EPFL, LAMS
« Industrial Automation » course

Dr. Yvonne-Anne Pignolet
yvonneanne.pignolet@epfl.ch

Dr. Jean-Charles Tournier
jean-charles.tournier@epfl.ch

ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Solution Sheet 3

1. What is the purpose and what are the main parts of an architecture?

   Purpose : guide implementation, support reasoning about elements, structure and behavior

   Parts: list important elements and their interaction

   In particular, for each layer the HW and SW components involved should be described on a high-level, in a way that ensures the reader understand the structure of the plant and how the functionality is mapped to the components and why choices have been made.

2. The CERN data center hosts all computing, administrative and scientific infrastructure of CERN, which includes more than 10,000 servers hosted in three rooms running 24/7. A remote extension of the data center is hosted at the Wigner Research Center for Physics in Hungary, which is connected to the main CERN campus through two independent and dedicated 100Gb/s fibre optic lines. The LHC experiment alone produces over 30 petabytes of data per year, and the center has more than 130 petabytes of stored data currently. To have such a massive computing center running, two main technical infrastructures are critical: power distribution and cooling and ventilation. The cooling and ventilation of the computing rooms is done through cold air introduced in the building via big pipes coming from the roof and going down to the floor. Three chillers on the building roof are responsible for pushing down the air into the building. The supervision is critical, since the servers will be stopped if the supervision system fails, to avoid overheating and fires.

   Consider yourself an employee of a system integrators bidding for the cooling automation system. Your colleague asks you for help with the following high-level architecture.

   “The system is distributed, with one PLC assigned to each room, and another one for the chillers. Server room PLCs communicate directly with the SCADA system and with the remote I/O deployed locally at each server box. We use CAN fieldbus connections. Each remote I/O (one per server box) regroup temperature sensors, the command and position of the grille. Room PLCs also receive data directly from the humidity sensor. The PLC dedicated to chillers could command the air flow, the air temperature and humidity. It also verifies that the air is at the desired temperature and the humidity level at the output of chillers.”

   To be able to give your colleague constructive feedback, answer the following questions

   a) Is the architecture complete with respect to your answer of question 1? If not, what is missing?

      The structure of the plant should be shown explicitly (what belongs to which layer), mapping of functionality, interactions, and justifications are missing,

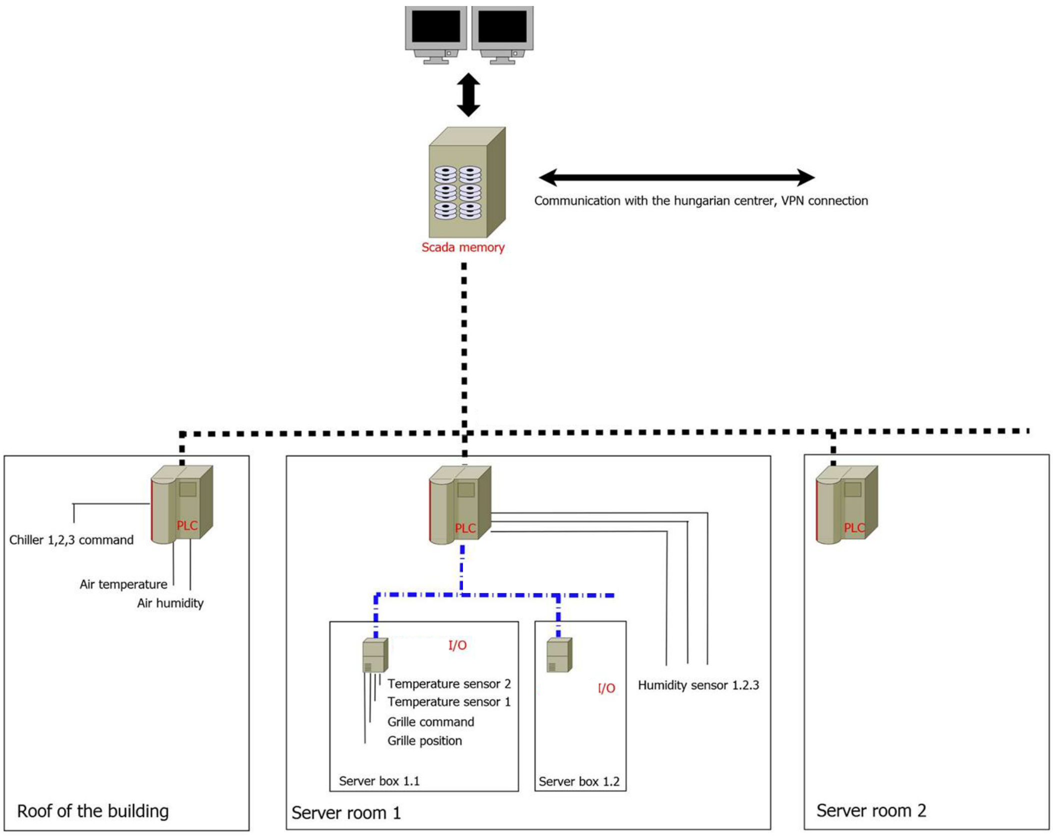

b) Which parts of the architecture figure and the description do you understand? Where do you have questions? What additional information on assumptions and justifications is necessary?

