WHITEPAPER 1.3

APR 16 2018

DAVIDE RIVOLA, VASCO MEDICI, LORENZO NESPOLI, GIANLUCA CORBELLINI, DAVIDE STREPPARAVA

HIVE POWER

Contents

1 Vision

> Consumers and communities will be empowered to actively participate in the electricity market and generate their own electricity, consume it or sell it back to the market while taking into account the costs and benefits for the system as a whole
>
> — *European Commission*

Our vision is to provide everyone with the possibility to create and manage electrical energy communities on the Ethereum blockchain, creating a world of shared energy for a brighter future. For prosumers participating to the energy communities, Hive Power will lower the energy tariffs and reduce the payback time of renewable energy sources. It will allow selling flexibility as a service to system operators, to help them balancing the grid.

1.1 The challenge

The power sector is facing a paradigm change, moving from a centralized structure with big power plants (hydro, coal, gas and nuclear) to a decentralized scenario of distributed energy resources (DER), such as solar and wind. Within this framework, a new actor is emerging, the prosumer, i.e. households or organizations which at times produce surplus energy and feed it into a distribution network; whilst at other times (when their energy requirements outstrip their own production) they consume energy from that grid. The transition to a prosumer-driven electrical grid can be quite bumpy [1]. The lack of centralized planning and the increase of intermittent electricity production in the lower levels of the grid due to the penetration of solar generation raise the stress on the electrical distribution grid. This can lead to severe power quality problems, especially voltage violations and line congestions, with which the distribution system operators (DSOs) need to cope. These problems are intensified by the increase of electricity consumption due to electrical heat pumps and electric vehicles, which further increase power oscillations within the distribution grid.

1.2 What needs to change

To overcome this problem significant investments in the grid infrastructure are expected. This could trigger a so called business "death spiral", or "grid defection" [2, 3] because prosumers reduce their regular energy consumption from the central grid in favor of energy they produce themselves. In addition, they can install batteries to further increase their energy independence. Most likely, these customers will still depend on the central grid for emergency or peak power use, so electric utility companies will have to operate their costly infrastructure and power-generating capabilities even as revenues from energy consumption decline.
This bleak scenario is not the future we want for ourselves or our children: a future with a widening gap between autonomous prosumers, almost totally disconnected from the rest of the grid, and simple energy consumers forced to pay for a more expensive and inefficient grid.
We envision a future electrical grid characterized by an increase in the exchange of energy between prosumers, consumers and electric utilities, optimizing the use of energy resources and infrastructure. In this context, enabling technologies like blockchain will allow decentralized prosumers to safely buy and sell electricity to each other at negligible transaction costs. New aggregators exploiting this technology can act as energy suppliers and compete in the global market. In this context, distributed energy storage systems (DESS) can also participate in the energy market and thanks to their high flexibility, they could quickly respond to dynamic price signals. This represents a big opportunity for substantial cost savings for the end user.

1.3 Our solution

In this paper we present the Hive Power platform, a decentralized autonomous organization (DAO). The goal of Hive Power is to create energy sharing communities where all participants are guaranteed to benefit from their participation, and at the same time achieve a technical and financial optimum for the entire community. This is achieved by devising a (mathematically sound) market mechanism that incentivizes the participants to collaborate with each other by coordinating their energy production and consumption. In contrast to other energy exchange market schemes, the Hive Power platform takes also into account technical aspects, such as cables, power rating, and voltage limits in order to provide an optimal solution that achieves multiple objectives. Hive Power is perfectly designed for the current grid configuration, enabling the electrical grid to operate safely and cost-effectively by ensuring a fair and resilient energy market for all actors involved.

2 Hive Power Platform

2.1 Concept

Hive Power develops a turnkey solution for the creation and management of local energy communities on the blockchain, called **Hives**. A Hive is a distributed energy market platform regulated through smart contracts where every prosumer can buy and sell electrical energy. Like in the real world, a Hive is composed by a set of **Workers**, who provide for the livelihood of the "colony", and a **Queen**, who rules and coordinates them (Fig. 1).

The Workers consist of blockchain-enabled smart meters, which allow their users to participate in the local market and interact with it through an **User App**.

Each Hive has an **Administrator** who sets it up and manages it. The management of Hives is performed using an **Admin App**, which provides access to a smart contract called the **Beekeeper**. An ERC20 utility token called the **Hive Token (HVT)** provides access to the Hive Power ecosystem (cf. 3).

In the following sections, we will describe more in detail the different components of the Hive Power platform.

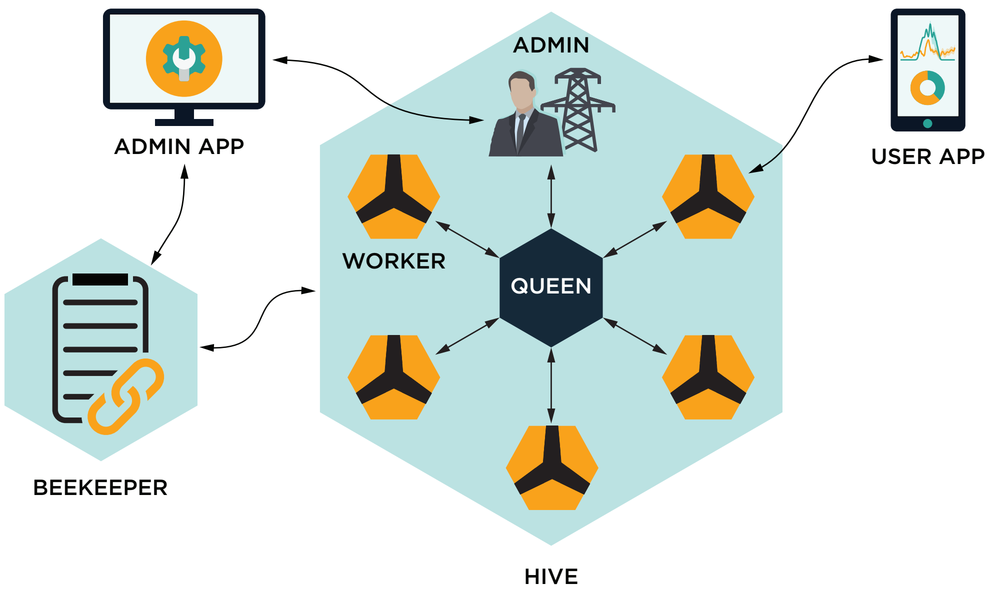

Figure 1: Hive concept

2.1.1 Worker

The basic component of a Hive is the Worker. A Worker is a blockchain-enabled electrical meter that measures energy production and consumption and is linked to a prosumer, e.g. a single appliance, a single-family household, a storage system, a solar power plant, an industrial facility, etc. A Worker measures and certifies its own energy production and consumption, and participates in the Hive energy market exchange. The Worker acts as the interface between the prosumer and the electrical grid and is equipped with a hardware oracle containing a cryptographically attestable anti-tampering sensor (cf. 7.3). The Worker also provides a forecasting service that predicts the energy consumption and production of the prosumer and an energy bidding system interacting with the Hive. A Hive can be composed of many Workers, while a Worker can only belong to one Hive.

2.1.2 Queen

Each Hive contains a Queen, which collects the energy forecasts and coordinates the energy production and consumption of the Workers in the Hive. It aggregates the data, and communicates it back to the Hive in an anonymous format. The Queen collects the energy consumption and production data of the Workers and handles the payments in the energy community.

2.1.3 Hive

A Hive is a distributed energy market platform regulated through smart contracts where every prosumer can buy and sell electrical energy. A Hive is able to interact with the external grid as a single entity, selling and buying energy and services. As shown in Figure 1, each Hive contains of a Queen who coordinates the Hive energy production and distribution, and a set of Workers. Workers in the Hive can exchange energy between each other via the Queen.

