

Trace-enabled Timing Model Synthesis for ROS2-based Autonomous Applications

Hazem Abaza^{*†}, Debayan Roy^{*}, Shiqing Fan[‡], Selma Saidi[†] and Antonios Motakis^{*}

[†]Technische Universität Dortmund, ^{*}Huawei Dresden Research Center, [‡]Huawei Munich Research Center
{hazem.abaza, selma.saidi}@tu-dortmund.de, {debayan.roy6, shiqing.fan, antonios.motakis}@huawei.com

Abstract—Autonomous applications are typically developed over Robot Operating System 2.0 (ROS2) even in time-critical systems like automotive. Recent years have seen increased interest in developing model-based timing analysis and schedule optimization approaches for ROS2-based applications. To complement these approaches, we propose a tracing and measurement framework to obtain timing models of ROS2-based applications. It offers a tracer based on extended Berkeley Packet Filter that probes different functions in ROS2 middleware and reads their arguments or return values to reason about the data flow in applications. It combines event traces from ROS2 and the operating system to generate a directed acyclic graph showing ROS2 callbacks, precedence relations between them, and their timing attributes. While being compatible with existing analyses, we also show how to model (i) message synchronization, e.g., in sensor fusion, and (ii) service requests from multiple clients, e.g., in motion planning. Considering that, in real-world scenarios, the application code might be *confidential* and formal models are unavailable, our framework still enables the application of existing analysis and optimization techniques. We demonstrate our framework’s capabilities by synthesizing the timing model of a real-world benchmark implementing LIDAR-based localization in Autoware’s Autonomous Valet Parking.

I. INTRODUCTION

Developing industry-strength autonomous applications requires teams of engineers with diverse backgrounds. Robot Operating System version 2 (ROS2) is a powerful middleware over which modular software components can be developed and composed easily to create autonomous applications. To leverage these benefits and the vast amount of open-source contributions, ROS2 is starting to get used even in time-critical systems such as self-driving cars [1].

In recent years, there have been efforts to develop timing analysis and optimization approaches for ROS2-based applications, e.g., [2]–[6]. These approaches typically assume that the application models are well-defined, i.e., the execution times of and the precedence relations between ROS2 callbacks are known. However, in many industry scenarios, such models are not provided by application developers. Further, during system integration, it is challenging to obtain many details—especially at the level of callbacks—due to confidentiality reasons.

In parallel to the model-based techniques, tracing ROS2-based applications have also gained interest. In this context, *ros2_tracing* provides a framework based on Linux Trace Toolkit: next generation (LTTng) [7]. It has tracepoints in ROS2 functions to identify callbacks and topics and also track them during runtime. *Autoware_Perf* [8] and *CARET* [9] add more tracepoints and use trace data to measure the response time of a callback, the communication latency between a pair

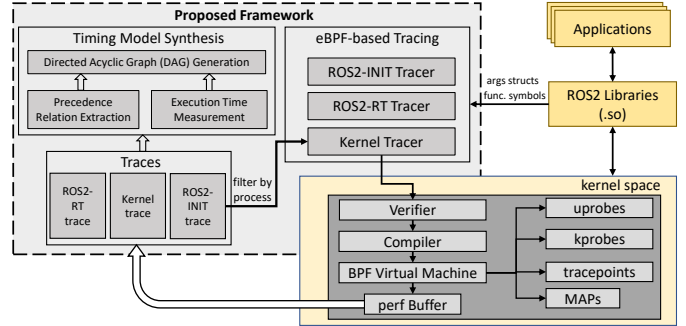


Fig. 1: Proposed trace-enabled timing model synthesis framework.

of callbacks, and the end-to-end latency of a callback-chain. Nevertheless, they do not show how to compute the execution times of callbacks and also provide insufficient information to extract precedence constraints in certain scenarios, e.g., in message synchronization.

Proposed timing model synthesis framework: This paper bridges the gap between tracing and timing analysis by synthesizing timing models of ROS2 applications. We use extended Berkeley Packet Filter (eBPF) [10] for tracing. Compared to LTTng, it does not require direct instrumentation and recompilation of standard ROS2 libraries and offers efficient trace filtering. Unlike LTTng, eBPF does not give access to local variables in functions and, hence, we could not re-use many tracepoints identified by [7]–[9]. Besides probing new ROS2 functions, we traverse complex data structures of their arguments to get desired information, e.g., topic names, callback IDs, and source timestamps of data. Our proposed framework, in Fig. 1, provides three tracers: (i) *ROS2-INIT* tracer logs the initialization of ROS2 nodes. (ii) During runtime, *ROS2-RT* tracer tracks the start and end of callbacks and data read from and write to Data Distribution Service (DDS) topics. Unlike [7]–[9], it also tracks client callbacks and message synchronization, which are found, e.g., in AUTOWARE’s Autonomous Valet Parking (AVP) [1]. (iii) *Kernel* tracer logs scheduler events in the operating system (OS) related to ROS2 callbacks only with the help of eBPF’s trace filtering.

