterns to generate a receive traffic trace for the NIC simulator. The simulator is written in Python using the event-based simulation library *SimPy*. It is configured for each experiment run (as explained in

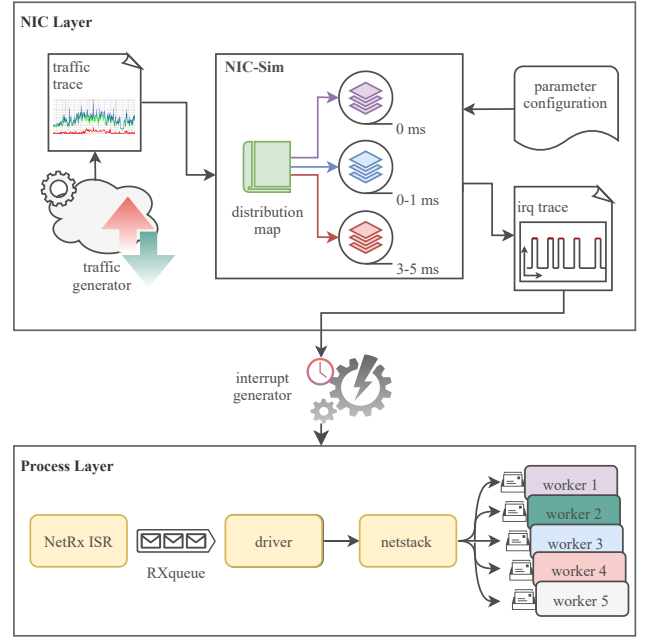


Figure 13: **Evaluation setup on two layers:** Interrupt/packet trace generation by a NIC simulator and process observation on a real-time IoT device.

Section 7.3) depending on relevant IP flows and processes. With the possibility to change NIC parameters, different interrupt traces can be created from the same network packet stream. These traces additionally contain packet metadata for use by the process layer.

8.3. Process Layer: Network Stack Implementation

The interrupt traces generated by the simulator are applied to the processing layer by an interrupt generator implemented on the ESP32. Due to the interrupt moderation, each interrupt notifies the network driver of a batch of one or more incoming packets. The NIC interrupt service routine (NetRx ISR in Figure 13) receives this batch of packet descriptors and appends them to an operating system queue to be fetched by the network driver. From here, each packet is processed by the network stack task. If the packet destination port has a socket registered to it, the packet descriptor is forwarded to the socket mailbox and the associated task is notified. Otherwise, the packet is dropped.

The receiving real-time worker processes get access to sockets through the socket API. Using a receive function, the processes can then read data from the socket mailbox. This approach implementation of Berkeley sockets corresponds to the mapping of IP flows to processes. The receiving processes are workers of different priority. Each incoming packet is tracked from its time of arrival at the NIC until processing in its worker process where it triggers the task workload.

³<https://www.espressif.com/en/products/socs/esp32>




Figure 14: Histogram of packets and caused interrupts over time with bin size of 5 seconds.

8.4. Experiment 1: Interrupt Generation

In all experiments, the IoT device runs four worker processes of different priority. The processes receive traffic using the widely used industrial communication protocol MODBUS/TCP⁴. As each of the processes binds their own socket, four queues with different interrupt moderation configurations are set up in the NIC.

Table 2 shows the queue configurations for the worker task flows. Queue 0 is configured to receive the IP flow of the critical task, hence no moderation is applied to this flow and packets are forwarded immediately upon arrival. Queues 1-3 are moderated with increasing delay values. Queue capacities are kept at a constant 128 packets per queue. Additionally, one baseline experiment is performed without any interrupt moderation. We observe the progression of interrupts generated in respect to packets received.

Table 2: Queue configurations for baseline process IP flows in milliseconds.

queue	absolute timer	packet timer
0	unmoderated	
1	30	20
2	40	30
3	80	70

Results

The number of interrupts generated depends on the number of packets received in each queue and their configuration. Figure 14 shows a comparison of packet and interrupt numbers for the baseline experiment without additional load. Queues 1 - 3 moderate interrupts in different time windows, so they generate fewer interrupts than queue 0, which is receiving packets for a critical task.