What are the sensors and actuators? Where are they, how many? Why? How and what do they communicate to the controllers? Why CAN for the fieldbus? What kind of controllers do you use? Where are they? Why? Distinction between field level and control level of communication? Why a distributed architecture? Why remote I/O? What do you mean by regroup? What about the SCADA system? Are there any constraints for interoperability, e.g. do the data collected by the supervision needs to be shared to other systems at CERN (e.g. high level monitoring or direct fire brigade monitoring), or is there already any legacy SCADA system which we need to connect to? ...

c) Improve the figure and the description based on your answers a) and b).

Server rooms of data centers must stay within certain temperature and humidity ranges to function optimally and to prevent hardware failure. If the temperature rises too high, gear will begin to malfunction or become damaged. Internal components begin to swell and pull away from each other (or simply burn-up). The ideal temperature range is 18-27 °C, see ASHRAE (American Society of Heating, Refrigeration, and Air-Conditioning Engineers) guidelines for data processing environments. Humidity must also be watched to ensure all gear work in the data center. If the air in the data center is too humid, water may condense on the internal components, resulting in shorts. If the air is too dry, a data center risks static build up, which can also result in electrical shorts. Typically, a data center will also want to maintain its humidity concentration within the 20-80% range. Thus cooling and environmental control is very important. If any part of the data center were to fail, it could result in a lot of damages and expensive repairs or replacements. In order to ensure that our data center's environment is optimal for server function, we will need at least to monitor the temperature, humidity, and airflow. Fire monitoring also needs be implemented but it is out of the scope of this work. Furthermore, automation layers above SCADA are not considered in this proposal.

Before explaining our control system design, we describe how the cooling of a data center works. Chillers on the building roof push cold the air into the building, air that will go down on the floor level, and will cool down the servers by coming out of holes in the ground located strategically between the racks. To maximize the cooling efficiency, a hot/cold aisle configuration: the front of the racks are facing each other on the 'cold' aisle and expel heat out in their backs to the 'hot' aisle. The doors and roofs placed in the cold aisles increase efficiency by only allowing cold air into the machines. The cold air comes out from the grills in the floor inside the 'cold' aisles. The cold air is introduced in the building through pipes coming from the roof and going down to the floor. Chillers on the building roof are responsible for cooling down the air[1].

[1] http://information-technology.web.cern.ch/about/computer-centre

We assume around 3300 servers are located in Hungary, while 7700 are located in Switzerland. The architecture for HW and SW are the same for the Hungarian and the Swiss data center. Furthermore, these two data centers will communicate with each other. The only difference between the two locations will be in the number of equipment involved.

The cooling automation system can be partitioned into four levels: a division between field level, control level, supervision level, and management level. The first level contains the actors and sensors, it is the interface between the exterior and the software. The next one is the control level where the PLCs are located. Then comes the supervisory level, where an operator can interact with the system through a Human-Machine Interface (HMI). The process data base is also located on that level. The last level is the enterprise one, where the high level orders are taken, such as a change in cooling regulation or a large number of additional servers added for example

We chose to implement twice the same hierarchical control system in the two different countries, resulting in optimized operation. The main goal of the control system is to ensure that the temperature and humidity stay within precise ranges. In order to change the values of those parameters, we will use the velocity of fans. Therefore, the field and control layer of our system will be based on the following components:

Sensors:
Temperature & humidity sensors (two each per hot aisle (redundancy), assuming 12 hot aisles per room, there are combined temperature and humidity sensors on the market (e.g., on Siemens website), chosen for simplicity of system)
Airflow sensor (one per room, examples on Siemens website)

Actuators:
Chillers (on the roof), fans (pushing the air into the server rooms, two per hot aisles (one each for in/out flow))

Remote IO:
Fieldbus connection for sensors and actuators, one per room, communicates to sensors and actuators with 4..20mA and Hart (the selected sensors don't support anything else).

Controllers:
Modular PLCs (full flexibility and extensibility, not special equipment necessary, two per location, primary/backup, can take over control if needed), communicate with remote IO via a fieldbus).
The PLCs evaluate the measurements received against the values for defined by the management team: in our case, the ranges of temperature and humidity levels. If the measured temperatures are lower than the smallest end of the ASHRAE's range (18°C), fans should decelerate to save energy. If they are higher than the largest end, fans should accelerate. In order to do so, the PLC will send a command to the actuators in order to modify the velocity. Conversely, we know that the larger the volume of fresh air brought in from the outside, the lower the relative humidity inside. So if humidity levels are lower than advised, fans should decelerate, and if humidity levels are higher fans should accelerate. Air circulation should increase and decrease according to the amount of activity from the servers. We make the assumption that activity is estimated by monitoring the number of servers that are functioning, in order for the fans to have an average value to follow. Then our control system will implement necessary (small) changes depending on the information received from the sensors.