As illustrated in Fig. 1, our proposed framework uses the collected traces to synthesize timing models of applications as directed acyclic graphs (DAGs). We model ROS2 callbacks as tasks (or vertices) and DDS communication between them using precedence relations (or edges). Further, we propose to model a service using n tasks if it is invoked by n different clients. Otherwise, DAG will show one vertex with n incoming

edges and n outgoing edges, i.e., $n \times n$ chains passing through the vertex, which is a wrong interpretation. Also, we propose to model a synchronization of m data, e.g., for sensor fusion, using $m + 1$ tasks, i.e., m tasks reading data each and then these m tasks feed data to another task that outputs the result.

Further, we show how to *combine* ROS2 and scheduler events to measure the execution time of a callback for each invocation. We get statistical information from these measurements, e.g., measured worst-case, best-case, and average values. We annotate the DAG with the obtained values.

The DAG we generate can serve as an input for analysis and optimization by, e.g., [2]–[6]. We employ our framework to synthesize the timing model of a *real-world benchmark* implementing LIDAR-based localization in AVP. We note that our framework does not need any modification in the application code unlike [11]. Also, it is *not tied* to any particular application because we probe functions in ROS2 middleware and not in applications directly.

II. PRELIMINARIES

A. ROS2 architecture and semantics

We study autonomous applications developed over ROS2. They comprise ROS2 *nodes* where each performs a particular function, e.g., localization and object detection. Nodes communicate using *topics*. A node publishes data on a topic, e.g., a “point cloud filter” node can publish “filtered point cloud” data on a topic. A node can subscribe to a topic and when there is new data on that topic, it will trigger a function, called *subscriber callback*, to handle the data. For example, the “point cloud filter” node subscribes to LIDAR data, and the filtering function is triggered by a new data coming from the LIDAR. Further, a node can run a *timer callback* that is triggered by a periodic timer signal. ROS2 also offers blocking and non-blocking remote procedural calls (RPC) via *services* and *actions*, respectively. Here, a function in a server node can be directly invoked from a client node. Services are implemented using topics, i.e., an RPC is made by publishing data on a *request* topic, and the results of the RPC are written on a *response* topic. We do not study actions as they do not fit well into real-time software implementations.

Overall, in software developed following ROS2 semantics, we can find timer, subscriber, service, and client callbacks. We assume that a *single-threaded executor* dispatches all callbacks in a node, which is also studied in existing works on ROS2 timing analysis [2]–[6]. In this case, a thread executes one callback at a time from start to end before looking again into the queue with ready-to-run callbacks [2]. For each callback cb_k , we measure its *execution time* $et_{k,u}$ at each invocation $u \in \mathbb{N}$. In Sec. IV, we provide a technique to measure it.

Data communication via a topic occurs through multiple abstraction layers in ROS2. For this work, we use ROS2 Foxy. Officially-supported client libraries like *rclcpp* and *rclpy* comprise APIs (application programming interfaces) to develop ROS2 applications in C++ and python, respectively. These libraries use the core C interfaces provided by the *rcl* library which implements ROS2 concepts. All communications

between ROS2 nodes are carried out using DDS mechanisms. We have used Eclipse Cyclone DDS by selecting the appropriate *rmw* (ROS MiddleWare) interface library. DDS layer is directly over Linux version 5.4 in our implementation. ROS2 layered architecture is illustrated in [7], [12], while [13] shows the communication between two nodes through these layers.

B. extended Berkeley Packet Filter (eBPF)

eBPF technology enables to run validated programs in kernel space without the need to recompile the kernel. *BPF Compiler Collection* (BCC) tool allows writing eBPF programs in restricted C that are converted into bytecode using Low-Level Virtual Machine (LLVM) Clang compiler. The bytecode is loaded into the kernel using *bpf()* syscall. eBPF programs can be attached to the user- and kernel-space functions using uprobes and kprobes, respectively, thereby enhancing system observability. In this work, we have used BCC version 0.26.0 and LLVM-clang version 10.0.

