INTERRUPT INSIDE

An in-depth magazine about IoT and embedded technology from Data Respons 1, 2018

Internet of Things:

Developing an emergency communication device for disaster relief work

- utilising mesh technology

3D Printing:
Improving motion control in a bipolar printer

Subsea:
Autonomous Cable Survey with Magnetometers

Methodology:
How optimal is your project approach?

data respons

WHAT'S **INSIDE**

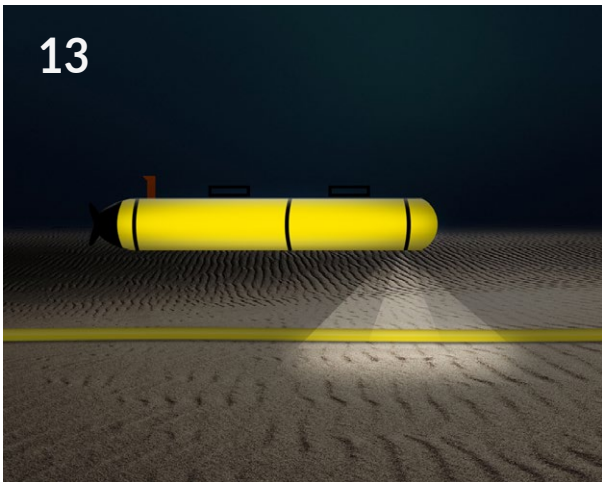

EDITOR

When we say that what we do is digitalisation of the industries of tomorrow, we really mean it. The way we use data and connected devices today is currently paving the way for a vast amount of improvements throughout the industries. And, there are winnings to be collected; not only do many of these innovations optimize our customers manufacturing and logistics, but they leave positive footprints by saving energy, cost, manpower and even lives.

In this issue of Interrupt Inside we will look at how technology and the internet of things can overcome the challenges of reaching people in areas of low connectivity. Our Danish team have helped develop the Reachi Device that uses Mesh technology to create network for the aid- and rescue workers enabling them to stay in touch and communicate in disaster areas where all communications have broken down. This is an area where we see clearly what IoT really is!

This issue also touches upon one of Data Respons' core strengths, our competence within effective and tailored project management. The article, How optimal is your project approach, gives you an overview of the typical methodologies and how they match the various characteristics for system development projects.

All articles presented in Interrupt Inside are written by our own specialist engineers and select guest writers. We welcome any feedback and suggestions from our readers.

Enjoy the reading!

KENNETH RAGNVALDSEN

Develo
COMMU
for

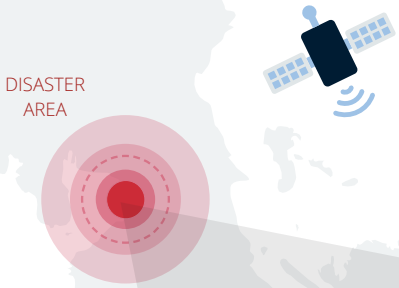

BY: Ole Larsen
Software Development Engineer
TechPeople

Co-writer: Thomas Halkier
CEO, Neocortec

Being both shock proof and water-tight, the Reachi device will stand up to extreme weather conditions. On top of being robust, the Reachi device is designed with user friendliness in mind. And the device is extremely energy effective, so that its battery can last as long as possible in areas where power supply is shaky or even completely absent.

CHOOSING THE RIGHT COMPONENTS

TechPeople has assisted LinkAiders in choosing the right components for the Reachi device, including rechargeable battery, micro controller etc. and stitching it all together to achieve the robustness needed.

The project has come quite far, with prototypes having been tested twice in the Philippines already. Also in the Philippines, a pilot with 1,000 devices is to be scheduled for sometime during 2018. At that time one of the key features of the Reachi device will be tested in a real life setting: The devices functioning together to form a dynamic and flexible Mesh network.

NEW NETWORK

As mentioned, the Reachi device has to function in spite of communication infrastructure in an area being destroyed. That requirement, combined with the need for ultra low power consumption, made LinkAiders turn to the Danish wireless mesh network company NeoCortec and their NeoMesh technology.
TechPeople took up the challenge to integrate NeoMesh into the Reachi device. NeoMesh is developed for IOT applications and allows for up to 65K mobile nodes. Moreover, it is able to handle dynamic topologies in real-time. As opposed to other technologies in the IoT workspace, NeoMesh allows any node in the network to dynamically change position. This feature makes NeoMesh ideal for the Reachi use case, as emergency workers equipped with a Reachi device move about and change position constantly, while the devices still can play their part in the network, receiving and transmitting data.

oping an emergency
UNICATION DEVICE
disaster relief work

In the immediate aftermath of a natural disaster, with all communication infrastructure destroyed, emergency workers need to send damage reports from their local communities to relief coordinators and other authorities. To do that, the Danish company LinkAiders is developing the Reachi device, a communication device for use under extreme conditions. LinkAiders cooperates with the Danish Red Cross, and the Reachi device will be pilot tested in the Philippines during 2018. The Danish tech consulting company TechPeople, owned by Data Respons, is helping LinkAiders design a solution that will function under the toughest conditions.

This flexibility is possible due to the NeoMesh Speed Routing Protocol. It replaces a central Network Manager with autonomous intelligent nodes, enabling all network nodes to link to each other automatically and dynamically, forming one single network that works, even if nodes change position or are replaced.

The NeoMesh routing protocol routes data seamlessly through the network and eliminates any concern in performance created by obstacles in the RF path, nodes being blocked or simply moving around within the network. Weak spots in a real life network can easily be fixed by just adding another node. Given it has the right network ID, it automatically becomes a part of the network.

FLEXIBLE COMMUNICATION INFRASTRUCTURE

Unlike other routing protocols the Speed Routing Protocol does not create the exact route from A to B in advance. Each NeoMesh node maintains a knowledge of which of the nearest nodes would be the best for the next hop. While data travels through the network this knowledge of the best next hop is used to decide the route of the data. The knowledge is kept up-to-date in real-time and the route is adjusted dynamically according to changes in topology and link quality.

Utilizing these network features the developers have been able to design an extremely flexible communication infrastructure. The network is divided into subsections, consisting of 1,000 Reachi device-equipped emergency workers. For this dedicated long-range deployment NeoCortec has developed a special radio module, with a line-of-sight range of >2,5 km.

1,000 SUB-NETWORKS

Each sub-network hosts a NeoMesh gateway transmitting data via a satellite uplink. Thus vital information can be sent to relief coordinators, even with a country's communication infrastructure completely damaged.

To cover the whole of the Philippines, the Reachi deployment plan envisions 1,000 sub-networks, each with its own satellite-enabled gateway. This means a million devices all in all, as each sub-network consists of 1,000 Reachi devices.
As the Philippines is one of the world's most disaster-prone countries, LinkAiders have chosen a tough environment to test their device. If successful it will not only provide help where help is desperately needed. The Reachi system will also set new standards for dynamic and flexible IOT network solutions.

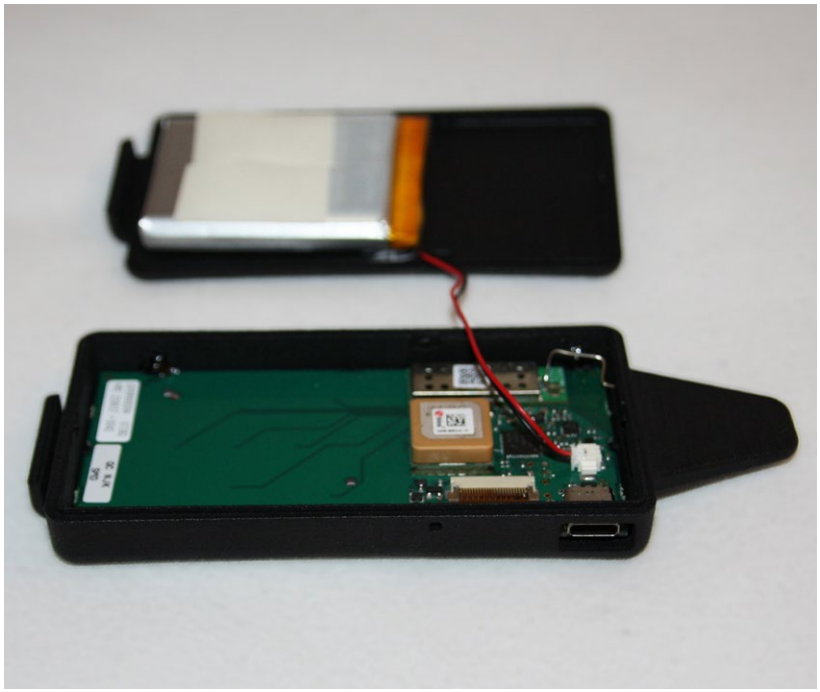

Smart from the inside: The Reachi device

Performance-aware energy-efficient data cache accesses

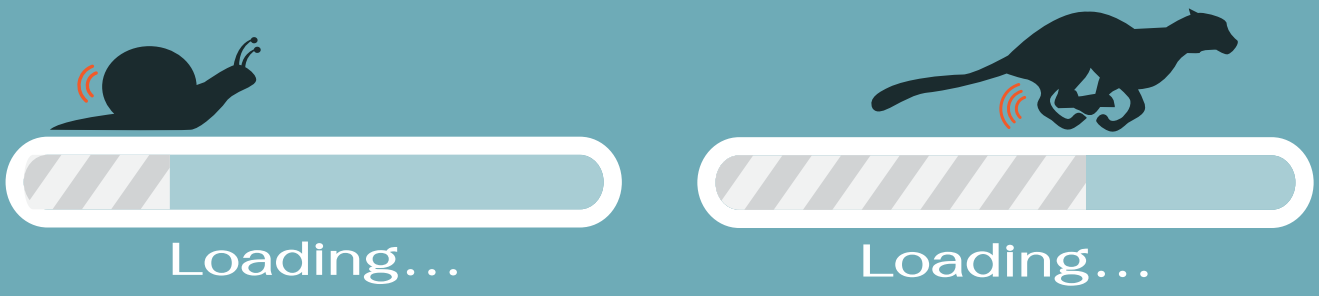

With the increased growth of Internet of Things (IoT) devices, the need for energy efficient computing systems is more important than ever. Many of these systems are battery operated and often in places where recharging or replacing the batteries is difficult. This is why researchers and the semiconductor industry are using a significant amount of resources in increasing energy efficiency by developing embedded systems such that it consumes as little power as possible while still increasing the performance.

