PX4: A Node-Based Multithreaded Open Source Robotics Framework for Deeply Embedded Platforms

Lorenz Meier, Dominik Honegger and Marc Pollefeys

Abstract-We present a novel, deeply embedded robotics middleware and programming environment. It uses a multithreaded, publish-subscribe design pattern and provides a Unixlike software interface for micro controller applications. We improve over the state of the art in deeply embedded open source systems by providing a modular and standards-oriented platform. Our system architecture is centered around a publishsubscribe object request broker on top of a POSIX application programming interface. This allows to reuse common Unix knowledge and experience, including a bash-like shell. We demonstrate with a vertical takeoff and landing (VTOL) use case that the system modularity is well suited for novel and experimental vehicle platforms. We also show how the system architecture allows a direct interface to ROS and to run individual processes either as native ROS nodes on Linux or nodes on the micro controller, maximizing interoperability. Our microcontroller-based execution environment has substantially lower latency and better hardware connectivity than a typical Robotics Linux system and is therefore well suited for fast, high rate control tasks.

I. INTRODUCTION

Micro aerial vehicles (MAVs) have been an active research topic for decades, but became of even more interest to the broader robotics community in recent years. While ground-based robots are deployed in applications today, they remain confined to particular work spaces. Micro air vehicles navigate much more freely outdoors, and due to their 3D motion are a suitable generalization for many robotics design challenges. Recently a number of large scale industry applications have been proposed which require higher levels of autonomy than deployed systems can offer today. As the field progresses and more sophisticated robotics problems are targeted, the software complexity of these systems increases rapidly. Similar to ground robotics recently, aerial robotics is reaching a level of complexity where individual research groups are unable to tackle the system design on their own. We address this problem by providing a very flexible, standards oriented and low-cost research platform supporting MAVs in the micro size scale but still offering a multithreaded industry standard programming environment. This allows path planning and other high-level navigation and control research topics to benefit from a fully featured Unix system even on platforms too small to carry a companion computer. With an object request broker (ORB) and a standardized API, it can be characterized as Robot Operating System (ROS) equivalent for small aerial or any other size and power constrained systems. The platform has also been successfully utilized in conjunction with Linux companion computers, when more processing power is necessary.

The contributions of this work are as follows: First we first discuss related works concerning existing deeply embedded systems. Second, we introduce the design criteria for such a system. Third, the general software architecture and implementation is presented. Last we show performance results and provide some possible applications for the framework.

II. RELATED WORK

As robotics research is currently tackling a wide range of problems from estimation and control to high-level planning, communication and middleware systems have emerged to support the increased complexity. Huang et al. created a publish-subscripe design pattern object request broker for robotics systems [1] which is particularly lightweight and low latency. Quigley et al. designed a complete robotics meta operating system which not only includes a middleware layer providing different communication methods including publish-subscribe patterns, but also a library of robotics packages [2]. Our previous work was concerned with a different layer and application domain and was not focused on the deeply embedded scope. It leveraged LCM [1] as middleware and added a set of computer-vision specific interfaces and algorithms [3]. While the presented work has much stronger similarities to Linux / Unix based robotics systems in terms of architecture or features, the most prominent existing research system in the micro air vehicle domain is PPZ [4][5]. It features a complete fixed wing and rotary autopilot suite, but does not support native multithreading or a native ROS integration. The authors of APM [6] have created a complete autopilot system for fixed wing, multicopters and traditional helicopters. While the system shows excellent flight performance, it utilizes an internal function-scheduling routine without support for preemption. It does support PX4 hardware and middleware and leverages the multithreading capabilities of the platform by running background worker threads. However the design, whilst using worker threads, does not decouple individual tasks and does not implement a general inter process communication scheme. Instead it utilizes different main loops for different vehicle types. The OpenPilot platform [7], which focuses on multicopters, implements a fully multithreaded solution, but uses an API that is custom to the operating system, therefore making the convergence and reuse between the deeply embedded and Unix platforms more challenging. One design feature of the platform is the linkage of inter process message formats to telemetry, which is making interoperability challenging. For

The authors are with the Computer Vision and Geometry Group, Institute for Visual Computing, Computer Science Department, ETH Zurich, 8092 Zurich, Switzerland lm@inf.ethz.ch

a complete discussion of open source systems, please refer to the excellent overview in [8]. Apart from pure onboard systems, hybrid onboard and off board research systems are widely used today. One of the first instances of this class, the system of Lupashin et al.[9] uses a hybrid system. It implements a custom off-board trajectory generation and position control, which is combined with on-board attitude and rate control. The architecture of Michael et al. [10] is very similar, but the system operates on top of the Robot Operating System (ROS) infrastructure.

