INTERRUPT

# INSIDE

An in-depth magazine about embedded technology from Data Respons 1, 2015

## Device specific power consumption control

Monitor usage and savings, both economically and environmentally.

## Industrial connected things

What challenges are there, and how can we meet them?

data respons
EMBEDDED SOLUTIONS

# WHAT'S **INSIDE**


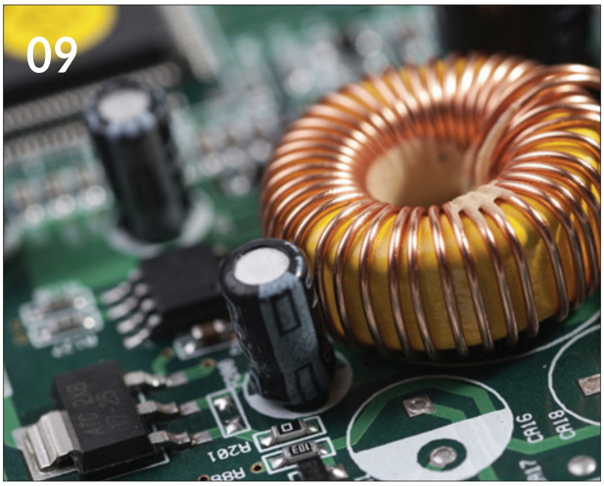




ARTICLE SERIES: **The impact of consumer electronics on development of high-reliability products**
This issues article: High density interconnect.
Next article: LTB - the dreaded Last Time Buy...

# WELCOME TO INTERRUPT INSIDE

Data Respons' vision, a smarter solution starts from inside, is our company's DNA described in one sentence. We truly believe that we can make the world smarter and we think that this starts from the inside - whether it be inside the heads of our specialist engineers or new technology embedded into the world's products and solutions.

IOT – Internet of Things – is one of the megatrends of our times. Everything surrounding us, not only devices and products, but also wearables and buildings, is getting smarter. In addition, modern wireless technology provides the connectivity in a more efficient and cost-effective manner. Many of the media articles about IOT are often related to consumer applications you control via a smartphone. However, IOT is about THINGS, not people. It is about the Industry- and technology companies integrating IOT and the ANALYTICS into their industrial products and applications. This is where Data Respons, our specialist engineers and our customers are meeting the world of IOT and the specific technology challenges.

In this magazine, we want to highlight relevant technology areas and thoughts related to IOT and connectivity. Additionally, this issue continues the series of articles discussing the impact of consumer electronics in developing high reliability products.

Our own specialist engineers and employees write all of the articles. We welcome any feedback and suggestions from our readers.
Enjoy the reading!

**KENNETH RAGNVALDSEN**

# DEVICE SPECIFIC POWER CONSUMPTION CONTROL

eWave is an energy display which goal is to raise awareness on the amount of energy spent in a household, given in kWh, currency and $CO_2$ emissions. The original eWave display was a passive display that only displayed the consumption of the main meter, but the latest version of the device can also actively control user-selected devices in a home to reduce unnecessary consumption. eWave is a product of the Sandefjord based company eWave Solutions, previously known as Miljøvakt.

BY: Andre Firing
Data Respons Alumni

## FROM AWARENESS TO CONTROLLED POWER SAVING

eWave started its life as an idea by entrepreneur and founder Gunnar Skalberg. Norway is currently among the top ten in the world when it comes to energy consumption per capita. The main reason for this is the cold weather, and the relatively cheap power prices. The main source of the power consumption in Norway is arguably heating of houses and water, and lighting. Households have been reported to be responsible for almost half of the $CO_2$ emissions in the world, and eWave Solutions' goal is to reduce this significantly.

The eWave tablet started out as a touch display, targeted at a scientific test project. The decision to go against the mainstream app-based energy displays combined with a gateway was a conscious choice. Previous studies conducted in the UK concluded that a physical device greatly increases awareness compared with a more abstract application on a smartphone. Since the main instrument for energy reduction in the original eWave was increased awareness of the issue, a concrete and visible tablet-based display would prove important. However, the eWave tablet has gone through several iterations, evolving from an awareness-based display to an active power saving home control device.

## SYSTEM AND FUNCTIONS

The latest version of the eWave tablet is currently capable of controlling switches wirelessly, and reading temperatures from wireless thermometers. It is also possible to use the eWave display to keep track of power consumption on several circuits, which makes the device perfect for tracking how much energy one single oven is using over time, with associated costs in both currency and $CO_2$ emissions. eWave can also display the current consumption of any circuit


with a high refresh rate. This feature gives the user instant feedback when they turn on or off an electric device.

The eWave tablet includes many features targeted at reducing energy consumption, including regular savings tips and overview of the current and historical power price from the energy providers. The tablet can also keep track of the household consumption. A savings account application is available in the tablet, which lets the user set up a saving goal per day, and keeps track of how much money the user saves over time. The user also sets up a yearly saving goal, which the tablet uses for feedback and status reports every day.

Furthermore, eWave Solutions and Data Respons is currently working on extending the eWave functionality further into the world of home automation with more smart control of household power consumption.

## HVALER PROJECT

In early 2014, eWave partook in a research project in Hvaler organized by Smart Energi Hvaler. The goal of the project was to test out new energy reducing technology and see the effects it had on the consumers. eWave proved to be one of the most influential devices participating in the project, resulting in a general consumption reduction of up to 20%. Some users were also able to use the eWave tablet to find electrical faults that increased energy consumption in their homes.

## SOFTWARE AND HARDWARE

The tablet used for the eWave project is a custom Android based tablet running on a dual core ARM Cortex-A9 CPU. The eWave application was developed using Qt for Android, which was in the early stages when the project started. Qt was chosen to ensure platform interoperability if another OS is chosen at a later stage. Using Qt with minimal Android support was somewhat of a gamble, but as the support grew better, the gamble paid off significantly. Qt is now supporting Linux, Android, iOS, OSX, Windows, WinRT, BlackBerry, Sailfish OS and more, which makes the eWave application highly portable.

For wireless applications, the eWave devices supports both Z-Wave and ZigBee. The Z-Wave API is currently only targeted at energy readings. For ZigBee, the tablet supports clusters for switches, thermometers, energy readings and a few others, but more clusters will be supported when needed. The tablet sets itself up as a ZigBee network coordinator, and automatically binds to all previously known devices that are nearby. This simplifies the user experience and makes the eWave easier to use.

All energy consumption data is stored both locally on the tablet, but also synchronized to a server maintained by eWave Solutions. This way, all historical data is safely stored and available if a device needs to be replaced. Having the consumption data stored on a remote server also enables large scale observation of consumer energy consumption habits. This data can be used as an indication for future improvements in the power infrastructure, which can be a major asset for the power company. The consumption data can also be used for research and commercial purposes.

## CONTRIBUTIONS

Data Respons has been responsible for the development of the application since December 2012, and is currently working closely with eWave Solutions' CTO on the home automation extensions for the product. The eWave project has been an important project for the R&D department in Asker over the past two years, with innovative development and exciting new technology. With the increasing focus on reduced energy consumption and $CO_2$ emissions, the eWave product projects a brighter future that R&D Services is proud to be a part of.

# INDUSTRIAL OR COMMERCIAL

# WIRELESS MESH TECHNOLOGIES

This article will try to answer two questions: Why would you choose an industrial wireless mesh technology compared to choosing a commercial wireless mesh solution? Secondly, which industrial technology works best for your application?

BY: Aksel Bondø
Senior Development Engineer
Data Respons

BY: Alexander Svensen
Development Engineer
Data Respons

## INDUSTRIAL VERSUS COMMERCIAL WIRELESS MESH TECHNOLOGIES

Industrial wireless and commercial wireless mesh technologies often use the same radio technology. For instance Zigbee, 6LoWPAN, WirelessHART, SmartMesh IP, and ISA100.11a use a 2.4GHz radio. This means we practically can use the same radio HW to run them all. Therefore there is not much gain in using industrial wireless mesh technologies over commercial solutions HW wise.

However, an industrial use case may need higher performance than commercial cases. For example:

- High reliability
- Low latency (close to real-time)
- Secure systems

This is where the industrial wireless mesh technologies can provide a better solution than the commercial solutions.

## MESH NETWORKS

What is wireless mesh compared to other wireless topologies?

A star topology is similar to what you might find in a home network where every end node is connected to the centralized router. It's main vulnerability is that if the router is disabled, then the whole network collapses. A cluster tree is similar to the star topology but with peripherals connected to the previous end devices (e.g. a DAQ (Data Acquisition) unit connected to a computer connected to the server). For the DAQ unit to be connected to the rest of the network, both the computer and the server needs to be operational. However, in the mesh topology every node can be both an end device or a router, meaning that each node has several links to the coordinator. This means that if one of the routers goes offline, most of the network is still intact by rerouting through the remaining routers.

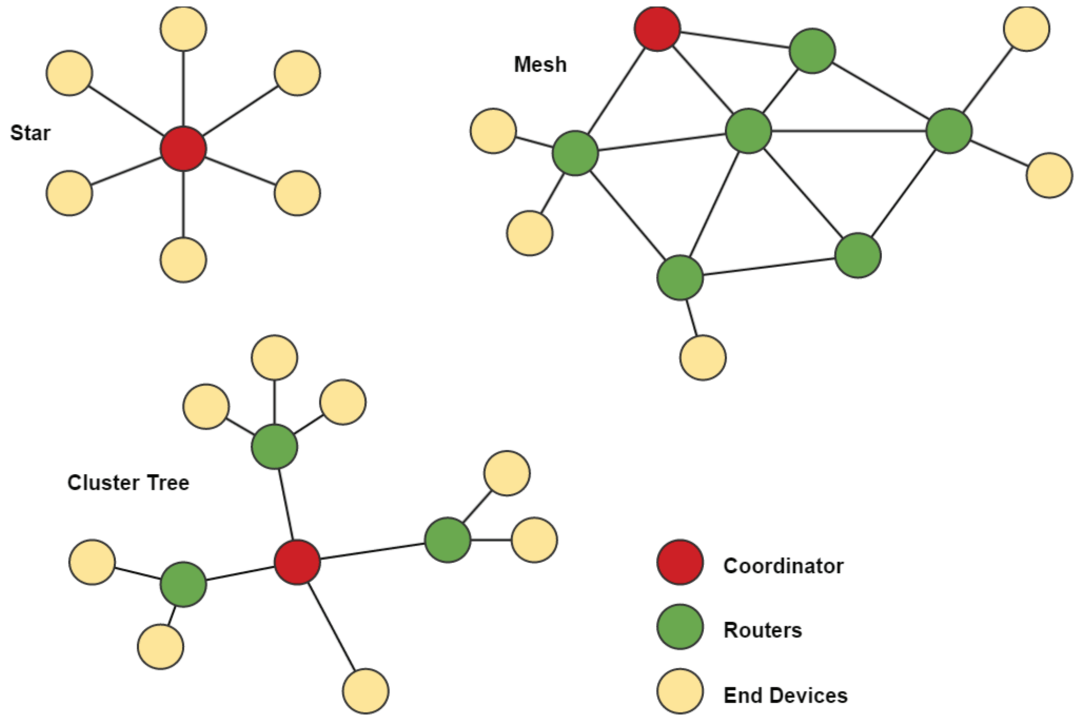


## DIFFERENT INDUSTRIAL WIRELESS MESH TECHNOLOGIES

There are three industrial wireless mesh technologies that we will go further into details on in this article:

- SmartMesh IP
- SmartMesh WirelessHART
- Nivis ISA100.11a

SmartMesh IP is the only one of these that is built for IP compatibility. However, they all provide resilience, reliability, scalability, security, and are made for low power applications.

