

A First Look at ROS 2 Applications Written in Asynchronous Rust

Martin Škoudlil ✉ 

Czech Institute of Informatics, Robotics and Cybernetics,
Czech Technical University in Prague, Czech Republic

Michal Sojka¹ ✉ 

Czech Institute of Informatics, Robotics and Cybernetics,
Czech Technical University in Prague, Czech Republic

Zdeněk Hanzálek ✉ 

Czech Institute of Informatics, Robotics and Cybernetics,
Czech Technical University in Prague, Czech Republic

Abstract

The increasing popularity of the Rust programming language in building robotic applications using the Robot Operating System (ROS 2) raises questions about its real-time execution capabilities, particularly when employing asynchronous programming. Existing real-time scheduling and response-time analysis techniques for ROS 2 focus on applications written in C++ and do not address the unique execution models and challenges presented by Rust's asynchronous programming paradigm. In this paper, we analyze the execution model of R2R – an asynchronous Rust ROS 2 bindings and various asynchronous Rust runtimes, comparing them with the execution model of C++ ROS 2 applications. We propose a structured approach for R2R applications aimed at deterministic real-time operation involving thread prioritization and callback-to-thread mapping schemes. Our experimental evaluation based on measuring end-to-end latencies of a synthetic application shows that the proposed approach is effective and outperforms other evaluated configurations. A more complex autonomous driving case study demonstrates its practical applicability. Overall, the experimental results indicate that our proposed structure achieves bounded response times for time-critical tasks. This paves the way for future work to adapt existing or develop new response-time analysis techniques for R2R applications using our structure.

2012 ACM Subject Classification Software and its engineering → Real-time systems software

Keywords and phrases ROS, Rust, Real-time, Response time

Digital Object Identifier 10.4230/LIPIcs.ECRTS.2025.1

Supplementary Material *Software (Source code)*: <https://zenodo.org/records/15446666>

Software (ECRTS 2025 Artifact Evaluation approved artifact):

<https://doi.org/10.4230/DARTS.11.1.2>

Funding This work was supported by the Technology Agency of the Czech Republic under the project Certicar CK03000033.

Zdeněk Hanzálek: This work was co-funded by the European Union under the project ROBOPROX (reg. no. CZ.02.01.01/00/22_008/0004590).

Acknowledgements We want to thank our reviewers and a shepherd for valuable comments that lead to improvement of this paper.

¹ corresponding author



1 Introduction

The Robot Operating System (ROS) [15] is an increasingly popular framework for developing robotic applications. Its second major version, denoted as ROS 2, was designed to fill the needs of industrial use cases, including support for real-time execution. ROS 2 supports writing applications in different programming languages by providing so-called client libraries. Out of the box, ROS 2 provides client libraries for C++ and Python; other languages, such as Rust, have various levels of community support.

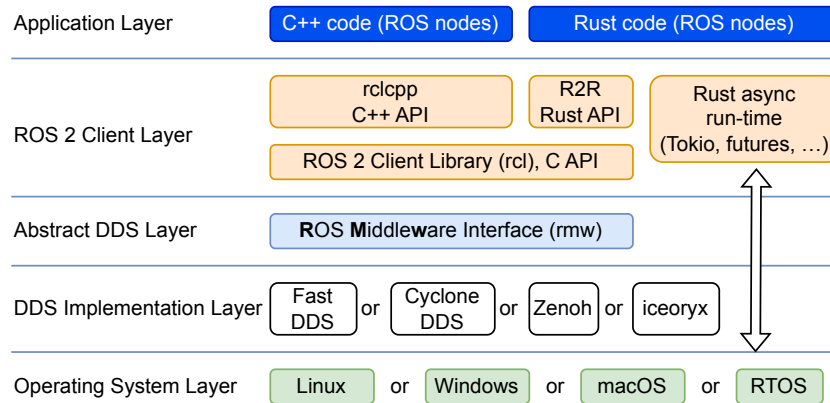
Despite its detailed design, early versions of ROS 2 did not deliver optimal timing predictability [3]. The discovered problems were since fixed and researchers started developing response time analysis techniques for various configurations of ROS 2 applications [22, 2, 10, 23]. All these techniques assume the use of C++ language, which is officially supported by the framework and is considered mature and suitable for real-time applications.

However, even C++ support in ROS is not without problems. Teper et al. [24] discovered that ROS multi-threaded executor is not starvation-free, making it unsuitable for analysis with existing techniques. Another problem often associated with C++ is its complexity. Writing reliable C++ applications becomes difficult even for professionals and even more so for less experienced users. That is one of the reasons why the Rust programming language gains in popularity. Its compiler can detect many types of common C++ errors, such as race conditions, at compile time. ROS supports Rust via several community-provided libraries. Among the most popular libraries are `rclrs` from the `ros2_rust` project [16] and R2R [5]. The former provides an API similar to its C++ counterpart `rclcpp`, and the latter offers so-called asynchronous (async in short) API, which allows multiplexing execution of concurrent tasks in a single operating system (OS) thread. While `rclrs` received some contributions from core ROS developers, R2R was developed independently as a part of the Sequence Planner framework [6].

Given the potential of the Rust language and its increasing popularity, it is essential to understand how it can be used to develop real-time robotic applications. In this paper, we look at how the async R2R library schedules and executes the code of ROS 2 applications and compare that with the approaches used in the C++ ROS client library. Since the async Rust applications offer high flexibility in how the application is executed, we propose a particular structure suitable for real-time applications that utilize a thread prioritization and callback-to-thread mapping scheme. We evaluate this structure by measuring end-to-end latencies in a synthetic application as well as in a more complex autonomous driving case study. With the synthetic application, we empirically compare different application structures utilizing different prioritization schemes and different async Rust runtime libraries and compare them with the C++ language. Our proposed structure achieves bounded response times for time-critical tasks, which is more suitable for real-time applications than the structures appearing in R2R documentation and example code. This opens the way for future work to either adapt existing response-time analysis techniques or design new ones to target R2R applications using our structure.

The specific contributions of the paper are:

1. We analyze the execution model of the async Rust R2R library and of several async Rust runtimes and compare them with the execution model of C++ ROS applications.
2. We propose how to structure R2R applications to be suitable for deterministic real-time operation.
3. We demonstrate that the response times of the synthetic application match the theoretical results of a uni-processor response-time analysis. Moreover, a more complex autonomous driving case study shows that deterministic timing is maintained even in processing chains involving more than two nodes and running concurrently with other chains.



■ **Figure 1** ROS 2 architecture for C++ and Rust R2R applications.

2 Background

This section describes the basic concepts of ROS 2, followed by an introduction to async programming in Rust and a description of scheduling execution in Rust async runtimes. Then, we compare features supported by the official C++ client library `rclcpp` and two Rust client libraries `R2R` and `rclrs`. Finally, we describe the execution model of the `R2R` library and compare it with the model of C++ ROS executors.

2.1 ROS 2

ROS 2 applications are composed of nodes, which can communicate with each other in a publish-subscribe manner. *ROS node* is an organizational unit for other *entities* such as timers, publishers and subscriptions. In this paper, we assume the traditional model with one ROS node per OS process, but ROS 2 also supports another model where multiple so-called composable nodes to run in a single process. A *publisher* allows sending messages to the associated *topic* and all *subscriptions* to the same topic then receive the message. Reception of the message or expiration of a timer results in invocation of an associated application *callback*. When the callbacks are executed depends on the *executor* associated with them. The C++ client library `rclcpp` provides a single-threaded executor, a multi-threaded executor, and an experimental events executor. Besides publish-subscribe communication, ROS 2 applications can also communicate via services and actions, which are internally built on top of publishers and subscribers. We do not explicitly address them in this paper, however their callback scheduling and time related aspects should not be much different from publishers and subscribers.

