Zeitschrift:	Geomatik Schweiz : Geoinformation und Landmanagement = Géomatique Suisse : géoinformation et gestion du territoire = Geomatica Svizzera : geoinformazione e gestione del territorio
Herausgeber:	geosuisse : Schweizerischer Verband für Geomatik und Landmanagement
Band:	118 (2020)
Heft:	9
Artikel:	On the performance of UAV laser scanning with lightweight sensors
Autor:	Vallet, Julien / Gressin, Adrien / Clausen, Philipp
DOI:	https://doi.org/10.5169/seals-905956

#### Nutzungsbedingungen

Die ETH-Bibliothek ist die Anbieterin der digitalisierten Zeitschriften. Sie besitzt keine Urheberrechte an den Zeitschriften und ist nicht verantwortlich für deren Inhalte. Die Rechte liegen in der Regel bei den Herausgebern beziehungsweise den externen Rechteinhabern. <u>Siehe Rechtliche Hinweise.</u>

### **Conditions d'utilisation**

L'ETH Library est le fournisseur des revues numérisées. Elle ne détient aucun droit d'auteur sur les revues et n'est pas responsable de leur contenu. En règle générale, les droits sont détenus par les éditeurs ou les détenteurs de droits externes. <u>Voir Informations légales.</u>

#### Terms of use

The ETH Library is the provider of the digitised journals. It does not own any copyrights to the journals and is not responsible for their content. The rights usually lie with the publishers or the external rights holders. <u>See Legal notice.</u>

**Download PDF:** 06.05.2025

ETH-Bibliothek Zürich, E-Periodica, https://www.e-periodica.ch

# On the performance of UAV laser scanning with lightweight sensors

The employment of LiDAR sensors has relatively recently expanded into UAV mapping. Under the influence of manufacture advertisement, the users may obtain an incorrect impression that the path for obtaining a high-quality mapping product with this technology is straightforward. In reality, mapping with LiDAR is more delicate than with a camera due to a lower level of redundancy within the process of georeferencing and a somewhat higher threshold on the size/weight per performance ratio with these sensors. This fact motivated us to present a practical benchmark evaluating a popular small LiDAR sensor in realistic conditions for intrinsic parameters such as noise or capacity to penetrate the canopy, as well as the «low-weight» inertial technology in terms of geometrical influences on the resulting point cloud. The presented analysis reveals such practical limitations, which are worse than those specified by manufactures or tested in laboratory conditions.

Der Einsatz von LiDAR-Sensoren hat sich in jüngster Zeit auf die Kartierung mit UAVs ausgeweitet. Durch den Einfluss der Herstellerwerbung können die Benutzer den falschen Eindruck gewinnen, dass der Weg zu einem qualitativ hochwertigen Kartierungsprodukt mit dieser Technologie unkompliziert ist. In Wirklichkeit ist aber die Kartierung mit LiDAR heikler als mit einer Kamera, da die Redundanz innerhalb des Georeferenzierungsprozesses geringer ist und es einen etwas höheren Schwellenwert für das Verhältnis von Grösse zu Gewicht pro Leistung mit diesen Sensoren gibt. Diese Tatsache hat uns motiviert, einen praktischen Benchmark-Test zur Bewertung eines beliebten kleinen LiDAR-Sensors unter realistischen Bedingungen zu erstellen. Die Analyse beinhaltet intrinsische Parameter wie Rauschen oder Durchdringungsfähigkeit der Baumkronen sowie den Gebrauch von «leichten» inertialen Navigationssystemen im Hinblick auf die geometrischen Einflüsse der resultierenden Punktwolke. Diese vorgestellte Analyse zeigt solche praktischen Einschränkungen, die in der Tat grösser sind, als jene von Herstellern angegeben oder unter Laborbedingungen getestet wurden.