## COMPARISON OF TECHNOLOGIES

The use case that the comparison is based upon is 40 nodes in each network, measuring environmental data, and sending data every second.

The compared technologies are SmartMesh WirelessHART node and manager modules, SmartMesh IP node and manager modules, and a custom HW designed by Data Respons with the Nivis ISA100.11a stack. For SmartMesh IP and SmartMesh WirelessHART, the cost is the cost of the modules, while for the custom ISA100.11a solution, the cost is a combination of development cost, stack license cost, radio certification cost, support cost, HW component cost, and PCB manufacture cost.

## MARKET SHARES

According to CDS, the market share between WirelessHART and ISA100.11a is about 50% for each in average considering Asia, US, and Europe.

- In the Asian market ISA100.11a is in favor.

| PARAMETERS | SMARTMESH WIRELESSHART | SMARTMESH IP | NIVIS ISA 100.11A |
|---|---|---|---|
| Indoor range | 100m | 100m | 100m |
| Manager price (100+) | (LTP5903-WHR) $654.78 | (LTC5800IWR-IPRA#PBF) $97.79 | (Custom) $44.1 |
| Maximum number of motes at a 1 second interval | 23 | 27 | >40 |
| Maximum number of motes at a 2 second interval | 48 | 65 | >65 |
| Mote price (100+) | (LTP5901IPC-WHMA1A2#PBF) $59.80 | (LTC5800IWR-IPMA#PBF) $25.30 | (Custom) $8.12 |
| Network Reliability | >99.999% | >99.999% | >99.999% |
| Network-Wide Reliability and Power Optimization | Yes | Yes | Yes |
| Number of channels | 15 | 15 | 16 |
| Number of nodes | 250 | 100 | 250 |
| Operating frequency | 2.4000 - 2.4835 GHz | 2.4000 - 2.4835 GHz | 2.360 GHz to 2.480 GHz |
| Per Transmission Frequency Hopping | Yes | Yes | Yes |
| Raw data rate | 250 kbps | 250 kbps | 250 kbps |
| Redundant Spatially Diverse Topologies | Yes | Yes | Yes |
| Temperature | –40°C to +105°C | –40°C to +105°C | -40°C to 105°C |
| Time Synchronized Network-Wide Scheduling | Yes | Yes | Yes |
| Wireless standards | IEC 62591 | 6LoWPAN and 802.15.4e | 802.15.4-2006, 6LoWPAN |
| Open source stack | No | No | Yes |

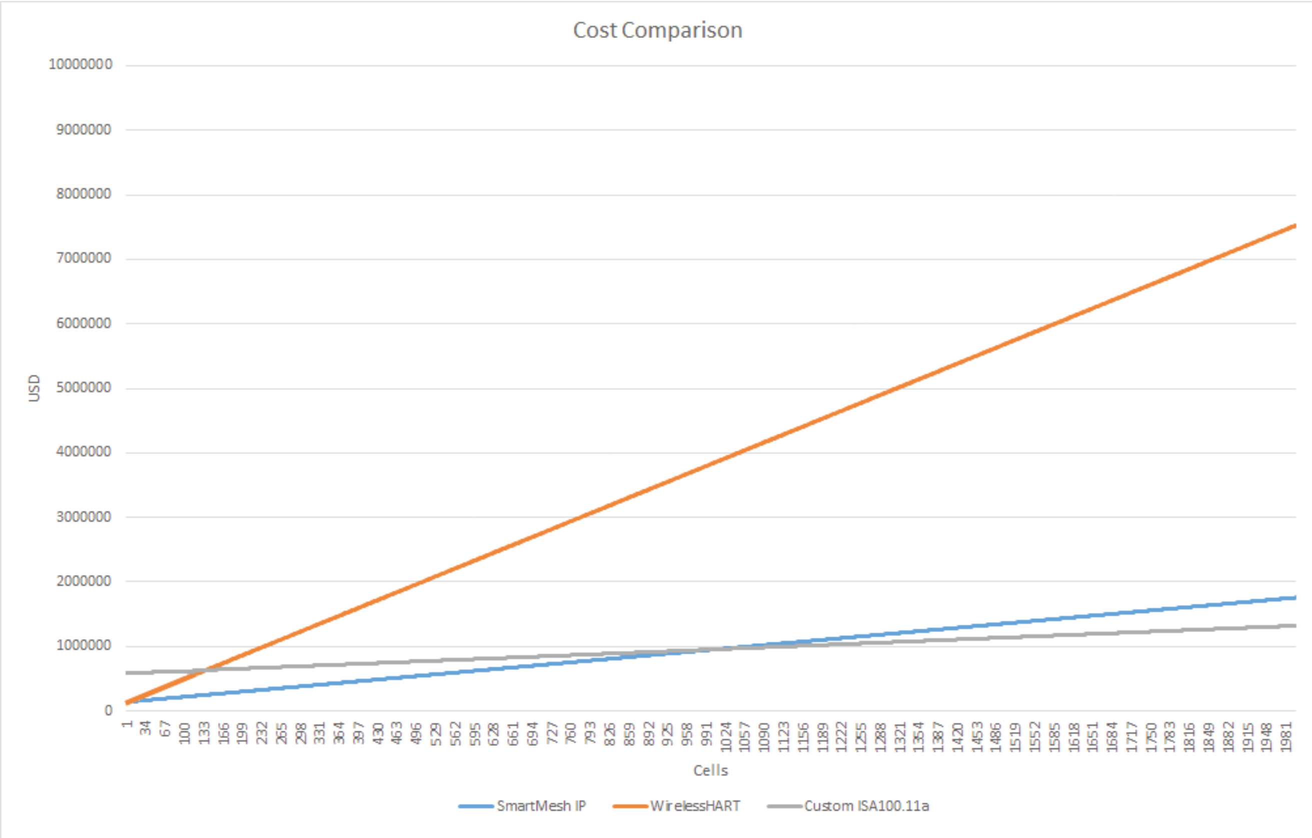


"The conclusion is therefore that the choice of technology really depends on the scale of your system.

- In the US market ISA100.11a and WirelessHART is split approximately 50% for each.

- In the European market WirelessHART is in favor.

SmartMesh IP is a fairly new protocol, so it has no considerable market share yet.

## CONCLUSION

As shown in the cost graph in the paragraph above, the custom ISA100.11a is cheaper for high volume production given the low cost per mote. SmartMesh IP is a valid option up until about 1000 managers, but has its strength in keeping a low cost for the mid sized systems.

The reason the SmartMesh IP is more expensive than the ISA100.11a for high volumes is due to the number of nodes supported, thus needing more managers. WirelessHART system is the most expensive but the most utilized mesh technology in Europe and the US.

The conclusion is therefore that the choice of technology really depends on the scale of your system. If you are setting up a small system, WirelessHART might be the way to go since you get a well-tested and used system .

If you are looking at a mid-sized system with up to 500 nodes, SmartMesh IP would be the logical choice. For larger systems you are better off choosing a custom ISA100.11a. This way you can keep the cost down and still have a large network of nodes. Note that this option requires more development cost than the other.