8.5. Experiment 2: Unfiltered Packet Flood

To observe the system under high traffic, it is subjected to packet floods ranging from 0 to 15000 packets per second. The worker setup on the device stays as defined in Experiment 1. To observe the effects of packet floods

when they are targeted at unregistered sockets they are subjected to a separate moderated NIC queue. We evaluate the compute load of the flood on the system. The experiments are repeated on four different absolute delay timer values for the added packet floods. The absolute timer values range from $800 \mu\text{s}$ to $3200 \mu\text{s}$ resulting in the designations *nomod* (for unmoderated flood traffic), *d800*, *d1600*, *d2400*, and *d3200*. As we are testing the system under higher than expected load, the packet timer can be disregarded for this experiment (cf. Section 2.5). Each experiment runs for a duration of 30 seconds. We observe the additional runtime of the processes incurred by the network traffic.

Results

The total rate of interrupts per packet ranged from 70% in the undisturbed experiment (Experiment 1) to 2% with high additional load of 15000 packets per second and $3200 \mu\text{s}$ absolute timer value. The absolute moderation timer is an effective tool to moderate the high load as more packets are coalesced into interrupts while the critical task is unaffected.




Figure 15: Additional runtime of critical process (queue 0) induced by increased network load.

Next, we observe the interrupt-induced runtime increase of the critical process. A significant mitigation of the harmful effects of packet floods can be obtained in all moderation configurations for the critical process. Processes of lower priority also benefit from the approach, as the CPU

⁴MODBUS/TCP traces provided by [18].

is freed up from unnecessary ISRs. Figure 15 shows the mitigating effects for the critical task under variable additional load. The visible linear increase continues throughout all experiments. Further, it can be seen that there is a scenario-specific optimal configuration between $d2400$ and $d3200$ due to the effects of packet burst processing (cf. Section 7.3.1). By increasing the delay parameters, more packets are coalesced for each interrupt, meaning that the networking tasks are confronted with larger bursts of packets per notification. This has negative effects on CPU load starting at a critical packet count. For the maximum depicted packet load of 5000 packets per second the additional runtime could be decreased by 80% (or 12 percentage points) resulting from the prevention of 93% of interrupts.

8.6. Experiment 3: Expected Load

Since the NIC retains packet notifications for low priority tasks, it causes an additional interrupt delay depending on the delay timer configurations. To investigate this delay, the second set of experiments was performed on a stable system with no unexpected traffic floods. The combination of absolute and packet delay timers spans a window for the period length of interrupts. Using the approximated incoming packet rate of a flow $1/t_P(f)$, developers can adjust the values to minimize the introduced interrupt delay as described in Section 7.3.2.

In this experiment, the four processes receive IP flows ranging from about 50 (for queue 0) to 200 (for queue 3) packets per second. We compare three different NIC configurations:

- *No moderation*: All queues trigger interrupts as soon as a packet arrives.
- *Medium moderation*: Queue 0 triggers an interrupt as soon as a packet arrives. Queues 1, 2, and 3 are configured to coalesce 2, 3, and 4 packets per interrupt, on average.
- *Strict moderation*: Queue 0 triggers an interrupt as soon as a packet arrives. Queues 1, 2, and 3 are configured to coalesce 5, 6, and 7 packets per interrupt, on average.

Results

The results, depicted in Figure 16, show that the critical process (queue 0) does not suffer from any additional delay. The delay added to processes of less priority is directly dependent on the chosen moderation parameters and process-specific load and can be chosen to fit the requirements of each process before or during runtime. At the same time, the delay window prevents interrupt frequencies to climb to a critical level. An increase in the OS-induced delay can be observed when too many packets are coalesced to one interrupt as the resulting bursts in packet processing increase the processing time for individual packets in the networking and worker processes.