III. DESIGN CRITERIA FOR DEEPLY EMBEDDED SYSTEMS

In this section we describe the design criteria for deeply embedded systems. We focus on the real-time aspects of the system. This is mostly defined by requirements for lowlatency sensor acquisition and control responses and specialized low level interfaces, such as I2C, SPI, CAN and PWM outputs. The PX4 software architecture is modeled to address typical estimation and control tasks in deeply embedded platforms with a modular approach. We describe the common layered control architecture of robotic platforms, reusability considerations, and interoperability concerns.

A. Controller Timing Requirements

Robot controllers can generally be modeled as a set of nested control loops. Each control loop has a reference set point and the current vehicle state as input. It generates the reference for the next inner loop. Even more advanced control structures often can be described as a set of nested control loops. The outer loops generally have less strict timing requirements compared to the innermost control loops, thus giving the system designer more flexibility on which platform to implement the outer layers. Fig. 1 illustrates the typical modeling and control of a generalised robot, following the notation of [10], but generalized to all vehicle types. For a complete example of this design scheme please refer to [11] for multi rotors and [12] for fixed wing aircraft.

B. Reusability

To date, when compared to open platforms, many research projects run custom hardware and software with very little testing and virtually no safety track record. An open and modular architecture is required to enable research groups to focus on their core research interests. We achieve a high level of reusability, while not sacrificing performance and safety by:

- Clear and clean layering of the software API from lowlevel to high-level controllers
- Multithreaded node-like architecture that decouples individual applications
- Availability of current sensors and actuator buses
- · Expansion bus on the hardware for new sensors
- · Safety hard-override in hardware at all times



Fig. 1: Independent of the vehicle type, the dynamical model and nested control architecture can be generalized. Different levels have different update rate and latency requirements.

C. Interoperability

The PX4 platform offers both compatibility and integration: The cross-platform API supports writing software packages which can be executed on the micro controller or on the ROS/Lnux system. It provides an interface between the micro controller based solutions and Linux companion computers to run them as a distributed system, leveraging the deeply embedded platform where real-time performance is necessary. The PX4 API allows to use the same code on ROS and on NuttX. For example the code in the listing below executes on both platforms and allows to build and run the same controller on the micro controller, on Linux, or in the Linux Software-in-the-Loop environment.

```
px4::Publisher<px4::rc_channels>
*pub = n.advertise<rc_channels>()
```

IV. SOFTWARE ARCHITECTURE

In the following section we describe the different layers of the developed software architecture. The software architecture is split into four main layers, as depicted in Fig 2: in the lower half, device drivers with device-specific code (e.g. for the particular microcontroller or bus type), and in the upper half, drivers which expose an interface for the system as device node. These are part of the operating system (NuttX [13]). The third layer is the micro object request broker which handles inter process communication efficiently and ensures data integrity between threads. The fourth layer is the application layer, which consists of individual applications (apps), such as flight control or state estimation modules.



Fig. 2: The different layers of the software architecture make the system horizontally and vertically modular.



Fig. 3: A single process can subscribe (consume) and publish multiple topics, allowing it to interface at different rates.

Our contribution consists of the PX4 middleware which provides devices drivers and a micro object request broker (μ ORB). The presented experimental results were obtained using the PX4 flight stack, which is a selection of estimators and controllers developed in close collaboration with the open source community. All hardware plans and the complete source code are available under a permissive BSD (software) and CC-BY-SA license (hardware) on the project website [14]. Although not formally certified, the system design is oriented towards several industry standards: The device drivers and operating system are modeled after the POSIX [15] interface standards. The off-board communication is using the commonly used MAVLink protocol [16]. The onboard networking is following the UAVCAN standard proposal [17].