Internally, ROS 2 implementation uses a layered architecture depicted in Figure 1. The figure also shows the main differences between C++ and Rust R2R nodes, which will be detailed below.

The communication between ROS nodes is handled by the ROS middleware (RMW) library, which supports different implementations of the actual communication services. Currently, ROS 2 provides several implementations of the Data Distribution Service (DDS) standard [14], with Fast DDS [9] being the default, and the latest ROS 2 release Jazzy Jalisco [15] adds experimental support for the Zenoh protocol [8].

ROS client libraries can receive information about events occurring on entities, e.g., timer expirations or receptions of messages on subscriptions, via the *wait set* data structure, which allows the client thread to wait for multiple entities simultaneously. After waiting, the wait

set reports which entities in the set are *ready*, i.e., one or more events occurred on them. The ready entities are reported in the same order in which the entities were added to the set. The process of obtaining ready entities is called *sampling*.

After a subscription is ready, its received message can be obtained from RMW with an operation called *take*, which is implemented by the `rcl_take` function.

2.2 Rust & asynchronous programming

Similarly to C++, Rust is a compiled programming language offering higher levels of abstraction than the C language but still giving the programmer detailed control over the usage of resources such as execution time and memory. An important difference from C++ is that the Rust compiler can guarantee memory safety, meaning that many classes of memory-related errors (race conditions, use-after-free) are detected and prevented at compile time.

The Rust language has support for asynchronous (*async* in short) programming, which is a form of concurrency where scheduling decisions are made in the application rather than by the OS, as is the case with thread-based concurrency. To support async programming, the Rust language defines two keywords `async` and `await`, and the `Future` trait², but leaves the implementation of the trait and related schedulers to independent libraries (crates in Rust terminology) called async runtimes. Popular async runtimes are Tokio³ and FUTURES⁴ (in small caps to avoid confusion with futures objects).

Async runtimes work with objects implementing the `Future` trait called *futures*. They represent a future execution of some code optionally producing a value. Applications can create futures in several ways: by calling a function annotated with the `async` keyword, by defining an `async` block, or by implementing the `Future` trait manually. Futures can be turned into *async tasks* by registering them with an executor, typically by calling its `spawn` method. *Executors*, which are usually provided by async runtimes and are conceptually similar to ROS C++ executors, then schedule and execute async tasks in one or more OS threads. Note that the overhead of async tasks is much lower than the overhead of OS threads. Async runtimes are known to easily handle millions of tasks. Futures can also be executed in the currently running task, without spawning a new one, by using the `await` keyword on them.

Async tasks can communicate by sending messages via async MPSC (Multi-Producer Single-Consumer) channels. Such a channel is typically implemented by a concurrent FIFO queue that allows asynchronous waiting on the receiving side until there is a message to dequeue. Technically, waiting is performed by calling a method on the receiving side that returns a future. Note that MPSC channel implementation and API can differ in different runtimes, but as long as they implement the `Future` trait, they are often compatible with other runtimes.

2.2.1 Execution model of futures runtime

We start our description of the Rust asynchronous execution model by describing the FUTURES runtime. Its implementation is simple, which allows building deterministic real-time applications on top of it. It allows mapping of async tasks to OS threads, but does not control scheduling parameters of the threads in any way, allowing the application to set them as appropriate.

² Traits in Rust are a concept similar to interfaces or abstract classes in other languages.

³ <https://tokio.rs/>

⁴ <https://crates.io/crates/futures>

2.2.1.1 Futures local executor

This section describes the current (v0.3.31) behavior of the `LocalPool` (local in short) executor within the `FUTURES` crate. The local executor is single-threaded and incorporates three primary data structures: 1) an incoming tasks vector, 2) a linked list of active tasks, and 3) a ready queue implemented by a concurrent linked list. When a future is spawned into the executor, a new async task is allocated in heap memory and added to the incoming vector. The executor repeatedly checks the incoming vector for new tasks and moves them to the list of active tasks, which incurs additional memory allocation for each task. Subsequently, the task is enqueued into the ready queue to execute the future or set up waiting for it. The executor ensures that each task is present in the ready queue only once at any given time. To execute a ready task, the executor removes the task from the ready queue and then invokes the `Future::poll` method on it. Whenever an active task (executing or waiting) becomes ready, it is always enqueued at the end of the ready queue.

2.2.1.2 Futures thread-pool executor

The thread-pool executor executes tasks in a pool of multiple worker threads. It uses only the ready queue data structure, which is implemented using the standard library's unbounded synchronous MPSC channel, i.e., concurrent linked list. When a waiting task is woken, it is added to the end of the ready queue. This operation can involve memory allocation. Threads in the thread pool dequeue ready tasks in FIFO order in mutually exclusive way and run them. If the queue is empty, the thread waits. The process of adding a running task back to the ready queue (without waiting) differs from the local executor in that the task is not added to the end of the ready queue but continues executing. This might result in starvation of other tasks. For example, if in a thread pool with N worker threads, there are N tasks that always become ready during their execution, other tasks will not be executed at all.

2.2.1.3 Grouping futures execution

The `futures::join!()` macro can be used to group multiple futures and wait for the completion of the group as a whole. The effect of joining the futures is that their code will not be executed in parallel. From the point of view of the executor, the group is treated as one task; if one future gets ready, the entire group becomes ready. When the group starts executing, it will poll all ready tasks in the group. The effect of this grouping in the executor is shown later in an example with ROS in Figure 4.

2.2.2 Tokio.rs runtime

Tokio is a popular multi-threaded runtime. It seems to be designed to maximize throughput rather than time determinism. Their scheduling policies are complex and difficult to understand due to the use of abstraction layers. The scheduling policies are partially described in the documentation [25] and can be outlined as follows: Each worker thread has a LIFO slot, which is used for dequeuing ready tasks at most three times in a row. We believe, this improves cache locality while also avoiding starvation. Then, tasks are dequeued from the local ready queue, which can hold up to 256 tasks. If the worker cannot dequeue tasks from the local or global (shared) ready queues (they are empty), it steals half of the tasks in the other worker's local queue. The victim of the theft is chosen as the first worker, with a non-empty queue when iterating workers with a random starting position. Such a behavior is clearly unsuitable for real-time applications. We will evaluate Tokio experimentally in Section 4.

2.3 ROS 2 C++ executors

Execution of callbacks in C++ ROS 2 applications is handled by ROS executors. Currently, ROS 2 includes a single-threaded executor, a multi-threaded executor, and an experimental events executor [21]. These are briefly described below.

Single-threaded executor first samples the associated entities and then executes callbacks of those that are ready. Callbacks are executed in the same thread and are ordered based on their type. Timers are first, followed by subscriptions, services, and clients. Callbacks of the same type were executed based on the order of their registration [2], but this has changed in ROS 2 Jazzy, where the order is no longer predictable⁵.

Multi-threaded executor executes callbacks in multiple threads. Callbacks are organized in callback groups, which can be of two types: *Mutually exclusive* or *Reentrant*. The executor threads access the wait set in a mutually exclusive manner; a thread that gets the access waits for the events, and once some entities are ready, their callbacks get executed by one or more threads subject to the policy of their callback group. Multi-threaded executor was found not to be starvation-free [24], which is problematic for real-time applications and their analysis.

Events executor is a recently added experimental executor that does not use wait sets but pushes events to the executor’s event queue directly from DDS callbacks. The main executor thread then dequeues the events and executes the associated callbacks in a loop. Recently, a multi-threaded version of the Events executor was proposed [13].