Figure 16: Median delay times of the four configurations. The IRQ delay specifies the average coalescing delays.

While the queue moderation can effectively reduce the impact of high packet loads, the tuning of moderation parameters requires care. Task deadlines, as well as packet loads and latencies for each process should be identified to reduce the impact of interrupt moderation under normal conditions while at the same time ensuring operability for critical tasks under unexpectedly high loads.

9. Towards a Unified Architecture for Real-Time Aware Packet Processing

In this section, we give an outlook on a combination of the two evaluated approaches, attending to challenges arising and necessary changes to be made to the individual designs. We furthermore discuss the compatibility of advantages gained through a unified hardware/software co-design for real-time packet reception.

9.1. Preliminary Considerations

To implement a unified design, the networking task, driver, and hardware components should be designed holistically and make available configuration parameters in a standardized manner. This design is subject to the same assumptions and requirements as the individual approaches. Most notably the placement of devices in a real-time IoT setting with no industrial grade real-time networking infrastructure. The preliminary considerations in Section 4 remain relevant here.

A unified design could combine the advantages of the two presented approaches while also mitigate some of the individual disadvantages. These shortcomings can be summarized as follows.

Hardware Approach.

- Only the amount of ISR runs is reduced in a prioritized manner. Once packets are received by the OS queue, no reordering can happen and packet processing can still lead to an inversion of priorities.
- The design creates bursts of low priority packets that, once inside the operating system, might block the processing of high priority packets and lead to priority inversion.




Figure 17: **Unified design:** Demultiplexing of packets and interrupt moderation in the NIC. Packet ordering, priority inheritance, and burst control in the network driver.

- Interrupt moderation and multiqueue parameters have to be set by developers apriori and for each process individually.

Software Approach.

- The number of interrupts can not be reduced other than by switching to polling mode.
- In the case of a flood of incoming unwanted packets the system can still be overwhelmed by IRQs and ISR runs or be forced into polling mode with the remaining packet processing overhead.
- The classification of packets and packet-wise cache invalidation adds workload to the (preempting) ISRs.

9.2. Overview

The combined design unifies the priority spaces of real-time tasks, packet processing, and network-generated interrupts. An abstract representation of a unified design can be seen in Figure 17. Depending on developer-defined tasks and their priorities incoming packets can be dropped and coalesced to fewer IRQs before any workload emerges in the RTOS. With the mapping of tasks to IP flows being performed in the NIC, the packet classification is shifted from the eager part of the driver to hardware. The moderated queues in the NIC remain configurable over the socket API and write their contents to the BD ring where the remaining part of the eager driver fetches them. The descriptors can then be placed in process specific packet queues enforcing the rate limitation and priority inheritance introduced by the software approach.

9.2.1. Changes to Individual Implementations

To accommodate the combined design the individual implementations have to be adapted. This section discusses the challenges of a new design and presents necessary changes.

Hardware Approach. Since the classification of packets should not have to be repeated in the driver, the packet descriptors written to the RTOS need to be extended to contain the priority of the expecting process. By furthermore adding a field for the IP-flow ID (i.e. the port number) flows of equal priority can be prevented from blocking each other in case one experiences a packet flood. This way, packet descriptors fetched from the ring buffer already contain priorities and flow IDs and only need to be enqueued accordingly by the ISR.

A more complex yet optional change to the original design could include the dynamization of queue configuration. As a means of reactive rate limitation the driver could be enabled to dynamically change NIC queue sizes during periods of high packet volume. This way, packets could be dropped in a priority aware manner before reaching the driver when the rate limitation is expected to lead to the same. The initial configuration via the socket API can remain unchanged.