Communication:
As a fieldbus we use Profibus[2], which provides robust communications and is compatible with the remote IO and PLCs we selected. Profibus can support up to 126 devices, which is sufficient for our purposes. As the physical processes involved are slow, it is not necessary to send sensor measurements at a high frequency. As a compromise between detecting potential threatening conditions early and not load the system to heavily, we propose to send measurements periodically, once per second. No demanding hard real-time guarantees need to provided, a latency of <1s is sufficient. Assuming a standard profibus message size of 250 B per sensor per second with all the necessary information inside results in 3.25 kBps per server room, which can easily be achieved with Profibus, even if we decide to add other sensors to the system. For the communication on higher automation layers we use Profinet and OPC UA due to their widespread use and proven performance.

SCADA:
If some fans are not working, or if sensors measure that humidity or temperature stay too long in a certain range, then the PLC should send a warning packet to the SCADA. The HMI of the SCADA is used to monitor and control the two PLCs. In case of an emergency scenario (emergency button pressed, critical levels measured by sensors, broken connections between hardware, lost server connections, fire) all the fans are put to full speed to guarantee a maximum air flux in the data center. When the alarm does not switch off after certain time threshold, a special group of experts within CERN's organization is automatically alerted. The system can adjust the target temperature value depending on the load of the servers. It also logs all the sensors/actuators values, as well as the alarms' states in a data base. Between the two data centers, communication will be possible thanks to the optic fiber lines. At the Supervision level, we can install a database coupled with a back-up database to be able to extrapolate knowledge from data: one could

[2] http://www.profibus.com/technology/profibus/

see how the seasons influence temperature and humidity values, in order to better detect when something goes wrong. The HMI is composed of two panels, one for each data center location. Both have a main area where the overall status is displayed, and where a "Sensors details" button allows a detail area with every sensors values to be visible or not. The panels also include a "Ventilation system" section so that the operator can interact with the fans by either starting/stopping them all, or by setting a particular speed value to each.

3. Industrial storage tanks are used in many industries to hold liquids or compressed gases for short or long term. Consider a simple installation of oil tanks, each instrumented by a controllable input valve, a controllable output valve, one sensor indicating the current oil level in the tank, five sensors indicating different threshold levels (very low, low, high, very high, spilling).

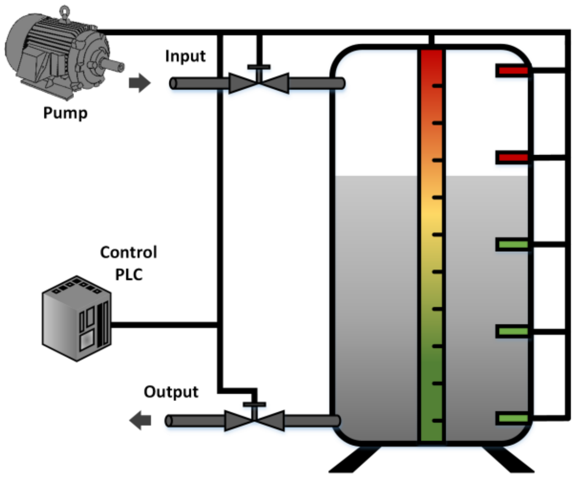

a) Does the instrumentation make sense?

Sensors: temperature and pressure could be added to enable better control, redundancy is provided by this kind of arrangement of level and threshold sensors.

Actuators: for safety-critical substances redundancy needs to be added

b) Do we need real-time communication for this plant?

Since safety is a critical issue, and pressure build-ups could occur suddenly, we need to communicate data with real-time guarantees. The required latency depends on the worst case time during which a combustion could occur and on the communication requirements of other parts of the automation system it is synchronized with.

c) Where in this plant would you use cyclic or event-driven communication?

Cyclic for oil level, temperature, pressure and the extreme threshold levels.

Event-driven for all threshold levels

e) Estimate the bandwidth requirements in this plant.

Assumption: one message ~250 bytes, all values fit in one message, 4 messages per second => 1kbps per tank

d) Is PROFIBUS a good solution for the field level network?

yes: bandwidth requirement supported, 10-15m distance, works with <126 devices, widely used, power for end devices, offers cyclic and event-driven communication

e) Could it make sense to use wireless communication for this installation?

+ could be cheaper (precise comparison would be needed) and easier, but we don't have that much equipment in this installation so cabling should be straightforward and distances are probably short as well.

- interference, in case of safety-critical substances not recommended