DATA CENTER ENVIRONMENTAL SENSOR

for safeguarding (not only!) the CERN Data Archive

CHEP2016

Julien Leduc, Quentin Genoud, German Cancio Melia, Daniele Francesco Kruse, Eric Cano, Steven Murray, Vlado Bahyl - CERN IT

Today's tape bits are now smaller than most airborne dust particles or even bacteria. Therefore tape media is now more sensitive to contamination from airborne dust particles. These can scratch the exposed magnetic substrate when mounted in the tape drive resulting in the loss of significant amounts of data. To mitigate this threat, CERN has prototyped and built custom environmental sensors hosted in the production tape libraries that sample the same airflow as the surrounding drives.

WHY DCES?

Tape libraries are highway for airflows that require continuous environmental monitoring of dust, temperature (<32C) and relative humidity (<80%).

Available realtime solutions are mainly derived from clean room monitoring, where rooms are specifically designed to minimize dust propagation and sensors are sized for full room monitoring. Such a solution has prohibitive aquisition and maintenance costs.

Another solution comes from environmental pollution sensors that are monitoring air quality in polluted areas of cities. Those sensors are measuring slowly evolving metrics and their measurement period is around 30 seconds.

Each tape librariy is absorbing 13 m³ per minute

Those environmental sensors also require regular maintenance and calibration.

Additionally most devices fitting our use case do not physically fit inside the libraries, which means that we would have to modify libraries to accomodate the sensor needs...

In our case the tape libraries impose the form factor and the specifications of the hosted sensors that must be validated by tape library vendors before integration.

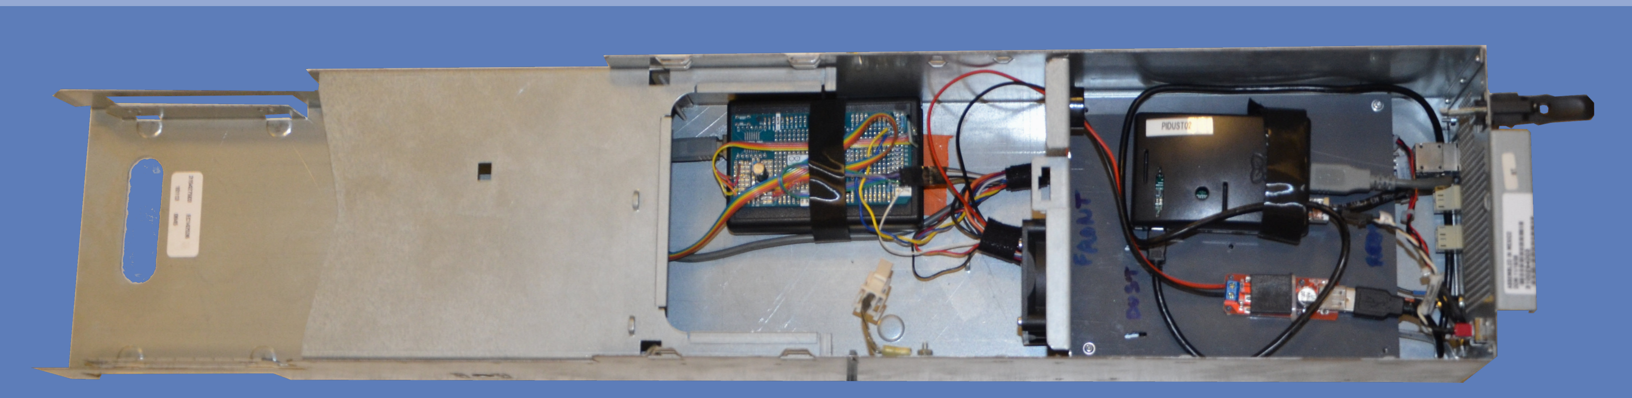

DCES sensor retro-fitted inside a standard tape drive tray

HARDENED DESIGN

DCES is composed of simple, robust, widely tested, cheap and widely available components.

Its brain is an arduino board that handle realtime ADC aquisitions and computations before sending those to an embedded computer: a Raspberry Pi 2 with no mechanical components (no fan, no HDD...).

It uses an industrial HVAC dust sensor coupled with a calibrated temperature and humidity sensor soldered on a generic PCB (arduino shield).

Its electronics has been successfully tested in a magnetic field up to 0.5 Teslas.