2.4 Feature comparison of rclcpp, R2R and rclrs

Before looking in detail at the R2R client library and its execution model, we provide a high-level comparison of features implemented in R2R and the other ROS Rust client library **rclrs**. As both are community-supported, they both lack some features implemented in the C++ client library **rclcpp**. The comparison of all three libraries is summarized in Table 1 and commented in more detail below.

R2R and **rclcpp** implement all communication styles implemented by ROS, but **rclrs** has only limited support for actions. Only the action message types are available for **rclrs** applications. If the application wants to use the actions implemented in another node, it has to implement all action logic and state machines itself.

All three libraries support a specific type of communication, built on top of services, that allows to work with node parameters. R2R does not yet fully support parameter ranges, which can be used to announce the permissible range of parameter values. Ranges are used by some GUI tools like **rqt** to provide “sliders” for changing the parameter values. Another difference in parameter handling is in how different libraries implement parameter locking to prevent concurrent accesses from the middleware and the application. **rclcpp** leaves locking up to the application, R2R uses a single lock per node, whereas **rclrs** has one lock for each parameter, potentially causing higher overhead in nodes with high number of parameters. R2R has a feature not available in other libraries, which simplifies working with parameters by using Rust’s **derive** macro to automatically generate parameter handling code for fields of an arbitrary structure.

With respect to time handling, a drawback of **rclrs** is the unavailability of timers. Time-based execution of user code has to be implemented by using standard Rust means, which prevents the correct function of such nodes with ROS simulated time. However, note

⁵ <https://github.com/ros2/rclcpp/issues/2532>

that simulated time is supported, but only for clocks and not for timers. Another time-related feature is support for tracing, which is being submitted to R2R as a result of this work [19]. Similar functionality is missing in `rclrs`.

The supported executors in R2R are detailed in the next section. Here, we just mention that `rclrs` supports only a single-threaded executor. Execution in multiple threads can be implemented by using multiple single-threaded executors or by sending work from callbacks to other threads via standard Rust means.

Neither Rust client library supports composable and lifecycle nodes. While the latter could be implemented with little effort, the former would require a deeper investigation of ABI compatibility between Rust and C++.

■ **Table 1** Feature comparison between `rclcpp`, R2R and `rclrs`.

Feature	<code>rclcpp</code>	R2R	<code>rclrs</code>
Communication			
Message generation	✓	✓	✓
Publishers and subscriptions	✓	✓	✓
Loaned messages (zero-copy)	✓	✓	✓
Tunable QoS settings	✓	✓	✓
Clients and services	✓	✓	✓
Actions	✓	✓	▲ ¹
Dynamic type support	✓	✓	✓
Parameters			
Parameter handling	✓	✓	✓
Parameter ranges	✓	▲	✓
Parameter locking	none	per-node	per-parameter
Derived parameters ²	✗	✓	✗
Time			
Timers	✓	✓	✗
Simulated time	✓	✓	✓
Tracepoints	✓	✓ ³	✗
Executors			
Single-threaded	✓	✓ ⁴	✓
Multi-threaded	✓	✓ ⁴	✗
Asynchronous programming style	✗	✓	✗
Other			
Composable nodes	✓	✗	✗
Lifecycle nodes	✓	✗	✗

Legend:

✓ Supported

✓ Supported with comment

▲ Partially supported

✗ Not supported

¹ See https://github.com/ros2-rust/ros2_rust/issues/244, https://github.com/ros2-rust/ros2_rust/pull/423, https://github.com/ros2-rust/ros2_rust/pull/410.

² Automatic generation of parameter handling code for fields of a structure.

³ Implemented in pull request <https://github.com/sequenceplanner/r2r/pull/117>, likely to be merged soon.

⁴ Executors are not a part of R2R, but are provided by asynchronous runtime libraries like `FUTURES` or `Tokio`. Hence, the exact types of supported executors depend on the selected library.

■ Listing 1 Registration of a subscription callback in R2R.

```
let subscription_future = subscription.for_each(|msg| async move {
    // Callback code for msg processing here
});
executor.spawn(subscription_future);
```

■ Listing 2 R2R setup equivalent to one single-threaded executor in C++.

```
1 local_executor.run_until_stalled();           // initialization
2 loop {
3     node.spin_once(Duration::seconds(1)); // sampling
4     local_executor.run_until_stalled();     // execution
5 }
```

2.5 R2R execution model

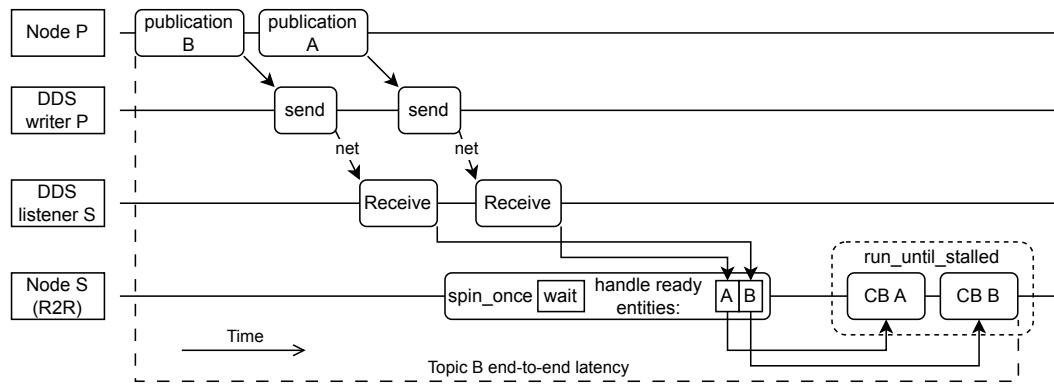
R2R [5] is one of ROS 2 client libraries for Rust. Its execution model differs from C++ because callback execution is managed by a Rust async runtime rather than by R2R itself as shown in Figure 1. R2R is only responsible for “sampling” the events from ROS entities and pushing them to the async runtime via MPSC channels.

To sample the events, R2R uses wait sets as the official C++ executors. Sampling is performed by function `Node::spin_once`. It creates a wait set with all entities in the ROS node (subscriptions, timers, clients, and services), and then it waits on it until one or more entities are ready. Until this point, the behavior is the same as in C++. Then, the ready entities are iterated over, received messages are taken (by calling `rcl_take`) from the RMW, and all events (messages and timer expirations) are pushed to async MPSC channels associated with their entity. R2R uses bounded asynchronous channels from the `FUTURES` crate, which are mapped 1:1 to entities. In the current implementation, the channels have a fixed capacity of 11 events. If the channel is full, new events are dropped, leading to unbounded response time (the dropped callback instances will never be executed).

The registration of entity callbacks differs from C++, where methods for creating timers or subscriptions take the callback as a parameter. In R2R, the corresponding methods return the receiving end of the associated MPSC channel, and the callback is bound to it by spawning an async task created from an async block, which periodically awaits the timer expiration or the message from the channel. See Listing 1 for an example. This way, if a subscription channel contains more than one message, the callback is consecutively executed for all messages available in the channel.

The execution of the callbacks is carried on by the executors. An example of the main R2R loop with the `FUTURES` local executor is shown in Listing 2. The call to `run_until_stalled` before the loop at line 1 is necessary if we want to avoid memory allocation in the loop. The call initializes the executor by moving the tasks from the incoming vector to the active list, which involves memory allocation. Since the `spin_once` function was not called before, no tasks were sampled, and no callback code is going to be executed.

When calling the `spin_once` function provided by R2R (line 3), it wakes up receiving tasks of all ready entity channels. Subscriptions first, then timers, etc. Note that this order differs from C++ executors. The order of waking tasks of the same entity type follows the creation order of the entities in the node.



■ **Figure 2** Example of execution of a chain of operations from publishing two messages in one node to processing it with R2R in another node. Horizontal lines represent different involved threads. The subscriptions in R2R were created in alphabetical order of topics.