BY: Salahuddin Asjad
Development Engineer
Data Respons

In the beginning of the 90s, the slowness of memory speed compared to processor speed was felt, known as *the memory gap*. This led to an introduction of cache memories, small amount of fast, but expensive memory, used to increase the performance of load operations by keeping required data closer to the processor than the main memory. However, because cache memories, L1 DC in particular, are optimized for performance rather than energy consumption, the energy consumed by cache memories can account for a significant amount of the total energy consumed by microprocessor-based architectures. Techniques such as Speculative Halt Tag Access (SHA) and Early Load Data Dependence Detection (ELD[3]), based on way-halting and sequential loads, respectively, can be used to reduce the energy dissipation without sacrificing the performance of L1 DC.

CACHE MEMORIES

To provide the Central Processing Unit (CPU) with necessary data as quick as possible, the most frequenctly used data is stored in the caches that are placed closer to the CPU. Figure 1a shows a typical memory hierarchy existing in today's computers. L1 cache is usually placed on-chip, such that it is possible to exploit locality by keeping data likely to be reused as close as possible to the CPU. If there is a cache miss in L1 DC, the search request will begin for L2 DC, which is often larger than L1 DC, thus results in higher latency. With each cache miss, the search proceeds to the next level memory until the requested data is found.

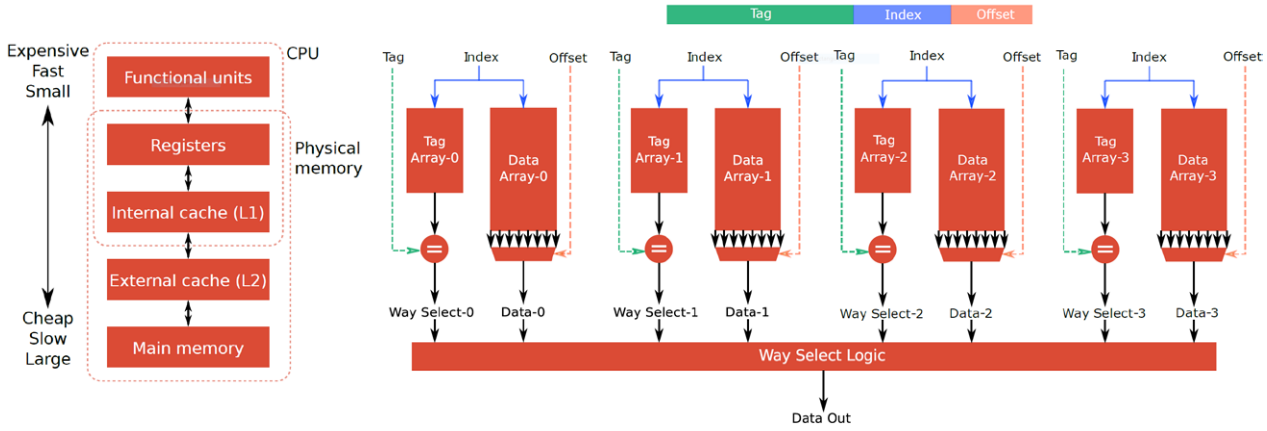

(a) A hierarchy of memories

(b) set associative L1 DC

Figure 1

In order to reduce the search time for data requests, the caches often have a restricted placement policy, known as cache associativity. Cache hits are then detected through an associative search of all tags, instead of searching through the entire cache. Conventional L1 DCs are often set-associative caches with low associativity, where the latency of load operations is optimized by accessing all ways with the same tag address in parallel, as shown in Figure 1b. However, this results in a signicant amount of wasted energy as only data from one way is used. To reduce the energy consumption, numerous cache architectures, such as way-prediction, way-shutdown and highly-associative have been proposed. However, these optimization techniques leads to increased latency and complexity, which makes them unattractive for L1 DCs.

SPECULATIVE HALT-TAG ACCESS (SHA)

Practical way-halting by speculatively accessing halt tags, is a cache architecture that can reduce the energy dissipation without increasing the latency and complexity. That is accompished by halting cache ways that cannot possibly contain the requested data, thus avoid accessing all ways with the same index unnecessarily. The technique is based on the observation that the displacement address often is small and usually only change the offset of the relative memory address, see Figure 2. This makes it possible to require the halt tags, low-order bits of the tag, using the base address, in parallel with the memory address calculation in the address generation stage. Since the base address and the displacement address is available in the address generation stage, a comparison of the tag and index bits of the base address and effective address can be done to determine if the displacement is small, before the data access stage.

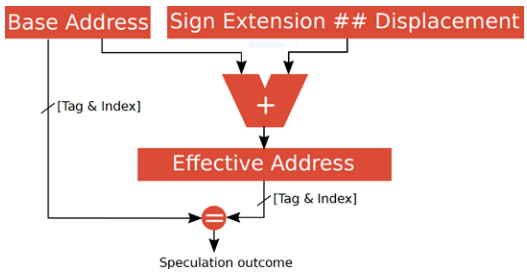

Figure 2:
Displacement check

When the displacement check succeeds, the halt tags can be accessed from the halt-tag cache, such that a halt tag check is performed before each L1 DC load operation. Halt tag bits of the base address are compared with the halt tag bits of each cache way accessed. If there is a match between the halt tag from the base address and the halt tag from a way, the bit corresponding to the way in the response vector is enabled to indicate that there is a halt-tag match, see Figure 3. Only the ways that have enabled bit are then accessed in the next pipeline.

The technique, has no performance penalty and adds very little complexity to a conventional processor core design. Although the displacement address often is small, the SHA technique can not be used for displacement addresses that will change the tag or index bits during the address calculation. When the displacement address is too large for the SHA technique, Early Load Data Dependence Detection (ELD[3]) can be used to reduce the energy dissipation.

EARLY LOAD DATA DEPENDENCE DETECTION (ELD[3])

Early Load Data Dependence Detection

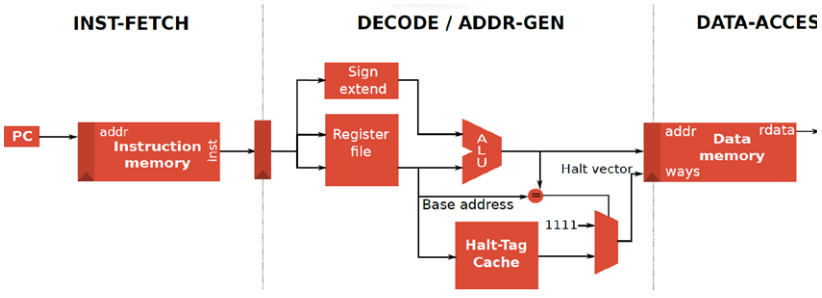

Figure 3:
SHA implemented in a three-stage pipeline.

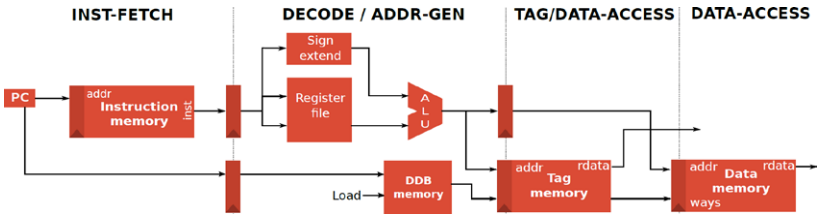

Figure 4: ELD³ implemented in a four-stage pipeline.

(ELD³) is an approach that can detect if the load operation has a data dependency with a following instruction that will cause a pipeline stall. If there is no data dependency between the load instruction and the following instructions, the load operation is performed sequentially where all tag ways are accessed, but only one data way in which the data resides, is accessed in the next clock cycle. However, if there is a data dependency, the data access is performed parallelly where both tag and data ways are accessed in the tag/data access stage, shown in Figure 4.

In order to decide whether the data ways should be accessed sequentially or in parallel, the information about data dependency between the load instruction and the following instructions must be available at the time of load operation. In a conventional in-order pipeline processor, the information must be available before the end of the address generation stage. Commonly, it is possible to check for data dependency by comparing the destination register of the load instruction with the source registers of the instruction that immediately follows it. However, it is not directly possible to check for data dependency between the load instruction in address generation stage and the second and third upcoming instructions, which is required by ELD³ technique. Therefore, a Data Dependency Bit (DDB) memory is implemented in the address generation stage that holds the dependency status for each instruction in level-one instruction cache (L1 IC). When a load instruction is detected after instruction fetch, the data dependency bit is accessed from the DDB memeory for the corresponding instruction. Figure 5 illustrates the DDB memory for a two-way L1 IC. The dependency bit will be correct as long as the cache line is not evicted from the L1 IC. Should a cache line be evicted from L1 IC, the load operation will still be executed correctly at the expense of an additional stall cycle. Moreover, the dependency bit for the load instruction will be updated during the writeback, such that the dependency bit is correct next time the load instruction is executed.

SHA+ELD³

By combining SHA and ELD³ together, the ELD³ technique can be used when the displacement address is too large

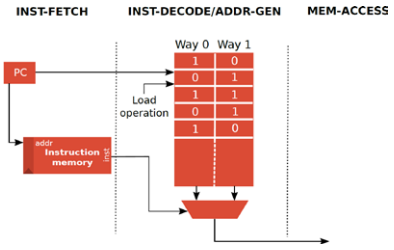

Figure 5:
DDB memory for a two-way L1 IC

for the SHA technique. A load operation can then be performed like this:

1: When the displacement is small, the halt tags are accessed but the DDB memory is not accessed. The tag and data ways are accessed in parallel, but SHA will halt both tag and data ways using the hit vector from halt tag access.

2: When the displacement is too large for SHA, the halt tags are not accessed, but the DDB memory is accessed. The outcome of DDB memory will decide the next step taken.

2a: If the DDB memory returns a dependency bit which is cleared, then all tag ways are accessed in parallel, but the data way is sequentially accessed.

2b: Or if the DDB memory returns a dependency bit which is set, the tag and data ways are accessed in parallel, such that the data can be forwarded to the following instruction and avoid a stall cycle.

RESULTS

The effectiveness of the SHA and ELD³ implementations were evaluated by running MiBench benchmark applications on a four-stage RISC-V RV32I 32-bit processor, implemented on Single-ISA Heterogeneous MAny-core Computer (SHMAC) framework. SHMAC is a research project initiated by the Energy Efficient Computing Systems (EECS) department at Norwegian University of Science and Technology (NTNU), that uses a tile-based architecture. The MiBench applications were compiled using the RISC-V GCC toolchain, and analyzed using the SHMACsim, a cycle-accurate simulator for the SHMAC framework.