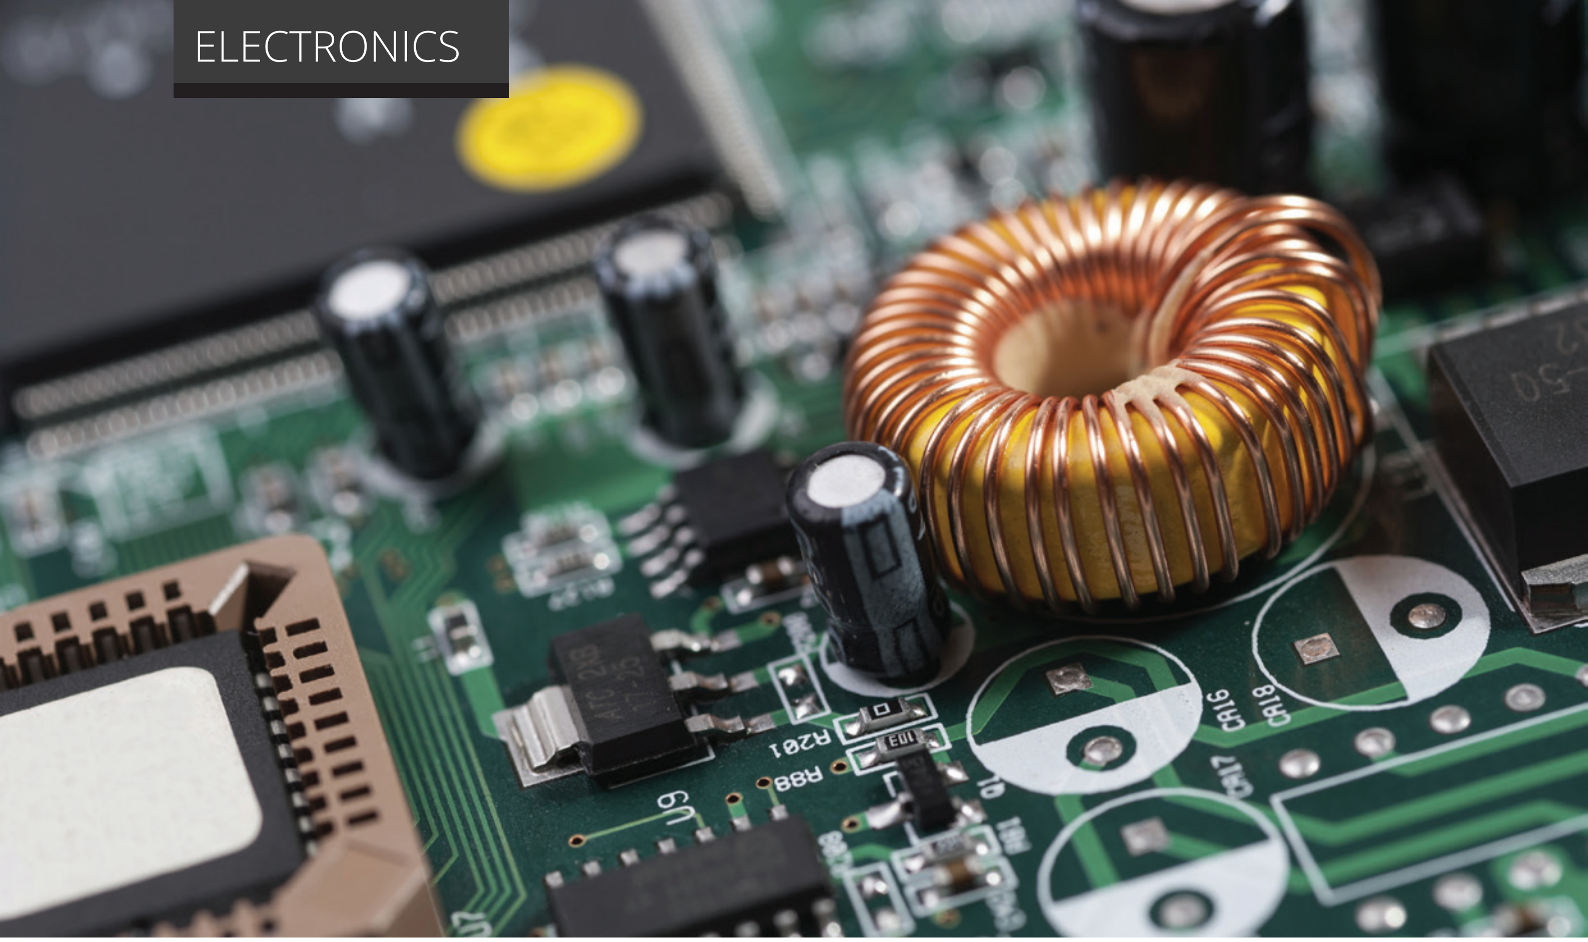


# VOLTAGE AND CURRENT MODE CONTROL

Switch Mode Power Supplies power-supplies (SMPS) or pulse width modulator (PWM) converters have been with us for many years. They came into serious play in the mid-sixties, just after the introduction of practical power transistors. Control of a PWM converter comes in two flavours; namely voltage-mode and current-mode. These two control methods have their own sets of advantages and disadvantages which will be presented in this article.

BY: Frode Sørengen
Senior Development Engineer
Data Respons

Switching power-supplies were not yet very popular outside military and space application when Silicon General introduced the first monolithic (IC-based) pulse width modulation controller, the SG1524, in 1975. Other semiconductor companies did not stand still when SG1524 made its debut. Competitors, like Motorola Semiconductors, Texas Instruments and Signetics, quickly responded and offered similar devices to the market. All these early devices were based on voltage-mode control.

The next major advancement was a legendary IEEE paper by Cecil Deisch in 1978, which introduced current-mode control in a practical and easy to understand circuit. This paper was not the first publication on current-mode control, but it was the one that made it known to industry. It was not until the early 1980s, when Unitrode introduced the first PWM controller with current-mode control, that this control method gained popularity.

>>

Soon after the introduction of current-mode PWM controllers, the industry adopted it as the preferred method for controlling PWM-converters.

## CONTROL TECHNIQUES OPERATION

### VOLTAGE MODE CONTROL

In the voltage-mode control scheme shown in Figure 1, the converter output voltage is sensed and subtracted from an external reference voltage with an error amplifier. The error amplifier produces a control voltage that is compared to a constant-amplitude sawtooth waveform. The comparator produces a PWM signal that is fed to drivers of controllable switches in the dc-dc converter. The duty cycle of the PWM signal depends on the value of the control voltage.

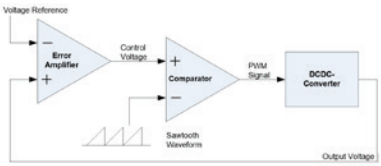


Figure 1. Main control scheme for voltage-mode control

### CONTROL TECHNIQUES - CURRENT-MODE CONTROL

The current-mode control scheme is presented in Figure 2. An additional inner control loop feeds back the inductor current signal. This current signal is converted into a proportional voltage, and is compared to the control voltage. The inner current loop turns the inductor into a voltage-controlled current source, effectively removing the inductor from the outer voltage control loop (see Figure 3). The great advantage relative to voltage-mode control is that it removes the double pole caused by the interaction between the output capacitor and the power stage inductor. With current-mode control, the output characteristics are reduced from double-pole to a single-pole which gives a significant stability improvement.

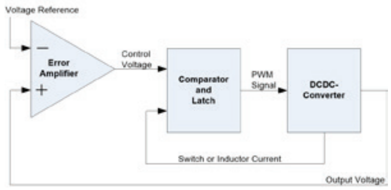


Figure 2.
Main control scheme for current-mode control.

## ADVANTAGES AND DISADVANTAGES

Figure 3 shows a simplified schematic of a buck-converter. The PWM converter power stage consists of the transistor Q (the switching element), diode D, inductor L and capacitor COUT. For a voltage-mode control, the output has a second-order low-pass filter characteristics. For current-mode control, the output has a first-order low-pass filter characteristics.

Figure 3.
Simplified schematic of a buck converter using current-mode control.

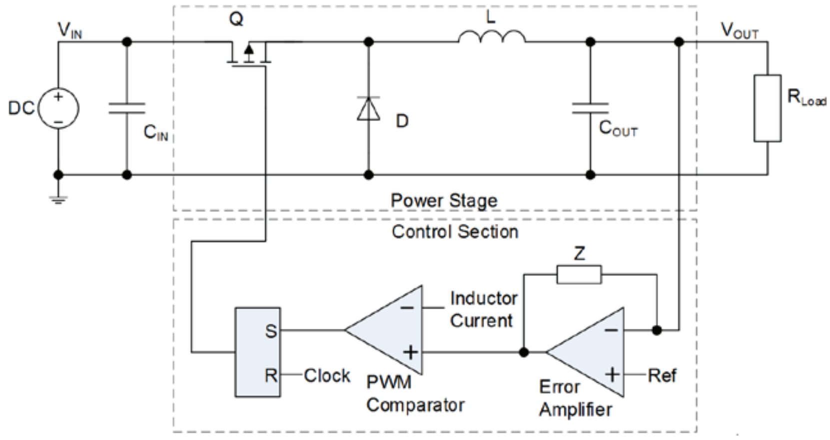


The advantages and disadvantages of the two PWM operated control modes are closely linked. Therefore, the advantages and disadvantages are only summarised for current-mode control in the following section. The summary for the two control methods is by no means complete, although the most important items should be included.

### CURRENT-MODE ADVANTAGES

The most attractive current-mode advantages are listed below.

**#1: Easy Compensation**

Voltage-mode frequency control has a sharp phase drop beyond the filter's resonant frequency (see Figure 4, blue trace) which requires a type three compensator to stabilize the system. Compensation is further complicated by the fact that the loop gain varies with input voltage.

Current-mode control looks like a single-pole system at low frequencies (see Figure 4, red trace), since the inductor is controlled by the current loop. This improves the phase margin, makes the converter much easier to control, and results in a higher gain bandwidth over a comparable voltage-mode circuit.

Figure 4 (next page) compares the power stage gain and phase of voltage-mode and current-mode, demonstrating how much easier the current-mode system is to compensate (because of its single pole characteristic).

### VOLTAGE MODE CONTROL

The difference between the desired and actual output voltages (error) adjusts the duty cycle to control the voltage applied across the inductor.

### CURRENT MODE CONTROL

The difference between the desired and actual output voltages (error) controls the peak inductor current.

### CCM AND DCM OPERATION

Continuous-conduction-mode (CCM) means that the current in the energy transfer inductor never goes to zero between switching cycles.

In discontinuous-conduction-mode (DCM) the current goes to zero during part of the switching cycle.

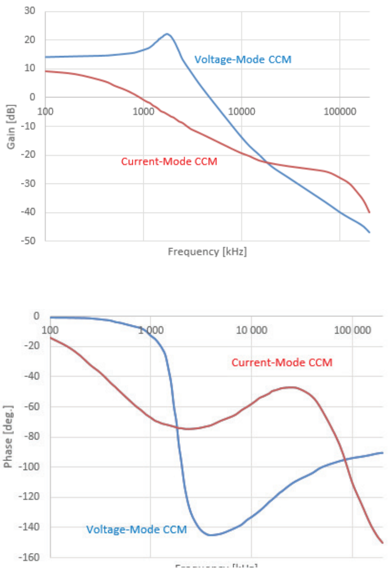


Figure 4.
Voltage- and current-mode gain and phase comparison

**#2: RHP Zero Converters**
Boost- and buck-boost-topology and all topologies based on these two topologies, including the commonly used flyback-converter, has a right half-plane (RHP) zero in the transfer function. With voltage-mode control, the transfer-function crossover must be well above the resonant frequency, or ringing will be introduced in the filter. In a converter where the crossover frequency is restricted by the presence of an RHP zero, this is not always possible. With current-mode control, it is not a problem to have a control loop crossover at or below the filter resonant frequency.

**#3: CCM and DCM Operation**
When moving from continuous-conduction mode (CCM) to discontinuous-conduction mode (DCM), the characteristics with voltage-mode control are drastically different, as shown in Figure 5. It is not possible to design a compensator with voltage-mode that can provide good performance in both regions. With current-mode control, crossing the boundary between the two types of operation is not a problem. The characteristics are almost constant in crossover region, as shown in Figure 6.

## BODE PLOT

A Bode plot is a graph of the gain (in dB) and phase of the transfer function versus log of frequency.

## LOOP GAIN DEFINITIONS

**Phase margin:**
Measure of phase angle above -180 degrees when gain crosses unity (0dB) for the final time. Phase margin is a measure of relative stability. Should be minimum 50 degrees for power-supply design.

**Gain margin:**
Measure of gain when phase is below -180degrees. Gain margin is a measure of sensitivity to parameter variation. Should be at least 10 dB.

## COMPENSATION

The compensator is normally included in the error amplifier: Its function is to shape the error signal to improve the closed-loop transfer function and overall performance.

## POLES AND ZEROES

Transfer functions poles and zeros are the frequencies for which the value of the denominator and numerator of transfer function becomes zero respectively.

## RHP ZERO

The right-half-plane (RHP) zero has the same 20 dB/decade rising gain magnitude as a conventional zero, but with 90° phase lag instead of lead.

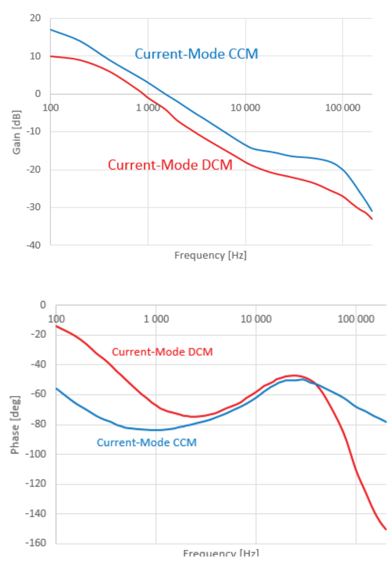


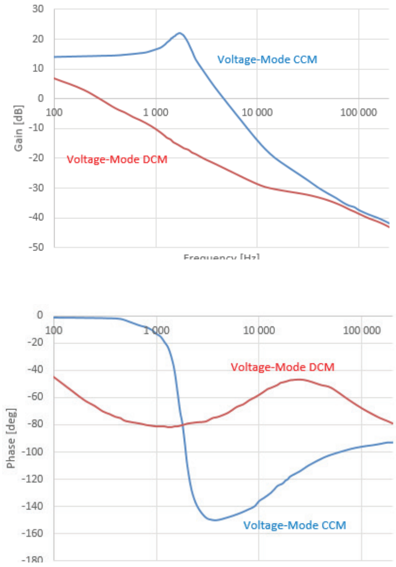


Figure 5.
Voltage –mode in CCM and DCM

**#4: Fast response to input voltage and load changes**
Current-mode control responds immediately to the input voltage changes and offers good line rejection. The fast response is also due to the single "pole" system (reduction in the converters dynamic order).