<code>spin_once</code> call	Ready entities after sampling	Executor ready queue state after sampling	Order of callback execution by the first call to <code>spin_until_stalled</code> after sampling.
1	B D	← [B] [D] ←	B ₁ , D
2	A B C	← [B] [D] [A] [C] ←	B ₁ , B ₂ , D, A, C

■ **Figure 3** Content of executor FIFO ready queue when two subsequent sampling calls return the same entity (B), producing two messages B₁ and B₂.

The second call to `run_until_stalled` inside the loop at line 4 executes the callbacks. Since the local executor schedules tasks in the FIFO manner, the tasks will run in the order in which they were woken up.

In such a setup, each channel will store at most one message because for each ready entity `spin_once` pushes only a single message to its channel. A subsequent call to `run_until_stalled` then processes all channels, leaving them empty. Therefore, the channel will never drop messages due to it being full.

Figure 2 shows the above described sequence of operations in time. Besides the R2R thread running the loop from Listing 2 (denoted as Node S), it shows publisher and DDS threads involved in the process.

Note that the callbacks may not always be executed in the FIFO order as described above. One such case can happen if `spin_once` is called multiple times without calling `run_until_stalled` in between. This is demonstrated in the example in Figure 3. The rightmost column shows in which order would be the callbacks called by a call to `run_until_stalled` following the `spin_once` calls.

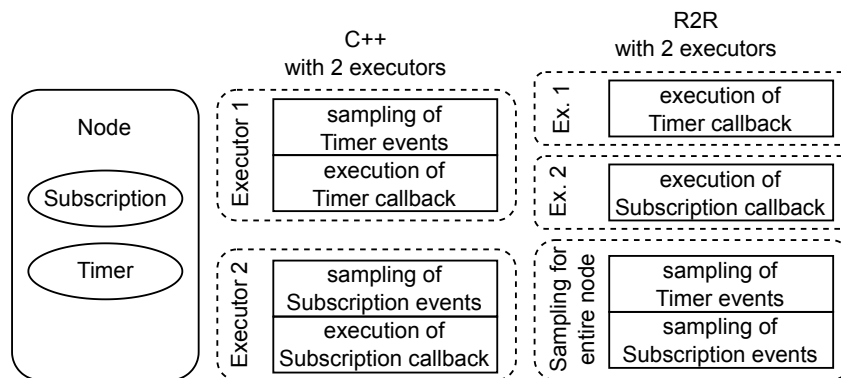
A similar effect can happen if multiple futures are joined as described in Section 2.2.1.3. An example of how the callbacks would be executed when grouping them via the `join` macro is shown in Figure 4.

2.6 Comparison of C++ and R2R execution models

The main difference between C++ and R2R execution models is summarized in Figure 5. R2R does not sample events in executors because waiting on a wait set is a synchronous blocking operation, which should not be executed from an async task.

Variant	Ready entities (in creation order)	Executor ready queue state after sampling with joined A, C, and E
1	A B C D E	← [A C E] [B D] ←
2	B C D E	← [B] [C E] [D] ←

■ **Figure 4** Execution order of the join group depends on the waking time of the joined tasks. In variant 1, the group is woken by waking of task A. In variant 2, task A is not ready, and the group is woken by task C.



■ **Figure 5** The sampling events of node's entities cannot be split into subsets in R2R, unlike in C++, but the callback execution can be assigned to different executors.

However, if the R2R application is organized as in Listing 2, the execution is equivalent to using a ROS C++ single-threaded executor with a minor difference that the priorities of subscription and timer callbacks are swapped. This means that the response-time analysis proposed in [2] can be applied to R2R as is while taking into account the difference in priorities.

To achieve the behavior of a C++ multi-threaded executor with R2R, sampling must be performed in a dedicated thread with a loop shown in Listing 3. The effect of different C++ callback group configurations can be accomplished as follows:

When all callbacks are executed in another dedicated thread with FUTURES local executor, e.g. by calling just `local_executor.run()`, this is an analogy to having all callbacks in the same mutually exclusive group.

When callbacks are executed by FUTURES thread-pool executor or Tokio runtime, this is equivalent to using ROS C++ multi-threaded executor with each callback in its own mutually exclusive group, i.e., each callback can run concurrently with other callbacks but not with itself. It is possible to combine tasks to create a larger mutually exclusive group by using the `join` macro.

■ **Listing 3** R2R sampling loop to mimic ROS C++ multi-threaded executor.

```
loop {
    node.spin_once(Duration::seconds(1));
}
```

To achieve the same behavior as the reentrant callback group of ROS C++ multi-threaded executor, where the tasks can be executed concurrently even with itself, the callback code should dynamically create tasks and spawn them to the Rust multi-threaded executor (e.g. `FUTURES ThreadPool` or `Tokio`).

3 Achieving deterministic real-time operation with R2R

To configure R2R for use in real-time applications, we propose to run the application under the Linux `SCHED_FIFO` scheduler and structure it as summarized by the following two rules:

- The main thread should run with the highest priority and should sample events from ROS, i.e., it should call `Node::spin_once` in a loop.
- The callbacks should be executed in lower-priority threads, each running one `FUTURES` local executors.

An example of the proposed structure is given in Listing 4. The priority of the main thread is set on lines 14–18, before creating ROS context and node at lines 20 and 21. A single subscriber is then created on line 23, followed by associating it with a callback implemented as an `async move` block (lines 24–26). Line 27 spawns a new OS thread with the given priority. The code executed by the thread is a closure defined at lines 5–10. It creates a new local executor (line 6), spawns to it a new async task from the passed future (line 8) and runs the executor loop (line 9).

Running the main thread with the highest priority serves two purposes. First, RMW/DDS threads, which are created internally during ROS initialization, inherit the same priority and, therefore, will never be delayed by callback execution. Second, sampling of entities for events will happen almost immediately after the event happens, without waiting for any callback to complete.

Execution of callbacks in lower-priority threads ensures that one can use schedulability analysis, which does not need to take into account ROS specifics and depends only on the scheduling policy of the OS scheduler. In Section 4, we demonstrate this by running each callback in a thread pinned to a single CPU and with priority set according to the rate-monotonic ordering and use uni-processor response-time analysis [1]⁶ to predict the response times.

Note that in the current R2R implementation, where all MPSC channels have the same fixed capacity, it could happen that a channel served by a callback in a low-priority thread could overflow, leading to message losses. Therefore, for reliable operation in all situations, R2R should be extended to make channel capacity configurable, and one should use a schedulability analysis to dimension the capacity of individual channels.

4 Experimental evaluation

To evaluate the proposed application structure and compare it with other alternatives, we designed a set synthetic benchmarks. These are described in Section 4.1. Evaluation on a more complex autonomous driving application follows in Section 4.2.

⁶ Application of more complex analysis techniques should be possible and is left for future work.