Figure 6 shows the average number of ways accesses when using SHA and ELD³

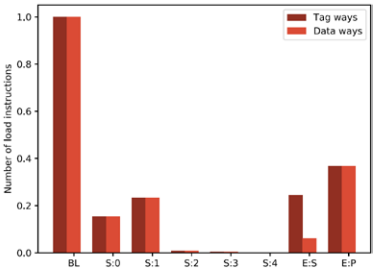

Figure 6:
Way distribution between SHA and ELD³

techniques in combination, relative to a conventional cache implementation. When the displacement is small and the SHA technique is used, only one tag and data way is accessed for most load instructions, shown as S:1. In addition, when there is a cache miss, zero tag and data ways are accessed. Load instructions that result in cache misses occur, as we can see from S:0, quite frequently. When the displacement is too large for SHA and there is no data dependency, we can see that the number of data ways accessed is reduced significantly by accessing the data ways sequentially using ELD³, shown as E:S.

CONCLUSION

Improving the energy-efficiency of computing is an important area of research, and there is a potential for reducing the energy dissipation of caches. This article shows that using the concept of practical way-halting and data dependency detection, it is possible to reduce the energy dissipation for L1 DC without reducing the performance.

IMPROVING MOTION CONTROL IN A BIPOLAR PRINTER

Sculpto – user friendly 3D printing for kids, school and leisure

No expert knowledge needed to use the Sculpto-printer. You can control it with your smartphone. Danish start-up company Sculpto makes 3D printing available for children, teachers and others. Developers from the Danish consulting company TechPeople, owned by Data Respons, have assisted Sculpto in developing a new, improved version of the printer.

BY: Brian Obel
Manager Aarhus,
TechPeople

Simon Breum Fisker and Jacob Lindeberg launched their Sculpto printer in 2016, and September 2017 they released a new, improved model called Sculpto+. The new printer is quicker and quieter. Also, a new "print engine" makes it even easier to find models on the internet, e.g. on thingiverse.com and print them directly from your phone.

But although extremely user friendly the machine is highly sophisticated as both the arm and the bed move in a polar coordinate system. For the Sculpto+, TechPeople applied some advanced printer engine control software, up until now only used in a few expensive high-end printers. Also TechPeople consultants improved the original PCB layout.

SOPHISTICATED DEVICE

The Sculpto+, though fairly cheap, user friendly and brightly coloured, it is a sophisticated little device. Developers from TechPeople have assisted Sculpto, among other things doing a PCB layout review. TechPeople's review expert made a number of recommendations for improvements. The new layout has eliminated engine noise. Now you can only hear a faint humming from the printer fan.

TechPeople also applied new printer engine control software. The control software makes acceleration and deceleration softer, resulting in higher speed and precision.

An improved motion control algorithm

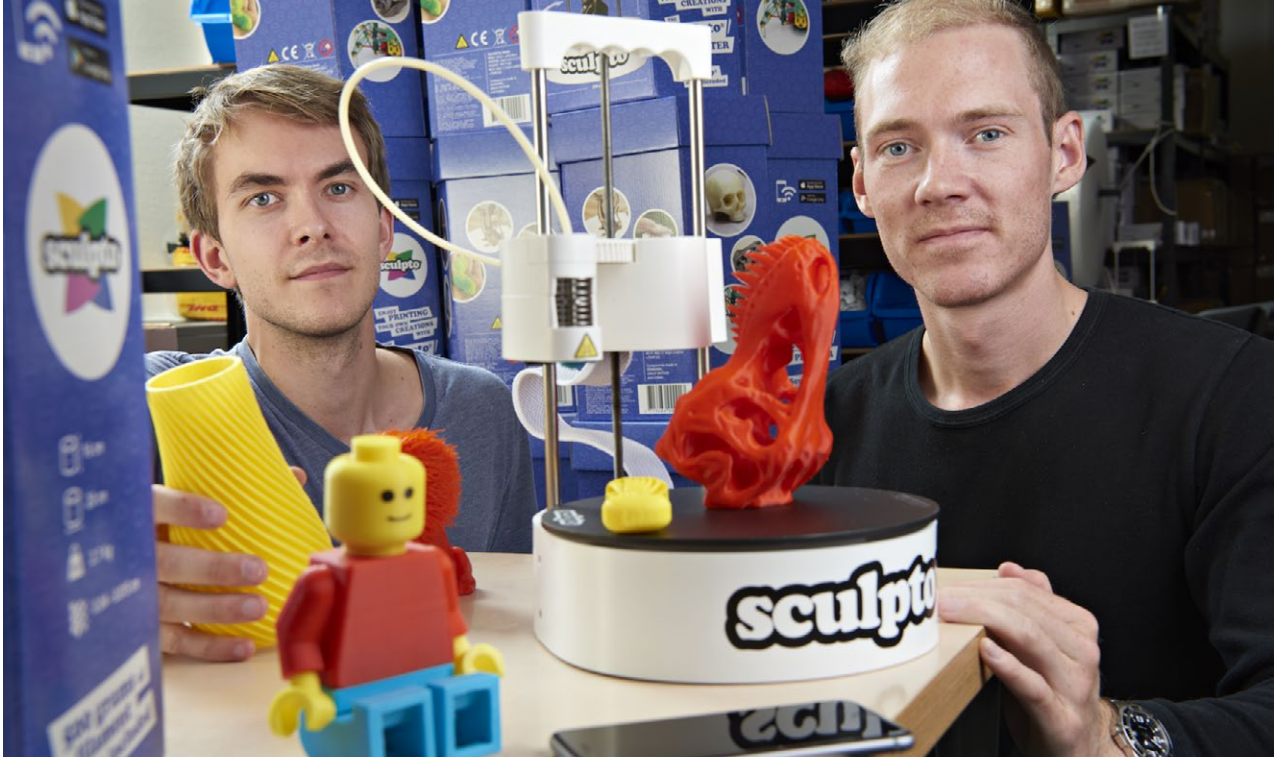

Simon Breum Fisker and Jacob Lindeberg, from Sculpto.

was developed and implemented for the printer. The algorithm ensures controlled acceleration of the printer's stepper motors and a smooth print pace. There were two major challenges regarding motion control for the 3D printer that needed to be considered.

PHYSICAL MOVEMENT

The first challenge was the physical movement of the print head and print plate. Due to the bipolar nature of the printer, the print plate is required to rotate with a speed that is dependent on print-head distance from center of the plate. Otherwise, the print-head will not keep a constant pace when moving across the print plate in a straight line. The constant pace is required to get a good print result.
The challenge is that this requires the print plate stepper motor to accelerate and decelerate quickly when crossing near the print plate center in order to turn the plate 180 degrees quickly, so that the print head can maintain a constant pace all the way across the plate.

Another issue related to physical movement of the print head is, that when it needs to change direction, this should be done as quickly as possible to avoid stopping too long at the same location and thus leaving too much print material in the edges of the print. Meanwhile, it is critical not to accelerate too fast because otherwise the stepper motors risk skipping steps with the result that the coordinate system of the 3D print becomes shifted in the middle and thus the print cannot be used.

MOTION PLANNING

The second challenge was motion planning based on a stream of print commands known as CNC g-codes. This means that the motion controller at any time only knows about a small part of the complete path for the 3D print.

Even with this limited information, the printer is required to keep a smooth pace by constantly updating its planning ahead route, in order for print material to be extruded in smooth layers. If the print-head at any time stops even for a short while, it will leave too much print material on the stopping location resulting in bumpy prints. The challenge here was that a lot of mathematics needed to be defined, so that print speed can be maximized while the motion is kept within several different constraints on speed and acceleration.

IMPROVED MOTION CONTROL

The solution for improved motion control consists of two major updates as well as a number of minor updates for stability, configuration, test and debug possibilities. The major updates are an updated motion planner algorithm that is based on constant jerk motion (jerk is the derivative of acceleration) and a stepper motor control that supports constant jerk motion and the microstepping step drivers used in the Sculpto printer.

The motion planner algorithm implements constant jerk motion governed by classical constant acceleration motion equations expanded to the constant jerk case. The equation that forms basis for all the optimization formulas derived and implemented in the motion planner is:

$$s(t) = s_0 + v_0 t + \frac{1}{2} a_0 t^2 + \frac{1}{6} j t^3$$

where,
s = distance
v = velocity
a = acceleration
j = jerk
t = time

Movement using constant jerk involves up to 7 steps in the motion as seen below

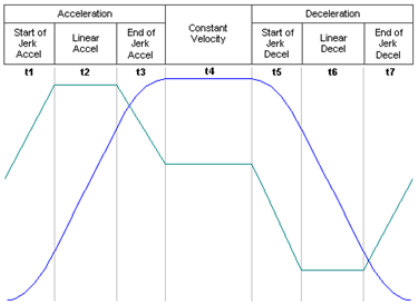

* (source at the bottom of the page)

Furthermore, the planner needs to handle cases where it is not possible to achieve maximum defined acceleration or velocity for each axis. These cases arise when the motion is limited either by total distance of movement or the dynamic velocity constraint. The dynamic velocity constraint is caused by 4-axis motion on the printer where each axis has an independent maximum velocity and acceleration since the overall move-

The control software makes acceleration and deceleration softer, resulting in higher speed and precision.

* Source: https://www.controleng.ca/servosoft/SSHelp1033/source/MotionProfile.htm

ment needs to follow point-to-point print commands.

UPDATED STEPPER MOTOR CONTROL

The updated stepper motor control supports constant jerk motion instead of the previously used constant acceleration motion. This pushed the microcontroller used in Sculpto closer to the limits, since the requirement for the solution was to use the existing microcontroller due to hardware cost considerations. Both memory and CPU resources was restricted.

A solution was implemented that found a compromise between memory usage by the pre-calculated stepper timings during motion planning and CPU load in the on-the-fly calculations during the 50 kHz stepper update interrupt routine. Furthermore, some of the calculations were converted from 32-bit fixed point to 64-bit fixed point in order to achieve the computation precision required by constant jerk motion equations.

SMOOTHER AND FASTER

All in all the software update enabled the Sculpto printer to use smoother and faster motions during printing, while keeping within the limits of their existing microcontroller. As a bonus Sculpto got access to and knowledge about more commands and settings in the internal motion controller. This enabled them to continue optimizing the printer settings and constraints as they see fit, even after the software update was completed.