A. µORB Middleware

The object request broker provides a data structure for data distribution. It follows the one-to-many publish-subscribe design pattern: A publisher wanting to share information advertises a topic. A topic is defined as a semantic message channel such as 'attitude' or 'position'. A subscriber can subscribe to a topic, and after the subscription is established ask at his own pace for new data (polling), or be woken from the thread sleep state at the instant new data is available. As Fig. 3 depicts a process can be both publisher and subscriber at the same time, and subscribe and publish to multiple topics.

Our implementation of this design pattern has particular strengths for realtime control applications:

- The topic handle is implemented as virtual file, allowing listeners to do blocking waits on interfaces and drivers (such as serial ports) and topics in parallel. This is commonly not supported by middleware solutions but saves a complete worker thread.
- The read-write lock of the publication allows efficient concurrency and ensures atomic reads and writes of the complete topic content.
- Subscribers can ask for a notification limit, allowing a subscriber to receive the topic only every N milliseconds. This is important for the efficiency of high-rate topics such as the 1KHz accelerometer updates.
- The asynchronous / blocking wait approach combined with the task priority setup of the operating system allows for minimal latency and deterministic scheduling in the control pipeline. Low-priority tasks and highpriority real time control tasks can be mixed.

B. Applications

Each state estimator and controller in the PX4 stack is implemented as standalone application, which is started with a main() function and then subscribes / publishes to different topics. Applications can be started and stopped at runtime.

C. Work Queue and HRT Callbacks

For applications that repeatedly only execute one function, such as device drivers, three different work queues to execute callbacks are available: The low priority and high priority work queues and the high resolution timer (HRT) callbacks. The two work queues execute in the normal application context, while the HRT callbacks operate in interrupt context for time-critical functions. Work queue entries are part of the normal scheduling and can access all operating system interfaces, HRT callbacks should be kept as short as possible and only support a subset of the OS API calls. However, HRT callbacks can publish to μ ORB topics.

D. Companion Computer

As this system is designed as deeply embedded system, the average robotic application will also provide a companion computer, commonly running ROS on Linux [18][19][3]. We not only offer a ROS interface for feedback and control, but go one step further: Our framework supports the native operation of nodes originally designed for the autopilot on ROS. This is feasible as the node centered design of the deeply embedded solution has the same architecture as on a Linux platform. Therefore we also build our software-in-the-loop simulation based on the ROS native port.

Figure 4 shows the architecture of the joint deeply embedded + Linux setup. Some components are exclusive to one of the platforms, e.g. the actuator drivers on the embedded platform or e.g. a simultaneous localization and mapping pipeline on the Linux system. Nodes that suit both environments can be executed on either platform. This has the particular benefit of allowing a proven version of a controller to run on the safety-critical deeply embedded controller, while testing a new version or different implementation on



Fig. 4: Interface to the companion computer and examples of applications which can be executed on either platform.

the Linux companion computer. Instead of encapsulating and hiding ROS in our environment, we run native and standard ROS nodes on Linux, allowing researchers already familiar with ROS to easily adopt the embedded codebase without having to learn a new API. Furthermore any time-critical nodes can be run on the deeply-embedded platform in realtime, whereas a Linux system does not offer that capability without using a real-time kernel on dedicated hardware. Our embedded solution is therefore suitable for low level control and allows to upgrade an existing ROS system to achieve real-time performance.

V. IMPLEMENTATION

In this section we describe an efficient implementation of the time critical driver layer and then introduce the major hardware components. The software platform is highly portable, but has been implemented for these results on the PX4 FMU hardware, with the main hardware features listed below and a schematic display of the inputs and outputs provided in Figure 6. The system offers a serial terminal interface to monitor the system status (for example system load via the 'top' command). System startup is managed through a set of shell scripts and parameters, which allows the full customization of the system startup if required. A MAVLink-enabled centralized parameter storage provides an easy to use (and GUI-supported) management of controller and general system parameters. Even non experienced users can therefore setup a fully customized system. Parameters can be changed during flight and stored in permanent storage (FRAM).