2.1.4 Hive Administrator

Each Hive has a unique Administrator. Typically, the Administrator also manages the local grid infrastructure and ensures that the Ethereum Meters are installed correctly. It could be for example the distribution system operator or a multi-unit building manager.
The Hive Administrator retains a small fee of the Workers stable token payments, in order to cover the costs of the infrastructure. Hive Power also receives a small fee, to support the maintenance of the Hive Power platform and finance its development. In the future, the Hive Administrator could be a smart contract with multiple owners.

2.1.5 Beekeeper

In the Hive Power framework the interaction between a Hive and its Administrator is not direct, but it is mediated by a smart contract called the Beekeeper, which will be exhaustively explained in the next paragraph.

2.2 Smart contracts

In order to provide an Ethereum-compliant platform and to manage the aforementioned Hives with the related components (Workes and Queen), a set of smart contracts will be developed and deployed on the Ethereum blockchain. In this chapter the two more significant contracts, named **Beekeeper** and **Hive**, will be explained in detail. Besides, others contracts will be provided to support the functionalities of the main ones (e.g. the management of a list containing the meters allowed to be added to a hive, the management of tariffs, etc.). Figure 2 shows the main interactions between Beekeeper, Hive and the auxiliary contracts.

2.2.1 Beekeeper contract

The Beekeeper contract provides a list of functionalities, the main related to hives management. Through it, a generic Hive Administrator is able to add, modify and delete hives. The Beekeeper exploits other auxiliary contracts to obtain mandatory informations: a **list of the meters** allowed to be used in a hive, a **HVT vault** where the HVTs needed to create a hive can be staked, a

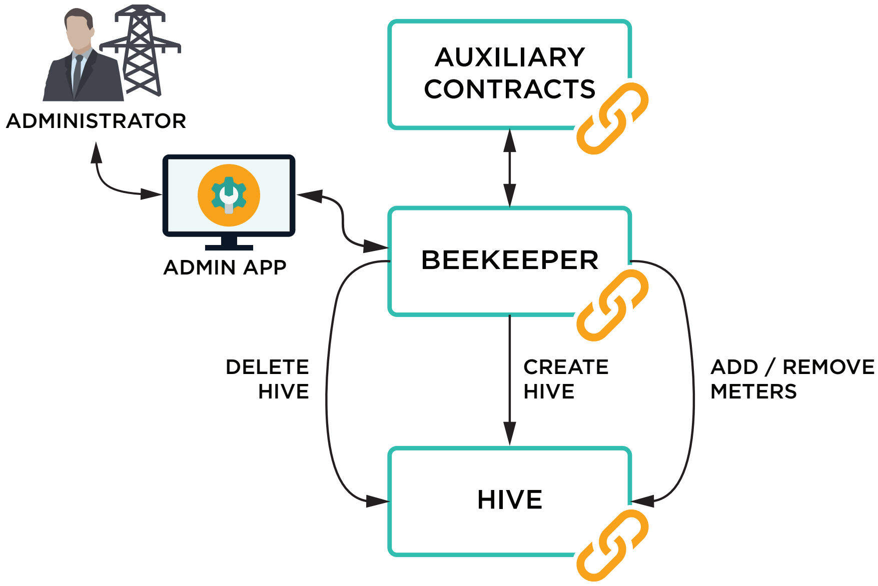

Figure 2: Interactions between the Hive Power contracts

token vault related to the stable coins exploited by each hive, a set of **tariffs** and a **list of hives** to maintain a track of all the created instances. The Hive Administrator is able to modify one of its hives typically adding or removing meters. Besides, Beekeeper has also to manage some of the aforementioned auxiliary contracts, e.g. adding new entries in the meters list.
The Beekeeper contract is upgradeable in order to be compatible with the addition of new functionalities and security updates. In the future Hive Power will set up a governance system to allow all HVT holders to provide input and feedback on the upgrades of the Beekeeper Contract, allowing them to address the development of this service, and as a consequence, to have meaningful influence on the entire project framework.

2.2.2 Hive contract

The Hive smart contract has the main functionality to manage a single hive. An instance of this contract is deployed on the blockchain by the Beekeeper if a sufficient staking of HVTs is provided. As Beekeeper, Hive exploits other contracts to properly work, e.g. the aforementioned meters list with the related features (adding/dropping meters). For each worker, Hive is able to send data about produced/consumed energy and consequently receive/pay stable coins using the related token vault. Similarly to Beekeeper, Hive contract is upgradeable in order to provide improvements in terms of new functionalities and security issues.

2.3 Scalability aspects

2.3.1 The problem

Currently blockchain technology is not exploitable for the management of all the acquired data in typical IoT applications such as Hive Power platform. The amount of measurements is too high to be reasonably stored in a public Ethereum blockchain. The reason is mainly the high gas cost to pay for each transaction on the chain. In the first releases, the Ethereum blockchain will be fundamentally used for administrative aspects (e.g. creating/deleting hives) and to store informations related to monthly production/consumption of energy and revenues/payments of the Workers.

2.3.2 The solution

In the first releases the data needed by the platform (e.g. the workers energy measurements, with a typical resolution of 15 minutes) will be saved in an off-chain custom framework provided by Hive Power. In the future, Hive Power will migrate the custom implementation to **State**

Channels technology, which provides secure, fast and economic micro-transactions and can be automatically connected to Ethereum blockchain. The migration will be actuated integrating Hive Power platform with generic State-channels frameworks like Raiden Network (https://raiden.network) or Liquidity Network (https://liquidity.network).

2.4 Apps

2.4.1 Admin App

The Admin App is a web application used by each Hive Administrator to maintain its own Hives. The access to the Admin App is granted by using the Hive Administrator's HVT wallet. The main features of the application are:

- **Beekeeper interface**: The Hive Administrator is able to interact with the Beekeeper Service for the general management of the Hive (e.g. creation, adding of a new meter/Worker, etc.).
- **Hive supervision**: The Hive Administrator can monitor the Hive status (e.g. statistics about the energy consumption/production of the entire Hive, comparison of the Workers' forecasts, chart of best performing Workers, etc.).
- **Fiat status and withdrawal**: The Hive Administrator can monitor the status of its stable token account and can purchase stable tokens to deposit into the account or can withdraw stable tokens from the account and exchange them for fiat currency.

2.4.2 User App

The User App is the preferred method to interact with the Hive. The main principle behind this app is "install and forget". After the initial setup of the User App, the user can "forget" about it, because the Hive mechanism automatically takes care of normal daily operations. The main features of the app are:

- **Cost savings**: Users are able to see the amount of local energy used by the Hive, and the cost savings compared to regular utility tariffs.
- **Energy stats**: Worker Owners are able to access useful energy statistics of their Workers.
- **Credit fill-in**: Users can register a credit card to automatically refill the Worker's wallet with currencies, using a third-party exchange service. Using an ERC20 stable token as fiat money will effectively protect the prosumers against price volatility.