■ Listing 4 Proposed structure of a real-time R2R application.

```

1  fn spawn_in_thread(future: impl Future, priority: ThreadPriority) {
2      let thread = ThreadBuilder::default()
3          .policy(RealTime(Fifo))
4          .priority(priority)
5          .spawn(move |_| {
6              let mut local_executor = executor::LocalPool::new();
7              let spawner = local_executor.spawner();
8              spawner.spawn_local(future).unwrap();
9              local_executor.run();
10         });
11 }
12
13 fn main() -> Result<(), Box<dyn Error>> {
14     thread_priority::unix::set_thread_priority_and_policy(
15         thread_priority::thread_native_id(),
16         ThreadPriority::try_from(MAIN_PRIORITY)?,
17         RealTime(Fifo),
18     )?;
19
20     let ctx = r2r::Context::create()?;
21     let mut node = r2r::Node::create(ctx, "example", "")?;
22
23     let subs = node.subscribe("/topic", QosProfile::default())?;
24     let future = subs.for_each(move |msg: Msg| async move {
25         // process msg
26     });
27     spawn_in_thread(future, ThreadPriority::try_from(CALLBACK_PRIORITY)?);
28
29     loop {
30         node.spin_once(SPIN_TIMEOUT);
31     }
32 }

```

4.1 Synthetic benchmarks

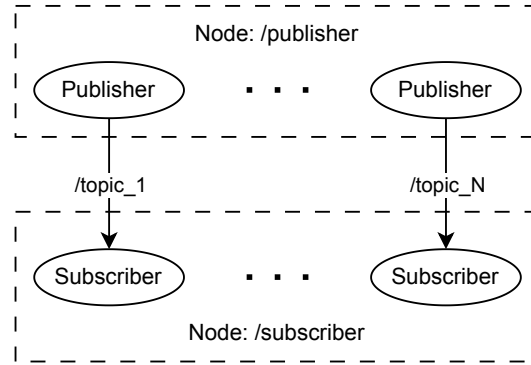
To evaluate the proposed application structure and compare response times obtained with different setups, we develop a benchmark application composed of five topics (see Figure 6). Messages are published periodically to the topics by a publisher node running in a dedicated process. The second node subscribes to the topics and executes callbacks with a specific fixed execution time. The publication periods and callback execution times are given in Table 2 and lead to CPU utilization of 90%.

We implement the subscribers in several variants using different execution strategies. The implemented variants are attached to the paper. We run the application and trace its execution with the help of LTtng⁷ toolkit. For C++ variant, we reuse the tracepoints already present in the ROS libraries, for R2R variants, we add tracepoints to R2R and also to subscription callbacks. The collected traces allow the measurement of intermediate or end-to-end latencies of individual messages and callback executions in the application. Then, we calculate various statistics from them.

⁷ <https://ltnng.org/>

■ **Table 2** Parameters of publishers and subscriptions in the benchmarking application.

Topic #	1	2	3	4	5
Publisher period [ms]	10	20	50	100	200
Subscription callback execution time [ms]	2	4	5	15	50



■ **Figure 6** ROS application used for evaluation with $N = 5$.

4.1.1 Execution environment

All experiments are executed on a laptop with AMD Ryzen 7 3700U CPU running Ubuntu 24.04 with the Linux kernel 6.8.0. The used ROS distribution is ROS Jazzy Jalisco. Processes that need their main thread to run with real-time priority (SCHED_FIFO) were started via the `chrt` command so that the RMW/DDS threads inherit the priority and policy from it. Other threads' priorities are set by calling `pthread_setschedparam` either directly in C++ or through the Rust `thread-priority` crate.

All threads in all experiments were run on a single isolated CPU core (hyper-threading is enabled, but both threads are set as isolated; only one of them is used for experiments) by using the `isolcpus` kernel command line argument. CPU affinity of the threads was set by starting the programs via the `taskset` command.

4.1.2 Publisher node implementation

The publisher node is implemented with R2R and Tokio runtime. Publications are triggered by absolute time timers implemented via the Linux `timerfd` system call and connected to Tokio. Each timer callback is an async task that publishes a message containing a single 64-bit integer. All threads are scheduled by real-time SCHED_FIFO scheduler. The priority of the main and RMW/DDS threads is set to 25, a single Tokio worker thread has priority of 24.

4.1.3 Implementation variants of the subscriber node

We implemented the following variants of the subscriber node that we compare. Below, we provide their names together with a short description.

futures. R2R-based subscriber node with FUTURES local executor running interleaved with the `spin_once` function in the same thread, as in Listing 2.

futures-join. R2R-based subscriber node with FUTURES thread-pool executor. `spin_once` is running alone in the main thread. All callback tasks are joined together; the single resulting task runs in the thread pool.

futures-rt. The variant proposed in Section 3, that is R2R-subscriber node with each callback running in a separate thread with real-time priority set according to the rate-monotonic (RM) priority assignment (the smaller period, the higher priority) in a FUTURES local executor. The `spin_once` function is called in a loop in the main thread with the highest priority.

futures-thread-pool. R2R-based subscriber node with a FUTURES thread-pool executor. `spin_once` is running in the main thread, and the thread-pool running the callbacks has two threads.

futures-2-threads. R2R-based subscriber node with the `spin_once` function is called in a loop in the main thread with the highest priority. The second thread executes all callbacks in the FUTURES local executor.

rcldcpp-rt. C++ subscriber node with each subscriber executed by its own single-threaded executor running on a dedicated thread with fixed real-time priority (RM). This is a similar scheme to the one proposed in [12].

rcldcpp-st. C++ subscriber node with the default single-threaded executor (started by `rcldcpp::spin(node)`) for all callbacks.

tokio. R2R-based subscriber node using Tokio to execute callbacks. All threads in the process inherit the same priority from the main thread (Tokio's default). The `spin_once` function is executed in a loop on a separate dedicated thread. Callbacks are executed by Tokio's async runtime.

tokio-rt. R2R-based subscriber node using Tokio to execute the callbacks. The `spin_once` function is executed in a loop on a separate dedicated thread. The worker threads running callbacks have all the same priority lower than the spin thread.

In all variants mentioned above, the main thread, along with all DDS threads, has priority 21, and if callback-running threads have different priorities, they use priorities from 20 down to 16, respectively. Tokio workers in *tokio-rt* have priority 20. All threads are executed with the `SCHED_FIFO` scheduler.

The only exceptions are variants prefixed with `nort`, in which the subscriber node with all its threads are scheduled by the default scheduler (`SCHED_OTHER`).

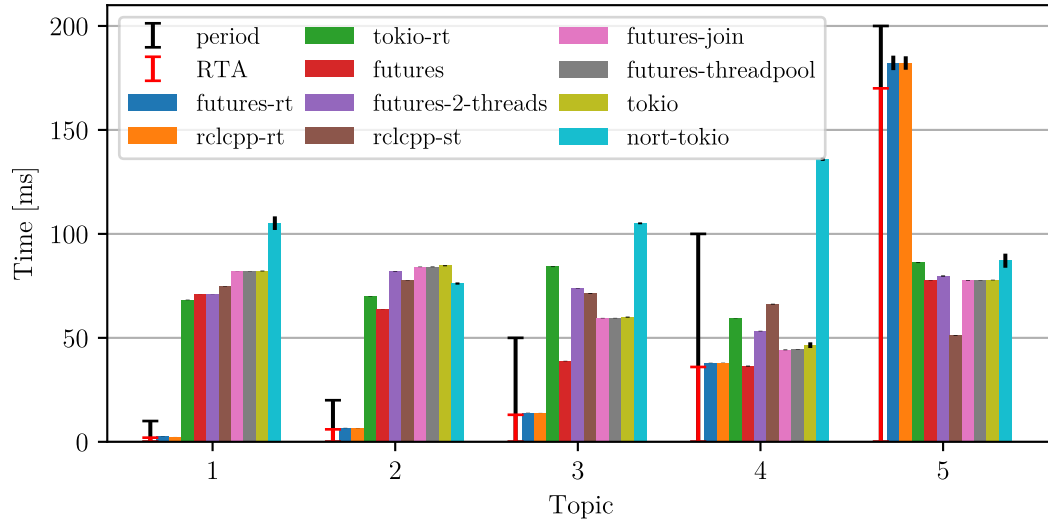
The callback work is emulated by a loop that executes for the given time. The elapsed time is measured by calling `clock_gettime()` with clock type `CLOCK_THREAD_CPUTIME_ID`.

Subscriptions are created with ROS Quality of service set to keep the last 100 messages because if we kept only one message, subscriptions with a response time greater than its period would lose messages. The value 100 was chosen because it is greater than what is actually needed so as not to lose any message.