A. Driver Layer

The system is clocked using a high resolution timer which supports callbacks on the interrupt level. These callback functions are ideally suited to read sensors efficiently and accurate at a high rate. The interrupt latency / jitter is below 4 microseconds. A wide variety of common peripherals is supported by the system, including MEMS sensors, external airspeed and pressure sensors, PWM, I2C and CAN motor controllers as well as PX4 specific peripherals, such as the flow sensor [20]. The current hardware supports up to 5



Fig. 5: PX4 FMU v1.7 open hardware controller board, measuring 50 x 36mm with a weight of 8g. The input/output board (or a custom research module) can be connected. The Pixhawk autopilot (FMU v2.4) is our 2nd generation.

serial ports, which can be used to communicate point-topoint or point-to-multipoint with radio modems. The support for commercial-off-the shelf digital radio control systems (S.BUS1/2, PPM, Spektrum) allows the use of low-cost but well-proven manual override solutions.

B. Hardware

The autopilot consists of two independent sub-modules, the flight management unit (FMU) and the input-output unit (IO). Fig. 5 shows the standalone flight management unit rev. 1.7 which is optimized for small scale research systems. We have made FMU + IO available as all-in-one board as well (Pixhawk). The main hardware features are:

- 168 MHz Cortex M4F, 256 KB RAM, 2 MB flash
- MPU6000 gyro/acc, L3GD20 gyro, LSM303D mag/acc
- 14 PWM (servo) outputs total (8 with hard override)
- Triple-redundant power supply inputs with failover
- 5 serial ports (2 with hardware flow control)
- 2 CAN ports, 1 I2C port, 1 SPI port, 3x ADC
- RC inputs: PPM, S.BUS1/2, DSM2/X, RSSI input

The rationale for the separation of both units is to allow a hard override to manual control using the safety processor (see Figure 6). This considerably improves safety, particularly in research setups.



Fig. 6: Input and output data streams of the system.



Fig. 7: Latency from publication of topics in one process to reception in a second process, exhibiting very low latency.

VI. RESULTS

In this section, we show first the performance of the platform while running an attitude estimator. We then give a reusability comparison among existing platforms and finally show the flexibility of our platform.

A. Performance

Despite running on a resource constrained system, our proposed architecture is capable of processing multiple sensors connected via different embedded bus systems at 1000 Hz or more each, with an average interrupt latency in the sensor readout of less than 4 microseconds. Context switches between tasks require only 25 microseconds. As depicted in Fig. 7 the inter process latency is very low. In contrast to other deeply embedded solutions, the sensor drivers directly publish to sensor topics, simplifying the development of higher-level nodes such as controllers, as no knowledge about the embedded interfaces is required. Compared to Linux, the interrupt latency and timings are lower and more consistent even under high load. The system design allows to log different sensors and controller outputs at variable rate when available (similar to a ROS bag), and Python and Matlab tools are provided to plot these logs for flight analysis or replay them in unit test harnesses for filter design.

B. Reusability

Recently a wide range of open aerial robotics platforms have become available. Due to their fast evolution, low cost and applicability to rovers and underwater vehicles, they represent general purpose robotic controllers. However, the reusability of these platforms depends on their modularity. Table I summarizes aspects relevant to the adaption and general reusability, including potential limits induced by the license. The BSD license does not limit the reuse in academic and industrial applications, while GPL licensed code is subject to certain restrictions. The column nodes describes whether one software module (e.g. a controller or estimator) is self-contained and can be easily exchanged against a different module without modifying the core system (equivalent to a ROS node). The column IPC describes if the system is multithreaded and offers a suitable generic inter process communication layer. The column ROS-IF (ROS interface) captures the ROS platform interface. The column ROS-N captures the native ROS support of flight control and guidance software. SITL refers to Software-in-the-loop, a pure software simulation mode. The license column describes the license model.

TABLE I: Platform reusability. The five selected platforms performed best out of all evaluated platforms.

System	Nodes	IPC	ROS-IF	ROS-N	SITL	Lic.
PX4	yes	yes	yes	yes	ROS	BSD
OpenPilot [7]	yes	yes	no	no	no	GPL
APM [6]	no	no	yes	no	yes	GPL
PPZ [4]	no	no	yes	no	yes	GPL
MultiWii [21]	no	no	yes	no	no	GPL

C. Flexibility

For research applications the ability to adapt the system to new vehicles and setups without requiring fundamental software or hardware changes is critical. Our platform has been utilised in various non-standard vehicle setups, including a novel spherical blimp design [22], but also in various other platforms across very diverse fields [23][24][25][26][27][28][29].