DCES electronics during magnetic field torture test

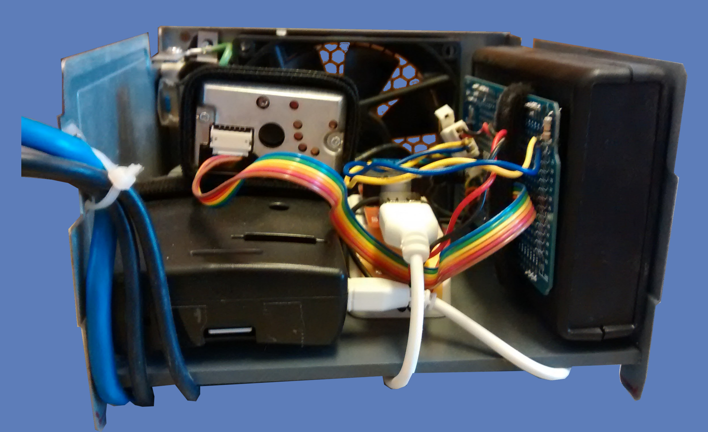

DCES retro-fitted inside a standard ATX PSU: portable WIFI DCES

DCES simple straight traversing airflow limits dust accumulation in the sensor.

Its fast measurements allows it to sustain strong airflows: in the tape libraries, the dust sensors are continuously sensing 1.1 m³ per minute.

Strong airflows allows to continuously blow dust outside of the sensor: it is continuously cleaning itself and divert dust flow from surrounding equipment.

4 DCES sensors have been running inside CERN tape libraries for the last 15 months without any maintenance.
They have gone through several power cuts, and have been sustaining stronger airflows than the surrounding tape drives. This allows them to ingest more dust than the tape drives, which means better sampling and additional dust protection for the tapes.

HOW DOES IT WORK?

An IR light source (850 nm wavelength LED) illuminates the sensor through hole allowing passing though dust particles to diffract light. A photodiole on the other side of the detector measures the scattered light intensity.

An Arduino microcontroller samples the scattered light intensity at 100Hz. It returns the average intensity and standard deviation of the 100 measurements performed during the past second.

Sensor through hole

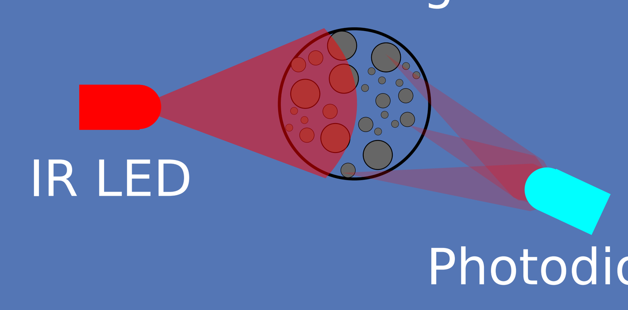

IR LED produces light scattered by traversing dust particles

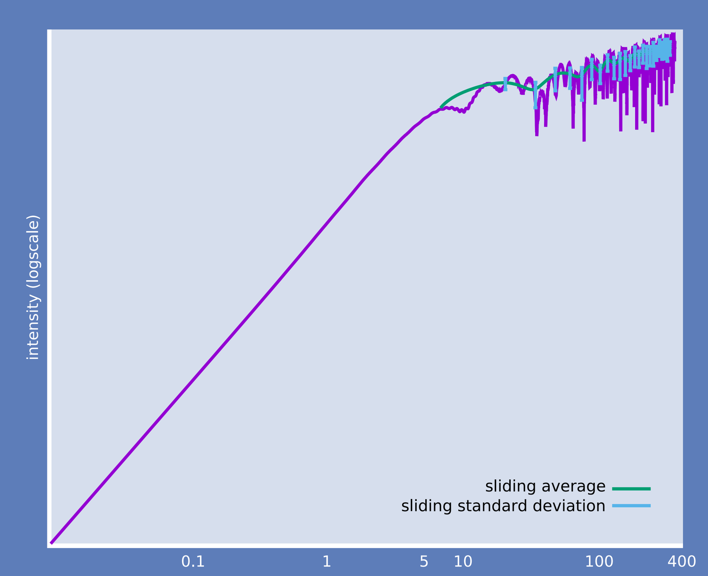

Measured intensity according to size factor relative to wavelength

Since dust particle diameters are similar to IR LED wavelength - up to a few microns - Mie scattering is valid for this domain.
For dust particles with a size between 0.5 and 4 microns: traversing particle size is proportional to the measured intensity. For those small particles, sliding window standard deviation is low.
For dust particles larger than 4 microns: particle size is still an increasing function of measured average intensity but not linear any more. For those large particles sliding window standard deviation is increasing with particle sizes.

During early prototyping DCES small dust particles detection sensibility and speed was compared against a commercial laser counter.

DCES was 12 times faster and was more precise than the commercial sensor because it could detect small dust particles peaks earlier and measure peaks at their highest intensity.

DCES also reports calibrated temperature and relative humidity measurements for completeness and dust calibration requirements.