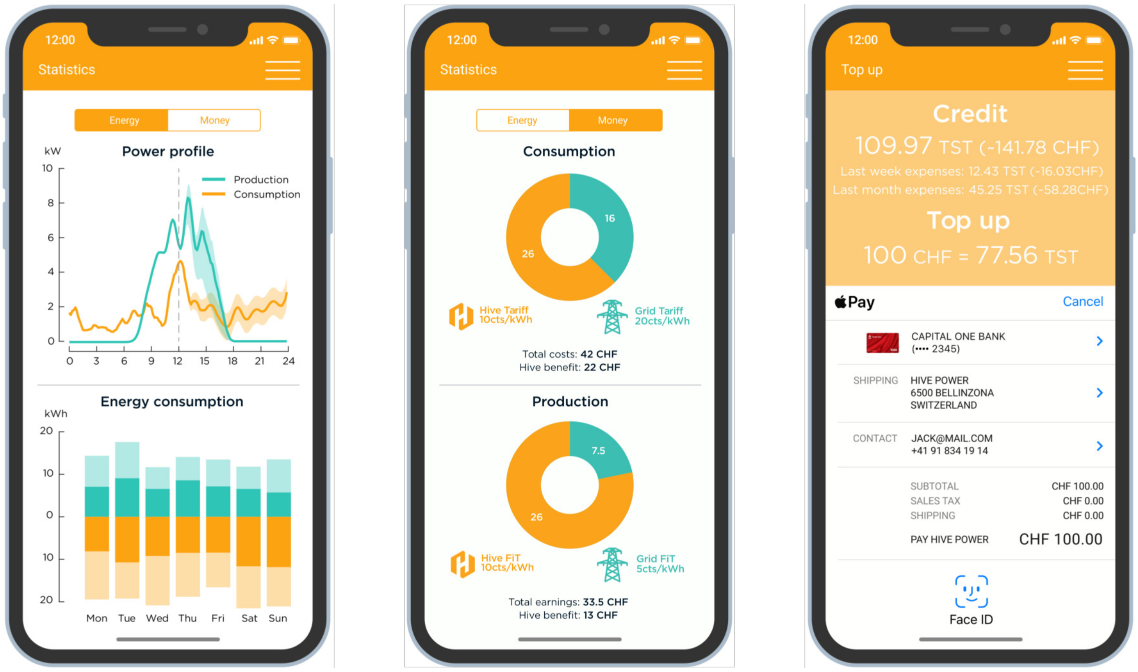

Figure 3: User App

3 Tokens

There are two types of tokens required to operate the Hive Power platform. The first is the HVT token, which is used to create, control, and manage the operations of the Hives and the technical governance of the platform. The second type of token (a third-party stable token) is used for all transactions and payments within the Hive Power platform.

3.1 Hive Token (HVT)

The Hive Token (HVT) is a standard ERC20 Ethereum token managed by a smart contract. The HVT token is used for the creation and management of Hives, including the registration of Ethereum Meters within a Hive, and participation of the Hive Administrator in the technical governance issues of the Hive Power platform. HVTs have a maximum supply (100’000’000) and will be created only once, during the upcoming crowdsale.

3.1.1 Hive management

An admininstrator requires HVTs in order to create a new Hive. To open a Hive, the admininstrator sends the required HVTs to the Beekeeper Contract. The Beekeeper Contract performs a “burn and stake” operation on the HVTs received, that is, 50% of the HVTs are “burned”, meaning taken permanently out of supply, and the remaining 50% are “staked” inside the contract. The burn function is a mechanism to discourage a Hive Administrator from performing unnecessary Hive management operations. It also serves to achieve and maintain the stable functioning of the Hive Power platform.
After the creation of a new Hive, the Hive Administrator can send additional HVTs to the Beekeeper Contract in order to register Ethereum Meters (Workers) in the Hive’s Meter List. When adding meters to the Hive, the Beekeeper Contract follows the same “burn and stake” mechanism as in the Hive creation. When a Hive is destroyed or a meter is removed from the Hive, the remaining 50% of the HVT tokens that were staked but not burned at the Hive creation, are sent back to the Hive Administrator’s HVT wallet.
The cost to create a Hive will be decided before the launch of Hive Power 1.0. The amount of HVT required to create a Hive will vary in order to ensure most equitable cost at all times.

3.1.2 Governance

The Hive Power platform is designed as a community platform for Hive Administrators. Hive Administrators, who stake HVTs in the Beekeeper Contract, will be allowed to participate in decisions regarding the technical governance of the Hive Power platform. More specifically, Hive Administrators can vote on proposals to upgrade the Beekeeper Contract. In order to exclude potential speculative voters, only Hive Administrators will be allowed to participate in the governance process. Moreover, the voting power of Hive Administrators will be weighted accordingly to the date they began staking HTVs in the Hive Power platform. That means that active, long-term Hive Administrators will have more voting power on technical governance issues within the platform.

3.2 Stable Tokens

The Hive Power platform is designed as an open framework that can be used for multiple purposes by any actor anywhere in the world. Although the platform is built on the Ethereum blockchain, the value of the Ethereum token is not stable enough to be used for Hive energy payments without the integration of external price oracles, which is a complex solution. Therefore, to protect against volatility, the energy payments within the Hive Power platform will be conducted via stable tokens that are pegged to national currencies or another type of stable asset. The Hive Administrator can decide at the time of Hive creation the stable token they want to use for energy transactions within the Hive. The Hive platform is capable of managing the interconnection of the Hives using different stable tokens. The interconnection will be ensured through decentralized token exchange protocols such as https://0xproject.com/ or https://swap.tech/.
Stable token frameworks are being actively developed. DAI (https://makerdao.com/) is the first stable ECR20 cryptocurrency pegged to the US Dollar (1 DAI = 1 USD). It automatically adjusts to changing market conditions in order to stabilize its value against major world currencies. Because of its stability, it will be the first stable token integrated into the Hive Power platform. As the Ethereum ecosystem matures, additional stable tokens will be added to the platform.

4 Market Design

4.1 The Problem

The current electrical energy market must be able to manage stochastic power injection from Distributed Energy Resources (DER). In order to reduce these injections, DSOs currently use a simple indirect method, consisting of a bi-level energy tariff, i.e. the price of buying energy from the grid is higher than the price of selling energy to the grid. This encourages individual prosumers to increase their self-consumption. However, this is inefficient in terms of regulating the aggregated power profile of all prosumers.
A more effective way to exploit prosumers' flexibility to increase their own energy production or decrease their energy consumption from the grid, would be to explicitly define a common energy target for the aggregated power profile of the prosumers. The prosumers should be provided with economic incentives to adhere to the target. Depending on the economic benefits associated with production and consumption options, prosumers could choose how to use their flexibility.
An unmet need in the energy field is the ability to accurately forecast energy production and consumption. An accurate forecast of the aggregated energy portfolio of prosumers would represent a huge benefit to energy providers, such as energy retailers, electricity suppliers, or entities responsible for balancing the energy grid, because it would reduce uncertainty when they bid to purchase energy in the energy market. To assure the accuracy of the forecast, energy providers could pay prosumers for their adherence to this forecast.
Being paid to adhere to their energy forecasts, is one revenue stream for prosumers. Another revenue stream comes from the possibility of creating Self-Consumption Communities (SCC), which buy/sell energy as a single entity on the grid. The calculation of financial benefit for SCCs is described below.

