Embedded Systems & IoT architecture

David Ngo

2021-02-25



Working with embedded systems since 2008, focus on embedded software development

About me



Mainly in medtech and industrial



Helping SKF to deliver Condition Monitoring products and be the market leader of rotating equipment performance.





SKF Condition Monitoring products

SKF Industry presence



Railway vehicles



Mining



Power Generation Steam Turbines



Paper Machines



Power Generation Steam Turbines



Cruising Ships





Wind

Printing

Agenda

- 1. What is IoT
- 2. IoT system infrastructure
- 3. IoT devices and choice of Operating System
- 4. Communication protocols
- 5. IoT products examples





Source: NIST



Source: NCTA

IoT System Infrastructure

- IoT Devices
- IoT Gateways
- Cloud & Big Data
- End-User Applications





IoT Devices

.

0

٠

•

0

Source: ARM

.

Embedded Systems

- SW & HW interconnected
- Constraints on HW
 - Power
 - Speed
 - Size
- HW cost vs SW cost
 - Volume
 - Need to get smart about software because that's where the costs are.

Figure 1 – Generalized Layers of Software



Architectural patterns

- Selecting the right architectural pattern will play a crucial role to the outcome of the project.
- Software architecture must be chosen wisely because once implemented it is not easy to change.
- Main forces
 - Cost
 - Time to Market
 - Quality



"Big ball of mud"

Open-source OS landscape

• Bare metal, no OS





- Embedded OS / RTOS
 - FreeRTOS, Mbed OS, Zephyr, etc.





- Full-featured OS
 - Linux, etc.



OS used in embedded development

Most Used	World	Americas	EMEA	APAC
Embedded Linux	31%	32%	31%	26%
FreeRTOS	27%	25%	24%	37%
Android	14%	12%	10%	26%

- Architecture style: Layer
 - Each layer is a group of modules that offer cohesive set of services.
 - Clear separation of concerns so that modules can be independently developed and maintained (portability, modifiability, reuse)
 - Layer n can use services of layer n-1 and not vice versa.
 - Bridging layers





- Software Architecture
 - Super loop
 - Event driven
 - State machine
- Typical task order
 - Read inputs
 - Process information (using state machine)
 - Output results
 - Delay until next loop execution
- Task examples
 - ADC samples
 - Blink LED
 - Read/write data to flash
 - Serial communication





- + Low complexity
- + Everything is processed in the super loop
- + Easy to see the which functions being executed and when

Quality

Attributes?

- + Timing is fixed at every loop, except for short interrupts (incoming events)
- + Functions do not execute at the same time, thus no race condition
- No function can execute too long to ensure response time
- Difficult to add complex functionality
 - e.g. pdf generation, TLS communication, UI, complex timeconsuming algorithms
 - Functions with long execution time could lead to heavy jitter and latency of other user inputs
 - Scheduling adjustment to fit new modules/libraries
 - Requires implementation with constraint resource



- When using bare metal?
 - Simple/ low complexity application
 - Real-time deadline is not a requirement
 - An application that does not require task/thread preemption
 - Lack of hardware resources to fulfill the need of an OS.
 - Don't want to use third-party libraries and drivers







- Speed
 - Built-in modules/libraries for connectivity, security and updates
 - Integrated cloud services
 - Broad support
- Complexity
 - Support many hardware and tools
- Cost
 - Free and open source with community support



Mbed OS 6 Conceptual Architectural



Developer Interface

Componentized, Layered Architecture



Zephyr OS



- Concurrent tasks/threads handling with priority and scheduling
- Inter-process communication and synchronization
- Deterministic behavior and timely response to critical tasks.
- Modest range of middleware and IoT libraries such as USB, TLS, IPv4, IPv6, WiFi, CAN, MQTT, etc).
- Tiny memory footprint, 100-200 KB
- Fast boot time, 50ms 100ms
- Long term maintenance and development costs decreased

- Each task has its own private stack and priority.
- Each task is a loop itself.
- When a task to be executed depends on each scheduling mechanisms (round-robin, preemptive-based scheduling)
- No need to break-down long execution function to meet time limitation









2. Label inputs/outputs



3. Identify 1st level tasks



4. Determine dataflow and dependencies



5. Identify application tasks



+ Concurrent programming paradigm. Splitting the software on tasks makes it modular, improve program structure, maintainability and performance.

- + Performance: Guarantee responsiveness to external events in a timely manner
- + Scalability, easier to add complex functionality
- + Portability: Move over to another HW vendor easily due to similar API

+ Portability and reusability, support various platforms. Migration cost decreased dramatically if moving to a different microprocessor.

- + Security: Some level of security through libraries
- + Modifiability: Easy to add new features, easy to modify without much affecting current system
- Requires precious resources on low-end embedded systems
- Overhead for handling OS features such as scheduling, context switching



Products Running Zephyr Today



Full-featured OS

- Open source
- Rich feature set
- Free or Low cost
- Hardware support

- HW Requirements
 - 32-bit and 64-bit processor architectures
 - 8 MB RAM
 - 4MB flash



Full-featured OS

- Bootloader
- Kernel
- Filesystem
- Services
- Application / Programs



Source: http://ccrs.hanyang.ac.kr/webpage_limdj/embedded/LinuxArchitecture.pdf

User space vs kernel space

Memory divided into two distinct areas:

• User space

- normal application software runs
- System calls
 - Examples: allocating memory (variables) or opening a file. Memory and files often store sensitive information owned by different users, so access must be requested from the kernel through system calls.
 - open(), read(), write(), close(), ioctl(), fork(), exit(), wait(), etc.