About Sculpto

Affordability and user friendliness

- 3D printing has a huge potential. But the typical printer is big, heavy and difficult to operate, say Simon Breum Fisker and Jacob Lindeberg, who founded Sculpto in 2015.

- We decided to develop a printer that was cheap, compact and user friendly. In particular we wanted to simplify the steps you have to go through from designing a model to actually printing it. That was really complicated and required a lot of technical knowledge. We wanted to automate that as much as possible and open up 3D printing to a much bigger group of users.

When Simon Breum Fisker and Jacob Lindeberg set out to develop the Sculpto printer, they decided to focus on the two aspects that they saw as the major barriers to 3D printing becoming a commodity: price and user friendliness.

- We took the printer control board and transformed it to the screen the user carries in his pocket – his smartphone. We developed an app for controlling the printer wirelessly. In that way we can continuously add value to our customers through the app and through online updates of the printer.

- Also, we chose bipolar printing because it makes the printer look nicer and more accessible. A bipolar printer extrudes the plastic at the intersection point of two circles, as opposed to a cartesian coordinate system, where the plastic is extruded at the intersection point of two straight lines.

AUTONOMOUS CABLE SURVEY

with magnetometers

An extensive network of conductors is needed to transmit electricity from power stations to the consumers. When these cables cross a large portion of water, they are usually buried beneath the seabed. The cables are here exposed to the marine environment and regular surveys are required. Unmanned subsea vehicles are usually used to carry out said surveys.

Most unmanned vehicles are remotely controlled by radio and their position is pinpointed via GPS. This is problematic for unmanned subsea vehicles, as electromagnetic waves do not propagate nearly as far in water as they do in air. Today, transmit speeds can come up to a couple of kilobytes per second when the range is about 200 meters. Alas, it's still insufficient for the videofeed needed for remote control.

BY: Aksel Stadler Kjetså
Development Engineer
Data Respons

ROVS AND AUVs

The most widespread unmanned subsea vehicles are ROVs (Remote Operated Vehicles). Said vehicles has a wired connection to the operating ship. This umbilical cord transmits video as well as the steering signals. To perform a cable survey the ROV is manually steered along a cable and inspected for damage. This is both tedious and expensive. Further, the need for an operating ship is a large drawback as it causes a significant cost increase. This is especially true for time consuming operations like cable surveys.

The alternative to ROVs are AUVs (Autonomous Underwater Vehicles). An autonomous system is more independent that an automatic. In short terms it's able to fulfill a mission entirely without human intervention. Ideally, an AUV is self-reliant and does not need an operating ship. This makes it more cost effective. Unfortunately, accurate IMUs (Inertial Measurement Unit) are expensive and the exact coordinates of the object to survey isn't always known. This might cause the AUV to search in the wrong area or even get lost entirely.

To prevent the latter, ships are often used to follow the AUV. Today, this is the one of the greatest criticism of AUVs as they, in practice, also need an operating ship in order to function properly.

MAGNETIC-BASED CABLE TRACKING

By using the magnetic field produced by a power cable it's possible to track along it with relative ease and high accuracy. Potentially, this can reduce the cost of cable surveys. Additionally, it enables tracking of buried cables, as well as a way to determine how deep they are buried. The latter is important as cables that are not properly buried are more prone to take damage. By mounting a camera on the AUV an operator can inspect unburied sections post-survey.

There are mainly three types of underwater power cables. The discussion here will be limited to three-phase AC power cables, where there is a separate cable for each phase. This is common for high power transmission lines crossing relatively small portions of water (< 100km). For longer cables, HVDC (High-Voltage DC) is used, as the loss from inductance to water outweighs the expenses of ACDC-converters.

The procedure for autonomous cable tracking can be roughly divided into three. Signal processing to extract the magnetic field generated by the cable,

Localization of the cable from its magnetic field and Steering the AUV along the cable. All these parts are performed live on the AUV; i.e. without any human intervention.

SIGNAL PROCESSING

To use the cables' magnetic field for autonomous tracking the AUV must be within range and isolated from other fields. The latter includes anomaly fields from magnetized rocks, the earth's geomagnetic field and the field induced by the AUV.

The range is proportional to the current in the cable and dependant on the accuracy of the magnetometers. As an example, the new 420Kv cable in Oslofjorden will be detectable at about 30m from the FLC3-70 magnetometer that costs about 1000NOK.

All AC currents in Norway are at 50Hz, so the magnetic field will oscillate at this frequency. Both the anomaly field and the magnetic fields are static. Lastly, the AUV induced field can be tuned. Consequently the sole field at 50Hz is produced by the cable. This makes it extractable from the sampled signal by the discrete Fourier transform, which is a well-known tool in engineering mathematics.

The purpose of the Fourier transform is to decompose a signal into its frequency part. Basically, it does this by calculating the covariance between the sampled signal and imaginary signals ranging from 0 to the sampling frequency. If the covariance is high the current frequency is a part of the signal. By using a three-axis magnetometer and implementing the Fourier transform, the amplitude and phase shift of the cable generated field can be isolated in each direction.

LOCALIZATION

After extracting the cables magnetic field it's possible to determine the relative heading and cross track distance to the cable. These are the parameters needed to autonomously steer along the cable and are notated in Figure 1 as ψ_{br} and Y.

The magnetometers are three-axis, meaning that they decompose the magnetic signal in x-y-z components relative to the AUV. We know the magnetic field is perpendicular to the cable, so the relative heading ψ_{br} between the AUV and the cable can be determined by trigonometric functions. As most AUVs are equipped with a compass, it's heading ψ_{AUV} is also known. From ψ_{br} and ψ_{AUV}, the heading of the cable, ψ_{cable}, can be determined.

The distance to the cable is a bit more complicated and requires two triaxial magnetometers. By triangulating the magnetic field at the two magnetometers the exact distance to the cable can be determined. The equations are rather lengthy and are omitted here. For readers with special interest, a derivation is given in Xiang, X. (2016).

STEERING

Autopilot

To track along the cable the AUV must be sufficiently close and the relative heading must be zero. The now known cable heading ψ_{cable} and cross track distance Y can be used to implement a steering algorithm.

Most AUVs use a small rudder to turn in yaw. Intuitively, it can be understood that the turn rate $\dot{\psi}$ is a depended on the AUVs rudder angle δ. A simple heading autopilot can therefore be implemented as:

$$\delta = K_p * (\psi_{desired} - \psi_{AUV})$$

If the desired heading equals the cable heading we obtain:

$$\delta = K_p * \psi_{br}$$

This is called a Proportional controller or P-controller, where K_p is a constant achieved by tuning. As seen from the equation, the rudder angle will be zero when actual heading equals the desired. If there is an ocean current, this controller is insufficient, as it will saturate. Therefore it might be wise to add integral action:

$$\delta = K_p * \psi_{br} + K_i \int \psi_{br}$$

This is called a Proportional-Integral controller (PI) and will sum the error over time, thus gradually suppress the effect from a constant opposing force. To limit overshoots, it's common to add a derivative term as well:

$$\delta = K_p * \psi_{br} + K_i \int \psi_{br} + K_d \frac{d}{dt} \psi_{br}$$
$$\delta(s) = K_p * \psi_{br} + K_i \psi_{br}/s + s * K_d \psi_{br}$$

This makes it a PID-controller, which is an extremely popular low-level controller. It's used in everything from AUV autopilots to regulating temperatures in ovens and so on. It has the edge over model-based controllers because it does not require a transfer function of the system, which might be difficult or even impossible to derive. Its drawback is the required tuning of the K-gains and that the controller will never be optimal.

An example of a model based heading autopilot is Nomoto's first order. It was first presented by Nomoto in 1957, but is still the basis of many heading autopilots used on marine crafts today. The transfer function is

$$\frac{\psi}{\delta}(s) = \frac{K}{s(1+Ts)}$$

Which is an integral part in cascade with a low pass filter. Or simpler put; the heading is sum of the rudder angles where too rapid changes are filtered out. K and T are here not achieved by tuning, but by a mathematical model of the craft.

For AUVs the mathematical model tend to be very complicated as it, compared to ships, has two extra controllable degrees of freedom (heave and pitch). Therefore, they are usually implemented with PID low-level controllers. Figure 2 shows the control loop of the system. In the AUV block the heading is measured and the relative heading is calculated from the magnetometer readings. In the PID block the rudder angle is calculated from the relative heading.

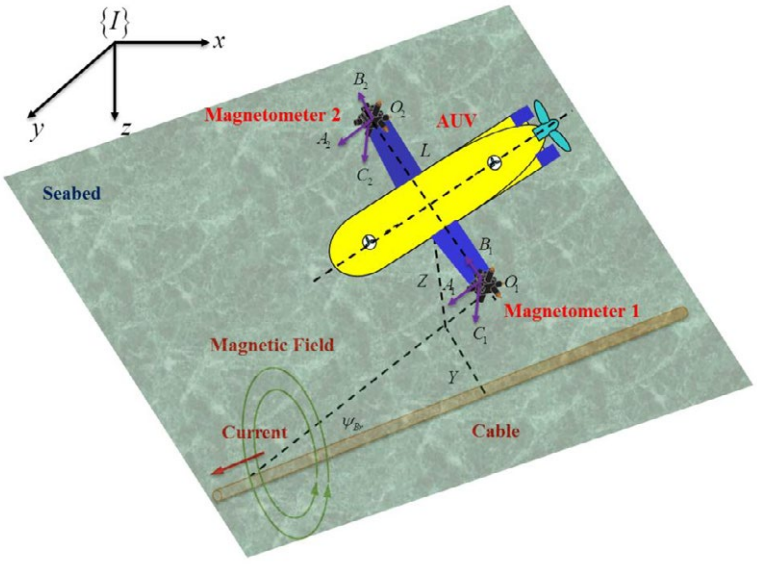

Figure 1: AUV and buried power cable

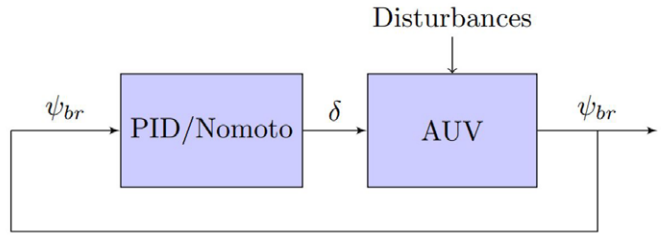

Figure 2:
Control loop for a heading autopilot

Guidance

Whether a tuned PID or a derived model-based controller is used, the AUV will be able to steer along the cable with the same heading. However, only using a heading autopilot does not ensure that the AUV moves closer to the cable. Instead of making a new low-level controller, a guidance system can be used.