The primary reasons to use eBPF instead of LTTng (as in *ros2_tracing* [7]) are as follows: (i) Using eBPF, we can attach probes to the entry and exit of ROS2 functions and read their arguments and return values, respectively, without modifying the standard libraries. Unlike with LTTng, there is no need for direct instrumentation and recompilation of ROS2 libraries. [9] has proposed to use *LD_PRELOAD* to redirect function calls to its tracing library instead of ROS2 libraries and then calls original ROS2 functions from there. However, this requires running several lines of code to update addresses to find the original functions, which adds significant tracing overheads without any additional capabilities. Also, such a redirection is not possible when a function is defined in the header. (ii) Trace filtering is more efficient in eBPF [14], which we use to obtain kernel events specific to ROS2 nodes only. Further, it is possible to filter events pertaining to one or more ROS2 nodes, which is useful for quick debugging. (iii) eBPF is safer and offers more programmability [10] that can be explored in the future, e.g., dynamic scheduling of ROS2 nodes to improve timing performance and secure execution of ROS2 callbacks.

III. PROPOSED TRACE POINTS

A. Trace points in ROS2

Using eBPF, we have attached *uprobes* and *uretprobes* to several ROS2 functions as shown in Table I. We distinguish between callbacks (CB) or topics during a ROS2 function call using the arguments passed on to it. Here, the main challenge is to understand complex ROS2 data structures and identify functions to probe so that we get the events that show what is happening at the application layer. Table I lists the information we extract using each probe. Due to limited space, we do not elaborate on how we traverse through the function arguments. We list below several important aspects of our tracer and concepts we have used while adding tracepoints.

- Using Linux *perf* tool, we obtain call graphs for ROS2 nodes while running different types of CBs. This helps to identify ROS2 functions that are called and can be probed.

TABLE I: Inserted probes in ROS2 Foxy.

No.	ROS2 lib	Function	Params/Purpose
\mathbb{P}_1	rmw_ cyclone dds_cpp	rmw_create_ node	shows node name and the PID of the thread that will execute node's CBs
\mathbb{P}_2	rclcpp	execute_timer	notifies timer CB starts
\mathbb{P}_3	rcl	rcl_timer_call	shows timer CB ID
\mathbb{P}_4	rclcpp	execute_timer	notifies timer CB ends
\mathbb{P}_5	rclcpp	execut_ subscription	notifies subscriber CB starts
\mathbb{P}_6	rmw_ cyclone dds_cpp	rmw_take_int	notifies a read event on a topic and shows subscriber CB ID, topic name and source timestamp of data
\mathbb{P}_7	message_ filters	operator	shows that a subscriber CB is used for data synchronization
\mathbb{P}_8	rclcpp	execute_ subscription	notifies subscriber CB ends
\mathbb{P}_9	rclcpp	execute_service	notifies service CB starts
\mathbb{P}_{10}	rmw_ cyclone dds_cpp	rmw_take_ request	notifies a service request received event and shows service CB ID, service name, and source timestamp of request
\mathbb{P}_{11}	rclcpp	execute_service	notifies service CB ends
\mathbb{P}_{12}	rclcpp	execute_client	notifies client CB starts
\mathbb{P}_{13}	rmw_ cyclone dds_cpp	rmw_take_ response	notifies a service response received event and shows client CB ID, service name, and source timestamp of response
\mathbb{P}_{14}	rclcpp	take_type_ erased_response	notifies if a client CB will be dispatched
\mathbb{P}_{15}	rclcpp	execute_client	notifies client CB ends
\mathbb{P}_{16}	cyclone dds	dds_write_impl	notifies a write event on a topic and shows the topic name and source timestamp of data/request/response

- Each event generated by our probes comprises (i) a timestamp for chronological ordering, (ii) a process ID (PID) to associate the event to a ROS2 node, and (iii) a probe name to indicate the type of information we can get from the event.
- We attach *uprobes* $\{\mathbb{P}_2, \mathbb{P}_5, \mathbb{P}_9, \mathbb{P}_{12}\}$ and *uretprobes* $\{\mathbb{P}_4, \mathbb{P}_8, \mathbb{P}_{11}, \mathbb{P}_{15}\}$ to *execute_** {timer, subscription, service, client} functions to get start and end times of CBs.
- We devise a technique to read the source timestamp (*srcTS*) of a particular data in a topic by probing *rmw_take_** {int, request, response} both at entry and exit $\{\mathbb{P}_6, \mathbb{P}_{10}, \mathbb{P}_{13}\}$. *srcTS* is passed by reference to this function and the value is not known when we enter the function since it is determined by calling other low-level DDS functions. Hence, we store the address of *srcTS* at the function entry in a BPF map and read the value from the stored address at the function exit.
- In ROS2, a service can be invoked from multiple client nodes. The response for a particular service request is sent to each client node and, hence, we get events for \mathbb{P}_{12} , \mathbb{P}_{13} , and \mathbb{P}_{15} . However, the client CB is dispatched only in the caller node and to distinguish that we probe *take_type_erased_response* at exit using \mathbb{P}_{14} and read its return value. If it returns 1 then the client CB will be dispatched.
- We have noticed a library named *message_filters* in AVP. It provides APIs to synchronize data arriving at a node in different topics—used for sensor fusion. We attach a probe \mathbb{P}_7 to identify a subscriber CB used for data synchronization, i.e., the probed function runs every time data is read from the topic and it needs to be synchronized.