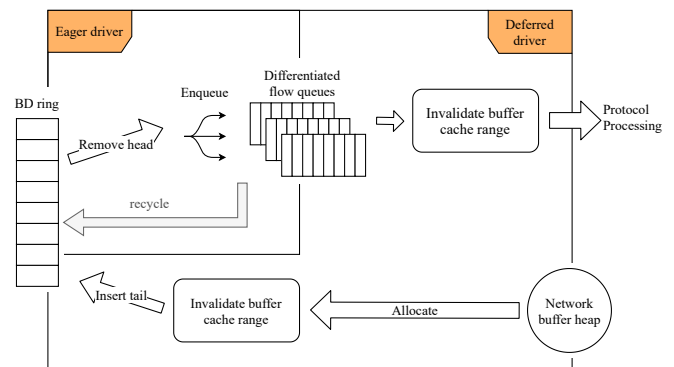


Figure 18: **Receive activity in adapted driver:** Classification, as well as cache invalidation of the header line is removed.

Software Approach. With the classification being shifted to hardware, the software side can save packet processing time at the eager driver in two manners: The ISR does not have to read the packet headers from memory in order to place the descriptors in the appropriate queue as the descriptors now contain priority and flow ID fields. With header inspection in the ISR being unnecessary, the time

consuming task of packet-wise cache line invalidation can also be eliminated from the eager driver, resulting in the receive workflow in Figure 18.

The due to the interrupt moderation bursty nature of packet reception through ISRs is not impeded by the combination of the two approaches. Packets, especially those of low priorities come in bursts, possibly resulting in immediate rate limitation by the driver. To prevent this from happening for each burst, queue sizes need to be adapted to each other. Furthermore, some of the unschedulable ISR runs will require more time, as bursts of packets need to be handled. However, this is somewhat mitigated as the comparison and raise of the network task priority only has to happen once at the end of the ISR and not for every packet as in the original design.

10. Conclusion

Unexpected floods of network traffic can delay process flows in real-time systems, putting critical real-time requirements at risk. This paper presented two individually evaluated approaches – one software- and one hardware-based – to mitigate the real-time violating effects of IP packet reception in constrained IoT devices.

The software approach is soft IP stack design to individually schedule packet processing for differently prioritized IP-flows after early demultiplexing. The issue of costly processing in the network driver is approached by integrating the possibility of deferred buffer processing into our architecture. On our test system, even when having to deal with packet buffers travelling CPU-caches, the CPU load caused by LP packets in an already occupied system is reduced by 86%, leading to a 7-fold speedup of concurrently running processes. Through limitation parameters, the software approach allows system designers to anticipate packet rates of certain soft real-time flows and derive an estimate for the respective processing WCET. Budgeting the same CPU resources to the processing of incoming packets, the networking subsystem can still process packets of a HP-flow for up to 600% higher overall traffic loads.

The hardware approach presented a multiqueue NIC design shifting the early demultiplexing to hardware. This way, the number of network-triggered IRQs can be moderated depending on packet priorities. By configuring the extended NIC via an adapted Berkeley socket API, currently active processes waiting for packets can register their priorities with the hardware, effectively mapping IP flows to processes when a socket is bound and changing queue parameters accordingly. We evaluated the hardware design using a NIC simulation and an IoT device running a real-time operating system. The results of these experiments show that our approach significantly reduces the impact of traffic floods on critical process runtimes by saving 93 % of interrupts and 80 % of processing delay under packet rates of 5000 per second while the configuration of multiqueue parameters requires knowledge about expected network traffic and real-time requirements.

The combination of the two approaches requires a hardware/software co-design of real-time NIC, network driver and IP stack implementation, further addressing the remaining weaknesses of the individual designs.

Acknowledgments

We sincerely thank the reviewers for their helpful comments which significantly improved this work. This research was supported by the German Academic Exchange Service (DAAD) as ide3a.