A common guidance law is LOS (Line of Sight). LOS was originally developed during the Cold War for military operations. It was used to intercept moving targets by predicting the meeting point. Specifically, it was used in surface-to-air missiles by constantly recalculating and adjusting the heading to ensure impact on the target aircraft. Assuming the cable will continue in a somewhat straight line, LOS guidance can be used in a less explosive manner to diminish the horizontal distance while holding a smooth course.

An illustration of the LOS steering law is shown in Figure 4. The desired heading is the heading of the cable added with the heading to the meeting point χ_r. The distance to this point is decided by tuning Δ.

The desired heading when using LOS is

$$\psi_{desired} = \psi_{cable} + atan\left(\frac{Y}{\Delta}\right)$$

Note that the meeting point is recalculated at each iteration, making the AUV steer in a gradual curve until the cross track error is diminished and it's directly above the cable. At this point, a mounted camera can be used to record video along the trajectory for inspections post-survey.

Figure 4 shows an illustration of the control loop when including the LOS-guidance in cascade with the autopilot.

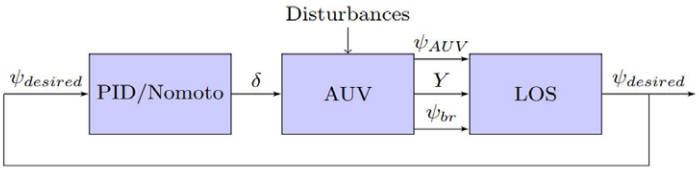

Figure 4: Control loop with autopilot and LOS guidance

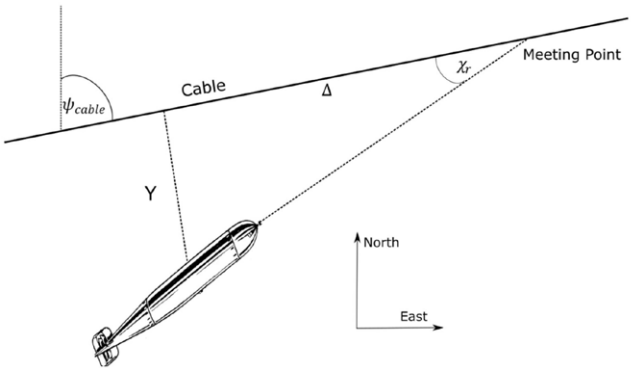

Figure 3:
LOS guidance

CONCLUSION

The subsea environment introduces many challenges for surveying installations. Today, ROVs are the most used solution. For a relatively small survey area, it may be the best solution. Yet, for cable surveys, which span a large distance, they have severe limitations both in cost and time. A proposal is to use AUVs in conjunction with magnetometers to autonomously track along power cables. This might dramatically reduce the cost of such surveys.

A great concern for AUVs are the complicated algorithms needed to find the survey area and to prevent getting lost.

By using the magnetic field of power cables, it's possible to determine points to steer after with relative ease. When this critical information is obtained, many existing control algorithms can be implemented.

Refrences

Fossen, T. (2011). Handbook of Marine Craft Hydrodynamics and Motion Control.

Kjetså, A. S. (2017). Localisation of Submarine Power Cables by Magnetometers on Remus 100AUV. Trondheim. From https://github.com/Akselk/Cable_Localization_by_magnetomers

Xiang, X. (2016). Subsea Cable Tracking by Autonomous Underwater Vehicle with Magnetic Sensing Guidance. Sensors.

A system development project involves different disciplines, and will always have a level of uncertainty. This uncertainty implies a degree of change, complexity, and risk. The chosen project approach will affect the success of the project. Why and how should a project approach be selected? This article addresses agility and various approaches for system development projects.

HOW OPTIMAL is your approach?

BY: Stig-Helge Larsen
Principal Development Engineer
Data Respons

SELECTION OF PROJECT APPROACH

The *project approach* is the way in which project deliverables will be realized [1]. A project approach consists of a set of applied methods, techniques or tools applied to satisfy expectations and needs. Project characteristics such as the level of uncertainty, the available resources and the project success criteria should influence development of a successful approach in each system development project.

Initially, before a project approach is selected, the relevant stakeholders are identified and their expectations are discussed. Their need for information and involvement during the project are also discussed and analyzed. Measurable project objectives and related success criteria must be identified, discussed and prioritized. All this is documented together with other initial high-level information, such as project purpose, prioritized high-level requirements, and project exit criteria. A question that should be answered by the customer (or sponsor) is in regard to project constraint priorities: Is the quality, time, functionality or cost the first priority? All the initial communication described above is useful input to find the approach that best fit with all priorities.

The prioritized success criteria are used as the guiding principle for the approach to be developed. Factors that may impact success must be identified. Further, through the approach development, the necessary success factors are selected to satisfy the success criteria. The chosen approach should also be based on a life cycle model with characteristics that match the project characteristics. An example of project characteristics is considerable uncertainty that implies a high rate of change, complexity, and risk of rework. For this example, an appropriate approach can be based on a life cycle model that allows the project to tackle a high amount of uncertainty, via small increments of work.

Prior to any detailed planning, an initial approach for a project may, among other things, consist of:

- A project life cycle model (e.g. an iterative model)
- Rules with respect to decision-making
- The way of gathering information and reporting
- The different meeting structures
- The responsibilities and authorities

In order to increase project success, the project approach must be customized because it depends on the success criteria, the available resources and the complexity. The chance of success can be increased by involving the people set to carry out the project work in adapting the approach. There are many different causes of complexity, for example a short duration, a large number of people involved, an insufficient budget, uncertainty in estimates, or external dependencies such as other projects. The project approach can also be changed during a project. It always depends on the situation.

CHARACTERISTICS OF PROJECT LIFE CYCLES

A *project life cycle* is all the phases/periods from idea to a final delivery. Characteristics of the following life cycle models will be addressed:

- Sequential (e.g. waterfall)
- Agile
- Iterative
- Incremental
- Hybrid

Traditional *waterfall* and *V-model* life

cycles, with their sequential processes, have a lot of planning upfront. They also have a lot of analysis and design before the build phase. The waterfall life cycle model, shown in figure 1, is most appropriate when the requirement are well known and fixed, in projects with stable teams and low risk. There is a review at the end of each phase to verify the work and to decide if the next phase can start. Testing starts only after the finished build phase because the phases do not overlap. The waterfall model can be problematic if used for long and complex projects, and for projects with many changing requirements, due to the sequential processes without several iterations.

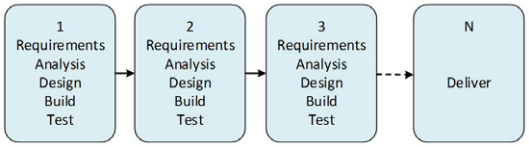

Figure 2.
An iterative life cycle allows feedback to improve and modify unfinished product or result through several iterations. It is optimized for learning.

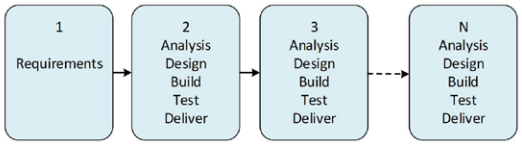

Figure 3.
An incremental life cycle delivers several small and usable subset of the whole solution. It is optimized for speed.

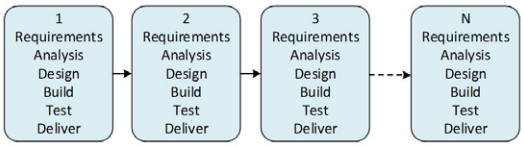

Figure 4.
An agile life cycle is both iterative and incremental. It has early and continuous delivery of a valuable product, and allows feedback and adaption to high degrees of change.

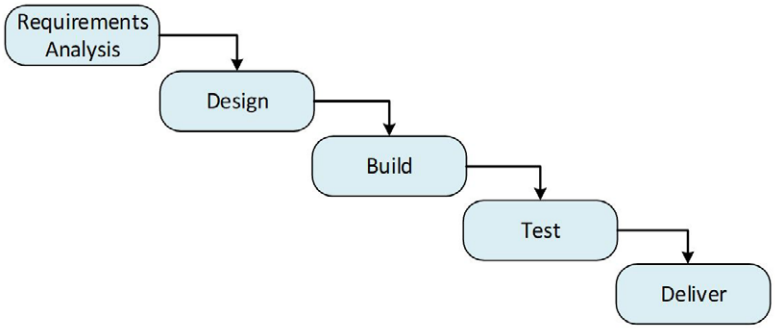

Figure 1: A waterfall life cycle takes advantage of things that are known and proven. Detailed requirements and plans are created at the beginning of the project. Each phase is completed fully before starting the next.

An *agile* life cycle can be an alternative to the waterfall life cycle. Project approaches based on an agile life cycle model are commonly used. Agile life cycles are both iterative and incremental. This means both repeated activities and frequent small deliveries, as shown in figure 4. The goal for agile approaches is to deliver a continuous flow of value to customers and achieve better business outcomes. Feedback on each delivery is used when planning the next iteration. Agile approaches will follow the principles of the *Agile Manifesto* [4].

Figure 2 shows an *iterative* life cycle. An iterative life cycle can be appropriate when the complexity is high and when frequent changes are expected.

Figure 3 shows an *incremental* life cycle. An incremental life cycle can be appropriate when the customer want frequent smaller deliveries with a subset of the complete solution because of business needs that cannot wait. Further, frequent reviews improves the quality. If an iteration-based agile approach is used, the team collaborates to finish the most important features in each iteration (each time-box).

Agile life cycles have several advantages for system development projects, but there are also some potential challenges to be aware of. Quantification of effort, time and cost is difficult at the beginning of an agile project life cycle because the team does not have all the upfront estimation and planning as in waterfall. However, the team can provide better estimates after a few iterations (sprints), when the team has established a reliable velocity (average amount of work completed in each iteration). Another challenge is risk for insufficient emphasis on necessary design and documentation. Further, a risk is also having only inexperienced engineers in an agile team – they should be combined with engineers or a project manager that has the experience needed to make the required decisions during the development process.