Algorithm 1 Extract callback attributes for a ROS2 node

Require: *PID* of the ROS2 node, *ROSEvents*, *SchedEvents*

```

1: CBlist = []
2: for event in ROSEvents(PID).SortByTime() do
3:   if event.type is CB_start ( $\mathbb{P}_2/\mathbb{P}_5/\mathbb{P}_9/\mathbb{P}_{12}$ ) then
4:     record CB.type  $\triangleright \mathbb{P}_2, \mathbb{P}_5, \mathbb{P}_9, \mathbb{P}_{12}$  denotes different CB type
5:     CB.start = event.time
6:   else if event.type is timer_call ( $\mathbb{P}_3$ ) and CB.start  $\neq \emptyset$  then
7:     CB.ID = event.ID
8:   else if event.type is take ( $\mathbb{P}_6/\mathbb{P}_{10}/\mathbb{P}_{13}$ ) and CB.start  $\neq \emptyset$  then
9:     CB.ID = event.ID
10:    if event.type is take_response ( $\mathbb{P}_{13}$ ) then
11:      CB.intopic = cat(event.topic, CB.ID)
12:    else if event.type is take_request ( $\mathbb{P}_{10}$ ) then
13:      CB.intopic = cat(event.topic, FindCaller(event, ROSEvents))
14:    else
15:      CB.intopic = event.topic
16:  else if event.type is dds_write ( $\mathbb{P}_{16}$ ) and CB.start  $\neq \emptyset$  then
17:    if event.topic is a service_request then
18:      top_out = cat(event.topic, CB.ID)
19:    else if event.topic is a service_response then
20:      top_out = cat(event.topic, FindClient(event, ROSEvents))
21:    else
22:      top_out = event.topic
23:    append top_out to CB.outtopic
24:  else if event shows will_not_dispatch_client ( $\mathbb{P}_{14}$ ) then
25:    CB.* =  $\emptyset$ 
26:  else if event shows sync_subscribe ( $\mathbb{P}_7$ ) and CB.start  $\neq \emptyset$  then
27:    CB.isSyncSubscriber = TRUE
28:  else if event.type is CB_end ( $\mathbb{P}_4/\mathbb{P}_8/\mathbb{P}_{11}/\mathbb{P}_{15}$ ) and CB.start  $\neq \emptyset$  then
29:    CB.end = event.time
30:    CB.ET = GetExecTime(CB.start, CB.end, PID, SchedEvents)
31:    CBlist.AddToCallback(CB)
32:    CB.* =  $\emptyset$ 
33: return CBlist

```

B. Trace points in the OS scheduler

We attach a tracepoint to *sched_switch*—an OS event that notifies when the scheduler gives a CPU to a new thread. Hence, from such an event, we get (i) the CPU where the switch is happening; (ii) the PID and the scheduling priority of both previous and new threads; and (iii) the state of the previous thread when the switch occurs. These events are necessary to measure a CB's execution time, considering that between its start and end (notified by ROS2 events), the thread running it might get preempted or need to wait for data or a signal. Also, we get the configured CPU affinity and scheduling priority of ROS2 nodes from these events. We note that if we record all *sched_switch* events, the memory footprint of the trace data will be too high, e.g., hundred megabytes of data per second. We reduce the memory footprint by an order of three or more by filtering these events based on the PIDs of ROS2 nodes. We get these PIDs using the probe \mathbb{P}_1 that are then shared with *sched_switch* event handler using BPF maps.

IV. TIMING MODEL SYNTHESIS

As mentioned in Sec. II-A, we use a single-threaded ROS2 executor. Hence, all events between a *callback_start* event and the next *callback_end* event related to a particular ROS2 node—identified by its PID—provide information about one execution of a specific CB, also referred to as a CB instance hereafter. Exploiting the above, we propose Alg. 1 to identify all CBs in a ROS2 node and extract their architectural and timing attributes. The algorithm requires: (i) the node's *PID*, (ii) all ROS2 events, *ROSEvents*, and (iii) all *sched_switch* events, *SchedEvents*. It traverses all ROS2 events pertaining

to the given *PID* in the chronological order (lines 2 – 33) and collects information as follows:

- It gets the CB's type and its start time from a *CB_start* event (lines 3 – 5). The type is identified from the probe name, e.g., a \mathbb{P}_2 's event notifies that a *timer* CB starts.
- It gets the CB's ID (lines 6 – 9) (i) from a *timer_call* event for a timer CB and (ii) from a *take* event for other CBs.
- The subscribed topic (if any) is obtained from a *take* event (lines 10 – 15).
 - For a client CB, the service response is read from the topic. To distinguish between responses to different clients, the CB's ID is *concatenated* to the topic name (lines 10 – 11).
 - For a service CB, the request is read from the topic. To distinguish between requests from different callers, the caller's ID is *concatenated* to the topic name (lines 12 – 13). To identify the caller, the algorithm first finds and saves the *PID* of the *dds_write* event with the same topic and source timestamp as *take*. Then, it finds the *timer_call* or *take* event with the same *PID*, chronologically preceding *dds_write* and after the last *CB_start*. This event then provides the ID of the caller CB.
- The published topic (if any) is obtained from a *dds_write* event (lines 16 – 23).
 - If the topic is used to request a service, it *concatenates* the CB's ID to the topic name (lines 17 – 18). When there are multiple callers of the same service, this helps to uniquely relate the topic to the caller CB.
 - If the topic is for service response, it finds the client CB and *concatenates* its ID to the topic name (lines 19 – 20). For this, the algorithm finds the *take_response* event with the same topic and source timestamp as *dds_write*. There can be n_{cl} such events when there are n_{cl} clients of the same service. For each such event, the algorithm finds the chronologically next *take_type_erased_response* event with the same *PID* to evaluate if the client CB will be dispatched. Now, in case it is evaluated to be true, the client CB ID is obtained from the *take_response* event.
- If a client CB is not dispatched (line 24)—as shown by the *take_type_erased_response* event—the information about the current instance is not stored (line 25).
- A subscriber CB is marked appropriately if it is used for data synchronization—as identified by a *sync_subscribe* event (lines 26 and 27).
- The end time of a CB instance is obtained from a *CB_end* event (lines 28 – 29). Also, its *execution time* is calculated using Alg. 2 (line 30).

At the end of a CB instance, the stored information is added to *CBlist*, a list of CBs (line 32). A new entry is added only if none of the existing entries identifies the same CB as the current instance otherwise the execution time is recorded and the published topic list is updated (if a new topic is encountered) in the matching entry. For all CBs except a service, the ID is used for matching, while for a service, both the ID and subscribed topic are matched. Note that, for

Algorithm 2 Compute execution time—*GetExecTime(...)*

Require: *start, end, PID, SchedEvents*

```

1: ExecTime = 0
2: last_start = start
3: for event in SchedEvents.SortByTime do
4:   if start < event.time < end then
5:     if event.prev_pid = PID then
6:       ExecTime = ExecTime + (event.time – last_start)
7:     else if event.next_pid = PID then
8:       last_start = event.time
9:     else if event.time > end then
10:      ExecTime = ExecTime + (end – last_start)
11:   return ExecTime

```

a service CB in *CBlist*, the updated topic name identifies a service request by a particular caller. After updating *CBlist*, all information related to the CB instance is deleted (line 32).

Alg. 1 returns the list of CBs and their attributes (line 34). Note that the statistical information (including the *worst-case*, the *best-case*, and the *average* measured values) can be obtained using the stored values of execution times of each CB instance. Also, for a timer CB, the difference between consecutive start times gives its approximate *period* of invocation.

Execution time measurement: Alg. 2 measures the execution time of a CB instance. It combines ROS2 and *sched_switch* events. It requires: (i) the *start* and the *end* times of the CB instance (obtained from ROS2 events); (ii) the *PID* of the ROS2 node that identifies the thread \mathcal{T} executing the CB; and (iii) all *sched_switch* events, *SchedEvents*. It finds all execution segments between the start and the end of the CB instance by chronologically traversing the events in *SchedEvents* and sum up their time lengths to compute the execution time (lines 3 – 11). We note that when *CB_start* event is generated, \mathcal{T} is running so the first execution segment starts with this event, hence, line 2. Following a similar reason, the last execution segment of the CB instance ends when we get the *CB_end* event, hence, lines 9 – 11. Further, if the previous thread in a *sched_switch* event is \mathcal{T} , it marks the end of an execution segment (lines 5 – 6). Conversely, if the next thread is \mathcal{T} then it marks the start of an execution segment (lines 7 – 8).