Here we present a vertical take-off and Landing (VTOL) vehicle setup as an example of a concrete system design, including a Linux companion computer. It is trivial to implement using the platform, despite no VTOL support designed into the system. As our platform design is airframeagnostic, it was possible to combine the fixed wing and multicopter controllers with a transition controller to obtain first flight results without changing the system architecture. As the evolution of the VTOL control logic progresses, it will later be trivial replace this transition scheme with a custom VTOL controller, as depicted in Figure 8. In fact the controllers can be hot-swapped in flight. As the controllers can be shifted between systems or even run in parallel, testflights are possible with experimental controllers in Linux, with proven controllers as fallback available on the deeply embedded system if the Linux system fails or the controllers implemented on it do not perform as desired.

VII. CONCLUSION

We have presented the first node architecture oriented, fully multithreaded modular robotics framework for deeply embedded platforms using a publish / subscribe design pattern. By design, our architecture is well suited for experimental setups. The ability to run individual nodes as native ROS nodes allows for a very high flexibility in the design process. Our system is highly extensible both in terms of hardware and software, from the addition of individual sensors to using alternate controllers, filters and estimators. Researchers also can just leverage the base system and run completely custom control and estimation pipelines. The



Fig. 8: Diagram of the PX4 VTOL architecture. The attitude controllers are present on the micro controller and Linux as they can be executed on either platform.

use of a POSIX programming model and the publish/subscribe API substantially lowers the barriers for researchers accustomed to ROS or similar non-embedded toolkits to implement deeply embedded solutions. The hardware of our platform is easily and internationally available as open hardware design from multiple vendors and has even sparked the creation of derivatives specialised for particular use cases.

ACKNOWLEDGMENT

As open source and open hardware platform PX4 has received numerous key contributions from the open source community, which are listed on our website. We would also like to thank all our open hardware partners, in particular 3D Robotics, who have adopted our designs (PX4FMU, PX4FLOW and Pixhawk) and made them physically available to the academic and open source communities.

REFERENCES

- A. S. Huang, E. Olson, and D. C. Moore, "LCM: Lightweight communications and marshalling," in *Intelligent robots and systems* (IROS), 2010 IEEE/RSJ international conference on. IEEE, 2010.
- [2] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "ROS: an open-source robot operating system," in *ICRA workshop on open source software*, vol. 3, 2009.
- [3] L. Meier, P. Tanskanen, L. Heng, G. H. Lee, F. Fraundorfer, and M. Pollefeys, "PIXHAWK: A micro aerial vehicle design for autonomous flight using onboard computer vision," *Autonomous Robots*, vol. 33, no. 1-2, pp. 21–39, 2012.
- [4] P. Brisset, A. Drouin, M. Gorraz, P.-S. Huard, and J. Tyler, "The paparazzi solution," *MAV2006, Sandestin, Florida*, 2006.
- [5] B. Remes, P. Esden-Tempski, F. Van Tienen, E. Smeur, C. De Wagter, and G. De Croon, "Lisa-s 2.8 g autopilot for gps-based flight of mavs," in *IMAV 2014: International Micro Air Vehicle Conference* and Competition 2014, Delft, 2014.
- [6] APM. (2014) Ardupilot. Website. [Online]. Available: http://www.ardupilot.com
- [7] OpenPilot. (2014) OpenPilot project. Website. [Online]. Available: http://openpilot.org
- [8] H. Lim, J. Park, D. Lee, and H. J. Kim, "Build your own quadrotor: Open-source projects on unmanned aerial vehicles," *Robotics & Automation Magazine, IEEE*, vol. 19, no. 3, pp. 33–45, 2012.
- [9] S. Lupashin, A. Schollig, M. Hehn, and R. D'Andrea, "The flying machine arena as of 2010," in *Robotics and Automation (ICRA), 2011 IEEE International Conference on*. IEEE, 2011, pp. 2970–2971.
- [10] N. Michael, D. Mellinger, Q. Lindsey, and V. Kumar, "The grasp multiple micro-uav testbed," *Robotics & Automation Magazine, IEEE*, vol. 17, no. 3, pp. 56–65, 2010.