References

- [1] V. Alcácer and V. Cruz-Machado. Scanning the industry 4.0: A literature review on technologies for manufacturing systems. *Engineering science and technology, an international journal*, 22(3), 2019.
- [2] J. E. Amiri and M. Kargahi. A predictable interrupt management policy for real-time operating systems. In *CSI Symposium on Real-Time and Embedded Systems and Technologies*. IEEE, 2015.
- [3] S. Bansal and D. Kumar. Iot ecosystem: A survey on devices, gateways, operating systems, middleware and communication. *International Journal of Wireless Information Networks*, 2020.
- [4] S. Becker, T. Pfandzelter, N. Japke, D. Bernbach, and O. Kao. Network emulation in large-scale virtual edge testbeds: A note of caution and the way forward. In *International Conference on Cloud Engineering (IC2E)*. IEEE, 2022.
- [5] I. Behnke, L. Pirl, L. Thamsen, R. Danicki, A. Polze, and O. Kao. Interrupting real-time iot tasks: How bad can it be to connect your critical embedded system to the internet? In *39th International Performance Computing and Communications Conference*. IEEE, 2020.
- [6] I. Behnke, P. Wiesner, R. Danicki, and L. Thamsen. A priority-aware multiqueue nic design for real-time iot devices. In *Proceedings of the 35th Annual ACM Symposium on Applied Computing*. ACM, 2022.
- [7] F. Bender, J. J. Brune, N. L. Keutel, I. Behnke, and L. Thamsen. Pieres: A playground for network interrupt experiments on real-time embedded systems in the iot. In *Companion of the ACM/SPEC International Conference on Performance Engineering*. ACM, 2021.
- [8] C. Blumschein, I. Behnke, L. Thamsen, and O. Kao. Differentiating network flows for priority-aware scheduling of incoming packets in real-time iot systems. In *25th International Symposium on Real-Time Distributed Computing (ISORC)*. IEEE, 2022.
- [9] P. Brooks. Ethernet/ip-industrial protocol. In *8th International Conference on Emerging Technologies and Factory Automation. Proceedings*, volume 2. IEEE, 2001.
- [10] D. Bruckner, M.-P. Stanica, R. Blair, S. Schriegel, S. Kehrer, M. Seewald, and T. Sauter. An Introduction to OPC UA TSN for Industrial Communication Systems. *Proceedings of the IEEE*, 107(6), 2019.
- [11] J. Cheng, W. Chen, F. Tao, and C.-L. Lin. Industrial iot in 5g environment towards smart manufacturing. *Journal of Industrial Information Integration*, 10, 2018.
- [12] R. Danicki, M. Haug, I. Behnke, L. Mådje, and L. Thamsen. Detecting and mitigating network packet overloads on real-time devices in IoT systems. In *Proceedings of the 4th International Workshop on Edge Systems, Analytics and Networking*. ACM, 2021.
- [13] P. Druschel and G. Banga. Lazy receiver processing (lrp): a network subsystem architecture for server systems. *ACM SIGOPS Operating Systems Review*, 30(SI), 1996.
- [14] A. Dunkels. Design and implementation of the lwIP TCP/IP stack. *Swedish Institute of Computer Science*, 2(77), 2001.