DAG synthesis: Based on *CBlists* generated by Alg. 1 for all ROS2 nodes, we can synthesize their timing model as follows:

- Each CB in a *CBlist* is a vertex in the DAG. For n_{cl} callers of a service, we have n_{cl} vertices because there are that many entries in *CBlist*. This is necessary because the same service is then part of multiple non-intersecting computation chains which shall be appropriately modeled.
- If the subscribed topic of a CB cb_k matches a published topic of another CB $cb_{k'}$, we draw an edge from $cb_{k'}$ to cb_k except when cb_k is part of data synchronization. Adding edges to and from a service CB follows the same rule as the request and response topic names are also updated for the caller and the client CBs.
- We have a divergence in the DAG when (i) a CB publishes on more than one topic or (ii) it publishes on a topic subscribed by more than one CB.
- If $cb_{k'}$ and $cb_{k''}$ publish on a topic that is subscribed by cb_k then we mark an 'OR' junction at cb_k , i.e., cb_k is triggered when either of $cb_{k'}$ and $cb_{k''}$ publishes data on the topic.

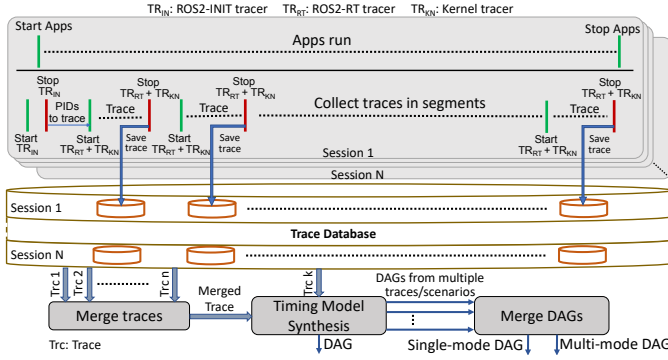


Fig. 2: Deployment of tracing and model synthesis framework.

- Two or more CBs in a ROS2 node can be used for data synchronization MS_α as marked in $CBlist$. The output of MS_α is computed only when new data has arrived on each of the subscribed topics of the CBs. Here, we add a vertex $cb_k^{\&}$ —mark as an ‘AND’ junction—with incoming edges from all CBs in MS_α . We draw an outgoing edge from $cb_k^{\&}$ to each subscriber CB of a topic on which the CBs in MS_α publish data. Note that when the input data to a CB in MS_α never arrives last during the synchronization, no published topic is found in the corresponding entry in $CBlist$.

V. DEPLOYMENT OF OUR PROPOSED FRAMEWORK

As shown in Fig. 1, we have built three tracers with the eBPF trace points. (i) ROS2-INIT tracer (TR_{IN}) uses \mathbb{P}_1 (see Table I) and traces the initialization of the ROS2 nodes. (ii) ROS2-RT tracer (TR_{RT}) uses $\mathbb{P}_2 - \mathbb{P}_{16}$ (see Table I) and records ROS2 events when the nodes are running. (iii) Kernel tracer (TR_{KN}) records *sched_switch* events for ROS2 nodes as identified by TR_{IN} . Fig. 2 shows how these tracers can be used in practice. TR_{IN} is activated before applications are started and it can be stopped after the initialization. After TR_{IN} identifies the threads of ROS2 nodes, TR_{RT} and TR_{KN} are activated to record the runtime ROS2 and scheduler events.

In this paper, we suggest synthesizing timing models based on measurements instead of formal worst-case execution time (WCET) analysis considering that the latter does not always scale for industry-strength autonomous applications. Hence, for accurate modeling, we need to collect large amounts of traces across several runs of the applications and our framework is compatible with this. Even when one run is long and the trace buffers are limited in size, we can stop TR_{RT} and TR_{KN} , store the traces in a database server, and then restart TR_{RT} and TR_{KN} with empty buffers. In the end, we might have a large number of traces in the server collected during multiple tracing sessions, as shown in Fig. 2.

There are several possibilities in which we can process the traces, as depicted in Fig. 2. (i) We can merge all traces and run our algorithms for DAG generation on the merged trace. (ii) We can generate a DAG for each trace and then they are merged, i.e., the vertices and edges in the final DAG will be a union of all input DAGs while for a CB’s execution time, the measured worst-case, best-case, and average values are calculated considering all input DAGs. We take this approach

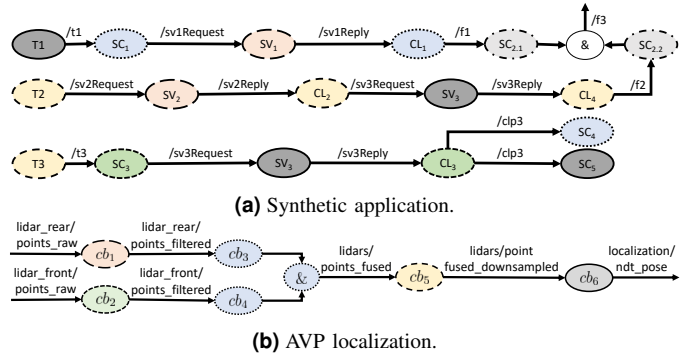


Fig. 3: Callbacks and the precedence relations between them.

for our experiments in this paper. (iii) We may have a combination of (i) and (ii), e.g., we merge traces collected over one run, generate DAG for each run, and then merge DAGs across multiple runs. (iv) Interestingly, it is also possible to generate multi-mode DAGs, e.g., if traces are merged per mode (or scenarios) such as city and highway driving.