4.1.4 Synthetic benchmark results

We executed all implemented variants for 20 seconds, leading to the publication of 100 to 2000 messages to the respective topics. All experiments were executed 10 times, and standard deviations of different runs were calculated and reported in graphs with error bars. The 99th percentiles of the measured end-to-end latencies (from the publication of a message to the end of callback execution) are reported in Figure 7. The graph also shows the publication periods and the expected worst-case response time calculated with a classical mono-processor response-time analysis (RTA) [1].

As can be seen, our proposed *futures-rt* achieves the same results as C++ *rcldcpp-rt*, and the results of both variants match the theoretical worst-case response time calculated by the RTA. Only topic 5, running in the lowest-priority thread, exhibits higher latencies than



■ **Figure 7** End-to-end latency (99th percentile) of various subscriber implementations. Black error bars represent standard deviations.

expected, which is caused by overheads of the sampling and RMW/DDS threads, which are not considered by our RTA⁸. Note that *futures-rt* and *relepp-rt* are the only variants, where all callbacks manage to complete before the deadlines given by the publication periods.

All other implemented variants fail to meet shorter deadlines for topics 1–3. This is expected for the default configurations of Rust async runtimes, which are not able to take into account the timing requirements of individual callbacks. This is especially visible for the *nort-tokio* variant, which has the highest latencies for three of five topics.

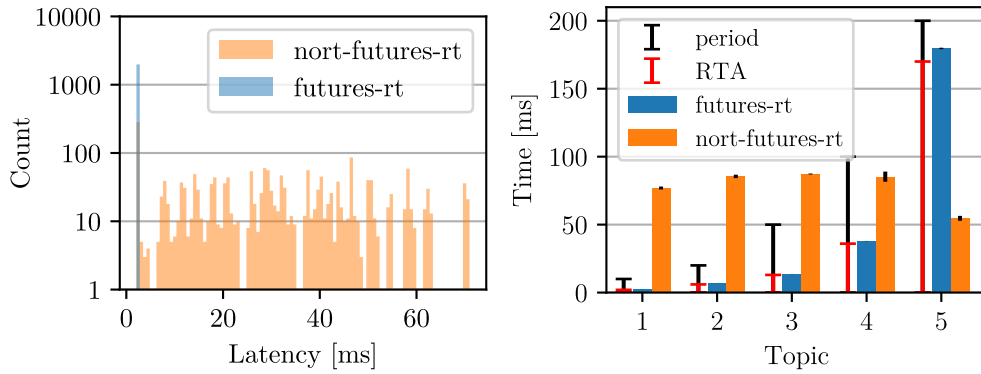
To investigate the effect of RMW/DDS thread priority, we compare the end-to-end latencies of our proposed *futures-rt* with and without giving the RMW/DDS threads the highest priority. The results in Figure 8 show that without setting the priority of RMW/DDS threads to be higher than the priorities of the callback threads, the callbacks fail to complete before the deadlines given by the publication periods.

4.2 Complex autonomous driving application

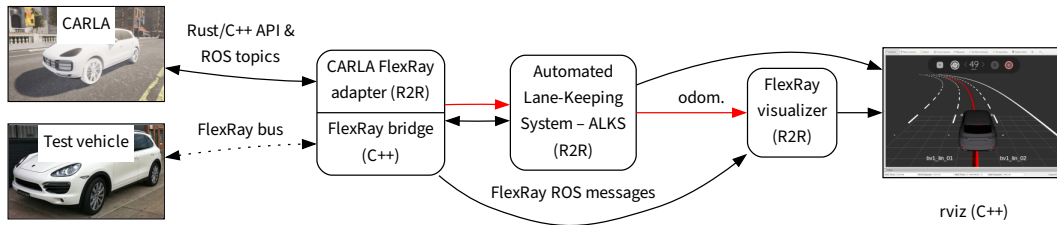
To evaluate the suitability of R2R for implementing more complex real-time applications, we used it to implement a simplified version of the Automated Lane Keeping System (ALKS) – a system that can automatically drive the car on highways with speed up to 130 km/h [26]. ALKS is the first autonomous driving system legally allowed in Europe that reaches SAE Level 3 of automation [17], meaning that driver supervision is not required.

Our ALKS implementation can drive a real Porsche Cayenne car by connecting to its FlexRay buses and communicating with the onboard control units. However, in this paper, we evaluate the version running against the simulated vehicle in the CARLA simulator [7]. The reason is that this version contains three ROS nodes written using R2R, whereas the real car version has only two R2R-based nodes. The high-level architecture of the application is depicted in Figure 9.

⁸ It should be possible to model the overheads as separate tasks or associate them to existing tasks, however, it would require a more detailed execution time analysis than what we do in this paper.



■ **Figure 8** Comparison of end-to-end latency with and without setting the real-time priority of RMW/DDS threads. Left: Histogram of latencies of the highest-rate topic 1. Right: 99th percentile of latencies of all topics.



■ **Figure 9** Architecture of the complex application when using the CARLA simulator or a real car. Red arrows represent the odometry chain evaluated below.

The first R2R node is the *CARLA FlexRay adapter*, which communicates with the CARLA simulator via its API (by using the `carla-rust`⁹ bindings) and ROS topics¹⁰. It publishes the information from the simulator to ROS in the same format as used when interfacing the real car via the C++-based *FlexRay bridge* node. The published information consists of eleven topics, which include quantities like speed, steering wheel angle, GPS position, odometry, Inertial Measurement Unit (IMU), information about other visible vehicles, and detected road lines. Most of the quantities are published with a frequency of 50 Hz, and information about vehicles and lines with a frequency of 25 Hz. Producing road line messages involves the execution of a non-trivial curve-fitting algorithm for the conversion of simulated lines to their parametric spline representation used on the FlexRay bus. In the opposite direction, these nodes convert the control commands for acceleration/braking and steering received via ROS to the simulator requests or the FlexRay messages.

The implementation of ALKS resides in the ALKS ROS node. The information received via eleven topics from the FlexRay is processed in the respective callbacks, where it is converted with simple calculations to the internal representation and stored in memory for later use. The main computation happens in a callback of a 50 Hz timer running in a separate thread, where all logic and PID controllers are calculated. Besides that, the computationally more demanding Model-Predictive Controller (MPC) that calculates optimal trajectory is

⁹ <https://github.com/jerry73204/carla-rust>

¹⁰ We need to combine ROS topics with CARLA API because CARLA 0.10.0 does not publish all needed information via ROS topics.

■ **Table 3** Overview of threads in R2R nodes of the ALKS application. All threads are scheduled using SCHED_FIFO and those running callbacks use `run*` methods of FUTURES local executor.

Node	Thread(s)	Callbacks	Priority	WCET ¹
FlexRay adapter	DDS	–	20	–
	loop { spin_once }	–	10	–
	run	all	10	0.3 ms [†]
	CARLA client threads	–	10	–
ALKS	DDS	–	20	–
	loop { spin_once; run_until_stalled }	subscribers	10	0.13 ms*
	run	timer	5	9.8 ms
FlexRay visualizer	DDS	–	20	–
	loop { spin_once; run_until_stalled }	all	10	0.26 ms*

¹ 99th percentile

[†] WCET of the odometry callback only.

* WCET of `run_until_stalled`, which can invoke one or more callbacks.

invoked from the same callback every 300 ms. At the end of every callback invocation the node publishes control commands for the vehicle as well as debugging and visualization topics. Topics related to the current trajectory are published whenever a new one is calculated.

The last R2R-based node is the *FlexRay visualizer*. It receives information from the real or simulated FlexRay bus and converts it to rviz markers, i.e., ROS messages, which can be visualized in 3D by the ROS `rviz` tool. It also receives preprocessed odometry information from ALKS, which is used for the visualization.