L'utilisation des capteurs LiDAR s'est relativement récemment étendue à la cartographie par Drone. Influencé par la publicité des fabricants, les utilisateurs peuvent avoir une fausse impression et ainsi penser que l'obtention d'un produit cartographique de haute qualité au moyen de cette technologie est simple. En réalité, la cartographie par LiDAR est plus délicate qu'avec une caméra en raison d'un niveau de redondance plus faible dans le processus de géoréférencement et un seuil un peu plus élevé pour le ratio taille-poids / performance de ces capteurs. Ce constat nous a motivés à présenter une étude pratique évaluant, dans des conditions réalistes un petit capteur LiDAR populaire sur des paramètres intrinsèques tels que le bruit ou la pénétration de la canopée, ainsi que l'influence de la technologie inertielle «légère» sur la géométrie du nuage de points résultant. L'analyse présentée ici révèle des limitations pratiques pires que celles spécifiées par les fabricants ou obtenues lors de tests dans des conditions de laboratoire. J. Vallet, A. Gressin, P. Clausen, J. Skaloud

## 1. Introduction

Sensor miniaturization supported the rise of the UAV industry [1], with the availability of smaller and lighter devices. Indeed, the process of data acquisition and production appear really easy with drones thanks to intuitive flight planning, reliable drone guidance, and subsequent automation of data processing. This process, however, remains guite challenging when using LiDAR-based products, where the quality of distance observations as well as georeferencing produced by lightweight sensors may vary rapidly in time and space. This fact motivates us to present a practical benchmark that sketches a realistic portrait of small airborne/mobile LiDAR and georeferencing sensors in terms of orientation and mapping performance. After introducing its design, we present a detailed comparison of sensors in terms of realistic accuracy as well as practical aspects related to their deployment and «mapping productivity».

## 2. Experimental setup

## Sensors

The aim of the experimental setup is to compare different sensors with respect to a reference during the same flight and thus acquisition conditions. Indeed, comparing sensors in different flights modifies the un-mastered parameters such as GNSS constellation or certain aspects of flight dynamics, which may bias the interpretation. The tested equipment comprises of a reference and a UAV-grade LiDAR that are embarked on the same airborne platform together with a reference and a small «UAV-type» camera, as well as reference and small inertial measurement units. The system is flown by helicopter over an area at a speed that is typical for small UAV's (multi-copter) to provide the same conditions for all sensors in terms of temperature, dynamics and height AGL (Above Ground Level). All sensors are rigidly

LiDAR sensor	Height AGL [m]	GSD [cm]	TOF [min]	# img	Speed [km/h]	Nominal Density [pst/m²]	Overlap Along/Across [%]	Beam Divergence [mrad]	Point Meas. Rate [kHz]	Scan Rate [Hz]	Return Mode
VQ480	230	3.7	10	50	50	70	75/35	0.27	200	100	multi
Puck	50	1.8	20	200	12	90–100	75/35	3	300	10 × 16beams	dual

Tab. 1: Flight and acquisition parameters for both LiDAR configurations (TOF - flight duration over zone).

mounted to the same assembly that is vibration dampened. Three different IMU's are used: (1) reference, navigation-grade INS (AIRINS, Ixblue); (2) UAVgrade MEMS-IMU (APX15, Trimble-Applanix); (3) 4 low-cost MEMS-IMU's (NavChip v1/2011, Thales) mounted on a modified Gecko board [2]. The GNSS signal provided by an airborne-grade dual frequency antenna is split between Javad (Delta TRE-3) and APX15 (Trimble) receivers, the former providing the time scale (1 PPS) to AIRINS, NavChip IMUs as well to two Li-DAR devices: (1) reference, medium-range



Fig. 1: Assembly of 3 IMU's, 2 LiDARs and 2 cameras.

VQ480 (Riegl); (2) an automotive LiDAR PuckLite (Velodyne). Two cameras, an IXAR180 (PhaseOne) with 80 megapixels and 42 mm lens, and an A6000 (Sony) with 20 mm lens and 24 megapixels, complete the assembly depicted in Figure 1.

#### Software

Trajectories were computed using the following software: APPS (IXBLUE) for integrating AIRINS data in loosely coupled manner with PPK (Post Processing Kinematic) obtained via Novatel's GrafNAV; POSPAC (Trimble/Applanix) for tightly coupled integration of APX15 observations, Posproc 2.0 (Applanix) with internally designed filter for loosely integrating NavChip (either raw, pre-calibrated or redundant) observations with PPK. The registration of point clouds was performed with the Riprocess suite (Riegl) and LIEO [3]. The orientation of imagery was computed from INS/GNSS trajectory with CAMEO [4].