An example of a specific life cycle is use of a model that group increments and/or iterations into several large phases, where each of the phases are divided into several smaller time-boxes. This enables high-level planning of one larger phase at a time, and more detailed planning for each time-box.

A commonly used *hybrid* life cycle is a combination of waterfall and agile. For example by using some agile methods such as short iterations (e.g. 2 weeks), backlog, frequent demonstrations, and retrospectives, but still follow other aspects such as considerable upfront estimation, analysis, and progress tracking according to waterfall approaches.

Use of both *Scrum* (including a board to visualize the flow of work) and elements of the *eXtreme Programming* (XP) method is a common blend of standard agile methodologies [2]. The Scrum framework provides guidance and description of concepts like product owner, scrum master, product backlog, sprint planning, daily scrum, sprint demonstration/review and sprint retrospective. Further, XP inspires engineering practices like continuous integration, refactoring, automated testing and test-driven development.

A *pragmatic* approach can be used together with waterfall, agile and hybrid approaches. A pragmatic approach will only use the practices that make sense for the individual team. The team will remove any unnecessary ritual, and focus on getting the quality and work done as quickly as possible. Agile is not what you do – agility is how you do it.

Attention to quality is a premise to release anything rapidly if an agile approach is used. Regression testing and testing at all levels are important – from unit testing to system and acceptance testing. This applies to both agile and sequential approaches. Several types of tests may be needed, for example stress, compatibility and usability testing, as well as load and performance testing. In addition, simulations are often useful for interim test of hardware and mechanical designs.

REQUIREMENT CHANGES

Figure 5 shows four life cycle categories and their characteristics related to degree of change and frequency of delivery. No life cycle is perfect for all projects. Instead, each project should find an optimal balance between the characteristics [2].

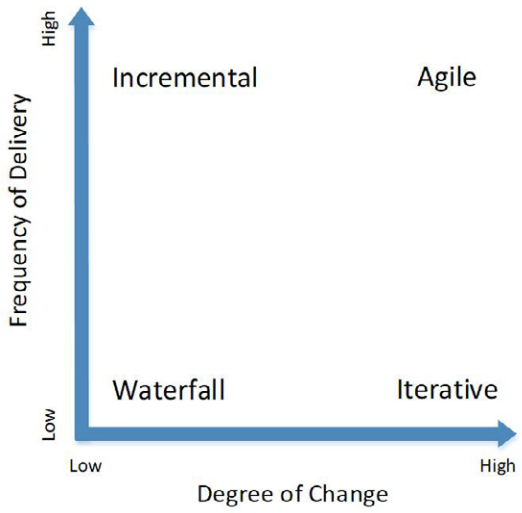

Figure 5
Four life cycle categories related to degree of change and frequency of delivery.

> An effective project manager can also help meet objectives and expectations, help respond to risk in a timely manner, help resolve project issues, help optimize use of resources, help manage changes and constraints, and help deliver the right products at the right time and cost.

Approaches based on a waterfall life cycle takes advantage of things that are known and proven. Detailed requirements and plans are created at the beginning of the project. This means that the sequential waterfall life cycle can be suitable for a small project with fixed requirements and low risk.

Agile approaches have early and continuous delivery of valuable products or results, and the projects can adapt to high rates of change. This can increase the customer satisfaction. Just-in-time requirement analysis means that a project starts with high-level requirements, and that the requirement specification is developed into more details during the project. Agile project teams should look for early and frequent deliveries to obtain feedback. When teams deliver small increments, they will better understand the true requirements. Software development is normally about learning while delivering value. Hardware development and mechanical development are similar in the design parts of the project. Therefore, an agile mindset can also be relevant for parts of hardware and mechanical development processes.

Magne Jørgensen at Simula Research Laboratory has recently performed a survey for 122 recently completed Norwegian IT projects: *"Requirement changes in IT projects: Threat or opportunity?"* [3].

The results indicated that it is useful to postpone adding details to the requirement specification if the project is large, has an agile approach, and/or has a time-and-materials contract. For non-agile approaches in IT projects of small or medium size and with a fixed price, a detailed requirement specification can be preferred and perhaps even necessary. The results also indicated that requirement changes during the projects due to learning contributed positively.

Requirement changes due to external changes, and imperfect early analysis, were negative for the successfulness of the projects. Half of the projects with a well-functioning agile process, time-and-materials contract, limited detailed requirements at project start-up, and frequent requirement changes during the project life cycle, were successful, and no projects in this group had a worse outcome than acceptable.

TEAMWORK SUCCESS FACTORS

Some examples of common teamwork success factors are clear objectives, joint responsibility (supporting one another), open and honest communication, mutual respect and trust between everyone, and flexibility (adapting to context and changes). Prioritizing, as well as rapid and transparent feedback, are common success factors for high-uncertainty projects that can imply high rates of change, complexity and risk.

A project needs several skills, and a team that has all the skills necessary to complete the work is a cross-functional team. The team members themselves should determine who will perform the work prioritized for the upcoming period. Empowered teams are more accountable and productive. Further, by limiting the work in progress, the cross-functional team members can collaborate more to deliver completed work. If team members are not 100 % allocated to a project, they can experience productivity loss because of task switching. Conversely, when every team member is 100 % allocated to a project, they can continuously collaborate and make the team more effective. The size of an agile team is also of importance. The PMI Agile Practice Guide [2] and the Scrum Guide [5] recommend a development team size between three and nine members.

Based on the project approach needs, a project manager may be desired. A project manager can add significant value in many situations, for example to facilitate a chartering process and collaboration, coach, give direction, help and advice. An effective project manager can also help meet objectives and expectations, help respond to risk in a timely manner, help resolve project issues, help optimize use of resources, help manage changes and constraints, and help deliver the right products at the right time and cost.

MEASUREMENT OF PERFORMANCE AND PROGRESS

Project measurement data is essential for improved forecasting, reporting and decision making. Two commonly used and recommended methods for empirical and value-based measurement of project performance and progress are *Earned Value* and *Burndown Charts*, as shown in figure 6 and figure 7. These two methods measure finished work. The measurements are based on what the team delivers, not what the team

predicts it will deliver. The Earned Value is the value of the work actually completed, accumulated at fixed time intervals – measured in either currency, work hours or story points. While the Burndown Chart shows work left to do (work hours or story points) versus time.

SUMMARY

An approach should be developed in order to optimize the project processes to achieve a successful project that satisfies expectations and needs. This means that the approach should provide the greatest chance of success. The development of a successful approach should be influenced by project characteristics, such as the level of uncertainty, the available resources and prioritized project success criteria. The uncertainty implies a degree of change, complexity, and risk. To satisfy the success criteria, necessary success factors are selected through the approach development.

Further, the developed approach should be based on a life cycle model with characteristics that match the project characteristics. Several aspects of working together, such as communication, responsibilities and decision-making, are established through the chosen project approach. The approach is always dependent on the situation, and should be open to changes during a project. It is not something you design on your own. The development of an approach is done in cooperation with important and influential key players, and the approach must be customized for each system development project.

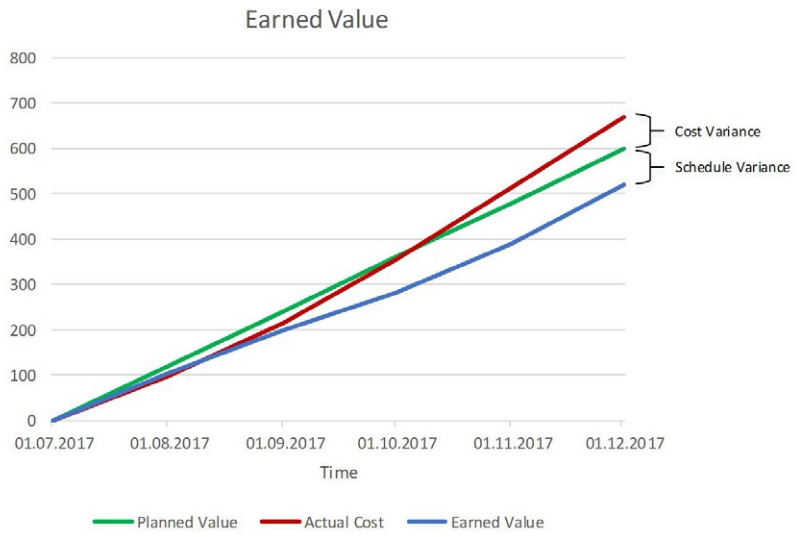

Figure 6: Earned Value is the value of the work actually completed. It measures project performance and progress in an objective manner.

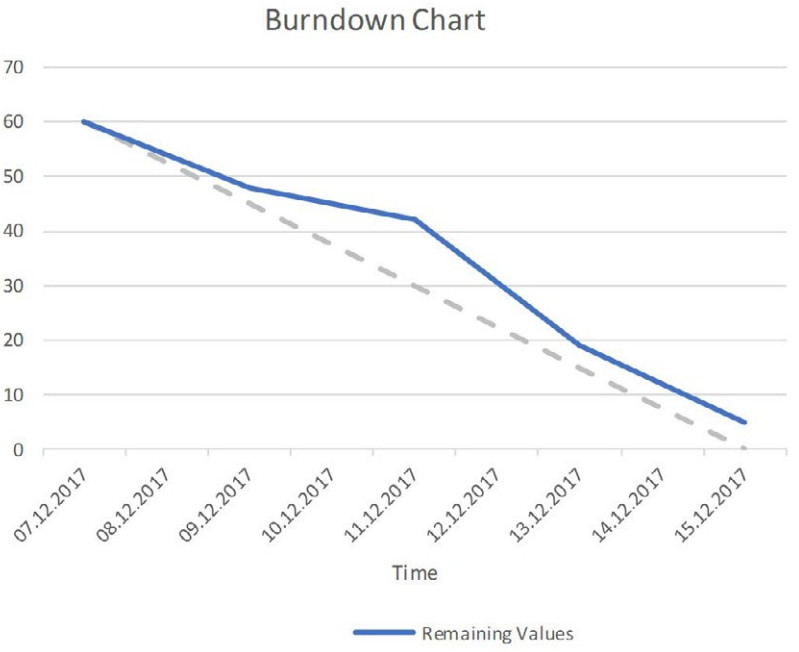

Figure 7: An iteration burndown chart shows work left to do in the iteration backlog.