- [15] N. Finn. Introduction to time-sensitive networking. *IEEE Communications Standards Magazine*, 2(2), 2018.
- [16] L. Foschini, V. Mignardi, R. Montanari, and D. Scotece. An SDN-Enabled Architecture for IT/OT Converged Networks: A Proposal and Qualitative Analysis under DDoS Attacks. *Future Internet*, 13(10), 2021.
- [17] F. Foukalas, P. Pop, F. Theoleyre, C. A. Boano, and C. Buratti. Dependable wireless industrial iot networks: Recent advances and open challenges. In *European Test Symposium*. IEEE, 2019.
- [18] I. Frazão, P. H. Abreu, T. Cruz, H. Araújo, and P. Simões. Denial of service attacks: Detecting the frailties of machine learning algorithms in the classification process. In *International Conference on Critical Information Infrastructures Security*. Springer, 2018.
- [19] S. Gebert, T. Zinner, S. Lange, C. Schwartz, and P. Tran-Gia. Performance modeling of software-defined network functions using discrete-time analysis. In *2016 28th International Teletraffic Congress*, volume 1. IEEE, 2016.
- [20] Y. Gilad and A. Herzberg. Fragmentation considered vulnerable: blindly intercepting and discarding fragments. In *Proceedings of the 5th Conference on Offensive Technologies*. USENIX, 2011.
- [21] T. Gomes, P. Garcia, F. Salgado, J. Monteiro, M. Ekpanyapong, and A. Tavares. Task-aware interrupt controller: Priority space unification in real-time systems. *IEEE Embedded Systems Letters*, 7(1), 2015.
- [22] C. Haar and E. Buchmann. Fane: a firewall appliance for the smart home. In *Federated Conference on Computer Science and Information Systems*. IEEE, 2019.
- [23] D. Henneke, L. Wisniewski, and J. Jasperneite. Analysis of realizing a future industrial network by means of Software-Defined Networking (SDN). In *World Conference on Factory Communication Systems (WFCS)*. IEEE, 2016.
- [24] W. Hofer, D. Lohmann, F. Scheler, and W. Schröder-Preikschat. Sloth: Threads as interrupts. In *30th Real-Time Systems Symposium*. IEEE, 2009.
- [25] W. Hofer, D. Lohmann, and W. Schröder-Preikschat. Sleepy sloth: Threads as interrupts as threads. In *32nd Real-Time Systems Symposium*. IEEE, 2011.
- [26] M. Honda, F. Huici, C. Raiciu, J. Araujo, and L. Rizzo. Rekindling network protocol innovation with user-level stacks. *ACM SIGCOMM Computer Communication Review*, 44(2), 2014.
- [27] B. Hubert et al. Linux advanced routing & traffic control howto. *Netherlabs BV*, 1, 2002.
- [28] D. Jansen and H. Buttner. Real-time ethernet: the ethercat solution. *Computing and Control Engineering*, 15(1), 2004.
- [29] T. Jeffree, P. Congdon, and M. J. Teener. Ieee standard for local and metropolitan area networks—audio video bridging (avb) systems. *IEEE Std 802.1BA-2021 (Revision of IEEE Std 802.1BA-2011)*, 2021.
- [30] T. Jeffree, P. Congdon, M. J. Teener, and C. Gunther. Ieee standard for local and metropolitan area networks—virtual bridged local area networks amendment 14: Stream reservation protocol (srp). *IEEE Std 802.1Qat-2010 (Revision of IEEE Std 802.1Q-2005)*, 2010.
- [31] C. A. Kent and J. C. Mogul. Fragmentation considered harmful. *SIGCOMM Comput. Commun. Rev.*, 17(5), 1987.
- [32] H. Laaki, Y. Miche, and K. Tammi. Prototyping a digital twin for real time remote control over mobile networks: Application of remote surgery. *Ieee Access*, 7, 2019.
- [33] C. K. Lee. Development of an industrial internet of things (iiot) based smart robotic warehouse management system. In *International Conference on Information Resources Management*. Association For Information Systems, 2018.
- [34] M. Lee, H. Kim, and I. Shin. Priority-based network interrupt scheduling for predictable real-time support. *Journal of Computing Science and Engineering*, 9(2), 2015.
- [35] M. Lee, J. Lee, A. Shyshkalov, J. Seo, I. Hong, and I. Shin. On interrupt scheduling based on process priority for predictable real-time behavior. *ACM SIGBED Review*, 7, 2010.