#### **Experimental site**

The chosen area has a surface of 60 ha (300 m  $\times$  2 km) and features different



Fig. 2: Experimental site near Aclens, Switzerland. Red dots represent Ground Control Points (GCPs) and blue dots the check points.

terrain types including urban and rural areas, forest, agricultural fields, roads, railroads and power lines. Twelve GCPs were distributed over the zone and their coordinates measured by static GNSS survey with an accuracy better than 2 cm (Figure 2). Additional check points have been deployed on some surfaces and measured by RTK GNSS over a short base. To calibrate the small optical sensors in terms of boresight and interior orientation parameters prior the test zone, a second, smaller urban area has been used within the same flight. This contained houses with roofs of different slope and orientation as well as a set of 10 GCPs. This area was flown over in clover leaf pattern.

#### Data acquisition

To ensure the same observation conditions as well as the compliance with flight regulation and safety over urban and high voltage power lines the flight was conducted with a helicopter and not with a UAV. The flight was performed during a winter period (Dec. 2017) at 2 different heights above ground level (AGL) to comply with range limitation of each scanner and to prevent interferences between the respective laser beams. The scanner and flight parameters are summarized in Table 1.

## 3. Comparison

The fact of having all sensors installed on the same assembly allows comparing them under the same flight/drive conditions (i.e., dynamic, GNSS constellation, and atmosphere). The evaluation focuses on three components, namely: (1) the quality of the trajectory (mainly the attitude) and its effect on the absolute accuracy of the point cloud, (2) the level of noise in the point cloud, (3) the capacity

	Gyro drift [°/h]	Acc. Drift [mg]	Accuracy PPK RP/H [°]	Accuracy PPK XYZ [m]
AIRINS	0.01	<0.5	0.002/0.005	0.03
APX15	10	2	0.08/0.03	0.03
NavChipV1	10 in-run			

#### Tab. 2: IMU manufacturer specifications.

to detect small objects and the ability to penetrate the high vegetation. As photogrammetry is also an airborne mapping technique, we compare over the same (open) area the point clouds derived from UAV-LiDAR and UAV-photogrammetry (performed the same day with a high-end Sony camera RX1RII). We present more details about UAV photogrammetric point clouds quality in a separate contribution [5].

#### **Orientation aspects**

Direct georeferencing (DG) or direct sensor orientation (DiSO), is a mandatory component of mobile LiDAR systems that has an important influence on the accuracy of the point cloud [6]. The error budget of a calibrated system with high-quality IMU flying at low altitude above the terrain is dominated by the positioning component, while the influence of attitude errors rises with the flying height and lower quality inertial sensors [6]. Irrespectively of the IMU quality, the attitude error in heading is generally 2-3 larger than in roll and pitch. The attitude quality between small IMU types can be evaluated using the navigation grade (the most accurate) system as reference. Indeed, AIRINS has one of the most accurate fibro optic gyros used within the airborne mapping market. The APX15 has small temperature-calibrated sensors that are aimed to professional UAVs. Table 2 shows the manufacturer specification of each tested IMU.

According to manufacturer specifications, the NavChip sensors are of a lower quality than those of APX15. However, as the Gecko4Nav board comprises of 4 Nav-Chip MEMS IMU, these can be used in a redundant way (SIMU). For that we adopt the method of synthetic low cost IMUs, described in [7], that involves switch-on bias determination, with the goal to reach – in terms of noise level – a comparable performance to Applanix APX15 while having smaller systematic effects. Indeed, this fusion method seems to be promising in reducing the cost of used IMUs on close-range sensors or inertial photogrammetry while shortening the duration of the in-flight calibration procedure.

#### Noise aspects

The use of automotive low-cost LiDARs for mapping is on one side interesting due to its price, but challenging on the other hand, because these sensors were designed for a different application. If the range's noise for mapping LiDAR is about 1-2 cm even at long distances (>500 m), the automotive sensor based on solid state LiDAR exhibits a noise level of 3-4 cm as assessed in laboratory conditions with rigid targets [8]. Hence, the characterization of noise level on various mapping surfaces is novel and relevant information. Practically, variability of the point cloud within a surface of known shape (e.g., planar) can be studied as an indicator in terms of amplitude and distribution (symmetric, biased, etc.). In absence of smooth surfaces, or as an alternative, it is also possible to analyze details of repetitive objects such as tile patterns on a roof or railroad tracks. This characteristic is less metric but indicates the smallest detection of detail(s) over a certain distance.

