Annex: Milestone #2 report

Coordination model, digital twins and design of distributed collaborative ML methods are implemented on an experimental edge-to-cloud architecture

EXPERA report: LASAGNE_AnnualReport_2022

Description

Milestone M2 concerns the design and implementation of the digital platform that will provide collaborative and distributed services supporting advanced user applications for the smart grid and energy management.

This platform is composed of the following elements:

1. **Coordination model**: Provide support for context-aware distributed interactions (information exchange, propagation or queries). We will adapt our most recent coordination middleware with distributed algorithms and mechanisms specific to smart grids (energy consumption, production, forecasting, and the applications developed in WP4).
2. **Digital twins**: We use the concept of intelligent digital twins, equipped with the above learning and autonomous capabilities, to represent the various entities (appliances, services and other Grid Edge Devices [GEDs]) within the coordination model. We define intelligent digital twins for collaborative learning (T3.1, T3.3), as well as specific digital twins for implementing the context-aware self-adaptive applications (WP4, T4.1).
3. **Distributed collaborative ML methods**: A microgrid, employing several power sources and appliances (heat pumps, EV charging stations, HVAC, etc.), involves several distributed sources of heterogeneous and imbalanced data that are difficult to collect timely. To avoid high-latency data transfers and processing at central locations (Cloud), we need optimal forecasting ML algorithms that run collaboratively among the distributed sources at the Edge.
4. **Experimental edge-to-cloud architecture**: Edge-To-Cloud (E2C) solutions will decrease the latency of communications, save bandwidth, lower data management costs, and enhance data privacy. GEDs can boost the performance and the security of applications by offering real-time AI-assisted processing and anonymising data. In LASAGNE, E2C solutions will provide mechanisms by which the edge intelligence can be automatically adapted over time.

Goals

Milestone M2 targets the original proposal's objective O2, here summarized – Develop a digital framework embedding: (1) collaborative machine learning based services to forecast power usage/production at low and middle voltage levels; and (2) collective interactions and

coordination among GEDs. This framework will be used as a toolset to develop end-user context-aware/self-adaptive applications and will support the foundation of the governance policies that regulate technical, contractual and financial interactions among the identified stakeholders. This objective covers two challenges of MICaLL2020: Energy & ICT infrastructure and Energy marketplaces & business models.

More specifically, M2 pursues two goals:

1. Develop a digital framework for collaborative learning among GEDs. Forecasting power usage/production faces noisy and unbalanced measurements, as well as IoT data dispersion through Edge solutions, where ML coordination and computation mechanisms are hindered by bandwidth limitations, latency and message loss. This needs a collaborative ML approach among the different data sources and computational units. To address this problem, we use a model that provides a coordination media and mechanisms, allowing the GEDs to exchange (via intelligent digital twins) information about production/consumption, so as to smoothly fulfill the requirements of context-aware/self-adaptive applications.
2. Setting up the digital marketplace E2C platform. We aim to integrate the outcomes of proposal's objectives 1, and 2 and test them in the concrete case of an end-user context-aware / self adaptive energy application. The LASAGNE framework will rely on the CLEMAP technology for the edge devices and the front-end Cloud infrastructure, backed by the SixSq Nuvla.io/NuvlaBox E2C platform.

Achievements

Regarding WP3, we summarize here the current status of the three planned tasks:

1. **Coordination model:** We deployed a working coordination model on multiple edge nodes in our "Les Vergers" living lab. The coordination model implementation supports various digital twins and algorithms (see point 2. below) that exchange information and interact together through the coordination platform. A prototype E2C infrastructure (see point 4. below) has been set up to gather real energy usage data and run preliminary experiments.
2. **Digital twins and algorithms**: We developed several digital twins working on behalf of consumers, producers, learning and regulator devices (peak shaving controllers):
    - Consumer and producer digital twins interact to perform contracts regarding the amount of energy exchanged (no pricing support yet), so as to optimize energy use.
    - Learning twins forecast production and consumption locally, through various techniques – Markov chain, or other ML mechanisms (see point 3. below).
    - A regulator twin controls the energy levels and stops any producer or consumer twin in case of peaks (anticipating or reacting on actual data).
    - Tragedy of the commons: preliminary dynamic pricing scenario for scheduling energy consumptions and thus avoiding tragedy of the commons where the energy consumption surges at the same time.
    - Game-theoretic framework for real-time energy trading.
3. **Distributed collaborative ML**: ML algorithms are running with real data.
    - By trialing several ML techniques applied to energy data, we have found that the most promising method is the LSTM.

- Centralised Federated Learning: preliminary results for a centralized federated learning of energy production, using an average aggregator of the models.
- Gossip-based Federated Learning: generic service providing decentralized federated learning, where the aggregation function is freely chosen by the applications. This service is provided by the coordination platform.

4. **Edge-to-cloud infrastructure**
    - Several CLEMAP smart-meters are deployed on the different school buildings of “Les Vergers” living lab. Data is stored in a cloud environment and retrieved locally so that we can replay the testbed scenarios with different edge nodes on our own lab premises.
    - Digital twins, which also use the retrieved living lab’s data, cannot yet send instructions to the edge devices located in “Les Vergers”. As a result, edge devices cannot react to simulated peak shaving or energy exchange scenarios. However, distributed learning algorithms will be used to solve this issue.
    - The abstract design of the full-fledged E2C infrastructure is being currently worked on. It involves a stack composed of: applications (for end-users, building managers, grid managers, etc.); LASAGNE services supporting those applications, including the coordination platform and the digital twins; smart meters providing data to the LASAGNE services; and the underlying network infrastructure for the communication among the edge nodes.