- [36] L. Leonardi, F. Battaglia, and L. L. Bello. Rt-lora: A medium access strategy to support real-time flows over lora-based networks for industrial iot applications. *IEEE Internet of Things Journal*, 6(6), 2019.
- [37] L. E. Leyva-del Foyo, P. Mejia-Alvarez, and D. de Niz. Integrated task and interrupt management for real-time systems. *ACM Transactions on Embedded Computing Systems*, 11(2), 2012.
- [38] Y. Li, J. Jiang, and S. H. Hong. Joint traffic routing and scheduling algorithm eliminating the nondeterministic interruption for tsn networks used in iiot. *IEEE Internet of Things Journal*, 2022.
- [39] C. Lin, Y. Bi, H. Zhao, Z. Liu, S. Jia, and J. Zhu. DTE-SDN: A Dynamic Traffic Engineering Engine for Delay-Sensitive Transfer. *IEEE Internet of Things Journal*, 5(6), 2018.
- [40] W. Liu, G. Nair, Y. Li, D. Nestic, B. Vucetic, and H. V. Poor. On the latency, rate, and reliability tradeoff in wireless networked control systems for iiot. *IEEE Internet of Things Journal*, 8(2), 2020.
- [41] A. Lonardo, F. Lo Cicero, A. Cotta Ramusino, I. Neri, F. Simeone, M. Sozzi, A. Biagioni, O. Frezza, P. S. Paolucci, R. Ammendola, et al. A fpga-based network interface card with gpudirect enabling realtime gpu computing in hep experiments. 2015.
- [42] S. Mäkinen, R. Iyer, P. Sarangam, D. Newell, L. Zhao, R. Illickal, and J. Moses. Receive side coalescing for accelerating TCP/IP processing. In *International Conference on High-Performance Computing*. Springer, 2006.
- [43] A. M. Mandalari, D. J. Dubois, R. Kolcun, M. T. Paracha, H. Haddadi, and D. Choffnes. Blocking without breaking: Identification and mitigation of non-essential iot traffic. In *Proceedings on Privacy Enhancing Technologies*, 2021.
- [44] C. W. Mercer and H. Tokuda. An evaluation of priority consistency in protocol architectures. In *16th Conference on Local Computer Networks*. IEEE, 1991.
- [45] M. Musleh, V. Pai, J. P. Walters, A. Younge, and S. Crago. Bridging the virtualization performance gap for HPC using SR-I/OV for InfiniBand. In *7th International Conference on Cloud Computing*. IEEE, 2014.
- [46] R. Müller, D. Danner, W. S. Preikschat, and D. Lohmann. Multi sloth: An efficient multi-core rtos using hardware-based scheduling. In *26th Euromicro Conference on Real-Time Systems*. IEEE, 2014.
- [47] L. D. Nguyen and A. Kortun. Real-time Optimisation for Industrial Internet of Things (IIoT): Overview, Challenges and Opportunities. *EAI Endorsed Transactions on Industrial Networks and Intelligent Systems*, 7(25), 2020.
- [48] M. Niedermaier, J.-O. Malchow, F. Fischer, D. Marzin, D. Merli, V. Roth, and A. Von Bodisco. You snooze, you lose: measuring PLC cycle times under attacks. In *12th Workshop on Offensive Technologies*. USENIX, 2018.
- [49] M. Niedermaier, J.-O. Malchow, F. Fischer, D. Marzin, D. Merli, V. Roth, and A. von Bodisco. You snooze, you lose: Measuring PLC cycle times under attacks. In *12th Workshop on Offensive Technologies*. USENIX, 2018.
- [50] M. Niedermaier, D. Merli, and G. Sigl. A secure dual-MCU architecture for robust communication of IIoT devices. In *8th Mediterranean Conference on Embedded Computing*. IEEE, 2019.
- [51] G. Parsons, J. Messenger, J. Farkas, C. Gunther, G. Garner, and S. Kehrer. Ieee standard for local and metropolitan area networks—timing and synchronization for time-sensitive applications. *IEEE Std 802.1AS-2020*, 2020.
- [52] G. Parsons, J. Messenger, M. J. Teener, and J. Farkas. Ieee standard for local and metropolitan area networks – bridges and bridged networks - amendment 25: Enhancements for scheduled traffic. *IEEE Std 802.1Qbv-2015*, 2016.
- [53] T. Qiu, J. Chi, X. Zhou, Z. Ning, M. Atiquzzaman, and D. O. Wu. Edge computing in industrial internet of things: Architecture, advances and challenges. *IEEE Communications Surveys Tutorials*, 22(4), 2020.
- [54] M. Schoeberl and R. U. Pedersen. tpIP: A Time-Predictable TCP/IP Stack for Cyber-Physical Systems. In *21st Interna-*