# Small object detection and vegetation penetration

Detection of small elements such as wires or penetration of vegetation is one of the main advantages of LiDAR technology with respect to passive optical sensing. For this reason, we compare the automotive sensor with the mapping standard also in these aspects. Detection of wires, thin objects and vegetation penetration is compared in the airborne scenario with the reference LiDAR VQ480 (Riegl). As the nominal point density is equivalent between both sensors, comparison can be done by counting ground points, or sub canopy points over a common area. As Puck (Velodyne) features only 2 echoes (first and last), chances to detect intermediate vegetation layer with this sensor are principally lower.

## 4. Results

#### **Orientation effects**

The first analysis represents the comparison between AIRINS and APX15. At each image projection center, we computed the direct georeferencing parameters (position and attitude) considering the calibrated boresight angles. For roll and pitch, APX15 offers an accuracy of 0.025° that corresponds to the specifications while for heading, the RMS is about 0.11°, so slightly worse than manufacturer specifications (0.08°). The RMS for Navchip SIMU are somewhat similar for roll and pitch while slightly better for heading (0.08°). More details on Navchip IMUs with continuous comparison to AIRINS are reported in [7].

To express the real and continuous effect of attitude quality on point clouds we generate them using the same high-precision LiDAR with different IMUs: AirINS (as reference), APX15, single Navchip without and with pre-calibration, and the SIMU. The comparison of different point clouds is depicted in Figure 3 for a single flight-line as an altimetric difference map (dZ). In all cases, we can notice that discrepancies increase at the edge of the swath. This is a typical amplification of angular errors, here at AGL > 200 m. For the Navchip, the performance increases significantly using pre-calibration with fusion.

Now we investigate two along-track profiles within the point cloud marked by 2 red lines on hillshade in Figure 3 (top left). We notice that the Z error is due to planimetric error at places where terrain is rising



Fig. 3: Reference hillshading (top left) generated with AIRINS and the differences in elevation for point clouds based on (1 – top right) APX15, (2 – bottom left) Navchip without precalibration and (3 – bottom right) Navchip-SIMU. The scale goes from 0 (green) to +/-10 cm (red/blue). White areas have errors greater than +/-10 cm.

(Figure 3 top left). This is mostly likely explained by a combined effect of pitch and heading error. Indeed, the heading error in APX15 in the range of  $0.1^{\circ}-0.2^{\circ}$  influences the edge of the swath (about 130 m from the swath center) by a displacement of 30-50 cm. In both cross-sections in Figure 4 we can notice for both Navchip and APX15 a planimetric error of about 30-45 cm. Further analysis reveals that on the cross strips (Figure 4), the synthetic Navchip provides less discrepancies.

#### DTM accuracy and noise

To estimate the DTM accuracy of the automotive LiDAR we have oriented Puck (Velodyne) and VQ480 (Riegl) observations with the same reference IMU (AI-RINS). Then, we extracted the ground with the same extraction parameters in Terrascan (Terrasolid) and controlled the results manually. The absolute Z error is obtained from GCPs (Table 3) and via a subtraction of regular altimetric grids (DTM) based on these classified point



Fig. 4: Cross-section 1 (top) and 2 (bottom) at edge of swath. Red cloud is AIRINS reference. Planimetric difference is the main error, not detected on flat surface but manifesting itself immediately with slopes. Units in both horizontal and vertical axes are meters.

clouds (Figure 5). From this study we can see that:

- The noise distribution of Puck is not symmetric and is centered below the real surface. The systematic component (average) is about 12 cm and the standard deviation is about 20 cm.
- The mean range bias affects the automated process of ground extraction within the Terrascan algorithm as that is based on triangulation of lowest points (after elimination of isolated outliers). Handling this issue would require employing denoising, however, most of the available algorithms work on symmetric distributions and therefore do not remove asymmetric noise [9]. This fact explains the observed systematic differences in Z between both point clouds.

The first two rows of Table 3 show the statistics of GCPs for various DTMs. The comparison with GCPs confirms the bias and higher noise of the Puck cloud. The 3<sup>rd</sup> line uses the same analysis for the Puck cloud when generated with APX15. The last line compares the photogrammetric DTM extracted form dense point matching of the RX1RII images. We see that the later provides a comparable performance to the DTM based on VQ480 LiDAR.