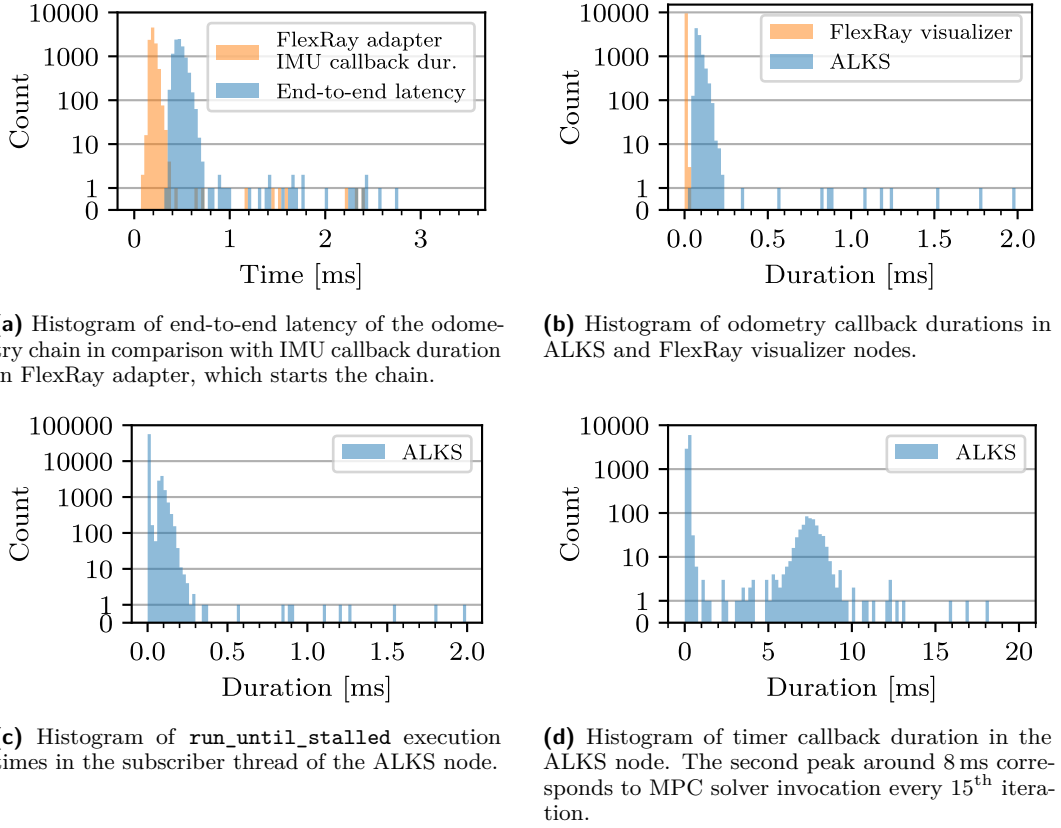
The overview of threads in the nodes and their SCHED_FIFO priorities is given in Table 3. Only the ALKS node runs callbacks in more than one thread.

4.2.1 Latency evaluation

To show that R2R is suitable for implementing practical real-time applications, we run the above-described ALKS application and trace it with LTTng. As opposed to the synthetic benchmark, here we did not restrict the threads to run on a single CPU core. From the traces, we obtain end-to-end latencies of the *odometry* chain (marked with red arrows in Figure 9). This chain goes through all three R2R-based nodes in the application. Specifically, it starts at the reception of the IMU message from CARLA in the FlexRay adapter. In the associated callback, we retrieve other information from CARLA using its API and publish the obtained data to the corresponding ROS topic. The ALKS node then receives (among others) the odometry message, converts it from relative to absolute values, and stores it for later use. It also publishes the converted absolute values for use in the FlexRay visualizer. Flexray visualizer just stores the received odometry values in memory for use in other callbacks.

Figure 10 shows histograms obtained from a trace recorded during approximately 3 minutes of execution of the ALKS application. We removed few initial samples from the trace to filter out outliers from the warm-up phase.

The histogram of measured odometry chain end-to-end latencies is shown in Figure 10a. In about 99% of cases, the latency is below 0.8 ms. The same figure also shows execution time (duration) of the first callback in the chain (IMU). As this callback communicates with the CARLA simulator over the network, its duration and jitter is dominated by this communication latency and this propagates further down the chain. Figure 10b shows



■ **Figure 10** Histograms of obtained from traces of the ALKS application.

execution time histograms of other callbacks in the chain. In the ALKS node, the majority of callback executions exhibits execution time jitter of about 0.2 ms, which is caused by waiting for mutexes protecting data shared with the timer callback running in another thread. A few outliers between 0.3 and 2 ms can be caused by the fact that the experiment was executed on a machine connected to the network and running other development tools, which might cause some interference. The jitter of the callback in the FlexRay visualizer node is very low, as the node is single-threaded and does not need to use mutexes in the callbacks.

To put this into the context, Figure 10c shows histogram of `run_until_stalled` call execution times in the ALKS node, whose 99th WCET percentile is mentioned in Table 3. This function calls the odometry callback as well as other subscriber callbacks in the node. Note that due to the higher number of its invocations when compared to the odometry callback, the 99th WCET percentile is smaller (0.13 ms) than the one of only the odometry callback (0.16 ms).

Figure 10d shows the execution time histogram of the timer callback in the ALKS node. The left peak represents the cases where the MPC planner was not invoked, whereas the peak around 8 ms includes the cases with planner invocations. Note that the end-to-end latency of the odometry chain is influenced by the timer callback only through short critical sections for accessing shared data. Long execution time of the MPC planner has no effect on the end-to-end latency of the odometry chain. This demonstrates that R2R with the FUTURES runtime and the structure proposed in this paper is capable of achieving deterministic end-to-end latencies of the chain of callbacks executing in different nodes.

5 Discussion

Designing a real-time ROS application with R2R is not much different from designing a C++ application with `rc1cpp`. One needs to carefully plan mapping of callbacks to OS threads and assign thread scheduling parameters according to the timing requirements. An advantage of the R2R approach is that sampling of ROS entities can be run in a thread independent from threads running application callbacks, decreasing time between samples to almost zero. This allows analyzing callback schedulability independently of ROS sampling.

An important difference between C++ and Rust ROS applications is how the application should structure shared data. While C++ does not impose many restrictions on the structure, Rust requires explicit association of data with their synchronization primitives. Structuring the data inappropriately causes problems with Rust borrow checker as well as with unwanted blocking of unrelated threads, negatively influencing the timing.

An advantage of using the asynchronous programming style is that it can simplify application code, especially if combining data from messages from multiple topics. In such a case, the Rust compiler automatically constructs the state machines, which would need to be implemented manually when using synchronous programming style. This simplified the design of our CARLA FlexRay adapter.

Another benefit of asynchronous programming is that applications can easily integrate reactions to events from different event sources. For example, ROS can be one event source and other ROS-unaware 3rd party libraries can provide other sources. Depending on the chosen asynchronous runtime, this can take advantage of using efficient event demultiplexing system calls like `epoll` or `io_uring`.

6 Related work

The study of real-time scheduling in ROS 2 began with Casini et al. [3], who modeled and analyzed mainline ROS 2. Their work led to corrections that enabled the development of formal schedulability analyses. Building on these foundations, researchers proposed schedulability analyses for the C++ single-threaded executor [2] or designed and analyzed a new custom executor [4]. Subsequent investigations into ROS multi-threaded executors [10, 20, 24] not only advanced schedulability analyses but also uncovered implementation flaws in the executor design. Beyond the scheduling in the executors, researchers have approached the latencies of DDS communication middleware from both experimental [11] and analytical [18] perspectives, providing a more comprehensive understanding of real-time performance in ROS 2 systems.