For voltage-mode, the situation is opposite since any change in line or load must first be sensed as an output change and then corrected by the feedback loop, which usually results in a slow response. Note that many modern voltage-mode controllers are implemented with a feed-forward technique which eliminates the effect of input voltage variations, but complicates the circuit analysis.

**#5. Current limiting on every cycle.**
Current mode control provides inherent current limiting on a cycle by cycle basis. Current limiting improves system reliability in response to current transients. The power component are protected from high peak currents due to this.

**#6: Current sharing with parallel modules**
A single error voltage can be used to control multiple converters in parallel. This results in the possibility for equal current sharing in modular converters.

>>

Figure 6.
Current –mode in CCM and DCM

## CURRENT-MODE DISADVANTAGES

The improvements offered by current-mode are impressive, but this technology also comes with its own unique set of problems which must be solved in the design process.

### #1: Current Sensing

Either the switch current or inductor current must be sensed accurately. This requires additional circuitry, and may result in power loss. In most isolated power-supplies, the switch current is sensed either with a resistor or current transformer. A wideband current sensing is required to accurately reconstruct the current signal.

Voltage-mode only requires voltage sensing which is easier than current sensing. This results in advantages as less noise, sometimes less power loss, less cost, and a higher resolution.

### #2: Sub-harmonic Oscillations Instability

Current-mode control may be unstable when the duty cycle of the converter approaches 50%. A compensating ramp is needed to fix the problem which may in turn introduce other complications.

### #3: Signal-to-Noise Ratio

The main problem in almost every current-mode supplies is noise on the current sense signal. In many power-supplies there is simply not enough signal to control the converter smoothly over the full range of operation. Voltage-mode control has a large-amplitude ramp waveform which provides good noise margin.

### #4: Extremely wide input voltage variation may be difficult to support.

A wide input voltage range creates design issues because large variations in PWM duty cycle ratios are required, and exceeding 50% duty cycle introduces issues with slope compensation.

Voltage-mode does not have any duty-cycle restriction.

## WHEN TO USE VOLTAGE-MODE CONTROL?

It may come as a surprise that voltage-mode control is still used in the industry when considering all the presented current-mode advantages above.

The main downside of current-mode control has always been the difficulty of implementation, due to the reduced signal-to-noise ratios in the control loop. It is the designer's job to make sure the current is sensed cleanly and accurately and it requires skill and effort by the designer.

Continual efforts to reduce the physical size pushes switching frequencies higher which in turn aggravates the noise sensitivity problem, so many integrated power supply makers have reverted back to voltage-mode control. This approach reintroduces inflexibility and lack of robustness as voltage-mode converters are more susceptible to component variation. It is possible to design a high performance power supply using voltage-mode control, although it will not achieve the same degree of ruggedness as using current-mode control. This too requires skill and effort.

> "...it difficult to imagine that IoT applications intended for use in the automotive or health care industry or any IoT device operating in harsh environments could accept to compromise on ruggedness. Skilled power supply designers are able to overcome current-mode control noise challenges when required.

Manufacturers in many segments would happily trade off robustness for smaller physical size or reduced unit cost. One common characteristic of these segments is a limited operating temperature range which eases the burden of designing a fairly robust power supply using voltage-mode control.

One hot example is the latest advancement of the Internet, namely Internet of Things (IoT). Many IoT applications will probably be implemented using voltage-mode control. Especially the applications intended for use in the commercial and cost sensitive markets. Examples of this could be wearables, building & home automation, lightning network, food production and retail.

It is, however, difficult to imagine that IoT applications intended for use in the automotive or health care industry or any IoT device operating in harsh environments could accept to compromise on ruggedness. Skilled power supply designers are able to overcome current-mode control noise challenges when required.

The conclusion is still that current-mode control is the first choice if rugged power supply design is priority.

This article is part of a **series of articles** regarding the impact of consumer electronics on development of high-reliability products.

# HIGH DENSITY
## INTERCONNECT

The previous article in this series discussed how a conventional PCB designed to high reliability requirements is incompatible with component packages with very small spacing between solder connections. Ball Grid Arrays (BGA), the dominant package for complex electronic components, feature sub-millimeter pitch driven forth by the consumer segment and their continual need for integration and miniaturization. Feature sizes demanded by high reliability electronic boards (IPC-6011, class 3) precludes use of packages with pitch smaller than 0.8mm using a conventional fabrication process, while popular components are offered with pitch of 0.65mm, 0.4mm and even less. Achieving Class 3 compliance when using these packages is only possible by extending the conventional fabrication process with features collectively known as High Density Interconnect or HDI.

BY: Haldor Husby
Principal Development Engineer
Data Respons

In conventional PCB fabrication, a feature called via accomplishes vertical connection between copper layers of the PCB. Vias are formed when holes drilled through a laminated board are copper plated forming a conductive barrel through the drilled hole. The barrel makes electrical contact to copper pads etched on the various layers and assures connection between them. Mechanically drilled, the holes extend through the entire thickness of the board, and there are practical limits to how small the drill diameter may be. Hole diameters smaller than 0.3mm becomes a cost driver.

### HDI CONSTRUCTION TYPES

The object of HDI is to achieve higher wiring density than conventional boards, and a central feature of the technology is the microvia, blind vias defined by

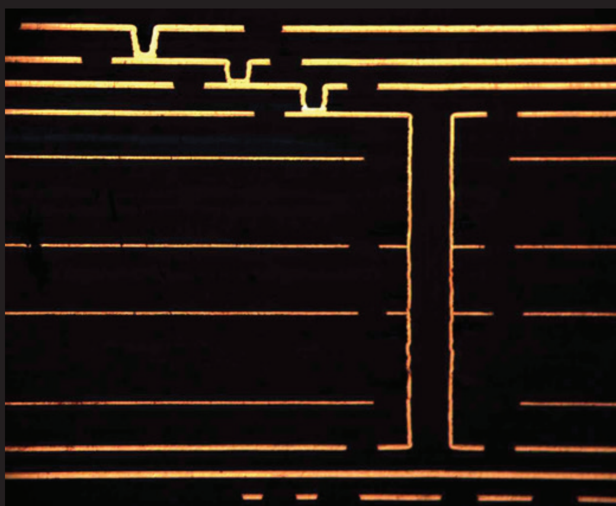

FOTO: ELMATICA

### DEFINITIONS

**HDI:** Printed circuit board with a higher wiring density per unit area than conventional printed circuit boards (PCB). They have finer lines and spaces (≤ 100 μm), smaller vias (≤ 150 μm), and capture pads (≤ 400 μm), and higher connection pad density (>20pads/cm2) than employed in conventional PCB technology

**MICROVIA:** A blind hole with diameter ( ≤ 150 μm) having a pad diameter ( ≤ 350 μm) formed by either laser or mechanically drilling, wet/dry etching, photo imaging or conductive ink formation followed by a plating operation for product development.

>>

hole diameters smaller than150µm and normally drilled by laser. The microvia only extends between two, or at most tree layers, and is usually used in combination with buried vias and regular conventional through-hole vias. A prime example of HDI technology is the BGA package itself which is a printed circuit board with extremely small features.

While the fabrication of conventional PCBs are standardized and well known, HDI boards are constructed in a wide variety of ways. IPC's sectional design standard for HDI (IPC-2226) details 6 general classes of construction from class I where microvias are formed in the surface of a conventional PCB, to type VI in which electrical interconnections and mechanical structure are formed simultaneously. One of the prevalent structures, type III, is shown in the illustration to the right.

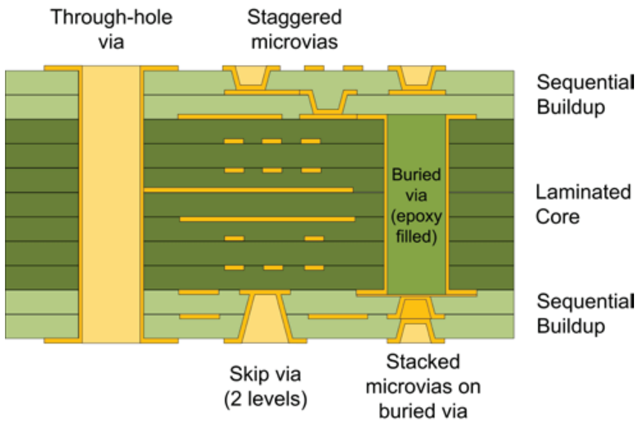


### HDI TYPE III STRUCTURE

A type III HDI structure has two levels of microvias, buried vias and regular through-hole vias. The fabrication process starts with a laminated PCB, which includes plated through-hole vias. Additional dielectric layers and copper foils are added to the board in sequential lamination cycles. After each cycle, small holes are laser drilled and plated to form microvias between pairs of layers. The through hole vias of the original laminated core form the buried vias in the finished structure.

### A COST ADDER ... AND SAVER

In low-volume segments, microvias were often added to a lay-up when routing a design using a conventional structure proved difficult or impossible. Adding microvias to a given lay-up adds cost, so microvias are often regarded as a cost driver to be avoided if possible. However, designers in the consumer segment use HDI and microvias aggressively to achieve cost reductions.

HDI's cost saving potential stems from its efficient use of space and area which in turn results in higher routing density. The via structure itself has smaller diameter and consumes less space. The primary advantage, however, comes from the small vertical extension of the microvia, which leaves larger routing channels on other layers. Judiciously employed it allows a reduction of the number of layers in the board for a given circuit, which more than offsets the cost added by the inclusion of microvias.

Published case studies claims PCB cost reductions in the order of 50% by layer reduction (from 18 layer to 10 in some examples), and reduction in overall board area. When shipping millions of boards the savings are well worth the extra effort.

None of this comes easy, however. An entirely different regime is required to make the right design decisions in terms of construction type, material choice, lay-up and the use of the chosen build before board routing may even start. The wider variety of construction types and methods demands close technical contact with a capable board shop or broker up front.

The board must be designed for the process it will be manufactured with, requiring a far deeper understanding of process and materials on the part of the designer than with conventional boards.

### SIGNAL AND POWER INTEGRITY

Microvias are also used effectively in improving signal and power integrity. Through-hole vias represent small capacitive loads and stubs which may cause a degradation of high-speed signals. Smaller, shorter microvias presents reduced parasitic loads and allows for routing without via stubs. Filled and capped microvias placed directly in component solder lands reduces inductance in the power distribution network. Microvia in pad is advantageous for high speed signaling as well.