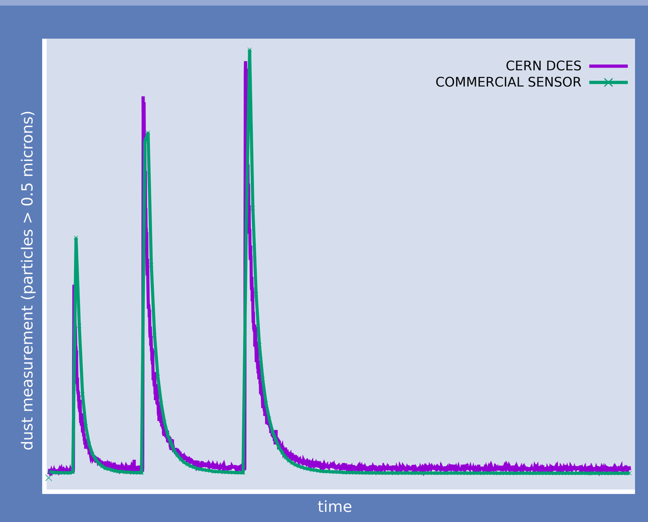

CERN DCES vs commercial sensor

EASY INTEGRATION

CERN DCES hardware design is available as opensource, under CERN Open Hardware License v1.2.
Everything is available on the open hardware repository website:
http://www.ohwr.org/projects/dces-dtrhf-ser1ch-v1
The PCB can be freely modified to fit a specific formfactor, add sensors or external communication interfaces...
Since it is arduino based many external libraries allows fast and easy integration of additional hardware.

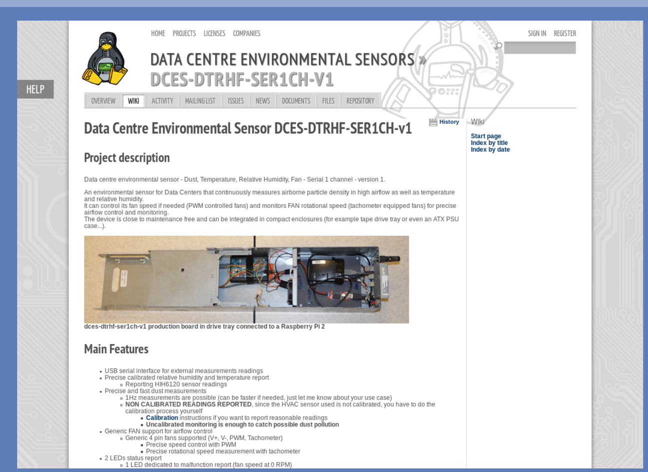

DCES project website on www.ohwr.org

SImilarly to hardware, the custom firmware is available as opensource under GPLv3.
A commercial license is available from CERN if GPLv3 is not suitable, please contact us if you are in that case.

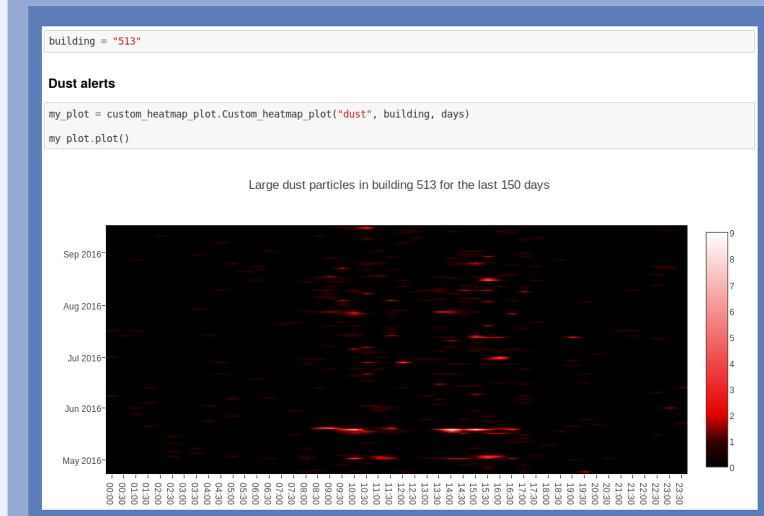

DCES Jupyter notebook integration for reporting

DCES easily integrates in the existing monitoring infrastructure. At CERN tape library dust sensors are puppetized and run a standard SNMP agent that sends SNMP traps when anormal levels of dust are reached.
It is sending dust information through a lemon sensor and into CERN elasticsearch instance.
For detailed analysis we are populating an influxdb timeseries database that allows custom analysis in Jupyter notebooks.

Integration possibilities are endless thanks to DCES being a fully open source project.
DCES does not rely on any CERN specific IT infratructure and has already been / is being integrated outside IT department - experiment and facility caverns - and CERN - scientific institutes, media companies... -.

DCES influxdb and Grafana integration

For more information Web: http://cern.ch/DCES Contact: julien.leduc@cern.ch