To deepen the analysis of the noise, we look at local variability on planar surfaces

## Photogrammétrie/Télédétection



Fig. 5: Grid differences between Puck and VQ480 DTM with a color scale and histogram (upper left) and cross-section (bottom right). In the cross-section, red dots represent VQ480 and green dots Puck. GCPs are the blue and yellow dots (blue = static GNSS, yellow = RTK GNSS).

	Mean dZ [m]	STD dZ [m]	Min/Max dZ [m]
VQ480-AIRINS	0.005	0.031	-0.065/0.054
Puck-AIRINS	-0.103	0.052	-0.235/0.033
Puck-APX15	-0.082	0.085	-0.265/0.063
Photogrammetry RX1RII	0.003	0.034	-0.082/0.074

Tab. 3: Altimetric statistics of DTM errors evaluated at GCPs.

of  $6 \times 6 \text{ m}$  such as roads/pavements, which is depicted in Figure 6. There the calculated standard deviation  $(1 \sigma)$  of Puck is 4 cm, while for both the VQ480 and photogrammetric point clouds it is about 1 cm.

Lastly, we analyze the level of noise in the Puck point cloud as a function of the range. The manufacturer specification states the noise level about 3-4 cm for ranges shorter than 30 m. As flying above a natural surface lower than 30 m is neither safe not practical, we took a flightline flown at 50 and 60 m AGL and measured the point cloud thickness at nadir direction. The results are shown in Figure 7. As horizontal distance from nadir increases, the range gets longer with a tangent of beam angle (to nadir) while the incidence angle of the beam decreases. The footprint of laser beam augments with the distance and inversely with the incidence angle on surface. This affects the level of noise that increases 4 to 5 times towards swath extreme compare to nadir direction.

#### Vegetation penetration

The main advantage of LiDAR over photogrammetry is its capability to make

	Difference to plane -8 cm (blue) to 8 cm (red)	Noise 1 <del>o</del> [m]	Min [m]	Max [m]	Point #	Pavement rugosity [m]
VQ480		0.011	-0.040	0.030	2635	
PUCK		0.041	-0.150	0.130	6147	<0.01
RX1RII (Photogr.)		0.011	-0.040	0.040	6318	

Fig. 6: Local variability on a planar surface for different point clouds represented by standard deviation of difference to plane. With both VQ480 and RX1, un-flatness of the real surface can be seen contrary to the Puck results. The second column in Figure 6 displays the micro relief of a pavement, which is visible in the VQ480 and photogrammetric clouds but not in the Puck cloud. The analysis of a micro relief is indeed relevant for detecting pavement deformation (ruts, potholes) or rail tracks.



Fig. 7: Puck point cloud noise on a flat surface as a function of horizontal distance from nadir at two flying heights.

observations under canopy. This includes the ground but also intermediate layers of the canopy. Such characteristics are shown in Table 4 for three forest areas composed of deciduous (Area 1), deciduous and conifer (Area 2) and conifer (Area 3) trees. For each area, we have computed the density of ground points. Puck penetrates canopy efficiently in deciduous area, but its penetration is poor in the conifer area. VQ480 provides 10 % more returns from ground than Puck despite having lower nominal sampling density. It is partly due to the fact that at 230 m AGL the VQ480 is far from its range limit whereas Puck at 50 m AGL is at the limit of its ranging capacity, hence the returns penetrating the canopy are likely too weak to be detected.

Figure 8 illustrates in cross-sections the rendering of point cloud canopy and penetration to the ground for both systems. In deciduous area, the Puck signal penetrates to the ground, but the resulting rendering of canopy is poor as the returns are mostly from trunks and large branches. This situation is opposite in conifer area: few returns from ground and well reconstructed canopy. This tends to confirm that the detected reflections of Puck come only from larger surfaces.

## 5. Conclusions

The study was motivated by very optimistic spotlight advertisements presented by the UAV mapping sector for the evaluated LiDAR sensor. We first designed a practical comparison between photogrammetry, LiDAR and inertial UAV sensors with respect to larger airborne equipment. The somewhat unique capability to test in a single flight all sensors was possible through a realization of a special mount holding all sensors on one rigid platform that is carried by a helicopter at altitudes and speeds mimicking UAV flights. This created the same physical input to all sensors and thus the possibility of fair comparison.