Large processor chips may pull several tens of Amperes and experience significant load steps. The grid of large through-hole vias under dense BGAs often perforate the power and ground planes that supply this current to an extent where the flow is restricted. Microvias, with smaller diameter and limited depth, reduces the perforation and helps limit ripple and voltage drops to an acceptable level.

### STAGGERED AND STACKED

When a signal must traverse several levels of microvias, vias on subsequent layers may be placed with a small offset, in a staggered fashion. It is, however, also possible to stack them right on top of each other and also on top of buried vias. Stacking microvias is more space efficient and makes for easier routing, but there is a cost in terms of reliability.

A single microvia by itself is the most reliable interconnect structure of all with staggered microvias as a close second. Stacked vias will experience greater

> ...the wider variety of construction types and methods demands close technical contact with a capable board shop or broker up front.

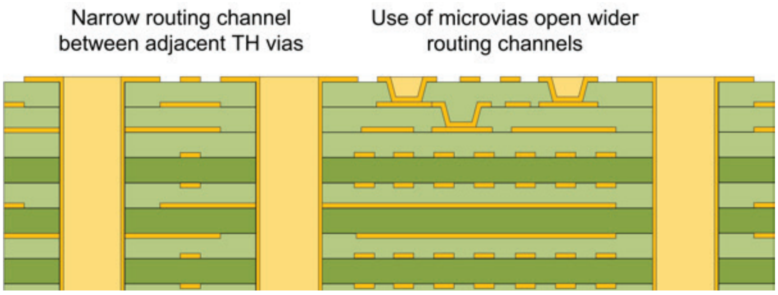


## LAYER REDUCTION

The escape routing from large BGAs usually determines the number of routing layers required in a complex PCB. Most connections must be channeled directly to an inner layer to be routed out of the BGA area. With a conventional PCB structure, the result is a dense grid of through-hole vias that obstruct the path for nets to escape. With 1mm pitch, only a single trace may pass between two adjacent vias resulting in a massive growth in the layer count as the array grows. HDI allows for smaller track and space, and with fewer vias extending through the board, broader routing channels may be established. The result is a higher routing density and significant reduction of the number of routing layers required.

Offset stacked layers of microvia

thermal stress during the solder reflow process and is generally less reliable than a conventional through-hole via. The buried via is considered the least reliable structure depending on the details of its formation. For high-reliability boards it is advisable to use staggered structures, and to take great care in specifying the buried vias.

### USE WITH CARE AND KNOWLEDGE

Adding HDI features to a board only after a designer has struggled and failed to route it using a conventional structure adds cost without benefit. No one will make good design decisions in desperation and on overtime. It is smarter to assess if a complex design could benefit from HDI and consider the various design options up front.

Realizing the full advantage of HDI technology takes careful planning and in-depth knowledge of construction techniques, processes and materials. And close cooperation with a capable board-shop or broker from the outset is mandatory.

## NEXT ARTICLE IN THIS SERIES

LTB must be the three most dreaded letters amongst those who manage electronic products with long market lifespans. The Last Time Buy notice informs the user of an electronic component that the fabricator will not accept order for the component after some specified date in the near future. Before this date the recipients of the notice must determine how many pieces of the component they need for all eternity and place their last order. Armed with uncertain forecasts they must balance the capital cost of large component stocks with the risk of lost upsides in the future. Companies with large portfolios receive LTB-notices every week and obsolescence management places a heavy burden on the organization. In the next issue of Interrupt Inside we look at the drivers behind escalating obsolescence, obsolescence management and design strategies for products with decades of market life using a supply chain geared for product lifespans of two or three years.

# FEM MODELLING

FEM modelling is an effective tool that reduces development time and cost. It can be applied to virtually any field of technology and this article will focus on piezoelectricity, structural dynamics and acoustics. When combined with signal processing the entire operation of e.g. a measurement system can be simulated, tested and modified before a real life prototype is produced and tested.

BY: Erik Asplund
Principal Development Engineer
Data Respons

FEM (Finite Element Method) is a method for solving the differential equations that describe e.g. a mechanical problem by subdividing the solution into a number of smaller parts. In modelling this means that an arbitrary structure is divided into a number of structural elements often shaped as rectangles or triangles (two dimensions) and rectangular boxes or tetrahedrons (three dimensions).

Examples of technology areas where FEM-modelling is used are structural mechanics, acoustics, fluid dynamics, heat flow, optics and electromagnetic fields. Data Respons has contributed with FEM-modelling in several projects dealing with various applications such as elastic wave propagation in steel bolts, acoustic noise pollution in offshore piling, sound propagation in district heating pipes and a fish tag.

## THE SOUND OF ICE

Some of us have perhaps noticed the peculiar and captivating sound that occurs when we skip a stone over thin, newly formed ice on a lake. There are a number of videos on Youtube illustrating this phenomenon and the sound the skipping makes. It is an interesting challenge to model the structural dynamic

The modelled simulation of the "Sound of Ice" can be experienced in the online version of this article at datarespons.com.

A simple structural dynamics FEM-modelling example showing a tuning fork vibrating and producing the musical note A (440 Hz).

"

It is interesting to notice that a bounce from the stone creates an essentially finite acoustic pulse in the air even though the ice sheet vibrates for a relatively long period of time.

and acoustic phenomena that cause the fascinating sound.

To mimic the proper conditions we need to model an ice sheet of e.g. 2 cm on top of a water volume. Above the ice is the air conducting the sound to our ears. The impact of the stone is modelled as a momentary point force acting perpendicular on the ice in the very centre of the model. In reality the listener is static and the stone hits the ice at progressively increasing distances from the listener. In the model we do the opposite and move the listener to different positions (radii) from the pounding stone. This way we can model the situation as a cylindrical 2D-model which greatly reduces the model complexity compared to a full 3D-model.

It is interesting to notice that a bounce from the stone creates an essentially finite acoustic pulse in the air even though the ice sheet vibrates for a relatively long period of time. However the sound pulse in the air gets progressively longer as the distance from the bouncing point increases. The initial part of the sound is emanating from the ice immediately surrounding the listener while the later part of the pulse is dominated by the sound that travelled through the air from the point where the stone hit the ice.

>>

Modelled acoustic pulses at a height of 1.5 m above the ice and at a distance of 10 m (blue) and 40 m (green) from the point where the stone hit the ice. To the human ear, the pulses created by the model sounds just like the real acoustic pulses.

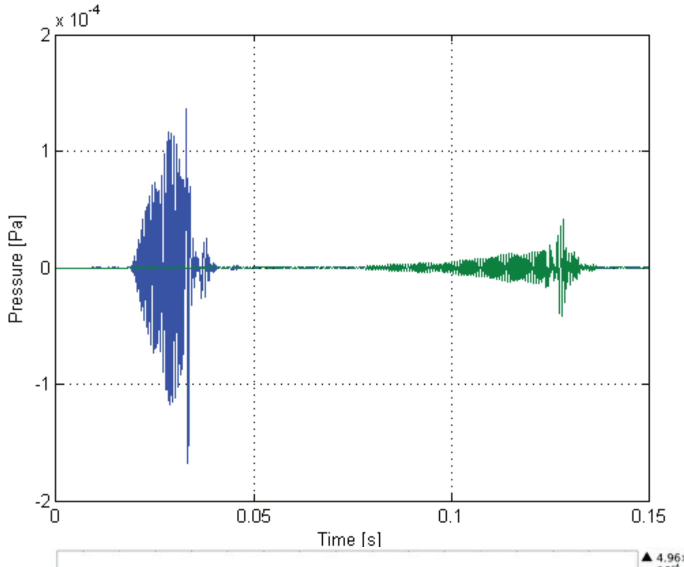


The picture shows the modelled acoustic pressure field (air cross section) in 5 m (y-axis) of air above the 2 cm ice sheet at 15 ms after the bouncing moment. It is evident from the picture that the sound reaching the human ear located somewhere along the x-axis (radius) is composed of both sound emanating directly from the ice surface beneath and a pulse travelling through the air from the point where the stone hit the ice. The sound coming from the ice below is the result of a horizontally propagating vertical vibration of the ice acting on the air above. The pitch of the pulses depends on the thickness of the ice. A 5 cm ice sheet results in a frequency peak at approximately 700 Hz and a 2 cm ice sheet results in a peak at approximately 1500 Hz.

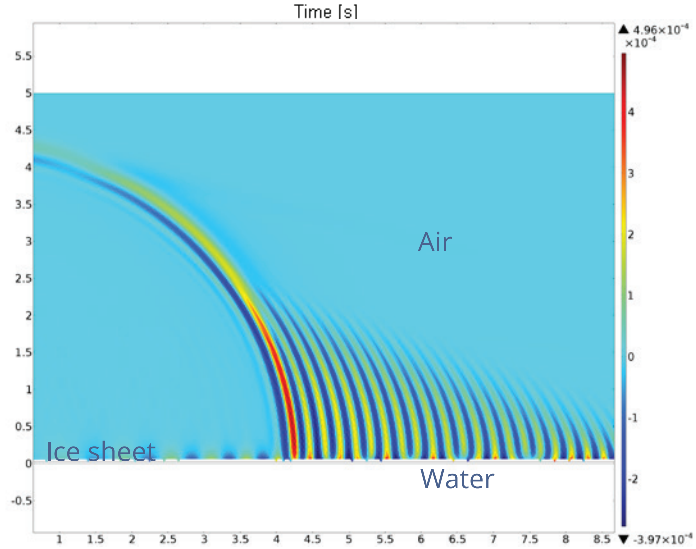


This picture shows the pulse (acoustic pressure) after 100 ms, propagating from the point of impact in the centre. The top surface of the model is a matched impedance surface allowing the sound to properly leave the model at 2 m above the surface of the ice. Below the ice is the water volume and the model is equally good at modelling what a diver in the water would hear from the skipping stone.

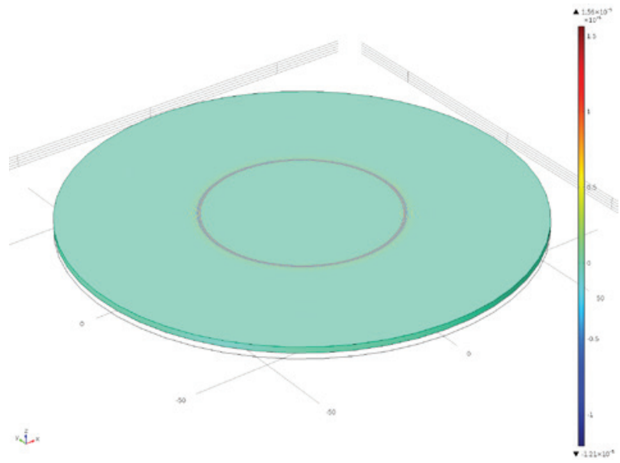

**ADVANTAGES WITH MODELLING**

A reduced need for practical tests and prototypes is not the only advantage with modelling. It also visualises phenomena otherwise impossible to see or measure and thereby increases the intuitive understanding of the modelled process. Examples are sound and mechanical stress fields that are made visual by the modelling tool.