VI. A CASE STUDY

Experimental Setup: Our setup comprises ROS2 Foxy, Eclipse Cyclone DDS, and Linux kernel 5.4.1 running on AMD Ryzen 9 3900X CPU with 46 GB RAM. We run two applications concurrently: (i) a part of Autoware’s AVP performing localization [1] and (ii) a synthetic application (*SYN*). To run (i), we directly use the *localization demo* provided by Autoware where the application runs for 80 seconds during which a car starts driving from a stationary position, parks, and then drives back to the same initial position. Further, we have developed *SYN* using six ROS2 nodes with different combinations of timer, subscriber, service, and client CBs. We run these applications 50 times, apply our DAG synthesis algorithms on traces collected for each run, and then merge these DAGs together, as explained in Sec. V.

Timing models: Fig. 3a shows the obtained DAG for *SYN*. CBs belonging to the same node in an application are marked with a distinct color and border. Also, the edges are annotated with the topic names. *SYN* covers different scenarios that are correctly identified by our framework. (i) It distinguishes between CBs of the same type in a ROS2 node, e.g., T2 and T3 are timer CBs; SC1 and SC4 are subscriber CBs; SV1 and SV2 are service CBs; and CL2 and CL4 are client CBs. (ii) It identifies different types of CBs in a node, e.g., T1, SC5, and SV3 are timer, subscriber, and service CBs, respectively. (iii) It finds multiple subscribers of a topic, e.g., /clp3 is subscribed by SC4 and SC5. (iv) It finds invocation of a service from different CBs, e.g., SV3 is invoked from SC3 and CL2. Also, we have two vertices for SV3, which is how we propose to model such a case. Imagine, if we had only one vertex for SV3 with incoming edges from SC3 and CL2 and outgoing edges to CL3 and CL4, then $SC3 \rightarrow SV3 \rightarrow CL4$ forms a sub-chain, which is incorrect. (v) It identifies data synchronization, e.g., data from /f1 and /f2 are synchronized and the output is /f3. For timing analysis, the vertex marked as “&” is a task that has a zero execution time and is an “AND” junction.

TABLE II: Execution times (in ms) of callbacks in AVP localization.

CB	Node	mBCET	mACET	mWCET
<i>cb</i> ₁	filter_transform_vlp16_rear	13.82	17.1	19.82
<i>cb</i> ₂	filter_transform_vlp16_front	23.31	27.07	30.5
<i>cb</i> ₃	point_cloud_fusion	0.41	3.1	3.97
<i>cb</i> ₄	point_cloud_fusion	0.38	0.62	3.36
<i>cb</i> ₅	voxel_grid_cloud_node	6.58	8.47	13.36
<i>cb</i> ₆	p2d_ndt_localizer_node	2.78	25.64	60.93


Fig. 3b shows the obtained DAG performing localization in AVP. It consists of 6 CBs in 5 ROS2 nodes. Node names are provided in Table II. The DAG shows that *cb*₁ and *cb*₂ are in different nodes and *filter LIDAR rear* and *front raw data*, respectively. The *filtered* data are sent to a *fusion* node where *cb*₃ and *cb*₄ are used for data synchronization, as identified by our framework. The *fused* data is *downsampled* by *cb*₅ in a *voxel_grid_cloud* node. Finally, *cb*₆ in a *p2d_ndt_localizer* node implements the localization algorithm on the *downsampled* data to determine the vehicle’s *pose*. Here, the DAG has only subscriber CBs and an “&” junction in the *fusion* node. Table II further reports the measured best-case, average, and worst-case execution times (mBCET, mACET, and mWCET, respectively) of each of the 6 CBs. They are measured over 50 runs. Here, the most computationally expensive CB, *cb*₂, has an average processor load of 27%—the LIDAR data arrives at 10 Hz. Such measurements are useful even for simple debugging and optimization, e.g., balancing load across processor cores or keeping the load below a certain threshold while determining core bindings of ROS2 nodes.

For each CB in *SYN*, we have used a constant computational load for a single run. By comparing the measured with the designed execution times, we have validated our framework’s ability to measure accurately. We change the load of each CB in *SYN* across runs to evaluate if the execution time profiles of AVP’s CB are sensitive to varying interfering loads. We note that for *SYN*, the measurement results over 50 runs are not important so we do not report them. Further, Fig. 4 shows that mACET and mBCET change negligibly for *cb*₂ with increasing runs, while mWCET increases by 10% over 23 runs and thereafter remains unchanged. Such an evolution of mWCET shows that our modeling accuracy improves with more traces. We point out that, if we assume that test cases can be generated with a high coverage, our framework can support accurate model synthesis using tracing and measurement.