- Kernel space
 - where the code of the kernel is stored and executes under.
 - Strictly reserved for running a privileged operating system kernel, kernel extensions, and most device drivers.



Source: RedHat

Embedded Linux

- D-Bus: IPC mechanism
- a medium for local communication between processes running on the same host

Process A



 $\ensuremath{\mathbb{C}}$ 2015 Javier Cantero - this work is under the Creative Commons Attribution ShareAlike 4.0 license

Process D

© 2015 Javier Cantero - this work is under the Creative Commons Attribution ShareAlike 4.0 license

Full-featured OS

+ Availability of file-systems, network connectivity, and UI support.

+ Native security features, constantly scanned for vulnerability

+ Versatile, scalable and has support for practically every device driver and protocol.

+ Modifiability, easy to add functionalities in the future, plenty of open-source applications available.

- Needs significant CPU resources, typically in range of few MB for both RAM and ROM footprint
- Soft real-time
- Booting time in several seconds
- High HW cost. Not suitable for 8 or 16-bit MCUs and low power devices.
- Steep learning curve, complex build systems, long compile time, etc.
- Open-source licensing obligation

Choosing the right one

- What is the target HW? What is the available resources on MCU?
- What is the required functionality and performance?
- Does it need real-time requirements?
- How much code do you need to develop? How much standard library it has?
- How complex are the algorithms involved?
- How hard will it be to maintain that code later on?
- Which one has been used in other products? Which one the team is familiar with?
- Languages Does it support Java, Python, C, C++
- Do you need safety-certified embedded OS?
- Maturity of OS, good technical support available



"Of course we can make fast decisions ... once we have considered the 4872 factors."



Source: https://www.pertech.co.il/lynxsecure



Main design constraints

- Power consumption
- Range
- Cost



Source https://industrytoday.com/best-uses-of-wireless-iot-communication-technology



- Bluetooth Low Energy
- LoRaWAN
- Wifi
- Many mesh network
 - o Zigbee
 - \circ Z-Wave
 - o Bluetooth mesh
 - Proprietary mesh

Source: https://www.fluidmesh.com/wireless-mesh-networks









- Low power
 - NB-IoT (5G)
 - LTE-M (5G)
- Plugged in
 - Wifi
 - Ethernet
 - 3G/4G/LTE





Constrained Application Protocol (CoAP), is a client-server protocol (one-to-one).

- Designed for interoperability with the HTTP and RESTful web through proxy.
- Request / response model

Message Queue Telemetry Transport (MQTT), is a **publish-subscribe protocol** that facilitates **many-to-many communication** mediated by brokers.

- Designed for bandwidth-efficient and high latency network
- Separation between producers and consumer of data.

MQTT

Space decoupling: Publisher and subscriber do not need to know each other (for example, no exchange of IP address and port).

Time decoupling: Publisher and subscriber do not need to run at the same time.

Synchronization decoupling: Operations on both components do not need to be interrupted during publishing or receiving.

Scalability: Easy to scale up system in number of publishers and subscribers.



MQTT QoS

- How hard the broker/client will try to ensure message is received.
- 0: Message delivered <u>at most once</u>, with no confirmation (fire-and-forget)
- 1: Message delivered <u>at least once</u>, with confirmation required.
- 2: Message delivered <u>exactly once</u> by using a 4step handshake.

Reliability	Latency / bandwidth
-------------	---------------------

MQTT useful features

- Persistent session
 - Avoid having to re-subscribing every time the connection is interrupted
- Retained message, 'last known good'
 - If a topic is only updated infrequently, then without a retained message, a newly subscribed client may have to wait a long time to receive an update.
- Last Will and Testament
 - 'Last Will' message that will be delivered to other clients when a client disconnects unexpectedly.
- Keepalive
 - Ensure the connection between client and broker is still open

Cloud & backend



Source: scnsoft.com

Cloud & backend

Intelligent lighting



AWS IoT Architecture



Examples of IoT product



Robot vacuum cleaner



Case 1: selection of HW/SW

- Small, wearable embedded device that perform
 measurement
- Collected data to be send to ______
 cloud for processing and storage
- Coin cell battery powered, minimum 1 year

Technology

• Low cost

• What is sample rate and duration required for measurement? How often?

- Which wireless network infrastructure can be used? Wifi? Bluetooth? Lora? 5G?
- How much data to be transferred?
- Which MCU has low power mode? Which communication method use low power?

SW

 Can we use single MCU for both measurement and wireless communication?

HW

Case 1: selection of HW/SW



Case 2: selection of HW/SW

- Industrial embedded device that perform measurement
- Data send to cloud for processing and storage
- ~100 devices per location

Technology

- Small battery powered, 3 year
- Low cost
- Hash environment

• What is sample rate and duration required for measurement? How often?

- Which wireless network infrastructure can be used? Wifi? Bluetooth? Lora? 5G?
- How much data to be transferred?
- Which MCU has low power mode? Which communication method use low power?

SW

 Can we use single MCU for both measurement and wireless communication?

НW









Case 3: selection of HW/SW

HW

- Industrial embedded device that perform measurements on running vehicles
- Collected data to be send to cloud for processing and storage
- Battery powered
- Wireless
- Easy installation

Technology

- What is sample rate and duration required for measurement? How often?
- Which wireless network infrastructure can be used? Wifi? Bluetooth? Lora? 5G?
- How much data to be transferred?
- Which MCU has low power mode? Which communication method use low power?

SW





Thank you