In practical tests and measurements there is always noise present, and it is sometimes difficult to independently change a single parameter in order to study its influence on the total system. In modelling, individual parameters can be altered at will and noise is not present unless it is deliberately added to the model.

## ULTRASONIC UNDERWATER DATA COMMUNICATION LINK

The performance of an ultrasonic underwater data communication link depends on several factors such as the transducer design, form of modulation and the acoustic propagation situation. The signal processing necessary for the modulation converting data bits to acoustic pulses can be modelled in Matlab and the performance of the transducer together with the acoustic circumstances can be modelled using FEM. Combined these two tools let us model and fine tune the entire system before moving on to prototype development and practical tests.

An ultrasonic transducer is often resonant and the efficiency is strongly dependant on the frequency of the carrier chosen for the data modulation. It is thus essential to know the resonance frequency of the transducer and also the sound radiation pattern as the output power varies with the radiation direction.

The piezo electric element vibrates and the vibrations spread through the entire transducer structure and perhaps further. The acoustic radiation field depend on the entire mechanical design of the transducer and modelling reveal how different parts of the transducer vibrate and affect the emitted field. FEM enable us to test and modify the mechanical design in order to reach an optimal acoustic output.

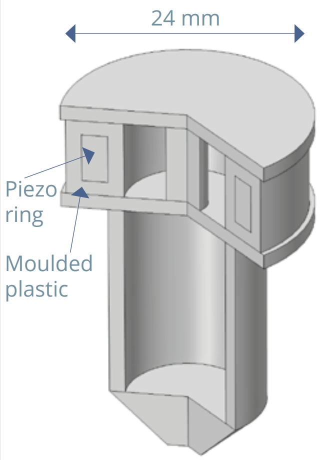


The picture shows a transducer model complete with piezoelectric ring and housing. Moulded plastic covers the piezo ring and protects it from the water. It is important to get the design correct and use a suitable material to achieve good mechanical impedance match between the piezo ring and the surrounding water.

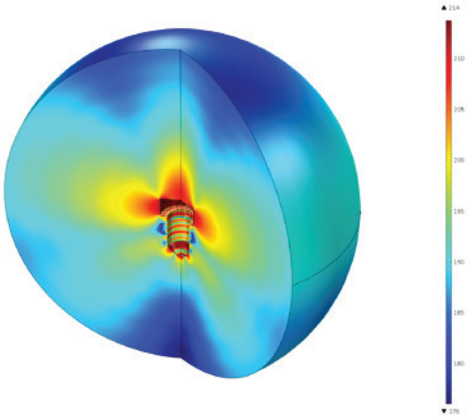

This is a visualisation of the acoustic level field in the water surrounding the transducer. It is obvious that the intensity of the sound at 50 kHz varies with both the aspect angle and distance and that the maximum communication range will depend on the orientation of the transducer. The colour scale is in dB rel. 1µPa.

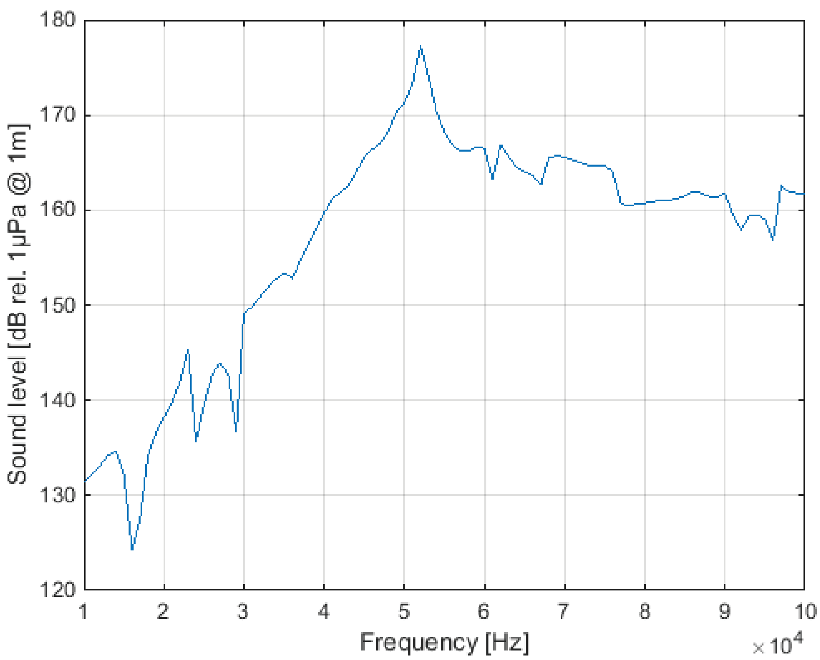


The graph shows the acoustic output level referred to a distance of 1 m in the horizontal plane of the transducer. Clearly the transducer is resonant and is suitable for a carrier frequency of approximately 50 kHz. The model includes the piezoelectric effect and the transducer is excited by 35 Vrms.

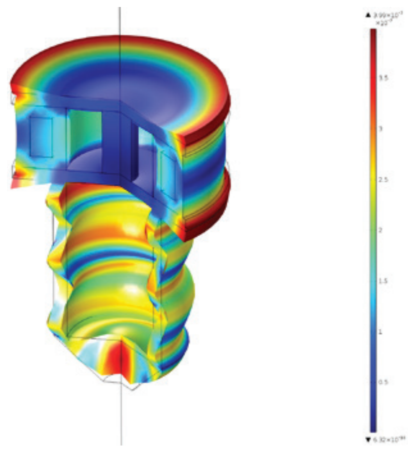

The picture shows (greatly exaggerated) how the entire transducer vibrates when the piezo ring is excited by the electric 50 kHz signal. This Comsol Multiphysics model includes piezoelectricity, structural dynamics and acoustics and thus models the complete chain from electric input to acoustic output.

# INDUSTRIAL CONNECTED THINGS

BY: Kristoffer Koch
Senior Development Engineer
Data Respons

In the industry, connecting large networks of sensors and actuators with smart logic is nothing new. While these networks are typically not Internet connected, they are however networked "things". Are there lessons to be learned from the industry when we build the Internet of Things for the future?

In a typical car bought the last decade, there are over 70 computers connected in an internal network. There are sensors that sense if you crash, triggering safety mechanism to save your life. There is also a network of sensors monitoring the engine in order to keep performance and emissions within the tolerated limits, or to control the charging of batteries in hybrid or electric cars. In modern ships, we find highly complex controller networks, so much that a modern ship is more like a floating factory, managing power, controlling ballast tanks and thrusters based on a number of inputs from the bridge and from sensors, and also providing the operator on the bridge with crucial information. The thrusters can be procured by one vendor, while the dynamic positioning systems can be bought by a different vendor, without sacrificing interoperability.

# What challenges are there, and how can we meet them?

### NETWORK OF THINGS

The key to success for these kinds of networks is that the machines speak the same language. Information a sensor gathers is typically not displayed directly to a human, but must rather be interpreted and "understood" by another machine. A climate control system e.g. needs to know what units temperature and humidity is measured and how this information is encoded on the network. For this to work, it is important that devices adheres to standards.

In automotive, cars typically use Controller Area Network (CAN) for communication. First, CAN was mostly a communication standard defining how raw information should be sent on the wire, but have later been extended with more standards that defines in detail behavior of specific applications. The CANOpen standards provides a wide range of device profiles for a myriad of applications, for everything from medical tomography to large crane installations. There is also a standard for writing a standard if none of the existing fits your application. The marine industry have also extended CANOpen by standardizing how high reliability redundant networks can be built for ships, maintaining control over the ship even in the events such as fire disabling parts of the ship.

There is a lot of good work that have been put into making such standards for interoperability, but how does it relate to a smarter and more connected Internet of Things? Without standards, we risk ending up with an Internet of incompatible things, or as Jean-Louis Gassée from the Apple initial alumni team put it, we end up with a "basket of remotes". Today we often see each Internet of Thing vendor providing their own app for controlling their devices, but they provide no way to integrate the different devices into doing new smart things in a seamless way. However, home automation standards such as ZigBee or Z-wave takes some of the same design decisions as industrial standards, and specifies how different kinds of devices should operate in order to be compliant with the standards.

### CHALLENGES

While industrial networks have typically been designed with safety and reliability in mind, security is another issue. Features such as authentication, authorizon and confidentiality are typically not subjects that are being addressed by industrial standards. If we are going to apply the experiences from industrial networked devices to the Internet of things, these are issues that needs to be addressed. The security expert Bruce Schneier compares the current situation to the general computer security of the mid-90s, when the Internet first saw widespread adoption, but without software or security practices ready for this revolution.

In 2010, security researchers studied the tire pressure sensors of a car. Since it is hard to make wired connections to a rotating wheel, the tire sensors were made wireless, and this was what interested the security researchers. By forging malicious data into the wireless receiver of the car, the researchers where able to take full control of the internal network of the car, and were able to monitor and control critical subsystems such as engine control and braking. That this breach was possible was mainly due to a design that did not take into account security attacks of this kind, but a design that was based on the assumption of security by isolation of the network.

A conceptually similar attack was brought against nuclear enrichment centrifuges in Iran, with the internet worm stuxnet. Even though the enrichment plant was not connected to the Internet, this worm also spread on USB thumb drives, and in the end, the attack succeeded in spinning the centrifuges into destruction.

A basis for making secure software for devices is to have a clear and well defined communication protocol. Typically it is seen that proprietary solutions have had less scrutiny and discussion than openly developed standards. One interesting case is the industrial bus HART, which have been extended to a wireless standard, WirelessHART. In this standard, in addition to the more traditional reliability and safety concerns, security is also addressed, keeping unauthorized devices out of the network, and keeping messages confidential and authenticated using encryption.

But even a good design can have a buggy implementation. While we have many good practices for achieving high software quality, it is sadly beyond the state of the art to implement perfect software without security holes. A device vendor that ships a product must acknowledge this in order to maintain a satisfactory level of security. Security, like hygiene in a hospital, must be viewed as an always ongoing process, embracing the whole lifetime of the products. In 2014, both General Motors and Tesla was ordered to do a recall because of a fire hazard when using damaged charging cables. GM had to physically bring in all the cars for repair, while Tesla performed an over-the-air software update to detect bad cables, and then limit the currents to safe levels.

### NEED FOR EXPERIENCE

While we are seeing a growth of new devices targeted for the consumer market, it is not clear how this myriad of devices should be connected in a meaningful way. The industry have been successful in defining standards and protocols for such use, but there is still work to be done with security before exporting these ideas to the mass market. Also, it is also wise for the industry to try to learn lessons from the development in the consumer market, where the innovation moves at a even faster pace.