Because an SCC is composed of a heterogeneous group of prosumers, the cost of the aggregated energy profile, C_h, is always smaller than the sum of the costs of the Workers, C_i. This virtual aggregation will generate a benefit in terms of money.

$$\sum_i C_i \geq C_h \tag{1}$$

$$\sum_i f_c(E_i) \geq f_c\left(\sum_i E_i\right) \tag{2}$$

$$\sum_{i\in\mathcal{C}} p_b E_i + \sum_{i\in\mathcal{P}} p_s E_i \geq f_c\left(\sum_i E_i\right) \tag{3}$$

where $\mathcal{C}$ and $\mathcal{P}$ are the sets of consumers and producers at a given time, $E_i > 0$ represents the incoming (bought) energy for the i_{th} prosumer and $E_i < 0$ represents the outgoing (sold) energy. The cost function for Workers exchanging energy with the DSO, is defined as:

$$f_c(E_i) = \begin{cases} p_b E_i, & \text{if } E_i \geq 0 \\ p_s E_i, & \text{if } E_i < 0 \end{cases}$$

where p_b and p_s are the prosumer buying/selling prices from/to the DSO, respectively, and $p_b > p_s$. The surplus is defined as $S = \sum_i C_i - C_h$

4.2 The solution in a non flexible setting

If Workers participating in the energy market are not flexible, i.e. they do not have deferrable loads or energy storages, the only way they can generate revenue is to create an SCC. The economic surplus S, resulting from equation 2, cannot be used as a driver for prosumers to coordinate their energy load. However, since $S \geq 0$, Workers will always have an economic benefit of participating in the energy market. In a non-flexible setting, the surplus revenue S is simply redistributed by applying a discount in the energy tariff in terms of [$/kWh]. In this case, no benefit is derived from the aggregated prosumer profile in terms of power quality within the Hive grid, or for the grid with which the Hive interfaces.

4.3 The solution in a flexible setting

If Workers are flexible in terms of energy production and storage, the surplus revenue S can be used as an incentive for prosumers to increase the overall self-consumption and the community welfare.
While maximizing the welfare of a community in a centralized way is conceptually simple, it has in practice many technical implications. The grid's central controller needs to know all of the Workers' local constraints and utility functions, which could include, for example, the desired internal temperature profile for a residential building, or the scheduled time for the electrical vehicle charge. Gathering this information for all of the Workers could be impractical and could raise concerns about privacy issues. If in addition grid constraints are taken into account, the central controller must retrieve the grid topology and cable parameters. Furthermore, the overall problem could be impossible to solve in a reasonable amount of time for a large number of Workers.

Many decentralized strategies to solve the energy load coordination problem have been proposed in recent years [4]. A distributed solution requires Workers to solve a local optimization problem, limiting the amount of shared information. Decentralized approaches to manage energy loads can be classified in three main categories: decomposition techniques, networked optimization, and non-cooperative games [5].

4.4 The HONEY algorithm

Hive Power will solve the coordination problem applying the Hierarchical Optimizer for Networked Energy (HONEY) algorithm. The base version of HONEY considers a single level hierarchy and is based on decomposition techniques. More specifically, it exploits a proximal algorithm called the Alternating Direction Method of Multipliers (ADMM), which is known for its good convergence

and stability properties [6]. Moreover, HONEY parallels the decomposed problems, reducing the time required for convergence, compared to classical Gauss-Seidel-like iterations.
The Workers will locally solve an optimization problem in order to maximize the Hive's well-being for the next 24 hours. Due to the intermittent nature of renewable energies and uncertainty in the Workers future production and consumption, setting the obtained strategy for the next 24 hours period could fail spectacularly, leading to very bad solutions. For this reason, a receding horizon strategy is adopted: only the first optimal action of the whole optimization period is implemented, and the process is repeated again each 15 minutes.
In the coordination phase, which occurs at the beginning of each 15-minute interval, the Workers solves their optimization problems and communicate their solutions to the Hive, which sends them back a coordination signal, allowing Workers to cooperate among themselves. The coordination phase ends when the iterative strategy has reached the optimal consumption/production scheduling for the Workers. Since the Workers do not possess precise information about their future energy production and consumption, the declaration phase must be followed by an actuation phase. The actuation phase begins right after the declaration phase and it ends when the market transactions are cleared every 15 minutes. At the end of the actuation phase, the Workers pay or get paid for the actual amount of energy they have consumed or produced. The surplus is redistributed to the Workers and a new declaration phase begins.
In the following paragraphs, the coordination and actuation phases, depicted in Figure 4, are explained in detail.

4.4.1 Coordination

In the coordination phase, the Workers continuously send their energy consumption and production forecasts to the Queen, who collects the data and communicates back to the Workers the aggregated energy profile of the Hive as well as other useful market signals (e.g. prices). An agreement on the consumption and production profiles of the Workers in the following time interval is reached. The Hive's energy consumption and production profiles are continuously coordinated in an iterative process.

4.4.2 Actuation

At the end of each time interval, the Workers communicate their certified consumption or production to the Queen. The Queen then compares them to those declared during the coordination phase and determines the cash flow of the specific time interval.
To balance the energy produced and consumed within the Hive, the Queen continuously controls if the Workers in the Hive diverge from their forecasted energy production and consumption values, and assigns each Worker a reputation. This reputation is used to reward those Workers who adhere to their forecasts and penalize those who do not.
Inside the Hive, payments are made using a stable token, which is defined at the moment of creation and is prepaid by the owners of the Workers (cf. 3.2). Every day the Queen collects the daily sum of the Hive's energy production and consumption and redistributes the correct number of stable tokens to the Workers and to the Hive Administrator.
Workers payments are always routed through the Queen; in general, Workers pay the Queen when they consume and are paid when they produce. The Queen is the only interface of the Hive towards the outside world. In case the Hive as a whole is consuming electricity and needs to pay for it, the Queen ensures that the Hive Administrator receives enough stable tokens in its wallet to pay the supplier. In case the Hive is producing energy or is providing a service to the grid for which it needs to be paid for, the Hive Administrator needs to prepay the Queen for it, so that the Queen can then automatically redistribute the earnings to the Workers.

4.4.3 HONEY extensions: grid constraints

If the network topology and grid parameters are known, voltage and power constraints can be easily integrated into the decomposition approach. However, network topology is often unknown in low-voltage networks. Thus, in the first implementations, we will assume the underlying network has enough capacity to avoid congestions. In this setting, all Workers have a direct link to the Queen and the connectors have a purely logical function, and are not mapped to the physical network.