Tracing overheads: For an experiment where we run *SYN* and *AVP localization* together for 60s, (i) we generate 9MB of trace data; and (ii) *bpftool* shows that our eBPF probes use 0.008 CPU cores on average, which is 0.3% of the computational load produced by the applications.

VII. CONCLUDING REMARKS

In this paper, we have presented a framework to trace ROS2-based autonomous applications using eBPF. It further implements algorithms to process traces and synthesize timing models of applications. Our framework can be extended to support other ROS2 versions and DDS implementations, while the proposed concepts can be applied to other software architectures (e.g., AUTOSAR) and operating systems (e.g., QNX).

**Fig. 4:** Estimation of timing attributes improve with more traces.

In the future, we would use the framework for debugging and optimization. We are logging the source timestamp of data on publisher and subscriber sides using which we can traverse data flow through a computation chain and calculate its end-to-end latency. We can add a tracepoint to *sched_wakeup* and compute the waiting time of a callback. Such measurements will help to debug the timings of a chain and optimize it as per requirements. Also, along the lines of [15], we can change the schedule configuration of ROS2 nodes online using eBPF to improve systems’ timing performance.

REFERENCES

- [1] Autoware. Autonomous Valet Parking Demonstration. [Online]. Available: <https://autowarefoundation.gitlab.io/autoware.auto/AutowareAuto/avpdemo.html>
- [2] D. Casini, T. Blaß, I. Lütkebohle, and B. Brandenburg, “Response-time analysis of ROS 2 processing chains under reservation-based scheduling,” in *Euromicro Conference on Real-Time Systems (ECRTS)*, 2019.
- [3] Y. Tang *et al.*, “Response time analysis and priority assignment of processing chains on ROS2 executors,” in *IEEE Real-Time Systems Symposium (RTSS)*, 2020.
- [4] H. Choi, Y. Xiang, and H. Kim, “PiCAS: New design of priority-driven chain-aware scheduling for ROS2,” in *IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, 2021.
- [5] A. A. Arafat, S. Vaidhun, K. M. Wilson, J. Sun, and Z. Guo, “Response time analysis for dynamic priority scheduling in ROS2,” in *ACM/IEEE Design Automation Conference (DAC)*, 2022.
- [6] T. Blaß, D. Casini, S. Bozhko, and B. B. Brandenburg, “A ROS 2 response-time analysis exploiting starvation freedom and execution-time variance,” in *IEEE Real-Time Systems Symposium (RTSS)*, 2021.
- [7] C. Bédard, I. Lütkebohle, and M. Dagenais, “ros2_tracing: Multipurpose low-overhead framework for real-time tracing of ROS 2,” *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 6511–6518, 2022.
- [8] Z. Li, A. Hasegawa, and T. Azumi, “Autoware_Perf: A tracing and performance analysis framework for ROS 2 applications,” *Journal of Systems Architecture*, vol. 123, 2022.
- [9] T. Kuboichi *et al.*, “CARET: Chain-aware ROS 2 evaluation tool,” in *IEEE International Conference on Embedded and Ubiquitous Computing (EUC)*, 2022.
- [10] L. Rice, *Learning EBPf: Programming the linux kernel for Enhanced Observability, networking, and security*. O’Reilly Media, 2023.
- [11] H. Teper, M. Günzel, N. Ueter, G. von der Brüggen, and J.-J. Chen, “End-to-end timing analysis in ROS2,” in *IEEE Real-Time Systems Symposium (RTSS)*, 2022.
- [12] M. Hansen, A. Blasdel, and C. Buscaron. ROS 2 Foxy Fitzroy: Setting a new standard for production robot development. [Online]. Available: <https://aws.amazon.com/blogs/robotics/ros-2-foxy-fitzroy-robot-development/>
- [13] T. Kronauer, J. Pohlmann, M. Matthé, T. Smejkal, and G. P. Fettweis, “Latency analysis of ROS2 multi-node systems,” *2021 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems (MFI)*, pp. 1–7, 2021.
- [14] S. D. Sharma and M. Dagenais, “Enhanced userspace and in-kernel trace filtering for production systems,” *Journal of Computer Science and Technology*, pp. 1161–1178, 2016.
- [15] T. Blass, A. Hamann, R. Lange, D. Ziegenbein, and B. B. Brandenburg, “Automatic latency management for ROS 2: Benefits, challenges, and open problems,” in *IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, 2021.