Refrences

1. Better Practices of Project Management

Based on IPMA Competences 4th revised edition 2017

By: John Hermarij

2. Agile Practice Guide. 2017

By: PMI

3. Requirement changes in IT projects
Computerworld Norge, week 47, 2017

By: Magne Jørgensen at Simula Research Laboratory

4. Manifesto for Agile Software Development, 2001
http://agilemanifesto.org/

5. The Scrum Guide, 2017
https://www.scrum.org/

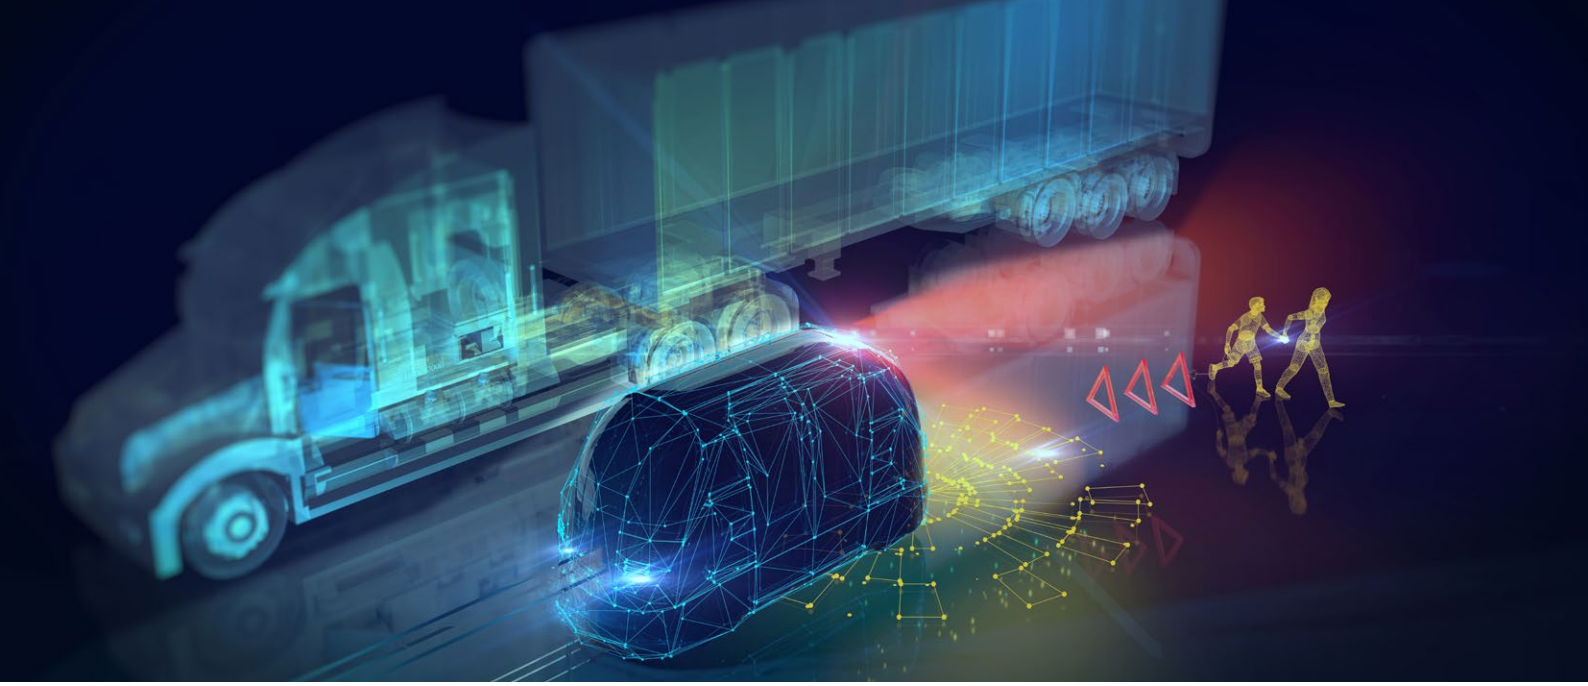

NEXT ISSUE:

Automotive: An industry in change

Our world is in a state of constant technological development. During the post-war period there have been major military technological achievements, much of which have spread to civil society. A rarely witnessed technological leap took place within the computer industry in the 1990s, processor speed doubled every 18 months and storage units rapidly became more effective. It became possible to develop the mobile telecommunications systems, which have subsequently been further developed in new generations' technology.

By: Crister Nilsson, Vehicle Industry Manager in Sylog AB

In the late 1990s electronics became more advanced in the vehicle industry and the number of computers started to increase substantially in vehicles to control individual functions. To develop more advanced functions, the computers in the vehicles needed to be connected together in networks for communication between different parts: engine, brakes and gearbox for functions such as automatic cruise control. The vehicle industry had long been regarded as an individual industry, with the power transmission chain's efficiency and environmental impact a major focus, but skills were now needed to be able to drive the development forward within electronics and software.

The challenges that the vehicle industry is now facing are no longer limited to the industry. Today connected service, energy and environmental impact, autonomy and AI are fundamental issues in all industries.

ELECTRIFICATION

Fully electrified power transmission chains are in their infancy, the intermediate stage of hybrid cars (electricity/combustion) still has a bigger share of the market. Tesla's presentation of the first Model S has highlighted the fact that it is possible to mass produce electrified vehicles for a wider market. Within the heavy vehicle industry companies are investing large sums in electrification of the power transmission chain. There are many reasons for this, but it is primarily minimisation of fossil fuels that is driving the development.

AUTONOMOUS VEHICLES

The development of self-driving vehicles is also in its infancy. The technology exists, and is used extensively in controlled environments such as mines.

The technology for self-driving vehicles will drive development of more advanced electronics in the form of new and faster sensors. Calculation speed in self-driving vehicles will need to be increased together with increased bandwidth in the communication between subsystems. Some parts of functions will also require the vehicle to be connected, partly to other vehicles and partly to back-end systems. The telecommunications industry will thus have become a part of the vehicle industry.

If the 1990s was the telecommunications industry's decade, there is a major chance that the 2020s will be the vehicle industry's decade. Views on ownership of vehicles will change and laws will need to be rewritten, nationally and internationally. The step of taking the technology out into open traffic will require a concentrated effort in terms of development in a number of different technological segments.

INSIDE WRITERS

STIG-HELGE LARSEN

Principal Development Engineer
Data Respons

Cand. Mag. Computer Science,
University College of Southeast Norway

AKSEL STADLER KJETSÅ

Development Engineer
Data Respons

Cybernetics NTNU, Master of Science,
Norwegian University of Science and Technology

SALAHUDDIN ASJAD

Development Engineer
Data Respons

(MSc) Computer Science,
Norwegian University of Science and Technology

OLE LARSEN

Software Development Engineer,
TechPeople A/S

Cybernetics NTNU, Master of Science,
Norwegian University of Science and Technology

CRISTER NILSSON

Consultant Manager & Automotive
Business Area responsible,
Sylog AB

B.Sc Computer Engineering

BRIAN HØRUP OBEL

Manager Aarhus,
TechPeople A/S

Master of Business Management & Programmer
Probana Business School, AABC - Aarhus

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,
Corp. Communications & Brand Manager,
Data Respons

Tel: +47 92 20 30 03
Email: ean@datarespons.no

TECHNICAL EDITORS:

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

TechPeople A/S,
Kim Farhenholtz, CEO
Email: kfi@techpeople.dk

Sylog AB,
Johan Jacobsson, CEO
Email: johan.jacobsson@sylog.se

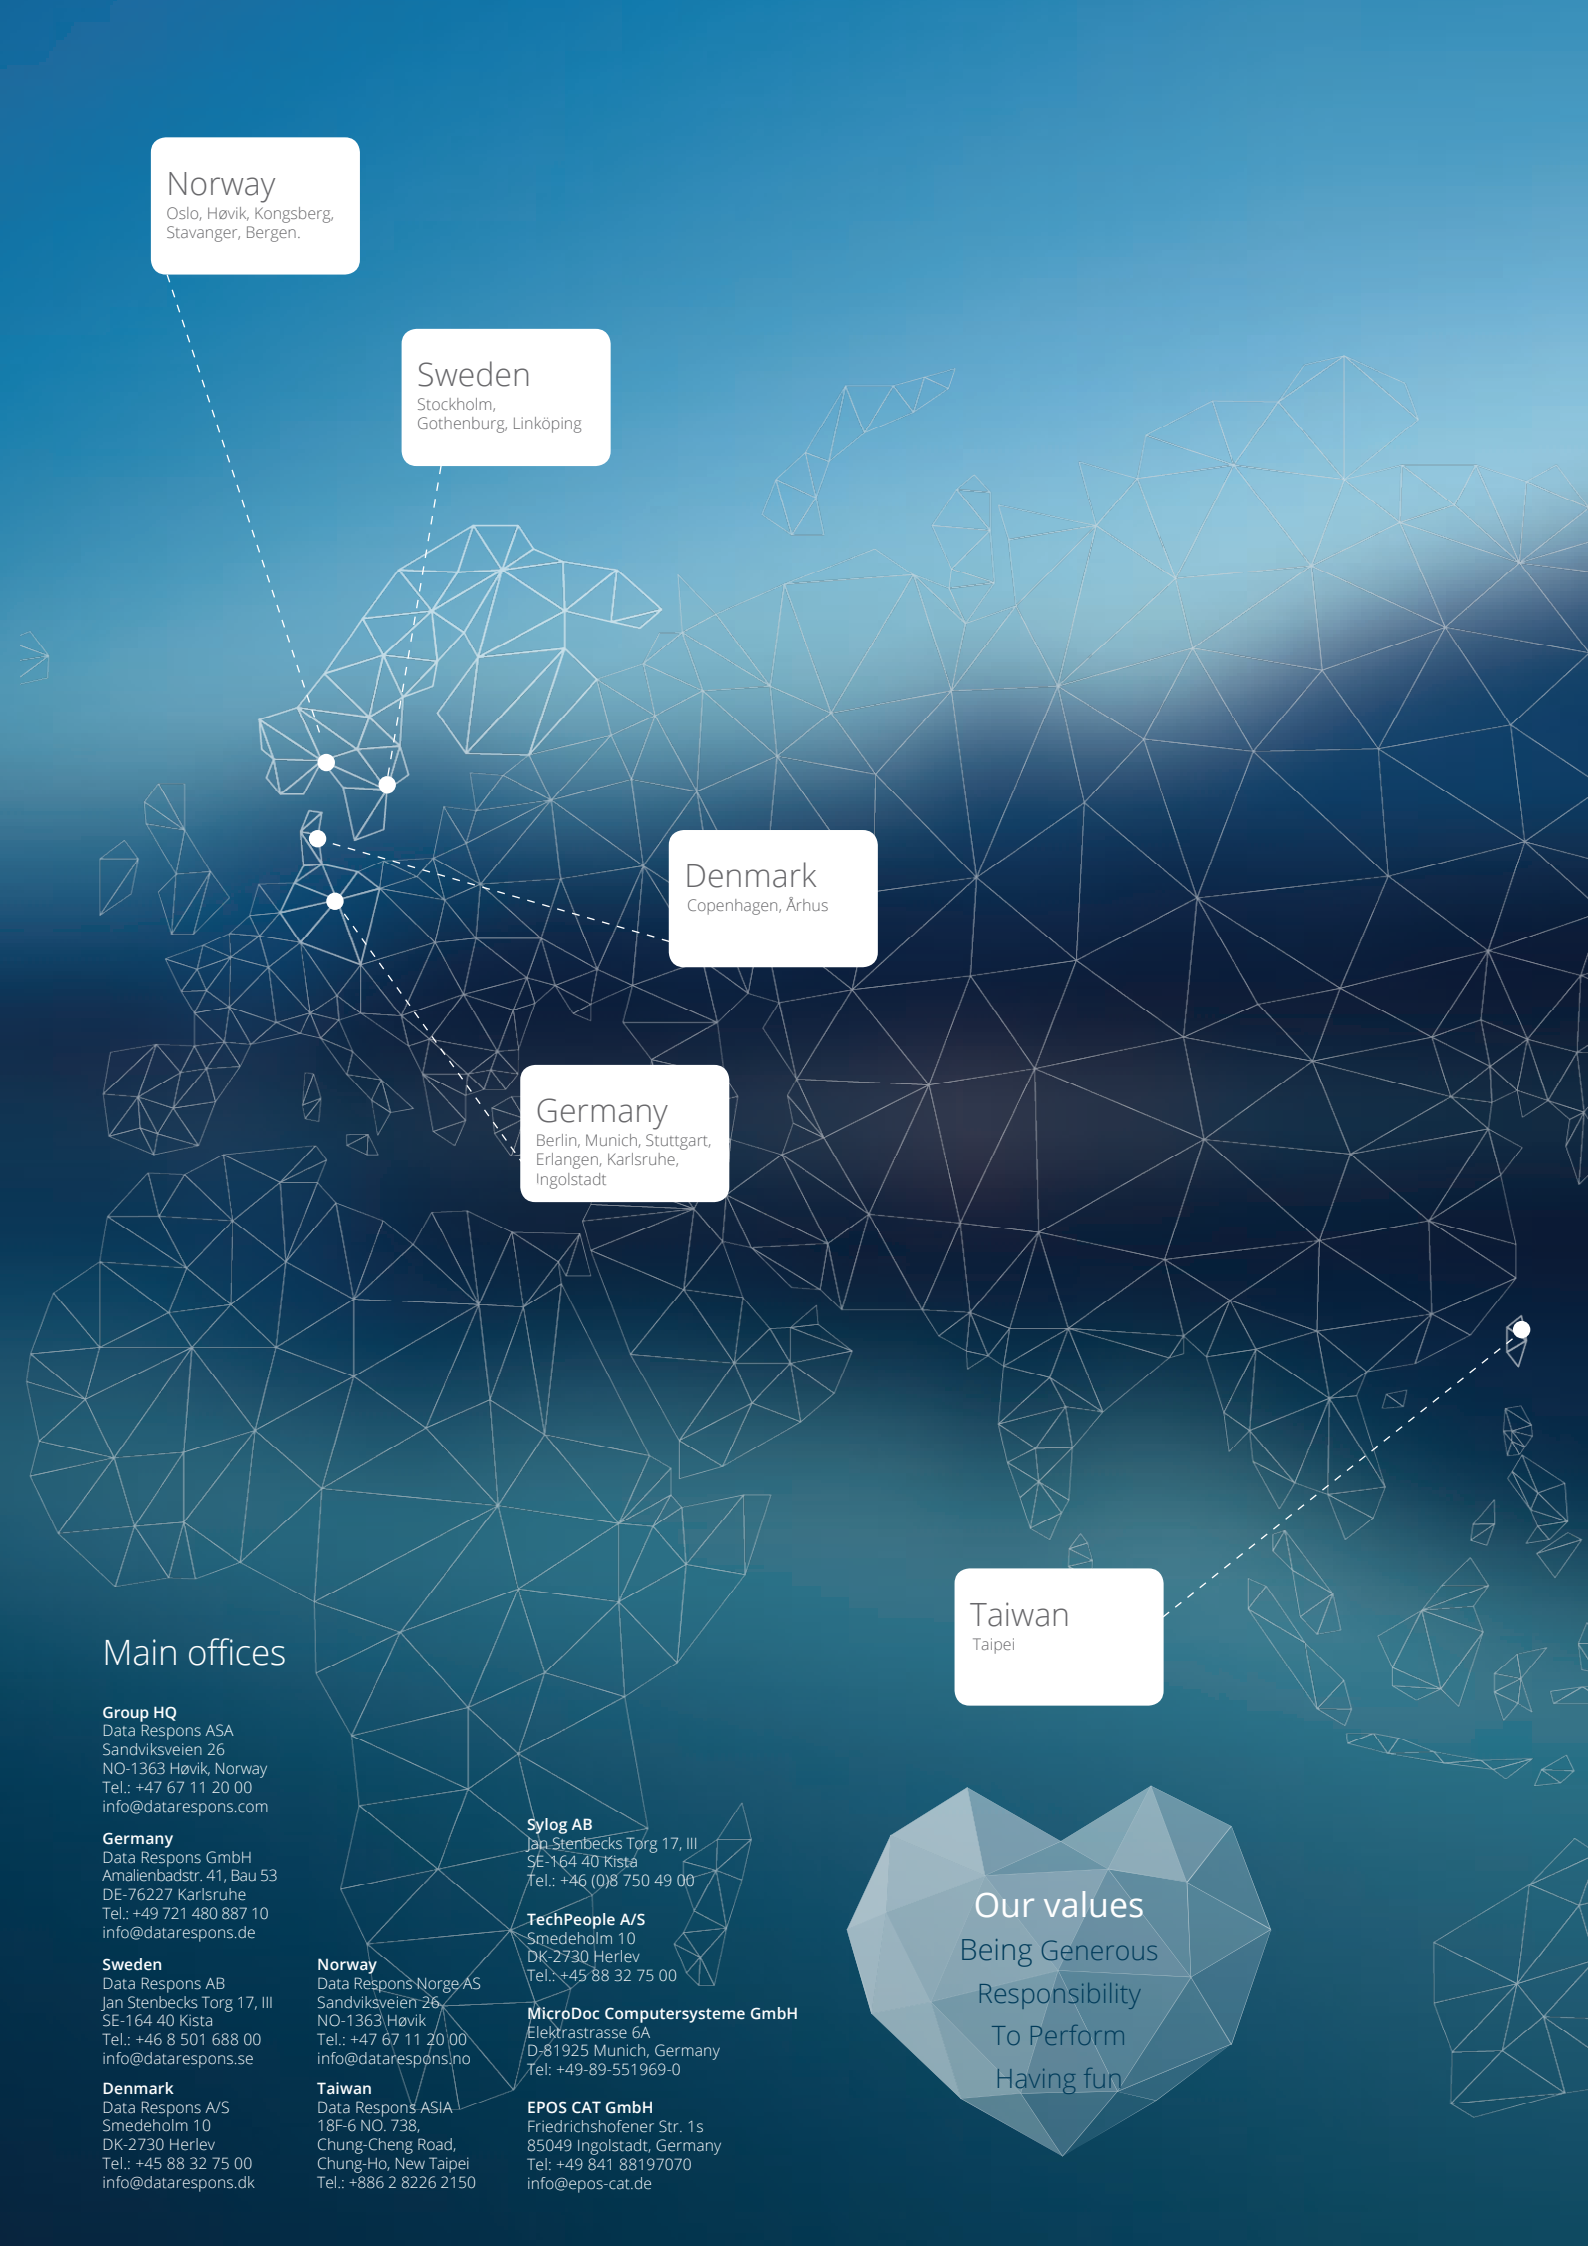

Main offices

Group HQ
Data Respons ASA
Sandviksveien 26
NO-1363 Høvik, Norway
Tel.: +47 67 11 20 00
info@datarespons.com

Germany
Data Respons GmbH
Amalienbadstr. 41, Bau 53
DE-76227 Karlsruhe
Tel.: +49 721 480 887 10
info@datarespons.de

Sweden
Data Respons AB
Jan Stenbecks Torg 17, III
SE-164 40 Kista
Tel.: +46 8 501 688 00
info@datarespons.se

Denmark
Data Respons A/S
Smedeholm 10
DK-2730 Herlev
Tel.: +45 88 32 75 00
info@datarespons.dk

Norway
Data Respons Norge AS
Sandviksveien 26,
NO-1363 Høvik
Tel.: +47 67 11 20 00
info@datarespons.no

Taiwan
Data Respons ASIA
18F-6 NO. 738,
Chung-Cheng Road,
Chung-Ho, New Taipei
Tel.: +886 2 8226 2150

Sylog AB
Jan Stenbecks Torg 17, III
SE-164 40 Kista
Tel.: +46 (0)8 750 49 00

TechPeople A/S
Smedeholm 10
DK-2730 Herlev
Tel.: +45 88 32 75 00

MicroDoc Computersysteme GmbH
Elektrastrasse 6A
D-81925 Munich, Germany
Tel: +49-89-551969-0

EPOS CAT GmbH
Friedrichshofener Str. 1s
85049 Ingolstadt, Germany
Tel: +49 841 88197070
info@epos-cat.de

Our values

Being Generous
Responsibility
To Perform
Having fun

A COMPLETE
TECHNOLOGY PARTNER
for smarter embedded and IoT solutions

Data Respons is a full-service, independent technology company and a leading player in the industrial IoT and the embedded solutions market. We provide R&D services and embedded solutions to OEM companies, system integrators and vertical product suppliers in a range of market segments such as transport & automotive, industrial automation, telecom & media, space, defence & security, medtech, energy & maritime, and finance & public sector.

data respons

DIGITALISATION OF THE INDUSTRIES OF TOMORROW