We have shown that the use of MEM's IMU such as Applanix APX15, commonly used in the UAV market, conforms to



Fig. 8: 0.8 m wide cross-section in canopy for Area 3 (top), 2 (middle) and 1 (bottom). Green=Puck, Red=VQ480.



2	Puck	102	86	17	0.17	2000
3	Puck	97	92	5	0.05	1660
1	VQ480	78	50	28	0.36	750
2	VQ480	98	72	26	0.27	2000
3	VQ480	70	60	10	0.14	750

Tab. 4: Density statistics for three forested areas (1 = deciduous, 2 = deciduous and conifer, 3 = conifer) for Puck and VQ480 Li-DAR.

manufacturer specifications in roll and pitch but its performance is slightly worse in heading. It showed also practically the limit of such IMUs when flying with LiDAR at higher altitude. Indeed, to keep the planimetric errors induced by attitude errors (mainly heading and pitch) below 10 cm, the flying height above terrain should be limited to 80 m. The same study showed also an important perspective for low cost redundant IMUs based on a synthetic fusion technique that is rivaling the APX15 performance at a fractional cost of hardware despite being based on relatively old sensors (manufactured in 2011).

As employment of Velodyne's Puck was a trend by many service providers at the time of this study, we addressed a deeper analysis on its potential and limits for airborne and mobile mapping. In this respect, we were a bit disappointed by its aerial performance, where its actual noise is far beyond the published specifications for distances > 50 m. Indeed, the resulting DTM has a noise of at least 5 cm (1  $\sigma$ ). On the practical aspects, the ground extraction from its observations is more time consuming because of the asymmetric noise for which the tested denoising techniques were not efficient. Nevertheless, the sensor can be used on small projects for forested areas or for wire / power-line detection at AGL of 30-40 m. In other words, there is a certain consistency in the price and performance ratio between the two LiDAR instruments for airborne applications (~8 k $\in$  vs. 120 k $\in$ ). In this respect, employment of a sensor like Riegl MiniVUX on UAVs has lower noise and higher ranging capacity, however, has larger volume, weight and price.

#### References:

[1] Colomina, I., Molina, P., 2014. Unmanned aerial systems for photogrammetry and remote sensing: A review. ISPRS Journal of Photogrammetry and Remote Sensing, 92: 79-97.

[2] Kluter T., 2012. Gecko4nav Technical reference manual v. 1.0., Bern University of Applied Sciences.

[3] Skaloud, J., 2017. LIEO – Lidar orientation and point cloud generation, Software description No. 6.0704, Swiss Federal Institute of Technology Lausanne (EPFL).

[4] Skaloud, J., Legat, K., 2006. CAMEO – Camera exterior orientation, Software description No. 6.0697, Swiss Federal Institute of Technology Lausanne (EPFL).

[5] Gressin, A., Vallet, J., Bron, M., 2020. About photogrammetric UAV-mapping: which accuracy for which application? The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, XLIII-B1-2020.

[6] Glennie, C. L., 2007. Rigorous 3D error analysis of kinematic scanning LIDAR systems, J. of Applied Geodesy, 1:147-157.

[7] Clausen, P., Skaloud, J., 2020. On the calibration aspects of MEMS-IMUs used in micro UAVs for sensor orientation. Proceedings of IEEE-ION Position Location and Navigation Symposium (PLANS), Portland, OR, USA.

[8] Glennie, C. L., Kusari, A., Facchin, A., 2016. Calibration and stability analysis of the VLP-16 laser scanner, International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences. XL-3/W4, 55–60.

[9] Mugner, E., Seube, N., 2019. Denoising of point clouds. The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Volume XLII-2/W17, 217–224.

Julien Vallet Helimap System SA Epalinges, Switzerland julien.vallet@helimap.ch

Adrien Gressin Insit Institute University of Applied Sciences Western Switzerland (HES-SO/HEIG-VD) adrien.gressin@heig-vd.ch

Philipp Clausen Jan Skaloud Geodetic Engineering Laboratory (TOPO) Swiss Federal Institute of Technology (EPFL) Lausanne, Switzerland philipp.clausen@epfl.ch jan.skaloud@epfl.ch