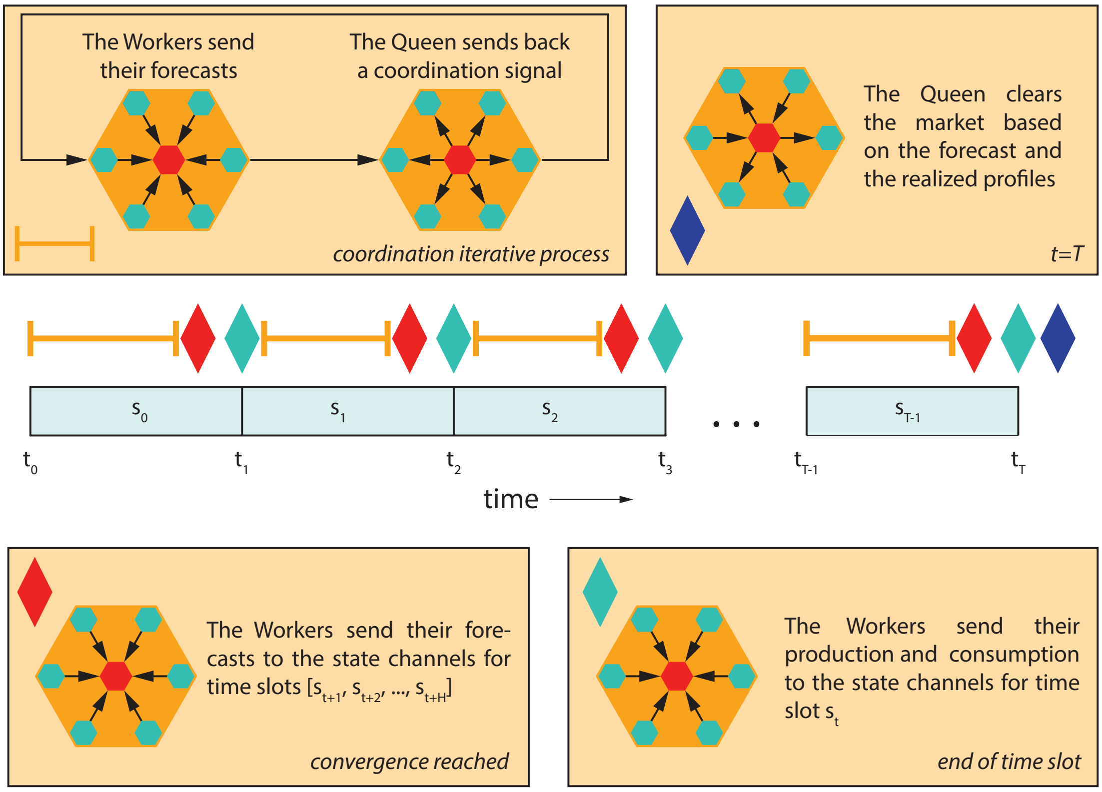

Figure 4: Hive Power market mechanism

One of the core activities of the Hive Power 1.1 phase will be the automatic modeling of the low-voltage electrical grid, based on collected data. Having a model for the electrical grid will allow the Workers to aggregate their power profiles while respecting electrical grid constraints. A feasible way to do this would be to use a linear approximation of the power flow equations [7], which works particularly well for low-voltage networks [4]. This will require an estimation of linear coefficients, linking voltage with active and reactive power measurements called 'voltage sensitivity coefficients'. It has been demonstrated that this can be done using distributed sensor networks of synchronized phasor measurement units (PMU) [8] or even smart meter data [9].
Another activity of the Hive Power 1.1 phase will be the extension of the proposed decomposition problem to a multilevel hierarchical setting.

4.4.4 HONEY extensions: multilevel aggregation

HONEY can be extended to a multilevel scenario (Fig. 5), allowing the creation of bigger prosumers pools. For example, prosumers in the low-voltage grid can be aggregated with prosumers in the medium-voltage grid. While this extension is conceptually easy, the accompanying optimization and communication strategies must be carefully designed in order to bound the required computational time.
A control problem in the electrical grid can be naturally decomposed geographically: for example, Workers connected to different low voltage transformers in a radial electrical grid can only influence their voltage levels indirectly by means of aggregated profiles at their transformers. This means that there will be no need for Workers to know all of the constraints of the grid they are connected to. Instead, they will only need to know the constraints they have influence over. Technically, this can be described as a rooted tree structure, with a Queen in each branching node and Workers in the leaf nodes. In this structure, the Queens' electrical constraints can only be influenced by her descendant nodes. This problem can be seen as a special class of networked optimization, which can be solved by the multilevel HONEY algorithm. From a practical point of view, the coordination phase can be solved with a forward-backward communication scheme. In the forward

passage each Queen, starting from the one located at the top level, sends both her reference signal, and the one received by her parent Queen, to her children nodes, which propagate it downwards through the hierarchy. At the same time, Workers in the leaf nodes solve their optimization problems as soon as they receive all the reference signals they require from the Queen. In the backward communication phase, Workers send their solutions back to their parent Queen, who collects them and sends the aggregated solution upward through the hierarchy. The market clearance can also be split across the branching nodes of the hierarchical structure, meaning all Workers can pay their DSO energy tariffs directly to their own Queens, who will send the aggregated amount to the Queens preceding them in the hierarchy. At the same time, the economic income generated by the system-level approach will be redistributed by the top-level Queen downwards to its children nodes.

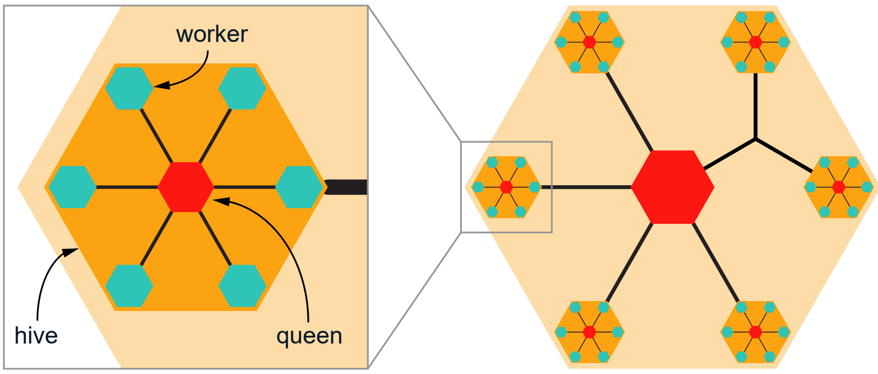

Figure 5: Multi level concept

4.5 Data modification attacks

During its initial phase, Hive Power will assume that Workers solve the decomposition problem in an honest manner. This seems a reasonable assumption, since solving the decomposition problem maximizes the social welfare of the entire energy community. It is however possible that some dishonest Workers could modify their local controllers, and assume a self-serving behavior, trying to increase their own benefits to the detriment of the other Workers. The self-serving Workers try to fool the financial redistribution mechanism if possible, or report false private information to increase their profits. Game theory can be used to analyse these kinds of opportunistic strategies, assuming the Worker's utility function is known.
In the following paragraph we analyse possible ways of deceiving the system, and we introduce a high-level description of the security measures that will be implemented during the project's Beta phase in order to prevent or mitigate data modification attacks.

4.5.1 Possible selfish strategies

We consider the following data modification attacks:

1. As previously stated, the declarative phase is followed by an actuation phase in which Workers are free to diverge from the previously obtained optimal strategy. If these phases are not modeled using game theory, Workers are allowed to cheat. However, this situation is very complex to model. Consider for example a market in which Workers pay a price proportional to their energy consumption and, if they consume more than the average, they pay an additional fee. If Worker A declares a very high consumption level for the following period, the other Workers could increase their consumption in the declarative phase, since they believe they are under-the-average consumers. If Worker A, after convergence, reveals that his needed consumption was in reality much less than the one he declared, he could now be an under-the-average consumer. The result is that Worker A averted the risk of paying the additional fee, to the detriment of the other Workers, who unnecessarily consumed more energy to remain average consumers based on Worker A's declarations.

2. It is likely that some Workers will try to solve a self-serving version of the distributed problem, even if by doing so they would not maximize the Hive's well-being . Consider for example the decomposition approach. If a Worker's utility function does not coincide with the decomposed problem, self-serving Workers can choose to solve another problem. For example, a typical formulation of the decomposed problem leads to the following Worker's optimization problem:

$$\arg\min_{\mathbf{u_i}\in\mathcal{U}_i} \|\mathbf{r}_i^\nu - (\hat{\mathbf{P}}_i^\nu + \mathbf{u}_i^\nu)\|_2^2 + \phi_g(\mathbf{u}_i^\nu) - \phi\left(\mathbf{u}_i^\nu, \sum \mathbf{u}_{-i}^\nu\right) \tag{4}$$

where $\mathbf{u}_i$ is the control action vector (for example an electrical battery input/output power) of the i_{th} Worker, $\mathbf{u}_{-i}$ represents the control actions of all the other agents, $\mathcal{U}_i$ is a set of Worker-specific problem constraints, $\mathbf{r}_\nu$ is a reference power profile provided by the aggregator at iteration ν, $\hat{\mathbf{P}}^\nu$ is the forecast worker's uncontrolled power profile, $\phi_g(\mathbf{u}_i^\nu)$ is the energy cost of the Worker and $\phi\left(\mathbf{u}_i^\nu, \sum \mathbf{u}_{-i}^\nu\right)$ is a case-specific objective, which depends on the aggregated profile of all the other Workers, $\sum \mathbf{u}_{-i}^\nu$. For example, in the case of SCC, $\phi\left(\mathbf{u}_i^\nu, \sum \mathbf{u}_{-i}^\nu\right)$ represents the difference in the economic surplus S, due to Worker i. In the first term, the coordination signal $\mathbf{r}^\nu$ induce cooperation among the Workers, allowing welfare maximization and to respect coupling constraints (e.g. bounded power flow in some points of the electrical grid). Being selfish, the worker could cheat solving

$$\arg\min_{\mathbf{u_i}\in\mathcal{U}_i} k\|\mathbf{r}_i^\nu - (\hat{\mathbf{P}}_i^\nu + \mathbf{u}_i^\nu)\|_2^2 + \phi_g(\mathbf{u}_i^\nu) - \phi\left(\mathbf{u}_i^\nu, \sum \mathbf{u}_{-i}^\nu\right) \tag{5}$$

where $k \leq 1$, or simply lying on the true value of $\hat{\mathbf{P}}_i^\nu$. In this way he maximizes his own utility, forcing the other prosumers to deal with the grid constraints, which decreases their economic benefits.