- [11] D. Mellinger, N. Michael, and V. Kumar, "Trajectory generation and control for precise aggressive maneuvers with quadrotors," *The International Journal of Robotics Research*, vol. 31, no. 5, 2012.
- [12] S. Park, J. Deyst, and J. P. How, "A new nonlinear guidance logic for trajectory tracking," in AIAA Guidance, Navigation, and Control Conference and Exhibit, 2004, pp. 16–19.
- [13] G. Nutt *et al.* (2014) Nuttx. Website. [Online]. Available: http://www.nuttx.org
- [14] PX4. (2014) PX4 project. Website. [Online]. Available: http://px4.io
- [15] Road vehicles Controller area network (CAN) Part 1: Data link layer and physical signalling, ISO Std. ISO 11898-1, 2003.
- [16] MAVLink. (2014) MAVLink micro air vehicle protocol. Website. [Online]. Available: http://mavlink.org
- [17] UAVCAN. (2014) UAVCAN project. Website. [Online]. Available: http://uavcan.org
- [18] A. S. Huang, A. Bachrach, P. Henry, M. Krainin, D. Maturana, D. Fox, and N. Roy, "Visual odometry and mapping for autonomous flight using an rgb-d camera," in *International Symposium on Robotics Research (ISRR)*, 2011, pp. 1–16.
- [19] M. Achtelik, S. Weiss, and R. Siegwart, "Onboard IMU and monocular vision based control for MAVs in unknown in-and outdoor environments," in *Robotics and automation (ICRA), 2011 IEEE international conference on*. IEEE, 2011, pp. 3056–3063.
- [20] D. Honegger, L. Meier, P. Tanskanen, and M. Pollefeys, "An open source and open hardware embedded metric optical flow cmos camera for indoor and outdoor applications," in *Intl. Conf. Robotics and Automation*, 2013.
- [21] Multiwii. (2014) MultiWii project. Website. [Online]. Available: http://www.multiwii.com
- [22] M. Burri, L. Gasser, M. Kach, M. Krebs, S. Laube, A. Ledergerber, D. Meier, R. Michaud, L. Mosimann, L. Muri, *et al.*, "Design and control of a spherical omnidirectional blimp," in *Intelligent Robots and Systems (IROS), 2013 IEEE/RSJ International Conference on*. IEEE, 2013, pp. 1873–1879.
- [23] J. Goppert, A. Shull, N. Sathyamoorthy, W. Liu, I. Hwang, and H. Aldridge, "Software/hardware-in-the-loop analysis of cyberattacks on unmanned aerial systems," *Journal of Aerospace Information Systems*, vol. 11, no. 5, pp. 337–343, 2014.
- [24] Q. Lin, Z. Cai, J. Yang, Y. Sang, and Y. Wang, "Trajectory tracking control for hovering and acceleration maneuver of quad tilt rotor uav," in *Control Conference (CCC), 2014 33rd Chinese*. IEEE, 2014.
- [25] P. C. Hickey, L. Pike, T. Elliott, J. Bielman, and J. Launchbury, "Building embedded systems with embedded dsls," in *Proceedings* of the 19th ACM SIGPLAN international conference on Functional programming. ACM, 2014, pp. 3–9.
- [26] C. Paulson and A. Sobester, Bespoke Balloon Launched Sensorcraft for Atmospheric Research Missions. American Institute of Aeronautics and Astronautics, 2014/09/28 2013. [Online]. Available: http://dx.doi.org/10.2514/6.2013-1270
- [27] G. James, L. Mejias, et al., "Towards the development of 1-to-n human machine interfaces for unmanned aerial vehicles," in *Proceedings* of the 2014 International Conference on Unmanned Aerial Systems (ICUAS'14). IEEE, 2014, pp. 211–221.
- [28] T. Nägeli, C. Conte, A. Domahidi, M. Morari, and O. Hilliges, "Formation flight of micro aerial vehicles without infrastructure," in *Intelligent Robots and Systems (IROS), 2014 IEEE/RSJ International Conference on.* IEEE, 2014.
- [29] R. D'Andrea, F. Augugliaro, D. Brescianini, L. Gherardi, M. Hehn, M. Mueller, and R. Ritz. (2014) Flying machine arena. Website. [Online]. Available: http://www.flyingmachinearena.org/how-to/