STM32F446

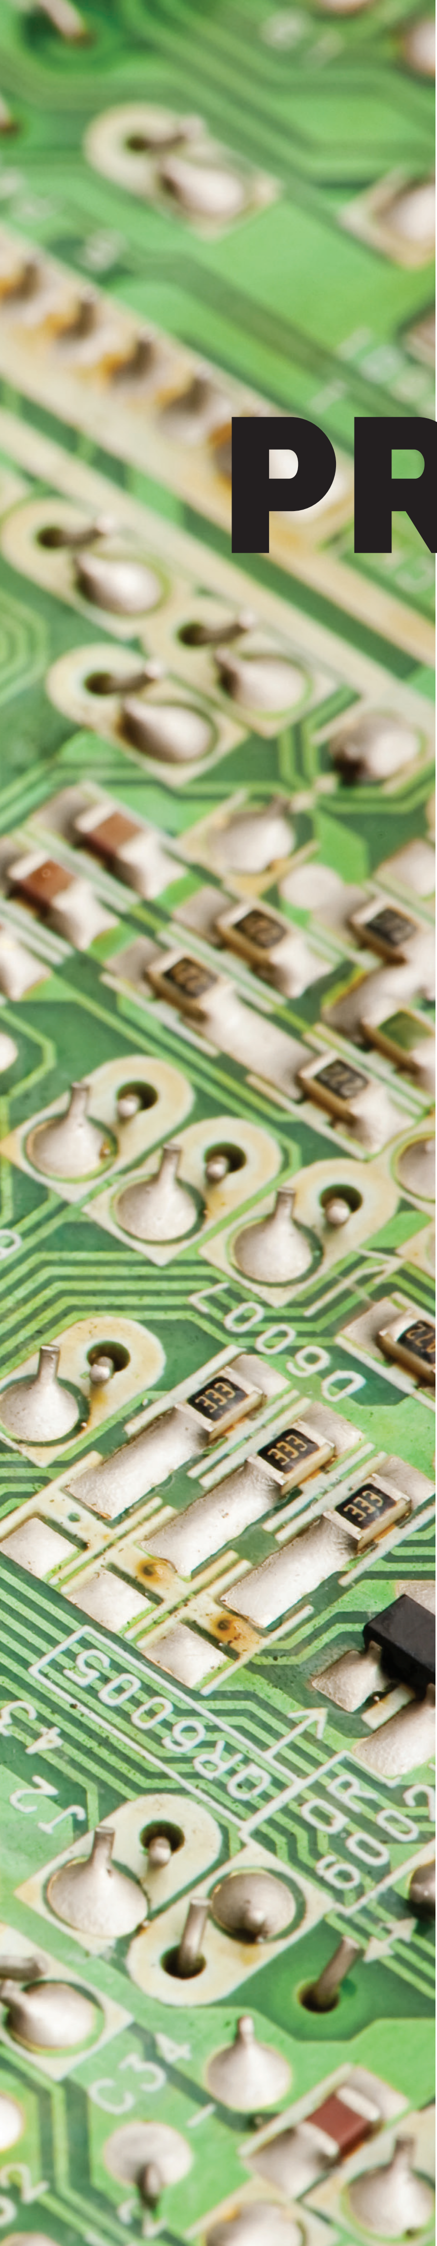

# PROS & CONS

## - OF USING STM32CUBEMX CODE GENERATION TOOL INSTEAD OF MANUALLY WRITING DRIVERS FOR AN ARM CORTEX-M MICRO-CONTROLLER.

A new trend is emerging from several microcontroller manufacturers. Driver code can now be configured and generated using provided tools. This article will take a closer look at a tool named STM32Cube-MX (from here on called Cube) from ST Microelectronics. It is made for their STM32, an ARM based family of microcontrollers. Cube is a graphical tool for selecting, configuring and generating project reports and code. It currently supports all STM32 microcontrollers, both the STM32F and STM32L series, but is only available for the Windows operating system so far. It can be run as a stand-alone application or as a plugin for the Eclipse integrated development environment (IDE).

BY: Patrick Hisni Brataas
Development Engineer
Data Respons

Cube is a great tool for narrowing down the possible choices when selecting a microcontroller. It let the user choose a microcontroller based on the required peripherals, family of microcontrollers, package type, flash size, ram size, minimum number of input/output pins and so on.

For getting started quick and easy with prototyping, a board support package and a lot of example project are readily available for the three IDEs: Keil uVision, IAR Embedded Workbench and Atollic TrueStudio.

### CONFIGURATION OF DRIVERS

After selecting the correct microcontroller for the application, the user interface has four main views. Following is a brief explanation of these four views.

### PINOUT VIEW

This view is shown in figure 1 (next page). It includes a visualization of the microcontroller and its pins and has a vertical side toolbar. This is where the required peripherals are selected.

>>

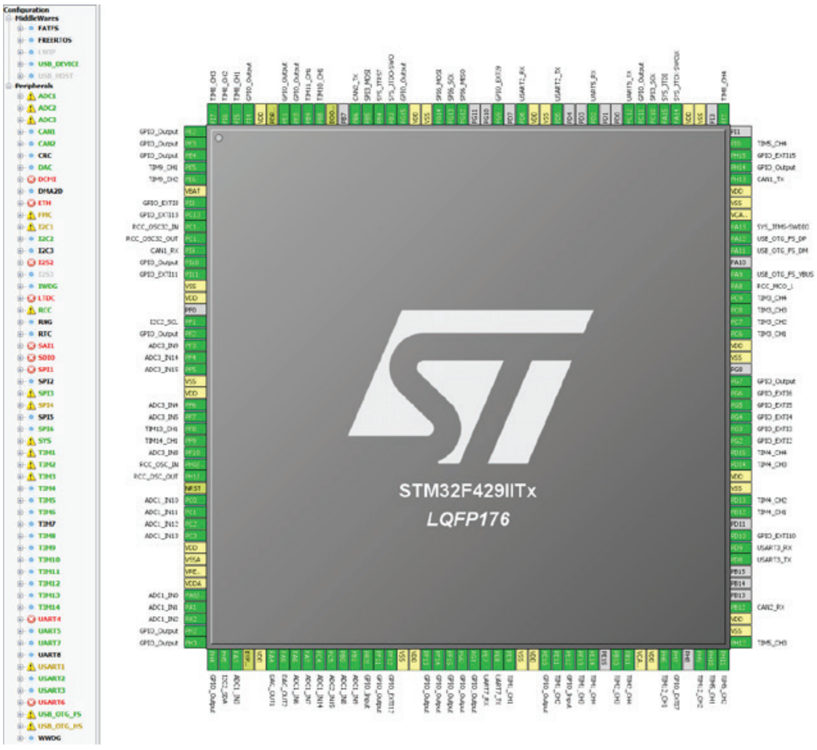


Figure 1

Selection of drivers can be made by choosing a pin that support the peripheral directly or by selecting the particular peripheral from the toolbar. The tool will automatically assign the peripheral to the appropriate pins. When using the toolbar, it will solve pin conflicts by moving a conflicted peripheral to unused pins that also support the peripheral. Sometimes this automatic conflict resolver might not be wanted and therefore it is possible to lock a peripheral to a pin if necessary. Both ways of selecting peripherals do not allow a combination that is not supported by the selected microcontroller.

This view has a focus on pinout and therefore only settings related to each pins possible configuration are set. For example, one can choose to enable SPI in either full duplex, receive only or transmit only. The three possible selections all enables the SPI peripheral, but they also have an effect on which pins are utilized. In this view, it is not possible to set the baud rate, data size, endianness, prescaler, clock polarity, etc., as they do not affect the pinout. Another example is enabling an ADC and ADC channel in the pinout view. This affects the pinout, but the selected sampling rate, data conversion mode, resolution, etc., does not.

The pinout view lets the developer enable and configure peripherals that affect the pinout. The only exception to this is that some middleware libraries can be enabled from this view even though they do not change the pinout. An example of this is FreeRTOS, which is a real-time operating system.
More settings related to the peripherals can be chosen in the configuration view.

## CONFIGURATION VIEW

The configuration view shows all the enabled peripherals and middleware libraries. In addition, it is possible to configure watchdog functionality, DMA transfers, enable the different interrupts and set additional clock and reset behavior.

In this view, the configuration of the peripherals is done. It is possible to set the baud rate, data size, endianness, prescaler, clock polarity, etc. of the SPI peripherals or the sampling rate, data conversion mode, resolution, etc., of the ADCs. An example is shown in figure 2.

Figure 2.

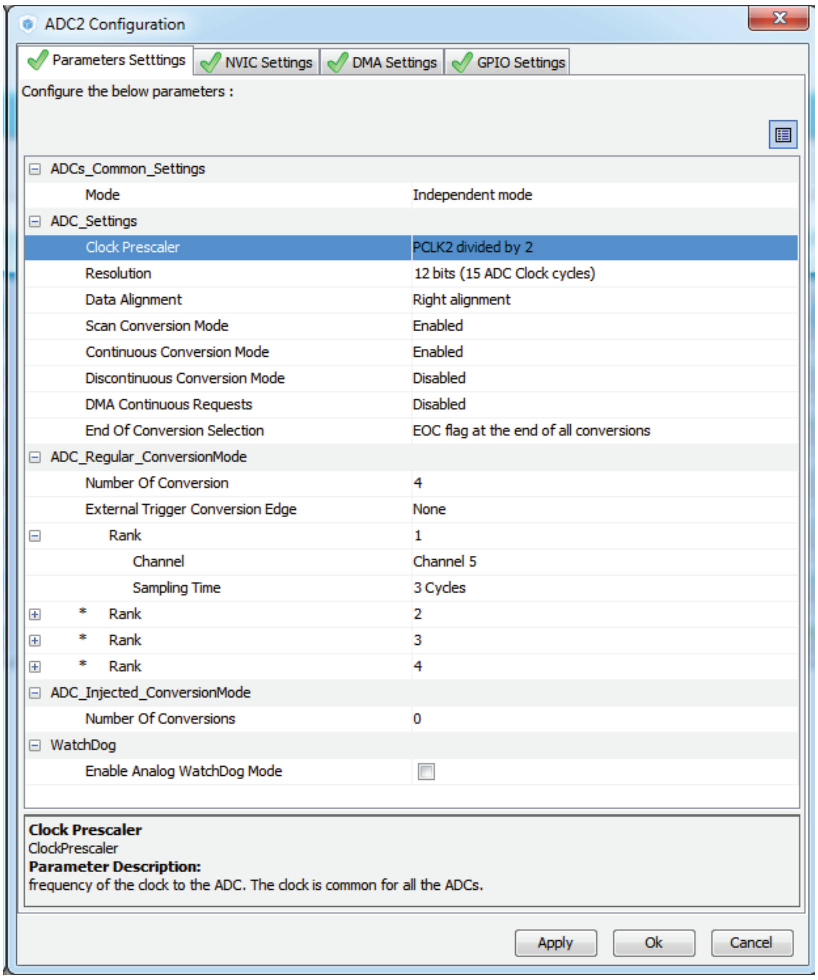

## CLOCK CONFIGURATION VIEW

All clock configurations should be done in the clock configuration view. This view is shown in figure 3 and provides a good overview of the clock tree.

This view enables the developer to choose between external and internal clock sources. Clock frequencies in the clock hierarchy are automatically calculated by setting the oscillator frequency used and adjusting the many prescalers and PLLs.

Invalid clock configurations are clearly shown in red, which makes it easy to discover and fix any incompatibilities.