4.5.2 Proposed solutions

We propose the following solutions to mitigate the aforementioned selfish strategies:

1. The first example above of a self-serving strategy is a threat only when the number of Workers in the Hive is low or if the self-serving Worker is a big energy consumer/producer. Nevertheless, we plan to implement a reputation control mechanism that prevents this kind of behavior. In essence, the utility function will be increased with a coefficient related to the difference between the declaration and the actuation phase. This coefficient is strictly related to the Worker's forecasting ability. With the introduction of this coefficient, we are at the same time promoting higher forecasting accuracy and discouraging self-serving behavior.

2. In order to avoid the second self-serving strategy, the Worker optimization problem must reflect the Worker's own utility. Considering the decomposition approach, it is clear that decomposed problem 4 does not reflect the economic utility of the Worker, as previously explained, due to the first value in the equation.
One possible way to force the Workers to solve problem 4 would be to make them pay for the first value, or part of it [10][11]. However, this can result in increasing energy tariffs for the Workers, who could opt-out of the energy market. A second option is to use different decomposition approaches to further decompose problem 4 into two sub-problems, to be solved iteratively. More specifically, the first sub-problem can be represented by the minimization of the first and third values of problem 4. Since these terms are known and common to all the agents, the minimization of this sub-problem can be carried out by the Queen, which will return $\mathbf{u}_i^{\nu+1/2}$ to the i_{th} Worker, where $\nu + 1/2$ indicates half iteration of the optimization process. In this way would be difficult for the Workers to reconstruct the other Workers power profile, since this information would be only implicitly available through $\mathbf{u}_i^{\nu+1/2}$. This means that, in order to maximize $\phi\left(\mathbf{u}_i^\nu, \sum \mathbf{u}_{-i}\right)$, Workers must accept to cooperate with each other to fulfill grid constraints.
An alternative solution is based on detecting Workers self-serving behaviour, through an intrusion detection system. Exploiting information gathered during the declaration phase, we can statistically detect if Workers are actually solving problem 4.If not, Workers can be temporary banned from the Hive energy market. From a probability standpoint, this can be seen as a penalty related to the Worker's utility function.

5 Use cases

The Hive Power platform can be used to optimize the behavior of Workers and the management of energy produced by every node of the network. In the following paragraphs, we describe three main user cases for the platform.

Figure 6: General components of a Hive Power platform

5.1 Self-consumption Communities (SCC)

Self-consumption Communities consist of a set of prosumers that exchange energy to maximize the **group's autarky**. Electrical generators and consumers in an SCC must be connected to low-voltage feeders on the same substation. They can internally optimize the synchronization of their energy production/consumption by exploiting electric storage and demand-side management. An example of a SCC is a **condominium** where the solar energy produced on the rooftop is consumed by the tenants.
Other members of the community, living in the same geographical district, also consume the solar energy produced from their own rooftops. When the solar system happens not to produce enough (e.g. at night or during cloudy days), the necessary electricity can be bought from the DSO grid. Similarly, the community can also sell excess solar power to the national grid (e.g. during summer days) and receive financial remuneration. Storage batteries can be used to optimize the self-consumption capabilities of the SCC network. The clear economic advantage of an SCC comes from the large gap between purchase prices of energy (around 0.20 USD/kWh) and the selling price of energy (around 0.05 USD/kWh) from the national grid.

5.2 Micro-Grid

As defined by United States Department of Energy, micro-grids are groups of interconnected energy loads and distributed energy resources within clearly defined electrical boundaries that act as a single controllable entity with respect to the grid. Several energy associations agree that micro-grids will enjoy a rapid deployment in the coming years as stated by the International Renewable Energy Agency in the IRENA Innovation Outlook: Renewable Minigrids:

> *Renewable mini-grids continue to gain momentum as energy solutions in areas where energy demand is not fulfilled, and where grid extension is not a cost-effective alternative. Renewable mini-grids are reaching maturity, as shown by their improved reliability, reduced environmental impact, enabling of increased local control over energy used, and sustained cost reductions. Renewable mini-grids represent a growing market that is potentially worth more than USD 200 billion annually. Renewables can be mixed with diesel-fueled capacity to convert between 50 and 250 gigawatts (GW) of capacity to hybrid mini-grids.*

We distinguish two kinds of micro-grids, which will be operated within the Hive Power platform according to different smart-pricing schemes:

- **Commercial and institutional (C&I) micro-grids**, aggregating existing on-site energy generation with multiple loads that are located in closely geography, and in which the Administrator is able to easily manage them. Examples are military bases and industrial micro-grids, where the security of the power supply security and its reliability are very important.

C&I micro-grids can connect and disconnect from the grid to enable it to operate in both grid-connected or island-mode. This case is similar to an SCC, where the energy generation capacity is large enough to cover most of the local energy demand.

- **Remote micro-grids** are never connected to the national grid and operate in island mode because of economic issues or geography location. In this case, it is fundamental to guarantee the power supply for critical loads (such as emergency rooms) and optimize the management of storage batteries. We aim to achieve this with smart pricing algorithms that force users to collaborate in order to guarantee optimal network performance.

5.3 Distribution Grid

The third use case concerns the low-voltage distribution grids, where an entire geographical district is connected to the national medium-voltage grid, involving a few or possibly hundreds of prosumers with different profiles:

- *Residential houses*, with domestic energy loads and eventually controllable loads such as heat pumps and water boilers
- *Commercial*, such as offices, small shops and malls
- *Industrial*, with small factories consuming and producing a relevant quantity of energy

Figure 7: Nodes of a Distribution Grid

The Hive Power platform users can employ energy generation systems, which are typically solar, and storage such as electric batteries. Innovative nodes of the networks can include district batteries (that can profit from economy of scale and higher efficiency), shared solar plants (e.g. in condominiums or parking lots) and electric vehicle charging stations. The end result is that all network users benefit from participating in the Hive Power platform, thanks to convenient tariffs (to buy and sell energy), so their assets (solar and batteries) become more profitable, providing economic benefits to both energy producers and consumers.

6 Business Model

Hive Power is a technology provider of an innovative platform that is fully open to existing and new energy actors. Hive Power's goal is to maintain a lean structure, focused on research and development activities. The revenue streams necessary to sustain the development and the maintenance of the platform will come from energy transaction fees and fees for additional services.

6.1 Energy transaction fees

On each energy transaction performed inside the Hive Power platforms, Hive Power retains a small fee that is proportionally sized in order not to significantly affect the benefits of the stakeholders, and collected during the settlement of the Hive Markets. This fee represents compensation for the benefits that Hive Power brings to the platform participants:

- Users are charged a discounted tariff (with respect to the actual energy prices) and can access locally-produced renewable energy. Energy produced by solar and other local production systems can be purchased at a bonus tariff.
- Distribution System Operators (DSO) avoid the risk of user customer loss (death spiral). By employing a prepaid billing system, they benefit from a reliable revenue stream.
- Hardware manufacturers can use the Hive Power platform to add blockchain capabilities with reduced R&D costs, and without paying licensing fees. Moreover, the deployment of SCCs powered by the Hive Power platform can enhance the distribution and sale of new hardware components.