7 Conclusion

In this paper, we analyzed the execution model of ROS 2 applications written in asynchronous Rust with the R2R ROS client library. The execution model is flexible, allowing the application designers to choose from many variants and asynchronous runtimes. We experimentally evaluated many of the variants and end-to-end latencies achieved with them. We proposed the *futures-rt* structure, which decouples sampling and callback execution to different threads, prioritizes sampling and DDS threads, and runs callbacks in dedicated lower-priority threads. This structure achieves the best results among all evaluated options and meets all deadlines. The measured latencies match well the theoretical response times calculated with the classical response-time analysis. Practical applicability of the proposed approach was demonstrated on an autonomous driving application implementing the Automated Lane Keeping System.

References

- 1 N. Audsley, A. Burns, M. Richardson, K. Tindell, and A.J. Wellings. Applying new scheduling theory to static priority pre-emptive scheduling. *Software Engineering Journal*, 8(5):284–292, September 1993. doi:10.1049/sej.1993.0034.
- 2 Tobias Blaß, Daniel Casini, Sergey Bozhko, and Björn B. Brandenburg. A ROS 2 Response-Time Analysis Exploiting Starvation Freedom and Execution-Time Variance. In *2021 IEEE Real-Time Systems Symposium (RTSS)*, pages 41–53, December 2021. doi:10.1109/RTSS52674.2021.00016.
- 3 Daniel Casini, Tobias Blass, Ingo Lütkebohle, and Björn B. Brandenburg. Response-Time Analysis of ROS 2 Processing Chains Under Reservation-Based Scheduling. In Sophie Quinton, editor, *31st Euromicro Conference on Real-Time Systems (ECRTS 2019)*, volume 133 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 6:1–6:23, Dagstuhl, Germany, 2019. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.ECRTS.2019.6.
- 4 Hyunjong Choi, Yecheng Xiang, and Hyoseung Kim. PiCAS: New Design of Priority-Driven Chain-Aware Scheduling for ROS2. In *2021 IEEE 27th Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 251–263, May 2021. doi:10.1109/RTAS52030.2021.00028.
- 5 Martin Dahl. R2R: Easy to use, runtime-agnostic, async Rust bindings for ROS 2. Sequence Planner, February 2025. URL: <https://github.com/sequenceplanner/r2r>.
- 6 Martin Dahl, Endre Erős, Kristofer Bengtsson, Martin Fabian, and Petter Falkman. Sequence Planner: A Framework for Control of Intelligent Automation Systems. *Applied Sciences*, 12(11):5433, January 2022. doi:10.3390/app12115433.
- 7 Alexey Dosovitskiy, German Ros, Felipe Codevilla, Antonio Lopez, and Vladlen Koltun. CARLA: An Open Urban Driving Simulator. In *Proceedings of the 1st Annual Conference on Robot Learning*, pages 1–16. PMLR, October 2017. ISSN: 2640-3498. URL: <https://proceedings.mlr.press/v78/dosovitskiy17a.html>.
- 8 Eclipse Foundation. Zenoh – The Zero Overhead, Pub/Sub, Store, Query, and Compute Protocol. URL: <https://zenoh.io/>.
- 9 eProsima. Fast DDS. URL: <https://www.eprosima.com/middleware/fast-dds>.
- 10 Xu Jiang, Dong Ji, Nan Guan, Ruoxiang Li, Yue Tang, and Yi Wang. Real-Time Scheduling and Analysis of Processing Chains on Multi-threaded Executor in ROS 2. In *2022 IEEE Real-Time Systems Symposium (RTSS)*, pages 27–39, December 2022. doi:10.1109/RTSS55097.2022.00013.
- 11 Tobias Kronauer, Joshua Pohlmann, Maximilian Matthé, Till Smejkal, and Gerhard Fettweis. Latency Analysis of ROS2 Multi-Node Systems. In *2021 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, pages 1–7, September 2021. doi:10.1109/MFI52462.2021.9591166.
- 12 G. M. Krukiewicz-Gacek. A Priority-Based Real-Time Scheduling Framework for ROS2. Master’s thesis, TU Delft, 2021. URL: <https://repository.tudelft.nl/record/uuid:4bfcfe05-d1a5-4125-b2a4-9187fb1d1236>.
- 13 Janosch Machowinski. The Callbackgroup Events Executor, October 2024. URL: https://web.archive.org/web/20241123211507/https://roscon.ros.org/2024/talks/The_Multithreaded_Events_Executor.pdf.
- 14 OMG. Data Distribution Service Specification Version 1.4. URL: <https://www.omg.org/spec/DDS/>.
- 15 Open Robotics. ROS 2 Jazzy, May 2024. URL: <https://docs.ros.org/en/jazzy/>.
- 16 ros2_rust. Rust bindings for ROS 2. ros2-rust, February 2025. URL: https://github.com/ros2-rust/ros2_rust.
- 17 SAE International. J3016_202104: Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles - SAE International, April 2021. URL: https://www.sae.org/standards/content/j3016_202104/.

- 18 Gerlando Sciangula, Daniel Casini, Alessandro Biondi, Claudio Scordino, and Marco Di Natale. Bounding the Data-Delivery Latency of DDS Messages in Real-Time Applications. In *DROPS-IDN/v2/Document/10.4230/LIPIcs.ECRTS.2023.9*. Schloss-Dagstuhl - Leibniz Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.ECRTS.2023.9.
- 19 Martin Škoudlil. Real-Time schedulability analysis of ROS 2 with nodes using asynchronous Rust. Master's thesis, Czech Technical University in Prague, January 2025. URL: <https://dspace.cvut.cz/handle/10467/120369>.
- 20 Hooria Sobhani, Hyunjong Choi, and Hyoseung Kim. Timing Analysis and Priority-driven Enhancements of ROS 2 Multi-threaded Executors. In *2023 IEEE 29th Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 106–118, May 2023. doi:10.1109/RTAS58335.2023.00016.
- 21 Alberto Soragna. The ROS 2 C++ Executors – General, June 2024. URL: <https://discourse.ros.org/t/the-ros-2-c-executors/38296>.
- 22 Yue Tang, Zhiwei Feng, Nan Guan, Xu Jiang, Mingsong Lv, Qingxu Deng, and Wang Yi. Response Time Analysis and Priority Assignment of Processing Chains on ROS2 Executors. In *2020 IEEE Real-Time Systems Symposium (RTSS)*, pages 231–243, December 2020. doi:10.1109/RTSS49844.2020.00030.
- 23 Harun Teper, Tobias Betz, Mario Günzel, Dominic Ebner, Georg Von Der Brüggen, Johannes Betz, and Jian-Jia Chen. End-To-End Timing Analysis and Optimization of Multi-Executor ROS 2 Systems. In *2024 IEEE 30th Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 212–224, May 2024. doi:10.1109/RTAS61025.2024.00025.
- 24 Harun Teper, Daniel Kuhse, Mario Günzel, Georg von der Brüggen, Falk Howar, and Jian-Jia Chen. Thread Carefully: Preventing Starvation in the ROS 2 Multithreaded Executor. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 43(11):3588–3599, November 2024. doi:10.1109/TCAD.2024.3446865.
- 25 Tokio.rs. Multi threaded runtime documentation. URL: <https://docs.rs/tokio/1.43.0/tokio/runtime/index.html#multi-threaded-runtime-behavior-at-the-time-of-writing>.
- 26 UNECE. Uniform provisions concerning the approval of vehicles with regard to Automated Lane Keeping Systems, January 2023. URL: <https://unece.org/transport/documents/2023/03/standards/un-regulation-no-157-amend4>.