- tional Symposium on Real-Time Distributed Computing*. IEEE, 2018.
- [55] S. Schriegel and J. Jasperneite. A Migration Strategy for Profinet Toward Ethernet TSN-Based Field-Level Communication: An Approach to Accelerate the Adoption of Converged IT/OT Communication. *IEEE Industrial Electronics Magazine*, 15(4), 2021.
- [56] P. Shinde, A. Kaufmann, K. Kourtis, and T. Roscoe. Modeling nics with unicorn. In *Proceedings of the Seventh Workshop on Programming Languages and Operating Systems*, 2013.
- [57] P. Shinde, A. Kaufmann, T. Roscoe, and S. Kaestle. We need to talk about nics. In *14th Workshop on Hot Topics in Operating Systems*. ACM, 2013.
- [58] L. Silva, P. Pedreiras, P. Fonseca, and L. Almeida. On the adequacy of SDN and TSN for Industry 4.0. In *22nd International Symposium on Real-Time Distributed Computing (ISORC)*. IEEE, 2019.
- [59] S. Spanos, A. Meliones, and G. Stassinopoulos. The internals of advanced interrupt handling techniques: Performance optimization of an embedded linux network interface. *Computer Communications*, 31(14), 2008.
- [60] B. Sprunt, L. Sha, and J. Lehoczky. Aperiodic task scheduling for hard-real-time systems. *Real-Time Systems*, 1(1), 1989.
- [61] B. Stephens, A. Akella, and M. Swift. Loom: Flexible and efficient NIC packet scheduling. In *16th Symposium on Networked Systems Design and Implementation (NSDI)*. USENIX, 2019.
- [62] J. K. Strosnider, J. P. Lehoczky, and L. Sha. The deferrable server algorithm for enhanced aperiodic responsiveness in hard real-time environments. *IEEE Transactions on Computers*, 44(1), 1995.
- [63] S. Sudhakaran, V. Mageshkumar, A. Baxi, and D. Cavalcanti. Enabling QoS for Collaborative Robotics Applications with Wireless TSN. In *International Conference on Communications Workshops (ICC Workshops)*. IEEE, 2021.
- [64] H. Tahaei, F. Afifi, A. Asemi, F. Zaki, and N. B. Anuar. The rise of traffic classification in iot networks: A survey. *Journal of Network and Computer Applications*, 154, 2020.
- [65] K. Tindell and J. Clark. Holistic schedulability analysis for distributed hard real-time systems. *Microprocessing and microprogramming*, 40(2-3), 1994.
- [66] S. Tripathi, N. Droux, T. Srinivasan, K. Belgaied, and V. Iyer. Crossbow: A vertically integrated qos stack. In *Proceedings of the 1st Workshop on Research on Enterprise Networking*. ACM, 2009.
- [67] S. Vitturi, C. Zunino, and T. Sauter. Industrial communication systems and their future challenges: Next-generation ethernet, iiot, and 5g. *Proceedings of the IEEE*, 107(6), 2019.
- [68] M. Wollschlaeger, T. Sauter, and J. Jasperneite. The Future of Industrial Communication: Automation Networks in the Era of the Internet of Things and Industry 4.0. *IEEE Industrial Electronics Magazine*, 11(1), 2017.
- [69] H. Zhu. NIC-based parallelism. In *Data Plane Development Kit (DPDK): A Software Optimization Guide to the User Space-Based Network Applications*. CRC Press, 2020.