6.2 Forecasting service

In addition to basic functionalities, Hive Power will provide an accurate forecasting service for the energy consumption/production of each Ethereum Meter (Worker). The service can be used by Workers to optimize their performance, and it will promote demand-side management of energy in storage batteries and thermal loads. A fee for this service will be charged to the Worker and will be deducted from the additional revenues generated to the Worker's owner by adherence to the projected forecast for energy production and consumption.

6.3 Energy load management service

In future editions of the Hive Power platform, the Worker will be able to control deferrable energy loads (such as heat pumps and boilers) and storage batteries to optimize its own performance (user's comfort and revenues). Energy load management features will also increase the revenues a Hive can generate in the Flexibility Market from their ability to be flexible in their energy supply and demand. The Worker's owner will be charged a fee to activate the energy load management service.

7 Roadmap

The Hive Power roadmap begins with the launch of a Token Generation Event, to finance the development and the implementation of an energy exchange platform for the aforementioned use cases. Our first goal is to develop an open platform that can be used by any energy actor (e.g. a DSO, a Self-consumption Community manager or a micro-grid manager).
Initially we will focus on the development, testing and validation of the Hive energy market design using an energy meter prototype, and then we will finalize the entire system design. The testing of the platform will take place in one of the largest pilot and demonstration projects in Europe. In the meantime, we will continue ongoing discussions with potential partners interested in adding value to their projects by incorporating the Hive Power platform into their hardware and software energy solutions.

7.1 Hive Power Beta (Q4 2017 - Q2 2018)

The Hive Power Beta version focuses on a simple user case of a single Self-consumption Community. The use-case consists of an apartment building owner who creates a Hive by staking HVTs and registers his Ethereum Meter prototypes in the Hive. The apartment tenants purchase tokens and use them to prepay the energy delivered to the apartment. By participating in the Hive, the tenants save money by paying for their electricity at lower costs that the normal utility tariffs. The Building owner automatically receives the tenant pre-payments and can use a portion of that to pay the DSO energy bill. By using the Hive Power platform, the tenants save on administrative costs, thanks to the automated billing system in the smart contracts. The building owner's financial risk is diminished because electricity tokens are prepaid by the tenants. The accuracy of the billing will be verifiable because the energy measured and certified by the meters will be stored on the blockchain.
The Ethereum Meter proof-of-concept will be run on the Hive Power Development Kit, a small open source computer connected to a triphase energy meter. The goal is to test the market design, and therefore payment channels will be implemented through a simple contract.

7.1.1 Demo Hive testbed

In order to develop a preliminary case of the aforementioned Self-consumption Community and the Ethereum Meter prototypes, in Q4 2017 Hive Power built its first demo testbed, called the **Demo Hive**, which is shown in Figure 8. The demo simulates a simple Hive, with two Workers (a producer and a consumer) and a Queen, who manages the interactions between the Hive components and the external grid. From a hardware point of view, the Workers are run on SmartPI boards, an acquisition board for the electrical measurements that is connected to a Raspberry Pi 3. The Queen of the Hive also runs on a Raspberry Pi 3. Regarding the software, the main purpose of the demo testbed is to maximize the Hive's autonomy by exploiting Hive Power tariffs, which are more economical than non-Hive tariffs. The Demo functions as follows: at the end of a day, (simulated as a period of 10 minutes), the produced/consumed energy of the Workers is calculated and tokenized using an ERC20 smart contract deployed on the Ropsten blockchain. The name of the token is DHT (Demo Hive energy Token). Each Worker owns a certain amount of DHTs and sends a part of them as payment to the energy producers (i.e. the other Workers, or through the Queen to the external grid if their energy consumption exceeds their energy production). The devices powering the Workers (Ethereum Meters) and the Queen are always con nected and synchronized with a Ropsten network to be able to transfer/receive token payments.

Figure 8: Demo Hive testbed

Figure 9: Hive Power Alpha Development Kit

7.2 Hive Power 1.0: Initial Release (Q2 2018)

At this stage, the Hive Power platform will be operational with regards to its basic functionalities. We will release:

1. **Smart contracts**: The Beekeeper, Hive and auxiliary contracts will be deployed on the Ethereum main net. DAI will be the only stable token available. All the contracts will be upgradeable.

2. **Admin and User Apps**: The initial version of the Dapps will be released. It will be accessible through the browser using Metamask or using mobile dapp browsers like e.g. cipher. Payments will be done with DAI only.

3. **Worker software**: The software running on a Worker will be released as source code and raspberry pi 3 image. It will be running the base version of HONEY algorithm. It will be able to read meters via Modbus TCP/IP interface and to control a battery via a RESTful web API.

4. **Queen server**: A server will be set up and its code will be released open-source. The Queen server will take care of coordinating the workers in the Hives. The Queen server will also host a time-series database in which the workers will have the possibility to store high resolution consumption and production data in order to visualize them in the User App.

7.3 Hive Power 1.1 (Q2 2018 - Q1 2019)

Pilots We will setup a Hive in a low voltage district in Switzerland, which will consist of single household Workers and multi-family buildings, coupled to a low voltage grid, interacting with the Hive Power ecosystem using the Hive User App. The Queen will be operated by the local DSO. The base version of the HONEY algorithm will be tested, while the multilevel hierarchical version will be released. At the same time, the security measures described in section 4.5 aiming at eliminating possible selfish behaviors and data modification attacks will be developed and tested in simulation.

Ethereum Meter An Ethereum Meter prototype will be developed in collaboration with a leading meter manufacturer. This energy meter will tokenize exchanged energy with a certificate of origin. The energy token used for transactions in this user case will not be exclusively tied to the Hive Power platform because they can also be used within other types of energy communities. The Ethereum Meter prototype will integrate an industrial system-on-chip. Lab tests are already on-going and we have already verified that the chip is able to run an Ethereum light client, a state channel client, and the Worker software.

Figure 10: Top: Current system-on-chip development kit. Bottom: Ethereum Meter tentative design.

Energy load management In Hive Power version 1.0, the Worker will gain the ability to control energy loads (such as heat pumps and boilers) and storage batteries to optimize its own performance (comfort and revenues).

7.4 Hive Power 2.0 (Q1 2019 - Q1 2020)

The Hive Power 2.0 platform will include the basic functionalities that were ready in version 1.0 and are now validated for operation in low-voltage district grids. Innovative features will be implemented to improve economic and technical performance. In parallel, a micro-grid testbed will be set up in a developing country.

Real-time operation Real-time operation of the Hive (inside the currently active time slot) is implemented via the off-chain high-speed Liquidity Network payment platform (http://liquidity.network/). Deviations from agreed energy production and consumption are settled by a reactive control algorithm leveraging the Hive's internal flexibility (energy storage, controllable load, or energy generation) or by bridging a Hive with larger and more flexible Hives. This process is managed through a Flexibility Pool where tokens allocated (or staked) for real-time operation are used. Staking tokens to the Flexibility Pool is a prerequisite to participate in the Hive's energy market.

Flexibility market In addition to participating in the Hive energy market, Workers can participate in the Hive Flexibility Market. Within the Flexibility Market, Hives interconnect with each other through virtual flexibility connectors, which allow them to exchange spare energy production and consumption flexibility. If so desired they can pool their extra energy storage to reach the required amount to be able to supply ancillary services to the grid in the form of primary and secondary frequency control, as well as congestion management. An energy load management feature will also increase revenues Workers receive from the Flexibility Market.