## POWER CONSUMPTION CALCULATOR VIEW

This view is used to calculate approximately how much current the microcontroller, with the selected peripherals and settings, will draw in different modes and in average. Control of the current consumption is important for low-power applications and this tool greatly helps the developer.

The view lets the developer set parameters as supply voltage, clock frequency, run/sleep/standby mode, RAM voltage, enabled peripherals, etc.

In lower power applications, the microcontroller will often sleep most of the time, only waking up periodically to check for events or when an interrupt occurs. This means that the microcontroller most likely has at least two modes with very different parameter values. The power consumption view makes it possible to add a sequence of steps, where one step equals a mode for a set time interval. When all the steps corresponding to the different modes have been added, a graphical representation is created. An example of this is shown in figure 4.

Figure 4.

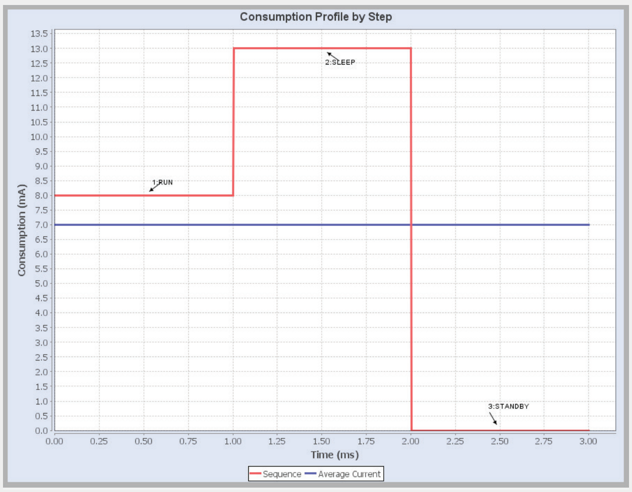


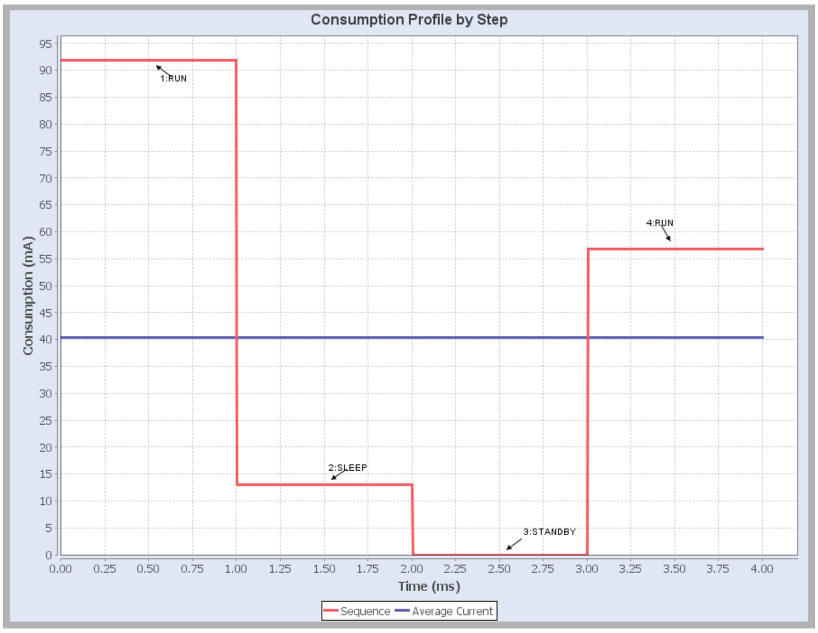


Figure 3.

In addition, the most regular batteries can be selected from a list and the approximate battery life can be estimated for a full charge.

## GENERATE PROJECT REPORTS

Reports can be created by a single click and contains much useful information.

The generated report contains information like:

- The selected microcontroller.
- What version of Cube and firmware package version that was used to generate the code.
- Which compiler is used and what version.
- An overview of microcontroller pins as shown in the pinout view.
- Pin list with mappings to package pin number, internal pin number/port, peripheral active on pin and a user-selected label.
- All power consumption calculations done in the power consumption calculator view.

## CODE GENERATION

After selecting the required peripherals in the pinout view and configuring the clocks and peripherals, it is possible to generate the initialization code. Not only is the code generated, but all necessary project files for a chosen IDE as well. In addition to this, template files are provided in the project, which gives some guidance on how one might structure the code.

In the generated files, there are commented sections where the custom code should be inserted. It is very important that the code the developer writes in generated files is written within these sections. Otherwise it will be gone when the project is regenerated. Even a minor change like changing the baud rate for a peripheral is recommended doing in Cube and not in the generated source file itself. Not because this is any faster (or especially slower), but it will prevent unnecessary possible errors and keep the generated documentation updated. In addition, if at a later time more major changes are to be done and the Cube project is not up to date with the generated code, one must remember to add the changes done to the source file. Always keep Cube updated!

Cube uses the STM32 hardware abstraction layer (HAL) library to create the initialization code, which makes it a lot easier to migrate between STM32 microcontrollers if needed.

By default, all generated code is put in a header and source file. The generated files can be separated into headers and

>>

source files for each type of peripheral to get a better overview by adjusting the settings in the top toolbar. This does not include middleware libraries.

### HOW TO USE THE GENERATED CODE AND HAL

As mentioned, the STM32 HAL library is used. After the code is generated, everything should be ready to use the HAL library to control the peripherals. The syntax of the HAL library is shown in the table below:

It is the function calls as shown first in the table that should be used to control the behavior of the peripherals. To start a basic timer the HAL_TIM_Base_Start() can be called or to send data over UART with DMA one could call HAL_UART_Transmit_DMA().

### PROS, CONS AND EXPERIENCES

The Cube is a new software with its initial release in February 2014 and it still has some bugs. Bugs can be something as trivial as a missing or repeated line in the generated code or an error in the user interface preventing the use of an actual valid setting. Despite this, it saves a lot of time. Additionally, six minor version updates have been released in one year, so it is rapidly becoming better.

It is easy to get a good overview of the peripherals used and which ones are still available. Making changes in the peripheral configuration is fast and easy, but keeping the default structure generated by Cube will probably save a lot of time when the code has to be regenerated.

| SYNTAX | EXPLANATION |
|---|---|
| HAL_PERIPHERALTYPE_Function(PERIHERAL) | Function call to perform a function on the peripheraltypes peripheral. |
| __HAL_PERIPHERALTYPE_FUNCTION(PERIPHERAL) | Macro to perform some function on the peripheraltypes peripheral register. |
| __ PERIPHERAL_FUNCTION() | Macro to perform some function on a peripheral. |

The second line in the table are macros that helps the developer change register values. The reason to use macros is that they are more portable and reduces the chance of setting the wrong bit. Depending on the application, these macros must sometimes be used. For example when changing the ADC sampling rate between two frequencies while the application is running.

The last line in the table can typically only enable and disable clocks to peripherals and functionality related to reset.

When new code is generated, one should check the initialization code that has been changed or added. A quick look can reveal missing lines or wrong settings that might otherwise be hard to detect. It is possible to add custom code to the initialization code that will not be removed when the code is regenerated. This makes it possible to correct many code generation errors.

All necessary callback functions are already prototyped and only the definitions have to be written. If more than one peripheral or similar can trigger one interrupt, there is already a handler determining where it was called from and the appropriate callback function is called.

If the macros provided by the HAL library is used, then the data sheet and reference manual are your friends to avoid errors.
It should be easy to set up the project in IDEs other than the three officially supported ones. Atollic TrueStudio uses the GCC compiler, which is also supported by several other IDEs.

Cube currently does not support generation of flash initialization code for enabling reading and writing to flash, but it might be supported in the future.

The pinout view automatically checks for conflicts and resolves them if possible. A pin list can be generated that is useful for hardware developers designing the custom hardware. Even if this tool is not used for code generation, it is useful for setting the pinout and determine if the selected combination of peripherals is valid.

The normal workflow is setting the required peripherals, configure the clock and then configure the peripherals. The power consumption calculator is optional, but if used it should be used when everything else is set and configured.
The Cube uses the HAL library and therefore ensures that the code can easily be ported to any other STM32 chip with minor effort, as long as the required hardware functionality is present.

# **INSIDE** WRITERS

## HALDOR HUSBY

Principal Development Engineer
Data Respons

MASc Electrical Engineering, University of Toronto
Siv.Ing. Electronics, Norwegian University of Science and Technology

## AKSEL BONDØ

Senior Development Engineer
Data Respons

MSc Systems Engineering, Stevens Institute of Technology

## FRODE SØRENGEN

Senior Development Engineer
Data Respons

Siv.Ing. Electronics, Norwegian University of Science and Technology

## ERIK ASPLUND

Principal Development Engineer
Data Respons

MSc. Electrical Engineering, Royal Institute of Technology (KTH), Stockholm, Sweden.

## KRISTOFFER KOCH

Senior Development Engineer
Data Respons

Master of Technology, Electrical Engineering, Digital Design
Norwegian University of Science and Technology (NTNU)

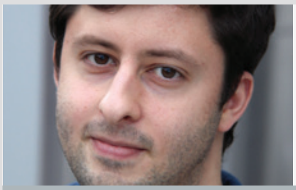

## PATRICK HISNI BRATAAS

Development Engineer
Data Respons

MSc Electronics and Computer Technology, Medical Technology, University of Oslo.

## ALEXANDER SVENSEN

Development Engineer
Data Respons

MSc Robotics and Intelligent Systems, University of Oslo.

## ANDRE FIRING

Development Engineer
Data Respons Alumni

## CONTRIBUTORS

### BENT FUREVIK

Development Engineer
Data Respons

### GUNNAR ANDREAS SKOGVOLD

Development Engineer
Data Respons


PUBLISHER:
**Data Respons ASA,**
Sandviksveien 26, 1363 Høvik
Tel: +47 66 11 20 00
info@datarespons.no

EDITOR-IN-CHIEF:
**Kenneth Ragnvaldsen**
CEO, Data Respons

EDITOR:
**Elisabeth Andenæs,**
Senior Communications Consultant
Data Respons
Tel: +47 92 20 30 03
Email: ean@datarespons.no

TECHNICAL EDITOR:
**Ivar A. Melhuus Sehm**
Director R&D Services, Data Respons
ise@datarespons.no

# MARITIME CERTIFIED PANEL PCs

Data Respons' latest series of rugged Panel PCs designed and tested for maritime and offshore environments.

## 8" Maritime Panel PC

The 8" Maritime Panel PC is a rugged low power multi-touch computer based on a Quad Core Intel® Atom™ Bay Trail processor and has a fanless, minimalistic design with a 8"panel with multi-touch sensor.

## 12" Maritime Panel PC

The 12" Bay Trail Maritime Panel PC is based on a Quad Core Intel® Atom™ Bay Trail processor with low power consumption and has a fanless, minimalistic design with a 12" panel with multi-touch sensor

## 24" Maritime Panel PC

The 24" panel multi-touch computer is based on an Intel® Core™ i5 (optional Core i7) Ultra Low Voltage CPU, SSD storage with no rotational parts and up to 16GB DDR3 memory.