Fraud detection To prevent fraudulent actions on the platform, each Hive Owner has to stake a pre-determined number of tokens for each Worker in the Hive. By comparing a Worker's payment transactions with real-time data, the Hive platform is able to detect fraudulent actions. Corrective measures will include loss of staked tokens, and temporary or permanent ban of the Worker from the Hive Power platform.

Hives flexibility Hives are virtual elements that have the flexibility to be reconfigured, merged and split.

7.5 Hive Power 3.0 (Q1 2020 and beyond)

After the release of Hive Power version 2.0 many advanced features will be evaluated.

Investment pool Workers and Hive Owners can decide to reinvest their profits in the Hive infrastructure (energy generation, storage, transport). The investments will increase profits by exchanging with nearby Hives.

Topology discovery When a new Worker is activated within a Hive, it will automatically receive information about the existence of other Hives within its geographical area.

Self-healing A Hive can automatically detect a Worker (Ethereum Meter) malfunction, tagging the Worker as defective. The Hive is programmed to adapt and overcome malfunctions in the platform by activating spare energy flexibility to cover the loss in production of a malfunctioning meter.

Multi-owner Hives Multiple entities will be allowed to be co-owners of a Hive, sharing decisions, revenues and risks.

8 People

8.1 Team

The founders of Hive Power have a wide range of academic and interdisciplinary backgrounds in the field of energy distribution grids/smart grids, ranging from micro-technical engineering to applied mathematics and informatics. Over the last seven years the team worked hard to develop a unique, effective and innovative algorithm based on a self-learning approach addressing the typical challenges of energy load management in smart grids. One of the largest utilities in Switzerland has integrated this approach in a new product to be sold on the Swiss market by since Autumn 2017 under the trademark GridSense.

8.1.1 Gianluca Corbellini

Gianluca Corbellini holds a M.Sc. in Mathematical Engineering from the Politecnico di Milano, focused on mathematical modelling, optimization and artificial intelligence. He has extensive experience working with multinational corporations in the energy field, having been asset manager for photovoltaic plants and research engineer in the oil and gas industry. In the University of Applied Science and Arts of Southern Switzerland (SUPSI) he is involved in the modelling of photovoltaic plants, and in the development of new business models for the optimization of smart grids. He was also a lecturer for the SUPSI course "Design of Energy Systems" regarding the design of micro-grids.

8.1.2 Davide Rivola

Davide Rivola is a senior researcher with a multi-disciplinary micro-engineering background. He is leading the Energy Systems research sector at SUPSI. Before his research activities he gained several years of industrial experience, designing industrial automation systems and developing real-time software for embedded electronics. During the last seven years he researched, developed and trialed in pilot projects fully decentralized energy management systems for self-consumption optimization and grid instability reduction. He is personally involved in Blockchain technology since 2013, with an enthusiasm that only grew during time.

8.1.3 Vasco Medici

Vasco Medici received a M.Sc. in Micro-Engineering from the Swiss Federal Institute of Technology in Lausanne and a Ph.D. in Neuroinformatics from the Swiss Federal Institute of Technology in Zürich. He previously worked in the development of real-time 3D video-based tracking applications. He currently leads the Intelligent Energy Systems Team at SUPSI, where he also teaches the "Introduction to Smart Grid" course. His main competences are system identification, algorithmics, modeling and simulation. He is the coordinator at SUPSI for the Swiss Competence Center for Energy Research on Future Swiss Electrical Infrastructure (SCCER FURIES). In close collaboration with industrial partners, his team runs a number of pilot projects in the field of demand side management applied to smart grids.

8.1.4 Lorenzo Nespoli

Lorenzo Nespoli received the M.Sc. degree in Energy Engineering from Politecnico di Milano in 2013. Since 2014 he works on multiphysics simulations and electric grid optimization at SUPSI, where he is lecturer for the "Introduction to Smart Grid" course. He is a Ph.D. candidate at the Swiss Federal Institute of Technology in Lausanne, where he is working on decentralized control algorithms and model-based forecasts for demand side management in the distribution grid, in the context of Swiss Competence Center for Energy Research - Future Swiss Electrical Infrastructure SCCER FURIES.

8.1.5 Davide Strepparava

Davide Strepparava is a researcher of the Intelligent Energy Systems Team at the Institute for Sustainability Applied to the Built Environment at SUPSI. He received a M.Sc. in Computer Science from the Politecnico di Milano. Before his academic career he worked for several years in the building automation and access control industries. He has a high level of experience in data science and database management. At SUPSI he is involved in research projects mainly related to the monitoring of solar plants and smart grids. Over the past two years he has worked extensively with blockchain technology, especially focused on the Ethereum platform, working on research projects related to decentralized and smart energy markets.

References

[1] I. Pérez-Arriaga, C. Knittle, and M. E. Initiative, "Utility of the future: An mit energy iniative response to an industry in transition." http://energy.mit.edu/publication/utility-future-report, 2016.

[2] P. Bronski, J. Creyts, L. Guccione, M. Madrazo, J. Mandel, B. Rader, D. Seif, P. Lilienthal, J. Glassmire, J. Abromowitz, *et al.*, "The economics of grid defection: When and where distributed solar generation plus storage competes with traditional utility service," *Rocky Mountain Institute*, pp. 1–73, 2014.

[3] The Economist Group Limited, "A world turned upside down." https://goo.gl/hWu16z, 2017. [Online; accessed 27-July-2017].

[4] D. K. Molzahn, F. Dorfler, H. Sandberg, S. H. Low, S. Chakrabarti, R. Baldick, and J. Lavaei, "A Survey of Distributed Optimization and Control Algorithms for Electric Power Systems," *IEEE Transactions on Smart Grid*, vol. 3053, no. c, pp. 1–1, 2017.

[5] B. Yang and M. Johansson, "Distributed optimization and games: A tutorial overview," *Lecture Notes in Control and Information Sciences*, vol. 406, pp. 109–148, 2010.

[6] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein, "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers," *Foundations and Trends® in Machine Learning*, vol. 3, no. 1, pp. 1–122, 2010.

[7] H. Almasalma, J. Engels, and G. Deconinck, "Dual-decomposition-based peer-to-peer voltage control for distribution networks," vol. 2017, no. 1, pp. 1718–1721, 2017.

[8] C. Mugnier, K. Christakou, J. Jaton, M. De Vivo, M. Carpita, and M. Paolone, "Model-less/measurement-based computation of voltage sensitivities in unbalanced electrical distribution networks," *19th Power Systems Computation Conference, PSCC 2016*, 2016.

[9] S. Weckx, S. Member, R. D. Hulst, J. Driesen, and S. Member, "Voltage Sensitivity Analysis of a Laboratory Distribution Grid With Incomplete Data," *IEEE Transactions on Smart Grid*, vol. 6, no. 3, pp. 1271–1280, 2015.

[10] D. Paccagnan, M. Kamgarpour, B. Gentile, F. Parise, J. Lygeros, and D. Paccagnan, "Distributed computation of Nash Equilibria in aggregative games with coupling constraints," in *2016 IEEE 55th Conference on Decision and Control (CDC)*, (Las Vegas, USA), pp. 6123–6128, 2016.

[11] S. Grammatico, "Dynamic Control of Agents Playing Aggregative Games with Coupling Constraints," *IEEE Transactions on Automatic Control*, vol. 62, no. 9, pp. 4537–4548, 2017.

CONTACTS

email:
support@hivepower.tech

web:
www.hivepower.tech

telegram:
t.me/hivepower